INTERMEDIATE MICROECONOMICS (EC201)

Course duration: 54 hours lecture and class time (Over three weeks)

Summer School Programme Area: Economics

LSE Teaching Department: Department of Economics

Lead Faculty: Dr Andrew Ellis (first-half) and Dr Francesco Nava (second-half) (Dept. of Economics)

Pre-requisites: Introductory Microeconomics, Calculus, Introductory Statistics.

Course Overview:

The aim of this course is to give students the conceptual basis and the necessary tools for understanding modern microeconomics at the intermediate level. In the context of this theoretical framework, the course explores a number of applied issues such as contract design, insurance, and ownership structures.

The course covers 6 broad areas:

- Consumer Theory
- The Theory of the Firm
- General Equilibrium
- Game Theory
- Oligopolistic Markets
- Information Economics

The theory of the consumer explores the demand side, while the theory of the firm discusses the supply side of the economy. General equilibrium puts the two parts together and discusses welfare implications, including in the presence of externalities.

The second part of the course introduces basic concepts in non-cooperative game theory, emphasising the strategic aspect of economic interaction. Game theory is then applied to analyse informational problems in economics, in particular problems of hidden information (adverse selection) and hidden action (moral hazard).

Whilst not all of the presentations will be as mathematical as that provided in the course text, knowledge of calculus is essential for the study of quantitative solutions to economic problems and, indeed, enhances one's understanding of the underlying concepts. In class on the first day of the course, differential calculus will be reviewed, and students will be introduced to the technique of constrained maximisation due to Lagrange.

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Lecture Plan:

Topic 1 - Consumer Theory

This part of the course studies consumers' preferences and budget constraints. It derives individual demand functions and analyses how these can be aggregated to build the market demand curve. Also the concepts of consumer surplus and price indexes will be discussed.

Topic 2 - The Theory of the Firm

This part of the course reviews the structure of production and studies the profit maximisation problem of the firm. It analyses how the firm responds to market stimuli both in the short and in the long run. The issues above are addressed for perfectly competitive firms as well as for monopolies. The market supply is also derived as the aggregate supply of firms that produce identical products.

Topic 3 - General Equilibrium and Welfare

The topic provides conditions for an economy to reach equilibrium and studies how equilibrium prices and quantities are determined. It identifies conditions under which the market equilibrium is efficient as well as those under which a central planner can implement an efficient allocation as a market-equilibrium.

Topic 4 - Game Theory

Game theory is used to study strategic interactions between agents and is a fundamental tool in modern economics. This topic analyses several general classes of games and defines relevant solution concepts in each of these. It begins by discussing static games of complete and incomplete information and by defining Dominant Strategy equilibria and Nash equilibria, in pure and mixed strategies. It proceeds by analysing dynamic and repeated games with complete information, and by introducing Subgame Perfection.

Topic 5 – Oligopolistic Markets

Two main game theoretic applications are considered. The first looks at the strategic behaviour of firms in a duopoly. The second looks at a model of entry-deterrence with pre-commitment strategies.

Topic 6- Information Economics

In many environments, agents involved in economic transactions have access to different information about profitability of trade between them. The final topic considers such scenarios: firstly, in adverse-selection and signalling models where one agent cannot observe another agent's characteristics (insurance market); secondly, in moral hazard models where one agent cannot observe another agent's action. The optimal design of contracts to provide incentives and elicit information is the main aim of the topic.

Suggested Reading:

The following text is recommended as additional reading to the lecture notes and class exercises.

Christopher Snyder and Walter Nicholson, *Microeconomic Theory: Basic Principles and Extensions*, (11th edition, International Edition), South-Western College Publishing (2011).

Please note that the textbook differs from previous editions as well as the American edition.

Formative Assessments:

1) Format: Hand-in Problem Set
   Date: Friday week one
   Results due: Tuesday week two

2) Format: Hand-in Problem Set
   Date: Tuesday of week three
   Results due: Thursday week three

Summative Assessments:

1) Format and Weight: Two Hour Midterm Exam (50%)
   Date: Wednesday of week two
   Results due: Monday of week three

2) Format and Weight: Two Hour Final Exam (50%)
   Date: Friday of week three
   Results due: Within a week

The precise time and location of the exams will be circulated during the programme.

Credit Transfer: If you are hoping to earn credit by taking this course, please ensure that you confirm it is eligible for credit transfer well in advance of the start date. Please discuss this directly with your home institution or Study Abroad Advisor.

As a guide, our LSE Summer School courses are typically eligible for three or four credits within the US system and 7.5 ECTS in Europe. Different institutions and countries can, and will, vary. You will receive a digital transcript and a printed certificate following your successful completion of the course in order to make arrangements for transfer of credit.

If you have any queries, please direct them to summer.school@lse.ac.uk

Intermediate Microeconomics EC201

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Introduction

The second part of the course discusses:

- Classical models of Choice Under Uncertainty
- Fundamental concepts and results in Game Theory:
    - Strategic Decision Making
    - Static and Dynamic Solution Concepts
    - Folk Theorems
- Classical game theortic applications:
    - Imperfect Competition
    - Adverse Selection
    - Signaling
    - Moral Hazard

Why Game Theory?

In many relevant economic environments the well being of individuals depends on decisions made by others

Game theory sheds light on behavior in strategic environments

Game theoretic models are commonly used to:

1. study oligopolistic competition [cartels, competition laws]
2. model insurance contracts [public option and socialized healthcare]
3. write incentive contracts [long term incentive contracts]
4. understand voting and political systems
5. design markets, mechanisms, and auctions
6. bid in auctions [ebay, spectrum auctions]
7. model externalities and public goods
8. make decisions on war strategies

Objectives of the Course

1. Introduce the classical model of choice under uncertainty.
2. Introduce models of behavior in strategic environments.
3. Understand how information affects strategic behavior.
4. Understand decisions and threats in dynamic environments.
5. Understand how reputation affects behavior in repeated interactions.
6. Introduce classical models of imperfect competition.
7. Understand screening or signalling in models of hidden information.
8. Understand incentive contracts in settings with moral hazard.

General Information

Lecturer: Francesco Nava, 32L Room 3.20

Email: f.nava@lse.ac.uk

Course Website: http://shortcourses.lse.ac.uk/

Enrollment Key: EC20119

Time: 2pm-5pm & morning classes

Location: TBD

Lectures: Week 2: Wed, Thur, Fri; Week 3: Mon, Tue, Wed

Classes: Week 2: Thur, Fri; Week 3: Mon, Tue, Wed, Thur

Daily Course Program

1. Choice Under Uncertainty & Static Games
   - Readings: SN7 & SN8; Supplementary: O2.1, O2.6-9
2. Mixed Strategy and Imperfect Competition
   - Readings: SN8 & SN15; Supplementary: O4.1-4, O3.1-2
3. Incomplete Information Games & Dynamic Games
   - Readings: SN8; Supplementary: O9.1-3 & O5.1-4
4. Dynamic Games & Repeated Games
   - Readings: SN8; Supplementary: O15.1-2
5. Asymmetric Information & Adverse Selection
   - Readings: SN18; Supplementary: S2, S3-2, S4
6. Signalling & Moral Hazard
   - Readings: SN18; Supplementary: S5.1 to S5.3.5

Course Requirements

- Classes follow lectures with 1 day lag
- The work for each class consists of one or more exercises
  Some answers to exercises will be posted
- A marked problem set will be due on Tuesday of week 3

Principal features of the final exam are:

- The final consists of 2 parts A and B
- Part A is worth 60% of the final grade
  - It consists of short questions
  - You will have to choose 3 out of 4 questions
- Part B is worth 40% of the final grade
  - It consists of long questions
  - You will have to choose 1 out of 2 questions

Materials

Main Textbook

Microeconomic Theory, Snyder & Nicholson, Thomson, 11th Edition [SN]

Slides

Include all materials required for the examinations

Slides are posted on the course website

Slides labeled "Extra" are not examinable

Supplementary Readings

An Introduction to Game Theory, Osborne, Oxford Press, 2003 [O]

Economics of Contracts, Salanie, MIT Press, 2005 [S]

Choice Under Uncertainty

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Choice Under Uncertainty:

- Statistics Review
- Definitions:
  - Lottery & Fair Lottery
  - Expected Utility
  - Risk Attitudes (Aversion, Neutrality, Loving)
  - Certainty Equivalent
- Measures of Risk Aversion:
  - Relative Risk Aversion
  - Absolute Risk Aversion
- Insurance, a first take:
  - Actuarially Fair Insurance
  - Under-insurance at Unfair Prices

Statistics Review

Statistics Review

A **random variable** X is a variable that records the possible outcomes x of a random event

Any random variable X is characterized by:

- the set of possible outcomes that can occur ($\mathbb{X}$) and by
- a probability distribution over the possible outcomes ($f : \mathbb{X} \to [0, 1]$)

Given a numerical random variable $\{\mathbb{X}, f\}$ (that is $\mathbb{X} \subseteq \mathbb{R}$):

- The probability of $X = x$ is denoted by $f(x)$
- The **expected value** of the RV X is denoted and defined by:

$$E(X) = \sum_{x \in \mathbb{X}} x f(x)$$

- The **variance** of the RV X is denoted and defined by:

$$V(X) = \sum_{x \in \mathbb{X}} (x - E(x))^2 f(x)$$

Expected Utility

Lotteries and Fair Lotteries

A **lottery** X is a random variable over **monetary outcomes**

- Any lottery is characterized by an outcome set and a probability distribution over monetary outcomes $\{\mathbb{X}, f\}$
- In general monetary outcomes can be negative

A lottery X is said to be **fair** if $E(X) = 0$

Examples:

- $\mathbb{X} = \{2, -1\}$, $f(2) = 1/3$, $f(-1) = 2/3$ is fair since:

$$E(X) = 2 * (1/3) - 1 * (2/3) = 0$$

- $\mathbb{X} = \{2, -1\}$, $f(2) = 1/2$, $f(-1) = 1/2$ is unfair since:

$$E(X) = 2 * (1/2) - 1 * (1/2) = 1/2$$

Such a lottery would be fair if an entry fee of 1/2 were charged

Expected Utility & Risk Preferences

Consider a decision maker that has preferences over monetary outcomes defined by a strictly increasing utility function $u : \mathbb{X} \rightarrow \mathbb{R}$

The **expected utility** of a lottery X is defined by:

$$E(u(X)) = \sum_{x \in \mathbb{X}} u(x) f(x)$$

Expected utility may differ from the utility of the expected value!!!

Preferences over lotteries:

- An individual is **risk averse** if $u'' < 0$
- An individual is **risk neutral** if $u'' = 0$
- An individual is **risk loving** if $u'' > 0$

Risk Preferences

Risk Preferences

Preferences over lotteries:

- A **risk averse** individual prefers $E(X)$ to the lottery X:
$$E(u(X)) < u(E(X))$$
- A **risk neutral** individual is indifferent between a lottery X and $E(X)$:
$$E(u(X)) = u(E(X))$$
- A **risk loving** individual prefers a lottery X to $E(X)$:
$$E(u(X)) > u(E(X))$$

Example, consider $\mathbb{X} = \{1, 9\}$ and $f(1) = f(9) = 1/2$:

- If preferences are concave, say $u(x) = x^{1/2}$, we get that:
$$E(u(X)) = 2 < \sqrt{5} = u(E(X))$$
- If preference are convex, say $u(x) = x^2$, we get that:
$$E(u(X)) = 41 > 25 = u(E(X))$$

Certainty Equivalent

The **certainty equivalent** x_{CE} of a lottery X is defined by:

$$u(x_{CE}) = E(u(X)) \quad \Leftrightarrow \quad x_{CE} = u^{-1}(E(u(X)))$$

The certainty equivalent is the amount of money x_{CE} that leaves the individual indifferent between the lottery X and the certain outcome x_{CE}

For any given lottery X we have that if:

- an individual is **risk averse then** $x_{CE} < E(X)$
- an individual is **risk neutral then** $x_{CE} = E(X)$
- an individual is **risk loving then** $x_{CE} > E(X)$

Risk Aversion

Consider a lottery $\mathbb{X} = \{a, b\}$, $f(a) = p$, $f(b) = 1 - p$ and a risk averse individual:

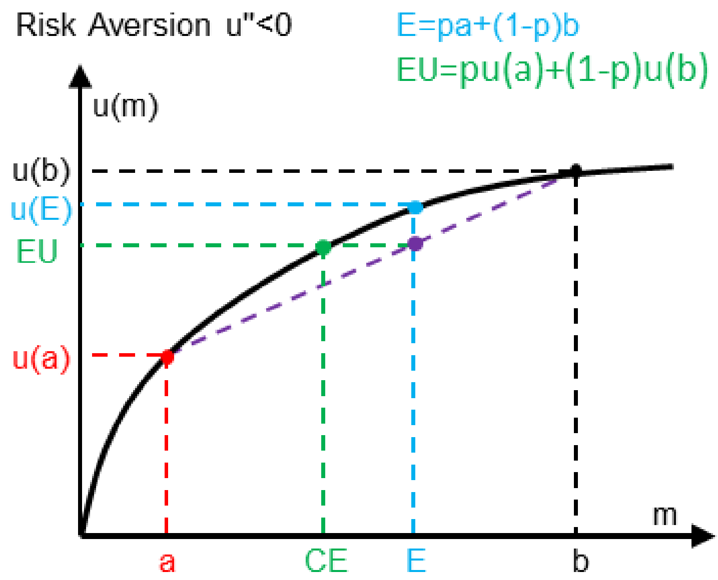

Risk Loving

Consider a lottery $\mathbb{X} = \{a, b\}$, $f(a) = p$, $f(b) = 1 - p$ and a risk loving individual:

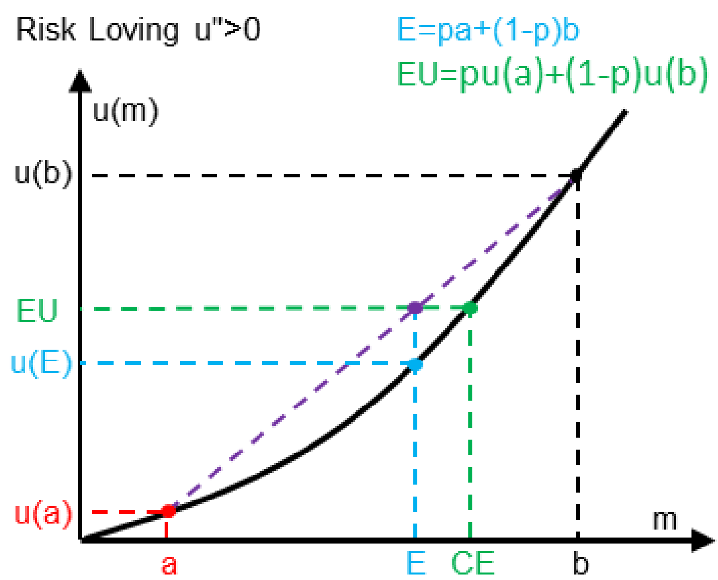

Relative Risk Aversion (Pratt)

The **coefficient of relative risk aversion** is defined by:

$$R(x) = -x\frac{u''(x)}{u'(x)}$$

It **is a measure of risk aversion** of individuals

Preferences u display constant relative risk aversion CRRA if $R' = 0$

Any CRRA preference takes the form:

$$u(x) = \alpha x^{\gamma} + \beta \text{ for } \alpha > 0,\ \gamma \in (0,1) \ \& \ \forall \beta$$

If preferences are CRRA then for any $k > 0$:

$$E(u(X)) = u(x_{CE}) \quad \Leftrightarrow \quad E(u(kX)) = u(kx_{CE})$$

Risk aversion does not change with proportional changes in the stakes

Absolute Risk Aversion

The **coefficient of absolute risk aversion** is defined by:

$$A(x) = -\frac{u''(x)}{u'(x)}$$

It **is a measure of risk aversion** of individuals

Preferences u display constant relative risk aversion CARA if $A' = 0$

Any CARA preference takes the form:

$$u(x) = -\alpha e^{-\gamma x} + \beta \text{ for } \alpha > 0,\ \gamma > 0 \ \& \ \forall \beta$$

If preferences are CARA then for any $k > 0$:

$$E(u(X)) = u(x_{CE}) \quad \Leftrightarrow \quad E(u(k+X)) = u(k+x_{CE})$$

Risk aversion does not change with additive changes in the stakes

A Model of Insurance

A Simple Insurance Model I

Consider the following decision problem faced by a risk averse individual:

- There are **two possible states** of the world $\{H, S\}$
- The individual can be healthy H or sick S
- The probability of being sick is p
- The income of an individual is:
  - Y if healthy
  - $Y - L$ if sick
- Let y denote the consumption if healthy and x if sick
- Preference satisfy $u'' \in (-\infty, 0)$ and:

$$pu(x) + (1 - p)u(y)$$

A Simple Insurance Model II

- Consumers can buy insurance coverage $z \in [0, L]$
- The **unit price** of insurance is q
- Therefore the **total premium** is qz
- If they do so, their consumption in the two states becomes:

$$\begin{aligned} y &= Y - qz \\ x &= Y - L - qz + z = Y - L + (1-q)z \end{aligned}$$

- If so, the problem of a consumer becomes:

$$\max_z pu(x) + (1-p)u(y)$$

- FOC with respect to z requires:

$$p(1-q)u'(x) = (1-p)qu'(y)$$

A Simple Insurance Model: Demand

- FOC can be written in terms of MRS as:

$$\frac{u'(x)}{u'(y)} = \frac{1-p}{p}\frac{q}{1-q}$$

- Thus a consumer of type t wants:

Full Insurance:	$z = L$	if	$q = p$
Under Insurance:	$z < L$	if	$q > p$
Over Insurance:	$z > L$	if	$q < p$

- This is the case because $u'' < 0$ implies:

$$q \left\{ \begin{array}{c} = \\ > \\ < \end{array} \right\} p \quad \Leftrightarrow \quad \frac{u'(x)}{u'(y)} \left\{ \begin{array}{c} = \\ > \\ < \end{array} \right\} 1 \quad \Leftrightarrow \quad x \left\{ \begin{array}{c} = \\ < \\ > \end{array} \right\} y$$

A Simple Insurance Model: Supply

- Suppose that an insurance company is selling the contract
- Its **profits on the contract** (q, z) are given by:

$$\pi(q,z) = (1-p)qz - p(1-q)z = (q-p)z$$

- The company' profits are:
  - positive if $q > p$
  - negative if $q < p$
  - zero if $q = p$
- The insurance price is **actuarially fair** if $q = p$

Full Insurance at Fair Prices

The consumer **fully insures at actuarially fair prices**.

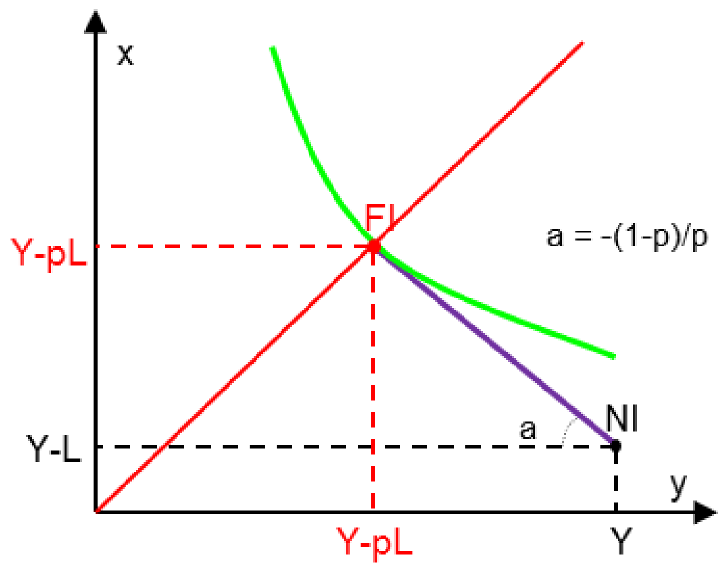

Under-Insurance at Unfair Prices

The consumer **under-insures at prices are not fair and** $q > p$.

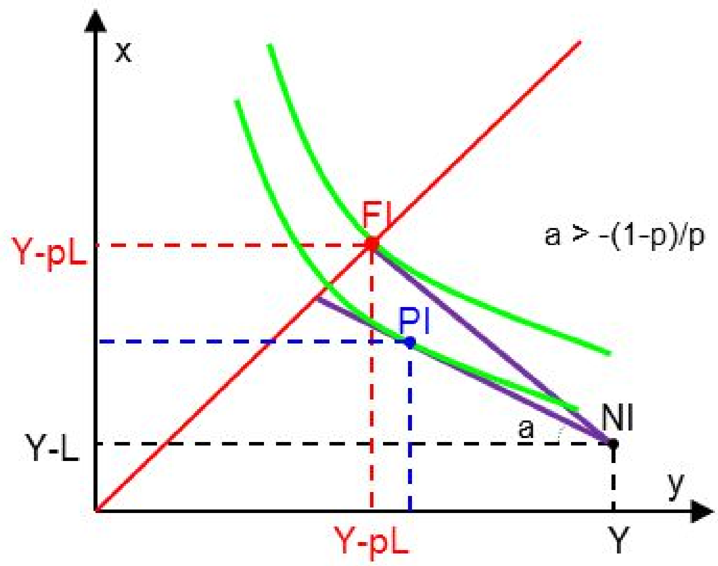

Static Complete Information Games

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2018

Summary: Games of Complete Information

- Definitions:
  - Game: Players, Actions, Payoffs
  - Pure Strategy
  - Best Response
  - Mixed Strategy
- Solution Concepts:
  - Dominant Strategy Equilibrium
  - Nash Equilibrium
- Properties of Nash Equilibria:
  - Multiplicity
  - Inefficiency
  - Non-Existence in Pure Strategies
  - Existence in Mixed Strategies
- Examples

Introduction to Games

Any environment in which the choices of an individual affect the well being of others can be modeled as a game.

What pins down a specific game:

- Who participates in a game [Players]
- The choices that participants have [Choices]
- The well being of individuals [Payoffs]
- The information that individuals have [Rules of the Game]
- The timing of events and decisions [Rules of the Game]

Introduction to Games

In all models discussed in the first part of the course:

- individual decisions did not affect the well being of others
- any dependence would just hinge from equilibrium prices

The next lectures discuss **complete information strategic form games**.

In such environments:

- Individuals know the environment
- All decisions take place at once
- Payoffs are interdependent

Static Complete Information Games

Complete Information Games

A complete information game G consists of:

- A **set of players**:
  - N of size n
- An **action set** for each player in the game:
  - A_i for player i's
  - An action profile $\mathbf{a} = (a_1, a_2, ..., a_n)$ picks an action for each player
- A **utility function** for each player mapping action profiles to payoffs:
  - $u_i(\mathbf{a})$ denotes player i's payoff of action profile $\mathbf{a}$

$B \backslash G$	s	m
s	5,2	1,2
m	0,0	3,5

Representing Simoultaneous Move Complete Info Games

Strategic Form

$1 \backslash 2$	s	m
s	5,2	1,2
m	0,0	3,5

Extensive Form

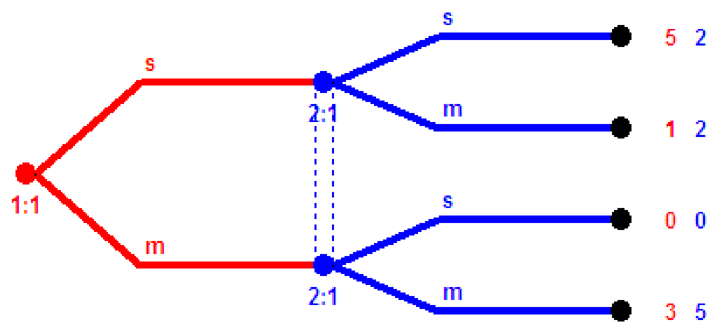

Representing Simoultaneous Move Complete Info Games

Strategic Form

1\2	s	m
s	5,2	1,2
m	0,0	3,5

Extensive Form

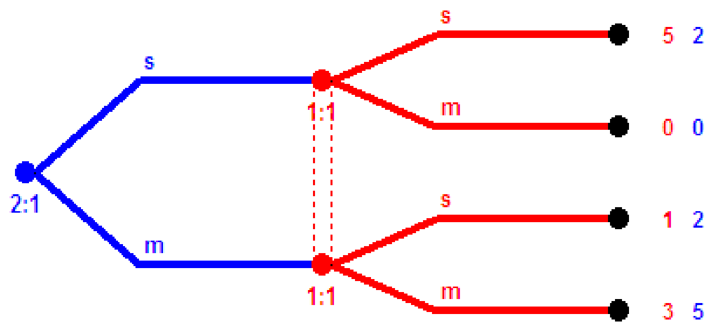

Feasible Payoffs and Efficiency

Strategic Form

1\2	s	m
s	5,2	1,2
m	0,0	3,5

Feasible Payoffs

Efficient Payoffs

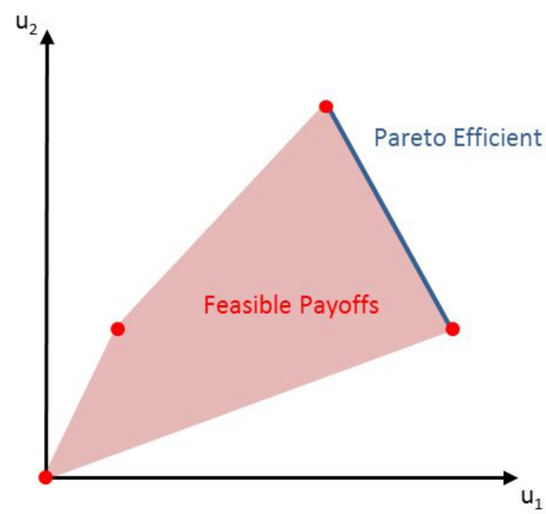

Strategies and Best Responses

Information and Pure Strategies

A strategy in a game:

- is a map from **information into actions**
- it defines a plan of action for a player

In a complete information strategic form game:

- players have no private information
- players act simultaneously

In this context **a strategy is an element of the set of actions**

For instance a (pure) strategy for player i is simply $a_i \in A_i$

Best Responses

Define a profile of actions chosen by all players other than i by $\mathbf{a}_{-i}$:

$$\mathbf{a}_{-i} = (a_1, ..., a_{i-1}, a_{i+1}, ..., a_n)$$

A strategy a_i is a **best response** to $\mathbf{a}_{-i}$ if and only if:

$$u_i(a_i, \mathbf{a}_{-i}) \geq u_i(a_i', \mathbf{a}_{-i}) \text{ for any } a_i' \in A_i$$

The **best response correspondence** of player i is defined by:

$$b_i(\mathbf{a}_{-i}) = \arg\max_{a_i \in A_i} u_i(a_i, \mathbf{a}_{-i}) \text{ for any } \mathbf{a}_{-i}$$

BR identifies the optimal action for a player given choices made by others.

$B \backslash G$	s	m
s	5,2	1,2
m	0,0	3,5

Dominance and DSE

Strict Dominance

- Strategy a_i **strictly dominates** a_i' if:

$$u_i(a_i, \mathbf{a}_{-i}) > u_i(a_i', \mathbf{a}_{-i}) \quad \text{for any } \mathbf{a}_{-i}$$

- a_i is **strictly dominant** if it strictly dominates any other a_i'
- a_i is **strictly undominated** if no strategy strictly dominates a_i
- a_i is **strictly dominated** if a strategy strictly dominates a_i

In the following example s is strictly dominant for B:

$B \backslash G$	s	m
s	5,-	2,-
m	0,-	1,-

Weak Dominance

- Strategy a_i **weakly dominates** a_i' if:

$$\begin{aligned} u_i(a_i, \mathbf{a}_{-i}) &\geq u_i(a_i', \mathbf{a}_{-i}) \quad \text{for any } \mathbf{a}_{-i} \\ u_i(a_i, \mathbf{a}_{-i}) &> u_i(a_i', \mathbf{a}_{-i}) \quad \text{for some } \mathbf{a}_{-i} \end{aligned}$$

- a_i is **weakly dominant** if it weakly dominates any other a_i'
- a_i is **weakly undominated** if no strategy weakly dominates a_i
- a_i is **weakly dominated** if a strategy weakly dominates a_i

In the following example s is weakly dominant for B:

$B \backslash G$	s	m
s	5,-	2,-
m	0,-	2,-

Dominance Examples

One strictly and one weakly dominated strategy:

1\2	L	C	R
T	-,1	-,2	-,1
B	-,0	-,1	-,3

One strictly dominant strategy:

1\2	L	C	R
T	-,1	-,2	-,3
B	-,0	-,1	-,2

One weakly dominant strategy:

1\2	L	C	R
T	-,1	-,2	-,2
B	-,0	-,0	-,1

Dominant Strategy Equilibrium

Definitions (Dominant Strategy Equilibrium DSE)

A strict **Dominant Strategy Equilibrium** of a game G consists of a strategy profile $\mathbf{a}$ such that for any $\mathbf{a}'_{-i}$ and $i \in N$:

$$u_i(a_i, \mathbf{a}'_{-i}) > u_i(a'_i, \mathbf{a}'_{-i}) \text{ for any } a'_i \in A_i$$

- For weak DSE change $>$ with $\geq$ ...
- A profile **a** is a DSE **iff** a_i is dominant for every player i.
- Formally, a profile **a** is a DSE **iff** $a_i \in b_i(\mathbf{a}'_{-i})$ for any $\mathbf{a}'_{-i}$ and $i \in N$.
- Example (Prisoner's Dilemma):

$B \backslash S$	N	C
N	5,5	0,6
C	6,0	1,1

Iterative Elimination of Dominated Strategies

To find dominant strategies eliminate dominated strategies from the game

If necessary repeat the process to possibly rule out more strategies

Consider the following example:

1\2	L	C	R
T	1,0	2,1	3,0
M	2,3	3,2	2,1
D	0,2	1,2	2,5

$\Rightarrow$

1\2	L	C	R
T	1,0	2,1	3,0
M	2,3	3,2	2,1
D	0,2	1,2	2,5

At the first instance only D is dominated for player 1

No strategy is dominated a priori for player 2

[Strategies in green in the table are dominated and thus eliminated]

Iterative Elimination of Dominated Strategies

Once D has been eliminated from the game:

Strategy R is dominated for player 2

No strategy is dominated for player 1

1\2	L	C	R
T	1,0	2,1	3,0
M	2,3	3,2	2,1
D	0,2	1,2	2,5

$\Rightarrow$

1\2	L	C	R
T	1,0	2,1	3,0
M	2,3	3,2	2,1
D	0,2	1,2	2,5

Once R has been eliminated from the game:

Strategy T is dominated for player 1

A final iteration yields (M, L) as the only surviving strategies

Dominance: Final Considerations

Dominance is often considered a benchmark of rationality:

- **Rational players never choose dominated strategies**
- **Common knowledge of rationality** means:
  players only employ strategies that survive iterative elimination

Dominance is a simple concept buy with important limitations:

- Often there is no dominant strategy even after iteration
- It often leads to inefficient outcomes

Thus a weaker notion of equilibrium needs to be introduced to model behavior especially for richer setups

Pure Strategy Nash Equilibrium

Nash Equilibrium: Introduction

Dominance was the appropriate solution concept if players had no information or beliefs about choices made by others

The weaker notion of equilibrium that will be introduced presumes that:

- players have **correct beliefs about choices made by others**
- players **choices are optimal given such beliefs**
- the environment is common knowledge among players

Such model allows for tighter predictions when dominance has no bite

Nash Equilibrium

Definition (Nash Equilibrium NE)

A (pure strategy) **Nash Equilibrium** of a game G consists of a strategy profile $\mathbf{a} = (a_i, \mathbf{a}_{-i})$ such that for any $i \in N$:

$$u_i(\mathbf{a}) \geq u_i(a_i', \mathbf{a}_{-i}) \quad \text{for any } a_i' \in A_i$$

- A profile $\mathbf{a}$ is a NE **iff** a_i is a best response to $\mathbf{a}_{-i}$ for any player i.
- Formally, a profile $\mathbf{a}$ is a NE **iff** $a_i \in b_i(\mathbf{a}_{-i})$ for any $i \in N$.

Properties:

- Strategy profiles are independent
- Strategy profiles common knowledge

More Examples

Games may have more NE's (Battle of the Sexes):

$B \backslash G$	s	m
s	5,2	1,2
m	0,0	3,5

Nash equilibria may not be efficient (Prisoner's Dilemma):

$B \backslash S$	N	C
N	5,5	0,6
C	6,0	1,1

Pure strategy Nash equilibria may not exist (Matching Pennies):

$B \backslash G$	H	T
H	0,2	2,0
T	2,0	0,2

Examples, Properties and Limitations

Games may have more NE's (Battle of the Sexes):

$B \backslash G$	s	m
s	5,2	1,2
m	0,0	3,5

Nash equilibria may not be efficient (Prisoner's Dilemma):

$B \backslash S$	N	C
N	5,5	0,6
C	6,0	1,1

Pure strategy Nash equilibria may not exist (Matching Pennies):

$B \backslash G$	H	T
H	0,2	2,0
T	2,0	0,2

Some Nash Equilibria are More Risky

Consider the Stag Hunt game:

$B \backslash G$	Stag	Hare
Stag	9,9	0,8
Hare	8,0	8,8

Best responses for this game are:

$B \backslash G$	Stag	Hare
Stag	9,9	0,8
Hare	8,0	8,8

Both players choosing to go for the stag is NE

Such NE involves greater risks of miscoordination than the NE in which both go for the hare

Three Player Example

- A game with more than 2 players:

3	L		R	
1\2	A	B	A	B
T	1,0,1	1,0,0	0,1,1	0,1,1
D	0,1,1	1,2,0	1,0,1	2,1,1

Three Player Example

- A game with more than 2 players:

3	L		R	
1\2	A	B	A	B
T	1,0,1	1,0,0	0,1,1	0,1,1
D	0,1,1	1,2,0	1,0,1	2,1,1

- To find all PNE check best reply maps:

3	L		R	
1\2	A	B	A	B
T	1,0,1	1,0,0	0,1,1	0,1,1
D	0,1,1	1,2,0	1,0,1	2,1,1

War of Attrition Example

Consider a game with two competitors involved in a fight:

- The set of players is $N = \{1, 2\}$.
- Competitors choose how much effort to put in a fight $A_i = [0, \infty)$.
- The value of winning the fight is for competitor $i \in N$ is v_i.
- The highest effort wins the fight and ties are broken at random.
- For each competitor the cost of fighting is simply $\min\{a_i, a_j\}$.
- The payoff of competitor i given their effort levels thus satisfy:

$$u_i(a_i, a_j) = \begin{cases} v_i - a_j & \text{if} \quad a_i > a_j \\ v_i/2 - a_j & \text{if} \quad a_i = a_j \\ -a_i & \text{if} \quad a_i < a_j \end{cases}$$

War of Attrition Example

As payoffs amount to:

$$u_i(a_i, a_j) = \begin{cases} v_i - a_j & \text{if} \quad a_i > a_j \\ v_i/2 - a_j & \text{if} \quad a_i = a_j \\ -a_i & \text{if} \quad a_i < a_j \end{cases}$$

Best response functions satisfy:

$$b_i(a_j) = \begin{cases} a_i > a_j & \text{if} \quad a_j < v_i \\ a_i = 0 \quad \text{or} \quad a_i > a_j & \text{if} \quad a_j = v_i \\ a_i = 0 & \text{if} \quad a_j > v_i \end{cases}$$

All Nash Equilibria of the game satisfy one of the following:

- $a_1 = 0$ and $a_2 \geq v_1$
- $a_2 = 0$ and $a_1 \geq v_2$

War of Attrition Example

The esiest way to find the NE in such games is plotting BRs:

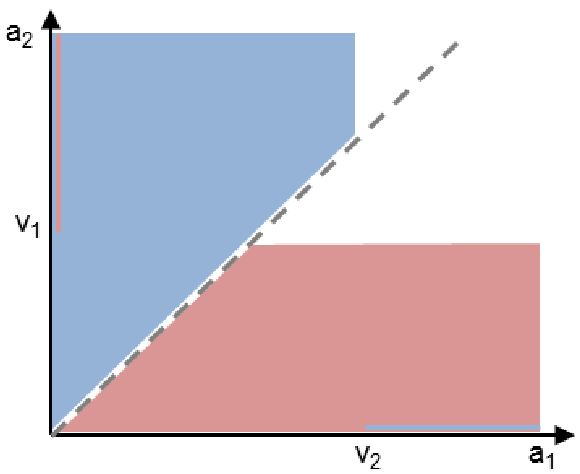

All Nash Equilibria of the game satisfy one of the following:

- $a_1 = 0$ and $a_2 \geq v_1$
- $a_2 = 0$ and $a_1 \geq v_2$

DSE implies NE

Fact

Any dominant strategy equilibrium is a Nash equilibrium

Proof.

If $\mathbf{a}$ is a DSE then $a_i \in b_i(\mathbf{a}'_{-i})$ for any $\mathbf{a}'_{-i}$ and $i \in N$.
Which implies $\mathbf{a}$ is NE since $a_i \in b_i(\mathbf{a}_{-i})$ for any $i \in N$. □

Mixed Strategies and Concealing

Introduction to Mixed Strategies

A problematic aspect of the solution concepts discussed in pure strategies was that equilibria did not always exist.

Technical reasons for the lack of existence were:

- Non-convexities in the choice sets;
- Discontinuities of the best response correspondences;

Intuitively the impossibility of concealing decisions was causing problems.

Such problems can be solved by introducing mixed strategies which guarantee the existence of at least one Nash equilibrium.

Mixed Strategy Definition

Consider complete information static game $\{N, \{A_i, u_i\}_{i \in N}\}$.

A mixed strategy for $i \in N$ is a probability distribution over actions in A_i.

Let $\sigma_i(a_i)$ denote the probability that player i chooses to play a_i.

Thus σ_i is a **mixed strategy** if:

- $\sigma_i(a_i) \geq 0$ for any $a_i \in A_i$;
- $\sum_{a_i \in A_i} \sigma_i(a_i) = 1$.

For instance $\sigma_1(B) = 0.3$ and $\sigma_1(C) = 0.7$ is a mixed strategy for 1 in:

1\2	B	C
B	2,0	0,2
C	0,1	1,0

Payoffs from Mixed Strategies

As in the last lectures often we denote:

$$\boldsymbol{\sigma}_{-i} = (\sigma_1, ..., \sigma_{i-1}, \sigma_{i+1}, ..., \sigma_N)$$

The **payoff to player** i from choosing σ_i when others follow σ_{-i} is:

$$\begin{aligned} u_i(\sigma_i, \boldsymbol{\sigma}_{-i}) &= \sum_{\mathbf{a}\in\mathbf{A}} \Pr(\mathbf{a}) u_i(\mathbf{a}) = \\ &= \sum_{\mathbf{a}\in\mathbf{A}} \prod_{j\in N} \sigma_j(a_j) u_i(\mathbf{a}) \end{aligned}$$

E.G. If players follow $\sigma_1(B) = \sigma_2(B) = 0.3$ in the game:

1\2	B	C
B	2,0	0,2
C	0,1	1,0

The payoff to player 1 is: $u_1(\sigma_1, \sigma_2) = (.09)2 + (.49)1 + (.42)0 = 0.67$

Best Responses

Denote the best response correspondence of i by $b_i(\boldsymbol{\sigma}_{-i})$

The map is defined by:

$$b_i(\boldsymbol{\sigma}_{-i}) = \arg\max_{\sigma_i} u_i(\sigma_i, \boldsymbol{\sigma}_{-i})$$

For instance consider the game:

1\2	s	m
s	5,2	1,2
m	0,0	3,5

If $\sigma_1(s) = 1$ then any $\sigma_2(s) \in [0, 1]$ satisfies $\sigma_2 \in b_2(\sigma_1)$

If $\sigma_1(s) < 1$ then only $\sigma_2(s) = 0$ satisfies $\sigma_2 \in b_2(\sigma_1)$

Dominated Strategies

- Strategy σ_i **strictly dominates** a_i if:

$$u_i(\sigma_i, \mathbf{a}_{-i}) > u_i(a_i, \mathbf{a}_{-i}) \quad \text{for any } \mathbf{a}_{-i}$$

- a_i is **strictly undominated** if no strategy strictly dominates it
- This allows us to rule out more strategies than before, eg:

1\2	L	C	R
T	6,6	0,2	0,0
B	0,0	0,2	6,6

- $\sigma_2(L) = \sigma_2(R) = 0.5$ strictly dominates C since:

$$u_2(\sigma_2, a_1) = 3 > u_2(C, a_1) = 2$$

Nash Equilibrium

Definition (Nash Equilibrium NE)

A **Nash Equilibrium** of a game consists of a strategy profile $\boldsymbol{\sigma} = (\sigma_i, \boldsymbol{\sigma}_{-i})$ such that for any $i \in N$:

$$u_i(\boldsymbol{\sigma}) \geq u_i(a_i, \boldsymbol{\sigma}_{-i}) \quad \text{for any } a_i \in A_i$$

Implicit to the definition of NE are the following assumptions:

- Each agent chooses his mixed strategy independently of others
- Each agent knows and believes which strategies the others adopt
- Each agent maximizes expected utility given his beliefs about others

Nash Equilibrium Computation Help

- A strategy profile σ is a Nash Equilibrium if and only if:

$$\begin{aligned} u_i(\sigma) &= u_i(a_i, \sigma_{-i}) \text{ for any } a_i \text{ such that } \sigma_i(a_i) > 0 \\ u_i(\sigma) &\geq u_i(a_i, \sigma_{-i}) \text{ for any } a_i \text{ such that } \sigma_i(a_i) = 0 \end{aligned}$$

- Intuitively, a player is indifferent between the actions he plays and prefers them to any other action.
- If a_i is strictly dominated, then $\sigma_i(a_i) = 0$ in any NE.
- If a_i is weakly dominated, then $\sigma_i(a_i) > 0$ in some NE then any profile of actions $\mathbf{a}_{-i}$ for which a_i is strictly worse occurs with 0 probability.

Examples

Games may have more NEs (Battle of the Sexes):

1\2	s	m
s	5,2	1,1
m	0,0	2,5

There are 2 PNE & a mixed NE in which $\sigma_1(s) = 5/6$ & $\sigma_2(s) = 1/6$:

$$u_1(s, \sigma_2) = 5\sigma_2(s) + (1 - \sigma_2(s)) = 2(1 - \sigma_2(s)) = u_1(m, \sigma_2)$$
$$u_2(m, \sigma_1) = \sigma_1(s) + 5(1 - \sigma_1(s)) = 2\sigma_1(s) = u_2(s, \sigma_1)$$

Games may have only mixed NE (Matching Pennies):

$B \backslash G$	H	T
H	0,2	2,0
T	2,0	0,2

There is a unique NE in which $\sigma_1(H) = \sigma_2(H) = 1/2$:

$$2(1 - \sigma_i(H)) = 2\sigma_i(H)$$

Nash Equilibrium Existence

Nash Equilibrium Existence

Theorem (NE Existence)

Any game with a finite number of actions possesses a Nash equilibrium.

Theorem (PNE Existence)

Any game with convex and compact action sets and with continuous and quasi-concave payoff functions possesses a pure strategy Nash equilibrium.

Assumptions in both theorems guarantee that best response maps are "continuous" on convex compact sets and thus existence...

Continuous Best Responses Imply Existence

A game without PNE and with a single NE:

$1\backslash 2$	B	C
B	2,0	0,2
C	0,1	1,0

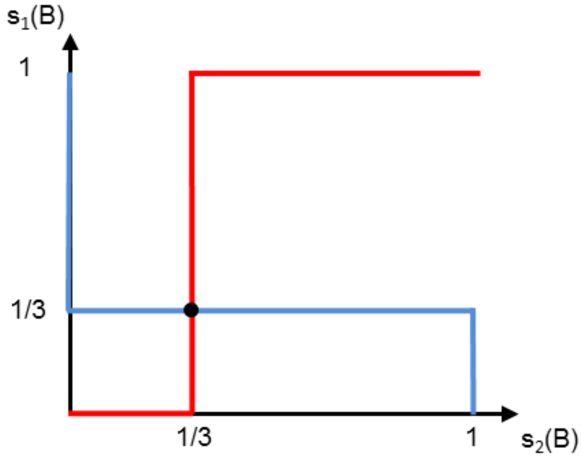

Oligopoly Pricing

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

The models of competition presented earlier in the course have explored the consequences on prices and trade of two extreme assumptions:

- Perfect Competition [Many sellers supplying many buyers]
- Monopoly [One seller supplying many buyers]

Today intermediate assumption is discussed:

- Oligopoly [Few sellers supplying many buyers]

Two models of competition among oligopolists are presented:

- **Quantity Competition** [aka Cournot Competition]
- **Price Competition** [aka Bertrand Competition]

Quantity Competition

A Duopoly

Consider the following economy:

- There are two firms $N = \{1, 2\}$
- Each firm $i \in N$ produces output q_i with a **cost function**
  - $c_i(q_i)$ mapping quantities to costs
- Aggregate output in this economy is $q = q_1 + q_2$
- Both firms face an aggregate **inverse demand** for output
  - $p(q)$ mapping aggregate output to prices
- The payoff of each firm $i \in N$ is its profits:

$$u_i(q_1, q_2) = p(q)q_i - c_i(q_i)$$

Profits depend on the output decisions of both

Cournot Competition: Duopoly

Competition proceeds as follows:

- All **firms** simultaneously **select their output** to maximize profits
- Each firm takes as given the output of its competitors
- Firms account for the effects of their output decision on prices

In particular the decision problem of player $i \in N$ is to:

$$\max_{q_i} u_i(q_i, q_j) = \max_{q_i} p(q_i + q_j)q_i - c_i(q_i)$$

[Historically this is the first known example of Nash Equilibrium – 1838]

[Example: Visa vs Mastercard]

Cournot Competition: Equilibrium

If standard conditions on primitives of the problem hold:

- a pure strategy Nash equilibrium exists
- the PNE is characterized by the FOC

If so, the problem of any producer $i \in N$ satisfies:

$$\frac{\partial u_i(q_i, q_j)}{\partial q_i} = \underbrace{p(q_i + q_j) + \frac{\partial p(q_i + q_j)}{\partial q_i} q_i}_{\text{Marginal Revenue}} - \underbrace{\frac{\partial c_i(q_i)}{\partial q_i}}_{\text{Marginal Cost}} \begin{cases} \leq 0 & \text{if} \quad q_i = 0 \\ = 0 & \text{if} \quad q_i > 0 \end{cases}$$

Marginal revenue accounts for the distortion in prices

Prices decrease if inverse demand is downward-sloping

FOC defines the best response (aka *reaction function*) of player i:

$$q_i = b_i(q_j)$$

Cournot Competition: Plot

The Cournot FOC of any produce graphically satisfies:

$$\underbrace{p(q_i + q_j) + \frac{\partial p(q_i + q_j)}{\partial q_i} q_i}_{\text{Marginal Revenue}} = \underbrace{\frac{\partial c_i(q_i)}{\partial q_i}}_{\text{Marginal Cost}}$$

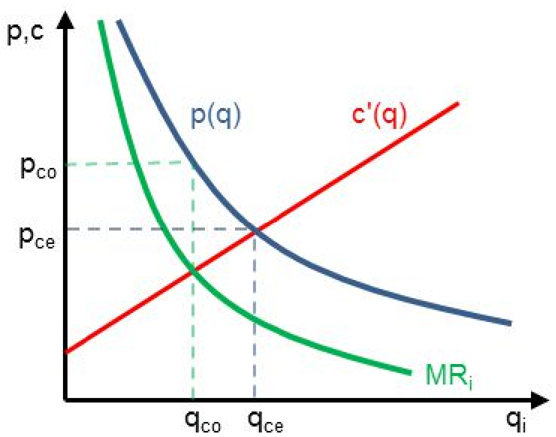

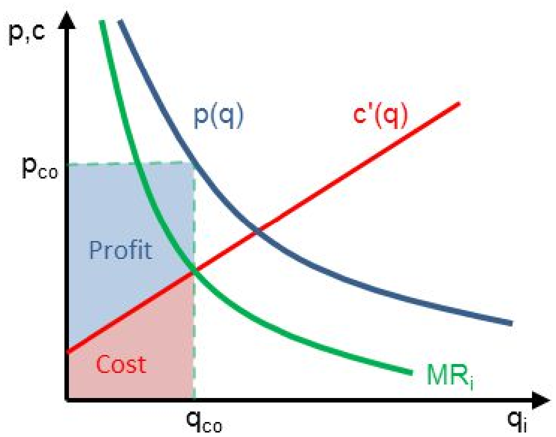

Cournot Competition: Example

Consider the following economy:

- $p(q) = 2 - q$
- $c_1(q_1) = q_1^2$ and $c_2(q_2) = 3q_2^2$

Firm i's problem is to choose production q_i given choice of the other q_j:

$$\max_{q_i}(2 - q_i - q_j)q_i - c_i(q_i)$$

The best reply map of each firm is determined by FOC:

$$\begin{aligned} 2 - 2q_1 - q_2 - 2q_1 &= 0 \Rightarrow q_1 = b_1(q_2) = (2 - q_2)/4 \\ 2 - 2q_2 - q_1 - 6q_2 &= 0 \Rightarrow q_2 = b_2(q_1) = (2 - q_1)/8 \end{aligned}$$

Cournot Equilibrium outputs are: $q_1 = 14/31$ and $q_2 = 6/31$

Perfect competition outputs are larger: $q_1^* = 3/5$ and $q_2^* = 1/5$

Collusion and Cartels

Collusion and Cartels

Suppose that the producers collude by forming a *cartel*

A **cartel maximizes the joint profits** of the two firms:

$$\max_{q_1,q_2} p(q)q - c_1(q_1) - c_2(q_2)$$

First order optimality of this problem requires for any $i \in N$:

$$\underbrace{p(q) + \frac{\partial p(q)}{\partial q} q}_{\text{Marginal Revenue}} - \underbrace{\frac{\partial c_i(q_i)}{\partial q_i}}_{\text{Marginal Cost}} \begin{cases} \leq 0 & \text{if} \quad q_i = 0 \\ = 0 & \text{if} \quad q_i > 0 \end{cases}$$

Aggregate profits are higher in the cartel

Players account for effects of their output choice on others

But the profits of each individual do not necessarily increase

Collusion: Plot

The Collusion FOC of any producer graphically satisfies:

$$\underbrace{p(q) + \frac{\partial p(q)}{\partial q} q}_{\text{Marginal Revenue}} = \underbrace{\frac{\partial c_i(q_i)}{\partial q_i}}_{\text{Marginal Cost}}$$

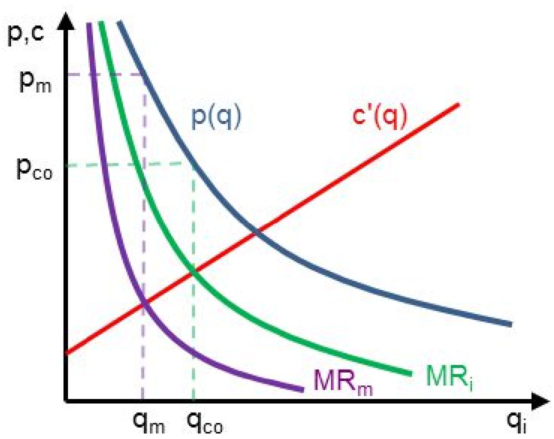

Collusion: Example

Consider the previous duopoly, but suppose that a cartel is in place

If so, FOC for the cartel production satisfy:

$$\begin{aligned} 2-2\overline{q}_1-2\overline{q}_2-2\overline{q}_1 &= 0 \Rightarrow \overline{q}_1=(1-\overline{q}_2)/2 \\ 2-2\overline{q}_2-2\overline{q}_1-6\overline{q}_2 &= 0 \Rightarrow \overline{q}_2=(1-\overline{q}_1)/4 \end{aligned}$$

Cartel outputs are: $\overline{q}_1 = 3/7$ and $\overline{q}_2 = 1/7$

Cournot outputs are larger: $q_1 = 14/31$ and $q_2 = 6/31$

Cartel profits are: $\overline{u}_1 = 3/7$ and $\overline{u}_2 = 1/7$

Cournot profits are: $u_1 = 392/961$ and $u_2 = 144/961$

Total profits are larger with a cartel in place, but not all firms may benefit

Collusion: Incentives to Defect

Suppose that the firm j produces the cartel output $\overline{q}_j$

If so, firm i may benefit by producing more than the cartel output since:

$$b_i(\overline{q}_j) > \overline{q}_i$$

In this scenario sustaining a cartel may be hard without output monitoring

In the example this was the case as firms preferred to increase output:

$$\begin{aligned} b_1(1/7) &= 13/28 > 3/7 \\ b_2(3/7) &= 11/56 > 1/7 \end{aligned}$$

If so the problem of sustaining the cartel becomes a Prisoner's dilemma

Price Competition

Bertrand Competition: Duopoly

Competition proceeds as follows:

- All firms simultaneously **quote a price** to maximize profits
- Each firm takes as given the price quoted by its competitors
- Firms account for the effects of their pricing decision on sales

Consider an economy with:

- Two producers with **constant marginal costs** c
- **Aggregate demand** for output given by $q(p) = (b_0 - p)/b$

Given the prices **demand** of output from firm $i \in N$ is:

$$q_i(p_i, p_j) = \begin{cases} q(p_i) & \text{if} \quad p_i < p_j \\ q(p_i)/2 & \text{if} \quad p_i = p_j \\ 0 & \text{if} \quad p_i > p_j \end{cases}$$

Bertrand Competition: Monopoly

If only one firm operated in such market it would choose the price to:

$$\max_p u(p) = \max_p q(p)(p-c)$$

Thus a profit maximizing monopolist would sell goods at a price:

$$\overline{p} = (b_0 + c)/2$$

With two producers the problem of each firm becomes:

$$\max_{p_i} u_i(p_i, p_j) = \max_{p_i} q_i(p_i, p_j)(p_i - c)$$

Bertrand Competition: Best Responses I

In the Bertrand model, i's optimal pricing is aimed at "maximizing sales"

In particular firm i would set prices as follows (for ε small):

- If $p_j > \overline{p}$, set $p_i = \overline{p}$ and capture all the market at the monopoly price
- If $\overline{p} \geq p_j > c$, set $p_i = p_j - \varepsilon$, undercut j and capture all the market
- If $c \geq p_j$, set $p_i = c$ as there are no benefits by pricing below MC

Such logic requires the best response of each player to satisfy:

$$p_i = b_i(p_j) = \begin{cases} \overline{p} & \text{if} \quad p_j \in (\overline{p}, \infty) \\ p_j - \varepsilon & \text{if} \quad p_j \in (c, \overline{p}] \\ \geq c & \text{if} \quad p_j = c \\ > p_j & \text{if} \quad p_j \in [0, c) \end{cases}$$

Thus in the unique NE both firms set $p_1 = p_2 = c$ and perfect competition emerges with just 2 firms!

Bertrand Competition: Best Responses II

To find PNE plot the best responses for $c = 1$ and $\overline{p} = 2$:

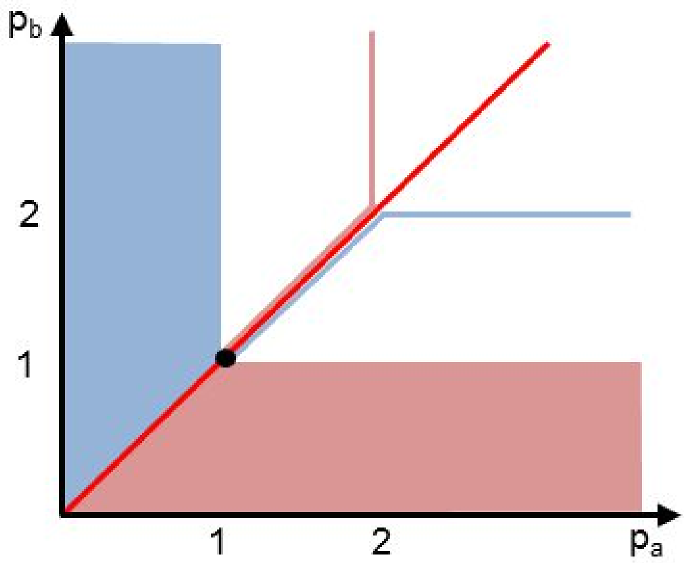

In blue the best response of player b, in pink that of player a.

Thus in the unique NE both firms set $p_a = p_b = 1$.

Capacity and Market Structure

Cournot vs Bertrand Competition

Bertrand model predicts that duopoly is enough to push down prices to marginal cost (as in perfect competition)

Cournot model instead predicts that few producers do not suffice to eliminate markups (prices above marginal cost)

In both models there are incentives to form a cartel and to charge the monopoly price

Neither model is intrinsically better

Accuracy of either model depends on the fundamentals of the economy:

- Bertrand works better when capacity is easy to adjust
- Cournot works better when capacity is hard to adjust

Games of Incomplete Information

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Games of Incomplete Information:

- Definitions:
  - Incomplete Information Game
  - Information Structure and Beliefs
  - Strategies
  - Best Reply Map
- Solution Concepts in Pure Strategies:
  - Dominant Strategy Equilibrium
  - Bayes Nash Equilibrium
- Examples

Static Incomplete Information Games

Incomplete Information Games

An **incomplete information game** consists of:

- N the set of players in the game
  - For convenience we consider only two player games
- A_i player i's action set
- $\mathbb{X}_i$ player i's set of possible signals
  - A profile of signals $x = (x_1, x_2)$ is an element $\mathbb{X} = \mathbb{X}_1 \times \mathbb{X}_2$
- f a distribution over the possible signals
- $u_i : A \times \mathbb{X} \to \mathbb{R}$ player i's utility function, $u_i(a|x)$

Bayesian Game Example

Consider the following Bayesian game:

- Player 1 observes only one possible signal: $\mathbb{X}_1 = \{K\}$
- Player 2's signal takes one of two values: $\mathbb{X}_2 = \{L, R\}$
- Probabilities are such that: $f(K, L) = 0.6$
- Payoffs and action sets are as described in the matrix:

$1\backslash 2.L$	C	D
A	1,2	0,1
B	0,4	1,3

$1\backslash 2.R$	C	D
A	1,3	0,4
B	0,1	1,2

Information Structure

Information structure:

- X_i denotes the signal as a random variable
- belongs to the set of possible signals $\mathbb{X}_i$
- x_i denotes the realization of the random variable X_i
- X_{-i} denotes a profile the signal of player $j \neq i$
- Player i observes only X_i
- Player i ignores X_{-i}, but knows f

With such information player i forms beliefs regarding the realization of the signals of the other players x_{-i}

Beliefs about other Players' Signals [Easy]

In this course we consider models in which signals are independent:

$$f(x) = \Pr(X_1 = x_1 \ \& \ X_2 = x_2) = f_1(x_1)f_2(x_2)$$

This implies that the signal x_i of player i is independent of X_{-i}.

Beliefs are a probability distribution over the signals of the other players.

Players form beliefs about signals received by others using Bayes Rule.

Independence implies that conditional observing $X_i = x_i$ the beliefs of player i are:

$$\begin{aligned} f(x_{-i}|x_i) &= \Pr(X_{-i} = x_{-i}|X_i = x_i) = \frac{\Pr(X_{-i} = x_{-i} \ \& \ X_i = x_i)}{\Pr(X_i = x_i)} = \\ &= \Pr(X_{-i} = x_{-i}) = f_{-i}(x_{-i}) \end{aligned}$$

Extra: Beliefs about other Players' Signals [Hard]

Also in the general case with interdependence players form beliefs about the signals received by the others by using Bayes Rule.

Conditional observing $X_i = x_i$ the beliefs of player i are:

$$
\begin{aligned}
f_i(x_{-i}|x_i) &= \Pr(X_{-i} = x_{-i}|X_i = x_i) = \\
&= \frac{\Pr(X_{-i} = x_{-i} \;\&\; X_i = x_i)}{\Pr(X_i = x_i)} = \\
&= \frac{\Pr(X_{-i} = x_{-i} \;\&\; X_i = x_i)}{\sum_{y_{-i} \in \mathbb{X}_{-i}} \Pr(X_{-i} = y_{-i} \;\&\; X_i = x_i)} = \\
&= \frac{f(x_{-i}, x_i)}{\sum_{y_{-i} \in \mathbb{X}_{-i}} f(y_{-i}, x_i)}
\end{aligned}
$$

Beliefs are a probability distribution over the signals of the other players.

Strategies and Best Responses

Strategies

Strategy Profiles:

- A strategy consists of a map from available information to actions:

$$\alpha_i : \mathbb{X}_i \to A_i$$

- A strategy profile consists of a strategy for every player:

$$\alpha(X) = (\alpha_1(X_1), \alpha_2(X_2))$$

- We adopt the usual convention:

$$\alpha_{-i}(X_{-i}) = \alpha_j(X_j) \quad \text{for } j \neq i \in \{1, 2\}$$

Bayesian Game Example Continued

Consider the following game:

- Player 1 observes only one possible signal: $\mathbb{X}_1 = \{K\}$
- Player 2's signal takes one of two values: $\mathbb{X}_2 = \{L, R\}$
- Probabilities are such that: $f(K, L) = 0.6$
- Payoffs and action sets are as described in the matrix:

$1 \backslash 2.L$	C	D
A	1,2	0,1
B	0,4	1,3

$1 \backslash 2.R$	C	D
A	1,3	0,4
B	0,1	1,2

- A strategy for player 1 is an element of the set $\alpha_1 \in \{A, B\}$
- A strategy for player 2 is a map $\alpha_2 : \{L, R\} \to \{C, D\}$
- Player 1 cannot act upon 2's private information

Dominant Strategy Equilibrium

- Strategy α_i **strictly dominates** α'_i if for any a_{-i} and $x \in \mathbb{X}$:

$$u_i(\alpha_i(x_i), a_{-i}|x) > u_i(\alpha'_i(x_i), a_{-i}|x)$$

- Strategy α_i is **dominant** if it dominates any other strategy α'_i.
- Strategy α_i is **undominated** if no strategy dominates it.

Definitions (Dominant Strategy Equilibrium DSE)

A **Dominant Strategy Equilibrium** of an incomplete information game is a strategy profile α that for any $i \in N$, $x \in \mathbb{X}$ and $a_{-i} \in A_{-i}$ satisfies:

$$u_i(\alpha_i(x_i), a_{-i}|x) > u_i(\alpha'_i(x_i), a_{-i}|x) \quad \text{for any } \alpha'_i : \mathbb{X}_i \to A_i$$

- That is, α_i is optimal independently of what others know and do.

Interim Expected Utility and Best Reply Maps

The **interim stage** occurs just after a player knows his signal $X_i = x_i$.

It is when strategies are chosen in a Bayesian game

The **interim expected** utility of a (pure) strategy profile α is defined by:

$$U_i(\alpha|x_i) = \sum_{\mathbb{X}_{-i}} u_i(\alpha(x)|x) f(x_{-i}|x_i)$$

With such notation in mind notice that:

$$U_i(a_i, \alpha_{-i}|x_i) = \sum_{\mathbb{X}_{-i}} u_i(a_i, \alpha_{-i}(x_{-i})|x) f(x_{-i}|x_i)$$

The **best response** correspondence of player i is defined by:

$$b_i(\alpha_{-i}|x_i) = \arg\max_{a_i \in A_i} U_i(a_i, \alpha_{-i}|x_i)$$

BR identifies which actions are optimal given the signal and the strategies followed by others.

Pure Strategy Bayes Nash Equilibrium

Definitions (Bayes Nash Equilibrium BNE)

A pure strategy **Bayes Nash Equilibrium** of an incomplete information game is a strategy profile α such that for any $i \in N$ and $x_i \in \mathbb{X}_i$ satisfies:

$$U_i(\alpha|x_i) \geq U_i(a_i, \alpha_{-i}|x_i) \quad \text{for any } a_i \in A_i$$

BNE requires **interim optimality**. That is, do your best given:

- the signal you have received;
- the strategy of other players.

BNE requires $\alpha_i(x_i) \in b_i(\alpha_{-i}|x_i)$ for any $x_i \in \mathbb{X}_i$ and $i \in N$.

Bayesian Game Example

Consider the following Bayesian game with $f(K, L) = p > 1/2$:

$1\backslash 2.L$	C	D
A	1,2	0,1
B	0,4	1,3

$1\backslash 2.R$	C	D
A	1,3	0,4
B	0,1	1,2

The best reply maps for both player are characterized by:

$$b_2(\alpha_1|x_2) = \begin{cases} C & \text{if} \quad x_2 = L \\ D & \text{if} \quad x_2 = R \end{cases} \qquad b_1(\alpha_2) = \begin{cases} A & \text{if} \quad \alpha_2(L) = C \\ B & \text{if} \quad \alpha_2(L) = D \end{cases}$$

The game has a unique (pure strategy) BNE in which:

$$\alpha_1 = A,\ \alpha_2(L) = C,\ \alpha_2(R) = D$$

DO NOT ANALYZE MATRICES SEPARATELY!!!

Transforming a Bayesian Game Example

Consider the following Bayesian game with $f(K,L) = p > 1/2$:

$1\backslash 2.L$	C	D
A	1,2	0,1
B	0,4	1,3

$1\backslash 2.R$	C	D
A	1,3	0,4
B	0,1	1,2

This is equivalent to the following complete information game:

$1\backslash 2$	CC	DC	CD	DD
A	$1, 3-p$	$1-p, 3-2p$	$p, 4-2p$	$0, 4-3p$
B	$0, 1+3p$	$p, 1+2p$	$1-p, 2+2p$	$1, 2+p$

You may then find BRs and BNEs in this modified table.

Extra: Relationships between Equilibrium Concepts

If α is a DSE then it is a BNE. In fact for any action a_i and signal x_i:

$$u_i(\alpha_i(x_i), a_{-i}|x) \geq u_i(a_i, a_{-i}|x) \;\; \forall a_{-i}, x_{-i} \;\; \Rightarrow$$

$$u_i(\alpha_i(x_i), \alpha_{-i}(x_{-i})|x) \geq u_i(a_i, \alpha_{-i}(x_{-i})|x) \;\; \forall \alpha_{-i}, x_{-i} \Rightarrow$$

$$\textstyle\sum_{\mathbb{X}_{-i}} u_i(\alpha(x)|x) f_i(x_{-i}|x_i) \geq \sum_{\mathbb{X}_{-i}} u_i(a_i, \alpha_{-i}(x_{-i})|x) f_i(x_{-i}|x_i) \;\; \forall \alpha_{-i} \Rightarrow$$

$$U_i(\alpha|x_i) \geq U_i(a_i, \alpha_{-i}|x_i) \;\; \forall \alpha_{-i}$$

BNE Example I: Entry Game

Consider the following market game:

- Firm I (the incumbent) is a monopolist in a market
- Firm E (the entrant) is considering whether to enter in the market
- If E stays out of the market, E runs a profit of 1$ and I gets 8$
- If E enters, E incurs a cost of 1$ and profits of both I and E are 3$
- I can deter entry by investing at cost $\{0, 2\}$ depending on type $\{L, H\}$
- If I invests: I's profit increases by 1 if he is alone, decreases by 1 if he competes and E's profit decreases to 0 if he elects to enter

$E \backslash I.L$	Invest	Not Invest
In	0,2	3,3
Out	1,9	1,8

$E \backslash I.H$	Invest	Not Invest
In	0,0	3,3
Out	1,7	1,8

BNE Example I: Exploiting the Transformation

Let π denote the probability that firm I is of type L.

- The original payoff matrix was:

$E \backslash I.L$	Invest	Not Invest
In	0,2	3,3
Out	1,9	1,8

$E \backslash I.H$	Invest	Not Invest
In	0,0	3,3
Out	1,7	1,8

- Its associated interim transformation amounts to:

$E \backslash I$	II	NI	IN	NN
In	$0, 2\pi$	$3\pi, 3\pi$	$3(1-\pi), 3-\pi$	$3, 3$
Out	$1, 7+2\pi$	$1, 7+\pi$	$1, 8+\pi$	$1, 8$

BNE Example II: Entry Game

Thus we have that:

- $\alpha_I(H) = \textit{Not Invest}$ is a strictly dominant strategy for $I.H$
- For any value of π, $\alpha_I(L) = \textit{Not Invest}$ and $\alpha_E = \textit{In}$ is BNE:

$$u_I(\textit{Not}, \textit{In}|L) = 3 > 2 = u_I(\textit{Invest}, \textit{In}|L)$$
$$U_E(\textit{In}, \alpha_I(X_I)) = 3 > 1 = U_E(\textit{Out}, \alpha_I(X_I))$$

- For π high enough, $\alpha_I(L) = \textit{Invest}$ and $\alpha_E = \textit{Out}$ is also BNE:

$$u_I(\textit{Invest}, \textit{Out}|L) = 9 > 8 = u_I(\textit{Not}, \textit{Out}|L)$$
$$U_E(\textit{Out}, \alpha_I(X_I)) = 1 > 3(1-\pi) = U_E(\textit{In}, \alpha_I(X_I))$$

$E \backslash I.L$	Invest	Not Invest
In	0,2	3,3
Out	1,9	1,8

$E \backslash I.H$	Invest	Not Invest
In	0,0	3,3
Out	1,7	1,8

BNE Example II: Exchange

A buyer and a seller want to trade an object:

- Buyer's value for the object is 3$
- Seller's value is either 0$ or 2$ based on the signal, $\mathbb{X}_S = \{L, H\}$
- Buyer can offer either 1$ or 3$ to purchase the object
- Seller choose whether or not to sell

$B \backslash S.L$	sale	no sale
3$	0,3	0,0
1$	2,1	0,0

$B \backslash S.H$	sale	no sale
3$	0,3	0,2
1$	2,1	0,2

- This game for any prior f has a BNE in which:

$$\alpha_S(L) = \text{sale}, \ \alpha_S(H) = \text{no sale}, \ \alpha_B = 1\$$$

- Selling is strictly dominant for $S.L$
- Offering 1$ is weakly dominant for the buyer

Dynamic Games

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Dynamic Games:

- Definitions:
  - Extensive Form Game
  - Information Sets and Beliefs
  - Behavioral Strategy
  - Subgame
- Solution Concepts:
  - Nash Equilibrium
  - Subgame Perfect Equilibrium
  - Perfect Bayesian Equilibrium
- Examples: Imperfect Competition

Extensive Form Games

Dynamic Games

- All games discussed in previous lectures were static. That is:
  - a set of players taking decisions simultaneously;
  - or not being able to observe the choices made by others.
- Today we relax such assumption by modeling the timing of decisions.
- In common instances the rules of the game explicitly define:
  - the order in which players move;
  - the information available to them when they take their decisions.
- A way of representing such dynamic games is in their Extensive Form.
- The following definitions are helpful to define such notion.

Basic Graph Theory

- A graph consists of a set of **nodes** and of a set of **branches**.
- Each branch connects a pair of nodes.
- A branch is identified by the two nodes it connects.
- A path is a set of branches:

$$\{\{x_k, x_{k+1}\} \,|k = 1, ..., m\}$$

where $m > 1$ and every x_k is a different node of the graph.

- A **tree** is a graph in which any two nodes are connected by exactly one path.
- A **rooted tree** is a tree in which a node designated as the **root**.
- A **terminal node** is a node connected by only one branch.

An Extensive Form Game

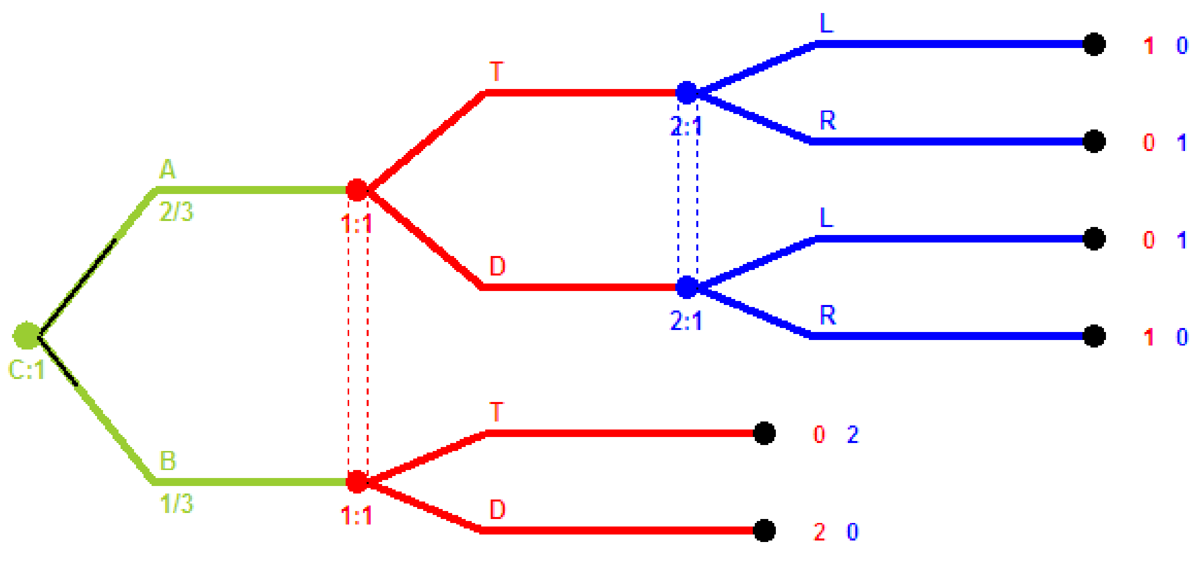

Extensive Form Games

An extensive form game is a rooted tree together with functions assigning labels to nodes and branches such that:

1. Each non-terminal node has a player-label in $\{C, 1, ..., n\}$:
   - $\{1, ..., n\}$ are the players in the game.
   - Nodes assigned to label C are **chance nodes**.
   - Nodes assigned to label $i \neq C$ are **decision nodes** controlled by i.
2. Each alternative at a chance node has a label specifying its probability:
   - **Chance probabilities** are nonnegative and add to 1.
3. Each node controlled by player $i > 0$ has a second label specifying i's **information state**:
   - Thus nodes labeled $i.s$ are controlled by i with information s.
   - Two nodes belong to $i.s$ **iff** i cannot distinguish them.

Extensive Form Games

4. Each alternative at a decision node has **move label**:
   - If two nodes x, y belong to the same information set, for any alternative at x there must be exactly one alternative at y with the same move label.
5. Each terminal node y has a label that specifies a vector of n numbers $\{u_i(y)\}_{i\in\{1,...,n\}}$ such that:
   - The number $u_i(y)$ specifies the **payoff** to i if the game ends at node y.
6. All players have **perfect recall** of the moves they chose.

Perfect Recall

With perfect recall information sets 1.2 and 1.3 cannot coincide:

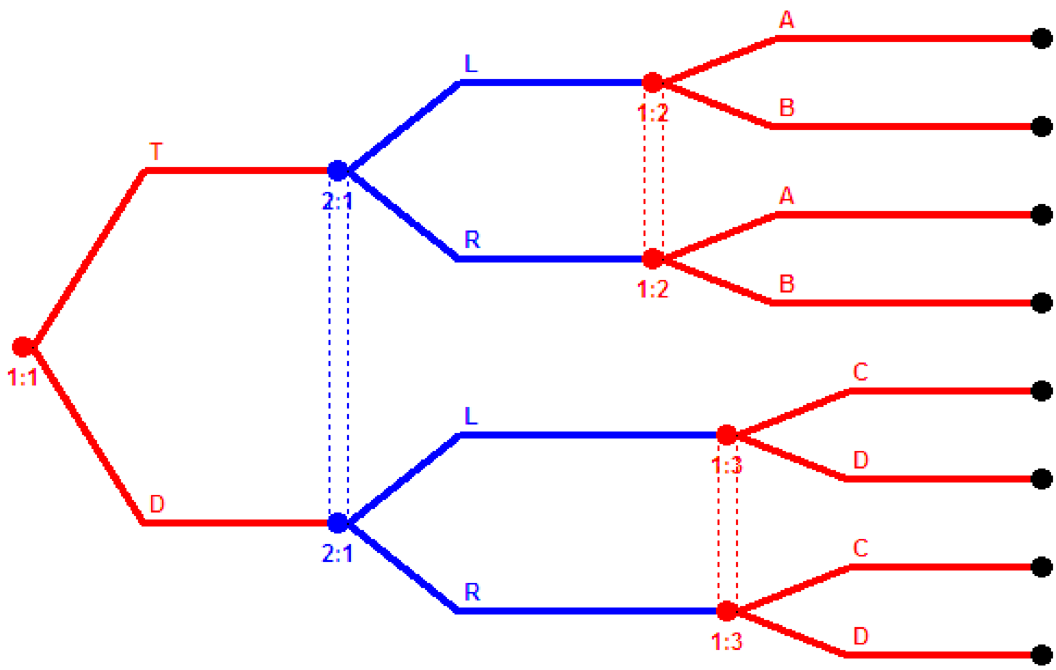

Without Perfect Recall

Without perfect recall assumption this is possible:

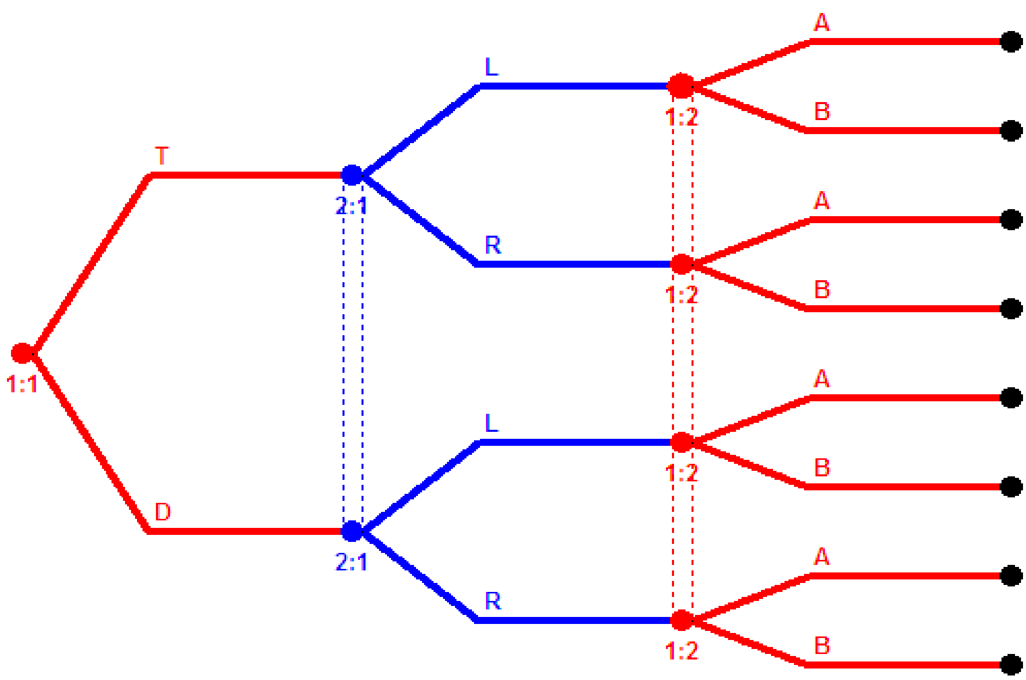

Perfect Information

An extensive form game has **perfect information** if no two nodes belong to the same information state.

With Perfect Information

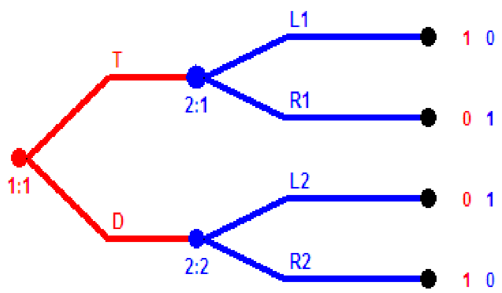

Without Perfect Information

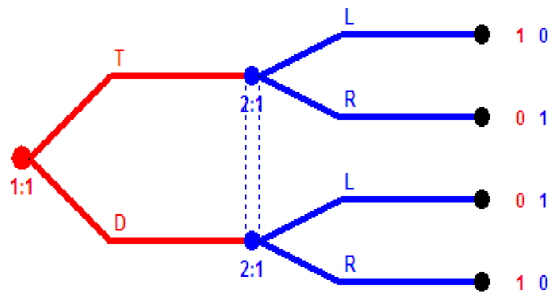

Behavioral Strategies

Throughout let:

- S_i be the set information states of player $i \in N$.
- $A_{i.s}$ be the action set of player i at info state $s \in S_i$.

A **behavioral strategy** for player i maps information states to probability distributions over actions.

In particular $\sigma_{i.s}(a_{i.s})$ is the probability that player i at information stage s chooses action $a_{i.s} \in A_{i.s}$.

Throughout denote:

- a behavioral strategy of player i by $\sigma_i = \{\sigma_{i.s}\}_{s \in S_i}$
- a profile of behavioral strategy by $\sigma = \{\sigma_i\}_{i \in N}$
- the chance probabilities by $\pi = \{\pi_{0.s}\}_{s \in S_0}$

Probabilities over Terminal Nodes

For any terminal node y and any behavioral strategy profile σ, let $P(y|\sigma)$ denote the probability that the game ends at node y.

E.g. in the following game $P(c|\sigma) = \pi_{0.1}(1)\sigma_{1.1}(R)\sigma_{2.2}(C) = 1/9$:

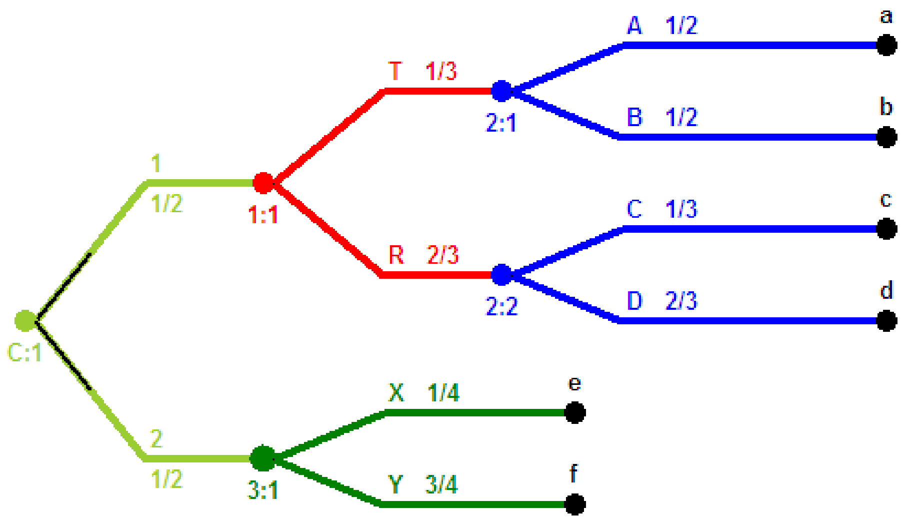

Expected Payoffs

If Ω denotes the set of end notes, the expected payoff of player i is:

$$U_i(\sigma) = U_i(\sigma_i, \sigma_{-i}) = \sum_{y \in \Omega} P(y|\sigma) u_i(y)$$

E.g. in the following game $U_1(\sigma) = 4/9$ and $U_2(\sigma) = 5/9$.

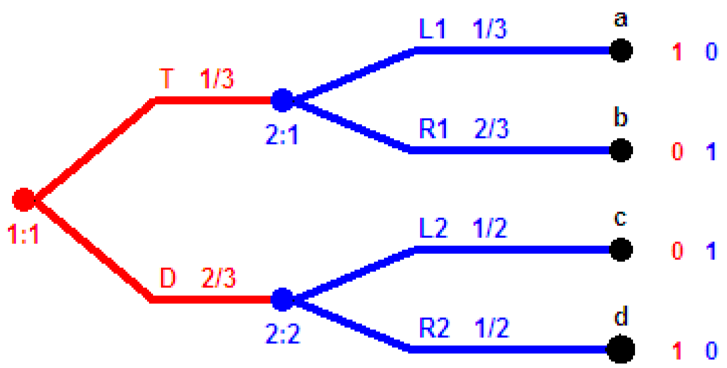

Nash Equilibrium

Nash Equilibrium

Definition (Nash Equilibrium – NE)

A Nash Equilibrium of an extensive form game is any profile of behavioral strategies such that:

$$U_i(\sigma) \geq U_i(\sigma'_i, \sigma_{-i}) \text{ for any } \sigma'_i \in \times_{s \in S_i} \Delta(A_{i.s})$$

Recall that σ'_i is any mapping from information sets to probability distributions over available actions.

The definition of NE is as in strategic form games.

What differs is the strategy (behavioral) that is expressed at every single decision stage and not on profiles of decisions for every individual.

A Simple NE Example

Consider the following three player game and the strategy:

- $\sigma_1(B) = 1$, $\sigma_2(D) = 1$, $\sigma_3(F) = 1$

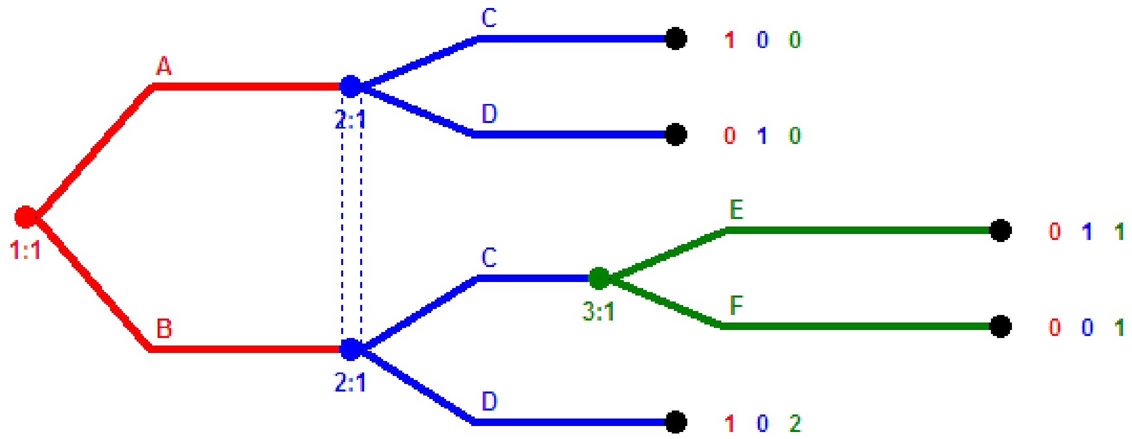

A Simple NE Example

Profile $\sigma_1(B) = 1$, $\sigma_2(D) = 1$, $\sigma_3(F) = 1$ is NE since:

$$\begin{aligned} U_1(B, \sigma_{-1}) &= 1 > 0 = U_1(A, \sigma_{-1}) \\ U_2(D, \sigma_{-2}) &= 0 \geq 0 = U_2(C, \sigma_{-2}) \\ U_3(F, \sigma_{-3}) &= 2 \geq 2 = U_3(E, \sigma_{-3}) \end{aligned}$$

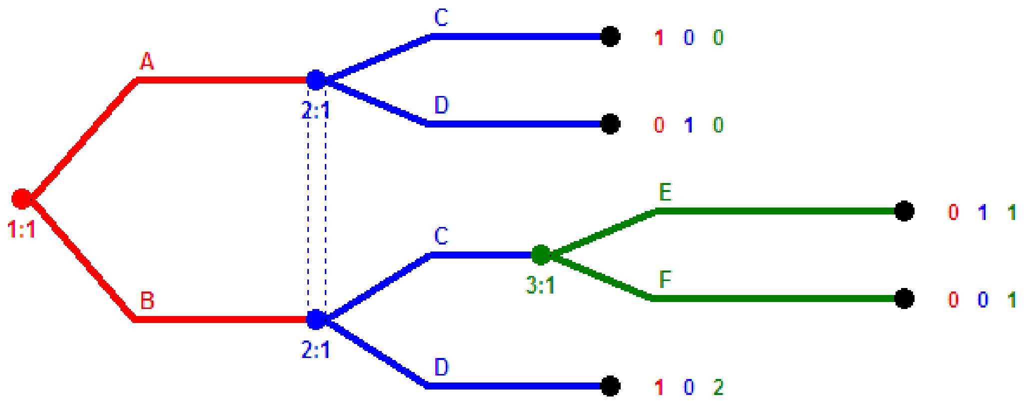

Ultimatum Game

This game possesses three types of NE, namely for any $a_{2.2}, a_{2.3}$.

NE1:
$\sigma_1 = [1] \ \& \ \sigma_2 = [A, a_{2.2}, a_{2.3}]$

NE2:
$\sigma_1 = [2] \ \& \ \sigma_2 = [R, A, a_{2.3}]$

NE3:
$\sigma_1 = [3] \ \& \ \sigma_2 = [R, R, A]$

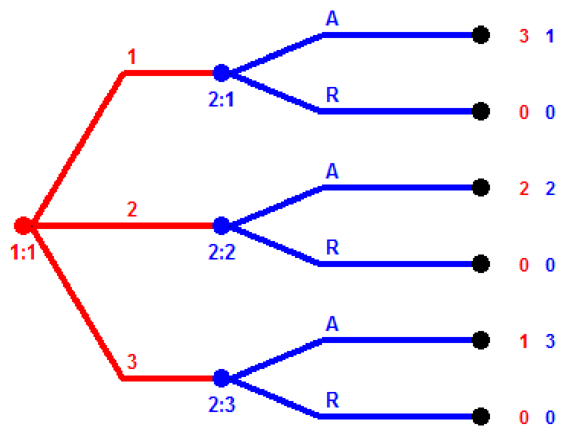

In the table, σ_1 and σ_2 denote the behavioral (pure) strategies of the two players.

Ultimatum Game

Strategy $\sigma_1 = [1]$, $\sigma_2 = [A, a_{2.2}, a_{2.3}]$ is NE since:

$$\begin{aligned} U_1(1, \sigma_2) &= 3 > U_1(a_1, \sigma_2) \ \text{ for any } a_1 \in \{2,3\} \\ U_2(\sigma_2, 1) &= 1 \geq U_2(a_2, 1) \ \text{ for any } a_2 \in \{A, R\}^3 \end{aligned}$$

Strategy $\sigma_1 = [1]$, $\sigma_2 = [R, A, a_{2.3}]$, is NE since:

$$\begin{aligned} U_1(2, \sigma_2) &= 2 > U_1(a_1, \sigma_2) \ \text{ for any } a_1 \in \{1,3\} \\ U_2(\sigma_2, 2) &= 2 \geq U_2(a_2, 2) \ \text{ for any } a_2 \in \{A, R\}^3 \end{aligned}$$

A similar argument works for the other proposed NE.

Only the first type of NE however, involves threats that are credible, since player 2 would never credibly reject an offer worth at least 1 when faced with an outside option of 0.

A Transformation to Find NE

The NE of the game can be found by looking at NE of the strategic form:

1/2	*AAA*	*RAA*	*ARA*	*AAR*	*RRA*	*RAR*	*ARR*	*RRR*
1	3, 1	0, 0	3, 1	3, 1	0, 0	0, 0	3, 1	0, 0
2	2, 2	2, 2	0, 0	2, 2	0, 0	2, 2	0, 0	0, 0
3	1, 3	1, 3	1, 3	0, 0	1, 3	0, 0	0, 0	0, 0

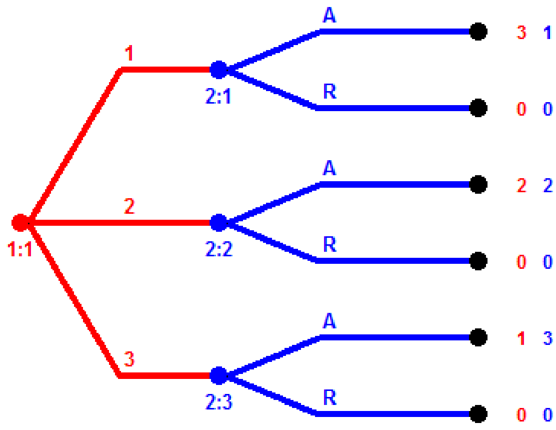

Subgame Perfect Equilibrium

Subgame Perfect Equilibrium

A **successor** of a node x is a node that can be reached from x for an appropriate profile of actions.

Definition (Subgame)

A *subgame* is a subset of an extensive form game such that:

1. It begins at a single node.
2. It contains all successors.
3. If a game contains an information set with multiple nodes then either all of these nodes belong to the subset or none does.

Definition (Subgame Perfect Equilibrium – SPE)

A *subgame perfect equilibrium* is any NE such that for every subgame the restriction of strategies to this subgame is also a NE of the subgame.

NE but not SPE

The following game has a continuum of NE but only one SPE:

- $\sigma_{1.1}(T) = 1$ and $\sigma_{2.1}(L) = 1$ is unique SPE.
- $\sigma_{1.1}(T) = 0$ and $\sigma_{2.1}(L) \leq 1/2$ are all NE, but none SPE.

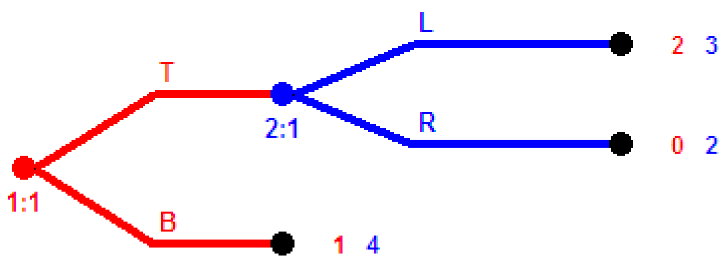

NE but not SPE

Strategy $\sigma_{1.1}(T) = 1$ and $\sigma_{2.1}(L) = 1$ is the unique SPE since:

$$U_{1.1}(T, L) = 2 \quad > \quad 1 = U_1(B, L)$$
$$U_{2.1}(L) = 3 \quad > \quad 2 = U_{2.1}(R)$$

Any strategy $\sigma_{1.1}(T) = 0$ and $\sigma_{2.1}(L) = q \leq 1/2$ is NE since:

$$U_1(T, \sigma_2) = 2q \quad \geq \quad 1 = U_1(B, \sigma_2)$$
$$U_2(L, B) = 4 \quad = \quad 4 = U_{2.1}(R, B)$$

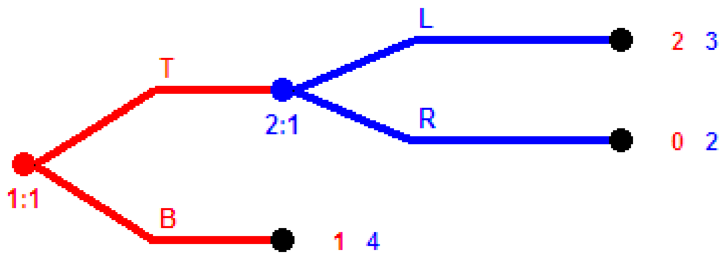

Computing SPE – Backward Induction

Definition of SPE is demanding because it imposes discipline on behavior even in subgames that one expects not to be reached.

SPE however is easy to compute in perfect information games.

Backward-induction algorithm provides a simple way:

- At every node leading only to terminal nodes players pick actions that are optimal for them if that node is reached.
- At all preceding nodes players pick an actions that optimal for them if that node is reached knowing how all their successors behave.
- And so on until the root of the tree is reached.

A pure strategy SPE exists in any perfect information game.

Backward Induction Example

SPE: $\sigma_{1.2} = \sigma_{1.3} = [A]$, $\sigma_{2.1} = [t]$, any $\sigma_{2.2}$ and $\sigma_{1.1} = [T]$.

NE but not SPE: $\sigma_{1.2} = \sigma_{1.3} = [A]$, $\sigma_{2.1} = [b]$, any $\sigma_{2.2}$ and $\sigma_{1.1} = [B]$.

Again NE may support empty threats.

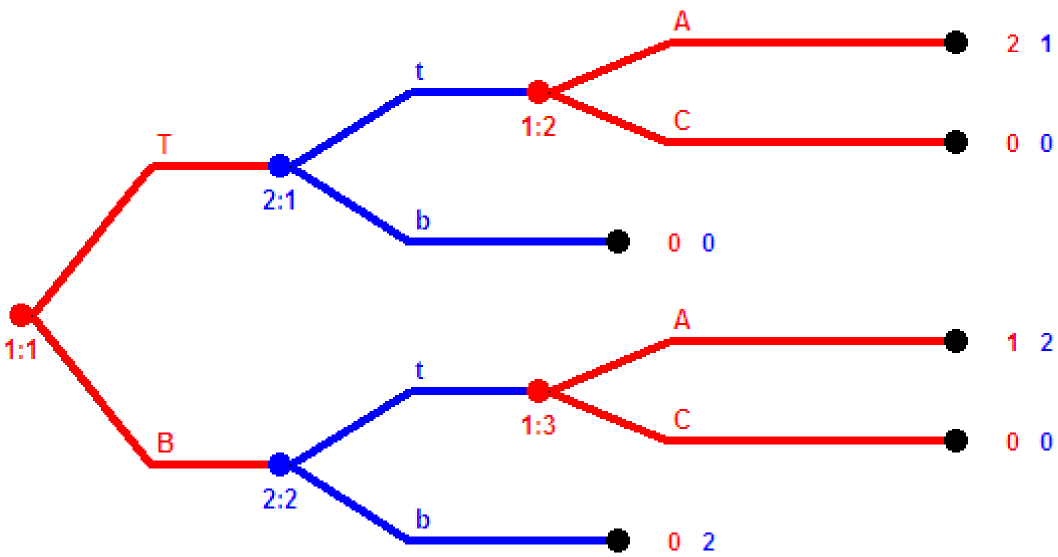

Dynamic Oligopoly

Duopoly: Stackelberg Competition

Implicit to both Cournot and Bertrand models was the assumption that no producer could observe actions chosen by others before making a decision.

In the Stackelberg duopoly model however:

- Players choose how many goods to supply to the market (as Cournot).
- One player moves first (the **leader**).
- While the other player moves after having observed the decision of the leader (the **follower**).
- Both players account for the distortions that their output choices have on equilibrium prices.

To avoid empty threats restrict attention to the SPE of the dynamic game.

Duopoly: Stackelberg Competition

The game is solved by backward induction:

- Consider the subgame in which the leader has produced q_L units.
- In this subgame (as in Cournot) the decision problem of follower is to:

$$\max_{q_F} p(q_L + q_F)q_F - c_F(q_F)$$

- Solving such problem defines the best response to the follower $b_F(q_L)$.
- By SPE the leader takes the follower's strategy into account when choosing his output.
- Thus the decision problem of the leader is as follows:

$$\max_{q_L} p(q_L + b_F(q_L))q_L - c_L(q_L)$$

Stackelberg Example

Consider the following economy:

- $d(q) = 2 - q$
- $c_F(q) = q^2$ and $c_L(q) = 3q^2$

The problems of both players are respectively defined by:

$$\max_{q_F} (2 - q_L - q_F) q_F - c_F(q_F)$$
$$\max_{q_L} (2 - q_L - b_F(q_L)) q_L - c_L(q_L)$$

Optimality of each firm is determined by FOC:

$$2 - 2q_F - q_L - 2q_F = 0 \Rightarrow q_F = b_F(q_L) = (2 - q_L)/4$$
$$3/2 - (3/2)q_L - 6q_L = 0 \Rightarrow q_L = 1/5$$

Stackelberg Equilibrium outputs are: $q_F = 9/20$ and $q_L = 1/5$

Cournot Equilibrium outputs are: $\overline{q}_F = 14/31$ and $\overline{q}_L = 6/31$

Duopoly: Market Entry

Consider the following game between two producers.

Firm 1 is the incumbent and firm 2 is the potential entrant.

Assume $P > L > 0$ and $M > P > F$.

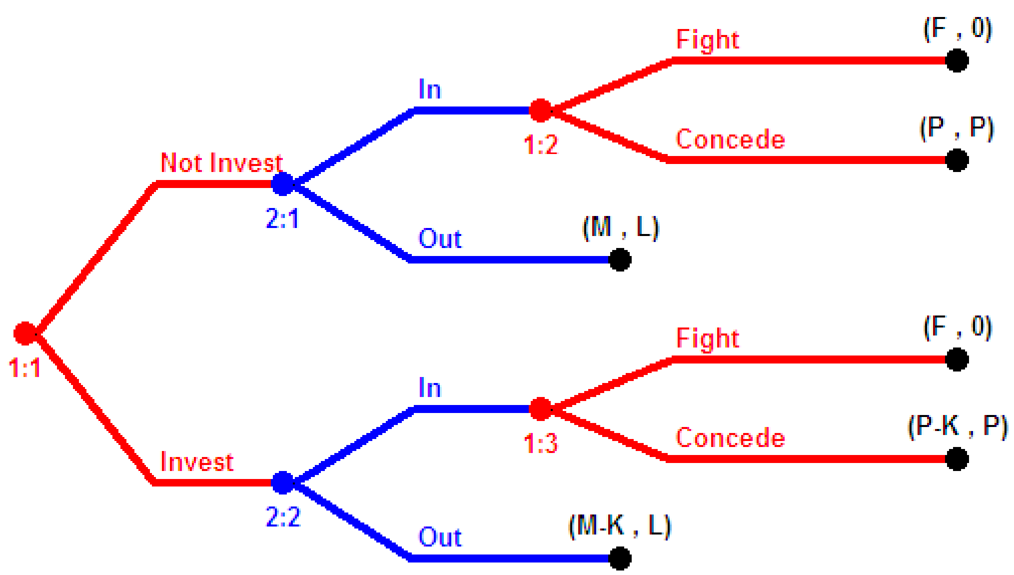

Duopoly: Market Entry

To find an SPE with successful deterrence, notice that:

- If 1 does not invest, it prefers to concede if entry takes place as $P > F$.
- Thus firm 2 prefers to enter if 1 does not invest as $P > L$.
- If 1 does invest it prefers to fight if entry takes place, provided that:

$$F > P - K$$

- If so firm 2 prefers to stay out if 1 has invested as $L > 0$.
- Thus firm 1 prefers to invest and deter entry if:

$$M - K > P$$

An SPE exists in which entry is effectively deterred if the cost satisfies:

$$M - P > K > P - F$$

Extra: Uncertainty and Dynamics

Extra: Dynamics and Uncertainty

If an extensive form game does not display perfect information, subgame perfection cannot be imposed at every information set, but only on subgames.

In such games a further equilibrium refinement may help to highlight the relevant equilibria of the game by selecting those which Bayes rule.

Definition (Perfect Bayesian Equilibrium – PBE)

A *perfect Bayesian equilibrium* of an extensive form game consists of a profile of behavioral strategies and of beliefs at each information set of the game such that:

1. strategies form an SPE given the beliefs;
2. beliefs are updated using Bayes rule at each information set reached with positive probability.

Repeated Games

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Repeated Games:

- Definitions:
  - Feasible Payoffs
  - Minmax
  - Repeated Game
  - Stage Game
  - Trigger Strategy
- Main Result:
  - Folk Theorem
- Examples: Prisoner's Dilemma

Feasible and Minmax Payoffs

Feasible Payoffs

Q: What payoffs are **feasible** in a strategic form game?

A: A profile of payoffs is feasible in a strategic form game if can be expressed as a weighed average of payoffs in the game.

Definition (Feasible Payoffs)

A profile of payoffs $\{w_i\}_{i \in N}$ is feasible in a strategic form game $\left\{N, \{A_i, u_i\}_{i \in N}\right\}$ if there exists a distribution over profiles of actions π such that:

$$w_i = \sum_{a \in A} \pi(a) u_i(a) \quad \text{for any} \quad i \in N$$

Unfeasible payoffs cannot be outcomes of the game

Points on the north-east boundary of the feasible set are Pareto efficient

Minmax

Q: What's the worst possible payoff that a player can achieve if he chooses according to his best response function?

A: The **minmax** payoff.

Definition (Minmax)

The (pure strategy) *minmax* payoff of player $i \in N$ in a strategic form game $\left\{N, \{A_i, u_i\}_{i \in N}\right\}$ is:

$$\underline{u}_i = \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i})$$

Extra: The mixed strategy minmax payoff satisfies:

$$\underline{v}_i = \min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma_i, \sigma_{-i})$$

The mixed strategy minmax is not higher than the pure strategy minmax.

Example: Prisoner's Dilemma

Minmax Payoff: $(1, 1)$

Feasible Payoffs: blue and red areas.

Individually Rational Payoffs: red area.

Stage Game

1\2	w	s
w	2,2	0,3
s	3,0	1,1

Payoffs

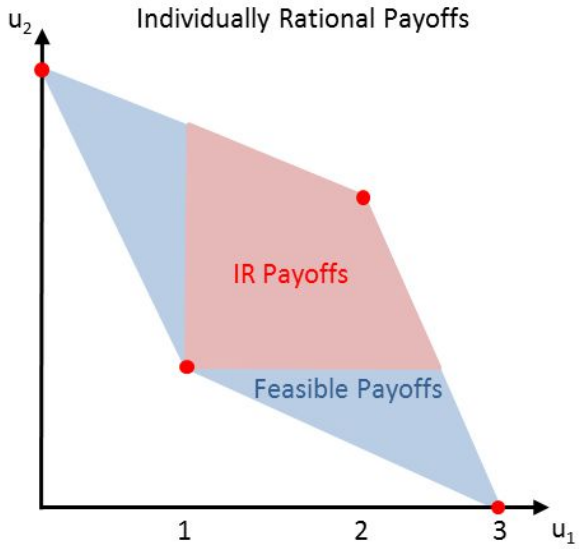

Example: Battle of the Sexes

Minmax Payoff: $(2, 2)$

Feasible Payoffs: blue and red areas.

Individually Rational Payoffs: red area.

Stage Game

1\2	w	s
w	3,3	1,0
s	0,1	2,2

Payoffs

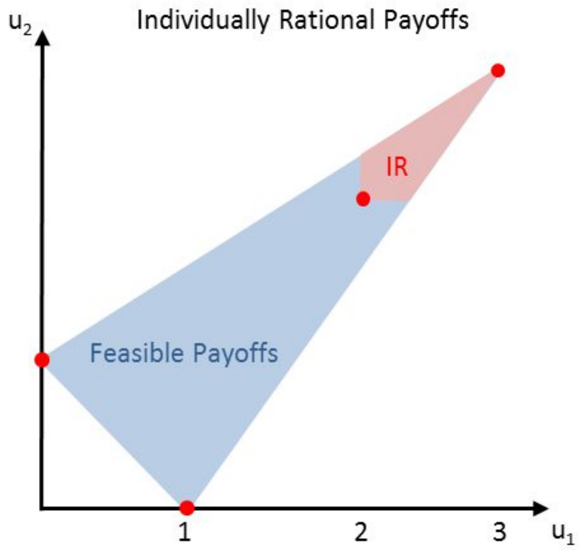

Repeated Games

Repeated Game: Timing

Consider any strategic form game: $G = \{N, \{A_i, u_i\}_{i \in N}\}$

Call G the **stage game**.

An infinitely **repeated game** is a strategic environment in which the stage game is played repeatedly by the same players infinitely many times.

Round 1 ... Round t ...

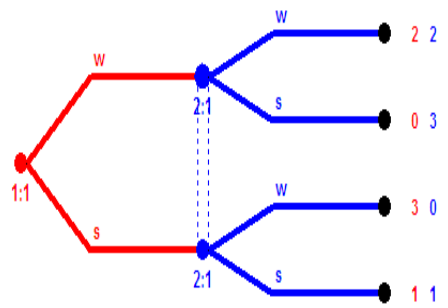

...

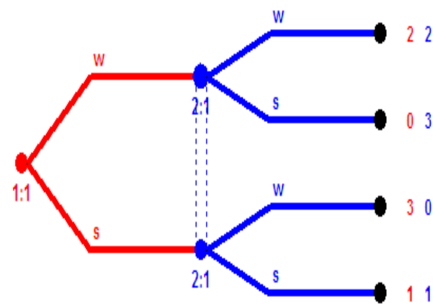

...

Repeated Games: Payoffs and Discounting

The value to player $i \in N$ of a payoff stream $\{u_i(1), u_i(2), ..., u_i(t), ...\}$ is:

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u_i(t)$$

The term $(1-\delta)$ amounts to a simple normalization,

... and guarantees that a constant stream $\{v, v, ...\}$ has value v.

Future payoffs are discounted at rate δ.

An infinitely repeated game can be used to describe strategic environments in which there is no certainty of a final stage.

In such view δ describes the probability that the game does not end at the next round which would result in a payoff of 0 thereafter.

Repeated Games: Perfect Information and Strategies

Today we restrict attention to perfect information repeated games.

In such games all players prior to each round observe the actions chosen by all other players at previous rounds.

Let $a(s) = \{a_1(s), ..., a_n(s)\}$ denote the action profile played at round s.

A **history** of play up to stage t thus consists of:

$$h(t) = \{a(1), a(2), ..., a(t-1)\}$$

In this context strategies map histories (ie information) to actions:

$$\alpha_i(h(t)) \in A_i$$

Sustaining Payoffs in Equilibrium

Prisoner's Dilemma Folk Theorem

Consider the prisoner's dilemma discussed earlier:

1\2	w	s
w	2,2	0,3
s	3,0	1,1

To understand how equilibrium behavior is affected by repetition, let's show why $(2, 2)$ is SPE of the infinitely repeated prisoner's dilemma.

Folk theorem shows that any feasible payoff that yields to both players at least their minmax value is a SPE of the infinitely repeated game if the discount factor is sufficiently high.

Grim Trigger Strategies

Consider the following strategy (known as **grim trigger strategy**):

$$a_i(t) = \begin{cases} w & \text{if } a(z) = (w, w) \text{ for any } z < t \\ s & \text{otherwise} \end{cases}$$

Such a strategy can be graphically represented as a 2-state automaton:

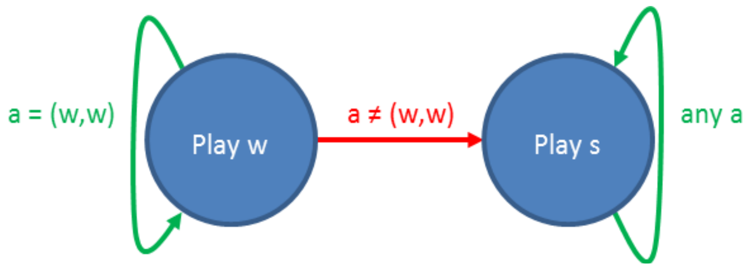

Grim Trigger SPE

Consider the grim trigger strategy:

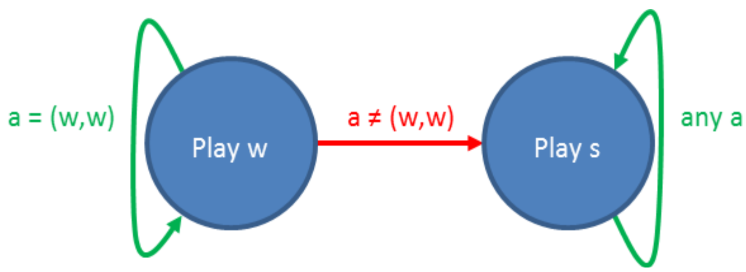

If all players follow such strategy, no player can deviate and benefit at any given round provided that $\delta \geq 1/2$.

In subgames in which $a(t) = (w, w)$ no player benefits from a deviation as:

$$(1-\delta)\left(3+\delta+\delta^2+\delta^3+...\right) = 3 - 2\delta \leq 2 \quad \Leftrightarrow \quad \delta \geq 1/2$$

In subgames in which $a(t) = (s, s)$ no player benefits from a deviation as:

$$(1-\delta)\left(0+\delta+\delta^2+\delta^3+...\right) = \delta \leq 1 \quad \Leftrightarrow \quad \delta \leq 1$$

Folk Theorem

Theorem (SPE Folk Theorem)

In any two-person infinitely repeated game:

1. *For any discount factor δ, the discounted average payoff of each player in any SPE is at least his minmax value in the stage game.*
2. *Any feasible payoff profile that yields to all players at least their minmax value is the discounted average payoff of a SPE if the discount factor δ is sufficiently close to 1.*
3. *If the stage game has a NE in which each players' payoff is his minmax value, then the infinitely repeated game has a SPE in which every players' discounted average payoff is his minmax value.*

Testing SPE in Repeated Games

Definition (One-Deviation Property)

A strategy satisfies the *one-deviation property* if no player can increase his payoff by changing his action at the start of any subgame in which he is the first-mover given other players' strategies and the rest of his own strategy.

Fact

A strategy profile in an extensive game with perfect information and infinite horizon is a SPE if and only if it satisfies the one-deviation property.

This observation can be used to test whether a strategy profile is a SPE of an infinitely repeated game as we did in the Prisoner's dilemma example.

Example: Different Punishments

Consider the following game – with minmax payoffs of $(1, 1)$:

1\2	A	B
A	0,0	4,1
B	1,4	3,3

Two PNE: (A, B) and (B, A) with payoffs $(1, 4)$ and $(4, 1)$.

Always playing (B, B) is SPE of the repeated game for δ high enough.

Consider the following grim trigger strategy:

$$a(t) = \begin{cases} (B,B) & \text{if } \; a(s) = (B,B) \; \text{ for any } \; s < t \\ (B,A) & \text{if } \; a(s) = \begin{cases} (B,B) & \text{for } \; s < z \\ (A,B) & \text{for } \; s = z \end{cases} \quad \text{for } \; z \in \{0, .., t-1\} \\ (A,B) & \text{if } \; a(s) = \begin{cases} (B,B) & \text{for } \; s < z \\ (B,A) & \text{for } \; s = z \end{cases} \quad \text{for } \; z \in \{0, .., t-1\} \end{cases}$$

Example: Different Punishments

1\2	A	B
A	0,0	4,1
B	1,4	3,3

If all follow such strategy, no player can deviate and benefit if $\delta \geq 1/3$.

When $a(t) = (B, B)$, no player player benefits from a deviation if $\delta \geq 1/3$:

$$(1-\delta)\left(4 + \delta + \delta^2 + \delta^3 + ...\right) = 4 - 3\delta \leq 3 \quad \Leftrightarrow \quad \delta \geq 1/3$$

When $a(t) = (B, A)$, no player player benefits from a deviation as:

$$\begin{aligned} (1-\delta)\left(0 + \delta + \delta^2 + \delta^3 + ...\right) &= \delta \leq 1 \quad \Leftrightarrow \quad \delta \leq 1 \\ (1-\delta)\left(3 + 4\delta + 4\delta^2 + 4\delta^3 + ...\right) &= 3 + \delta \leq 4 \quad \Leftrightarrow \quad \delta \leq 1 \end{aligned}$$

When $a(t) = (A, B)$, no one wishes to deviate for symmetric reasons.

Example: Different Equilibrium Behavior

Next show why $(1.5, 1.5)$ is also an SPE of the repeated PD as $\delta \to 1$.

Consider the following pair of strategies:

$$a(t) = \begin{cases} (w, s) & \text{if } t \text{ is even } \textbf{and } \; a(z) \notin \{(s,s),(w,w)\} \; \text{ for any } z < t \\ (s, w) & \text{if } t \text{ is odd } \textbf{and } \; a(z) \notin \{(s,s),(w,w)\} \; \text{ for any } z < t \\ (s, s) & \text{otherwise} \end{cases}$$

If $\delta \geq 1/2$, no player can profitably deviate from the strategy.

When $a(t) = (s, w), (w, s)$, no player benefits from a deviation as:

$$1 \leq (1-\delta)\left(3\delta + 3\delta^3 + 3\delta^5 + ...\right) = 3\delta/(1+\delta)$$

$$(1-\delta)\left(2 + \delta + \delta^2 ...\right) = 2 - \delta \leq (1-\delta)\left(3 + 3\delta^2 + 3\delta^4 ...\right) = 3/(1+\delta)$$

When $a(t) = (s, s)$, no player benefits from a deviation as:

$$(1-\delta)\left(0 + \delta + \delta^2 + \delta^3 + ...\right) = \delta \leq 1 \quad \Leftrightarrow \quad \delta \leq 1$$

Adverse Selection

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Adverse Selection:

1. Hidden Characteristics
2. Uninformed party moves first

- Monopoly:
  - One type of consumer
  - Multiple types of consumer
- Competition
- Definitions:
  - Pooling Equilibrium
  - Separating Equilibrium
- Market for Lemons

Monopoly Surplus Extraction

Adverse Selection: Monopoly Setup

- Consider an economy two goods, x and y.
- A firm produces good x using y at constant marginal cost c.
- The firm sells **bundles** (X, P) which cost P dollars and contain X units.
- An alternative pricing schemes considered in the literature are:
  - **Uniform tariff**
  $$P(x) = px$$
  - **Two-part tariff**
  $$P(x) = p_0 + p_1 x \quad \text{if} \quad x > 0$$

Complete Information: One Consumer Type

Begin by looking at the complete information benchmark:

- All consumers are all identical.
- Endowments given by $(e_x, e_y) = (0, Y)$.
- The budget constraint of an individual then requires that:
$$y(x) = \begin{cases} Y - P & \text{if} \quad x = X \\ Y & \text{if} \quad x = 0 \end{cases}$$
- Preferences over the two goods are given by:
$$U(x, y(x)) = u(x) + y(x) = \begin{cases} u(X) + Y - P & \text{if} \quad x = X \\ Y & \text{if} \quad x = 0 \end{cases}$$
- Assume: $u(0) = 0$, $u_x > 0$, $u_{xx} < 0$.

Consumer & Firm Decision Problems

Given such a bundle, consider the decision problem of the consumer:

- A consumer purchases the bundle (X, P) only if

$$U(X, y(X)) - U(0, Y) = u(X) - P \geq 0 \qquad \text{(PC)}$$

- Such constraint is known as **Participation Constraint**.

Given such demand consider the decision problem of the firm:

- A firm chooses the bundle (X, P) to maximize profits

$$\max_{X,P} P - cX \quad \text{subject to PC}$$

- PC must hold with equality or else the firm could increase profits by raising the price while holding fixed the quantity

$$P = u(X)$$

Solving the Firm Decision Problem

- First order conditions of the firm's problem then require that:

$$c = \lambda u_x(X) \text{ and } \lambda = 1$$

  where λ denotes the Lagrange multiplier on the PC.

- Defining $\varphi = u_x^{-1}$, the solution of the problem simply amounts to

$$(X, P) = (\varphi(c), u(\varphi(c)))$$

- Unlike in the standard monopolist problem, the solution of this problem is **efficient as prices equal marginal costs**.

- It is still exploitative however because buyers are left at their reservation utility:

$$U(X, y(X)) - U(0, Y) = 0$$

Extra: Alternative Implementation

- The same conclusion obtains with two-part tariffs.
- The optimal two-part tariff $P(x) = p_0 + p_1 x$ sets:

$$p_1 = c \;\; \& \;\; p_0 = u(\varphi(c)) - p_1 \varphi(c)$$

- Given the tariff demand amounts to $x = \varphi(c)$.
- The fixed fee p_0 is set so to have PC binding.
- The firm effectively chooses x by changing $P(\cdot)$ exploiting FOC.
- Thus choosing a two-part tariff and choosing a bundle are equivalent.

Complete Information: Multiple Consumer Types

Suppose that consumers have multiple types:

- Let $t \in \{L, H\}$ denote the type of a consumer with $H > L$.
- Let $\pi(t)$ denote the proportion of types t in the population.
- The monopolist knows the type of every consumer.
- Preferences of a consumer of type t are:

$$U(x, y) = tu(x) + y = \begin{cases} Y + tu(X) - P & \text{if} \quad x = X \\ Y & \text{if} \quad x = 0 \end{cases}$$

- Setup meets the **single crossing condition** which requires indifference curves of the two types to cross only once.
- Consumers cannot resell the units purchased.

Complete Information: Multiple Consumer Types

With more types and complete info not much changes:

- The firm **price discriminates** by selling two bundles $(X(t), P(t))$.
- The participation constraint of each type t becomes

$$U(x, y|t) - U(0, Y|t) = tu(X(t)) - P(t) \geq 0 \qquad \text{(PC(t))}$$

- As before PC(t) holds with equality $P(t) = tu(X(t))$.
- Using these two facts the problem of the monopolist's becomes:

$$\max_{X(H),X(L)} \sum_t \pi(t)[tu(X(t)) - cX(t)] \Rightarrow tu_x(X(t)) = c$$

- The resulting equilibrium bundles amount thus to

$$(X(t), P(t)) = (\varphi(c/t), tu(\varphi(c/t)))$$

Adverse Selection

Incomplete Information: Multiple Consumer Types

If the firm cannot recognize the two types and knows only $\pi(t)$:

- Firm may still offer several bundles $(X(t), P(t))$...
  ... but cannot guarantee that type t purchases $(X(t), P(t))$.
- Each consumer decides which type he reports to be...
  ... and buys bundle $(X(s), P(s))$ if he reports to be type s.
- The net-payoff of a consumer of type t claiming to be s is:

$$V(s|t) = tu(X(s)) - P(s)$$

- If the firm keeps offering the complete information $P(t)$...
  ... both types of consumers purchase $P(L)$ since:

$$\begin{aligned} V(L|H) &= (H-L)u(X(L)) > 0 = V(H|H) \\ V(L|L) &= 0 > (L-H)u(X(H)) = V(H|L) \end{aligned}$$

Incomplete Information: No Pooling

Offering the same contracts however is not optimal for the firm:

- Consider decreasing $P(H)$ to $\bar{P}(H) > P(L)$ so that:

$$V(H|H) = Hu(X(H)) - \bar{P}(H) = V(L|H)$$

- Such a change would increase the firm's profits as:

$$\pi(H)\bar{P}(H) + \pi(L)P(L) \geq P(L)$$

Theorem (No Pooling)

It is not optimal for the firm to offer contracts that lead consumers to pool.

Incomplete Info: Participation & Incentive Constraints

The previous remark implies that the firm wants to satisfy both:

- The **participation constraint** for any type $t \in \{L, H\}$:

$$V(t|t) \geq 0 \qquad \text{(PC(t))}$$

- The **incentive constraint** for any type $t \neq s \in \{L, H\}$:

$$V(t|t) \geq V(s|t) \qquad \text{(IC(t))}$$

The problem of the firm can now be written as

$$\max_{X(\cdot),P(\cdot)} \sum_{t\in\{L,H\}} \pi(t)[P(t) - cX(t)] \text{ subject to}$$
$$V(t|t) \geq V(s|t) \text{ for any } t \in \{L, H\}$$
$$V(t|t) \geq 0 \text{ for any } t \in \{L, H\}$$

Incomplete Information: Binding Constraints

Prior to solving the problem, notice that:

- PC(L) holds with equality (otw firm can increase profits raising $P(L)$):

$$V(L|L) = Lu(X(L)) - P(L) = 0$$

- IC(H) holds with equality (otw firm can increase profits raising $P(H)$):

$$V(H|H) = Hu(X(H)) - P(H) = Hu(X(L)) - P(L) = V(L|H)$$

- PC(H) is strict (by the previous two equalities and $H > L$):

$$V(H|H) = Hu(X(H)) - P(H) > 0$$

- IC(L) is strict (by no pooling theorem as otw $X(H) = X(L)$):

$$V(L|L) = Lu(X(L)) - P(L) > Lu(X(H)) - P(H) = V(H|L)$$

Incomplete Information: Binding Constraints

Graphically the previous observations follow as:

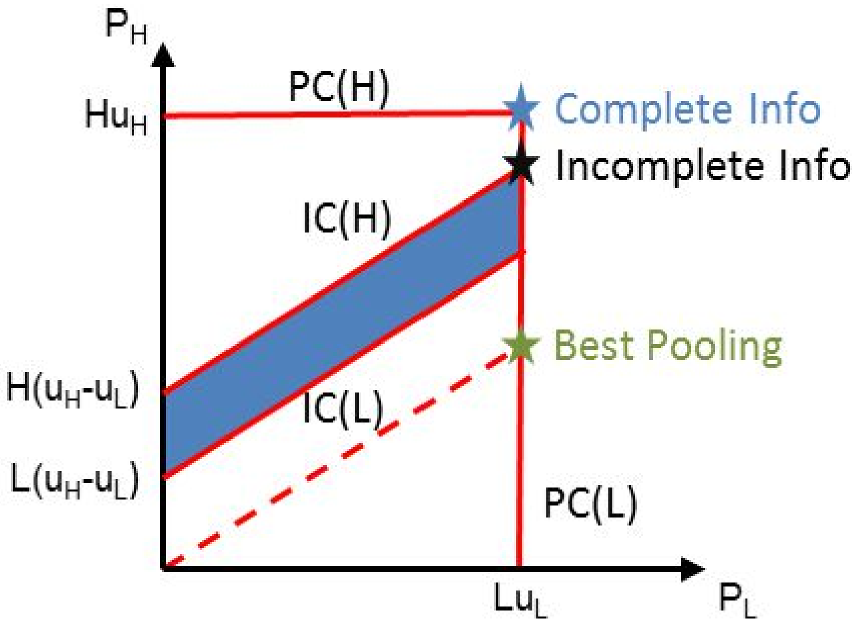

Incomplete Information: Optimal Pricing

The previous remarks simplify the firm's problem to:

$$\max_{X(t),P(t)} \left[\sum_{t\in\{L,H\}} \pi(t)[P(t) - cX(t)]\right] + \lambda V(L|L) + \mu[V(H|H) - V(L|H)]$$

First order optimality requires:

$$\begin{aligned} -\pi(H)c + \mu H u_x(X(H)) &= 0 && (\text{x(H)}) \\ -\pi(L)c + \lambda L u_x(X(L)) - \mu H u_x(X(L)) &= 0 && (\text{x(L)}) \\ \pi(H) - \mu &= 0 && (\text{P(H)}) \\ \pi(L) - \lambda + \mu &= 0 && (\text{P(L)}) \end{aligned}$$

Notice that $\mu = \pi(H)$, $\lambda = 1$ and thus:

$$\begin{aligned} Hu_x(X(H)) &= c \\ Lu_x(X(L)) &= \frac{c}{1 - (\pi(H)/\pi(L))\,[(H/L) - 1]} \end{aligned}$$

$P(H)$ and $P(L)$ are pinned down by the two binding constraints.

Incomplete Information: No Distortion at the Top

Notice that the optimality conditions for $x(t)$ require that:

$$\begin{aligned} MRS_{xy}(H) &= MRT_{xy} = c \\ MRS_{xy}(L) &> MRT_{xy} = c \end{aligned}$$

This principle carries over to more general setups and requires:

Theorem (No Distortion at the Top)

In the **second-best optimum** *for the firm, the* **high valuation** *types are offered a* **non-distortionary** *(efficient)* **contract**.

In general (if the single-crossing condition is met) second-best optimum $X_{SB}(t)$ when compared to full-information optimum $x_{FB}(t)$ satisfies:

$$\begin{aligned} X_{SB}(H) &= X_{FB}(H) \\ X_{SB}(L) &< X_{FB}(L) \\ X_{SB}(L) &< X_{SB}(H) \end{aligned}$$

Incomplete Information: Competition

With competition and free entry firms do not run positive profits.

Or else entering firms would profit by offering $\bar{P}(t) \in [cX, P(t))$.

As they would sell to all buyers $\Longrightarrow$ competition requires $P(t) = cX$.

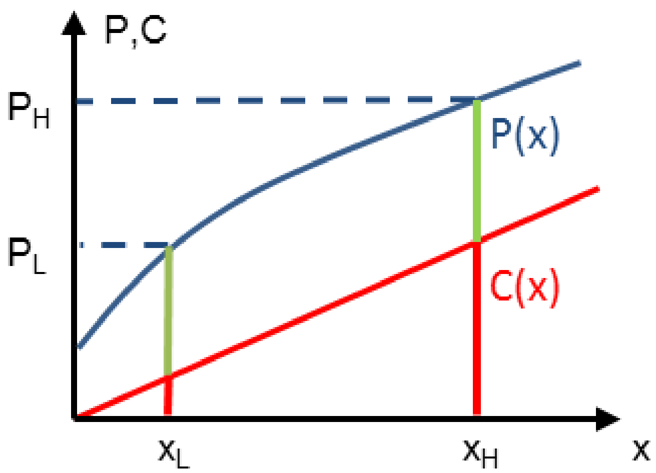

In blue P(x), in red C(x), in green monopoly profits on each type.

A Market for Lemon

Competition, Costs and Adverse Selection

A Market for Lemons

Consider the following economy:

- There are **two qualities** of goods $\{H, L\}$.
- There number of firms **selling one unit** of quality t is N_t.
- The cost of producing quality t is c_t and $c_H > c_L$.
- There is a **large number of buyers**, $N_B > N_H + N_L$.
- Every buyer who **wants to buy one unit** of the good.
- The value of buying quality t is u_t and $u_H > u_L$.
- **Gains from trade are certain** as $u_t > c_t$.
- Prices are determined by **competitive equilibrium**.
- We consider both scenarios in which quality is:
  - observable by buyers;
  - unobservable by buyers.

A Market for Lemons: Observable Quality

If the quality is observable, quality t sells at a quality-specific price p_t.

For prices (p_H, p_L) demand and supply satisfy:

- if the buyer demands quality H, then

$$u_H - p_H \geq \max\{u_L - p_L, 0\}$$

- if the buyer demands quality L, then

$$u_L - p_L \geq \max\{u_H - p_H, 0\}$$

- if the buyer does not demand any quality, then

$$0 \geq \max\{u_H - p_H, u_L - p_L\}$$

- if a seller supplies quality t, then

$$p_t \geq c_t$$

  where the converse holds if they do not supply.

A Market for Lemons: Observable Quality

For prices (p_H, p_L) to form an equilibrium **markets must clear**.

In a competitive equilibrium:

- the payoff of buyers equals zero

$$U(p) = \max\{u_H - p_H, u_L - p_L, 0\} = 0,$$

  since $U(p) > 0$ implies that the market does not clear as

$$D_H(p) + D_L(p) = N_B > N_H + N_L \geq S_H(p) + S_L(p);$$

- sellers of quality t must weakly prefer to sell

$$p_t \geq c_t,$$

  as $p_t < c_t$, implies $U(p) > 0$.

Hence, a **competitive equilibrium** $(p_H, p_L) = (u_H, u_L)$ always **exists**.

A Market for Lemons: Unobservable Quality

When quality is unobservable, sellers can claim to sell any quality.

Several types of competitive equilibria can arise:

- **separating equilibria** where qualities sell at different prices (p_H, p_L):
  - in some of these at most one quality is traded;
  - in others both qualities can be traded.
- **pooling equilibria** where qualities sell at a common price $\bar{p}$:
  - in some of these no quality is traded;
  - in others both qualities can be traded.

For convenience, denote the fraction of sellers of type t by

$$\pi_t = \frac{N_t}{N_H + N_L}.$$

A Market for Lemons: Separating Equilibria

In separating equilibria, the payoff of buyers equals zero

$$U(p) = \max\{u_H - p_H, u_L - p_L, 0\} = 0,$$

as players know the quality in any separating equilibrium as before.

If **at most one quality is traded**, there exists a quality t such that

$$u_t \leq p_t \quad \Rightarrow \quad c_t < p_t.$$

But, this **cannot be a separating equilibrium** as

$$D_t(p) = 0 < N_t = S_t(p).$$

If **both qualities are traded** at different prices, then $p_t = c_t$ and all sellers would claim to sell quality H.

So, this **cannot be a separating equilibrium.**

A Market for Lemons: Pooling Equilibria

In pooling equilibria, the payoff of buyers also equals zero

$$U(\bar{p}) = \max\{\pi_H u_H + \pi_L u_L - \bar{p}, 0\} = 0,$$

as $U(p) \geq 0$ by construction and since $U(p) > 0$ implies everyone buys

$$\bar{D}(\bar{p}) = N_B > N_H + N_L \geq \bar{S}(\bar{p}).$$

If **no quality is traded**, then

$$\pi_H u_H + \pi_L u_L < \bar{p} \quad \Rightarrow \quad c_L < \bar{p}.$$

But if so, this **cannot be a pooling equilibrium** as

$$\bar{D}(\bar{p}) = 0 < N_L = \bar{S}(\bar{p}).$$

If **both qualities are traded**, then $p_t \geq c_t$ for all t.

So, $\bar{p}$ **can be a pooling equilibrium if and only if**

$$c_L < c_H \leq \bar{p} \leq \bar{u}.$$

A Market for Lemons: Unravelling Summary

If quality is unobservable and $c_H > \bar{u}$:

- no separating equilibria (with or without trade) exist;
- no pooling equilibria (with or without trade) exist;
- **so no competitive equilibrium exists**!

If quality is unobservable and $c_H > \bar{u}$:

- no separating equilibria (with or without trade) exist;
- no pooling equilibria without trade exist;
- **a pooling competitive equilibrium exists** in which

$$\bar{p} = \pi_H u_H + \pi_L u_L = \bar{u}.$$

A Market for Lemons: Intuition

With unobservable quality, **equilibria may fail to exist**.

Intuitively, in such instances there is no equilibrium as:

- in any pooling equilibrium H sellers want to separate;
- in any separating equilibrium L sellers want to pool.

This can happen, **even if gains from trade are certain** – that is $u_t > c_t$.

Similar models can explain **rationing in competitive insurance markets**.

The key difference relative to the earlier setting is that **private information directly affects the payoff of the uninformed party**.

A Market for Lemons: Example

Consider an economy in which:

- $N_B = 20,\ u_H = 9,\ u_L = 6;$
- $N_H = 3,\ c_H = k > 4;$
- $N_L = 6,\ c_L = 4.$

If so, a competitive pooling equilibrium exists when

$$c_H = k \leq 7 = \bar{u}.$$

But, no competitive equilibrium exists when $k > 7$.

Signaling

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Signaling:

1. Hidden Characteristics
2. Informed party moves first

- Costly Signals:
  - Educational Choice
  - Pooling Equilibria
  - Separating Equilibria
- Costless Signals

Spence Signaling Model

A Signaling Model

Consider the following educational choice model:

- There are two types of workers $\{g, b\}$
- Type t having probability π_t
- Workers can signal their type by acquiring education
- Different types have different costs to acquire education
- Firms can distinguish workers only by their education and ...
  ... compete on wages to hire workers given such information

Timing:

1. Nature determines the type of each worker
2. Workers decide how much education to acquire
3. Firms simultaneously make wage offers (Bertrand)
4. Workers decide whether or not to accept an offer

Signaling: Educational Choice Model

In particular consider the following model:

- Individuals can acquire any level $e \in [0, 1]$ education
- The cost of acquiring level e education for type t is $c(e|t)$
- Suppose that costs satisfy:

$$c(0|t) = 0 \ \ \& \ \ c_e(e|t) > 0 \ \ \& \ \ c_{ee}(e|t) > 0$$
$$c_e(e|g) < c_e(e|b)$$

- Suppose that firms offer a wage schedule $w(e)$
- If so the payoff of a worker of type t with education z is:

$$u(e|t) = w(e) - c(e|t)$$

- Assumptions on costs and preference imply that the single crossing condition is met (ie indifference curves cross once)

Educational Choice Model: Firms

In this economy firms are modeled as follows:

- Firms know that the productivity of a worker of type t is $f(t)$.
- If firms knew the type of a worker, they'd pay type t exactly $f(t)$.

  Bertrand competition implies that wages equal productivity.

  Recall that Bertrand $\Rightarrow$ price equals marginal cost.
- Initially firms know only the probability that a worker is of type t, π_t.
- Firms can observe the educational decisions of workers.
- After workers have made their educational decision, firms:
  1. update their beliefs on the basis of this new information;
  2. pick a wage schedule $w(e)$ that depends on education.

Equilibria and Signaling

Educational Choice Model: Types of Equilibria

This model has two different types of equilibria:

- **Separating equilibria** in which the two types of worker:
    1. choose different education levels;
    2. are paid different wages;
    3. prefer not to mimic the other type.
- **Pooling equilibria** in which the two types of worker:
    1. choose the same education level;
    2. are paid the same wage;
    3. prefer not to be separated from the other type.

Educational Choice Model: Separating Equilibria I

In a separating equilibrium:

- Workers of different type choose different education levels e_t.
- Firms recognize either type t by his education e_t.
- Firms pay each type its marginal productivity $w(e_t) = f(t)$. [Bertrand competition $\Rightarrow$ wages equal productivity]
- For types to reveal themselves at wage $w(e_t)$ it must be that:

$$\begin{aligned} f(g) - c(e_g|g) &\geq f(b) - c(e_b|g) && \text{(IC(g))} \\ f(b) - c(e_b|b) &\geq f(g) - c(e_g|b) && \text{(IC(b))} \end{aligned}$$

- Since education has no effect on productivity and since bad workers are identified by their education it must be that $e_b = 0$. [In a model with more types only the lowest acquires no education]

Educational Choice Model: Separating Equilibria II

- By the last observation, incentive constraints can be written as:

$$\begin{aligned} c(e_g|g) &\leq f(g) - f(b) \\ c(e_g|b) &\geq f(g) - f(b) \end{aligned}$$

- IC conditions require that $e_g \in [\underline{e}, \overline{e}]$ with boundaries defined by:

$$c(\overline{e}|g) = f(g) - f(b) = c(\underline{e}|b)$$

- If the two types of workers chose in equilibrium education levels $e_g \in [\underline{e}, \overline{e}]$ & $e_b = 0$ the ex-post beliefs of a firm are:

$$\pi_g(e) = \begin{cases} 1 & \text{if} \quad e = e_g \\ 0 & \text{if} \quad e = e_b \end{cases}$$

- Beliefs don't have to be pinned down for $e \neq e_g, e_b$.

Educational Choice Model: Separating Equilibria III

- To guarantee that no worker chooses $e \neq e_g, e_b$ suppose that:

$$w(e) = \begin{cases} f(b) & \text{if} \quad e < e_g \\ f(g) & \text{if} \quad e \geq e_g \end{cases}$$

- If such are the wages no good worker prefers to deviate since:

$$\begin{aligned} f(g) - c(e_g|g) &\geq f(g) - c(e|g) \text{ for } e > e_g \\ f(g) - c(e_g|g) &\geq f(b) - c(e|g) \text{ for } e < e_g \end{aligned}$$

- Moreover no bad worker prefers to deviate since:

$$\begin{aligned} f(b) &\geq f(g) - c(e|b) \text{ for } e > e_g \\ f(b) &\geq f(b) - c(e|b) \text{ for } e < e_g \end{aligned}$$

- Other wages schedules achieve the same result, but complicate the analysis unnecessarily.

Educational Choice Model: Separating Equilibria IV

There is a multiplicity of separating Perfect Bayesian equilibria:

- They are characterized by the education levels satisfying:

$$e_g \in [\underline{e}, \overline{e}] \quad \& \quad e_b = 0.$$

- Workers receive the efficient wage, namely their productivity.
- But no equilibrium is efficient since good workers lose resources to signal their type by investing in education.
- The Pareto dominant equilibrium is the one in which $e_g = \underline{e}$ since the cost of acquiring education is the lowest.
- The multiplicity is due to the unspecified beliefs for $e \neq e_g, e_b$.

Educational Choice Model: Separating Equilibria V

There is a multiplicity of separating Perfect Bayesian equilibria:

$$e_g \in [\underline{e}, \overline{e}] \equiv [e_L, e_H] \quad \& \quad e_b = 0.$$

Graphically these equilibria can be seen in the plot:

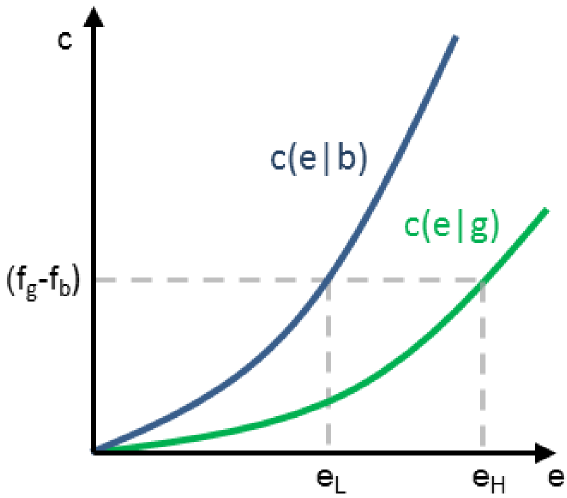

Educational Choice Model: Separating Equilibria VI

The multiplicity disappears for appropriately chosen beliefs:

- The **intuitive criterion** for out of equilibrium beliefs says:

$$e > \underline{e} \implies \pi_g(e) = 1.$$

- This criterion is reasonable as bad workers prefer $e = 0$ to $e > \underline{e}$:

$$f(b) > f(g) - c(e|b)$$

  which holds by definition of $\underline{e}$.

- If firms' beliefs meet the intuitive criterion the only PBE that survives is the one in which $e_g = \underline{e}$ and $e_b = 0$.
- Indeed if $e_g > \underline{e}$, good workers prefer to switch to $\underline{e}$ since:

$$f(g) - c(e_g|g) < f(g) - c(\underline{e}|g).$$

The **Pareto dominant equilibrium** is the only PBE that **survives** the intuitive criterion and involves the lowest education levels.

Educational Choice Model: Pooling Equilibria I

In a pooling equilibrium:

- All workers choose same education levels e^*.
- Firms cannot recognize workers by their education e^* and pay all workers their expected productivity:

$$w(e^*) = \pi_g f(g) + \pi_b f(b)$$

- To guarantee that no worker chooses $e \neq e^*$ suppose that:

$$w(e) = \begin{cases} f(b) & \text{if} \quad e < e^* \\ \pi_g f(g) + \pi_b f(b) & \text{if} \quad e \geq e^* \end{cases}$$

- For types not to reveal themselves at wage $w(e^*)$ it must be that:

$$w(e^*) - c(e^*|b) \geq f(b)$$

- Such condition requires $e^* \leq \widehat{e}$ where $\widehat{e}$ is defined by:

$$c(\widehat{e}|b) = \pi_g[f(g) - f(b)]$$

Educational Choice Model: Pooling Equilibria II

There is a **multiplicity of pooling** Perfect Bayesian **equilibria**:

- They are characterized by an education level $e^* \in [0, \widehat{e}]$.
- Workers wages are inefficient and differ from their productivity.
- Again the multiplicity is due to the unspecified off equilibrium beliefs.

No pooling equilibrium meets the intuitive criterion:

- Consider an effort level E defined by:

$$w(e^*) - c(e^*|b) = f(g) - c(E|b).$$

- When $e > E$, firms should then set $\pi_g(e) = 1$ and pay $w(e) = f(g)$.
- If so, good workers choose $e = E + \varepsilon$ while bad ones do not since

$$f(g) - c(e|g) > w(e^*) - c(e^*|g) > w(e^*) - c(e^*|b) > f(g) - c(e|b).$$

Educational Choice Model: Pooling Equilibria III

There is a **multiplicity of pooling** Perfect Bayesian **equilibria**:

$$e^* \in [0, \hat{e}] \equiv [0, e_M].$$

Graphically these equilibria can be seen in the plot:

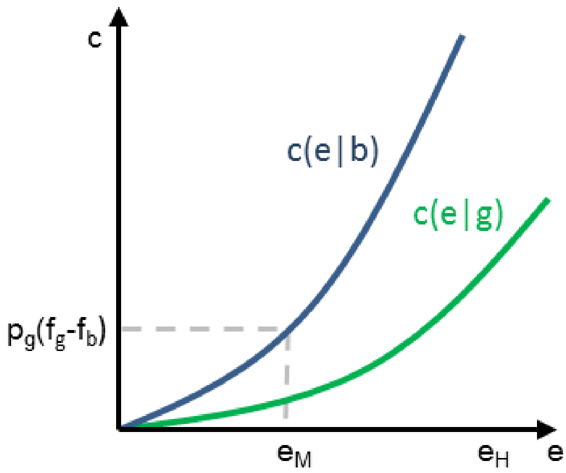

Comparison: Signaling vs Adverse Selection

The results on signaling **differ** from those on adverse selection since:

- There is a multiplicity of pooling equilibria.
- There is a multiplicity of separating equilibrium.
- The incentive constraint neither type may bind in a pooling equilibrium.

However most result **coincide** when the intuitive criterion is applied:

- There are no pooling equilibria.
- There is a unique separating equilibrium.
- The incentive constraint of the bad type binds.
- The incentive constraint of the good type does not bind.
- Inefficiencies arise to provide incentives to the good type.

Educational Choice Model: Example

Consider a setting in which:

- $c(e|g) = e^2$ and $f(g) = 46$;
- $c(e|b) = 4e^2$ and $f(b) = 10$;
- $\pi_g = 1/9$ and $\pi_b = 8/9$.

Thus, for separating equilibria, we have that

$$\begin{aligned} c(\overline{e}|g) &= \overline{e}^2 = 36 = f(g) - f(b) \quad \Rightarrow \quad \overline{e} = 6; \\ c(\underline{e}|b) &= 4\underline{e}^2 = 36 = f(g) - f(b) \quad \Rightarrow \quad \underline{e} = 3. \end{aligned}$$

While, for pooling equilibria, we have that

$$c(\underline{e}|b) = 4\underline{e}^2 = 4 = \pi_g(f(g) - f(b)) \quad \Rightarrow \quad \widehat{e} = 1.$$

Extra: Costless Signaling

Extra: Costless Signals I

The theory on costless signals is more involved. Some more result can be understood from the following example:

- There are N risk-neutral individuals
- All of them can participate in the production of a public good
- In particular player i chooses his effort $e_i \in \{0, 1\}$
- The public good is produced only if all exert $e_i = 1$
- The cost of exerting effort c_i is private information and costs are uniformly distributed on $[0, 1]$ – ie $\Pr(c_i < b) = b$
- The preferences of player i with cost c_i are – for $a < 1$:

$$u_i(e|c_i) = a \prod_{j \in N} e_j - c_i e_i$$

- The unique BNE of this game has all players exerting no effort
- This follows since there is positive probability that $c_i > a$ for some i

Extra: Costless Signals II

If a signaling stage is introduced prior to effort decision such that:

- Each agent can announce his willingness to exert effort
- In particular each agent can say $\{Yes, No\}$ to him exerting effort

When such signaling stage is added, then there is a BNE in which:

- Any agent announces *Yes* if and only if $c_i \leq a$
- Each agent chooses $e_i = 1$ if and only if all said *Yes*
- This is a BNE since individuals no longer risk wasting their effort

However, many other BNE are exist in which no information is exchanged at the communication stage. Such BNE are known as babbling equilibria.

Moral Hazard

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Summary

Hidden Action Problem aka:

1. **Moral Hazard** Problem
2. **Principal Agent** Model

- Outline
- Simple Moral Hazard Model:
  - Complete Information Benchmark
  - Hidden Effort
  - Agency Cost

Principal-Agent Models

Outline: Moral Hazard Problem

The basic ingredients of a moral hazard model are as follows:

- A **principal** and an **agent**, are involved in bilateral relationship.
- Principal wants Agent to perform some task.
- Agent can choose how much effort to devote to the task.
- The outcome of the task is pinned down by a mix of effort and luck.
- Principal cannot observe effort and can only motivate Agent by paying him based on the outcome of the task.

Timing:

1. Principal chooses a wage schedule which depends on outcome.
2. Agent chooses how much effort to devote the task.
3. Agent's effort and chance determine the outcome.
4. Payments are made according to the proposed wage schedule.

A Principal-Agent Model

Consider the following simplified model:

- A task has two possible monetary outcomes: $\{\underline{q}, \overline{q}\}$ with $\underline{q} < \overline{q}$.
- Agent can choose one of two effort levels: $\{e_1, e_2\}$ with $e_1 < e_2$.
- The probability of the high output given effort e_i is:

$$p_i = \Pr(q = \overline{q} | e_i)$$

- Assume that $p_1 < p_2$ – ie more effort $\Rightarrow$ better outcomes.
- Principal chooses a wage schedule w.
- Agent is risk averse and his preferences are:

$$U(w, e) = E[u(w)] - c(e)$$

- Principal is risk neutral and his preferences are:

$$V(w) = E[q - w]$$

Observable Actions

Principal-Agent Model: Complete Info I

Begin by looking at the complete information benchmark:

- Principal can observe the effort chosen by Agent.
- Principal picks a wage schedule w_i that depends on Agent's effort.
- Agent's reservation utility is $\underline{u}$ – utility from resigning.
- The **participation constraint** of Agent requires:

$$U(w_i, e_i) = u(w_i) - c(e_i) \geq \underline{u}$$

- By picking wages appropriately Principal chooses e_i and w_i.
- The problem of Principal thus is to:

$$\max_{e_i, w_i} E[q|e_i] - w_i + \lambda[u(w_i) - c(e_i) - \underline{u}]$$

- Recall that $E[q|e_i] = p_i\overline{q} + (1 - p_i)\underline{q}$.

Principal-Agent Model: Complete Info II

- Recall the problem of Principal:

$$\max_{e_i, w_i} E[q|e_i] - w_i + \lambda[u(w_i) - c(e_i) - \underline{u}]$$

- The lowest wage v_i that induces effort e_i from Agent is:

$$u(v_i) - c(e_i) = \underline{u}$$

- Thus Principal chooses to induce effort e_* if and only if:

$$e_* \in \arg\max_{e_i \in \{e_1, e_2\}} E[q|e_i] - v_i$$

- Principal then induces such effort choice by offering wages:

$$w_{i*} = \begin{cases} v_i & \text{if} \quad e_i = e_* \\ v_i - \varepsilon & \text{if} \quad e_i \neq e_* \end{cases}$$

- Complete info implies that FOC for the wage requires $MC = Price$:

$$1/u_w(w_i) = \lambda$$

Hidden Actions

Principal-Agent Model: Incomplete Info I

Next consider the case in which effort is unobservable for Principal.

If so, Principal can only condition wage $w(q)$ on outcome q:

$$w(q) = \begin{cases} \underline{w} & \text{if} \quad q = \underline{q} \\ \overline{w} & \text{if} \quad q = \overline{q} \end{cases}$$

If Principal prefers Agent to exert high effort e_2:

- Agent's **participation constraint** at e_2 requires:

$$p_2 u(\overline{w}) + (1 - p_2) u(\underline{w}) - c(e_2) \geq \underline{u}$$

- Agent's **incentive constraint** guarantees that they pick high effort:

$$p_2 u(\overline{w}) + (1 - p_2) u(\underline{w}) - c(e_2) \geq p_1 u(\overline{w}) + (1 - p_1) u(\underline{w}) - c(e_1)$$

Principal-Agent Model: Incomplete Info II

It is easy to rewrite IC and PC as follows:

$$\begin{array}{ll}(PC) & p_2 u(\overline{w}) + (1-p_2)u(\underline{w}) \geq c(e_2) + \underline{u} \\ (IC) & (p_2 - p_1)(u(\overline{w}) - u(\underline{w})) \geq c(e_2) - c(e_1)\end{array}$$

Graphically, the constraints amount to

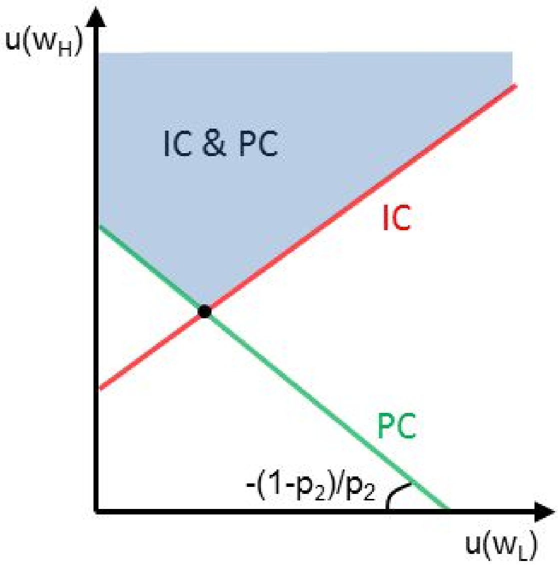

Principal-Agent Model: Incomplete Info III

Formally, Principal who wants Agent to exert e_2:

- maximizes its profits by choosing $w(q)$ subject to:
  1. Agent's Participation Constraint at e_2
     [Agent prefers to exert high effort than to resign]
  2. Agent's Incentive Constraint
     [Agent prefers to exert high effort than low effort]
- the Lagrangian of this problem amounts to

$$\max_{\underline{w},\overline{w}} \quad p_2[\overline{q} - \overline{w}] + (1-p_2)[\underline{q} - \underline{w}]$$
$$+\lambda[p_2 u(\overline{w}) + (1-p_2)u(\underline{w}) - c(e_2) - \underline{u}]$$
$$+\mu[(p_2 - p_1)(u(\overline{w}) - u(\underline{w})) - c(e_2) + c(e_1)]$$

Principal-Agent Model: Incomplete Info IV

Writing out Lagrangian explicitly the Principal's problem becomes:

$$\begin{aligned}\max_{\underline{w},\overline{w}} \quad & p_2[\overline{q}-\overline{w}]+(1-p_2)[\underline{q}-\underline{w}] \\ & +\lambda[p_2 u(\overline{w})+(1-p_2)u(\underline{w})-c(e_2)-\underline{u}] \\ & +\mu[(p_2-p_1)(u(\overline{w})-u(\underline{w}))-c(e_2)+c(e_1)]\end{aligned}$$

First order conditions for this problem are:

$$\begin{aligned}-p_2+\lambda p_2 u_w(\overline{w})+\mu(p_2-p_1)u_w(\overline{w}) &= 0 \\ -(1-p_2)+\lambda(1-p_2)u_w(\underline{w})-\mu(p_2-p_1)u_w(\underline{w}) &= 0\end{aligned}$$

Solving for λ and μ we find that:

$$\lambda = \frac{p_2}{u_w(\overline{w})}+\frac{1-p_2}{u_w(\underline{w})} > 0$$

$$\mu\left[\frac{p_2-p_1}{p_2}\right] = (1-p_2)\left[\frac{1}{u_w(\overline{w})}-\frac{1}{u_w(\underline{w})}\right] > 0$$

Principal-Agent Model: Incomplete Info V

The previous argument establishes that:

1. both μ and λ are positive if u is increasing and concave;
2. the incentive constraint binds since $\mu > 0$;
3. the participation constraint binds since $\lambda > 0$.

Thus, wages $\overline{w}, \underline{w}$ are found by solving the two constraints IC & PC:

$$\begin{aligned}u(\underline{w}) &= \underline{u} - c(e_2)\frac{p_1}{p_2-p_1}+c(e_1)\frac{p_2}{p_2-p_1} \\ u(\overline{w}) &= \underline{u} + c(e_2)\frac{1-p_1}{p_2-p_1}-c(e_1)\frac{1-p_2}{p_2-p_1}\end{aligned}$$

As firms prefer low wages, PC binds; as agents are risk averse, so does IC.

Principal-Agent Model: Incomplete Info VI

If Principal wants Agent to exert e_1:

- it pays v_1 with certainty;
- at such wage worker chooses e_1;
- Principal's profit amounts to:

$$p_1\overline{q} + (1 - p_1)\underline{q} - v_0.$$

To conclude the principal chooses to enforce effort e_2 if

$$p_2[\overline{q} - \overline{w}] + (1 - p_2)[\underline{q} - \underline{w}] > p_1\overline{q} + (1 - p_1)\underline{q} - v_0.$$

Principal-Agent Model: Key Ideas

With incomplete information:

- Agent is not fully insured.
- Volatility in payoffs is required for him to exert effort.
- As actions are unobservable compensation must increase with output.

If Principal was more risk averse than the agent:

- It would sell the company to Agent.
- It would face no risk and extract the full surplus.

Principal-Agent Model: Example I

Example: $e \in \{0,1\}$, $u(w,e) = 2w^{1/2} - e$, $\underline{u} = 1$,
$\overline{q} = 4$, $\underline{q} = 0$, $p_1 = 3/4$, $p_0 = 1/4$

Complete Info: what are w_1, w_0, e^*?

- Wages w_1 and w_0 are found by PC(e):

$$\begin{aligned} 2w_1^{1/2} - 1 &= 1 \Rightarrow w_1 = 1 \\ 2w_0^{1/2} - 0 &= 1 \Rightarrow w_0 = 1/4 \end{aligned}$$

- Optimal effort $e_* = 1$ is found by comparing profits:

$$\begin{aligned} \frac{3}{4}\overline{q} + \frac{1}{4}\underline{q} - w_1 &> \frac{1}{4}\overline{q} + \frac{3}{4}\underline{q} - w_0 \\ \frac{3}{4}4 - 1 = 2 &> \frac{3}{4} = \frac{1}{4}4 - \frac{1}{4} \end{aligned}$$

- Thus the agent is fully insured by the principal

Principal-Agent Model: Example II

Incomplete Info: what are $\overline{w}$, $\underline{w}$, if principal wants $e_* = 1$?

- Wages $\overline{w} = 25/16$ and $\underline{w} = 1/16$ are found by solving PC(1) and IC:

$$\begin{aligned} \frac{3}{4}(2\overline{w}^{1/2} - 1) + \frac{1}{4}(2\underline{w}^{1/2} - 1) &= 1 \\ \frac{3}{4}(2\overline{w}^{1/2} - 1) + \frac{1}{4}(2\underline{w}^{1/2} - 1) &= \frac{1}{4}(2\overline{w}^{1/2}) + \frac{3}{4}(2\underline{w}^{1/2}) \end{aligned}$$

- If principal wants $e_* = 0$, a wage $w_* = 1/4$ satisfying PC(0) suffices:

$$2w_*^{1/2} - 0 = 1$$

- The principal, however, prefers $e_* = 1$ since:

$$\begin{aligned} \frac{3}{4}(\overline{q} - \overline{w}) + \frac{1}{4}(\underline{q} - \underline{w}) &> \frac{1}{4}\overline{q} + \frac{3}{4}\underline{q} - w_* \\ \frac{3}{4}(4 - \frac{25}{16}) + \frac{1}{4}(-\frac{1}{16}) = \frac{29}{16} &> \frac{3}{4} = \frac{1}{4}4 - \frac{1}{4} \end{aligned}$$

- The principal cannot fully insure the agent with incomplete information since it would undermine the incentives to exert effort.

Extra: Competitive Insurance

Intermediate Microeconomics – Summer School

Francesco Nava

London School of Economics

June 2019

Competition in Insurance Markets

Insurance Markets

Consider the following economy:

- Individuals have two types $\{H, L\}$.
- The fraction of individuals of type t is π_t.
- Any individual can be healthy or sick.
- The probability of type t being sick is σ_t.
- Assume that $\sigma_H > \sigma_L$.
- The income of an individual is Y if healthy and $Y - K$ if sick.
- Let y_t denote the consumption of type t if healthy & x_t if sick.
- Preference of type t satisfy:

$$\sigma_t u(x_t) + (1 - \sigma_t) u(y_t)$$

Insurance Markets

- The insurance market is competitive (free entry).
- Consumers can buy insurance coverage $z_t \in [0, K]$...
  ... at a unit price p_t [ie total premium $p_t z_t$].
- If they do so their consumption in the two states becomes:

$$\begin{aligned} y_t &= Y - p_t z_t \\ x_t &= Y - K - p_t z_t + z_t = Y - K + (1 - p_t) z_t \end{aligned}$$

- If so the problem of a consumer becomes:

$$\max_{z_t \in [0,K]} \sigma_t u(x_t) + (1 - \sigma_t) u(y_t)$$

- Thus, FOC with respect to z_t requires for type t:

$$\sigma_t (1 - p_t) u'(x_t) = (1 - \sigma_t) p_t u'(y_t)$$

Insurance Markets

- FOC can be written in terms of MRS as:

$$\frac{u'(x_t)}{u'(y_t)} = \frac{1-\sigma_t}{\sigma_t}\frac{p_t}{1-p_t}$$

- Thus a consumer of type t wants:

$$\begin{array}{lll} \text{Full Insurance:} & z_t = K & \text{if} \quad p_t = \sigma_t \\ \text{Under Insurance:} & z_t < K & \text{if} \quad p_t > \sigma_t \\ \text{Over Insurance:} & z_t > K & \text{if} \quad p_t < \sigma_t \end{array}$$

- The profits of an insurance company are given by:

$$\textstyle\sum_t \pi_t z_t (p_t - \sigma_t)$$

  thus a company does not run a loss provided that $p_t \geq \sigma_t$.

Competition in Insurance Markets: Full Info

Assume that insurance companies can distinguish the two types.

If so, the companies set a different price for each type.

Since the markets are perfectly competitive insurance companies:

- Offer price $p_t = \sigma_t$ to type $t \in \{H, L\}$.
- At such prices all consumers fully insure.
- And each firm makes zero profits.

No entrant could benefit from offering competing policies.

Competition in Insurance Markets: Incomplete Info

If insurance companies cannot distinguish the two type:

- Offering the complete information contracts is suboptimal... ... as all players claim to be of type L to pay $p_L = \sigma_L < p_H$.
- This cannot be optimal for a firm since it would run a loss:

$$\pi_H K(p_L - \sigma_H) + \pi_L K(p_L - \sigma_L) < 0$$

- Alternatively a firm may not attempt to distinguish consumers... but may offer a price that would lead to break-even if all fully insure:

$$p = \pi_H \sigma_H + \pi_L \sigma_L$$

- If so, low risk type L wants to under insure as $\sigma_L < p$ and... high risk type H wants over insurance as $\sigma_H > p$ and picks $z_H = K$.

Competition in Insurance Markets: No Pooling

If, however, different types respond to p as detailed above, the firm:

- can tell types apart as only type H buys full insurance;
- prefers to raise prices on those individuals to σ_H.

Consumer H thus prefers to mimic type L by buying at price p as many units as type L, that is z_L.

If so, the firm profits by offering a policy $(\underline{p}, \underline{z})$ such that:

- is preferred by type L consumers but not by type H
- with a lower price $\underline{p} \in (\sigma_L, p)$ and a lower quantity $\underline{z} < z_L$.

Competition in Insurance Markets: No Pooling

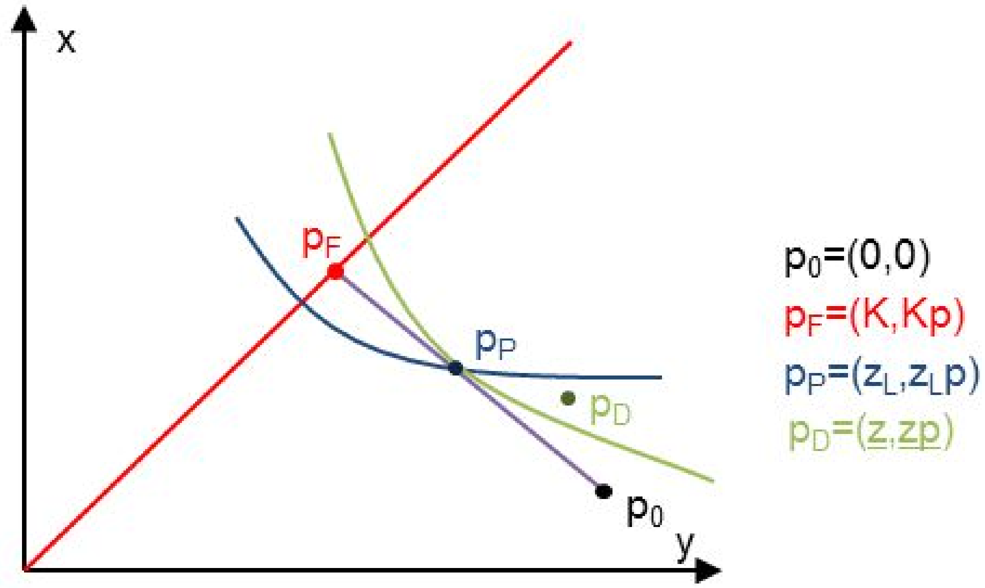

Theorem (No Pooling)

There is no pooling equilibrium in a competitive insurance market.

Insurance Markets: Separating Equilibria may Not Exist

Thus firms must offer separating contracts if an equilibrium is to exist:

- Consider $(p_H, z_H) = (\sigma_H, K)$ and $(p_L, z_L) = (\sigma_L, w)$
- For players of type H to choose (p_H, z_H), IC requires:

$$u(Y - \sigma_H K) \geq \sigma_H u(Y - K + (1 - \sigma_L)w) + (1 - \sigma_H)u(Y - \sigma_L w)$$

- For players of type L to choose (p_L, z_L), IC requires:

$$\sigma_L u(Y - K + (1 - \sigma_L)w) + (1 - \sigma_L)u(Y - \sigma_L w) \geq u(Y - \sigma_H K)$$

- PROBLEM: if π_L is high enough both contracts are dominated... ... by pooling contract $(p', z') = (p + \varepsilon, K - \varepsilon)$.
- If so a competitive insurance market may have no equilibrium.
- Cause: Profits from each type depend directly on hidden info!

Insurance Markets: Separating Equilibria may Not Exist

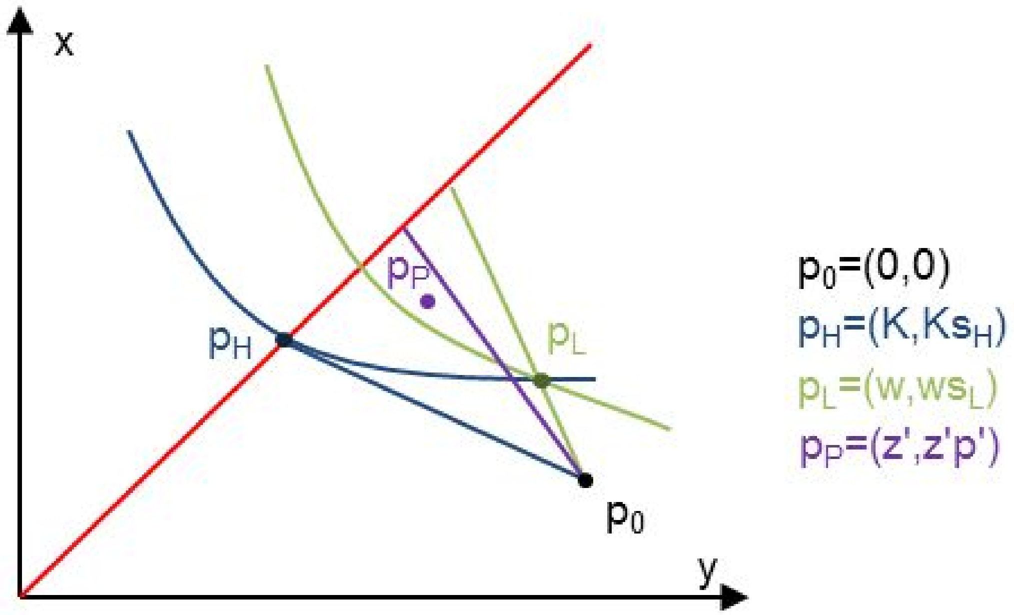

Theorem (No Equilibrium)

No equilibrium may exist in a competitive insurance market.

Insurance Markets: Separating Equilibria may Not Exist

The magenta region (left plot) identifies the pooling contracts that are profitable if purchased by both types and that are accepted by both types:

- if such region is non-empty (left plot) no equilibrium exists;
- if the region is empty (right plot) a separating equilibrium exists.

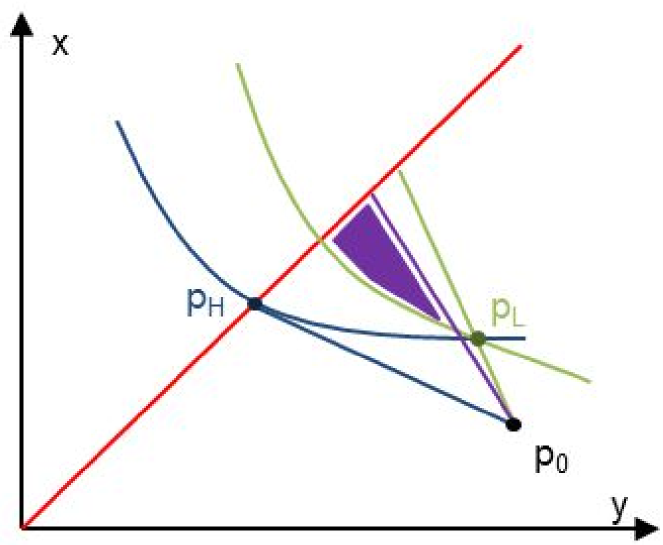

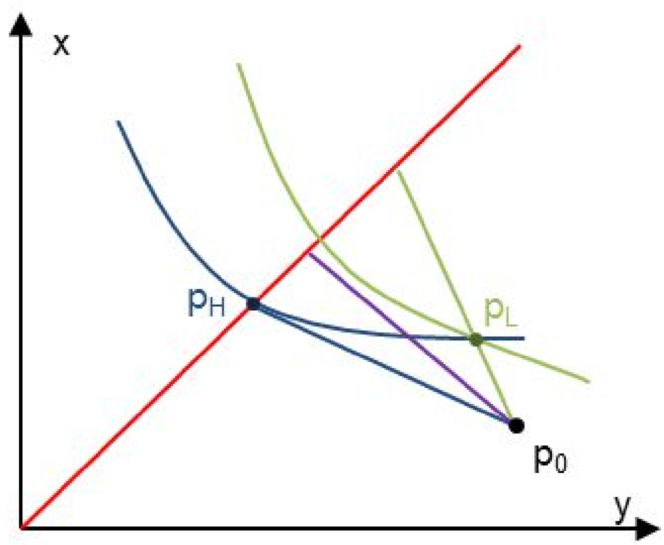

EC201 - Intermediate Microeconomics
SUMMER SCHOOL 2019 – FRANCESCO NAVA

Homework Assignments

The list of daily assignments follows. I suggest that you attempt at least some of them prior to the class. Numbers refer to exercises from the textbook. The problems labeled "extra" are not required, but are good practice.

1. Due Thursday June 28
   - Problem 1 (below)
   - Normal: 7.1, 7.4, 7.7
   - Extra: 7.3, 7.10
2. Due Friday June 29
   - Problem 2 (below)
   - Normal: 8.1, 8.2, 8.4
   - Extra: 8.9
3. Due Monday July 2
   - Problem 3 (below)
   - Normal: 15.1, 15.2, 8.7(a)
   - Extra: 8.5, Problem 4 (below)
4. Due Tuesday July 3
   - **Hand-In Problem Set**
   - 8.3, Problem 5 (below)
   - Normal: 8.6, 15.7(a)
5. Due Wednesday July 4
   - Problem 6 (below)
   - Normal: 18.2, 18.3, 18.4
   - Extra: 18.5 (a-b)
6. Due Thursday July 5
   - Problem 7 and 8 (below)
   - Normal: 18.7
   - Extra: 18.1

Problem 1 (Uncertainty) Rick is considering whether to spend 5 dollars betting on Republicans winning the next election. If Republicans were to win the election, Rick would be paid 4 dollars for any dollar that he has bet. The utility that Rick derives from a (positive or negative) cash transfer of x dollars is determined by the following utility function,

$$u(x) = (475 + 75x)^{1/2}.$$

Rick believes that the probability of republicans winning the next election is 1/3.

1. Find the expected value of such a lottery.

2. Find Rick's expected utility of taking such a gamble. Would he accept it? Or would he reject it and get $x = 0$?
3. What's the certainty equivalent of such a lottery.

Problem 2 (Static Games) Consider the following static two-player game:

$1\backslash 2$	A	B
A	$1, 2$	$3, 7$
B	$7, 3$	$2, 2$

Player 1 is the row player, and his payoff is the first to appear in each entry. Player 2 is the column player and his payoff is the second to appear in each entry.

1. Find the pure strategy Nash equilibria of the game, and show that they are equilibria.
2. Find the mixed strategy Nash equilibrium of the game.
3. Derive the mixed strategy best responses.

Problem 3 (Bayesian Games) Consider the following Bayesian game played by two players (1 and 2) who are deciding whether to cooperate, C, or defect, D. Two states are possible, *Good* and *Bad*. Suppose that Player 2 knows the state, while Player 1 thinks that the state is *Good* with probability p. Payoffs in each state respectively satisfy

State *Good*:

$1\backslash 2$	C	D
C	$0, 0$	$1, 1$
D	$1, 1$	$0, 0$

State *Bad*:

$1\backslash 2$	C	D
C	$0, 0$	$0, 1$
D	$1, 0$	$3, 3$

Player 1 is the row player, and his payoff is the first to appear in each entry. Player 2 is the column player and his payoff is the second to appear in each entry.

1. What is the set of possible strategies for the two players in this game?
2. Find the pure strategy Bayes-Nash equilibria for all values of $p \in (0, 1)$.

Problem 4 (Cournot Uncertainty) Two firms compete to sell a good. Firm 1 has total costs of production $C_1(q_1) = (q_1)^2 + 2q_1$ and its costs are known to Firm 2. The total costs of Firm 2 depends on its type. If Firm 2 is of type L, its costs are $C_L(q_L) = 2q_L$. If Firm 2 is of type H, its costs are $C_H(q_H) = 2(q_H)^2$. Firm 2 knows its type. But Firm 1 only knows that Firm 2 can have either cost structure with equal probability. The inverse demand for the output produced by the two firms in this market satisfies:

$$p(q_1 + q_2) = 10 - 2(q_1 + q_2)$$

Firms choose how much output to produce in order to maximize their profits. Find the Bayes-Nash equilibrium of this game. Characterize the equilibrium output strategies for both firms. Find the market price for each of the two possible cost configurations.

Problem 5 (Repeated Games) Consider the following asymmetric Prisoner's Dilemma:

$1\backslash 2$	C	D
C	$3, 4$	$1, 6$
D	$4, 0$	$2, 2$

1. Find the minmax values of this game.
2. Then, consider the following "trigger" strategy: any player chooses C provided that no player ever played D; otherwise any player chooses D. Write the two incentive constraints that if satisfied would make such a strategy a NE. Then, write the two additional incentive constraints that if satisfied would make such a strategy a SPE. What is the lowest discount rate for which such strategy satisfies all the constraints.

Problem 6 (Adverse Selection) Consider an economy with a monopolistic electricity supplier. Assume that the costs of producing a unit of electricity are 1$. There are only two goods in this economy namely money, y, and electricity, x. All consumers in this economy are endowed with 100$ in money and no electricity. There are two types of buyers in the economy: type H has high value for electricity, while type L does not. In particular assume that preferences satisfy:

$$\begin{aligned} u(x,y|H) &= 8x^{1/2} + y \\ u(x,y|L) &= (9/2)x^{1/2} + y \end{aligned}$$

1. If the monopolist can recognize the type of any individual, find the optimal bundles sold to both types. Why is this outcome efficient?
2. Suppose that 1/8 of all individuals in the population are of type H. If the monopolist cannot recognize the type of any individual, find the optimal bundles sold to both types in equilibrium. Why is the outcome inefficient?

Problem 7 (Signaling) Consider Spence's signalling model. A worker's type is $t \in \{0,1\}$. The probability that any worker is of type $t = 1$ is equal to 2/3, while the probability that $t = 0$ is equal to 1/3. The productivity of a worker in a job is $(t+1)^2$. Each worker chooses a level of education $e \geq 0$. The total cost of obtaining education level e is $C(e|t) = e^2(2-t)$. The worker's wage is equal to his expected productivity.

1. Find pooling equilibrium education levels.
2. Find separating equilibrium education levels.

Problem 8 (Moral Hazard) Consider the Principal-Agent model discussed in the slides. Suppose that the effort exerted by the agent can take one of three values $e \in \{1/3, 2/3\}$. Also suppose that the Agent's preferences are given by $u(w,e) = 2w^{1/2} - e$. Leisure yields to the Agent a reservation utility $\underline{u} = 1$. The principal's problem can have one of two outcomes: success or failure $\{\overline{q}, \underline{q}\}$. The payoffs to the Principal in these two events are: $\overline{q} = 4$ if the outcome is a success and $\underline{q} = 0$ if the outcome is a failure. The probability of a success is: $\overline{p}_{1/3} = 1/3$ if the Agent chooses the low effort; and $\overline{p}_{2/3} = 2/3$ if the agent chooses the high effort.

1. Let effort be observable. Compute the full-information wages at each effort level. What is the profit maximizing effort for the Principal?
2. Now suppose that the Principal cannot observe effort. For each effort level find output dependent wages that induce the Agent to exert such an effort.
3. Which effort level maximizes the profits of the principal if he cannot observe effort? Which wage schedule should he set to induce the agent to exert such effort.

EC201 Hand-In Problem Set

Short Questions (17 MARKS EACH)

1. Ann is deciding whether to bet on a tennis match. A friend offers to give her 20 dollars if the lower ranked player wins, while she has to pay him 12 dollars otherwise. The utility that she derives from a (positive or negative) cash transfer of x dollars is determined by the following utility function,

$$u(x) = (16 + x)^{1/2}.$$

   Ann believes that the probability of the lower ranked player winning the match is p.

   (a) Find the expected value of this lottery. For what values of p is the expected value positive? (5 marks)

   (b) Find Ann's expected utility when betting on the match. For what values of p would she accept the bet? (6 marks)

   (c) Find Ann's certainty equivalent for this lottery when $p = 3/4$. (6 marks)

2. Consider the following Bayesian game played by two players 1 and 2. Two states are possible, A and B. Suppose that player 2 knows state, while player 1 deems both states equally likely. Payoffs in each state respectively satisfy

State A:

$1\backslash 2$	s	p
s	$1,1$	$0,2$
p	$0,0$	$2,0$

State B:

$1\backslash 2$	s	p
s	$0,2$	$1,1$
p	$2,0$	$0,0$

.

   Player 1 is the row player, and his payoff is the first to appear in each entry. Player 2 is the column player and his payoff is the second to appear in each entry.

   (a) What is the set of possible strategies for either player in this game? (7 marks)

   (b) Find a pure strategy Bayes Nash equilibrium of the game. (10 marks)

3. Suppose that two friends have split 5 indivisible cookies according to the following protocol. Player 1 gets to divide the cookies between two dishes and Player 2 gets to choose which of the two dishes to consume (while the remaining dish is consumed by Player 1). First, assume that all cookies are alike, so that both players value every cookie equally, and care only about consuming more cookies.

   (a) Set up this problem as an extensive form game. [4 marks]

   (b) Find a Subgame Perfect equilibrium of this game. Write the behavioral strategy for both players, and check that no deviation is profitable. [6 marks]

   (c) Does the game possess any Nash equilibrium that is not Subgame Perfect? [7 marks]

LONG QUESTION (49 Marks)

4. Two firms compete to sell a good. Total costs of the two firms respectively satisfy

$$C_1(q_1) = 2q_1 \quad \text{and} \quad C_2(q_2) = q_2^2.$$

The total output produced in the economy is $Q = q_1 + q_2$. The inverse demand for the total output produced by the two firms in this market satisfies

$$p(Q) = \begin{cases} 10 - 2Q & \text{if} \quad Q \leq 5 \\ 0 & \text{if} \quad Q > 5 \end{cases}.$$

Remember that the inverse demand curve identifies the highest price for which all the units supplied to the market are purchased.

(a) First assume that firms compete on quantities. Find the Cournot equilibrium output at the two firms, the equilibrium price, and the profits at each of the two firms. (20 marks)

(b) Then suppose that the two firms form a cartel. Find the cartel price and the output produced at the two firms? Compare your results with part (a). Does the cartel make higher profits than the two firms in part (a)? Does it produce more output on aggregate? (15 marks)

(c) Finally assume that firms choose their output taking prices as given. Find the perfect competition price and the output produced at the two firms? Again, compare your results with parts (a) and (b). (14 marks)