A (micro) course in microeconomic theory for MSc students

Gaudeul, Alexia

22 May 2009

A (Micro) Course in Microeconomic Theory for MSc Students

Alexia Gaudeul*

May 23, 2009

Abstract

Those lecture notes cover the basics of a course in microeconomic theory for MSc students in Economics. They were developed over five years of teaching MSc Economic Theory I in the School of Economics at the University of East Anglia in Norwich, UK. The lectures differ from the standard fare in their emphasis on utility theory and its alternatives. A wide variety of exercises for every sections of the course are provided, along with detailed answers. Credit is due to my students for 'debugging' this material over the years. Specific credit for some of the material is given where appropriate.

JEL codes: A1, A23, D0.

Keywords: Economics, Microeconomics; Utility Theory; Game Theory; Incentive Theory; Online Textbook; Lecture Notes, Study Guide, MSc.

*Please direct comments and report mistakes to me at a.gaudeul@gmail.com. My website is at http://agaudeul.free.fr.

1 Programme

The programme of this course is divided in three parts; choice under uncertainty, game theory and incentive theory. The whole of the course can be covered in sixteen hours of teaching, along with eight hours of workshops, over eight weeks. This is an intensive program that is designed both to cover the basics in each area and progress quickly to more advanced topics. The course is thus accessible to students with little background in economics, but should also challenge more advanced students who can focus on the later sections in each parts and concentrate on the suggested readings. Exercises covering each part increase gradually in difficulty and are often of theoretical interest on their own. Detailed answers are provided.

This course may be complemented with lectures on consumer and firm theory, and on general equilibrium concepts. Those are not covered in those notes.

The main concepts that are covered in each part are listed below:

Choice under uncertainty (2×2 hours lectures, 1×2 hours workshop): Rationality and axiomatic theories of choice. Expected utility theory: its foundations and applications. Experimental tests of expected utility theory. Alternatives to EUT: prospect theory, rank dependent expected utility theory, cumulative prospect theory, regret theory.

A large part of this topic is not covered in standard text books and this topic thus requires independent reading.

Game Theory (4×2 hours lectures, 2×2 hours workshops): Normal (or strategic)-Form Games. Dominance, Nash equilibrium, and mixed-strategy extensions. Extensive-Form Games. Information sets, subgame-perfect equilibria, and backward induction. Incomplete-information games, Bayesian Games, Bayesian-Nash equilibria, forward induction. Purification of mixed strategy equilibria. Auctions. Repeated games and the Folk theorem.

Incentive Theory (2×2 hours lecture, 1×2 hours workshop): Principal-agent models. Adverse selection (hidden information). Moral hazard (hidden action). Second best and efficiency. Liability constraints. Risk aversion. Commitment. Repeated adverse selection. Hold up problem. Informed principal. Revelation principle. Direct Revelation Mechanism.

2 Textbooks:

The main texts for the course are:

- Kreps D.M., 1990, A course in Microeconomic Theory, FT/Prentice Hall (hereafter ‘Kreps’).
- Varian H.R., 1992, Microeconomic Analysis, Norton, 3d edition (hereafter ‘Varian’).

- Mas-Colell A., M.D. Whinston and J.R. Green, 1995, Microeconomic Theory, Oxford University Press (hereafter 'Mas-Colell').

  All of these cover, more or less, the entire module, except for choice under uncertainty. Each has its strengths, though these are different. Mas-Colell is the most up-to-date and comprehensive, but also the most advanced. It is advised for those with a good background in microeconomics and with good mathematical skills, and for those who wish to pursue further studies in economics.

Contents

List of Figures

Part I

Choice under uncertainty

3 Introduction

This section will give tools to think about choice under uncertainty. A model of behavior when faced with choices among risky lotteries will be presented – The von Neumann-Morgenstern ("vNM") Expected Utility Theory ("EUT") (1944).[1] The consequences of EUT in terms of prescribed or predicted behavior will be analyzed, and this will be compared with empirical and experimental data on the behavior of agents. That data will be shown not to conform with EUT in some cases, which means EUT may not be an adequate descriptive model of behavior. Alternative models that take better account of the reality of the patterns of decision making of economic agents will thus be presented and discussed.

4 Readings

4.1 Textbook Readings

- Kreps, Ch. 3
- Varian, Ch. 11
- Mas-Colell, Ch. 6.

4.2 Other general readings

The following are useful survey papers, though note that they partly overlap:

- Machina M., 1987, Choice under uncertainty: problems solved and unsolved, Journal of Economic Perspectives 1(1), 121-154.
- Starmer C., 2000, Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk, Journal of Economic Literature 38(2), 332-382.

The following website at The Economics New School is well designed and informative:

- Choice under risk and uncertainty, http://cepa.newschool.edu/het/essays/uncert/choicecont.htm

The following articles are well written and motivating:

- Chakrabortty A., Why we buy what we buy, The Guardian, May 20th 2008.
- Do economists need brains?, The Economist, July 24th 2008.

[1] von Neumann J. and O. Morgenstern, 1944, Theory of Games and Economic Behavior, Princeton University Press.

4.3 Articles[2]

- Bateman I.J., A. Munro, B. Rhodes, C. Starmer and R. Sugden, 1997, A Test of the Theory of Reference Dependent Preferences, Quarterly Journal of Economics 112(2), 479-505.
- Camerer C., 1995, Individual Decision-Making, in J.H. Kagel and A. Roth, Handbook of Experimental Economics, Princeton University Press, Ch. 8, esp. pp. 617-65.
- Hey J. and C. Orme, 1994, Investigating Generalizations Of Expected Utility-Theory Using Experimental-Data, Econometrica 62(6), 1291-1326.
- Kahneman D. and A. Tversky, 1979, Prospect theory: an analysis of decision under risk, Econometrica 47(2), 263-291.
- Kahneman D. and A. Tversky, 1999, Choices, values and Frames, Cambridge University Press.
- Machina M., 1982, Expected Utility analysis without the independence axiom, Econometrica 50, 277-323.
- Rabin M., 2000, Risk Aversion and Expected-Utility Theory: A Calibration Theorem, Econometrica 68(5), 1281-1292.
- Shefrin H. and M. Statman, 1985, The Disposition to Sell Winners Too Early and Ride Losers Too Long - Theory and Evidence, Journal of Finance 40(3), 777-790.
- Shogren J.F., S.Y. Shin, D.J. Hayes and J.B. Kliebenstein, 1994, Resolving differences in willingness to pay and willingness to accept, American Economic Review 84(1), 255-270.
- Sugden R., 1985, New developments in the theory of choice under uncertainty, Bulletin of Economic Research 38(1), 1-24.
- Sugden R., 2002, Alternatives to Expected Utility: Foundations, in P.J. Hammond, S. Barberá and C. Seidl (eds.), Handbook of Utility Theory Vol. II, Kluwer: Boston.
- Tversky A. and D. Kahneman, 1991, Loss aversion and riskless choice: a reference dependent model, Quarterly Journal of Economics 106(4), 1039-1061.

5 Basic tools and notations

5.1 The objects of preference and choice

'Lotteries' or 'gambles' or 'prospects' are situations where the **outcome** of one's **action** are **uncertain**.

[2]Some articles in this list were contributed by previous teachers in MSc Economic Theory 1 at the UEA.

Actions can consist in buying, selling, going out, staying in, buying an umbrella, etc.

Uncertainty may be due to the incomplete information about the action of others, the inability to predict complex events (weather), limitations in one's capacity to process complex data, etc.

Outcomes are defined in terms of utility. Outcomes will be influenced by your choice of action and the realization of random events on which you have no influence. For example, a chicken that decides to cross a road will either die if a car happens to come by, or live if no car happens to come by. We assign 'utility' to those two possible outcomes. At its most basic, utility is defined in terms of preferences, for example, do you prefer sun or rain?

The agent is supposed to know what acts she can choose, i.e. she knows her options. She is also supposed to know the set of all possible 'state of nature' that may prevail, as well as be able to evaluate the probability, objective or subjective, of the occurrence of each possible state of nature. She is also supposed to be able to evaluate the utility of all possible consequences of each of her actions (outcomes). Consequence functions are defined as what happens if you chose an act and such or such state of nature occurred.
Consider for example a chicken faced with the decision whether to cross the road. The consequence function of crossing the road combined with the event that a car is on the road is what happens in that case, i.e. the chicken is run over and dies. This can be denoted as follows:

$$c(\underbrace{\text{car on the road}}_{\text{event}},\underbrace{\text{cross the road}}_{\text{action}}) = \underbrace{\text{death}}_{\text{outcome}}$$

Under this setting, you are supposed to know the utility of the outcome of your whole set of action under any possible circumstances, e.g. the chicken is supposed to be able to anticipate any possible event that may occur when he crosses the road, and the outcome of his action under any of those events. You know this in advance even though you may never have experienced such combination of circumstances.

Obviously, your choice would be greatly simplified if a function allowed you to calculate the expected utility of an action based on the expected probability of the outcomes that will result from that action and an evaluation of the utility of each of those outcomes, so as to obtain automatically your expected utility from an action.

Why would you want to predict the utility of an action? This is because you are constantly having to make decisions under uncertainty, and want to take the decision that will maximize utility. We will see in the following that subject to some assumptions on how you make basic decisions and rank the utility of events, then one obtain a very simple way to evaluate the utility of any outcome. Let us first present what is a lottery and how they can be represented graphically.

5.2 Lotteries

Lotteries are diagrammatic representations of the choice you are facing. Simple lotteries assign probabilities to outcomes, which can be numerical, monetary or other (e.g. a trip to the moon or a teddy bear).

1. Consider one agent A, for example, agent A is 'chicken'.
2. $(1, ..., s, ..., S)$ is the set of states of nature, for example (car on the road, no car on the road)
3. $(1, ..., x, ..., X)$ is the set of actions that are available to you, for example (cross the street, don't cross the street)
4. $c(s, x)$ is a consequence function or outcomes, which is dependend both on the state of nature s and your action x. For example, $c($ car on the road, cross the street$)$ is 'death', as explained above.
5. $p(s)$ is a probability function, that gives out the probability of each state of nature. For example, $p($ car on the road$) = 0.7$.
6. Most of the time, one will take action x as given and denote $c(s, x)$ in short hand as c_s and $p(s)$ as p_s. Then, the lottery L that results from action x is defined by the vector $C = (c_1, ..., c_s, ..., c_S)$ of possible consequences of x and its vector $P = (p_1, ..., p_s, ..., p_S)$ of the probabilities of each consequences in C.
7. Most of the time, one will take the set of consequences as a given, so that lottery L will be denoted by its vector P of probabilities associated to each consequences.
8. Take the set of consequences $C = ($ death, life on this side of the road, life on the other side of the road $)$. Consider L_1 the lottery which results from the action 'cross the street' and L_2 the lottery which results from the action 'don't cross the street'. In short-hand, one can denote L_1 as $(0.7, 0, 0.3)$ and L_2 as $(0, 1, 0)$.
   (a) If the chicken prefers lottery L_1 to L_2, one will denote this as $L_1 \succ L_2$ and this guarantees the chicken crosses the road.
   (b) If the chicken is indifferent between lottery L_1 and lottery L_2, one will denote this as $L_1 \sim L_2$, and the chicken may or may not cross the road.
   (c) If the chicken considers the lottery L_1 to be at least as good as L_2, one will denote this as $L_1 \succeq L_2$ and say that the chicken has a weak preference for L_1. In that case, one does not know whether the chicken is indifferent between L_1 and L_2 or if it actually prefers L_1 over L_2. One knows however that the chicken does not prefer L_2 over L_1. From an observational point of view, if the chicken crosses the road, one knows that $L_1 \succeq L_2$ but nothing more.
9. If the chicken crosses the road, then that means that $L_1 \succeq L_2$. However, saying that $L_1 \succeq L_2$ does not give any explanation for the behavior of the chicken, but merely translates an observation of its behavior. I do not know why the chicken crossed the road, I merely know it did.

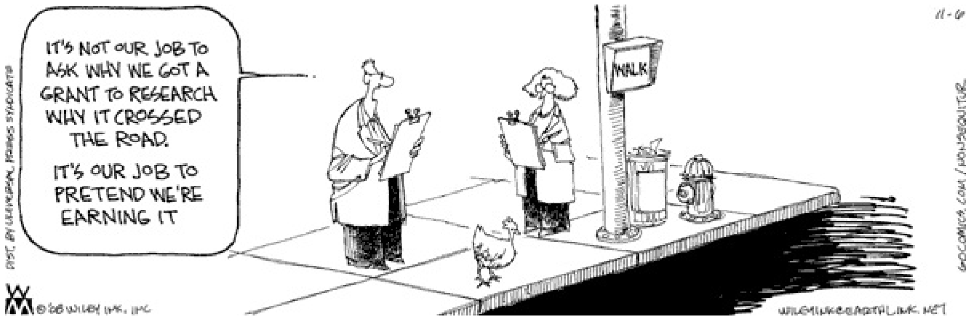

Figure 1: Why did the chicken cross the road? (c) Wiley Miller

5.3 The Marschak-Machina Triangle

Denote $C = \{A, B, C\}$ a set of consequences and $P = (p_A, p_B, p_C)$ a vector of probability defined over C. Note that one will always have $p_A + p_B + p_C = 1$, as an event will always occur ('no event' is itself an event...). The set P of all possible vectors of probabilities is depicted below in the *Marschak-Machina Triangle* ('M-M Triangle').[3]

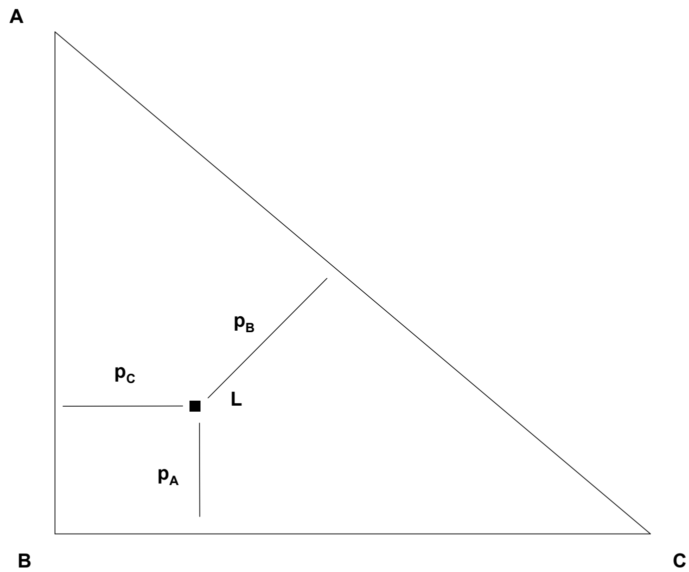

Figure 2: The Marschak-Machina Triangle

[3]For more details, read Machina M.J., 1987, Decision Making in the Presence of Risk, Science 236, 537-543. This expository device was introduced in Marschak J., 1950, Rational behavior, uncertain prospects, and measurable utility, Econometrica 18, 111-141.

The corners represent the certainty cases, i.e. A is such that $p_A = 1$, B is such that $p_B = 1$ and C is such that $p_C = 1$. A particular lottery, $L = (p_A, p_B, p_C)$ is represented as a point in the triangle. By convention, $c_A \succ c_B \succ c_C$: the top corner is preferred to the origin which is preferred to the right corner. Any lottery over three events can be represented. p_A is measured along the vertical axis, p_C is measured along the horizontal, leaving p_B to be measured from the other vertex. Thus at the origin, B, where p_A and p_B are both equal to 0, then $p_B = 1$. The hypotenuse is the range of lotteries for which $p_B = 0$, i.e. the hypotenuse is all lotteries that combine outcome A and C only.

We will be interested in the shape of utility indifference curves, which connect all lotteries over which the agent is indifferent. The Marschak-Machina triangle will allow us to show utility indifference curves in a probability space, represented by the triangle. The representation of lotteries in the Marschak-Machina triangle, as well as the representation of preferences over lotteries in this same M-M triangle, will repeatedly be used to illustrate theoretical proposition introduced in those lectures.

Exercises:

1. Represent the following three lotteries in the M-M triangle: $L_1 = (0, 1, 0)$, $L_2 = (0.5, 0, 0.5)$, $L_3 = (1/3, 1/3, 1/3)$
2. Represent the lotteries the chicken is facing when crossing the road, assuming the ordered set of consequences is (death, life on this side of the road, life on the other side of the road) and the probability a car is on the road is 0.7.
3. Represent one possible utility indifference curve representing the preferences of the chicken if it is found to cross the road. What properties must the utility indifference curve have?

6 Expected Utility Theory

6.1 The axioms of von Neumann and Morgenstern's Expected Utility Theory

von Neumann and Morgenstern ("vNM") Expected Utility Theory ("EUT") offers four axioms that together will guarantee that preferences over lotteries can be represented through a simple utility functional form.

6.1.1 Completeness

Completeness axiom: For all L_1, L_2, either $L_1 \succeq L_2$ or $L_2 \succeq L_1$

This axiom guarantees that preferences are defined over the whole set of possible lotteries. Graphically, in the M-M triangle, this axiom guarantees that any two points in the triangle are either on the same indifference curve or on two different curves. Indeed, according to the axiom, either the consumer is indifferent between L_1 and L_2, in which case both lotteries are on the same indifference curve, or they have a strict preference over L_1 and L_2, in which case they are on different indifference curves.

6.1.2 Transitivity

Transitivity axiom: For any L_1, L_2, L_3, if $L_1 \succeq L_2$ and $L_2 \succeq L_3$ then $L_1 \succeq L_3$

This axiom guarantees there are no cycles in preferences, i.e. a situation where I prefer bananas to apple, apple to oranges and oranges to bananas is not possible... Graphically, this axiom guarantees that indifference curves in the M-M triangle do not cross. Indeed, if there is intransitivity, then indifference curves must cross inside the triangle. For example, below, I represent indifference curves such that $L_1 \succ L_2$ and $L_2 \succ L_3$ (note that crossing outside the triangle is not a problem). Now, you can check that if I want to represent an indifference curve such that $L_3 \succ L_1$ (which is intransitive), then it must cross one or the other or both of the two previous indifference curves inside the triangle (Figure below). Conversely, if indifference curves cross inside the triangle, then there may be situations where transitivity does not hold.

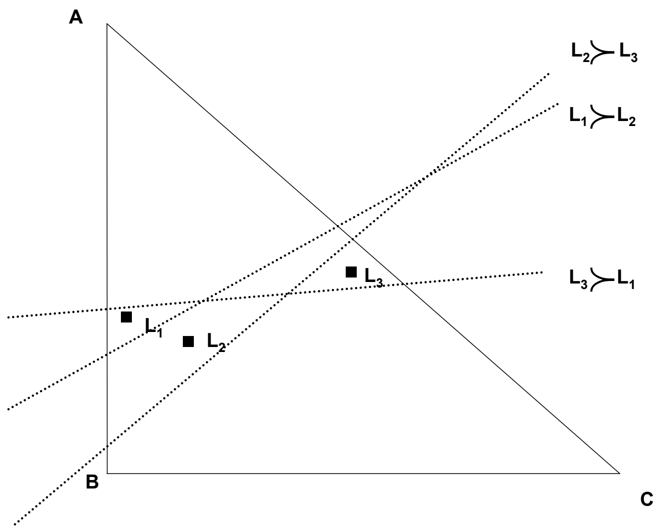

Figure 3: Violation of transitivity

6.1.3 Continuity or Archimedean axiom

Continuity axiom: For any $L_3 \succ L_2 \succ L_1$,
there exists a unique α, $0 \leq \alpha \leq 1$ such that $\alpha L_3 + (1-\alpha)L_1 \sim L_2$.

Uniqueness of α guarantees the indifference curves are continuous. This is because this axiom guarantees that any point in the triangle (any lottery) has an equivalent along the hypotenuse, and that this equivalent is unique. Suppose indeed that L_3 is the top corner of the triangle and L_1 the right corner of the M-M triangle, and consider L_2 any point in the triangle. The axiom says that for some α in $[0, 1]$, L_2 will be equivalent in terms of preferences to $\alpha L_3 + (1-\alpha)L_1$. But $\alpha L_3 + (1-\alpha)L_1$ is a point on the hypotenuse of the triangle. Therefore, any lottery has an equivalent along the hypotenuse, and conversely. This means that there are no spaces either on the hypotenuse or in the triangle that would have no equivalent (there are no 'jump' in the indifference curves).

6.1.4 Substitutability or Independence Axiom

Independence axiom: For any L_1, L_2 and L_3 such that $L_1 \succ L_2$, then for any $\alpha \in (0, 1)$, $(1-\alpha)L_1 + \alpha L_3 \succ (1-\alpha)L_2 + \alpha L_3$.

The independence axiom guarantees that indifference curves are parallel straight lines in the M-M triangle. Indeed, represent L_1, L_2 and L_3, three lotteries, in the M-M triangle. $(1-\alpha)L_1 + \alpha L_3$ and $(1-\alpha)L_2 + \alpha L_3$ are parallel translations of L_1 and L_2 into the M-M space, so a line that links L_1 and L_2 will be parallel to a line that links $(1-\alpha)L_1 + \alpha L_3$ and $(1-\alpha)L_2 + \alpha L_3$. Suppose $L_1 \sim L_2$ then I must have $(1-\alpha)L_1 + \alpha L_3 \sim (1-\alpha)L_2 + \alpha L_3$. Therefore, indifference curves will be parallel across the whole of the triangle (it is just a matter of choosing α and L_3) (Figure below).

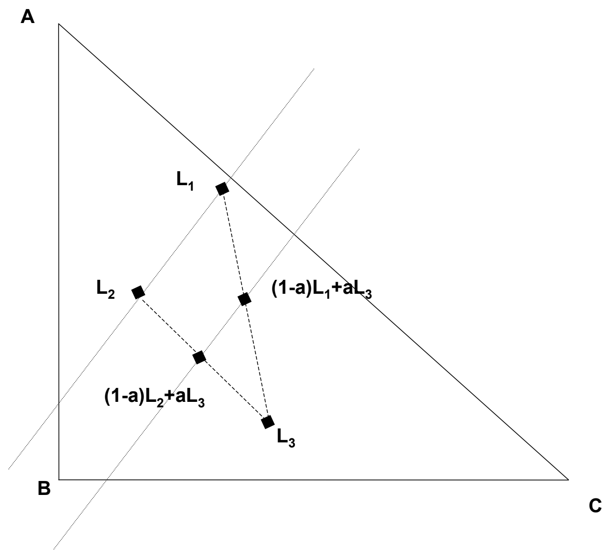

Figure 4: Independence axiom

Exercise: The following exercise adresses a common misconception over what *independence* means. Suppose you are in Sydney and you are offered lotteries over the outcome space C =(train ticket to Paris, train ticket to London). Do you prefer $p = (1, 0)$ or $q = (0, 1)$, i.e. a ticket to Paris or a ticket to London? Suppose now the outcome space is C =(ticket to Paris, ticket to London, movie about Paris). Do you prefer $p' = (0.8, 0, 0.2)$ or $q' = (0, 0.8, 0.2)$?

Answer: Many students who say they prefer p to q will also say they prefer q' to p'. It may be that this type of choice is the result of improper understanding of what is a lottery rather than something more basic, so agents do not want p' because "there is no point watching a movie about Paris if I go there anyway". Indeed, students may not understand that each lottery will result in only one of the outcomes being realized, not two combined outcomes together. An alternative explanation may be that students do not want to face, under p', the prospect of watching a movie about Paris which would make them regret even more not having won the ticket to Paris. In any case, this type of reversal of preferences is a violation of the independence axiom, as adding a common consequence to the original lotteries should not change preferences among lotteries.

6.2 A Representation Theorem

Theorem 1. *[Representation theorem] If the four axioms presented above hold, then there exists a utility index u such that ranking according to expected utility accords with actual preference over lotteries. In other words, if we compare two lotteries, L_1 and L_2 represented by the probability vectors $P = (p_s)_{s=1,...,S}$ and $Q = (q_s)_{s=1,...,S}$ over the same set of outcomes $\mathcal{S} = (1, ..., S)$ then*

$$L_1 \preceq L_2 \Leftrightarrow \sum_{s=1}^{s=S} p_s u(c_s) \leq \sum_{s=1}^{s=S} q_s u(c_s)$$

Proof: See pp.176-178 of the Mas-Colell or p. 76 of Kreps. The easiest proof assumes there exist a worst, w, and a best lottery, t, and defines the utility of any lottery p as the number such that $p \sim u(p)t + (1 - u(p))w$. By continuity, that number exists and is unique. p itself is defined over the set of consequence $C : (a, b, c)$, and each of those consequences can be ascribed an utility according to the above method. Define thus $u(a)$ the utility of consequence a for example. We have $p \sim p(a)a + p(b)b + p(c)c$ (by reducibility). We have $a \sim u(a)t+(1-u(a))w$, and similarly b and c. By the independence axiom, $p \sim (p(a)u(a)+p(b)u(b)+p(c)u(c))t+(1-p(a)u(a)-p(b)u(b) - p(c)u(c))w$. Therefore, $u(p) = p(a)u(a) + p(b)u(b) + p(c)u(c)$.

Notes and implications:

- The utility function is an ordinal measure of utility, not a cardinal measure. This means the specific value $u(c_s)$ is not what is important, rather it is the ranking of lotteries which must be translated in a correct way by the utility function. This means you do NOT have to translate outcomes into one common measure, such as for example money.
- Said in another way, the theorem is a *representation theorem* – in other words it means we can represent preferences using the utility function, but that doesn't mean that individuals gain 'utility' from outcomes.
- To make the point further, $u(.)$ is unique only up to a positive, linear transformation. This means that $u(.)$ and any $U(.) = a + bu(.)$ such that $b > 0$ represent the same utility function. This means utility numbers have no meaning per se.
- This utility functional has a long history dating back to Bernouilli. The contribution of vNM was to show this was the only type of utility functional that respected the above series of four normatively reasonable axioms.
- EU is linear in probabilities. vNM's EUT makes it possible to obtain preferences between complex lotteries through a simple adding up of the utility of each of the components of the lottery weighted by their respective probabilities.
- EU indifference curves in the Marschak-Machina triangle are represented by parallel lines, sloping upward. Utility increases north-west.

Summing up, one had to remember the following points:

- Axioms of EUT are intuitively 'reasonable' assumptions about preferences. Whether they fit real choice (positive axioms) or are good guides for action (normative axioms) is up to one's perspective. The axioms, while reasonable, are not necessarily prescriptive or necessarily backed up or drawn from experimental evidence. One will see indeed that agent's actions do not necessarily fit with EUT, which means one or more of the axioms are not respected.
- Together, the axioms imply EUT while EU representations imply the axioms are verified.
- You need to know, understand and be able to use the axioms.

Application to lottery comparisons

EUT provides a simple ways to compare lotteries. Indeed, consider the EU of lottery L over x_1, x_2 and x_3 :

$$EU(L) = p_1U(x_1) + p_2U(x_2) + p_3U(x_3) \tag{1}$$

Rewrite this as

$$EU(L) = p_1U(x_1) + (1 - p_1 - p_3)U(x_2) + p_3U(x_3) \tag{2}$$

If we differentiate that equation with respect to probability, then we obtain

$$dEU(L) = -dp_1(U(x_2) - U(x_1)) + dp_3(U(x_3) - U(x_2) = 0 \tag{3}$$

for any two points on the same indifference curve. Rearranging,

$$\frac{U(x_2) - U(x_1)}{U(x_3) - U(x_2)} = \frac{dp_3}{dp_1} \tag{4}$$

which is a constant (check indeed that the ratio is independent of the specific normalization chosen for the utility functional). Therefore, $\frac{dp_3}{dp_1}$, which is the slope of indifference lines in the M-M triangle, is always the same, no matter where we are in the Marschak-Machina triangle. This means that once we have found two points in the M-M triangle which give the same utility to a particular person, then we can predict how that person will choose between any two point in the M-M triangle. It is this remarkable feature of expected utility theory which makes it so straightforward to test: you need only find two (non-degenerate) indifferent lotteries to know the preferences over the whole set of lotteries.

Checking that EUT holds means one will have to check that indifference curves are indeed straight parallel lines, by asking agents for their preferences over lotteries.

Application to risk

Consider an agent with wealth \$10,000 and utility normalized to u(\$x)=ln(x) defined over monetary outcomes, i.e. for example, \$10 provides utility u(\$10)=ln(10) $=$ 2.3026. Suppose this agent is offered a lottery $(1/2, 1/2)$ over the set of consequences (-\$100,\$101). Will the agent accept the bet?

- If the agent prefers not betting and keeping her \$10,000, she gets utility $\ln(10000) = 9.2103404$.
- If she bets, then she loses \$100 with probability $1/2$ and gains \$101 with probability $1/2$, so her expected utility is $1/2\ln(10000 - 100) + 1/2\ln(10000 + 101) = 9.2103399$.
- This is less than ln(10000), so the agent rejects the bet.

So far, so good.

- Now, suppose the agent is offered a lottery $(1/2, 1/2)$ over the set of consequences (-\$800,\$869). Then you can check she rejects this as well. Indeed, $\ln(10000) = 9.2103404 < 1/2\ln(10000 - 800) + 1/2\ln(10000 + 869) = 9.2103144$.
- Suppose now she is offered a lottery $(1/2, 1/2)$ over the set of consequences (-\$8000,\$38476). Then she rejects this as well. Indeed, $\ln(10000) = 9.2103404 < 1/2\ln(10000 - 8000) + 1/2\ln(10000 + 38476) = 9.1948633$.

Is this reasonable? For more on this issue, read Rabin (2000).[4]

7 Critique of the axioms of EUT

7.1 The Allais Paradox and fanning out

The Allais paradox (1953)[5] was first expressed in the context of the *common ratio* effect, and was then generalized to include the *common consequences* effect.

7.1.1 The common ratio effect

Consider the quartet of distributions (p_1, p_2, q_1, q_2) depicted below in a Marschak-Machina triangle and which, when connected, form a parallelogram. Let them be defined over outcomes (consequences) $x_1 = \$0, x_2 = \$50$ and $x_3 = \$100$.

1. p_1 is $(0\%, 100\%, 0\%)$: : \$0 with 0% chance, \$50 with 100% chance, \$100 with 0% chance.

[4]Rabin M., 2000, Risk Aversion and Expected-Utility Theory, Econometrica 68(5), 1281-1292.

[5]Allais M., 1953, Le comportement de l'homme rationel devant le risque, critique des postulats et axiomes de l'école américaine, Econometrica 21(4), 503-546.

2. p_2 is $(1\%, 89\%, 10\%)$: \$0 with 1% chance, \$50 with 89% chance, \$100 with 10% chance.

3. $q_1 = (89\%, 11\%, 0\%)$: \$0 with 89% chance, \$50 with 11% chance, \$100 with 0% chance.

4. $q_2 = (90\%, 0\%, 10\%)$: : \$0 with 90% chance, \$50 with 0% chance, \$100 with 10% chance.

When confronted with this set of lotteries, there are people who choose p_1 over p_2 and choose q_2 over q_1. This contradicts the independence axiom of expected utility, as we are going to prove, both *diagrammatically* and *analytically*, that an expected utility maximizer who prefers p_1 to p_2 ought to prefer q_1 to q_2. This contradiction is called the "Allais Paradox".

Diagram (see figure below): Consider the diagram below where the four lotteries above are represented. Note how the line that connects p_1 and p_2, and the one that connects q_1 and q_2, are parallel. If indifference curves are parallel to each other, as in EUT, then it should be if $p_1 \succ p_2$ then $q_1 \succ q_2$. We can see this diagrammatically by considering an indifference curve which translates the fact that $p_1 \succ p_2$: it ought to separate p_1 from p_2, with p_1 above and p_2 below. A parallel indifference curve will divide q_1 from q_2, with q_1 above and q_2 below, from which one can conclude that $q_1 \succ q_2$.

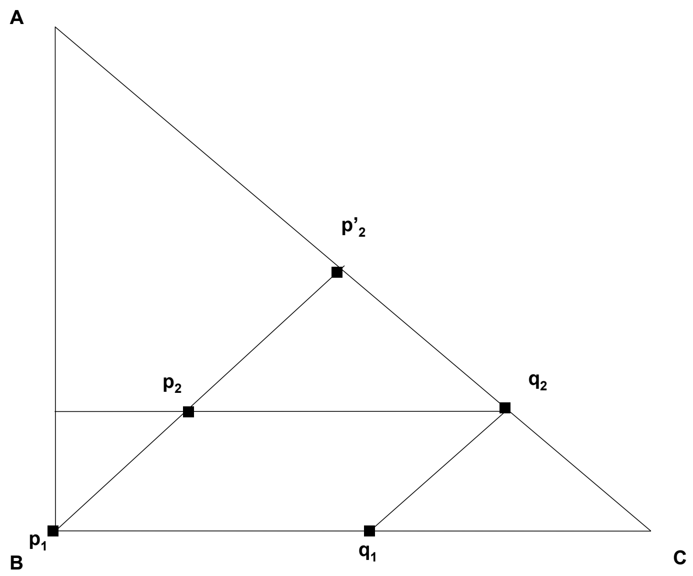

Figure 5: Common ratio and common consequences

Analysis: An analytical proof would go as follows: As $p_1 \succ p_2$ then by the von Neumann-Morgenstern expected utility representation, there is some elementary utility function u such that:

$$u(\$50) > 0.1u(\$100) + 0.89u(\$50) + 0.01u(\$0) \tag{5}$$

But as we can decompose

$$u(\$50) = 0.1u(\$50) + 0.89u(\$50) + 0.01u(\$50) \tag{6}$$

then subtracting $0.89u(\$50)$ from both sides, the first equation implies:

$$0.1u(\$50) + 0.01u(\$50) > 0.1u(\$100) + 0.01u(\$0) \tag{7}$$

Adding $0.89u(\$0)$ to both sides:

$$0.1u(\$50) + 0.01u(\$50) + 0.89u(\$0) > 0.1u(\$100) + 0.01u(\$0) + 0.89u(\$0) \tag{8}$$

Combining the similar terms together, this means:

$$0.11u(\$50) + 0.89u(\$0) > 0.1u(\$100) + 0.90u(0) \tag{9}$$

which implies that $q_1 \succ q_2$, which is what we sought.

7.1.2 The common consequence effect.

With reference to the graph above, tests of the common ratio effect involve pairs of choices like (p_1 or p_2) and (q_1 or q_2). Tests of the common consequence effect involve pairs of choices like (p_1 or p_2') and (q_1 or q_2).

Example: Consider the choice between $p_1 = (0\%, 100\%, 0\%)$ and $p_2' = (9\%, 0\%, 91\%)$. Suppose the agent prefers p_1 to p_2' but prefers q_2 to q_1. Graphically, one can see that this contradicts the independence axiom too.

The 'common consequence' effect is less robust than the 'common ratio' effect, i.e. the contradiction in choice is less often observed in that case. This may be due to the simpler nature of common consequences lotteries, which as the name indicates involve comparison between lotteries defined across a maximum of two consequences only, rather than a maximum of three in the common ratio effect. Part of the issue with the common ratio effect may thus be due to how complicated that type of choice is, rather than to some inherent behavior pattern.

7.2 Process violation

A common example of violation of transitivity is the "P-bet, \$-bet" problem:

Consider the following two lotteries:

1. P-bet: \$30 with 90% probability, and zero otherwise.
2. \$-bet: \$100 with 30% probability and zero otherwise.

An agent is first asked which lottery they would prefer playing, and then asked what price they would buy a ticket to play the P-bet or the \$-bet. Notice that in this example, the expected payoff of the \$-bet is higher than that of P-bet, but the \$-bet is also more risky (lower probability to win). This may be why people tend to choose to play the P-bet over the \$-bet. Yet, the same people are ready to pay *more* for a ticket to play the \$-bet than for a ticket to play a P-bet. Said in another way, although when directly asked, they would choose the P-bet, they are willing to pay a lower certainty-equivalent amount of money for a P-bet than they do for a \$-bet. For example, they might express a preference for the P-bet, but be ready only to pay \$25 to play the P-bet while being ready to pay \$27 for the \$-bet.

Many have claimed that this violates the transitivity axiom. The argument is that one must be indifferent between the certainty-equivalent amount ("price") of the bet and playing it, so that in utility terms, taking the example above again, I would have $U(P-bet) = U(\$25)$ and $U(\$-bet) = U(\$27)$. By monotonicity, since more money is better than less money, $U(\$27) > U(\$25)$ and so we should conclude that $U(\$-bet) > U(P-bet)$. Yet, when asked directly, people usually prefer the $P-bet$ to the $\$-bet$, implying $U(P-bet) > U(\$-bet)$. Thus, the intransitivity.

However, the question is whether this "intransitivity" is not simply due to overpricing of \$-bets, i.e. agents being unable to price bets correctly. For example, I once offered to a student a ticket for a bet giving $\frac{1}{2}$ chance of \$100 and $\frac{1}{2}$ chance of \$0. I was offered \$12 in exchange for that ticket. When I asked the same person how much they would sell me a ticket for this bet, they quoted a price of \$43... Was the discrepancy due to improper understanding by the student, to lack of experience, to an exaggerated fear of getting it "wrong" in front of others, or to the fact this was not real money? Would the discrepancy have survived if the student had been given the opportunity to change his bids, or if a negotiation process had been put in place? One could design a series of alternative design for the experiment, but it illustrates a Willingness to Accept / Willingness to Pay disparity that survives whatever the experimental set-up (on the discrepancy WTA/WTP, see 'endowment effect')

7.3 Framing effect and elicitation bias

In order to determine if the P-bet, \$-bet anomaly is due to the procedure by which the preference between lotteries is elicited, or to true intransitivity in preferences, consider an experiment mentioned by Camerer (1995).[6] In this experiment, the P-bet offers \$4 with probability $\frac{35}{36}$, 0 else, while the \$-bet offers \$16 with probability $\frac{11}{36}$, 0 else.

The subject is asked to choose between the two, and then is asked how much they are prepared to pay for each of them:

[6] Camerer C., 1995, Individual decision making, in The handbook of experimental economics, J.H. Kagel and A.E. Roth (Eds), Princeton University Press, p. 660.

- If the subject chose the P-bet and quoted a higher price (e.g. \$3.5) for it than for the \$-bet (e.g. \$3), then the preferences were judged to be coherent.

- If the subject chose the P-bet but then quoted a lower price (e.g. \$3.5) for it than for the \$-bet (e.g. \$4.5), then this was judged to potentially indicate a reversal of preference. In order to determine if that potential reversal was merely due to overpricing of the \$-bet, experimenters then asked the subject to choose between each bet and a stated amount, \$4 in this case.

  - If the subject chooses \$4 rather than the \$-bet, then that means they previously overpriced the \$-bet and there was thus a framing effect at play rather than a true violation of transitivity. Indeed, the only reason there was an apparent violation of transitivity is merely that in one case the subject was asked to make a choice between lotteries, and in the second case, he was asked to price lotteries. Those are rather different mental processes.
  - If on the other hand the subject indeed choose the \$-bet over the \$4 and the \$4 over the P-bet, then that meant there was true violation of transitivity, independent of any framing effect. Indeed, the question was framed the same way (choice between lotteries) and led to a contradiction of transitivity.

Loomes, Starmer and Sugden (1991)[7] contend that violations of transitivity may occur because agents use rather less sophisticated techniques than EUT to evaluate lotteries. They posit that agents tend to choose lotteries with the larger probability as long as the payoffs are close, and choose lotteries with the larger payoff is the payoffs are far away. This can lead to intransivities. For example, suppose an agent is asked to choose between 60% probability of getting \$8 (L_1) and 30% probability of getting \$18 (L_2). Payoffs are close, so most agents choose the first as it has higher probabilities to win. Suppose then the agent is offered \$4 for sure against 60% probability of getting \$8. Again, payoffs are close, so most people agents choose the first. Finally, however, most agents would choose 30% probability of getting \$18 rather than \$4 with certainty, as payoffs are far apart so the lottery with the higher payoff is preferred. This leads to a violation of transitivity. Indeed, the first choice implies $u(L_1) > u(L_2)$. The second choice implies $u(L_1) < u(\$4)$. The third choice implies $u(L_2) > u(\$4)$. Combining the two last choices, $u(L_2) > u(L_1)$, which contradicts the original choice between lotteries. As above, the advantage of this design is that the agent is always asked for their choice among lotteries and are never asked to evaluate them individually.

This can be generalized to say that agents will try to minimize the number of informations on which to draw their decisions rather than taking into account all the parameters in the decision. E.g. when choosing between two brands

[7] Loomes G., C. Starmer and R. Sugden, 1991, Observing Violations of Transitivity by Experimental Methods, Econometrica 59(2), 425-439.

of foods, you will not look all the ingredients that came into their making but rather will choose salient dimensions such as sugar content and price to make your decision. For a very interesting account of the kind of reasoning people do when buying, see Viswanathan, Rosa and Harris (2005).[8] For more on designing experiments to avoid 'framing' subjects, i.e. to avoid obtaining results that are dependent on the design on the experiment rather than on true subject behavior, read Machina (1987).[9]

7.4 Endowment effects

The endowment effect is such that people are less willing to pay for an object than they would ask as payment if they owned it. This is the difference between the WTP (Willingness To Pay) and the WTA (Willingness To Accept). This 'endowment effects' is not explained by wealth effects, but may be due to loss aversion (which will come up in 'regret theory'). For example, if losses are more painful than equally sized gains are pleasurable, then one will offer to pay only $12 to play a lottery with 50% chance of getting $100 so as to minimize loss in case of bad luck (when the lottery draws $0, so the loss is $12), while one will be prepared to sell the bet only at the much higher $43 to minimize loss in case of bad luck, in this case, when the lottery draws $100 (the loss is $57).

Note that the disparity between WTP and WTA is explained in some measure within EUT, but not to the extent it appears in reality. For more on this, see exercise 5 of Choice under Uncertainty.

7.5 Discussion

Those experiments that contradict the predictions of EUT are interesting because they do not require to estimate utility functions for individuals, and they allow for a direct test of the axioms. The drawback is that they do not allow one to know how 'badly' inconsistent the choices of the agents are. This would require many more experiments. The only information from those experiments is the percentage of individuals whose choices violate EUT, and in what specific way the choice is violated. An obvious critique of the above experiments is that the choices asked from subjects in those experiments are rather too complicated and unintuitive compared with the type of choices they are facing in real life. The issue is then to present the problem in an intuitive form, or to allow the subjects to get acquainted with the way the experiment is set up and the way its consequences will affect them (i.e. let them play many rounds).

Note however that some of the experiments exposed above were adapted to rats using food as a currency and those experiments showed violations of the independence axiom as well. This points to a possible evolutionary benefit of

[8]Viswanathan M., J.A. Rosa and J.E. Harris, 2005, Decision Making and Coping of Functionally Illiterate Consumers and Some Implications for Marketing, Journal of Marketing 69(1), 15-31.

[9]Machina M.J., 1987, Choice Under Uncertainty: Problems Solved and Unsolved, The Journal of Economic Perspectives 1(1), 121-154.

behaving in ways that are different from those of an expected utility maximizer. Some experiments have been done comparing how farmers in traditional farming communities evaluate lotteries, and have indeed shown their behavior, while different from urbanized people, may have evolutionary benefits in a context where droughts and famines are likely (Humphrey and Verschoor (2004).[10]

8 Alternatives and generalizations of EUT

Contradictions within the framework of EUT led to the development of various alternatives to EUT, of which we will present some. In a first part, we consider prospect theory and its successors, in a second part we consider regret theory. Starmer (2000)[11]provides good further reading on this topic.

8.1 Prospect Theory, Rank Dependent EUT and Cumulative Prospect Theory

In this section, by order of difficulty and chronology, we consider prospect theory ('PT') by Kahneman and Tversky (1979),[12] rank dependent utility theory ('RDEUT') (Quiggin, 1982),[13] and Cumulative Prospect Theory ('CPT') (Kahneman and Tversky, 1992).[14] Those utility representations differ from EUT by considering subjective probabilities as a function of objective probabilities, that correspond to how people estimate probabilities in reality (probability weighting functions). They also differ by considering an "editing phase", whereby one divides the set of consequences of a lottery into either 'gains' or 'losses'. Different probability weighting functions are assigned to probabilities of a consequence depending on whether the consequence is a gain or a loss.

8.1.1 Prospect Theory

Consider the set of consequences $C = (a, b, c, d)$ such that $a \prec b \prec c \prec d$ and consider lottery $L = (p_a, p_b, p_c, p_d)$. Under prospect theory, its utility function is of the form:

$$U(L) = \pi(p_a)u(a) + \pi(p_b)u(b) + \pi(p_c)u(c) + \pi(p_d)u(d) \tag{10}$$

$\pi(p)$, the probability weighting function, takes the following form:

[10]Humphrey S.J. and A. Verschoor, 2004, Decision-making Under Risk among Small Farmers in East Uganda, Journal of African Economies 13(1), 44-101.

[11]Starmer C., 2000, Developments in Non-expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk, Journal of Economic Literature 38(2), 332-382.

[12]Kahneman D. and A. Tversky, 1979, Prospect Theory: An Analysis of Decision under Risk, Econometrica 47(2), 263-291.

[13]Quiggin J.C., 1982, A theory of anticipated utility, Journal of Economic Behavior and Uncertainty 3(4), 323-343.

[14]Kahneman D. and A. Tversky, 1992, Advances in prospect theory: cumulative representations of uncertainty, Journal of Risk and Uncertainty 5(4), 297-323.

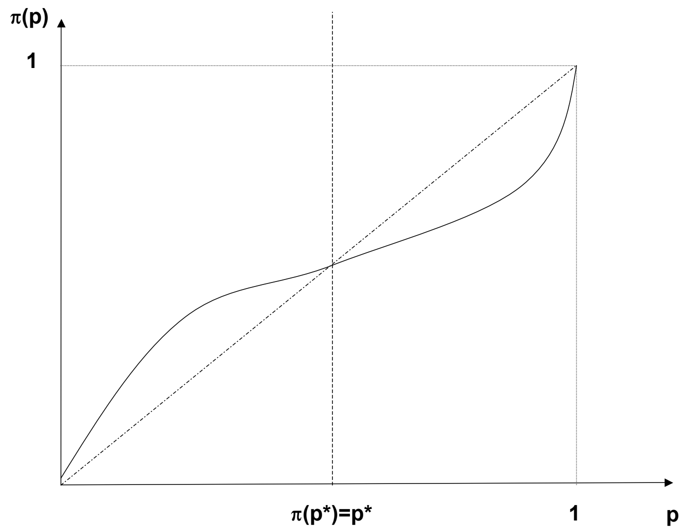

Figure 6: Probability weighting function

The reason for offering that type of representation is that there tends to be a difference between 'psychological probability' and 'objective probability': agents are quite able to estimate lotteries with probability between around 1/3 and 2/3 but overestimate low occurrence events and underestimate high occurrence events. Below p^* in the graph above, the agent weighs events above their statistical probability, while this is the opposite for events with statistical probability above p^*. Intuitively, agents overweigh the probability of events with low statistical probabilities, maybe because they are not used to estimating the probability of their occurence and they feel anxious about them. The other side of the coin is that they will under-estimate the probability of events that are statistically quasi-given. If for example I over-estimate the probability of being run down by a car (low probability event), then this means I underestimate the probability that I will cross the road safely (high probability event). p^*, which is the probability that agents are able to estimate correctly, is variously evaluated between 0.3 and 0.5.

Those systematic "errors" in the perception of probabilites mean that for example, a bet offering 5% chance of $100 may sell for $6, 50% chance may sell for $50 while 95% chance may sell for $93. EUT would find it difficult to explain how risk aversion would vary so much depending on the bet (the agent is first risk loving, then risk neutral, then risk averse). CPT would explain this variation with probability weighting functions. One would thus have:

$$\pi(5\%)u(100) + \pi(95\%)u(0) = u(6) \tag{11}$$
$$\pi(50\%)u(100) + \pi(50\%)u(0) = u(50) \tag{12}$$
$$\pi(95\%)u(100) + \pi(5\%)u(0) = u(93) \tag{13}$$

Normalizing $u(0) = 0$, then one finds that

$$\pi(5\%) = u(6)/u(100) \tag{14}$$
$$\pi(50\%) = u(50)/u(100) \tag{15}$$
$$\pi(95\%) = u(93)/u(100) \tag{16}$$

Suppose the agent is risk neutral, then

$$\pi(5\%) = 6/100 \tag{17}$$
$$\pi(50\%) = 1/2 \tag{18}$$
$$\pi(95\%) = 93/100 \tag{19}$$

CPT thus does not require one to assume anything other than risk-neutrality to explain the pricing of the different lotteries above.

8.1.2 Rank Dependent EUT

RDEUT differs from PT in that it requires that the sum of probabilies over a whole set of events be equal to one. This is doe as follows: If one orders events from the least to the most preferred, ($a \prec b \prec c \prec d$), then one writes the utility of lottery $L = (p_a, p_b, p_c, p_d)$ as:

$$U(L) = f(p_a)u(a) + f(p_b)u(b) + f(p_c)u(c) + f(p_d)u(d) \tag{20}$$

with

$$f(p_a) = \pi(p_a) \tag{21}$$
$$f(p_b) = \pi(p_a + p_b) - \pi(p_a) \tag{22}$$
$$f(p_c) = \pi(p_a + p_b + p_c) - \pi(p_a + p_b) \tag{23}$$
$$f(p_d) = 1 - \pi(p_a + p_b + p_c) \tag{24}$$

Compared with PT, RDEUT guarantees that the sum of the probability weights assigned to each events sum up to 1. Check indeed that $f(p_a) + f(p_b) + f(p_c) + f(p_d) = 1$. In RDEUT, $f(.)$ is a cumulative probability function, with outcomes added in the order of their preferences.

8.1.3 Cumulative Prospect Theory

CPT differs from RDEUT by (arbitrarily) dividing the set of consequences into losses and gains and assigning different probability weighting function to them.

Suppose for example that I consider a and b as losses and c and d as gains. Defining f_G the probability weighting utility function applied to gains, and f_L the probability weighting function applied to losses, I then have

$$U(L) = f_L(p_a)u(a) + f_L(p_b)u(b) + f_G(p_c)u(c) + f_G(p_d)u(d) \tag{25}$$

with

$$f_L(p_a) = \pi_L(p_a) \tag{26}$$
$$f_L(p_b) = \pi_L(p_a + p_b) - \pi_L(p_a) \tag{27}$$
$$f_G(p_c) = \pi_G(p_c + p_d) - \pi_G(p_d) \tag{28}$$
$$f_G(p_d) = \pi_G(p_d) \tag{29}$$

with $\pi_G(.)$ defined over gains and $\pi_L(.)$ defined over losses, both increasing. In order to guarantee the sum of probabilies is one as in RDEUT, I must have that $\pi_L(p_a+p_b) = 1-\pi_G(p_c+p_d)$, that is, the cumulative probabilities of loss events is the complement of the cumulative probabilities of gain events. CPT differs from RDEUT by taking into account whether the event is considered as a loss or as a gain. This is intuitively justified by saying that losses affect the agent more than gains, and are thus overweighted. The distinction between 'gain' and 'losses' is of course subjective, and must be calibrated depending on the individual and the situation. An alternative way to consider gains and losses is to assign different utility functionals to gains compared to losses, whereby agents are almost risk neutral with respect to gain (for example, a 50/50 chance to gain \$100 would be evaluated at \$49), while being very risk averse with respect to losses (for example, a 50/50 chance to lose \$100 would be evaluated at $-$\$60, i.e. the agent is ready to pay \$60 not to play the lottery).

Under CPT and RDEUT, indifference curves take the following form:

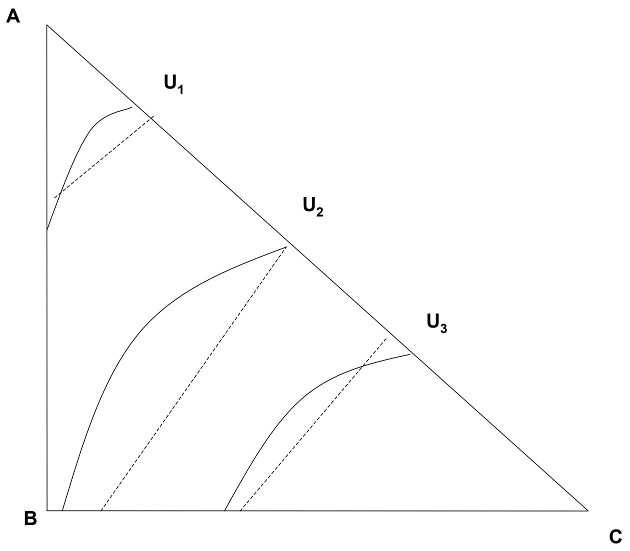

Figure 7: RDEUT and CPT curves

At this point, given the complexity of the arguments, it is rather difficult to make intuitive reasonings to justify the precise shape of the indifference curves. Consider however the indifference curve corresponding to utility U_2. If the agent was an EU maximizer, then its indifference curve would take the form of the dotted line. On the hypothenuse, one considers a lottery with equal probability between A and C, which the agent evaluates as in EUT (if one assumes $p^* = 0.5$). On the line joining B and C, one has a lottery with a better outcome with high probability, which is going to be under-evaluated, which explains the indifference curve is above the EU curve with the corresponding utility. However, the under-evaluation is not so high since the probability of the best event is 0. As the probability of that best event increases, so does the discrepancy, until a point where the discrepancy decreases again as the lotteries involve closer outcomes. The indifference curves are thus concave.

Exercises:

1) Try to justify the compared shape of the EU and RDEU curves at level of utility U_1 and level of utility U_3.

2) How does CPT change the shape of the curve compared to RDEUT?

8.2 Regret theory

Regret theory differs from the above theories because it does not obtain a 'utility function' but rather, a gain/loss function that evaluates by how much choosing a lottery over another and then comparing their outcome is expected to make one happy/unhappy. To the difference of other theories, regret theory explicitly takes into account the opportunity cost of making a decision. Decisions are not taken in a vacuum; making one decision precludes making another one. The theory is exposed in Loomes and Sugden (82, 86)[15][16] and Sugden (85).[17] Agents thus faced with alternative lotteries do not seek to maximize expected utility but rather to minimize expected regret (or maximize expected rejoicing) from their choice.

Formally, suppose lottery p has probabilities $(p_1, .., p_n)$ while lottery q probabilities $(q_1, .., q_n)$ over the same finite set of outcomes $x = (x_1, ..., x_n)$. Expected rejoice/regret from choosing p over q is:

$$E(r(p, q)) = \sum_i \sum_j p_i q_j r(x_i, x_j) \tag{30}$$

where $p_i q_j$ is the probability of outcome x_i in lottery p and outcome x_j in lottery q_j. Lottery p will be chosen over lottery q if $E(r(p, q))$ is positive and q will be chosen over p if $E(r(p, q))$ is negative.

Note that if $r(x_i, x_j) = u(x_i) - u(x_j)$ then $E(r(p, q))$ is simply the difference in expected utility between lottery p and q and regret theory leads to exactly the same decision as EUT.

Under regret theory:

1. $r(x, y)$ is increasing in x so that the higher the good outcome, the higher the rejoicing.

2. $r(x, y) = -r(y, x)$, so that regret/rejoice is symmetric: Getting the good outcome x rather than the bad outcome y produces the same amount of rejoicing than the amount of regret induced by getting the symmetric outcome. The expected rejoice at a gain is the same as the expected regret at a same sized loss,

3. $r(x, y) > r(x, z) + r(z, y)$ when $x > z > y$, that is the rejoicing increases more than proportionately with the difference in outcome. I rejoice more if I gain \$100 rather than \$0 than the sum of rejoicing if I gain \$50 rather than \$0 and \$100 rather than \$50, even though the result is the same. Consider for example how French people would react if their rugby team beat Australia (pride, celebration), compared to how they would react if France beat New Zealand (the usual...), and New Zealand beat Australia (no one cares, at least in France, except maybe in Toulouse...).

The advantage of the regret/rejoice model is that the "indifference curves" over lotteries derived from it can be intransitive, i.e. yield up preference reversals. From the beginning of the course, we know that this means indifference curves

[15] Loomes G. and R. Sugden, 1982, Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty, The Economic Journal 92(368), 805-824.

[16] Loomes G. and R. Sugden, 1986, Disappointment and Dynamic Consistency in Choice under Uncertainty, The Review of Economic Studies 53(2), 271-282.

[17] Sugden R., 1985, Regret, recrimination and rationality, Theory and Decision 19(1), 77-99.

can cross *inside* the M-M triangle. Exercise 6 in 'Choice under Uncertainty' shows that not only does RT allow for intransitive preferences, but it also mandates one specific direction in which preferences can be intransitive, i.e. while it would allow $A \succ B \succ C \succ A$, it would not allow $A \prec B \prec C \prec A$. This means that a test of RT is to check that intransitiveness, when it occurs, occurs in one direction only. Experiments tend to bear this out.

The main issue with RT is that while comparisons between two lotteries are relatively easy under its setting, comparisons between several lotteries are much more involved. For example, while it is easy to model the choice between (marriage, no marriage), it is more difficult to model the choice between (marriage, civil partnership, none of the above). Of course, one could divide the comparison between those three lotteries into three comparisons between two lotteries, but one then loses most of the point of modeling the effect of regret on lottery choices.

Part II
Game theory

The structure of this lecture as well as the notations used are drawn from Christopher Wallace's lectures on game theory at Oxford. Additional material on auctions and repeated games can be found in Gibbons R., 1992, A Primer in Game Theory, FT Prentice Hall.

9 Introduction

This section aims to present essential tools for predicting players' actions in a range of strategic situations. By order of complexity, we will study strategic form games, where players choose actions at the same time so there is no conditioning of actions based on the action of others, extensive form games where players choose actions in succession but know what action was taken by others previously so there is no uncertainty, and finally Bayesian games where players choose action in succession but are not certain what the other played previously. We will also cover repeated games under full information, where players condition their action on their observation of what was played in a previous stage of the game.

The lecture builds on the analysis of one single game, a coordination game, which is made progressively more complex so as to introduce new concepts. Those include the concept of dominant strategy, Nash equilibrium, mixed strategy Nash equilibrium, backward induction, subgame perfect Nash equilibrium, Bayesian Nash equilibrium and forward induction. A few other games are introduced, includings games of auctions (to illustrate the use of Bayesian Nash equilibrium concepts) and the Prisoners' dilemma (in relation to infinitely repeated games).

This lecture does not introduce many applications of game theory, as it is expected the student will encounter applications relevant to his or her area of specialization further on in the course of his or her MSc. Rather, the lecture aims to give good mastery of notations and techniques for solving a wide range of games.

10 Readings:

10.1 Textbook reading:

- Kreps, Chs. 11-14
- Varian, Ch. 15
- Mas-Colell, Chs. 7-9

Among the many specialist text books which are currently available, by degree of difficulty, one finds:

- Carmichael F., 2005, A Guide to Game Theory, FT Prentice Hall (less advanced)
- Gibbons R., 1992, A Primer in Game Theory, FT Prentice Hall (more advanced)
- Fudenberg D. and J. Tirole, 1991, Game Theory, MIT Press (very advanced)

10.2 Articles[18]

- Cho I.-K. and D. Kreps, 1987, Signalling Games and Stable Equilibria, Quarterly Journal of Economics 102, 179-222.
- Cho I.-K. and J. Sobel, 1990, Strategic Stability and Uniqueness in Signalling Games, Journal of Economic Theory 50, 381-413.
- Geanakoplos J., 1992, Common Knowledge, The Journal of Economic Perspectives 6, 53-82.
- Goeree J.K. and C.A. Holt, 2001, Ten Little Treasures of Game Theory and Ten Intuitive Contradictions, American Economic Review 91(5), 1402-1422.[19]
- Mailath G.J., L. Samuelson and M. Swinkels, 1993, Extensive form reasoning in normal form games, Econometrica 61, 273-302.
- Palacios-Huerta I., 2003, Professionals play minimax, Review Of Economic Studies 70(2), 395-415.
- Rubinstein A., 1991, Comments on the Interpretation of Game Theory, Econometrica 59, 909-924.
- Shaked A., 1982, Existence and Computation of Mixed Strategy Nash Equilibrium for 3-Firms Location Problem, Journal of Industrial Economics 31, 93-96.
- Van Damme E., 1989, Stable equilibria and forward induction, Journal of Economic Theory 48(2), 476-496.

11 Strategic form games

Definition: A strategic form game is defined by:

1. Players: The set $N = 1, ..., i, ..., n$ of agents who play the game, for example: Adelina and Rocco.
2. Strategies: For each $i \in N$, I define the set of strategies S_i with typical element s_i available to agent i, for example, {Cooperate, Defect}.
3. Payoffs: Denote $S = (S_1, ..., S_i, ..., S_n)$ the set of all strategies available to all the players, for example ({Cooperate, Defect},{Cooperate, Defect}).

To each strategy profile $s = (s_1, ..., s_i, ..., s_n)$ in S, for example (Cooperate, Defect), one associates payoff $u_i(s)$ corresponding to that combination of strategies. $u = (u_1, ..., u_i, ..., u_n)$ is the set of payoffs of the game, defined for all s in S.

[18]Some articles in this list were contributed by previous teachers in MSc Economic Theory 1 at the UEA.

[19]This is an article that is entertaining and easy to read about how what people play may detract from Nash equilibria, though in predictable ways.

$G = N, S, u$ defines a strategic form game.

Notation: One will denote $s_{-i} = \{s_1, ..., s_{i-1}, s_{i+1}, ..., s_n\}$ the set of actions taken by agents other than i in the strategy profile s and $S_{-i} = \{S_1, ..., S_{i-1}, S_{i+1}, ..., S_n\}$ the set of strategies available to players other than i.

Example: The Prisoners' Dilemma

1. Players $N = 1, 2$.
2. Strategies $S_i = C, D$, $i = 1, 2$
3. Payoffs $u_1(C, D) = u_2(D, C) = -6$, $u_1(C, C) = u_1(C, C) = -1$, $u_1(D, C) = u_2(C, D) = 0$ and $u_1(D, D) = u_2(D, D) = -4$.

The game can be represented in normal form as follows:

	C	D
C	-1,-1	-6,0
D	0,-6	-4,-4

11.1 Dominance

The following definitions introduce the concept of a strictly dominant strategy equilibrium. In a strictly dominant strategy equilibrium, players play their strategy irrespective of the action of others.

Definition: Strategy $s_i \in S_i$ strictly dominates strategy $s'_i \neq s_i$ in S_i for player i if $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all s_{-i} in S_{-i}.

If s_i strictly dominates another strategies s'_i, then that strategy is strictly dominated and can be eliminated from consideration by i.

Definition: Strategy s'_i is strictly dominated if there is an $s_i \in S_i$ that strictly dominates it.

If s_i strictly dominates all other strategies, then it is strictly dominant. Note that s_i is by definition unique.

Definition: $s_i \in S_i$ is strictly dominant for i if it strictly dominates all $s'_i \neq s_i$ in S_i.

An equilibrium in strictly dominant strategies exists if all players have a strictly dominant strategy. Note that from the above, a strictly dominant strategy is necessarily unique.

Definition: $s^* \in S$ is a strictly dominant strategy equilibrium if $u_i(s^*_i, s_{-i}) > u_i(s_i, s_{-i})$ for all players $i \in N$ and for all strategy profiles $s_{-i} \in S_{-i}$, that is, if all elements of s^* are strictly dominant strategies.

Example: The strictly dominant strategy equilibrium in the Prisoners' Dilemma is {D,D}. Indeed, D is a strictly dominant strategy for player 1 as $u_1(D,C) > u_1(C,C)$ and $u_1(D,D) > u_1(C,D)$. The same holds for player 2.

Remark: One can also define weak dominance as follows: Strategy $s_i \in S_i$ weakly dominates (or simply "dominates") strategy $s_i' \neq s_i$ in S_i for player i if $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_{-i} in S_{-i}.

11.1.1 Iterated delection of strictly dominated strategies

A strategy s_i' that is strictly dominated will not be part of any equilibrium of the game. One can therefore exclude strictly dominated strategies from the set of available strategies.

Example: Consider the following example:

	L	M	R
U	2,3	4,1	-1,6
D	5,3	4,-5	7,4

M is strictly dominated by L. L is strictly dominated by R. Both L and M can thus be eliminated. U is only weakly dominated by D, as it obtains the same payoff as D when 2 plays M. It cannot thus be eliminated from the set of available strategies.

One obtains the following remaining game:

	R
U	-1,6
D	7,4

In that remaining game, U is strictly dominated by D. The dominant strategy equilibrium of the game, obtained by iterated deletion of strictly dominated strategies, is thus (D, R).

Games with a strictly dominant strategy equilibrium are of limited interest. Indeed, in so far as players' actions are not dependent on others' action, those games cannot properly be called strategic. We introduce in the following part a concept, that of Nash equilibrium, which is of interest in proper strategic games.

11.2 Nash equilibrium

Definition: A Nash equilibrium ('NE') is a strategy profile s^* such that for every player i, $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$ for all $s_i \in S_i$.

Note how different the NE concept is from the dominant strategy concept: the NE concept does not require that s_i^* be dominant for all s_{-i} in S_{-i}, but only for s_{-i}^*. That is, *taking s_{-i}^* as given*, agent i must not strictly prefer to deviate to an action other than s_i^*.

Example: Nash Equilibrium in a Coordination Game

1. Players N =French (F), British (B).
2. Strategies S_i =Coffee (C), Pub (P), i =F,B
3. Payoffs $u_F(C,P) = u_B(C,P) = 1$, $u_F(C,C) = u_B(P,P) = 4$, $u_F(P,C) = u_B(P,C) = 0$ and $u_F(P,P) = u_2(C,C) = 3$.

The game can be represented in normal form as follows:

		B	
		C	P
F	C	4,3	1,1
	P	0,0	3,4

Neither strategies is strictly (or weakly) dominated for either players. However, suppose that F plays C. Then B is better off playing C. Suppose that F plays P. Then B is better off playing P. Formally, $u_F(C,C) \geq u_F(P,C)$ and $u_B(C,C) \geq u_B(C,P)$. Similarly, $u_F(P,P) \geq u_F(C,P)$ and $u_B(P,P) \geq u_B(P,C)$. This means that both $\{C,C\}$ and $\{P,P\}$ are Nash equilibria of the game.

Note how we found two Nash equilibria of the game above. This generalizes to saying that Nash equilibria are not necessarily unique, unlike dominant strategy equilibria. However, from the definition, one can check that any dominant strategy equilibrium is also a Nash equilibrium. There is no way in the game above to choose which of the Nash equilibria is more likely to be chosen. However, the Nash equilibrium concept allows one to say that players will play either one or the other Nash equilibria. The Nash equilibria we found are such that players choose actions in a deterministic way, that is, if for example the Nash equilibrium is {C,C}, then both players play C. Those are called *pure strategy Nash equilibria* ('PSNE'). We will see below there exists a third Nash equilibrium of this game, where players choose actions at random according to pre-defined probability. Those are called *mixed stategy Nash equilibria* ('MSNE').

A common way to find Nash equilibria is to use the concept of Best Response Function, which is particularly useful when players' action set is continuous (such as quantity or price in a game of competition).

Definition: The best-response function ('BRF') for player i is a function B_i such that $B_i(s_{-i}) = \{s_i | u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \text{ for all } s'_i\}$. The BRF states what is the best action for i the whole range of possible profile of actions of other agents.

Definition: s^* is a Nash equilibrium if and only if $s^*_i \in B_i(s^*_{-i})$ for all i. This means that elements of s^* must be best responses to each other.

Example: In the example above, $B_F(C) = C$ and $B_F(P) = P$. Similarly, $B_B(C) = C$ and $B_B(P) = P$.

11.2.1 Nash equilibrium in mixed strategies

In the following, we consider the possibility for players to choose actions at random according to pre-defined probabilities. For the sake of easy modeling, we assume agents have access to a randomizing device (such as a coin for example), which allows them to randomize over their actions. For example, an agent who decides to play C with probability 1/8 in the game above can do so by saying he will play C whenever three throws of the coin all give out 'tail'. We will see later whether agents can randomize in an accurate and rational way without access to such a randomizing device.

Example: Mixed strategy Nash equilibria ('MSNE') in the Coordination Game: We did not consider above the case of mixed strategies. Consider thus strategies s_i of the form: i plays C with probability p_i and play P with probability $1 - p_i$. Strategy s_i can be denoted in short as p_i. Suppose players F and B play strategies p_F and p_B respectively. The payoff to F of playing C is then $u_F(C, p_B) = 4p_B + 1(1 - p_B) = 3p_B + 1$. Similarly, $u_F(P, p_B) = 3 - 3p_B$. I also obtain that $u_B(p_F, C) = 3p_F$ and $u_B(p_F, P) = p_F + 4(1 - p_F) = 4 - 3p_F$.

Therefore, F will play C whenever $3p_B + 1 > 3 - 3p_B$, that is, whenever $p_B > \frac{1}{3}$. She will play P whenever $p_B < \frac{1}{3}$ and will be indifferent between the two actions whenever $p_B = \frac{1}{3}$. Similarly, B will play C whenever $3p_F > 4 - 3p_F$, that is if $p_F > \frac{2}{3}$. He will play P whenever $p_F < \frac{2}{3}$, and will be indifferent between the two actions whenever $p_F = \frac{2}{3}$.

One thus has three Nash equilibria: $\{p_F = 0, p_B = 0\}$ and $\{p_F = 1, p_B = 1\}$ as before, and a mixed strategy Nash Equilibrium $\left\{\frac{2}{3}, \frac{1}{3}\right\}$ whereby each player plays its own favorite option with probability $\frac{2}{3}$, and the other option with probability $\frac{1}{3}$. Under that MSNE, the payoff for F is 2 and the payoff for B is 2 as well.

Graphically, one can represent the best response functions of both players as follows:

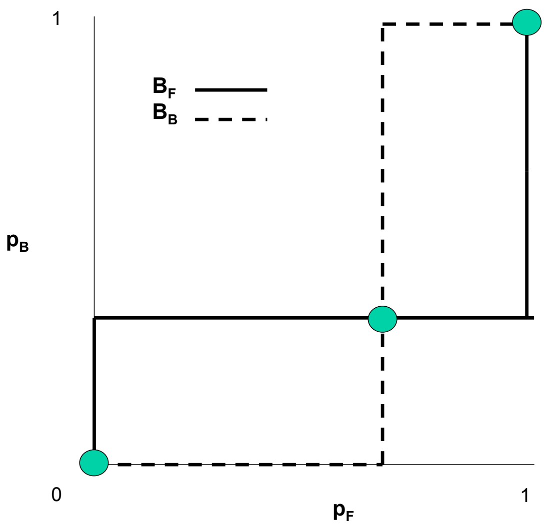

Figure 8: Best response functions in the Coordination Game

The graph can be read as follows: for any $p_F < \frac{2}{3}$, B prefers to play P, so $p_B = 0$. For $p_F = \frac{1}{3}$, B is indifferent between P and C, so p_B can be anywhere between 0 and 1. For $p_F > \frac{2}{3}$, then B prefers to play C, so $p_B = 1$.

You can check that all Nash equilibria are at crossing points of the best response functions, as implied by the definition of Nash equilibria in terms of BRF.

Palacios-Huerta (2003)[20] is a good introduction to the debate over whether agents are actually able to play mixed strategies. This is particularly important when the only NE is in mixed strategies, as in the game under study in the article (penalty strikes). Most of the evidence in experimental settings is not encouraging (agents find it difficult to choose actions randomly). However, Palacios-Huerta (2003) shows that professional footballers, who have sufficient experience and incentives, appear to be able to use mixed strategies in their choice of which side of the goal to strike a ball.

12 Extensive form games with perfect information

We consider in the following settings in which agents choose actions in succesion and observe what the other players played before choosing their own action.

[20]Palacios-Huerta I., 2003, Professionals play minimax, Review Of Economic Studies 70(2), 395-415.

Definition: An extensive form game with perfect information consists of:

1. Players: A set of players $N = 1, ..., i, n$ with typical member i.
2. Histories: A set of histories $H = (H_1, \ ..., \ H_n)$ with typical member $h = (h_1, ..., h_n)$. h is a sequence of actions by individual players, with typical member h_i which denotes all actions taken by i in the past.
   The start of the game is denoted as $\emptyset \in H$ is . If $h \in H$, but there is no $(h, a) \in H$ where a is an action for some player, then h is a terminal history. That means there is no further action to be taken by any of the players; the game is finished. The set of terminal histories is denoted as $Z \subset H$.
3. Player function: A function $P : H \backslash Z \to N$, assigning a player to each non-terminal history. This player function indicates when is the turn of each player to play.
4. Payoffs: vNM payoffs for each $i \in N$ are defined over terminal histories, $u_i : Z \to \mathbb{R}$. This assigns a payoff to each terminal histories of the game.

An extensive form game G is thus defined by N, H, P and $\{u_i\}_{i\in\mathbb{N}}$

To illustrate the above, suppose one plays a prisoners' dilemma game twice. h_1 at stage 1 may be (C), that is, player 1 played C in the first stage. h_1 at stage 2 may be (C,C), that is, player 1 played C in stage 1 and in stage 2. The set of terminal histories, over which payoffs will be defined, is:

{(C,C),(C,C)},{(C,C),(C,D)},{(C,C),(D,D)},{(C,C),(D,C)},

{(C,D),(C,C)},{(C,D),(C,D)},{(C,D),(D,D)},{(C,D),(D,C)},

{(D,C),(C,C)},{(D,C),(C,D)},{(D,C),(D,D)},{(D,C),(D,C)},

{(D,D),(C,C)},{(D,D),(C,D)},{(D,D),(D,D)},{(D,D),(D,C)}. For example, in the last history, player 1 played D in both stages, while player 2 played D in the first stage and C in the second.

Example: Game of entry deterrence

1. Players: $N =$Entrant (E), Incumbent (I).
2. Histories: $H : \{(\emptyset)$, (Stay out), (Enter), (Enter, Fight), (Enter, Accomodate)$\}$.
   The set of terminal histories is $Z : \{$(Stay out), (Enter, Fight), (Enter, Accomodate)$\}$
3. Player function: $P(\emptyset) = E$ and $P(Enter) = I$. This means E plays first, and I may play only if E enters, otherwise, I does not have anything to do.
4. Payoffs u_I(Stay out) $= 2, u_E$(Stay out) $= 0$, $u_I(Enter, Fight) = u_E(Enter, Fight) = -1$, u_I(Enter, Accomodate) $= u_E$(Enter, Accomodate) $=$ 1. Note how payoffs are defined over all possible histories/

The game can be represented in extensive form as follows:

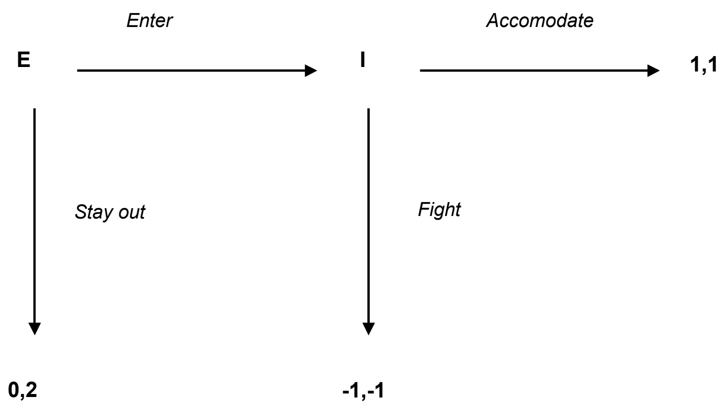

Figure 9: Extensive form representation of the game of entry deterrence.

The game can be represented in strategic / normal form as well, by getting rid of the information about histories and player functions.

Strategies $S : (S_E, S_I)$ are S_I : {Fight, Accomodate} and S_E : {Stay out, Enter} . Payoffs are defined over the terminal histories that result from a specific set of strategies, so I have to define equivalence, such that u(Stay out, Fight)=u(Stay out), u(Stay out, Accomodate)=u(Stay out), u(Enter, Fight)=u(Enter, Fight), u(Enter, Accomodate)=u(Enter, Accomodate).

One thus obtains the following strategic form representation of the game:

	Fight	Accomodate
Stay out	0,2	0,2
Enter	-1,-1	1,1

Once the game has been put into its strategic form, it is easy to find its pure strategy Nash Equilibria: {Stay out, Fight} and {Enter, Accomodate}. In order to find its mixed Nash equilibria, denote p_E the probability for the entrant to stay out, and p_I the probability for the incumbent to fight. Then

- u_E(Stay out,p_I) $= 0$,
- u_E(Enter,p_I) $= -p_I + (1 - p_I) = 1 - 2p_I$
- u_I(p_E,Fight) $= 2p_E - (1 - p_E) = 3p_E - 1$
- u_I(p_E,Accomodate) $= 2p_E + (1 - p_E) = p_E + 1$.

From those payoffs, the entrant enters whenever $p_I < \frac{1}{2}$, stays out whenever $p_I > \frac{1}{2}$ and is indifferent between entering or not when $p_I = \frac{1}{2}$. The incumbent fights whenever $3p_E - 1 > p_E + 1$, which never happens, is indifferent between fighting or not when $p_E = 1$, and accomodates for any $p_E < 1$.

There are therefore an infinity of Nash equilibria of the entry game, of which one PSNE, $\{\text{Enter}, p_I = 0\}$ and an infinite range of mixed strategy Nash equilibria denoted $\left\{\text{Stay out}, p_I \geq \frac{1}{2}\right\}$. Note that the PSNE $\{\text{Stay out}, p_I = 1\}$ is simply an extreme form of a MSNE where the incumbent fights with probability 1.

As in the above, the usual way to find equilibria of an extensive form game is to reduce it to a normal form game and find the NE of that normal form game. The equilibria of the extensive form game will be a subset of the NE of the normal form game, as the normal form game neglects important information about the succesion of actions in the game. We see below how the set of Nash equilibria can be parsed down through a process of *backward induction*.

12.1 Backward induction

From the above, the entrant stays out because she expects the incumbent to fight upon entry with probability at least half. However, since the entrant stays out, this is a completely arbitrary belief as there is no way for the entrant to verify that prediction. Is that arbitrary belief reasonable? When the entrant enters, then the incumbent will always find it best to accomodate. Therefore, the only credible belief for the entrant is that when the entrant enters, then the incumbent would accomodate. Therefore, it is not reasonable for the incumbent to believe that the incumbent would fight with any positive probability if she were to enter, because it is not credible for the incumbent to actually fight upon entry. One can thus state that the only credible equilibria of the game are those that survive a process of backward induction where, beginning from the last stage, actions that players will not choose in that stage are eliminated from consideration.

For example, in the game of entry deterrence, the incumbent will not fight when called upon to choose its action. Eliminating 'fight upon entry' from consideration, the game that remains is thus a choice for the entrant between staying out, which gives payoff 0, and entering, which, by backward induction, will give a payoff of 1. The entrant thus chooses to enter. The only equilibrium that survives this process of backward induction is $\{\text{Enter, Accomodate}\}$.

Definition: *Backward Induction:* A NE of a game survives a process of *backward induction* if it does not involve any non-credible threat, that is, if all strategies played at each stage (histories) in the game are best-response to each other.

For example, 'fight' is not best response to 'enter', so that any equilibrium that involves 'fight' in response to 'entry' does not survive backward induction.

Exercise: Consider the following *centipede game*:

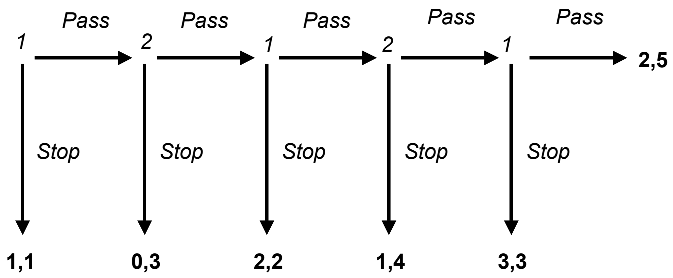

Figure 10: Centipede game

a) What NE of this game survive backward induction?

b) How would you play the game if you were not sure the other player follows a process of backward induction?

13 Extensive form games with imperfect information

Definition: An extensive form game with imperfect information consists of:

1. Players: As before
2. Histories: As before
3. Player function: As before
4. Payoffs: As before
5. Information: A player's information set defines what the player knows about the previous history of the game when he is called upon to play.

Definition: A singleton information set is such that a player knows the history of the game that led to the stage he/she is at. A subgame is a game that follows a singleton information set. A game can have many subgames, and a subgame can have subgames as well.

Example: In the Prisoners' Dilemma, no one of the players knows what the other player played so there is no subgame over than the whole game which follows from history $H = \emptyset$. In the entry deterrence game, the incumbent, once it is its turn to play, knows the entrant entered. Therefore, the game after entry where the incumbent has to choose whether to enter or not is a subgame. The entry deterrence game has two subgames: the one that follows from history $\{\emptyset\}$ and the one that follows from history {Enter}.

Definition: A Nash equilibrium is subgame-perfect ('SPNE') if the players' strategies constitute a Nash equilibrium in every subgame.

Example: The subgame perfect Nash equilibrium ('SPNE') in the prisoners' dilemma is also its NE. The subgame perfect Nash equilibrium in the entry deterrence game is {Enter, Accomodate}. The outcome of the SPNE in the centipede game is {Stop,Stop} (based on the SPNE {(Stop,Stop,Stop),(Stop, Stop)}).

In all those cases, the SPNE coincides with the NE obtained by backward induction, either because the game is one-stage only (Prisoners' Dilemma), or because there is perfect information, i.e. each decision node is a singleton information set, that is, players know with certainty the history that led to the current stage in the game when it is their turn to play.

From the above examples, one can say that the set of NE that survive backward induction and the set of *subgame perfect NE* coincide in *extensive form games with perfect information.*

13.1 Forward induction

The following example will allow us to introduce the process of forward induction, which is a process of elimination of candidate SPNE to further restrict the set of reasonable NE beyond the capabilities of the process of backward induction. A SPNE that is robust to the process of backward induction may not be robust to the process of forward induction.

Example: The following example introduces imperfect information in the Coordination game. The French player has the option to stay home, where she gets payoff 3.5 from eating a madeleine while reading a novel. The English player gets payoff 0 in that case. If the French player goes out, then the British player knows this (for example he calls her home and does not get an answer). However, he does not know if she went to the pub or to the café.

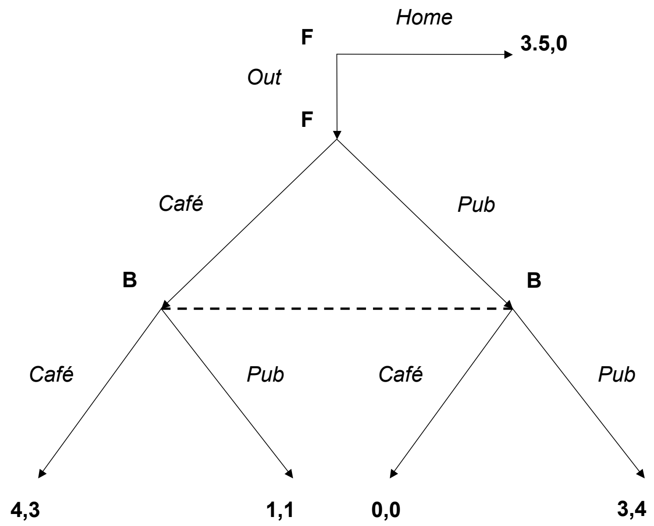

Figure 11: Coordination game with an outside option

Note that the payoff for F of staying at home makes her willing to go out only if her preferred PSNE obtains, but not anytime else (when the other PSNE is played, or when the MSNE is played). Intuitively, this means she will go out only if she believes the British player will go to the café, in which case she goes to the café, so the British player should go to the café if he realizes she went out. We are going to see if this intuition is borne out in the following process of finding out reasonable equilibria of the game:

This game that has three subgames, the one starting at $\{\emptyset\}$ (the whole game), the one starting at $\{\text{Home}\}$ and the one starting at $\{\text{Out}\}$. As seen before, there are two PSNE of the subgame starting at $\{\text{Out}\}$ as well as one MSNE. Consider thus all possible SPNE candidates once only NE of the subgame starting at $\{\text{Out}\}$ are retained.

1. $\{(Home, Café), Café\}$
2. $\{(Home, Pub), Pub\}$
3. $\left\{(Home, \frac{2}{3}), \frac{1}{3}\right\}$
4. $\{(Out, Pub), Pub\}$
5. $\{(Out, Café), Café\}$
6. $\left\{(Out, \frac{2}{3}), \frac{1}{3}\right\}$

Note that, for example in the first case, while F stays home, we still express what would happen if she did not. Backward induction eliminates the first potential equilibrium: if F expects $\{Café, Café\}$ to be played in the subgame starting at $\{Out\}$, then she would go out. In the same way, the fourth and the sixth equilibrium can be eliminated as well: if she expects {Pub,Pub} to be played in the subgame starting at Out, then she is better off staying in. 2, 3 and 5 are thus the only SPNE of the game.

In order to select among those, one will have to refine the concept of a SPNE by using a process of *forward induction*: the second equilibrium would require B to think that when F goes out, she goes to the pub, which would not be rational of F since she would have done better staying home in that case. Similarly, the third equilibrium requires B to think F plays a mixed strategy that obtains payoff of 2, even though in that case she would have done better staying at home. The only equilibrium that thus survives forward induction is $\{(Out, Café), Café\}$. In that equilibrium, F does not wish she could change any of her action at any stage of the game (*i.e.* once at the café, she is happy to have gone out). This confirms our intuition as stated at the beginning of the analysis of this game.

At this stage, we can define the process of forward induction:

Definition: *Forward induction:* To the difference of backward induction that starts from the end of the game and assumes that behavior at any stage of the game will be rational, forward induction starts from the beginning of the game and assumes that present behavior is optimal given what is expected to happen later. An equilibrium will survive the process of backward induction when it is such that the behavior of a player proves to have been optimal given the succeding strategies played in the game.

Remark: The 'Intuitive Criterion' of Cho and Kreps (1987)[21] uses forward induction to restrict the set of NE by eliminating those that are based on unreasonable out-of-equilibrium beliefs, i.e. those out-of-equilibrium beliefs that imply a player would play against his or her own interest if players deviated from the equilibrium. The problem with this approach is that, since actions that are out of equilibrium are not played, there is in principle nothing that could restrict the beliefs of someone once they are played. In fact, it may be rational to believe that a player that plays out of equilibrium is actually irrational, and thus hold out any beliefs about how he might play or what are his beliefs!

14 Bayesian Games

14.1 Example

A Bayesian game differs from a game with imperfect information in that, while in games with imperfect information the players may not know what was played before, in Bayesian games players may not know who they are playing against. In practice, this means they will not know what the other player will play in

[21] Cho I-K. and D.M. Kreps, 1987, Signaling games and stable equilibria, Quarterly Journal of Economics 102(2), 179-221

advance of the game, which is quite different from not being perfectly informed of what is played during the game. Before going on to define further what is a Bayesian game, let us study a specific Bayesian game, extended from our usual Coordination Game:

Example: The British person likes the French person with probability $\frac{1}{2}$, and dislikes her with probability $\frac{1}{2}$. The British person knows his own type, B_l (likes) or B_h (hates). The British person is very private and does not display his feelings openly, so the French person does not know the type of the British person and assigns probability $\frac{1}{2}$ to the British person liking her.

If B likes F, then payoffs are as follows:

		B	
		C	P
F	C	$4,3$	$1,1$
	P	$0,0$	$3,4$

If B does not like F, then payoffs are as follows:

		B	
		C	P
F	C	$4,0$	$1,4$
	P	$0,3$	$3,1$

The game can be put in strategic form version as follows:

- Players: $i \in N = \{F, B_l, B_h\}$
- For each i, $s_i \in \{C, P\}$

Payoffs can be shown as follows, with the first payoff the payoff of the French person, the second the payoff of the British person of type l and the third payoff the payoff of the British person of type h. Columns show all possible combinations of actions that B of type l or h might take.

		B			
		C_l, C_h	C_l, P_h	P_l, C_h	P_l, P_h
F	C	$4,3,0$	$\frac{5}{2},3,4$	$\frac{5}{2},1,0$	$1,1,4$
	P	$0,0,3$	$\frac{3}{2},0,1$	$\frac{3}{2},4,3$	$3,4,1$

This table is thus interpreted as follows: F can choose between C and P, and B can choose either C whether it is of type l or h, or C if it is of type l and P if it is of type h, or P if it is of type l and C if it is of type h, or P if it is of type l and P if it is of type h. In each combination of strategies, the expected payoff of F is first, then the payoff for B_l is second and the payoff for B_h is third.

By underlining best responses, one finds one pure strategy Nash Equilibrium, $\{C, C_l, P_h\}$. There also are Mixed Strategy Nash Equilibria: Suppose F mixes between C and P with probability p and $1-p$ respectively, while B_l mixes between C and P with probability p_l and $1-p_l$ respectively, and B_h mixes between C and P with probability p_h and $1-p_h$ respectively. Under those conditions,

- the expected payoff for F in playing C is $4p_lp_h + \frac{5}{2}p_l(1-p_h) + \frac{5}{2}(1-p_l)p_h + 1(1-p_l)(1-p_h)$
- the expected payoff for F in playing P is $0p_lp_h + \frac{3}{2}p_l(1-p_h) + \frac{3}{2}(1-p_l)p_h + 3(1-p_l)(1-p_h)$
- the expected payoff for B_l in playing C is $3p + 0(1-p)$
- the expected payoff for B_l in playing P is $1p + 4(1-p)$
- the expected payoff for B_h in playing C is $0p + 3(1-p)$
- the expected payoff for B_h in playing P is $4p + 1(1-p)$

B_l thus plays C s.t. $p > \frac{2}{3}$, B_h plays C s.t. $p < \frac{1}{3}$. There are thus five situations:

- If $p < \frac{1}{3}$, then B_h plays C and B_l plays P, but then F is better off playing C : this is not a Nash equilibrium.
- If $p = \frac{1}{3}$, then B_h is indifferent between playing C or P while B_l will play P. This is an equilibrium only if B_h 's strategy makes F indifferent between playing C or P, so I must have $p_h = \frac{2}{3}$.
- If $\frac{1}{3} < p < \frac{2}{3}$, then B_h plays P and B_l plays P. Then F would play P, but then B_h would play C. This is not a Nash equilibrium.
- If $p = \frac{2}{3}$, then B_h plays P and B_l is indifferent between C and P. This is an equilibrium only if B_l 's strategy makes F indifferent between playing C or P, so I must have $p_l = \frac{2}{3}$.
- If $p > \frac{2}{3}$, then B_h plays P and B_l plays C. Then F plays C, which is a Nash equilibrium.

There are thus two MSNEs, $\left\{\frac{1}{3}, P, \frac{2}{3}\right\}$ and $\left\{\frac{2}{3}, \frac{2}{3}, P\right\}$ in addition to the PSNE $\{C, C_l, P_h\}$.

14.2 Definition

In the following, we provide a formal definition of a Bayesian game.

Definition: A Bayesian game of incomplete information consists of

1. Players, labelled $i \in N = \{1, ..., n\}$
2. Types, for each player i, there is a set T_i of possibles types, labeled t_i. Player i might be of different types that have different payoffs, but all possible types of player i play the same role in the game.
3. Actions, for each player i, there is a set A_i of possible actions, labeled a_i. This means that whatever i's type, he has access to the same range of action.
4. Beliefs, for each pair (i, t_i), a probability distribution over T_{-i}, written $p_i(t_{-i}|t_i)$.This denotes what player $i$of type t_i believes the type of other player are.
5. Payoffs, for each player i, a vNM utility function $u_i : A \times T \to \mathbb{R}$. As said above, payoffs depend not only on the identity of the player, but on his type.

A Bayesian game can thus be written $\Gamma : \{N, \{T_i\}_i, \{A_i\}_i, \{p_i\}_i, \{u_i\}_i\}$

Remark: Beliefs are the beliefs of each player of each different type over the possible types of each other player. In the previous game, for example, F believed B could be of two types, l or h, with equal probability. Since there was only one type of F, one did not need to specify different beliefs for different types of F. Also, both B_l and B_h had the same belief about the type of F, since F could be of only one type.

Exercise: Formulate the Bayesian Coordination Game studied above as a Bayesian game of incomplete information according to the notations exposed above.

14.3 Examples

In the following, we consider two Bayesian games and find their Nash Equilibria. The first example is designed to show that if players differ in their types, then even though each type of players play pure strategies, what is observed by an external observer are mixed strategies. The second example considers auction mechanisms, where we show that auctions can be used as a tool to get players to reveal their own type.

14.3.1 Harsanyian Purification

In this part, we consider how MSNEs may arise when there is some uncertainty over the type of the agents playing the game. Consider the following Noisy Coordination Game which is a modification of the Coordination Game: The French person is not sure how much the British person likes going to the pub, and the British person is not sure how much the French person likes going to the café. Payoffs are as follows:

		B	
		C	P
F	C	$4+f, 3$	$1+f, 1+b$
	P	$0, 0$	$3, 4+b$

F is not sure about b, and believes it is distributed according to the uniform probability distribution over $(0, a) : U[0, a]$. Similarly, B is not sure about c, and believes it is distributed according to the uniform probability distribution over $(0, a) : U[0, a]$. Both F and B know their own type, f and b respectively.

Consider the following Bayesian strategy, such that B plays P s.t. $b > b^*$, and F plays C s.t. $f > f^*$. This is a very simple kind of threshold strategy whereby depending on their type, players will decide either to play C or B. Note that this is a considerable simplification from for example assuming that B would vary the probability with which he would play C or B depending on his type. Here, he plays pure strategies, either P or C. However, from the perspective of B, the probability that F plays C is then $\frac{a-f^*}{a}$ while from the perspective of F, the probability that B plays P is then $\frac{a-b^*}{a}$. Therefore, while players play pure strategies depending on their type, this is observationaly equivalent to mixed strategies from the point of view of the other player.

- The best strategy for F of type f is to play C s.t. $\frac{b^*}{a}(4+f)+(1-\frac{b^*}{a})(1+f) > (1-\frac{b^*}{a})3$, so F plays C s.t. $f > 2-6\frac{b^*}{a}$.
- Similarly, the best strategy for B of type b is to play C s.t. $(1-\frac{f^*}{a})3 > (1-\frac{f^*}{a})(1+b)+\frac{f^*}{a}(4+b)$, so B plays C s.t. $b < 2-6\frac{f^*}{a}$.

Both players are indifferent between playing C or P when $f = f^*$ and $b = b^*$, which leads to the following system of equations:

$$\begin{aligned} f^* &= 2-6\frac{b^*}{a} \\ b^* &= 2-6\frac{f^*}{a} \end{aligned}$$

which is solved for $f^* = b^* = \frac{2a}{6+a}$.

- From the point of view of F, the probability that B goes to the pub is then
$$\Pr(b \geq \frac{2a}{6+a}) = 1-\frac{2}{6+a}$$
- From the point of view of B, the probability that F goes to the café is
$$\Pr(f \geq \frac{2a}{6+a}) = 1-\frac{2}{6+a}$$

One will notice that as $a \to 0$ (uncertainty disappears), then each player can expect the other player to play her preferred choice with probability $\frac{2}{3}$, which corresponds to the MSNE of the game. This means that 'purification' of the pure-strategy Bayesian equilibrium of this game resembles the MSNE of the unperturbed game. Note however that the PSNE of the unperturbed game are also Bayesian equilibria of the perturbed game. This means that purification of pure-strategy Nash equilibria does not exclude any NE of the original game with no uncertainty.

14.3.2 Auctions

Auctions are a specific example of Bayesian games, whereby one or many buyers and one or many sellers must reach an agreement over the exchange of one or many goods. This agreement may be about exchange price, quantity exchanged, and the identity and allocation to sellers and buyers.

We will examine two very simple mechanisms:

1. One is such that a seller must decide to which of two buyers to sell one indivisible good in its possession, and at what price.

2. The other is such that one buyer is faced with one seller with one indivisible good, and must reach an agreement over the price at which that one indivisible good will be exchanged.

There are a variety of auction mechanisms which could be applied to either of the two situations, and there are also a variety of other situations where auction mechanisms could be applied. We will examine only the first price sealed bid auction as applied to the first situation, and the double auction as applied to the second situation.

1) First Price Sealed Bid Auction

Two buyers, i and j, submit sealed bids to a seller. The highest bidder wins the object, and pays his/her bid to the seller. The object is worth v_i to buyer i and v_j to buyer j. i knows v_i but not v_j, and j knows v_j but not v_i. i believes v_j is distributed uniformly over $[0,1]$, j believes v_i is distributed uniformly over $[0,1]$.

Strategies are of the form $b_i : [0,1] \to \mathbb{R}$, which for any v_i assigns a bid $b_i(v_i)$.

Payoffs can be represented as follows:

$$u_i(b_i; v_i) = \begin{cases} 0 \text{ if } b_i < b_j \\ \frac{v_i - b_i}{2} \text{ if } b_i = b_j \\ v_i - b_i \text{ if } b_i > b_j \end{cases}$$

We are going to looks for a linear bidding strategy, s.t. $b_i(v_i) = \alpha + \beta v_i$. Neglecting the case where $b_i = b_j$ (an event of probability measure 0), and assuming $b_j \in [\alpha, \alpha + \beta]$, the expected payoff to player j who is bidding b_j is then

$$\begin{aligned} E(u_j(b_j; v_j)) &= \Pr(b_j \geq \alpha + \beta v_i)(v_j - b_j) \\ &= \Pr(\frac{b_j - \alpha}{\beta} \geq v_i)(v_j - b_j) \\ &= \frac{b_j - \alpha}{\beta}(v_j - b_j) \end{aligned}$$

This is a concave function which attains a maximum for $b_j = \frac{v_j + \alpha}{2}$.

Suppose j indeed chooses $b_j = \frac{v_j + \alpha}{2}$. Then

$$\begin{aligned} E(u_i(b_i; v_i)) &= \Pr(b_i \geq \frac{v_j + \alpha}{2})(v_i - b_i) \\ &= \Pr(2b_i - \alpha \geq v_j)(v_i - b_i) \\ &= (2b_i - \alpha)(v_i - b_i) \end{aligned}$$

This is a concave function that attains its maximum for $b_i = \frac{\alpha + 2v_i}{4}$. Combining the two equations for b_i and b_j, one obtains $\alpha = 0$ and $\beta = \frac{1}{2}$. Each bidder thus bids half his/her own valuation for the good.

Remark: This linear bidding strategy is the unique symmetric Bayesian equilibrium if valuations are uniformly distributed. There may be other, non-symmetric Bayesian strategies, and if valuations are not uniformly distributed, then the linear bidding strategy is not necessarily optimal.

Exercise: Consider the three following alternative bidding rules, and solve for the optimal bidding strategy:

- The second price auction, whereby the highest bidder wins the good and pays the bid of the second highest bidder.
- The 'all pay' auction, whereby the highest bidder wins the good and both bidders pay their own bid.
- The 'loser pays' auction, whereby the highest bidder wins the good and the second highest bidder pays the bid of the highest bidder.

2) Double Auction

A buyer with valuation $v_b \sim U[0,1]$ for a good is faced with a seller with valuation $v_s \sim U[0,1]$ for a good. Both have to submit their bids (offered price for the buyer, asked price for the seller), simultaneously. If $p_b < p_s$ then there is no trade. If $p_b \geq p_s$, then they split the difference, so the trade price is $p = \frac{p_s+p_b}{2}$.

Payoffs are thus as follows:

$$u_b(p_b;v_b) = \begin{cases} 0 \text{ if } p_b < p_s \\ v_b - \frac{p_s+p_b}{2} \text{ if } p_b \geq p_s \end{cases}$$

$$u_s(p_s;v_s) = \begin{cases} 0 \text{ if } p_b < p_s \\ \frac{p_s+p_b}{2} - v_s \text{ if } p_b \geq p_s \end{cases}$$

As before, one will consider linear bidding strategies s.t. $p_b(v_b) = \alpha + \beta v_b$ and $p_s(v_s) = \gamma + \delta v_s$.

Consider the strategy of the buyer. The expected payoff to the buyer who is offering p_b is

$$E(u_b(p_b;v_b)) = \Pr(p_b \geq p_s)(v_b - \frac{E(p_s \mid p_b \geq p_s) + p_b}{2})$$

with

$$\begin{aligned} E(p_s \mid p_b \geq p_s) &= E(\gamma + \delta v_s \mid p_b \geq \gamma + \delta v_s) \\ &= \gamma + \delta E(v_s \mid \frac{p_b - \gamma}{\delta} \geq v_s) \\ &= \gamma + \delta \frac{p_b - \gamma}{2\delta} \\ &= \frac{1}{2}\gamma + \frac{1}{2}p_b \end{aligned}$$

so

$$E(u_b(p_b;v_b)) = \Pr(\frac{p_s-\gamma}{\delta} \geq v_s)(v_b - \frac{\frac{1}{2}\gamma + \frac{1}{2}p_b + p_b}{2})$$
$$= \frac{p_b-\gamma}{\delta}(v_b - \frac{\gamma+3p_b}{4})$$

This is a concave function that is maximized for $p_b = \frac{\gamma+2v_b}{3}$. Similarly, the optimal bidding strategy for the seller is to ask $p_s = \frac{\alpha+\beta+2v_s}{3}$. Combining the two equations for p_b and p_s with the assumed expression of their value, one obtains $\beta = \delta = \frac{2}{3}$. while $\alpha = \frac{\gamma}{3}$ and $\gamma = \frac{\alpha+\beta}{3}$. This is solved for $\gamma = \frac{1}{4}$ and $\alpha = \frac{1}{12}$. This means that both buyer and seller will "shade" their valuation by a factor of $2/3$ and there is a range of values of v_s and v_b for which no trade ever takes place.

Further work: When does trade take place? When is it optimal for trade to take place? Is the above mechanism efficient? Can you think of a more efficient mechanism for bilateral trade?

15 Repeated games

In this part, we will consider contexts in which the same game is played a repeated number of times, and we will show that repeating a game over several time periods allows players to sustain equilibria that would not be equilibria of the one-stage game and may improve per-period outcomes for both of them.

15.1 Finitely repeated games

In this part, we consider settings in which all players know when the game will end.

15.1.1 Example

Consider the following game:

		2		
		L	C	R
	T	$3,1$	$0,0$	$5,0$
1	M	$2,1$	$1,2$	$3,1$
	B	$1,2$	$0,1$	$4,4$

a) Suppose the game is played only once. Find all its Nash Equilibria.

Answer: This game has two PSNEs, (T, L) and (M, C). It also has mixed strategy equilibria as follows:

Suppose player 2 mixes between L, C and R with probabilities l, c and $1-c-l$ respectively. Then expected payoffs for 1 can written as follows:

$$u_1(T) = 3l + 5(1-c-l)$$
$$u_1(M) = 2l + c + 3(1-c-l)$$
$$u_1(B) = l + 4(1-c-l)$$

and one can see that $u_1(T) > u_1(B)$ so player 1 never plays B if player 2 plays a mixed strategy.

Consider now player 1 who mixes between T and M with probabilities t and m respectively. Then expected payoffs for 2 can be written as follows:

$$u_2(L) = t + m$$
$$u_1(C) = 2m$$
$$u_1(R) = m$$

If $m = 0$ then 2 would play L, while if $m > 0$, then $u_1(C) > u_1(R)$, so R is never played in a mixed strategy equilibrium. One is thus left with the following game to examine in order to find MSNEs of this game:

		1	
		L	C
2	T	$3,1$	$0,0$
	M	$2,1$	$1,2$

A MSNE of this game must be such that

$$3l = 2l + (1-l)$$
$$1 = 2(1-t)$$

so the unique MSNE of this game is such that $l = \frac{1}{2}$, $c = \frac{1}{2}$, $t = \frac{1}{2}$ and $m = \frac{1}{2}$.

b) Suppose now the game is repeated twice, that the players can observe the outcome of the first stage before the second stage begins and the per-period discount factor is δ. Show that the payoff $(4,4)$ can be achieved in the first stage in a pure-strategy subgame-perfect Nash equilibrium, and describe the conditional strategies which achieve this.

Answer: Consider 2's conditional strategy such that if 1 plays B in the first stage, 2 plays L in the second stage, and if 1 plays T in the first stage, 2 plays C in the second stage.

This strategy can be explained as such: 2 plays R in the first stage, hoping that 1 will play B. If 1 plays B indeed, then he is rewarded by 2 playing L in the second stage, when 1 will play T and get 3. If 1 plays T instead (best response to R), then he is punished by 2 playing C in the second stage, when 1 will play M and get only 1.

1 will conform to 2's expectations whenever the discounted payoff to 1 under this equilibrium, $4 + 3\delta$ is more than the deviation payoff of $5 + \delta$. This is true for $\delta \geq 1/2$.

Now, suppose that as under this equilibrium, 1 plays B in the first stage. Then the best response for 2 is to play R. There is thus no incentive under this equilibrium for 2 to deviate from the first stage action.

Therefore, there is a PSNE of the two-stage game such that:

- 1 plays B and 2 plays R in the first stage,
- If 1 played B in the first stage, then 2 plays L in the second stage (so 1 plays T in the second stage).
- If 1 played T in the first stage then 2 plays C in the second stage (so 1 plays M in the second stage).

This equilibrium is supported for any $\delta \geq 1/2$. One can check there is no incentive for players to deviate from its prescribed course of action.

15.1.2 Definition

Definition: A finitely repeated game with discounting of the stage game G over T periods is defined by

1. Strategic form game $G = \left\{N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N}\right\}$.
2. T, the number of periods over which the stage game G is played.
3. s^t the strategy profile played in period t.
4. $H = \bigcup_{t=0}^{t=T} S^t$, the set of possible histories in the game, with $S^0 = \emptyset$ the initial history.
5. $U_i = \sum_{t=1}^{t=T} \delta^{t-1} u_i(s^t)$, the set of payoffs of the game, with s^t the strategy profile played in period t, and δ the per-period discount factor.

$\boxed{\Gamma = \left\{G, H, \{U_i\}_{i \in N}\right\} \text{ defines the finitely repeated game based on } G}$

Note: A finitely repeated game is merely a special form of an extensive form game with perfect information where the player function is the same at every stage of the game and includes all players ($P(h) = N$ for any h in H).

Example: In the game given as our first example, then:

- $N = \{1, 2\}$,
- $S_1 = \{T, M, B\}$, $S_2 = \{L, C, R\}$, while
- $\{u_i\}_{i \in N}$ are as shown in the normal form representation of the game (for example, $u_2(T, L) = 1$).

The twice repeated game is such that

- $H = \overset{t=T}{\underset{t=0}{\cup}} S^t$, the set of possible histories in the game, has got elements such as for example $\{(T, L), (M, C)\}$.
- $U_i = \sum_{t=1}^{t=T} \delta^{t-1} u_i(s^t)$ with s^t the strategy profile played in period t. For example, in the example above, $U_1 = u_1(T, L) + \delta u_1(M, C) = 3 + \delta$.

15.1.3 Exercise

Consider the following game:

	W	X	Y	Z
A	$2,2$	$x,0$	$-1,0$	$0,0$
B	$0,x$	$4,4$	$-1,0$	$0,0$
C	$0,0$	$0,0$	$0,2$	$0,0$
D	$0,-1$	$0,-1$	$-1,-1$	$2,0$

where $x > 4$.

a) Assume that the game is played only once. Find all its pure strategy Nash equilibria.

Answer: There are three PSNEs: (A, W), (C, Y) and (D, Z).

b) Assume now that the game is repeated N times, that the players can observe past outcomes before the current stage begins and that the discount factor is δ.

i) Suppose that $N = 2$. For what values of x can the payoff $(4, 4)$ be achieved in the first stage in a pure-strategy subgame-perfect Nash equilibrium. Describe the strategies which achieve this.

Answer: (B, X) can be supported in the first stage if 1 does not deviate to A to get x and 2 does not deviate to W to get x. The best punishment if 1 deviates is for 2 to play Y in the second stage, and the best punishment if 2 deviates is for 1 to play D in the second stage. The reward if both conform is to play A and W respectively in the second stage. This equilibrium is supported if both 1 and 2 conform, which happens if $4 + 2\delta \geq x$, which is rewritten as $\delta \geq \frac{x-4}{2}$. Since $\delta \leq 1$, this is possible only subject to $x \leq 6$.

ii) Suppose that $x = 9$. What is the smallest number of N such that the payoff $(4, 4)$ can be achieved in the first stage in a pure-strategy subgame-perfect Nash equilibrium. Describe the strategies which achieve this.

Answer: Since $x = 9 > 6$ then (B, X) cannot be supported in the first stage of a two stage game (see (ii)). Now, the minimum N that supports (B, X) in the first stage must be such that $4+2\delta+2\delta^2+2\delta^3+\ldots 2\delta^{N-1} \geq 9$. If this is supported for $\delta = 1$, then this is sufficient for our proof. Now,

$4+2+2+2 > 9 > 4+2+2$, so we need only 4 periods to support (B, X) in the first period. If (B, X) is played in the first period, then this will be followed by (A, W) in all the following periods. If 1 deviates to A in the first period, then (C, Y) will be played in all the following periods, while if 2 deviates, then (D, Z) will be played in all the following periods.

15.2 Infinitely repeated games

Consider in this part situations in which the time horizon for both players is infinite, so they do not know when the game will end. In that case, solving by backward induction does not work, since there is no final period to start the backward induction from.

15.2.1 Example

Consider the following prisoners' dilemma game in normal form:

	C	D
C	$5,5$	$0,7$
D	$7,0$	$1,1$

The one-stage PSNE of this game is (D, D). This is the unique NE of this game. Consider now a finitely repeated game of this stage game, with end period T. One can check that by backward induction, the only SPNE of this stage game is to play (D, D) in all periods. Does this translate into the infinitely repeated version of the prisoners' dilemma game? One will see that under some conditions on the discount factor, the outcome (C, C) can be obtained in every stages of the repeated game if T is infinite.

Definition: *Convex Hull:* Define the Convex Hull as the range of payoffs that can be attained by playing pure-strategy combinations of the stage game, then the convex hull of the game can be represented as follows:

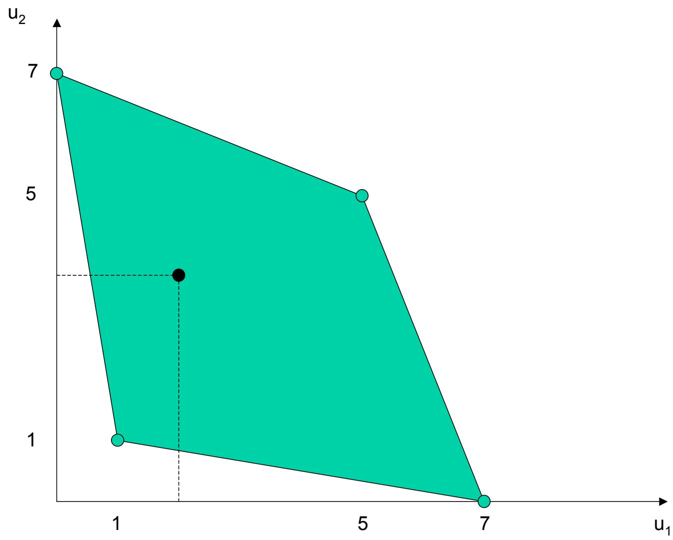

Figure 12: Convex hull for the prisoners' dilemma

For example, if players play (C, D) with probability $\frac{1}{2}$, (D, C) with probability $\frac{1}{4}$ and (D, D) with probability $\frac{1}{4}$, then 1 will get expected payoff of 2 and 2 will get expected payoff of $\frac{15}{4} = 3.75$. This is the point that is represented inside the hull.

15.2.2 The Nash-Threats Folk Theorem

Nash-Threats Folk Theorem: Every payoffs within the convex hull that are higher than the Nash equilibrium payoff profile for both players can be achieved as a SPNE of the infinitely repeated game subject to δ being high enough.

Remark: Note the proviso that δ be high enough. This means that some payoffs that may be achieved for some δ may not be achievable for some lower δ, and conversely. There are thus two ways to apply the theorem: either see what is the maximum payoff that may be achieved for a given δ, or see what level of δ is needed to attain a given payoff.

Example: The payoffs that are attainable can be represented as follows in the PD studied above:

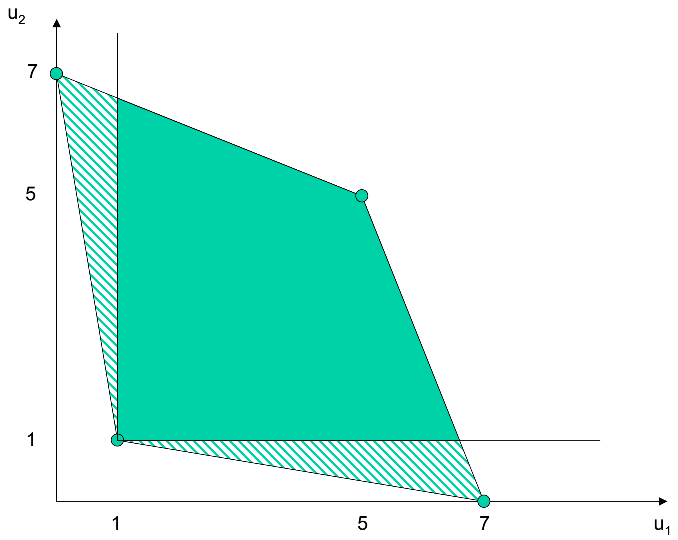

Figure 13: Set of attainable payoffs in the prisoners' dilemma

The convex hull minus its hashed areas is the set of attainable payoffs. We will see below how such payoffs may be achieved.

Example: In the PD under study, let us for example study how the payoff $(5, 5)$ can be attained.

Consider the following 'tit-for-tat' conditional stage game strategy: Play C in the first period. In the t^{th} period, and if the other player played C in all previous periods, then play C as well. If the other player played D in any of the previous periods, then play D.

Playing C will then be optimal only if in all previous periods both players always played C, and if playing C in this period obtains higher payoff than playing D.

Let us see therefore what is the payoff to playing D: If one plays D in this period, then one makes 7 this period, and 1 in all subsequent periods. If one plays C in this period, then one will also play C in the next period (the incentives are the same next period as in this period), so the expected payoff of playing C this period is 5 forever. Therefore, playing C this period is optimal s.t.

$$5 + 5\delta + 5\delta^2 + 5\delta^3 + ... = \frac{5}{1-\delta} \geq 7 + \frac{\delta}{1-\delta}$$

which translates in

$$\delta \geq \frac{1}{3}$$

As claimed in the theorem, one can attain payoff $(5,5)$, which is inside the convex hull and more than what is obtained in the NE of the stage game for both agents, s.t. δ high enough, in this case, $\delta \geq \frac{1}{3}$.

Extension: Suppose $\delta < \frac{1}{3}$, for example, $\delta = \frac{1}{4}$. What is the maximum payoff that can be attained? The theorem tells us that a payoff within the non-shaded area in the convex hull on the graph above can be attained, but how, and what payoff?

The 'how' is rather complicated, and involves assuming that players have access to publicly observable randomizing devices so that for example, they know when to play C and when to play D. This public randomizing device would draw their strategy at random according to pre-determined probabilities. It would be possible for either player to check the other player played what they were supposed to play. We can prove that the greatest payoff that may be achieved subject to a given δ is such that both players play symmetric strategies, i.e. both play C with the same probability. The probability p with which both 1 and 2 are asked to play C determines their payoff a. Indeed, I will have $a = 5p^2 + 7p(1-p) + (1-p)^2$.

1 will not deviate from playing C when told to do so s.t. $5p + \delta\frac{a}{1-\delta} \geq 7p + 1(1-p) + \frac{\delta}{1-\delta}$ (the constraint for 2 is the same). I can thus conclude, replacing a by its expression in p and rewriting the above equation, that I will need that p be set such that $\frac{a-1}{1+p} \geq \frac{1-\delta}{\delta}$. Note that, as stated previously, for any $\delta \geq \frac{1}{3}$ then $p = 1$ and full cooperation can be sustained. For $\delta = \frac{1}{4}$, no cooperation is sustainable ($p = 0$).

Part III

Incentive theory

Notations used in this lecture are drawn from Laffont J-J. and D. Martimort, 2002, The Theory of Incentives, Princeton University Press.

16 Introduction

In this lecture, we consider situations with asymmetric information. A principal is faced with an agent with whom he has to enter into a contractual relationship. He does not know either the level of ability of the agent (adverse selection, whereby inept agents are mixed with able ones) or he cannot control the level of effort expended by the agent into the delivery of the contracted performance (moral hazard then arises from the part of the agent, who may be tempted to slack). Issues such as those give rise to incentive (or contract) theory, whose aim is to examine how best to organize the principal-agent relationship in contexts of asymmetric information. Domains of application for contract theory are outlined, and some typical agency problems are presented, as well as their solution.

Asymmetric information is different from imperfect information. Instead of all agents on the market holding the same set of (possibly imperfect) information on every other agents in the market, some agents know more than others. Information that can be found out at some cost is not asymmetric information. Asymmetric information arises when agents are not motivated to reveal information they hold, for example because that information provides them with some advantage in a relationship. We will see how a principal can set up contractual relationships such as to motivate the agent to reveal that information, at some cost to the principal.

There are many different types of asymmetric information. Asymmetric information on the type of the agent leads to adverse selection problems. In that setting, for example, the fact you are healthy or not depends on outside factors and not on your habits. Asymmetric information on the effort of the players leads to moral hazard. In that case, for example, the fact you are healthy or not depends on your habits. In the first case, you will want to screen agents, in the second, you will want to monitor their habits or give them incentives to behave 'correctly', i.e. in your own interest.

Some examples of situations with moral hazard or adverse selection follow:

Example 1: When buying a car, the buyer does not know whether the quality of the car is good or bad. The seller/owner knows the quality of his car. Absent any way to prove the quality of his car, the seller of a good quality car will suffer a discount due to the presence of bad quality car sellers in the market. This might be an explanation for why cars lose so much value even right after being bought new.

Example 2: When selling car insurance, the insurer does not know whether the insured is careful or not in driving. The insured knows it, and those who drive badly will be more motivated to buy insurance than others. In the limit, the insurer should never sell insurance because only if it underestimated the risk of the insured would it sell. The same type of problem arises in borrowing and in health insurance, as well as in many used goods markets.

Example 3: In education, is a bad grade due to low effort by the student (moral hazard on the part of the student), to his/her intellectual limitations (type of the student) or to bad teaching (which can be either due to low effort on the part of the teacher or to his/her low ability)?

Example 4: In the current credit crunch, is the poor performance of banks due to bad management of the banks, opportunistic behavior on the part of bankers, bad macroeconomic policies on the part of the governement, or poor regulatory oversight? The following Dilbert character seems to go for the later option:

Figure 14: Poor regulatory oversight (c) Scott Adams

It is often difficult in each of those example to agree on who has the most information (e.g. health). There are a variety of contracts that are signed between agents in those settings, which are meant to alleviate asymmetric information. For example, bonus/malus systems for car insurance are used in France: if you have an accident, your insurance price increases, which deters drivers from driving carelessly. US insurance companies impose medical checkups to screen their insurees. Companies offer return guarantees so as to prove their trust in the quality of their goods.

17 Readings

17.1 Textbook Readings.

- Kreps, Chs. 16-17
- Varian, Ch. 25
- Mas-Colell, Chs.13-14

The following textbooks might also prove useful to you:

- Laffont J-J. and D. Martimort, 2002, The Theory of Incentives; the Principal-Agent Model, Princeton University Press.
- Salanié B., 2005, The Economics of Contracts, MIT Press.

17.2 Articles[22]

- Akerlof G.A., 1970, The market for "lemons": Quality uncertainty and the market mechanism, The Quarterly Journal of Economics 84(3), 488-500.
- Grossman S.J. and O.D. Hart, 1983, An analysis of the Principal-Agent problem, Econometrica 51(1), 7-45.
- Puertz R. and A. Snow, 1994, Evidence on Adverse Selection, Equilibrium Signalling and Cross-Subsidization in the Insurance Market, Journal of Political Economy 102(2), 236-257.
- Rees R., 1985, The Theory of Principal and Agent, Parts I and II in Hey, J. & Lambert P. eds, Surveys in the Economics of Uncertainty, Blackwell; also in Bulletin of Economic Research.
- Spence A.M., 1973, Job Market Signalling, The Quarterly Journal of Economics 87(3), 355-374.
- Taylor C.R., 1999, Time-on-the-Market as a Signal of Quality, Review of Economic Studies 66, 555-578.

The theory of games with incomplete information and the theory of asymmetric information can be applied to a range of economic problems, as follows:

a) Regulation

- Laffont J. J. and J. Tirole, 1993, A Theory Of Incentives In Procurement And Regulation, The MIT Press. (This is a comprehensive book on the theory of regulation.).

b) Auctions

- Klemperer P., 2000, Why Every Economist Should Learn Some Auction Theory, Microeconomics 0004009, Economics Working Paper Archive at WUSTL, available on the web.

c) Theory of the firm

- Fama E.F., 1980, Agency problems and the theory of the firm, Journal of Political Economy 88(2), 288-307.
- Holmstrom B.R. and S.N. Kaplan, 2001, Corporate Governance and Merger Activity in the U.S.: Making Sense of the 1980s and 1990s, MIT Dept. of Economics Working Paper No. 01-11.
- Jensen M.C. and W.H. Meckling, 1976, Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure, Journal of Financial Economics 3(4), 305-360.

[22]Some articles in this list were contributed by previous teachers in MSc Economic Theory 1 at the UEA.

- Milgrom P. and J. Roberts, 1992, Economics, Organization and Management, Prentice Hall (This book is well written, accessible and motivating and of great help to future managers).

d) Labour Economics

- Green J. and N.A. Stokey, 1983, Comparison of Tournaments and Contests, Journal of Political Economy 91, 349-364.
- Hart O.D., 1983, Optimal Labour Contracts under Asymmetric Information: An Introduction, Review of Economic Studies 50(1), 3-35.
- Holmstrom B., 1982, Moral Hazard in Teams, Bell Journal of Economics 13(2), 324-340.
- Laffont J-J. and M.S. Matoussi, 1995, Moral Hazard, Financial Constraints and Sharecropping in El Oulja, Review of Economic Studies 62, 381-399.

e) Banking and Insurance

- Freixas X. and J-C. Rochet, 1997, Microeconomics of Banking, The MIT Press, second print 1998 (An excellent review for those interested).
- Rothschild M. and J.E. Stiglitz, 1976, Equilibrium in competitive insurance markets: An essay on the economics of imperfect information, Quarterly Journal of Economics 90(4), 629-649.
- Stiglitz J.E. and A. Weiss, 1981, Credit Rationing in Markets with imperfect information, American Economic Review 71, 393-410.

f) Health Economics

- Arrow K.J., 1963, Uncertainty and the welfare economics of medical care, American Economic Review 53(5), 941-973.
- Zechauser R., 1970, Medical Insurance: A case study of the trade off between risk spreading and appropriate incentives, Journal of Economic Theory 2, 10-12.

18 Agency and adverse selection.

In this part, a typical model of adverse selection will be presented and the timing of the contractual relationship will be outlined. The first best outcome (which arises in contexts of perfect information) will be determined graphically and compared with the second best outcome (with asymmetric information). The role of incentive and participation constraints will be explained and the maximization program of the principal will be determined as well as its graphical solution. That solution will be compared with the first best.

18.1 The model

A principal gets utility $S(q)$ from obtaining quantity q of a product that is produced by an agent at cost $C(\theta, q) = \theta q$. θ is the type of the agent, known to the agent and unknown to the principal, who expects $\theta = \theta_H$ with probability $1 - v$ and $\theta = \theta_L$ with probability v, with $\theta_H > \theta_L$.

The contracting variables C are (q, t), t being the (money) transfer from the principal to the agent and q the quantity produced by the agent. Contract C is enforceable under law, that is, a third party (a Court) can check a contract was signed and can check whether t and q that were determined in the contract were paid and/or produced or not.

Under contract C, the principal's utility is then $S(q) - t$. Agent of type θ obtains utility is $t - \theta q$.

The timing of the contractual relationship is as follows:

- At $t = 1$, the agent learns its type θ.
- At $t = 2$ the principal offers a (menu of) contract(s) C (C_1, C_2, C_3, ...).
- At $t = 3$ the agent chooses a contract or chooses not to enter a contractual relationship ($C = C_0 = (0, 0)$ is always an option, as no one can be forced into signing a contract).
- At $t = 4$ the contract is executed: the agent produces the quantity as agreed in the contract and the principal pays her the transfer as agreed in the contract.

The contracting timeline can thus be represented as follows:

$t = 1$	$t = 2$	$t = 3$	$t = 4$
Agent learns type θ	Principal offers contracts C	Agent chooses contract	q is produced t is paid

18.2 The first best (θ known)

In a 'first best' world, the principal knows θ; he can therefore ask any quantity from the agent provided the agent accepts. Knowing the type θ of the agent, the program of the principal is:

$$\max_{q,t} S(q) - t \tag{31}$$
$$\text{s.t. } t - \theta q \geq 0$$

The principal thus maximizes the following Lagrangian:

$$\max_{q,t} (S(q) - t) + \lambda(t - \theta q) \tag{32}$$

which is maximized for

$$S'(q) - \lambda\theta = 0 \tag{33}$$
$$-1 + \lambda = 0$$
$$t - \theta q = 0 \tag{34}$$

so that

$$\lambda = 1 \tag{35}$$
$$S'(q) = \theta$$
$$t = \theta q \tag{36}$$

Since $\lambda > 0$, the rent left to either type is 0 so type H will be offered contract C_H such that $t_H - \theta_H q_H = 0$ while agent of type L will be offered contract C_L such that $t_L - \theta_L q_L = 0$.

Given that $\lambda = 1$, the objective of the principal is then:

$$\max_q S(q) - \theta q \tag{37}$$

When faced with an agent of type H, the principal will set q_H s.t.

$$S'(q_H) = \theta_H \tag{38}$$

and when faced with an agent of type L, he will set q_L s.t.

$$S'(q_L) = \theta_L \tag{39}$$

He will thus offer contract

- $C_H = (q_H, \theta_H q_H)$ to agent of type H and contract
- $C_L = (q_L, \theta_L q_L)$ to agent of type L,

with $q_L = S'^{-1}(\theta_L)$ and $q_H = S'^{-1}(\theta_H)$.

The outcome of the maximization of the principal's objective function can be represented graphically as follows:

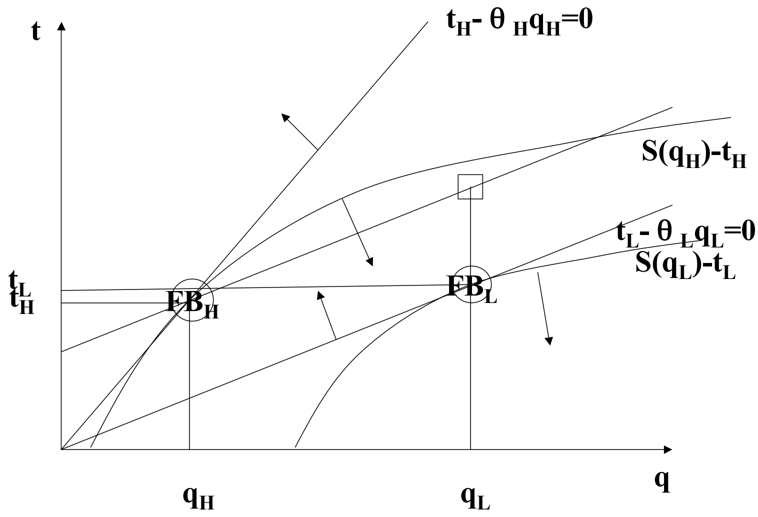

Figure 15: The graphical derivation of the first best outcome

Interpretation of the graph: The arrows indicate the direction in which utility of the agent and of the principal, respectively, are maximized. Utility P for the principal is represented by the iso-utility convex function such that $S(q) - t = P$, which determines all combinations of q and t such that utility P is achieved. Utility A for the agent is represented by iso-utility curve $t - \theta q = A$. In the equilibrium of the first best, A is equal to 0 for both types of agents. Utility for the principal is maximized at the point where the iso-utility convex function of the principal is tangent with the iso-utility function of the agent, i.e. the point such that $t - \theta q = 0$ and $S(q) - t$ is maximized.

18.3 The second best

18.3.1 Incentive compatibility

In the previous graph, one can observe that the utility of the good type is higher if she chooses the first best contract of the bad type rather than the contract that was destined for her. Therefore, the first best is not implementable if the type of the agent is not known to the principal. Under the second best, each type of agent must choose willingly the contract that is destined for him/her, so that the following incentive constraints ('IC') must be verified:

$$t_H - \theta_H q_H \geq t_L - \theta_H q_L \tag{IC$_H$}$$
$$t_L - \theta_L q_L \geq t_H - \theta_L q_H \tag{IC$_L$}$$

Under those conditions, type H chooses contract C_H and type L chooses contract C_L.

18.3.2 Participation constraints

We must also have

$$t_H - \theta_H q_H \geq 0 \tag{IR$_H$}$$
$$t_L - \theta_L q_L \geq 0 \tag{IR$_L$}$$

if we want both types of agents to participate. Those are the individual rationality ('IR') constraints.

Notations: Denote $U_H = t_H - \theta_H q_H$ and $U_L = t_L - \theta_L q_L$ and denote $\Delta = \theta_H - \theta_L > 0$.

18.3.3 The program of the principal

The program of the principal is then

$$\max_{t_L, q_L, t_H, q_H} v(S(q_L) - t_L) + (1 - v)(S(q_H) - t_H) \tag{40}$$

s.t.

$$U_H \geq U_L - \Delta q_L \tag{41}$$
$$U_L \geq U_H + \Delta q_H \tag{42}$$
$$U_H \geq 0 \tag{43}$$
$$U_L \geq 0 \tag{44}$$

How can the program be solved?

- First, note that adding up the participation constraint for both types, one obtains:

  $$U_H + U_L \geq U_L + U_H + \Delta(q_H - q_L) \tag{45}$$

  which means that

  $$q_H \leq q_L \tag{46}$$

  The high type will produce less than the low type in the second best outcome.

- Second, note that if the participation constraint on type H is verified and the incentive constraint for type L is verified, then the participation constraint for type L is verified. The participation constraint for type L is thus superfluous.
- Third, suppose now the incentive constraint for type L is binding. Then $U_L = U_H + \Delta q_H$, which, translated into the incentive constraint for type L, means I must have $U_H \geq U_H + \Delta(q_H - q_L)$, which is always the case as $q_H \leq q_L$. Therefore, the incentive constraint for type L is also superfluous.

From this, I can conclude that the participation constraint for the high type will be binding ($U_H = 0$) while the incentive constraint for the low type will be binding ($U_L = \Delta q_H$). This ensures all constraints are verified while minimizing the rent given out to each type. Taking the two above inequalities and replacing t_H and t_L by their expression in terms of q_L and q_H, the program of the principal can thus be simplified into:

$$\max_{q_H, q_L} v(S(q_L) - \theta_L q_L) + (1-v)(S(q_H) - \theta_H q_H) - v\Delta q_H \tag{47}$$

Maximizing the above with respect to q_L and q_H one obtains

$$S'(q_L) = \theta_L \tag{48}$$

and

$$S'(q_H) = \theta_H + \frac{v}{1-v}\Delta \tag{49}$$

Note how the result is distorted compared to the first best maximization program: There is no distortion for the good type who will produce the same as in the first best, while the bad type will produce less than in the first best.

The outcome of the maximization of the principal's objective function can be represented graphically as follows:

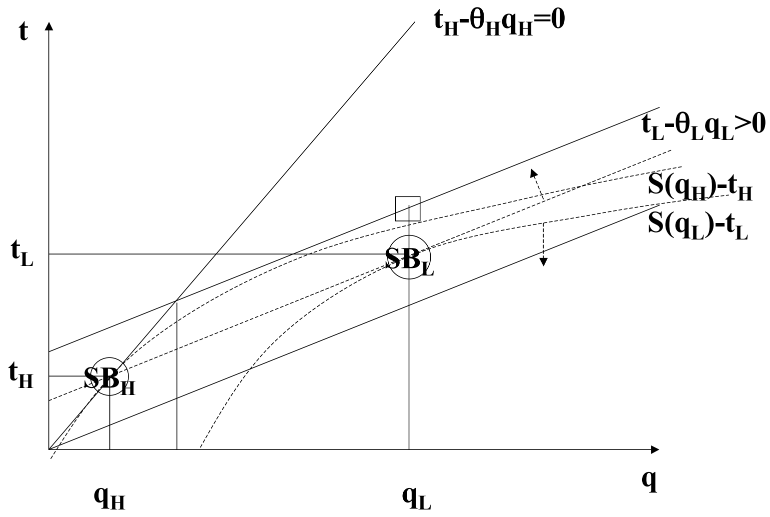

Figure 16: The graphical derivation of the second best outcome

Interpretation of the graph: No rent is left to the bad type, who produces less than the optimal level. The good type earns as much as what she would earn if she chose the contract for the bad type, and she produces the first best (optimal) level. The principal trades off between lowering the rent extracted by the good type and still producing close to the efficient level. By lowering the quantity asked from the bad type, the principal can decrease the rent that the good type must be given. This is how we go from the transfers and quantities that would be necessary to attain the first best, to the transfers and quantities the principal asks for at the second best optimum.

19 Agency and moral hazard

In situations with moral hazard,

- The uncertain outcome is endogenous to the situation at hand (success or failure depend on effort which depends on the contract that is signed).
- The result from the contract is a noisy signal for effort, i.e. it is not possible to know what effort was exerted from the result that was obtained.
- Transfers from the principal to the agent are constrained by the fact the agent may have limited liability for its actions (i.e. there are limits on what punishment you can impose in case of failure).

19.1 The model

Consider an agent and a principal. The agent is asked to execute a task by the principal, and in that task the agent may exert effort e either 0 (no effort) or 1 (effort). The cost of effort to the agent is $\psi(0) = 0$ for no effort and $\psi(1) = \psi > 0$ for effort. The utility to the agent from getting transfer t and exerting effort e is

$$U(t, e) = u(t) - \psi(e) \tag{50}$$

Production, which is what the principal is interested in, can be either q_H or q_L, with $q_H > q_L$, and the situation is such that

$$\Pr(q = q_H | e = 0) = p_0 \tag{51}$$
$$\Pr(q = q_H | e = 1) = p_1 \tag{52}$$

Success probabilities depends on effort and I will assume $p_1 > p_0$ so that a higher effort brings about a higher probability of success – otherwise, why encourage effort?.

The principal's utility function is

$$V(q, t) = S(q) - t \tag{53}$$

with $S(q)$ convex. The principal offers contracts such that the transfer will be conditioned on the result:

- $t(q_H) \equiv t_H$ will be paid to the agent if q_H is observed,
- $t(q_L) \equiv t_L$ will be paid if q_L is observed.

The timing of the relationship is as follows:

- At time $t = 1$, the principal offers contracts (t_H, t_L).
- At time $t = 2$, the agent accepts or reject the contract(s). If the contract(s) is/are rejected, the game ends.
- At time $t = 3$, the agent decides whether to exert effort or not.
- At time $t = 4$, the outcome is realized
- At time $t = 5$ transfers are made according to the contract and the realization of the outcome.

The contracting timeline can thus be represented as follows:

$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$
Principal offers t_H, t_L	Agents accepts or rejects	Agent exerts effort or not	Outcome is realized	Payment is made

19.2 The first best outcome (perfect information on effort)

Suppose the principal can observe effort and thus condition payment not on the outcome but on whether effeot was exerted or not.

His objective function is:

$$\max_{q_H, t_H, q_L, t_L} V = p_1(S(q_H) - t_H) + (1 - p_1)(S(q_L) - t_L) \tag{54}$$

$$\text{s.t. } p_1 u(t_H) + (1 - p_1)u(t_L) - \psi \geq 0 \tag{55}$$

if he wants the agent to exert effort (the constraint ensures the agent's ex-post anticipated utility from accepting the contract and exerting effort is positive), and

$$\max_{q_H, t_H, q_L, t_L} V = p_0(S(q_H) - t_H) + (1 - p_0)(S(q_L) - t_L) \tag{56}$$

$$\text{s.t. } p_0 u(t_H) + (1 - p_0)u(t_L) \quad \geq \quad 0 \tag{57}$$

if instead the principal wants the agent to exert no effort (this can happen if compensating for a high effort is too expensive).

If effort is desired, and denoting λ the Lagrange multiplier, the Lagrangian can be written:

$$\begin{aligned} L = p_1(S(q_H) - t_H) + (1 - p_1)(S(q_L) - t_L) \\ +\lambda[p_1 u(t_H) + (1 - p_1)u(t_L)] \end{aligned} \tag{58}$$

Maximizing with respect to t_H and t_L, one obtains:

$$-p_1 + \lambda p_1 u'(t_H^*) = 0 \tag{59}$$

$$-(1 - p_1) + \lambda(1 - p_1)u'(t_L^*) = 0 \tag{60}$$

This means that $t_H^* = t_L^* \equiv t_1^*$ and t_1^* is set such that the agent gets 0 utility so $u(t_1^*) = \psi$.

The same result is obtained if the principal chooses not to induce effort, but t_0^* is set such that $u(t_0^*) = 0$.

The principal will chooses to induce effort only if:

$$p_1 S(q_H) + (1 - p_1)S(q_L) - t_1^* \geq p_0 S(q_H) + (1 - p_0)S(q_L) - t_0^* \tag{61}$$

which can be rewritten:

$$(p_1 - p_0)(S(q_H) - S(q_L)) \geq t_1^* - t_0^* \tag{62}$$

19.3 The second best outcome

Suppose the agent has infinite wealth, so any t_L (punishment in case of bad outcome) can be asked for. Suppose also the principal cannot observe the agent's effort or equivalently, that he cannot prove effort was low. Then the agent must be induced to choose the desired level of effort, so incentive constraints must be verified as well as participation constraints. Suppose high effort is desired. The agent's incentive constraint is then such that

$$p_1 u(t_H) + (1 - p_1)u(t_L) - \psi \geq p_0 u(t_H) + (1 - p_0)u(t_L) \tag{63}$$

and the participation constraint is such that

$$p_1 u(t_H) + (1 - p_1)u(t_L) - \psi \geq 0 \tag{64}$$

Suppose for simplicity that $u(t) = t$ (the agent is risk neutral).

Solving the principal's maximisation program, the two constraints (Incentive and Participation) will be binding. This gives a system of two equalities with two unknown,

$$p_1 t_H + (1 - p_1)t_L - \psi = p_0 t_H + (1 - p_0)t_L \tag{65}$$
$$p_1 t_H + (1 - p_1)t_L - \psi = 0 \tag{66}$$

which has an unique solution:

$$t_L^* = -\frac{p_0 \psi}{p_1 - p_0} \tag{67}$$
$$t_H^* = \frac{(1 - p_0)\psi}{p_1 - p_0} \tag{68}$$

So the expected payment for the principal is

$$p_1 t_H^* + (1 - p_1)t_L^* = \psi \tag{69}$$

This is the same as would be paid if effort was observable, i.e. the principal merely compensates the agent for her effort. One would also see that the same expected payment (0) as in observable effort would be made in case no effort was needed. This means that even with asymmetric information on effort, the first best level of effort is implemented.

Note however this holds only when the agent is risk neutral and has infinite wealth. Let us now see what happens if there is a liability constraint or if the agent is risk averse.

19.3.1 Liability constraint

Suppose the limited liability constraint is L, so the agent cannot lose more than L. One will then have to impose limited liability constraints such that $t_H > -L$ and $t_L > -L$. L can be interpreted as the initial wealth of the agent.

If $-\frac{p_0\psi}{p_1-p_0} > -L$, then the problem is unchanged (the first best outcome can be implemented).

If $-\frac{p_0\psi}{p_1-p_0} < -L$, then the outcome without liability constraint cannot be achieved, and one obtains that

$$t_L^* = -L \tag{70}$$

$$t_H^* = -L + \frac{\psi}{p_1 - p_0} \tag{71}$$

As in the previous program, $t_H^* = t_L^* + \frac{\psi}{p_1-p_0}$, but t_L^* is now raised to $-L$. The agent will get a limited liability rent, as her expected utility is now

$$EU = p_1 t_H^* + (1-p_1)t_L^* - \psi \tag{72}$$

$$= p_1 \frac{\psi}{p_1 - p_0} - L - \psi \tag{73}$$

$$= (p_1 - p_0)\frac{\psi}{p_1 - p_0} - L - \psi + p_0 \frac{\psi}{p_1 - p_0} \tag{74}$$

$$= -L + p_0 \frac{\psi}{p_1 - p_0} > 0 \tag{75}$$

The principal's utility is then

$$V = p_1(S(q_H) - t_H^*) + (1-p_1)(S(q_L) - t_L^*) \tag{76}$$

$$= p_1(S(q_H) - \frac{\psi}{p_1 - p_0}) + (1-p_1)S(q_L) - L \tag{77}$$

which is less than the first best. Because of limited liability constraints in moral hazards problems, the principal may not want to induce effort even when this would be optimal in the first best with perfect observation of effort.

19.3.2 Risk aversion

We assumed up to now that $u(t) = t$, that is, the agent was risk neutral. If the agent is risk averse, then $u(t)$ is strictly concave. Then the principal will want to limit the difference between t_L and t_H because this then makes the contract less risky, and thus more acceptable to the agent. Risk aversion thus induces a distortion from the first best, that is, there are situations where the principal will not want to induce effort (it is too costly to do so) even though this would be optimal under full information, and would occur under risk neutrality.

20 Extensions

This part examines some hidden assumptions made in the previous parts. In particular, we assumed that:

1. Agents are able to commit to fulfilling the terms of the contract they sign. What happens then if agents cannot commit not to renegotiate a contract once the type of the agent (adverse selection) or the outcome of his effort (moral hazard) is known?

2. A court of law could be called on to enforce the terms of the signed contracts. What happens then if the principal may renege on the contract?

We will therefore look at those issues and see how standard theory applies in those alternative settings when the principal cannot commit. We will also examine other possible extensions, such as what happens if the principal is better informed than the agent, rather than the opposite.

20.1 Repeated adverse selection.

The following model with repeated adverse selection illustrates what happens if the principal cannot commit not to renegotiate after a contract is accepted or after an action is taken by the agent (chapter 9, section 3, Laffont and Martimort, 2002). Consider thus the standard adverse selection model, and assume it is repeated twice with the same agent and the same principal both period.

$$V = S(q_1) - t_1 + \delta(S(q_2) - t_2) \tag{78}$$

is the total utility of the principal over the two periods (the second period payoffs are discounted by δ).

$$U = t_1 - \theta q_1 + \delta(t_2 - \theta q_2) \tag{79}$$

is the total utility of the agent over the two periods, with the second period payoff discounted by δ.

Assume the principal can commit not to renegotiate. Then the optimal contract is to repeat the optimal one-period contract in the two period. However, this then calls for the bad type to produce $q_H^{SB} < q_H^*$. The principal would then gain from renegotiating after the first period, once agents revealed their type through their choice of contract, so as to get the bad type, H, to produce the first best level q_H^*. However, if the principal did this, the good type would then have to be compensated for not taking the contract of the bad type. Therefore, if the principal cannot commit not to renogotiate type H 's contract, then she must leave a higher surplus to the good type than if it could commit.

Denote $q_{\tau t}$ and $t_{\tau t}$ the quantity and transfers asked from and paid to type $\tau = H, L$ in period $t = 1, 2$. For good type (L) not to choose the contract designed for the bad type if the bad type's contract is due to be renegotiated in the second period so the bad type would produce q_H^*, one must have:

$$t_{L1} - \theta_L q_L^* + \delta(t_{L2} - \theta_L q_L^*) \geq t_H - \theta_L q_{H1} + \delta(t_H - \theta_L q_H^*)$$

which can be rewritten as

$$U_L \geq U_H + \Delta(q_{H1} + \delta q_H^*)$$

with U_L (U_H) the intertemporal utility of the good (bad) type, the right hand side of the equation the utility of the good type if she chose the contract designed for the bad type, and as usual, $\Delta = \theta_H - \theta_L$.

This is to be compared with the constraint with full commitment (if the principal can commit not to renegotiate)

$$U_L \geq U_H + \Delta(q_{H1} + \delta q_{H2})$$

Since $q_H^* > q_{H2}$ the first constraint is more stringent than the second, which means the principal loses from his inability to not to renegotiate. There may be a loss of efficiency from that situation if δ is high. Indeed, the principal may then prefer not to know the types of the agents in the first period (and thus offer a pooling contract in the first period, whereby both types choose the same contract), so as not to learn the type of the agent in the first period, and thus not have any basis for a renegotiation in the second period, when the optimal one period contract will be offered.

20.2 The hold up problem

In the following, we explore what happens in the more serious problem where the principal cannot commit not to renege on a contract. This may be for example because there is no court of law that can force him not to do so. In a moral hazard setting, this leads to a hold up problem (chapter 9, section 4, Laffont and Martimort, 2002). Take indeed the standard moral hazard problem. The principal would like to commit to pay more when the result is good. However, after the effort is made and a high result is achieved, the principal's self interest is to renege on the contract and not compensate the agent for his effort. This problem can happen in a wide variety of settings even when an efficient court of law is present. For example, this can happen if no formal contract was signed beforehand or if the result to be achieved was ill-defined. This hold-up problem is particularly prevalent for the State, especially if the judiciary is not independent of the executive. The cost of forcing the State's to abide by its commitments can indeed be prohibitive. It is for example too easy for the State to change its taxation policy to overturn previous pro-business commitments, or to withdraw financial incentives to firms that adopted socially or environmentally-responsible practices once firms have changed their way of doing business.

In those cases, the result of the principal's inability to commit not to renege on a contract is even worse for the principal than in the case where there was inability to commit not to renegotiate: no contract can be signed with the agent and agents exert only the lowest effort level. The problem may be mitigated if rather than signing contract, the agent and the principal bargain after the state of nature is revealed. It may also be mitigated if the agent knows what is the state of nature before signing the contract and exerting effort.

20.3 Other extensions

20.3.1 Informed principal

Situations with an informed principal are those in which it is the type of the principal that matters, and this type is not known to either the principal or the agent before contracting (we will see later about the case where the principal knows the type). The situation can be modeled as follows:

- $V = S(q, \theta) - t$, is the principal's payoff.
- $U = u(t - \theta q)$, is the payoff of the agent. As can be seen, it depends on the type of the principal.

The contracting timeline can be represented as follows:

$t = 1$	$t = 2$	$t = 3$	$t = 4$
Principal offers contracts C	Agent chooses contract	Principal learns type θ	q is produced t is paid

Note that to the difference of the standard incentive problem, the principal must now have an incentive to reveal his own type truthfully. Once the principal reveals his own type, the agent produces q and receive t as agreed in the contract for when the principal is of that specific type. The situation is reversed compared to the standard, informed agent setting: if the principal is of a bad type, then the first best level of production will be asked, while if the principal is of a good type, then there will be more production than first best optimal (chapter 9, subsection 1, Laffont and Martimort, 2002). Note however that if the agent is risk neutral, then the first best levels of production can be achieved, while in the standard case, even if the agent was risk neutral, there was still a distortion.

In the above, we focused on a case where, at the time of offering the contract, the principal did not know his own type. Myerson (1983)[23] examines the more complicated case where the principal knows his own type before offering the contract. This type of situation can happen for example in franchise contracts. A potential MacDonald franchisee may know less about her business than the franchiser (MacDonald), who indeed would know more about local market conditions or the current success of its offering. This type of situation may also apply to a job applicant when negotiating an employment contract with a firm. The agent will then wants to learn the type of the firm (close to bankruptcy or not, for example) from the type of contract the firm offers. Complex issues of signaling, as in the 'lemons problem' for example, arise, as the principal's choice of contract may reveal his own type to the agent. There are two main cases from the point of view of the theory: ones when the type of the principal only enters the principal's utility function,[24] and ones when the type of the principal also enters the agent's utility function.[25] That later case is the one we looked at above.

[23] Myerson R.B., 1983, Mechanism Design by an Informed Principal, Econometrica 51(6), 1767-1797.

[24] Maskin E. and J. Tirole, 1990, The Principal-Agent Relationship with an Informed Principal: The Case of Private Values, Econometrica 58(2), 379-409.

[25] Maskin E. and J. Tirole, 1992, The Principal-Agent Relationship with an Informed Principal, II: Common Values, Econometrica 60(1), 1-42.

20.3.2 Mixed models

Seldom does one have pure adverse selection or pure moral hazard. In most situations, there is either adverse selection followed by moral hazard or moral hazard followed by adverse selection (chapter 7, Laffont and Martimort, 2002). Insurance is a typical example: the effort by the agent in avoiding accidents may depends on its type, her tendency for recklessness for example. In this case, adverse selection (type) precedes moral hazard (effort).

The adverse selection problem will be exarcebated by moral hazard (compounding factor), as the bad, reckless type will do less effort to avoid accidents, thus increasing the discrepancy in terms of probability of a bad outcome between good and bad type. In the opposite case, where moral hazard is *followed by* adverse selection (for example, when effort to get an education also changes one's type in a way that is favorable), then the moral hazard problem may be weakened by adverse selection, as agents may exert more effort in order to acquire a good type, and not only just to increase the probability of a good outcome.

20.3.3 Limits to contractual complexity

Contracts are generally simpler than what theory would call for. Laffont and Matoussi (1995)[26] examine real world sharecropping contracts and assess the loss of efficiency due to their relative simplicity. Are simple contracts signed because there is too much of a cost to write complicated contracts? What then does that cost consist in? Tirole (2009)[27] is a recent contribution exploring this topic.

20.3.4 The information structure

It was assumed that the agent knew its own type (adverse selection), or that it was able to choose its own level of effort (moral hazard). What happens now if the participants can choose what information they have access to and at what stage? This calls for a design of the contracts informational structure as well as of the contracts. It then becomes necessary to motivate information acquisition by the agent and/or the principal.

21 The revelation principle.

In adverse selection problems we assumed up to now that the principal would offer one contract (or less) for each different type of agent. For example, with two types, the principal would offer either:

- Two different contract that would be such that each type would choose a different one (the contracts are 'separating' both types), or it would offer only

[26]Laffont J-J. and M.S. Matoussi, 1995, Moral Hazard, Financial Constraints and Sharecropping in El Oulja, Review of Economic Studies 62, 381-399.

[27]Tirole J., 2009, Cognition and Incomplete Contracts, American Economic Review 99(1), 265-294.

- One contract that both types would choose (the contract is 'pooling' both types).
- One contract such that only one type would accept a contract (the contract is excluding one type). The other type chose the 'no contract' option $C = (0, 0)$ that gave it its reservation utility.

In the case of separating contracts, as each type chose a different contract, the principal gained information on the agents' types from their choice of contract.

This section re-examines the very simple setting introduced in the lecture on adverse selection to examine whether it could ever be beneficial for the principal to offer a third or more contracts, offer more options than simply quantity and transfer (q, t) contracts or asking the agent for more information than simply their choice of contract.

Stating the above in another way, we restricted ourselves up to now to *Truthful Direct Revelation Mechanisms ('DRM')* whereby:

1. The principal offered a menu of contract $C(\theta) = (q(\theta), t(\theta))$, depending on θ, the type announced by the agent. This is the *mechanism.*
2. Agents were only required to *directly reveal* a type, so agent of type θ who would announce θ' would produces $q(\theta')$ and get $t(\theta')$. Her utility was then $U(\theta, \theta') = t(\theta') - \theta q(\theta')$.
3. The principal designed her contracts such that agent of type θ would announce her type was θ (*truthful*) and got contract $C(\theta)$. This was guaranteed by the incentive constraint:

$$U(\theta, \theta) = t(\theta) - \theta q(\theta) \geq U(\theta, \theta') = t(\theta') - \theta q(\theta') \tag{80}$$

   for any $\theta' \neq \theta$

   In the case where there was for example 'bunching' or pooling of type θ and θ' , there was no loss for the agent in announcing θ, rather than for example θ', as $C(\theta) = C(\theta')$. This means there was at least weak incentive to reveal one's type.

This part will prove that the above type of solution, in terms of *truthful DRM* is general, that is, it is never beneficial to offer more contracts than there are types (principle of economy), there is no point in asking the agent for more than to just announce a type (or for more than just choosing a contract), and there is no loss in generality in looking only at DRM where agents announce their type truthfully.

Obviously, if the only communication between the agent and the principal is to tell one's type (or another type), then there can be only so many contracts as there are types.[28] In order to test the principle of economy, let's thus consider

[28] Unless, to the announcement of a type, the principal offers a contract chosen at random between a number of contracts. But then, among those contracts, there must be one that gives higher utility to the principal than another, given the announced type, so the principal would be better off not choosing randomly.

a set of messages M that can be transmitted by the agent to the principal. Those messages can be the choice of contract, the type of the agent, or any other information. The mechanism is then of the form $C: M \to A$, with A the set of allocations $(q,\, t)$. This means that to any message m in M, the function C associates an allocation (contract) *s.t.* $C(m) = (q(m),\, t(m))$. Agent θ will then rationally choose to communicate the message $m^*(\theta)$ that maximizes her surplus, that is

$$m^*(\theta) = \arg\max_{m'}\{t(m') - \theta q(m')\}. \tag{81}$$

The *revelation principle*, which we introduce below, tells us that there is no need for such complication in the mechanism, so (1) each agent can limit her communication to announcing her type (rather than some more general message m), and (2) there is also no need to consider mechanisms whereby an agent would announce a type different from her own.

Denoting Θ the set of types, (1) is shown by saying that instead of having a mechanism such that one type send her optimal message and the message maps into an allocation A:

$$\Theta \overset{m^*(\theta)}{\to} M \overset{C(m)}{\to} A \tag{82}$$

one can simply consider $C'(\theta) = C \circ m^*(\theta)$ whereby the principal replaces C by the equivalent menu of contract C' that maps types directly into allocations:

$$\Theta \overset{C'(\theta)=C\circ m^*(\theta)}{\to} A \tag{83}$$

(2) is shown by saying we can limit ourselves to *truthful direct revelation mechanism*, whereby agent announce their own type. Indeed, replacing m' by $m^*(\theta')$ in equation (81) one obtains:

$$m^*(\theta) = \arg\max_{m'}\{t(m') - \theta q(m')\} \tag{84}$$

$$= \arg\max_{\theta'}\{t(m^*(\theta')) - \theta q(m^*(\theta'))] \tag{85}$$

and since $t(m^*(\theta)) = t'(\theta)$ under the menu of contract C' introduced above, this means that in the same way as the agent's message truthfully revealed the type of the agent under contract C, the type announced under contract C' is truthful.

Intuitively, suppose the contracts offering were such that an agent of type θ is better off announcing she is of type $\theta' \neq \theta$ (all other agents announce truthfully). Then this contract offering is equivalent to another contract offering that gets rid of the contract that is assigned to an agent that announces she is of type θ. Under that new contract offering, agents θ and θ' are pooled into the same contract. This new contract offering is equivalent to a contract offering such that whether you announce your type is θ or θ', you are assigned the same contract. And in this case, agent of type θ may as well announce she is of type θ. Therefore, by extension, any contract offering that is not a truthful DRM has an equivalent that is a truthful DRM.

The following introduces an example (auctions) where the revelation principle can be used to great effect to simplify and solve apparently intractable problems. We also consider limitations to the truthful DRM in voting problems, where there are many agents and agents know not only their own type, but also know others' types.

21.1 Application: Auctions

We saw in the game theory part of this course that offering agents to play second price auctions resulted in a DRM where each agent announced her type (i.e. her valuation for the good). If she won (her valuation was the highest), she paid the second highest announced valuation. First price auction however did not result in a DRM: agents, rather than announcing their own valuation, would announce half their own valuation (indeed, each agent would optimally bid half her valuation for the good). However, first price auctions are implementable through the following DRM: agents announce their own type (valuation for the good), and pay half of that valuation if they win. More generally, the revelation principle is useful in the context of auctions as it shows there is no point in studying such and such specific mechanism (first price, second price, loser pays, etc.), as one can limit oneself to simple, direct revelation mechanism, knowing that they are the equivalent of the more complicated designs.

21.2 Voting mechanisms and limits to truthful DRMs

What we said up to now was that a principal who wished to make a decision when faced with informed agents could limit himself to considering only mechanisms (ways to reach a decision) that are truthful DRMs. However, this does not mean that such mechanisms will have any desirable properties, or even that they will exist.

In order to illustrate this point, consider the following collective decision problem. There are three agents, a, b and c who have to collectively choose between option, x, y, z or w. Suppose agents rank options according to the following table:

		Preferences
Agents	a	$x \succ y \succ z \succ w$
	b	$x \succ y \succ z \succ w$
	c	$z \succ y \succ x \succ w$

In this example, agent b ranks option x as best, y as second best, z as second worst and w as worst.

1) Suppose a principal *knows the agents' preferences* and wishes to design a mechanism such that the chosen option will maximize collective welfare. In this case, this can be interpreted as minimizing the sum of the rank of the option chosen across all agents. This is the *Borda rule*, named for Jean-Charles de Borda, who devised the system in 1770 for elections to the French Academy of

Sciences.[29] Choosing x sums up to 5. y sums up to *6*. z sums up to *7* and w sums up to *12*. The optimal decision would thus be x.

2) Suppose now however that the principal *does not know the ranking of each option* by each agent. Suppose however that the agents know their own rankings, and the rankings of others. One could then offer agents to choose according to the Borda rule. Each agent should announce (secretly, that is, only to the principal) a ranking of alternatives. The principal should choose the option that minimizes the sum of the rank of the alternatives across all agents.

Suppose thus that under this mechanism, all agents announce their rankings of options truthfully. Then as seen above, option x would be chosen. However, agent c could change to announcing $y \succ z \succ w \succ x$ instead of $z \succ y \succ x \succ w$, thus lowering the overall rank of x and increasing the rank of y. The announcements would then be as follows:

		Preferences
Agents	a	$x \succ y \succ z \succ w$
	b	$x \succ y \succ z \succ w$
	c	$y \succ z \succ w \succ x$

y now scores *5*, while x scores 6, so that by this manipulation, c gets its second preferred choice adopted, rather than its third choice.[30] Anticipating this, a or b could change their choice, leading to a cycle of changes away from the truth.

It is possible to generalize from this, and to say that one cannot implement a truthful DRM in such a voting situation, *except for the dictatorial rule* (one agent is given the role of a dictator and chooses for others. He then chooses his most preferred option. This is a DRM). Obviously, this dictatorial rule is very inefficient, since it fits the preference of only one person.

This is the subject of the Gibbard-Satterthwaite theorem (Chapter *23.C* of the Mas-Colell),[31] which is similar (if not identical)[32] to the very interesting Arrow's impossibility theorem (Chapter *21.C* of the Mas-Colell).[33]

This means that there is no hope for a rule that would lead the principal to choose the best option for all possible rankings of the options by all agents. This is because any rule could be subject to manipulation under some circumstances; some agents would prefer to lie about their type (ranking of alternatives).

[29] de Borda J-C., 1781, Mémoires sur les élections au scrutin, Histoire de l'Académie Royale des Sciences, Paris.

[30] To the credit of Borda, he knew his scheme was subject to manipulation, saying 'My scheme is intended only for honest men'.

[31] Gibbard A., 1973, Manipulation of Voting Schemes: A General Result, Econometrica 41, 587-601.
Satterthwaite M.A., 1975, Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions, Journal of Economic Theory 10, 187-217.

[32] Reny P.J., 2000, Arrow's Theorem and the Gibbard-Satterthwaite Theorem: A Unified Approach, University of Chicago Working Paper.

[33] Arrow K.J., 1950, A Difficulty in the Concept of Social Welfare, Journal of Political Economy 58(4), 328-346.

While this is true if agents know each other's type (ranking), there are truthful DRM in Bayesian strategies if agents know their own type (ranking), but do not know the type (ranking) of the other agents (Chapter *23.D* of the Mas-Colell). In that case there are mechanisms such that the agents announce their type truthfully given their *expectation* (the average) of what the other agents' types will be. Note that to the different of truthful DRM that we looked at up to now, where the agent would play the same whathever his belief about other agents' types (truthful DRMs in dominant strategies), truthful DRM in Bayesian strategies lead the agent to play differently depending on his belief about other agents' type. Truthful DRM in Bayesian strategies are thus less robust than truthful DRMs in dominant strategies, as their outcome will depend on agents' beliefs, which may be arbitrary and out of step with reality.

Part IV

Exercises

22 Choice under uncertainty[34]

1. The figure shows indifference curves for a particular individual over lotteries involving three possible outcomes, x_1, x_2, and x_3, such that x_3 is strictly preferred to x_2 which in turn is strictly preferred to x_1.

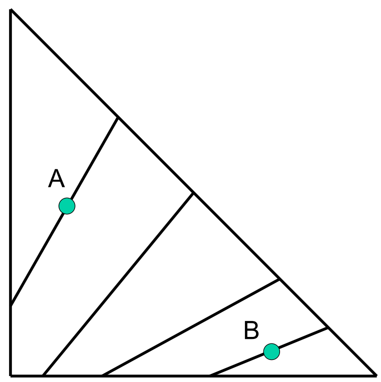

Figure 17: Two lotteries and indifference curves in the Marschak-Machina triangle

Justifying your answer, explain:

a) how their attitudes towards risk differs between point A and point B in the figure and

b) which of the axioms of expected utility theory are broken by these indifference curves.

2. A six-sided die is to be rolled. Consider the following four lotteries:

- Lottery A pays £10 on all numbers.
- Lottery B pays £24 on number 1-3, and nothing on 4-6.
- Lottery C pays £10 on numbers 1-4, and nothing on 5-6.

[34] Exercises 1, 2 and 3 were contributed by previous teachers in MSc Economic Theory 1 at the UEA.

- Lottery D pays £24 on numbers 5-6, and nothing on 1-4.

Alexia prefers A to B. She also prefers D to C.

a) Are these preferences consistent with expected utility theory?

b) Are they consistent with prospect theory?

c) Are they consistent with regret theory?

In each case justify your answer.

3. Consider the following six lotteries:

- Lottery A pays £10 with probability 1.
- Lottery B pays £30 with probability 0.4, otherwise nothing.
- Lottery C pays £10 with probability 0.25 otherwise nothing.
- Lottery D pays £30 with probability 0.1, otherwise nothing.
- Lottery E pays £12 with probability 0.5, otherwise nothing.
- Lottery F pays £30 with probability 0.2, £12 with probability 0.6, otherwise nothing.

Adelina's preferences over lotteries satisfy the axioms of expected utility theory. She prefers lottery C to D. What are her preferences over the following pairs of lotteries?

a) C versus E

b) D versus E

c) E versus F

d) A versus B

In each case provide careful justification for your answer.

4. Consider an individual with utility function $U(x) = ln(x)$, where x is income. Suppose the individual rejects a 50-50 chance of losing £10 and winning £11 at an initial wealth level w.

a) Show that this implies that $ln(w - 10) + ln(w + 11) \leq 2ln(w)$.

b) Deduce from a) that $w \leq 110$ (Reminder: $ln(ab) = ln(a) + ln(b)$ and $cln(a) = ln(a^c)$)

c) Suppose the individual's wealth is $w = 110$ and the individual is offered a 50-50 chance of losing £100 and winning £Y. What is the minimum Y such that the individual accepts?

d) Suppose now the individual's wealth is $w = 200$ and the individual is offered a 50-50 chance of losing £100 and winning £Z. What is the minimum Z such that the individual accepts?

e) Comment on the above results. Do you find the predictions from c) and b) reasonable? Justify your answer.

5. Consider a lottery that gives outcome G with probability q and outcome L with probability $1 - q$. Suppose the initial wealth of the agent is w and the agent is an expected utility maximizer.

a) Suppose the agent offers others to play the lottery and prices the ticket for the lottery at p_S. What is the minimum price that the agent will set? Write down the condition on p_S in terms of the agent's utility function $u(.)$.

b) Suppose now the agent wishes to buy a ticket for the lottery. What is the maximum price p_B that the agent would be ready to pay for that ticket?

c) Suppose that $G = 10$, $L = 2$, $w = 10$ and $q = 0.5$. Suppose also that $u(x) = ln(x)$. Compute the values of p_S and of p_B.

d) Do you find that $p_S = p_B$? Why or why not?

6. [From Loomes and Taylor (1992)][35]

Consider a world with three possible states, S_1, S_2 and S_3, that occur with probabilities p_1, p_2 and p_3 respectively. An agent has the choice between actions A, B and C, such that action A obtains outcome a if S_1 occurs, and d otherwise; action B obtains outcome e if S_3 occurs, and b otherwise; and action C obtains outcome c in any state of the world. Outcomes are expressed in monetary terms and $a > b > c > d > e$.

Suppose the agent's behaviour fit the axioms of Expected Utility Theory.

a) Represent A, B and C as lotteries over outcomes a, b, c, d and e. What are the expected utilities of A, B and C?

b) Is it possible for an agent to prefer A to B, B to C, and C to A? Explain your answer.

[35] Loomes G. and C. Taylor, 1992, Non-transitive preferences over gains and losses, The Economic Journal 102(411), 357-365.

Suppose now the agent exhibits regret aversion. She associates regret $r(x, y)$ when she would have obtained outcome x in the lottery she chose while she would have obtained outcome y in the lottery she did not choose. When comparing two lotteries, L_1 and L_2, she chooses L_1 if the expected regret of choosing L_1 is positive. Assume that: $r(x, y)$ is increasing in x; $r(x, y) = -r(y, x)$; and $r(x, y) > r(x, z) + r(z, y)$ when $x > z > y$.

c) Interpret the above assumptions concerning the regret function.

d) Write down the conditions under which the agent chooses: B over A; C over B; A over C. Is this consistent with regret theory? Comment. (Hint: sum up the three conditions you identified and simplify using the properties of the regret function).

e) Write down the conditions under which the agent chooses: A over B; B over C; C over A. Is this consistent with regret theory? Conclude.

23 Game Theory 1 [36]

1. In the following game, find all pure strategy Nash equilibria:

	L	R
L	6,6	0,8
R	0,3	2,2

2. Consider the following game:

	W	X	Y	Z
A	4,6	0,4	8,4	9,2
B	4,2	2,6	6,4	8,4
C	4,2	2,8	4,9	6,0
D	0,5	0,3	1,7	2,6

a) Apply the procedure of iteratively deleting strictly dominated strategies.

b) Apply iterative deletion of weakly dominated strategies to the game that remains after completion of (i).

c) Identify all pure and mixed strategy Nash equilibria of the game.

3. Consider the following game in extensive form

[36] Exercises 1 to 5 were contributed by previous teachers in MSc Economic Theory 1 at the UEA.

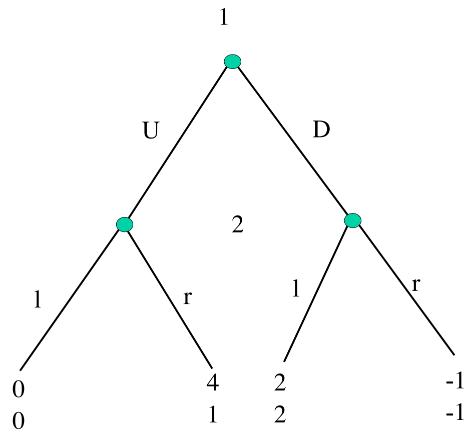

Figure 18: Extensive form game for exercise 3, Game Theory 1

a) Convert the game into a game in normal form.

b) Find all Nash and subgame perfect Nash equilibria.

4. Consider the following game in extensive form:

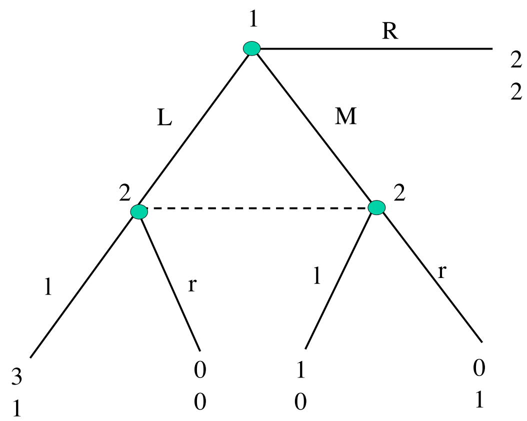

Figure 19: Extensive form game for exercise 4, Game Theory 1

a) Convert the game into a game in normal form.

b) Find all Nash and subgame perfect Nash equilibria.

c) Comment on the plausibility of these.

5. Consider the following game in extensive form:

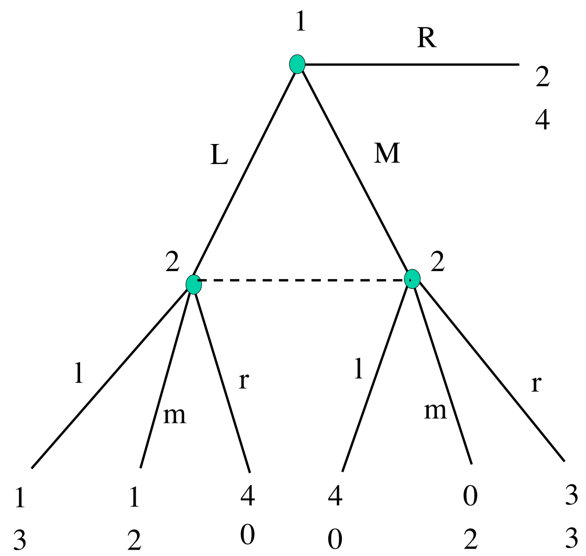

Figure 20: Extensive form game for exercise 5, Game Theory 1

a) Convert the game into a game in normal form.

b) Find all Nash and subgame perfect Nash equilibria.

c) Comment on the plausibility of these.

6. Consider the following game:

	Stop	Left	Middle	Right
Up	10,2	-10,1	0,0	-10,-10
Down	-10,-10	10,-5	1,0	10,2

a) Suppose that players are able to agree on their move before playing. What would they choose?

b) From now on, suppose players 1 and 2 are not able to communicate and must choose their moves independently. Suppose you are player 2. Intuitively, what seems like a reasonable way to play when you are not sure what player 1 will choose? Explain.

c) What are the pure strategy Nash equilibria of this game?

d) Consider mixed strategy equilibria such that player 1 mixes between Up and Down with probability q and $1 - q$ respectively.

i) Show that player 2 will never play Left for any value of q.

ii) Show that if player 2 plays mixed strategies and plays Stop with some positive probability, then it plays Right with zero probability.

iii) Show that player 2 will play Middle whenever q belongs to the interval [1/6, 5/6].

e) Rationalize your choice of action in (a) taking into account (d)(iii).

24 Game Theory 2

1. Consider the following game in extensive form:

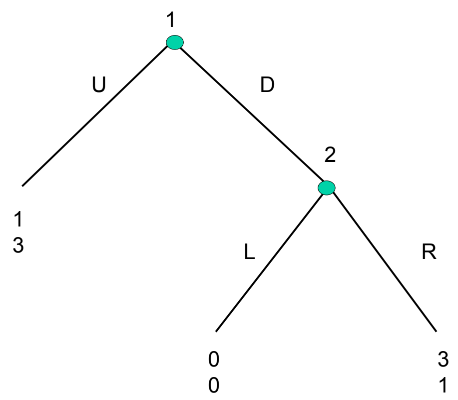

Figure 21: Extensive form game for exercise 1, Game Theory 2

a) Write the game in normal form and find all of its Nash equilibria.

b) Find the Subgame Perfect Nash Equilibrium of this game. Explain the intuition for this equilibrium.

c) Suppose player two can pre-commit to spend $c > 0$ in case it has to play R. The decision of player two whether to pre-commit or not is know to player 1.

i) Draw the new game in extensive form, where in a first stage player 2 decides to pre-commit or not.

ii) Find the Subgame Perfect Nash Equilibrium of this new game for every possible value of $c > 0$.

iii) When does player 2 pre-commit? Explain.

2. Consider a buyer and a seller. The seller has only one unit to sell and values the good at £50 while the buyer values it at £60 if it is of a high type and £55 if it is a low type. The seller does not know the type of the buyer, but knows there is proportion h of high type. The buyer knows its own type.

a) Suppose the seller asks for a price p for the good and p is a take it or leave it offer. The buyer can only either accept or reject the price offer. Find out the seller expected payoff as a function of p. What price will the seller set?

Suppose now the game is played over two periods, and the buyer's discount factor is $t < 1$ (the buyer is impatient) while the seller's discount factor is 1 (the seller is patient). Suppose also the seller can change prices between periods: p_1 is the price asked in period 1 while p_2 is the price asked in period 2. Both periods involve the same buyer and the same seller, and prices p_1 and p_2 are announced in the first period. The seller can commit to prices p_1 and p_2.

b) Draw the extensive form of the game, with N a play by nature determining the type of the buyer.

c) Suppose $60 > p_1 > 55 > p_2 > 50$. What is the expected payoff for a high type buyer from accepting in the first period? From waiting? What about the low type?

d) Show there exist p_1 and p_2 such that high type buyers buy in the first period and low type buyers buy in the second. What conditions must p_1 and p_2 satisfy? What is the expected payoff of the seller under that type of equilibrium?

e) Suppose the seller cannot commit on p_2 in period 1. Is the separating equilibrium of question d) still sustainable? Why or why not?

3. Consider two players, A and B, who are engaged in a public goods game whereby the public good is produced and has value 1 to both if either or both of the players contribute to it, while it has value 0 if neither contributes. This means the good is produced and its value is available to all even if only one player contributes to it. The costs of contributing are c_a and c_b for players A and B respectively. The term c_a can take one of the two values: $C > 1$ with probability $1 - \lambda$ and $c < 1$ with probability λ; c_b independently follows the same distribution. We call a player whose cost of contributing is C the inefficient type, while if that cost is c we call that agent the efficient type. Each player knows his/her type, but not the other player's type. Players choose at the same time whether to contribute and do not observe whether the other player contributes.

a) Explain why a player of type C is called inefficient and a player of type c is called efficient.

b) Explain why it is important that contributions be made at the same time and players cannot observe whether the other player contributes.

c) Suppose $\lambda = 0$ which means both players are efficient. Draw the game in normal form. What are the Nash equilibria of the game? What if $\lambda = 1$?

d) Suppose now $1 > \lambda > 0$. Draw the game in extensive form (Hint: Represent the draw of players' type as a move by nature).

e) Depending on the type of the agents, what is the Pareto optimal outcome of this game?

f) Show that in a Bayesian equilibrium of this game, inefficient player A never contributes, while efficient player A contributes with some probability p_a. Express p_a and p_b as a function of the parameters of the game. Be careful to explain the notations you use and the reasoning you are following.

g) Suppose now the game is repeated twice, and players observe each other's action after the first period and before playing the game again. Denote $p_a^t(c_a)$ the probability that player A of type c_a contributes in period t. Show again that a player of a low type never contributes. How does player B interpret a contribution by player B in period 1?

25 Incentive theory[37]

1. Consider an agency relationship in which an agent can take two efforts, either low or high. The agent's effort might result in two results, 10 or 100 depending on the state of nature.

[37] Exercises 1, 3 and 4 were contributed by previous teachers in MSc Economic Theory 1 at the UEA.

We will consider three different agency relationships. Each relationship is summarised by one of the following matrices. In these matrices, each entry is defined as the probability that the result is the COLUMN result given that the effort is the ROW effort. So, for instance, in CASE A, 0.3 is the probability of the result being 100, given effort is LOW.

CASE A

effort/result	10	100
LOW	0.7	0.3
HIGH	0.1	0.9

CASE B

effort/result	10	100
LOW	0.1	0.9
HIGH	0.9	0.1

CASE C

effort/result	10	100
LOW	0.6	0.4
HIGH	0	1

Each matrix represents a different situation but in only one of them is there a real moral hazard problem. Discard two of these matrices and carefully argue why there is no moral hazard problem in them.

2. [Simplified from Myerson (2008)][38]

Consider a one-time production project that requires an initial capital input worth $K = 100$, and then returns revenue worth $R = 240$ if the project is a success, or returns no revenue (0) if the project is not a success.

The project's probability of success depends on the manager's hidden action. If the manager diligently applies good effort to managing the project, then probability of success is $p_G = 1/2$. On the other hand, if the manager behaves badly and abuses her managerial authority in the project, then the probability of success is $p_B = 0$, but the manager gets hidden private benefits worth $B = 30$ from such abuse of power.

The manager owns no personal assets and does not invest anything in the project.

a) Show that the expected returns from the project (including possible private benefits) are greater than the cost of its capital inputs only if the manager chooses to be good.

[38] Myerson R.B., 2008, Perspectives on Mechanism Design in Economic Theory, American Economic Review 98(3), 586-603.

b) Suppose the manager is paid wage $w > 0$ in case of success, and nothing in case of failure. Under what condition on w does the manager undertake to be good?

c) Suppose the investor chooses w such that the manager undertakes to be good. What is then the highest possible return of the project to the investor? Does the investor invest? Comment.

3. Consider a monopolist who has N customers, of which a proportion p is of type 1, and the remaining proportion $(1 - p)$ is of type 2. The payoff for a consumer of type $i = 1, 2$ from consuming quantity q of the good is $U_i = q[t_i - q/2] - R(q)$, with $i = 1, 2$, $t_2 > t_1$ and $R(q)$ denotes the payment made to buy q. The consumers' reservation payoff (or consumer surplus) is zero.

The monopolist wishes to discriminate between customers; would he know the customers 'types' he would extract the full consumer surplus from them. However, he is unable to do so, and manages to price discriminate only up to the second degree. This means he will offer two contracts destined for each type, with contract $i = 1, 2$ to be chosen by type i denoted as $C_i = (q_i, R(q_i))$. The monopolist has a constant marginal cost c and no fixed cost so the firm's type-specific profit is $S_i = R(q_i) - cq_i$.

Your task is to determine the monopolist's optimal quantity-payment schedule that accomplishes second-degree price discrimination. Follow the questions in sequence.

a) Derive the optimal quantity-payment schedule when the monopolist can distinguish between the types of the consumers. Show your answer graphically (on a quantity- payment plane) and explain your findings.

b) Introduce asymmetric information and show how the first best scheme becomes sub-optimal using a graph.

c) Set the monopolist's profit maximization problem formally with appropriate self-selection and participation constraints. Identify which constraints will bind, and which will not.

d) What is the meaning of "no distortion at the top" and "no rent at the bottom" in this context?

e) Finally, solve for the optimal quantity-payment schedule and show it on the graph.

4. Imagine that a hospital is contracting out its cleaning service. The contracted firm is to supply a level I of cleanliness, which is observable. There are two types of firms, which differ in their costs of producing cleanliness. For the first type, the cost of providing I level of cleanliness is I^2, while for the second one it is kI^2, with $k > 1$.

The payment that is made to the chosen contractor is P, which may be made dependent on I. The contractor's profit is thus $S = P - I^2$, or $S = P - kI^2$ depending on its type. Assume that its reservation profit is 10. The hospital's utility is $U = AI - P$, where A is a positive parameter.

a) Which firm is the "top" (efficient) type, which one is the "bottom" (inefficient) type?

b) Assume that the hospital can observe the type of the firm. Write the maximisation problem for the hospital if it contracts with the top type. Find the first order conditions and the relationship between P and I at the solution.

c) Assume that the hospital cannot observe the type of the firm. Write the maximisation problem for the hospital if it wants the two types of firms themselves to select different types of contract. Assume that 1/2 of catering firms are of each type.

d) Identify the constraints that will bind in the solution, and solve for the optimal (P, I) schedule.

e) Present your answer to d) graphically.

5. Consider an agent with utility function $U = u(q) - c(e)$ where $u(.)$ is defined over output and $c(e)$ is the cost of effort e. Effort e can be either 0 or 1 and determines the probability of success in producing output $Q > 0$. The cost of effort is $C > 0$ if effort is 1 and 0 if effort is 0. p_1 is the probability with which the agent exerting effort $e = 1$ is successful and produces output $Q > 0$, while $p_0 < p_1$ is the probability with which the agent exerting effort $e = 0$ is successful and produces output $Q > 0$. If the agent is not successful, then he produces no output.

Consider a paternalistic employer that seeks to maximize the welfare U of the agent, and can pay wages such that the agent that is successful receives wage W and the agent that is not successful receives w. Note that wages can be negative in which case the agent pays the employer. Note also that the employer cannot make losses in expected terms.

a) Consider the first best scenario in which effort is verifiable.

i) Write down the agent's participation constraint (the condition under which expected payoffs for the agent is positive, taking into account wages from the employer) and the agent's incentive compatibility constraint (the condition under which the agent exerts effort).

ii) Write down the employer's objective function and maximization program.

iii) Write down the employer's necessary and sufficient Kuhn and Tucker optimality conditions with respect to W and w. How much does the employer pay in case of success and in case of failure? Comment.

b) Consider now the second best scenario in which effort is non-verifiable.

i) Write down the employer's objective function and maximization program.

ii) Write down the employer's first-order condition for optimality with respect to W and w. Show that $W \neq w$. Comment.

6. A firm that faces productivity shocks has the following profit function: $\pi = xh - w$ with $x = x_H$ with probability q (favourable shock) and $x = x_L$ with probability $1 - q$ (adverse shock), $x_H > x_L > 0$, h the number of hours worked by the employee of the firm and w the wage paid to the worker.

The worker's utility is $u(w) - f(h)$, with $u(.)$ concave, $f(.)$ convex, w the wage paid by the firm to the employee and h the number of hours worked by the employee. Suppose the firm chooses the terms of the worker's employment contract, and the worker's reservation utility is 0. The contract is signed before realization of the shock; the productivity shock is observed after contracting. Contract variable can be made dependent on the realization of the shock as its realization is publicly observable and the contract can thus be enforced by the Courts.

a) Who is the principal? Who is the agent? What variables are the firm and the worker going to contract upon?

b) Draw the timeline of contracting.

c) Write the principal's objective function. Write the agent's individual rationality (IR) constraint.

d) Find the first order condition for the maximization of the principal's objective function subject to the agent's IR constraint. (Reminder on constrained maximization: The solution(s) to a constrained maximization program, with $f(x, y)$ the objective function and $u(x, y) > 0$ the constraint, is obtained by taking the derivative of that program's Lagrangian $f(x, y) - vu(x, y)$, *w.r.t.* x, y and v, and equating it to *0*, with v the Lagrange multiplier).

e) Prove that the agent will receive the same wage regardless of the realization of the shock. When does the agent work more?

f) Suppose the shock is not publicly observable, is known only to the firm and the firm can lie about the realization of the shock. Suppose it offers the contract determined in *e)* above. Will the principal lie about the realization of the shock? When? Why?

g) Write down the conditions under which the principal has no incentive to lie. Using only those conditions, compare the wage and the hours worked in case of an adverse and in case of a favourable shock.

Part V

Correction of exercises

26 Choice under uncertainty

1.

a) The steeper the indifference line, the more risk averse the agent is. See page 126, note 6 of Machina (1987).[39] Intuitively, suppose for example that the vertices of the triangle are *£0*, *£50* and *£100*. If the agent is risk neutral, then she values a lottery *(1/2, 0, 1/2)* the same as *(0, 1, 0)* and the indifference line linking those two points thus represents the preferences of a risk neutral agent. A risk averse agent will prefer the second lottery to the first, so her indifference line will be steeper. The opposite holds for an agent that is risk loving.

b) The independence axiom requires that indifference lines be parallel, but here they are not parallel, so the independence axiom is not respected. Other axioms are not broken, as there is no crossing of indifference curves within the triangle, indifference curves are continuous and they span the whole area of the triangle.

2.

a) According to EUT, A preferred to B translates in $u(10) \geq \frac{1}{2}u(24) + \frac{1}{2}u(0)$. Multiplying by $\frac{2}{3}$ both sides and adding $\frac{1}{3}u(0)$ on both sides, one finds $\frac{2}{3}u(10) + \frac{1}{3}u(0) \geq \frac{1}{3}u(24) + \frac{2}{3}u(0)$, which translates into C preferred to D.

Therefore, the combination A preferred to B and D preferred to C is not consistent with EUT.

b) For A preferred to B and D preferred to C to be consistent with prospect theory we must have both:

$$\pi(1)u(10) \geq \pi(\frac{1}{2})u(24) + \pi(\frac{1}{2})u(0) \tag{86}$$

$$\text{and} \tag{87}$$

$$\pi(\frac{2}{3})u(10) + \pi(\frac{1}{3})u(0) \leq \pi(\frac{1}{3})u(24) + \pi(\frac{2}{3})u(0) \tag{88}$$

Normalizing $u(0) = 0$ (this is allowed as the utility function is unique up to a linear transformation) this can be rewritten into the following condition:

$$\frac{\pi(1)}{\pi(\frac{1}{2})} \geq \frac{u(24)}{u(10)} \geq \frac{\pi(\frac{2}{3})}{\pi(\frac{1}{3})} \tag{89}$$

I will show that this can occur under some specification of the function $\pi(.)$. Suppose $\pi(.)$ is a continuously increasing function,

[39] Machina M., 1987, Choice under Uncertainty: problems solved and unsolved, Journal of Economic Perspective 1(1), 121-154.

$\pi(\frac{1}{3}) = \frac{1}{3}$, $\pi(x) > x$ for any $x < \frac{1}{3}$ and $\pi(x) < x$ for any $x > \frac{1}{3}$ (this fits common assumption about the shape of the $\pi(.)$ function and corresponds to some empirical findings), and suppose $u(x) = x$ for any x (risk neutral agent). Note also that $\pi(1) = 1$. The condition above can then be rewritten as

$$\frac{1}{\pi(\frac{1}{2})} \geq \frac{24}{10} \geq 3\pi(\frac{2}{3}) \tag{90}$$

so I must have and This is consistent with prospect theory, as $\pi(\frac{2}{3}) \leq \frac{2.4}{3}$ from the fact $\pi(x) < x$ for any $x > \frac{1}{3}$, while $\pi(\frac{1}{2}) \leq \frac{10}{24}$, while more stringent than what we assumed, does not contradict PT.

c) Under regret theory, A preferred to B translates in

$$\frac{1}{2}r(10,0) + \frac{1}{2}r(10,24) \geq 0 \tag{91}$$

while D preferred to C translates in

$$\frac{1}{3}r(24,0) + \frac{2}{3}r(0,10) \geq 0 \tag{92}$$

In regret theory, one will generally expect $r(x,y) > r(x,z) + r(z,y)$ with $z \in]y,x[$ (rejoice increases more than proportionally in the size of the difference between bad and good outcome). This means that $r(24,0) > r(24,10) + r(10,0)$ (condition A).

One will also expect that $r(x,y) = -r(y,x)$ (symmetry between rejoice and regret) and allows one to rewrite the two equations as

$$r(10,0) \geq r(24,10)$$

$$r(24,0) \geq 2r(10,0)$$

I can accept that the first equation is verified (the feeling of rejoicing when getting 24 rather than 10 may be less than from getting 10 rather than nothing). If that first equation is verified, the second equation is more stringent than condition A. However, while more stringent than usual assumptions of RT, the second equation does not contradict it in any way. Therefore, such preferences are consistent with regret theory.

3. We will use the same logic as the one used in 2.b), by making appropriate substitutions and using dominance arguments.

a) E is preferred to C because it gives more weight to a higher outcome (First Order Stochastic Dominance)

b) Since E is preferred to C which is preferred to D, then E is preferred to D (Transitivity)

c) F preferred to E, again by First Order Stochastic Dominance.

d) C is preferred to D so A is preferred to B (EUT, as in 2.b).

4.

a) The utility of rejecting the lottery is $\ln(w)$ (the individual is left with the same wealth as before), while if the individual plays the lottery, then she either loses £10, in which case wealth is reduced to $w-10$, or she gains £10, in which case her wealth is $w+11$. The expected utility of playing the lottery is then $\frac{1}{2}\ln(w-10)+\frac{1}{2}\ln(w+11)$, and rejecting the lottery translates by EUT in $\frac{1}{2}\ln(w-10)+\frac{1}{2}\ln(w+11) \leq \ln(w)$. A straightforward transformation obtains the result.

b) Transforming the inequality in *a*),

$$\ln(w-10)+\ln(w+11) \leq 2\ln(w) \tag{93}$$
$$(w-10)(w+11) \leq w^2 \tag{94}$$
$$w \leq 110 \tag{95}$$

c) The individual accepts subject to

$$\ln(109-100)+\ln(109+Y) \geq 2\ln(109) \tag{96}$$
$$Y \geq \frac{109^2}{9} - 109 \tag{97}$$
$$Y \geq 1211.111... \tag{98}$$

d) The individual accepts subject to

$$\ln(200-100)+\ln(200+Z) \geq 2\ln(200) \tag{99}$$
$$Z \geq \frac{200^2}{100} - 200 \tag{100}$$
$$Z \geq 200 \tag{101}$$

e) Whether an individual accepts or rejects a lottery depends on his wealth as much as on the shape of his utility function. Losses that are big with respect to one's wealth require high compensation. In *c*), the individual requests a very high Y to compensate for the risk to lose 100. The expected value of the lottery is then required to be at least £555. It would seem irrational to reject it, but then, one must put oneself in the situation where one risks to lose 92% of one wealth, vs. the chance to multiply it by 11. Would you necessarily accept it?

5.

a) Suppose the agent owns the lottery and prices the ticket for the lottery at p_S. The price p_S at which the agent will sell must be such that $qu(w+G)+(1-q)u(w+L) \leq u(w+p_S)$

b) Suppose the agent does not own the lottery. The maximum price p_B that the agent would be ready to pay for a ticket to play the lottery must be such that $qu(w-p_B+G)+(1-q)u(w-p_B+L) \geq u(w)$.

c) I will have p_S such that

$$
\begin{aligned}
0.5\ln(10+10)+(1-0.5)\ln(10+2) &= \ln(10+p_S)\\
p_S &= \sqrt{(10+10)(10+2)}-10\\
p_S &= 5.4919
\end{aligned}
$$

and I will have p_B such that

$$
\begin{aligned}
0.5\ln(10-p_B+10)+0.5\ln(10-p_B+2) &= \ln 10\\
(10-p_B+10)(10-p_B+2) &= 10^2\\
p_B^2-32p_B+140 &= 0
\end{aligned}
$$

which has got two solutions, $p_1 = \frac{32+\sqrt{464}}{2} = 26.7$ and $p_2 = \frac{32-\sqrt{464}}{2} = 5.2297$. Obviously, p_1 is not reasonable, so $p_B = 5.2297$.

d) p_B is lower than the selling price p_S so we have some explanation from EUT for the apparent paradox of the WTA/WTP disparity. However, in reality, the disparity one finds between WTA and WTP is much higher than what is obtained here, or that obtained under any reasonable functional form for $u(.)$, so EUT is not sufficient to explain the observed disparities.

6.

a) A table that summarizes the situation is as follows:

	S_1	S_2	S_3
A	a	d	d
B	b	b	e
C	c	c	c

A can thus be represented as a lottery over outcomes a,b,c,d,e as $(p_1,0,0,p_2+p_3,0)$. B is $(0,p_1+p_2,0,0,p_3)$ and C is $(0,0,1,0,0)$.

The expected utility of A is $u(A) = p_1u(a)+(1-p_1)u(d)$. Similarly, $u(B) = (1-p_3)u(b)+p_3u(e)$ and $u(C) = u(c)$.

b) Such a pattern of preference is not possible under EUT since it is intransitive.

c) $r(x, y)$ increasing in x means that the higher the good outcome, the higher the rejoicing.

$r(x, y) = -r(y, x)$ means that regret/rejoice is symmetric. Getting the good outcome x rather than the bad outcome y produces the same amount of rejoicing than the amount of regret induced by getting the symmetric outcome.

$r(x, y) > r(x, z) + r(z, y)$ (condition A) means that the rejoicing increases more than proportionately with the difference in outcome. I rejoice more if I gain £100 rather than £0 than if I gain £50 rather than £0 and then £100 rather than £50, even though the result is the same.

d) The agent chooses B over A s.t. $p_1 r(b, a) + p_2 r(b, d) + p_3 r(e, d) \geq 0$.

The agent chooses C over B s.t. $p_1 r(c, b) + p_2 r(c, b) + p_3 r(c, e) \geq 0$

The agent chooses A over C s.t. $p_1 r(a, c) + p_2 r(d, c) + p_3 r(d, c) \geq 0$

Summing up the three equations, one obtains:

$$p_1(r(b,a)+r(c,b)+r(a,c))+p_2(r(b,d)+r(c,b)+r(d,c))+p_3(r(e,d)+r(c,e)+r(d,c)) \geq 0$$

Each of the terms are positive by inequality (A), so the inequalities are consistent with regret theory, meaning that nothing in regret theory precludes that this pattern of intransitive choice could happen (which does not mean it must happen under regret theory!).

e) In the same way as in (d), write the inequalities and sum them up, and you will see that the sum cannot be more than 0 under regret theory. This means that while regret theory would allow a cycle whereby B is preferred to A which is preferred to C which is preferred to B, it does not allow for the opposite cycle whereby A is preferred to B, B is preferred to C which is preferred to A. This type of prediction can then be tested experimentally; verifying such a pattern would give credence to regret theory if it is shown that other alternatives to EUT that allow for intransitivity do not however mandate that cyclical preferences have a specific direction.

27 Game theory 1

1.

There are no pure strategy equilibria of this game. (Note: Any provision of mixed strategy equilibria would be a waste of time and would potentially evidence lack of knowledge of the concept of pure strategy Nash equilibrium).

2.

a) D is strictly dominated by B and also by C. Once D is eliminated, Z is strictly dominated by X.

b) C is weakly dominated by B. Y is weakly dominated by X. In the game that remains, A is weakly dominated by B. Knowing player 1 plays B, player 2 plays X. The equilibrium obtained by iterated elimination of weakly dominated strategies is thus (B, X).

c) NE do not involve any strictly dominated strategies, but may involve weakly dominated ones, so the game to consider involves only A, B and C and W, X and Y.

There are two PSNE of this game, (A, W) and (B, X).

Denoting a, b and c the probabilities for player 1 to play A, B and C respectively, and w, x and y the probabilities for player 2 to play W, X and Y respectively, MSNEs must be such that player 1 must be indifferent between playing A, B or C. I must thus have

$$u_1(A) = 4w + 8y = u_1(B) = 4w + 2x + 6y = u_1(C) = 4w + 2x + 4y$$
$$w + x + y = 1$$

This is possible only if $y = 0$ and $x = 0$ so that $w = 1$.

That equilibrium can hold only if player 2 plays W with probability 1 so one must have:

$$u_2(W) = 6a + 2b + 2c \geq u_2(X) = 4a + 6b + 8c$$
$$u_2(W) = 6a + 2b + 2c \geq u_2(Y) = 4a + 4b + 9c$$
$$a + b + c = 1$$
$$a \geq 0, b \geq 0, c \geq 0$$

This can be rewritten:

$$(6 - 2b)/8 \leq a$$
$$(7 - 5b)/9 \leq a$$
$$a + b + c = 1$$
$$a \geq 0, b \geq 0, c \geq 0$$

or $a \geq \max((6-2b)/8, (7-5b)/9)$ and $a+b+c=1$ and $a \geq 0, b \geq 0, c \geq 0$.

Now, $(6-2b)/8 \geq (7-5b)/9$ when $b \geq 1/11$, so we can summarize by saying that if $b \geq 1/11$ then I must have $a \geq (6-2b)/8$ and $c = 1-a-b$ and $a \geq 0, b \geq 0, c \geq 0$ while if $b \leq 1/11$ then I must have $a \geq (7-5b)/9$ and $c = 1-a-b$ and $a \geq 0, b \geq 0, c \geq 0$.

This is represented in the graph below with b on the horizontal axis and a on the vertical axis. For $b \leq 1/11$, I must have a higher than $(7-5b)/9$ while for $b \geq 1/11$ I must have a higher than $(6-2b)/8$. I must also have $a \leq 1-b$ to ensure that $c \geq 0$.

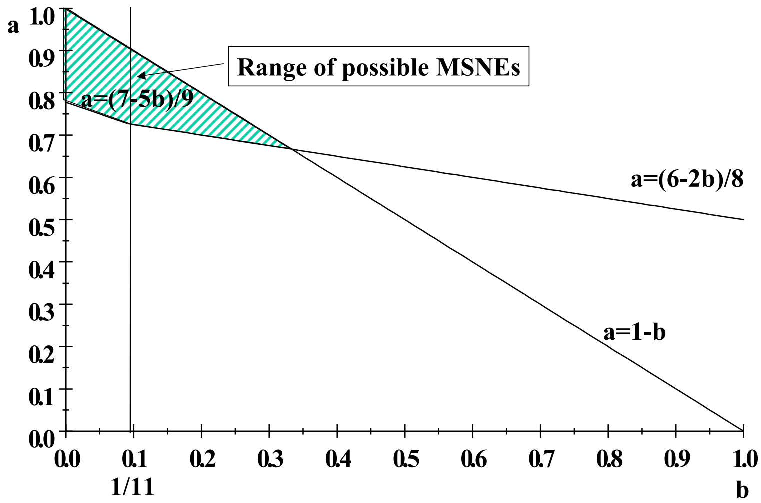

Figure 22: Range of possible MSNEs in Exercise 2, Game Theory 1

This thus defines a range of possible mixed strategies defined by the constraints represented above. Note that the PSNE such that $a = 1$ and $w = 1$ fits the constraints above.

3.

a) The game in normal form is as follows:

	Left	Right
Up	$0, 0$	$4, 1$
Down	$2, 2$	$-1, -1$

b) There are two PSNE in this game, (U, R) and (D, L), and a MSNE s.t.

$$\begin{aligned} u_1(U) &= 4r = u_1(D) = 2l - r \\ l + r &= 1 \\ u_2(L) &= 2d = u_2(R) = u - d \\ u + d &= 1 \end{aligned}$$

with solution $\left[u = \frac{3}{4}, d = \frac{1}{4}, l = \frac{5}{7}, r = \frac{2}{7}\right]$

The only SPNE, obtained by backward induction, is (U, R) as if 1 plays U, then 2 plays R so 1 gets 4 while if 1 plays R, then 2 plays L and 1 gets only 2. Anticipating this, 1 plays U, and this induces 2 to play R.

4.

a) The game in normal form is as follows:

	l	r
L	$3,1$	$0,0$
M	$1,0$	$0,1$
R	$2,2$	$2,2$

b)

Nash equilibria:

Note that M is strictly dominated by R, and can thus be eliminated from consideration.

There are two PSNE in the remaining game, (L, l) and (R, r).

MSNEs must be such that 2 must be indifferent between l and r so it must be that

$$\begin{aligned} p(L) + 2p(R) &= 2p(R) \\ p(L) + p(R) &= 1 \end{aligned}$$

so

$$\begin{aligned} p(L) &= 0 \\ p(R) &= 1 \end{aligned}$$

1 must at least prefer R to L so

$$\begin{aligned} 3p(l) &\leq 2p(l) + 2p(r) \\ p(l) + p(r) &= 1 \end{aligned}$$

The solutions are thus such that $p(l) \leq 2/3$ so the range of MSNEs is such that $p(L) = 0$ and $p(R) = 1$ while $p(l) \leq 2/3$.

Subgame Perfect Nash Equilibria:

There is one SPNE such that 1 plays R.

Consider now SPNE such that 1 plays L or M. Denote the beliefs of 2 as $\mu(R)$, $\mu(L)$, $\mu(M)$, i.e. 2, when it comes its turn to play, believes 1 played R with probability $\mu(R)$. Obviously, when it comes for 2 to play, $\mu(R) = 0$. The expected payoff to 2 of playing l is $\mu(L)$, and of playing r is $\mu(M)$. Therefore,

- If $\mu(L) > \mu(M)$, then 2 plays l and then 1 is better off playing L. By backward induction, under those beliefs, 1 will indeed play L since this gets it 3 rather than 2 if it played R, so there is a sustainable belief such that $\mu(L) = 1$ and $\mu(l) = 1$.
- If $\mu(L) < \mu(M)$, then 2 plays r and 1 is indifferent between playing L or M since both get it payoff 0. But then, by backward induction, 1 is better off playing R so the belief that $\mu(R) = 0$ is not sustainable. We can thus eliminate the possibility that a belief such that $\mu(R) = 0$ and $\mu(L) < \mu(M)$ could be sustained under a SPNE.
- If $\mu(L) = \mu(M) = \frac{1}{2}$ then 2 is indifferent between playing l or r. 1 must be indifferent between playing L or M so I must have $3\mu(l) = \mu(l)$, or $\mu(l) = 0$. But if $\mu(l) = 0$, then by backward induction 1 is better off playing R so the belief that $\mu(R) = 0$ is not sustainable. We can thus eliminate the possibility that a belief such that $\mu(R) = 0$ and $\mu(L) = \mu(M)$ could be sustained under a SPNE.

There are thus two pure SPNE, one such that 1 plays R, the other such that 1 plays L and 2 plays l.

c) In the first SPNE, the payoff to 1 is 2. In the second SPNE, its payoff is 3. By forward induction therefore, the second SPNE will be the equilibrium of the game. Indeed, when it comes the turn for 2 to play, it must assume 1 played L since this is the only way it could get a better payoff than by playing R. Therefore, the only plausible SPNE, obtained by forward induction, is such that 1 plays L and 2 plays l.

5.

a) The game in normal form is as follows:

	l	m	r
L	$1,3$	$1,2$	$4,0$
M	$4,0$	$0,2$	$3,3$
R	$2,4$	$2,4$	$2,4$

b)

Nash equilibria:

There are no strictly dominated strategies in this game. There is one PSNE of this game, (R, m).

MSNEs must be such that player 2 is indifferent between l, m and r. In shorthand, denote p_R as R, p_M as M and p_L as R

$$\begin{aligned}
3L + 4R &= 2L + 2M + 4R \\
3L + 4R &= 3M + 4R \\
L + M + R &= 1
\end{aligned}$$

This means I must have $L = M$ (second equation), which, translated into the first, means $L = M = 0$ and from the third, that $R = 1$.

In that equilibrium, player 1 must at least prefer playing R, so I must have

$$\begin{aligned}
2l + 2m + 2r &\geq l + m + 4r \\
2l + 2m + 2r &\geq 4l + 3r \\
l + m + r &= 1 \\
l \geq 0, m &\geq 0, r \geq 0
\end{aligned}$$

which can be rewritten as

$$\begin{aligned}
3m &\geq 2 - 3l \\
3m &\geq l + 1 \\
l + m + r &= 1 \\
l \geq 0, m &\geq 0, r \geq 0
\end{aligned}$$

or $m \geq \max[(2 - 3l)/3, (l + 1)/3]$ and $l + m + r = 1$ and $l \geq 0, m \geq 0, r \geq 0$.

Now, $(2 - 3l)/3 \geq (l + 1)/3$ when $l \leq 1/4$ so we can summarize by saying that if $l \leq 1/4$ then I must have $m \geq (2 - 3l)/3$ and $r = 1 - l - m$ and $l \geq 0, m \geq 0, r \geq 0$ while if $l \geq 1/4$ then I must have $m \geq (l + 1)/3$ and $r = 1 - l - m$ and $l \geq 0, m \geq 0, r \geq 0$.

This is represented in the graph below with l on the horizontal axis and m on the vertical axis. Note the constraint that I must also have $m \leq 1 - l$ to ensure that $r \geq 0$.

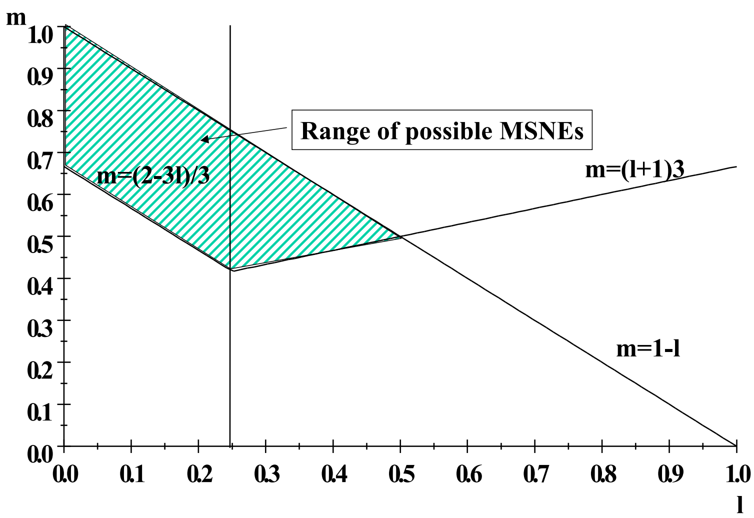

Figure 23: Range of possible MSNEs in Exercise 5, Game Theory 1

Subgame Perfect Nash Equilibria:

1 plays R is a SPNE. Consider now if there are other SPNE that involve 1 playing L or R. Denote $\mu(L)$, $\mu(M)$ and $\mu(R)$ the beliefs of player 2 once it comes its turn to play, *i.e.* its estimated probabilities for 1 to play L, M or R respectively. Obviously, $\mu(R) = 0$ since 2 knows that since it is its turn to play then it must be that 1 did not play R.

Consider the best strategy for 2 given its belief that L is played with probability $\mu(L)$. The expected payoff of playing l is $3\mu(L)$, the expected payoff of playing m is 2, and the expected payoff of playing r is $3\mu(M)$.

- If $3\mu(L) > \max(2, 3\mu(M))$, then 2 plays l, but in that case, by backward induction, 1 is better off playing M, therefore the belief of 2 is not sustainable (it should be $\mu(M) = 1$).
- If $2 > \max(3\mu(L), 3\mu(M))$ then 2 plays m, but in that case, by backward induction, 1 is better off playing R, therefore the belief of 2 is not sustainable.
- If $3\mu(M) > \max(2, 3\mu(L))$ then 2 plays r, but in that case, by backward induction, 1 is better off playing L, therefore the belief of 2 is not sustainable.
- Suppose now $\mu(L) = \frac{2}{3}$ so 2 is indifferent between playing l or m, and does not play r. For this belief to be sustainable, 1 must

be indifferent between playing L or M, and cannot believe that $\mu(r) > 0$. So, denoting $\mu(l)$, $\mu(m)$ the beliefs of 1, I must have

$$\begin{aligned} \mu(l) + \mu(m) &= 4\mu(l) \\ \mu(r) + \mu(l) + \mu(m) &= 1 \\ \mu(r) &= 0 \end{aligned}$$

so I must have

$$\begin{aligned} \mu(l) &= \frac{1}{4} \\ \mu(m) &= \frac{3}{4} \end{aligned}$$

There is thus a SPNE such that 1 plays L with probability $\frac{2}{3}$ and 2 plays l with probability $\frac{1}{4}$.

- Suppose now $\mu(M) = \frac{2}{3}$ so 2 is indifferent between playing m or r, and does not play l. For this belief to be sustainable, 1 must be indifferent between playing L or M and cannot believe $\mu(l) > 0$. So, denoting $\mu(m)$ and $\mu(r)$ the beliefs of 1, I must have

$$\begin{aligned} \mu(m) + 4\mu(r) &= 3\mu(r) \\ \mu(r) + \mu(l) + \mu(m) &= 1 \\ \mu(l) &= 0 \end{aligned}$$

This system does not admit solutions.

There are thus two SPNE: R and a SPNE such that $p(L) = \frac{2}{3}$, $p(M) = \frac{1}{3}$, $\mu(l) = \frac{1}{4}$, $\mu(m) = \frac{3}{4}$.

c) The expected payoff to 1 of playing R is 2 while the expected payoff of playing the other SPNE is 1, so that by forward induction player 1 should play R. The only plausible SPNE is thus for 1 to play R.

6.

a) Players would either agree on (Up, Stop) or (Down,Right), however, it would be very hard for them to reach an agreement on which one to play, so they would probably also agree on using a randomizing device to choose which equilibrium to choose. This randomizing device should guarantee each player the same payoff in expectation, so one would need for this randomizing device to lead with equal probability to one or the other equilibrium being chosen. A coin toss would do the trick.

b) It would seem reasonable for 2 to play Middle. Indeed, while it could make more by playing Stop if 1 plays Up, there is too much of a risk to lose -10 in case 1 plays Down (which may indeed happen if *1* bet on 2 playing Right).

c) There are two PSNE, (Up, Stop) and (Down, Right). The 'reasonable' play by 2 evoked above is thus not a PSNE.

d) Denote $p(\text{Up}) = q$.

i) The expected payoff for 2 if playing Left is $u_2(L) = q - 5(1-q)$, while the payoff of playing Stop is $u_2(S) = 2q - 10(1-q)$, the payoff of playing Right is $u_2(R) = -10q + 2(1-q)$ and the payoff of playing Middle is $u_2(M) = 0$.

Left is thus played subject to $6q - 5 \geq \max(12q - 10, 2 - 12q, 0)$, that is for $q \geq \max((12q-5)/6, (7-12q)/6, 5/6)$, while I must also have $q = q$. There is thus ONE point where L might be played, and that is when $q = 5/6$. It is therefore false to say that L will never be played for any q.

Note: Arguing that Left was never strictly preferred to either of the other actions got the student full marks, with a bonus for those exploring the case where $q = \frac{5}{6}$.

ii) The payoff for 2 of playing Stop is $u_2(S) = 2q - 10(1-q)$ while the payoff of playing Right is $u_2(R) = -10q + 2(1-q)$. Stop is thus preferred to Right when $2q - 10(1-q) > -10q + 2(1-q)$ that is when $q > \frac{1}{2}$, and conversely. Now, if $q = \frac{1}{2}$, *2*'s payoff of playing R and S are negative, so player 2 prefers to play Middle (note the payoff for Left is also negative). In summary, this means that whenever player 2 might be indifferent between S and R, it actually prefers playing something else than either S or R, so that in any situation where either S or R might be played, it must be that the other action is not preferred.

iii) 2 will play Middle subject to $0 \geq \max(q - 5(1-q), 2q - 10(1-q), -10q + 2(1-q))$ or

$$\frac{5}{6} \geq q \geq \frac{1}{6}$$

e) The range where Middle is optimal is large, so that if one assigns equal probability to q being any number in the interval $(\frac{1}{6}, \frac{5}{6})$ (uniform distribution), then the probability to be right in playing Middle is $\frac{2}{3}$, which is more than the probability Stop is optimal $(\frac{1}{6})$ or the probability Right is optimal $(\frac{1}{6})$. Therefore, in the absence of any reasonable belief over the possible actions of 1, it is best to play M as a way to reduce the risk of making a mistake.

28 Game theory 2

1.

a) The game in normal form is as follows:

		2 L	2 R
1	U	$1, 3$	$1, 3$
	D	$0, 0$	$3, 1$

This game has two PSNE, (U, L) and (D, R). Its MSNEs much be such that 2 is indifferent between L and R, so

$$3u = 3u + d$$
$$u + d = 1$$

so that I must have $d = 0$ and $u = 1$. 1 must at least weakly prefer U to D, so I need

$$1 \geq 3r$$
$$l + r = 1$$

so there is a range of MSNE such that $r \leq \frac{1}{3}$ and $u = 1$.

b) Suppose 1 plays U, then its payoff is 1. Suppose now 1 plays D, then 2 will play R and the payoff to 1 will be 3. Therefore, by backward induction, the unique SPNE of the game is (D, R).

c) Player 2 can precommit (see end of "Game Theory 2" for the case where player 2 cannot pre-commit).

i)

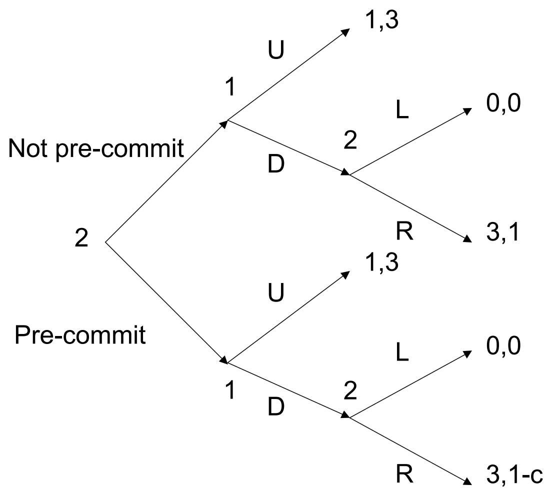

Figure 24: Extensive form game in Exercise 1, Game Theory 2

ii) Put the game in normal form (not necessary for the resolution, but helpful). Denote C the decision to pre-commit and NC the decision not to pre-commit, x the probability with which 2 commits, l the probability with which she plays l and u the probability with which 1 plays U.

	NC, L	NC, R	C, L	C, R
U	$1, 3$	$1, 3$	$1, 3$	$1, 3$
D	$0, 0$	$3, 1$	$0, 0$	$3, 1 - c$

There are three PSNE, $\{U, (NC, L)\}$ and $\{U, (C, L)\}, \{D, (NC, R)\}$.

MSNEs must be such that 2 is indifferent between (NC, L), (NC, R), (C, L) and (C, R) so I must have

$$3u = 3u + (1-u) = 3u = 3u + (1-c)(1-u)$$

for which the unique solution is $u = 1$.

For u to indeed be equal to 1, then 1 must at least weakly prefer U to D so I must have

$$1 \geq 3(1-x)(1-l) + 3x(1-l)$$

so I have a range of MSNEs $\left\{U, l \geq \frac{2}{3}\right\}$.

1. Suppose $1 - c > 0$. Then the player 2 strictly prefers playing R whenever it comes its turn to play. By backward induction, this means player 1 will play D whether 2 pre-commits or not. Knowing this, player 2 does not pre-commit. The SPNE is thus $\{D, (NC, R)\}$.
2. Suppose now $1 - c < 0$ and 2 pre-commits. Then the player 2 strictly prefers playing L if player 1 plays D, in which case 1 gets 0. 1 thus prefers playing U, and 2 thus gets 3. Suppose 2 does not pre-commit. Then, as before, (D, R) is played, and 2 gets only 1. Therefore, 2 prefers to pre-commit and 1 plays U. The SPNE is thus $\{U, (C, L \text{ if pre-commit}, R \text{ if no pre-commit})\}$, equivalent to $\{U, (C, L)\}$.

Exercise: Consider the case where $1 - c = 0$.

iii) 2 will pre-commit s.t $1 - c < 0$. The only effective pre-commitment is the one that is expensive enough that it changes the expected action of 2 in a credible manner, so it is indeed best for 2 to play L when it comes it turn to choose between L and R and it has pre-commited. Note also that the pre-commitment was all the more effective as it was public (known by all). In the variant in appendix A, one will examine how important it is for the pre-commitment to be observable.

2.

a) In the following, I will assume that when the buyer is indifferent between accepting or rejecting the price offer, then it accepts.

1. If $p > 60$ then the seller does not sell. Profit for the seller is 0.
2. If $60 \geq p > 55$, then the seller sells with probability h and its expected profit is $h(p - 50)$.
3. If $55 \geq p > 50$ then the seller always sells and its expected profit is $(p - 50)$.
4. If $p \leq 50$ then the seller always sells and its expected profit is $(p - 50)$

The seller maximizes its payoff. It will set either $p = 60$, in which case profit is $10h$, or $p = 55$, in which case profit is 5. It will choose one over the later s.t. $10h \geq 5$, or $h \geq \frac{1}{2}$.

b)

i)

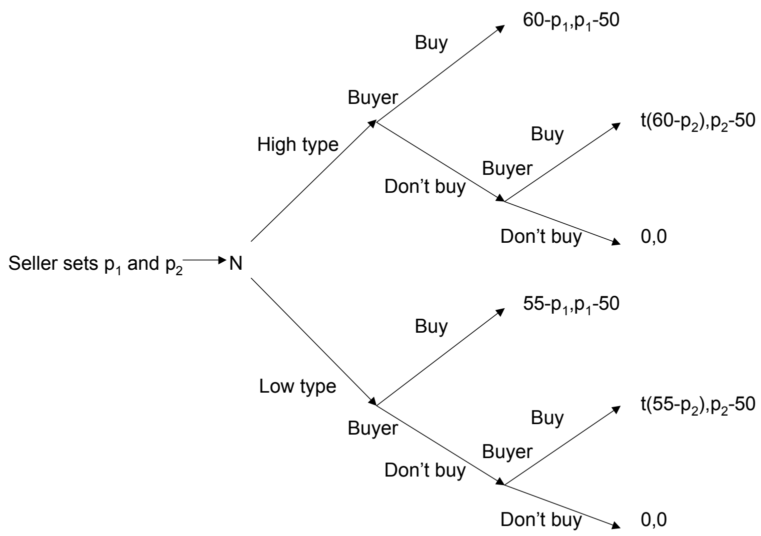

Figure 25: Extensive form game in Exercise 2, Game Theory 2

ii)

1. If the high type accepts in the first period, then its payoff is $60 - p_1 > 0$.
2. If the high type accepts in the second period, then its payoff is $t(60 - p_2) > 0$
3. The low type does not accept in the first period, as its payoff would be $55 - p_1 < 0$.
4. If the low type accepts in the second period, then its payoff is $t(55 - p_2) > 0$.

iii) The conditions for such a pattern of acceptance are such that

$$60 - p_1 \geq t(60 - p_2)$$
$$60 - p_1 \geq 0$$
$$55 - p_1 < 0$$
$$55 - p_2 \geq 0$$

The expected payoff for the seller is then $hp_1 + (1-h)p_2 - 50$, which it maximizes in p_1 and p_2 under the conditions above. This leads to setting $p_2 = 55$ and $p_1 = 60 - 5t$, so her expected payoff is $h(60 - 5t) + (1 - h)55 - 50 = 5(1 + h - ht)$.

iv) It is sustainable, as the buyer of a high type would still be better off buyig first, while the price p_2 cannot be raised to more than the buyer of a low type's willingness to pay, 55, so there is no possibility of exploiting the knowledge that the buyer who does not buy in the first period is of a low type.

3.

a) A player of type $\bar{c}$ is called inefficient because it costs it more to produce the good than the good is valuable. Similarly, the player of type $\underline{c}$ is called efficient because it costs it less to produce the public good than the public good is valuable.

Contributions must be made at the same time else there will be a waiting game whereby the players wait for the other to contribute.

Similarly, it is important they cannot observe what the other player contributes else they could devise strategies such that their contribution is dependent on what the other player contributes. Note this does not require they contribute in succession: each can simply announce their menu of contribution as a function of the contribution of the other and then, having observed each other's menu, choose what to contribute.

Those two assumptions therefore considerably simplify the game to be analyzed.

b) If both players are of a low type, the normal form of the game is as follows:

	C	NC
C	$1-\bar{c}, 1-\bar{c}$	$1-\bar{c}, 1$
NC	$1, 1-\bar{c}$	$0, 0$

Not contributing is a dominant strategy for both players, so the unique PSNE of the game is (NC,NC).

If both players are of a high type, the normal form of the game is as follows:

	C	NC
C	$1-\underline{c}, 1-\underline{c}$	$1-\underline{c}, 1$
NC	$1, 1-\underline{c}$	$0, 0$

There are two PSNE of the game, (NC, C) and (C, NC) and a MSNE such that, denoting c_a the probability a contributes and c_b the probability that b contributes then

$$1-\underline{c} = c_a$$
$$1-\underline{c} = c_b$$

which means both players contribute with probability $1-\underline{c}$.

c)

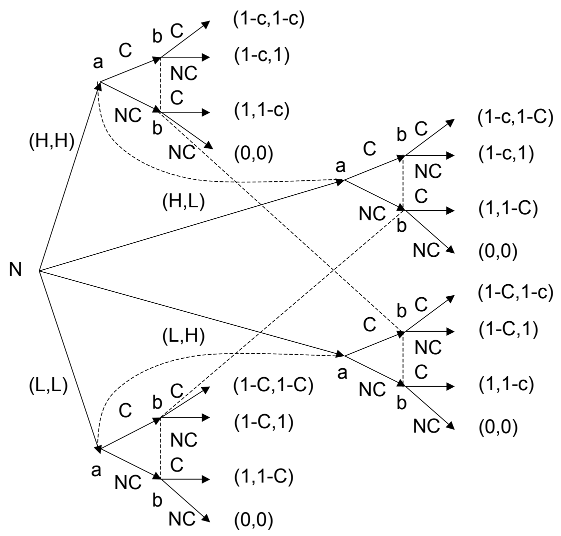

Figure 26: Extensive form game in Exercise 3, Game Theory 2

d) If both agents are of a low type, then none should contribute. If one is of a low type and the other of a high type, then the high type should contribute. If both are of a high type, only one should contribute.

e) Suppose $s_a(c_a)$ takes the form of a mixed strategy, indicating with what probability player a contributes depending on its type c_a. Similarly for $s_b(c_b)$. Strategies must be at least weakly optimal for each players of each different type. Under those conditions, the expected payoff to player a of high type contributing is $1 - \underline{c}$ while the expected payoff of not contributing is $\lambda s_b(\underline{c}) + (1-\lambda)s_b(\overline{c})$, so for the player a of a high type to contribute, I must have

$$1 - \underline{c} \geq \lambda s_b(\underline{c}) + (1-\lambda)s_b(\overline{c})$$

Similarly, player a of a low type will contribute with some probability s.t.

$$1 - \overline{c} \geq \lambda s_b(\underline{c}) + (1-\lambda)s_b(\overline{c})$$

However, the second equality cannot be verified since $1 - \overline{c} < 0$, so I must have $s_a(\overline{c}) = 0$. By symmetry, $s_b(\overline{c}) = 0$. This means that $s_a(\underline{c}) = s_b(\underline{c}) = \min(\frac{1-\underline{c}}{\lambda}, 1)$. Note how this generalizes findings in *b*).

f) A player of a low type again never contributes as contributing is a dominated strategy for that type of player. Player b will interpret a contribution by player a in period 1 as a signal that a is of a high type, since only high types contribute. Therefore, and because a will know b knows it is of a high type, b will not contribute in the second period if a contributes in the first. This means that presumably, in the optimal strategy, a contributes with lower probability in the first period than in the game with only one period in order to avoid revealing its own type as often. (Note: It was not necessary and not required to solve the game).

Variation on 1.c):

Let us consider a variant of exercise 1 section c, when player 2 cannot precommit to spend $c > 0$ when she has to play R : 1 does not know whether 2 precommits or not. This variant will allow us to examine how important it is for the pre-commitment to be observable.

i)

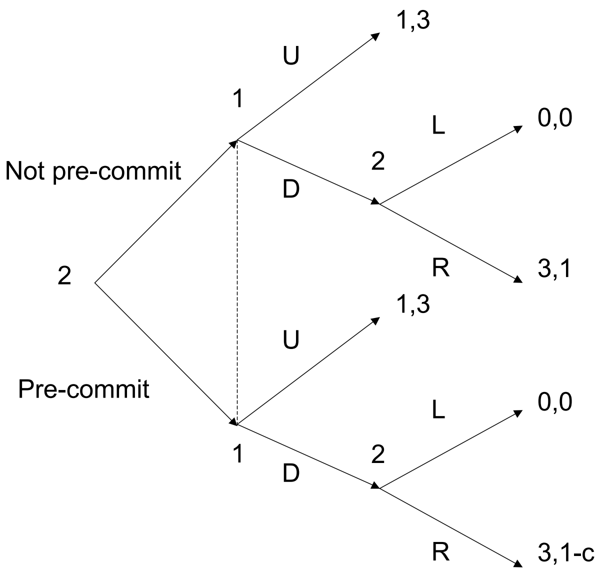

Figure 27: Extensive form game in variant of Exercise 1, Game Theory 2

ii) As before, there are three PSNE, $\{U, (NC, L)\}$, $\{U, (C, L)\}$, $\{D, (NC, R)\}$, and a range of MSNEs $\left\{U, l \geq \frac{2}{3}\right\}$.

1. Suppose $1 - c > 0$, then the player 2 strictly prefers playing R whenever it comes its turn to play. By backward induction, this means player 1 will play D whether it believes 2 pre-committed or not. Knowing this, player 2 does not pre-commit. The SPNE is thus $\{D, (NC, R)\}$.
2. Suppose now $1-c < 0$, then the player 2 strictly prefers playing L if player 1 plays D and 2 pre-committed, in which case 2 gets 0, while it strictly prefers playing R if player 1 plays D and 2 did not pre-commit, in which case 2 gets 1.

   Denote μ the belief of 1 regarding whether 2 did pre-commit. The expected payoff to 1 of playing D is then $3(1-\mu)$ while its expected payoff of playing U is 1. Therefore, if $\mu < \frac{2}{3}$, then 1 plays D while if $\mu > \frac{2}{3}$ then 1 plays U, and if $\mu = \frac{2}{3}$ then 1 is indifferent between U and D. Note that if one denotes l the probability 2 plays l, 1 indifferent between U and D translates in $l = \frac{2}{3}$. The SPNEs are then $\left\{D, (\mu \leq \frac{2}{3}, L \text{ if pre-commit, } R \text{ if no pre-commit})\right\}$ and $\left\{U, (\mu \geq \frac{2}{3}, L \text{ if pre-commit}, R \text{ if no pre-commit})\right\}$.

iii) If $1-c < 0$, 2 will precommit s.t. $\mu > \frac{2}{3}$, so for any $\mu > \frac{2}{3}$ then 2 pre-commits and thus the only reasonable belief by forward induction is $\mu = 1$.

Otherwise, if $\mu < \frac{2}{3}$, then 2 does not pre-commit and thus the only reasonable belief by forward induction is $\mu = 0$.

Finally, if $\mu = \frac{2}{3}$, then 2 commits with probability $\frac{2}{3}$, so $\mu = \frac{2}{3}$ is a reasonable belief.

There are thus only three reasonable SPNEs obtained by forward induction when $1 - c < 0$, which are:

1. $\{D, (\mu = 0, L \text{ if pre-commit, } R \text{ if no pre-commit})\}$, equivalent to $\{D, (NC, R)\}$
2. $\{U, (\mu = 1, L \text{ if pre-commit}, R \text{ if no pre-commit})\}$, equivalent to $\{U, (C, L)\}$
3. $\left\{U, (\mu = \frac{2}{3}, L \text{ if pre-commit}, R \text{ if no pre-commit})\right\}$, equivalent to $\left\{U, l = \frac{2}{3}\right\}$

Whether one or the other obtains is not known, or more precisely, depends on the players' beliefs, while in the original exercise where precommitment was observed, 2 was able to force the belief of player 1 according to her own interest.

If $1-c > 0$, as said above, the SPNE is $\{D, (\mu = 0, L \text{ if pre-commit, } R \text{ if no pre-commit})\}$, equivalent to $\{D, (NC, R)\}$.

The possibility to pre-commit thus adds the possibility of a reasonable SPNE that obtains 2 a higher payoff than without pre-commitment. However, whether it obtains or not depends on what 1 believes 2 believes 1 believes 2 will do.

Note how the process of forward induction allowed one to restrict the set of NEs and SPNEs.

29 Incentive theory[40]

1.

Of the three situations given in the question, case A is the only one where the real moral hazard problem arises.

- First consider case B. Here Low effort gives the principal an expected payoff of $0.1 * 10 + 0.9 * 100 = 91$, whereas High effort gives $0.9 * 10 + 0.1 * 100 = 19$. The principal would prefer the agent to give Low effort. The agent on his part always prefers to give low effort, because it involves low disutility. So in this case there is no moral hazard.
- Next, consider case C. Here, High effort gives 100 for sure. So if the principal ever observes an outcome of 10, he will conclude that the agent must have given Low effort. This helps the principal write a contract in which an outcome of 10 is always punishable. The punishment can be made appropriately large so that the agent will be deterred from giving low effort. So here though the moral hazard problem potentially exists, it can be solved easily as effort is in effect observable.
- Now consider case A. Here Low effort gives a lower expected payoff (37), and High effort gives a higher expected payoff (91). The principal prefers High effort, but the agent prefers Low effort, and the principal cannot determine the agent's effort from the outcome observed. Hence, there is a moral hazard problem.

2.

a) $p_G R = 120 > K = 100$ while $p_B R + B = 30 < K = 100$

b) Moral hazard constraint: $p_G w > B$

Participation constraint: $p_G w > 0$

Limited liability constraint is ensured as $w > 0$

Therefore, I need $w > \frac{B}{p_G} = 60$

c) The investor will choose w as small as possible to ensure good behaviour, so $w = 60$. In that case, the return to the investor is $p_G(R - w) - K = 90 - 100 = -10$. The investor does not invest due to moral hazard on the part of the manager. The investor may try to get a manager with some capital to invest in the project, or try to limit the private benefits B the manager can get when she behaves badly.

3.

[40] Answer keys for exercises 1, 3 and 4 were originally written by Dr Bibhas Saha.

a) Full information solution: The monopolist's problem is:

$$\max \pi_i = R_i - cq_i$$
$$\text{Subject to } U_i = q_i[\theta_i - q_i/2] - R_i \geq 0$$

This comes down to maximizing $\pi_i = q_i[\theta i - q_i/2] - cq_i$. The solution is $q_i = \theta_i - c$, for $i = 1, 2$. The monopolist will sell this much quantity to each type and will charge $R_i = \frac{(\theta_i^2 - c^2)}{2}$.

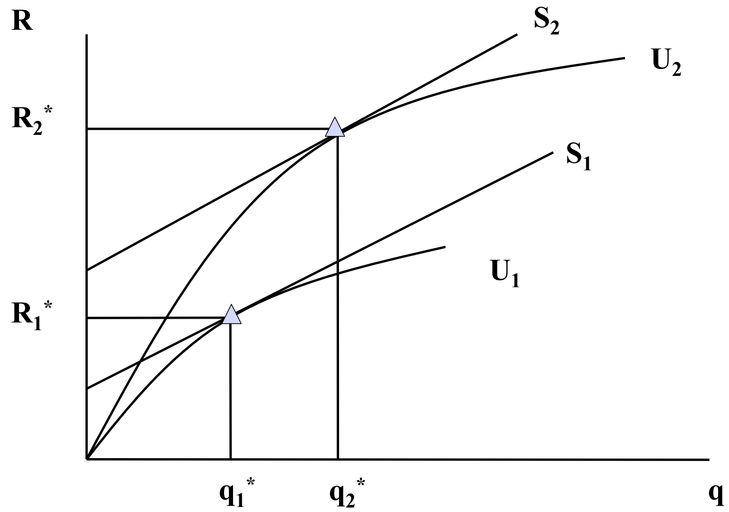

Figure 28: First best solution in Exercise 3, Incentive Theory

In the above figure we show the full information solution as it occurs at the tangency point of the indifference curve and iso-profit curve for each type. Note that on the (q, R) plane, higher indifference curve means lower utility and higher iso-profit curve means higher profit. Also the indifference curve of type 2 is steeper than that of type 1. In the figure we have drawn the indifference curves at the reservation utility level (which must pass through origin).

b) See the graph below:

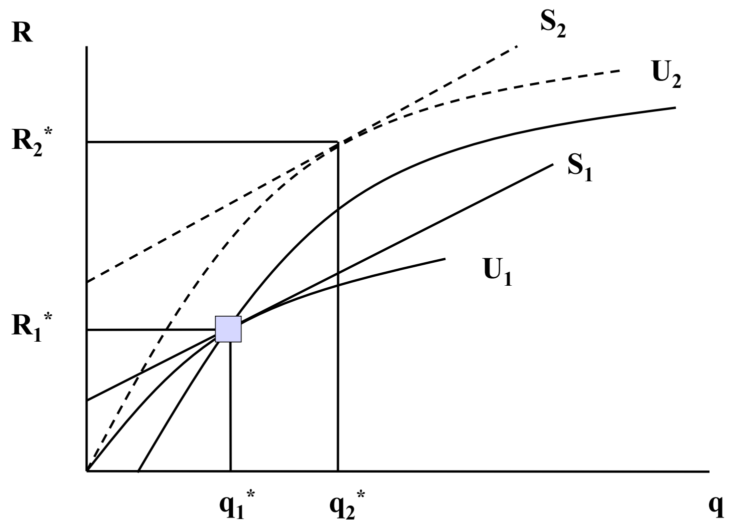

Figure 29: Outcome of the first best contracts under asymmetric information in Exercise 3, Incentive Theory

If the monopolist does not have full information and yet it offers (q_1, R_1) and (q_2, R_2) both, both types of consumers will choose (q_1, R_1). As evident from the above graph, type 2 consumers' utility increases (thick indifference curve below the dotted one), while the type 1 consumers cannot choose anything else. Hence, the monopolist will sell only (q_1, R_1) bundle and get $N(R_1 - cq_1)$ as profit.

c) The monopolist's problem is now modified as;

- Maximize $E\pi_i = N[\alpha\{R_i - cq_1\} + (1-\alpha)\{R_2 - cq_2\}]$

Subject to

Self-selection constraints:

- $U_1 = q_1[\theta_1 - q_1/2] - R_1 \geq q_2[\theta_1 - q_2/2] - R_2$
- $U_2 = q_2[\theta_2 - q_2/2] - R_2 \geq q_1[\theta_2 - q_1/2] - R_1$

and participation constraints

- $U_1 \geq 0$
- $U_2 \geq 0.$

Of these the second self-selection constraint will bind and the first participation constraint will bind. This is easy to check. Just assume otherwise and see that you will run into contradictions.

No distortion at the top means that $q_2 = \theta_2 - c$ must hold. The efficiency condition will hold for the top type. No rent at the bottom means that $U_1 = 0$ the reservation utility, which is confirmed at the binding of the first participation constraint.

Substitute into the objective function the expressions for R_1 and R_2 obtained by setting the second self-selection constraint and the first participation constraint with equality, and then maximise it with respect to q_1 and q_2. You should get the following:

- $q_2 = \theta_2 - c$

and

- $q_1 = \theta_1 - c - \{(1-\alpha)/\alpha\}[\theta_2 - \theta_1]$.

Associated payments (R_1, R_2) can be easily calculated.

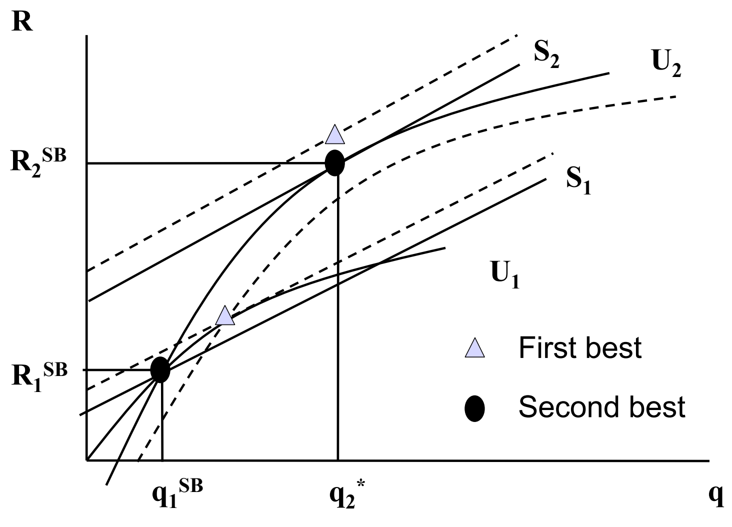

Figure 30: Second best solution in Exercise 3, Incentive Theory

In this figure we have shown the optimal solution. Optimal q_1 has now fallen as the monopolist's choice is now given at the point of intersection of the two solid indifference curves. This follows from the fact that the self-selection constraint must bind for the top type. He should be indifferent between choosing his own bundle and the one meant for the bottom type. Choice of q_2 is unaffected, but R_2 falls, implying that concession has been given to the top type.

4.

a) The second firm is the top type. It is the efficient type as it can supply a given level of cleanliness at a lower cost.

b) Full information solution. The hospital's problem is to

- Maximise $U = AI - P$
- Subject to $P - kI^2 \geq 10$.

It is straightforward to obtain: $I = \frac{A}{2k}$ and $P = 10 + \frac{A^2}{4k}$.

c) The hospital's problem is to

- Maximise $EU = (1/2)[AI_1 - P_1] + (1/2)[AI_2 - P_2]$

Subject to,

Self-selection constraints:

- $P_1 - I_1^2 \geq P_2 - I_2^2$
- $P_2 - kI_2^2 \geq P_1 - kI_1^2$

Participation constraints:

- $P_1 - I_1^2 \geq 10$
- $P_2 - kI_2^2 \geq 10$.

d) The second self-selection constraint and the first participation constraint will bind. By setting equality to these two constraints obtain expressions for P_1 and P_2, substituting them into the objective function and maximising it, one obtains the following:

- $I_1 = A/[2(2-k)] < A/2$, because $k < 1$, and
- $I_2 = A/(2k)$.

You can derive P_1 and P_2 accordingly.

e) See the following graphs. In the top panel we have depicted the full information solution. The bottom panel shows how the optimal solution will look like under asymmetric information. P_2 rises above the full information solution, while I_1 is below its full information level. If these two bundles $\{(I_1^{SB}, P_1^{SB}), (I_2^{SB}, P_2^{SB})\}$ are offered together, the top type will select (I_2^{SB}, P_2^{SB}) as intended.

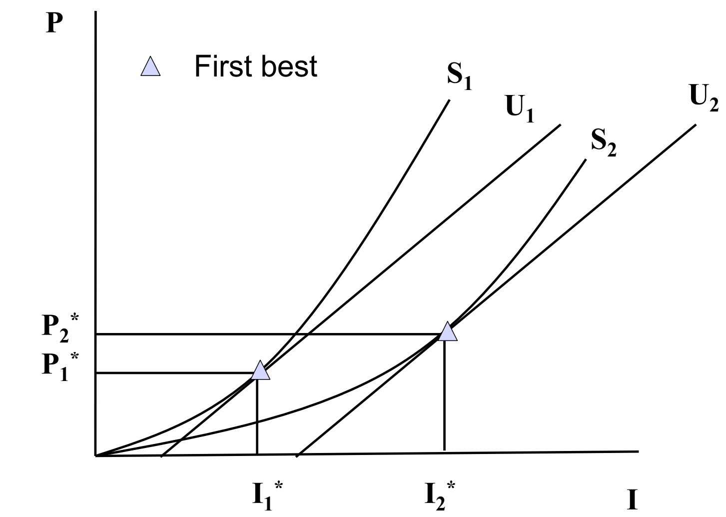

Figure 31: First best solution in Exercise 4, Incentive Theory

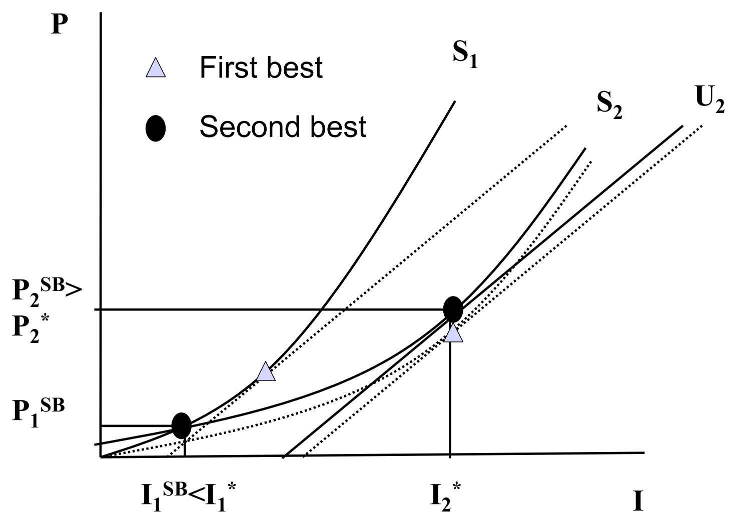

Figure 32: Second best solution in Exercise 4, Incentive Theory

5.

a)

i) The agent that does not exert effort will have expected payoff:

$$EU_{e=0} = p_0 u(Q+\overline{w}) + (1-p_0)u(\underline{w})$$

The agent that exerts effort will have expected payoff:

$$EU_{e=1} = p_1 u(Q+\overline{w}) + (1-p_1)u(\underline{w}) - C$$

Since effort is monitorable, the agent's incentive compatibility constraint ($(p_1-p_0)(u(Q+\overline{w})-u(\underline{w})) \geq C$) does not come into account, and one only has to check that $EU_{e=1} > 0$.

ii) The employer will maximize its objective function $EU_{e=1}$ (the utility of the agent) s.t. $p_1\overline{w} + (1-p_1)\underline{w} \geq 0$ (this is the budget constraint for the employer)

$$\begin{aligned} &\max p_1 u(Q+\overline{w}) + (1-p_1)u(\underline{w}) - C \\ \text{s.t. } &p_1\overline{w} + (1-p_1)\underline{w} \geq 0 \end{aligned}$$

iii) The budget constraint will be binding, so $p_1\overline{w} + (1-p_1)\underline{w} = 0$ and the government maximizes

$$\max_{\bar{t}} p_1 u(Q+\overline{w}) + (1-p_1)u\left(\frac{-p_1\overline{w}}{1-p_1}\right) - C$$

so that at the optimum

$$u'(Q+\overline{w}) = u'\left(\frac{-p_1\overline{w}}{1-p_1}\right)$$

so

$$u'(Q+\overline{w}) = u'(\underline{w})$$

so s.t. monotonicity conditions on u', then $Q+\overline{w} = \underline{w}$ (this implies that $\overline{w} < 0$). This is explained by the wish of the government to minimize the variability of the income of the agent as the agent is risk averse. Note that if effort wasn't verifiable, then the agent would not do effort as income is the same either way.

b)

i) If effort is not verifiable, then the agent must be induced to choose to do effort, so the program of the principal is

$$\begin{aligned} &\max p_1 u(Q+\overline{w}) + (1-p_1)u(\underline{w}) - C \\ \text{s.t. } &p_1\overline{w} + (1-p_1)\underline{w} \geq 0 \text{ (budget constraint)} \\ &(p_1-p_0)(u(Q+\overline{w}) - u(\underline{w})) \geq C \text{ (incentive constraint)} \end{aligned}$$

ii) The Lagrangian of the optimization problem is

$$\mathcal{L} = p_1 u(\overline{w}) + (1-p_1)u(\underline{w}) - C - \mu(p_1\overline{w} + (1-p_1)\underline{w}) - \lambda(C_1 - (p_1-p_0)(u(Q+\overline{w}) - u(\underline{w})))$$

The first order conditions for maximization are as follows:

$$p_1 u'(\overline{w}) - \mu p_1 + \lambda(p_1 - p_0)u'(Q + \overline{w}) = 0$$
$$(1 - p_1)u'(\underline{w}) - \mu(1 - p_1) - \lambda(p_1 - p_0)u'(\underline{w}) = 0$$

The first condition can be rewritten as:

$$u'(\overline{w}) + \lambda(1 - \frac{p_0}{p_1})u'(Q + \overline{w}) = \mu$$

From this we can already see that $\mu > 0$ as long as λ is positive (since $u'(\overline{w}) \geq 0$ (increasing utility function)).

Replacing this expression of μ into the second condition, we find

$$(1 - \lambda\frac{p_1 - p_0}{1 - p_1})u'(\underline{w}) = u'(\overline{w}) + \lambda(1 - \frac{\pi_0}{\pi_1})u'(Q + \overline{w})$$
$$u'(\underline{w}) - u'(\overline{w}) = \lambda[\frac{p_1 - p_0}{1 - p_1}u'(\underline{w}) + (1 - \frac{\pi_0}{\pi_1})u'(Q + \overline{w})]$$

Then, $\lambda \geq 0$ if $u'(\underline{w}) - u'(\overline{w}) > 0$ which is the case as long as $\underline{w} \leq \overline{w}$. In that case, both λ and μ will be positive at the optimum so the two constraints, budget and incentive, will be binding:

$$\pi_1\overline{w} + (1 - \pi_1)\underline{w} = 0$$
$$C - (p_1 - p_0)(u(\overline{w}) - u(\underline{w})) = 0$$

This means that in order to induce effort, the government will have to differentiate payment between the high and the low payoff in order to induce the agent to seek to increase the probability of the high payoff by exerting effort. This comes at the expense of making the payoff to the agent more uncertain.

6.

a) The principal is the firm as it draws up the contract, while the agent is the employee. The contract will be drawn based on two variables, w the wage and h the number of hours of work, and the wage and number of hours worked may be made dependent on the result of the shock (positive or negative).

b)

- At time $t = 1$, the principal offers a contract such that $C_H = (w_H, h_H)$ is to be paid and worked in case of a favourable shock and $C_L = (w_L, h_L)$ is to be paid and worked in case of an adverse shock.
- At time $t = 2$, the agent accepts or reject the contract. If the contract is rejected, the game ends.

- At time $t = 3$, the outcome is realized.
- At time $t = 4$, transfers are made according to the contract and the realization of the outcome.

c) The principal's objective function is its expected profit $E\pi$:

$$E\pi = q(x_H h_H - w_H) + (1-q)(x_L h_L - w_L)$$

The agent's individual rationality constraint is such that its expected utility EU is positive:

d) The Lagrangian of the optimization problem is

$$\mathcal{L} = q(x_H h_H - w_H) + (1-q)(x_L h_L - w_L) + \lambda(q(u(w_H) - f(h_H)) + (1-q)(u(w_L) - f(h_L)))$$

with λ the Lagrange multiplier for the rationality constraint.

Maximization in λ obtains that

$$q(u(w_H) - f(h_H)) + (1-q)(u(w_H) - f(h_H)) = 0$$

so the rationality constraint is binding.

Maximization in h_H, h_L, w_H and w_L obtains the following first order conditions:

$$\begin{aligned} x_H - \lambda f'(h_H) &= 0 \\ x_L - \lambda f'(h_L) &= 0 \\ -1 + \lambda u'(w_H) &= 0 \\ -1 + \lambda u'(w_L) &= 0 \end{aligned}$$

e) From the above equalities, I obtain that

$$\begin{aligned} \frac{x_H}{f'(h_H)} &= \lambda \\ \frac{x_L}{f'(h_L)} &= \lambda \\ \lambda u'(w_H) &= 1 \\ \lambda u'(w_L) &= 1 \end{aligned}$$

so $u'(w_H) = u'(w_L)$ which means that $w_H = w_L$. The agent is paid the same wage irrespective of the shock because the agent is risk averse. Now, from the two first equalities, I have that $\frac{x_H}{f'(h_H)} = \frac{x_L}{f'(h_L)}$ so

$$f'(h_L) = \frac{x_L}{x_H} f'(h_H)$$

Since $x_L < x_H$ this means that

$$f'(h_L) < f'(h_H)$$

so the agent works less in case of a negative shock (this is because f' is increasing as f is convex).

f) Under the contract above, the principal has to pay the same wage whether the agent works h_H hours or h_L hours. This means that in case of a negative shock, the principal will want to lie and pretend there was a positive shock so the agent has to work more.

g) One must include incentive constraints for the principal so it does not lie, so one needs:

$$x_H h_H - w_H \geq x_H h_L - w_L$$

so the principal pays w_H and asks the worker to work h_H hours in case of a positive shock, and

$$x_L h_L - w_L \geq x_L h_H - w_H$$

so the principal pays w_L and asks the worker to work h_L hours in case of a negative shock.

Adding up those two conditions, one finds that

$$\begin{aligned} x_H h_H - w_H + x_L h_L - w_L &\geq x_H h_L - w_L + x_L h_H - w_H \\ (x_H - x_L) h_H &\geq (x_H - x_L) h_L \\ h_H &\geq h_L \end{aligned}$$

so the agent will be asked to work more in case of a positive shock. Then, looking at the second incentive constraint:

$$w_H - w_L \geq x_L (h_H - h_L)$$

so the agent will be paid more in case of a positive shock than in case of a negative shock. This means that the first best outcome cannot be achieved if the principal can lie. The possibility of lying works to the detriment of the principal.