ECN103: The Economics of Uncertainty
Lecture Notes

Andrés Carvajal[1]

Fall quarter, 2019

[Version: December 6, 2019]

[1] University of California–Davis, Department of Economics, acarvajal@ucdavis.edu

A. Carvajal

Contents

Syllabus

I. Instructor: Andrés Carvajal
Office number: SSH1101
e-Mail address: **acarvajal@ucdavis.edu**
Office hours: Mon, 10h00 – 12h00, or by appointment
Lectures: Tue and Thu, 13h40 – 15h00, at Olson 118

II. Teaching Assistant: Keisuke Teeple
Office number: SSH120
e-Mail address: **kwteeple@ucdavis.edu**
Office hours: Wed, 16h00 – 18h00
Discussion: Thu, 18h10 – 19h00 at SocSci 90 and 19h10 – 20h00, at Wellman 202

III. Goals: Students will become comfortable with the most important results of the theory of decision-making under uncertainty, including portfolio theory and counter-intuitive comparative statics; and with the theory of incentives and imperfect information. Importantly, the students will acquire the analytical abilities to understand the existing literature as well as modelling abilities to apply these tools to new settings.

IV. Evaluation: The mark for the course will based on weekly problem sets, a midterm and a final exam. The problem sets will be due at the end of the lectures on Thursday of each week; the midterm will be on *October 31*; the final on *December 11 at 18h00*. The average mark of the problem sets will be 20% of the final mark; the midterm exam will count 30%, and the remaining 50% of the mark will be the final exam.

V. Pre-requisites: Students are required to know basic calculus, probability and intermediate microeconomics. In particular, it is assumed that the students have a good understanding of classical demand theory [JR §3.2 to §4.3] and producer theory [JR §5].

VI. Academic dishonesty: Please see the attached code of conduct. For the protection of the integrity of our institution and the honest students, any student found violating the code of conduct will be reported and given the most severe consequences we have the authority to impose. This will amount to an F in the entire course.

VII. Outline:

Week 1: Review of probability theory: random variables; distribution, probability and density functions; expectation and variance.

Weeks 2–6: Decision theory under uncertainty: Preferences under uncertainty; Expected Utility Theory; Risk attitudes; Optimal savings; Optimal output; Stochastic dominance; Optimal portfolio; Non-EU theories. [Lecture notes, JR §2.4 and W §4.1 to §4.4, §5.2 and §6.2]

Weeks 7–10: An introduction to the economics of information: Insurance markets, à la Rotschild–Stiglitz; Akerlof's markets for lemons; Capital markets with adverse selection; Spence's Screening problem; Credit rationing in financial markets, à la Stiglitz–Weiss; Investment and financial markets, à la Bernanke–Gertler. [Lecture notes, JR §10.1, FR §2.3 and §5.4, and W §9]

Week 11: An introduction to the economics of incentives: Agency theory and moral hazard.[Lecture notes, JR §10.2, W §11, FR §2.5, §4 and §5.5]

VIII. References: Lecture notes are available for all the content of the course. The references above are for:

FR: Freixas, X., and J.C. Rochet (2008). *Microeconomics of banking*, 2/e. MIT Press.

JR: Jehle, G.A., and P.J. Reny (2011). *Advanced microeconomic theory*, 2/e. Pearson: Prentice Hall.

W: Wolfstetter, E. (1999). *Topics in microeconomics: Industrial organization, auctions and incentives*. Cambridge University Press.

None of these books is mandatory, but the closest one to our material, both in terms of content and of level, is JR.

Part I

Decision Theory

Note 1

Preferences

We, human beings, are complicated. Our behavior is difficult to model. Here, we consider the problem of a decision-maker who has to choose from a set of alternatives. The decision-maker could be a person, or a group of people (for instance a family), but for simplicity we will treat the decision-maker as a person. After a brief general treatment, we will consider a set of alternatives that is natural for the case of uncertainty.

In order to make the problem tractable, we will abstract from the question of how the set of alternatives is determined, and will model the person's behavior through just one key element: what she likes.

1.1. Choice space and preferences

Consider a situation in which a person faces a nonempty set $\mathcal{D}$ of alternatives. We refer to $\mathcal{D}$ as the *choice space*. The problem we study in decision theory is how the person makes her choice, when she is allowed to pick one alternative from $\mathcal{D}$, or perhaps from a subset of it. The key step in our analysis of the person's choice is to model "what she wants." For us, the individual's preferences are subjective judgments about the relative desirability of the available choices: given two alternatives, preferences are defined by her answer to the question "Is the first alternative at least as good as the second one?" Formally, then, the decision-maker's *preferences* are a binary relation $\succsim$ defined on the choice set: given a pair of alternatives p and p', we write $p \succsim p'$ if, according to the person's tastes, p is at least as good as p'. We take the person's preferences as exogenous, in the sense that we do not explain where they come from. Instead, we concentrate on the problem of studying the individual's behaviour given her preferences, under the assumption that these preferences will *not* be affected by the person's choices.

We start by studying properties that the individual's preferences may (but need not) satisfy. We first study properties of a binary relation under which it makes sense to identify this relation with someone's preferences: binary relation $\succsim$ is *complete* if for any p and any $p' \neq p$, either $p \succsim p'$ or $p' \succsim p$; it is *reflexive* if for any p, $p \succsim p$; and it is *transitive* if $p \succsim p'$ and $p' \succsim p''$ imply $p \succsim p''$. If the relation complete, reflexive and transitive, we say that it is *rational*.

Reflexivity is imposed for consistency with our interpretation of 'weak' preference. The other two properties are less innocuous. Decision-makers with incomplete preferences may find instances in which they are unable to choose: they simply cannot make a value judgment about the relative (subjective) quality of two alternatives. People with non-transitive preferences are open to full rent extraction, as a person could find a cycle of choices for which the person is willing to pay a positive premium at each step.

In economics, one usually assumes that the decision-maker under consideration has rational preferences, although in some cases —e.g. very complicated problems— it may be reasonable to consider that individual's preferences are incomplete; also, some cases of nontransitive preferences are sometimes observed in real life. In any case, from now on we fix a rational binary relation $\succsim$, and define the following (induced) binary relations on the choice set: the *strict preference* relation, $\succ$, by saying that $p \succ p'$ if it is not true that $p' \succsim p$; and the *indifference* relation, $\sim$, by saying that $p \sim p'$ if it is true that $p \succsim p'$ and $p' \succsim p$.

EXAMPLE 1.1. Argue that $\succ$ is transitive, but not reflexive. Could this relation be complete?

Answer: For transitivity, note that $x \succ y$ and $y \succ z$ imply that $x \succsim y$ and $y \succsim z$. The latter implies that $x \succsim z$, by transitivity of $\succsim$. If it was not true that $x \succ z$, then, by definition and completeness of $\succsim$, it would have to be true that $z \succsim x$, and hence, by transitivity of $\succsim$, that $z \succsim y$, which contradicts the fact that $y \succ z$.

To see that $\succ$ is *not* reflexive, note that if for some x it was true that $x \succ x$, then it would not be true that $x \succsim x$, by definition of $\succ$, but this would contradict reflexivity of $\succsim$.

Relation $\succ$ can be complete: imagine the relation $>$, derived from $\geq$ over the real line. In many cases in economics, though, $\succ$ is incomplete. The best example is given by consumer theory under certainty. □

EXERCISE 1.1. Argue that $\succ$ is transitive, but not reflexive, and that $\sim$ is reflexive and transitive. Could this relation be complete? Could it be rational?

It is most usual in economics to represent a decision-maker's preferences by a function that gives a higher value the more the person likes an alternative. That is, we say that binary relation $\succsim$ is *represented by function* $U : \mathcal{D} \to \mathbb{R}$ if $U(p) \geq U(p')$ occurs when, and only when, $p \succsim p'$. We say that $\succsim$ is *representable* if there is some function U that represents it.

The function U that represents $\succsim$ is called *utility function*. Notice that if a preference relation is representable, then there are infinitely many different utility functions that represent it. All these representations will have the same level sets (i.e. the same ordinal information), but may give nontrivially different utility levels (i.e. different cardinal information). It is for this reason that interpersonal comparisons of utility are problematic.

EXERCISE 1.2. Argue that representability implies rationality. Do you think that rationality implies representability?

Note that in some cases the existence of a utility function that represents an individual's preferences is very easy to establish. For example, if $\mathcal{D}$ is finite, then any complete preference relation on it will be representable. But there are also well-known cases of preference relations that cannot be represented by utility functions, because, in some sense, they partition the choice space into "too many" equivalence classes.

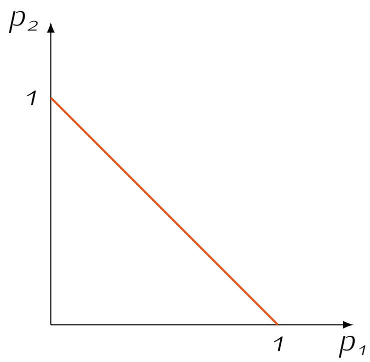

Figure 1.1: Δ when $X = 2$

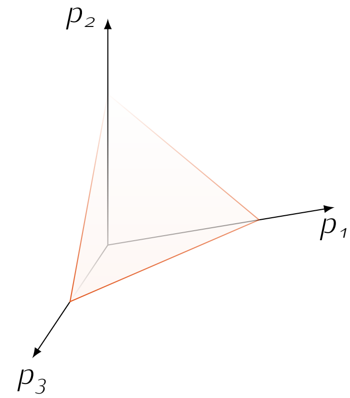

Figure 1.2: Δ when $X = 3$

1.2. Uncertainty and expected utility

The main interest in these lectures will be the case in which the consequences of the decision-maker's choices are not fully determined by her, and are subject to uncertainty. For this purpose, it is useful to endow our problem with a probabilistic framework. Let $\mathcal{X} \neq \emptyset$ be a set of *outcomes*. This could be an abstract set, or, if you would like more definiteness, a set $\mathcal{X} \subseteq \mathbb{R}$ of monetary values. Also, let Δ be the space of all probability distributions over $\mathcal{X}$. For example, if there are only finitely many possible outcomes, we can write $\mathcal{X} = \{1, \ldots, X\}$, and the set of probability distributions is $\Delta = \{p \in \mathbb{R}^X_+ \mid \sum_x p_x = 1\}$. Figs. 1.1 and 1.2 represent, respectively, the space Δ of lotteries when $X = 2$ and $X = 3$. For spaces with more than three elements, it is not possible to have a graphical representation.

In the current setting of uncertainty, we may want to have special properties on the utility function that represents the person's preferences: we say that $\succsim$ has an *expected utility (EU) representation* if there exists a function $u : \mathcal{X} \to \mathbb{R}$ such that for any pair of lotteries p and p', we have that $p \succsim p'$ if, and only if, $\mathrm{E}_p(u) \geq \mathrm{E}_{p'}(u)$. In this case, we can define the utility function over lotteries $U(p) = \mathrm{E}_p(u)$, and it is immediate that U represents $\succsim$.

When a preference relation has EU representation, *any* function U that represents them is known as *expected utility* or *von Neumann–Morgenstern function*. Given a U, the associated function u is known as its *Bernoulli index*.

Example 1.2. A decision-maker faces uncertainty over the set of outcomes $\mathcal{X} = \{1, 2, 3, 4\}$. Her preferences over lotteries are given by

$$p \succsim p' \text{ if, and only if, } p_1 + p_2 \geq p'_1 + p'_2.$$

Find a Bernoulli index that gives an EU representation of these preferences.[1]

Answer: Any utility index $u : \{1, 2, 3, 4\} \to \mathbb{R}$ such that $u(1) = u(2) > u(3) = u(4)$ will do; the index with $u(1) = u(2) = 1$ and $u(3) = u(4) = 0$ is the most immediate choice. □

Example 1.3. A decision-maker faces uncertainty over the set of outcomes $\mathcal{X} = \{1, 2\}$. Her preferences over lotteries are represented by the following function:

$$U(p) = \begin{cases} p_1 & \text{if } p_2 < 1/2; \\ 0, & \text{otherwise.} \end{cases}$$

Can these preferences have an EU representation?

Answer: No, they cannot. To see this, suppose, by way of contradiction, that there exists an EU representation, and let its Bernoulli index be u. Then, $u(1) = U((1, 0)) > U((0, 1)) = u(2)$, which implies that, for any $0 < p \leq 1$, $pu(1) + (1 - p)u(2) > u(2)$. This would imply that $p \succ (0, 1)$ whenever $p_2 < 1$, which is not the case. □

Example 1.4. A decision maker faces uncertainty over the set of outcomes $\mathcal{X} = \{1, 2\}$. Her preferences over lotteries are that $p \succsim \tilde{p}$ when, and only when, $p_1 p_2 \geq \tilde{p}_1 \tilde{p}_2$.

a. Is she rational?

b. Draw her indifference map.

c. Do her preferences admit an EU representation?

Answer: a. Yes: any preference relation that can be represented by a utility function is complete, reflexive and transitive; in this case, function $U(p) = p_1 p_2$ represents the preferences.

b. Fig 1.3 shows three indifference curves: one of them is the red dot, located at $(1/2, 1/2)$; another one is given by the two blue dots, located at $(1/4, 3/4)$ and $(3/4, 1/4)$; and the third one is given by the two green dots, which are located at $(0, 1)$ and $(1, 0)$. To the right of $(1/2, 1/2)$ the direction if improvement is up and to the left; to the left, it is down and to the right. This is indicated by the six arrows in the figure.

c. No, they do not. According to her preferences $(1, 0) \sim (0, 1)$. If these preferences had an EU representation, it should satisfy $u(1) = u(2)$, in which case all lotteries would be indifferent for her, and her map would consist of only one indifference curve. This is not the case, as seen in part (b). □

Example 1.5. A decision maker faces uncertainty over the set of outcomes $\mathcal{X} = \{1, 2, \ldots, X\}$. Her preferences over lotteries are that $p \succsim \tilde{p}$ when, and only when,

$$\alpha \max\{p_1, p_2, \ldots, p_X\} + \beta \min\{p_1, p_2, \ldots, p_X\} \geq \alpha \max\{\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_X\} + \beta \min\{\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_X\},$$

where $\alpha \geq \beta$ are fixed coefficients.

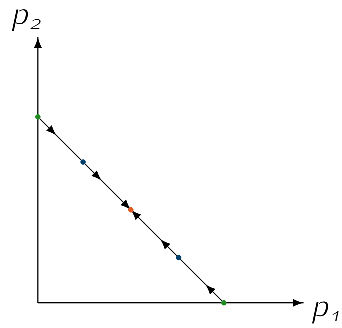

Figure 1.3: Indifference map for Example 1.4

a. Is she rational?

b. Assuming that there are only two states, draw her indifference map and determine conditions under which her preferences admit an EU representation.

c. Do the conditions you just gave in part (b) suffice when there are three states of the world?

Answer: a. Yes: function

$$U(p) = \alpha \max\{p_1, p_2, \ldots, p_X\} + \beta \min\{p_1, p_2, \ldots, p_X\}$$

represents these preferences, and any representable relation is rational.

b. Below the 45° line, $p_1 > p_2$ so the utility level is $U = \alpha p_1 + \beta p_2$; above it $p_2 > p_1$ so the utility level is $U = \alpha p_2 + \beta p_1$; at the line, $U = (\alpha + \beta)/2$. Fig 1.4 shows three indifference curves for the case $\alpha > \beta$: one of them is the red dot, located at $(1/2, 1/2)$; another one is given by the two blue dots, located at $(1/4, 3/4)$ and $(3/4, 1/4)$; and the third one is given by the two green dots, which are located at $(0, 1)$ and $(1, 0)$. To the left of $(1/2, 1/2)$ the direction of improvement is up and to the left; to the right, it is down and to the right. This is indicated by the four arrows in the figure.

When $\alpha = \beta$, the utility is $U = \alpha + \beta$ everywhere, so the map consists of a single indifference curve covering the whole space of lotteries.

For the case $\alpha > \beta$, the map is inconsistent with EU. If $\alpha = \beta$, the preferences have an EU representation with $u(1) = u(2)$.

c. The condition is insufficient. Suppose that $\alpha = \beta = 1$, and $X = 3$. By definition, $(1, 0, 0) \sim (0, 1, 0) \sim (0, 0, 1)$, so an EU representation would require that $u(1) = u(2) = u(3)$. This would imply that $(1, 0, 0) \sim (1/3, 1/3, 1/3)$, but

$$\max\{1, 0, 0\} + \min\{1, 0, 0\} = 1 \neq \frac{2}{3} = \max\left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\} + \min\left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\}.$$

For $X \geq 3$, EU would require $\alpha = \beta = 0$.

□

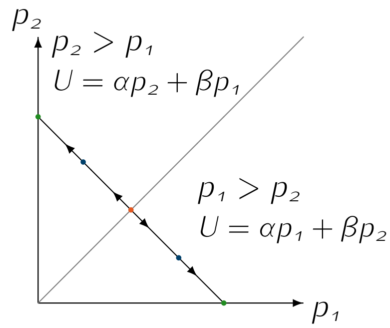

Figure 1.4: Indifference map for Example 1.5, when $\alpha > \beta$

Exercise 1.3. A decision-maker faces uncertainty over the set of outcomes $\mathcal{X} = \{1, 2, 3\}$. Suppose that her preferences are represented by $U(p) = \max\{p_1, p_2\} + ap_3$, where $a > 0$ is a parameter.

a. Rank the following lotteries according to this person's preferences: $p^1 = (0, 1, 0)$, $p^2 = (1/2, 1/2, 0)$, $p^3 = (2/3, 1/3, 0)$, and $p^4 = (0, 0, 1)$.

b. Could these preferences have an EU representation?

c. If this decision-maker could have any lottery p that she wants, would she ever choose one with $p_1 > 0$ and $p_2 > 0$? How about one with $p_3 \neq 0$?

Exercise 1.4. Suppose that there are three possible states of the word, and an individual's preference relation over lotteries is such that $(1, 0, 0) \sim (0, 1, 0) \sim (0, 0, 1)$. Argue that if this person's preferences admit an EU representation, then she is indifferent between all possible lotteries.

Exercise 1.5. Suppose that a person's income is a random variable X defined on the interval $[0, 4]$, and her preferences are as follows: for any pair of lotteries p and p', $p \succsim p'$ if, and only if, the probability of earning at least \$2 under lottery p is greater than or equal to the probability of earning at least \$2 under lottery p'.

a. Argue that these preferences admit an EU representation.

b. Argue that this person strictly prefers to have her income distributed uniformly in $(0, 3)$ than to have \$1.5 for sure.

Note that, by definition, any increasing transformation of a utility function represents the same preference relation as the function itself. For this reason, in economics, we often say that utility functions have only "ordinal meaning." In the case of preferences over lotteries the same is true, but one has to be very careful: any increasing transformation of function U will represent the same preferences, but the same is true only for some very restricted transformations of u!

Exercise 1.6. Argue that if an individual's preferences have an EU representation with Bernoulli index $u(x)$, then the index $\tilde{u}(x) = au(x) + b$, for any numbers $a > 0$ and b, also gives an EU representation of $\succsim$.[2]

The exercise tells us that the expectations of increasing *affine* transformations of u do represent the same preference relation.[3] The following proposition, which we state with no proof,

[2] This is, it is true that $p \succsim p'$ if, and only if, $\sum_{x=1}^{X} p_x \tilde{u}(x) \geq \sum_{x=1}^{X} p'_x \tilde{u}(x)$?

[3] Function f is affine if for some numbers α and β, we have that $f(u) = \alpha + \beta u$.

says that any other transformations of u will change the preferences they represent. For this reason, one often refers to the Bernoulli index as "cardinal utility index."

Proposition 1.1. *Suppose that U and $\tilde{U}$ are EU representations of $\succsim$. Let u and $\tilde{u}$ be their respective Bernoulli utility indices. Then, there exist numbers α and $\beta > 0$ such that $\tilde{u}(x) = \alpha + \beta u(x)$ for every x.*

Reading. The following is an excerpt of an article published in *The New York Times*. Discuss it critically.

> *Las Vegas uses flashing lights and ringing bells to create an illusion of reward and to encourage risk taking. Insurance company offices present a more somber mood to remind us of our mortality. Every marketer knows that context and presentation influence our decisions.*
>
> *For the first time, economists are studying these phenomena scientifically. The economists are using a new technology that allows them to trace the activity of neurons inside the brain and thereby study how emotions influence our choices, including economic choices like gambles and investments.*
>
> *For instance, when humans are in a "positive arousal state," they think about prospective benefits and enjoy the feeling of risk. All of us are familiar with the giddy excitement that accompanies a triumph. Camelia Kuhnen and Brian Knutson, two researchers at Stanford University, have found that people are more likely to take a foolish risk when their brains show this kind of activation.*
>
> *But when people think about costs, they use different brain modules and become more anxious. They play it too safe, at least in the laboratory. Furthermore, people are especially afraid of ambiguous risks with unknown odds. This may help explain why so many investors are reluctant to seek out foreign stock markets, even when they could diversify their portfolios at low cost.*
>
> *If one truth shines through, it is that people are not consistent or fully rational decision makers. Peter L. Bossaerts, an economics professor at the California Institute of Technology, has found that brains assess risk and return separately, rather than making a single calculation of what economists call expected utility.*

From: Cowen, T., *The New York Times*, Economic Scene (April 20, 2006) "Enter the Neuro-Economists: Why Do Investors Do What They Do?"

Note 2

Risk aversion

WE NOW WANT TO MODEL THE DECISION-MAKER'S TASTE FOR RISK.[1] We do this by studying the decision-maker's ranking of any risky lottery and the (degenerate) lottery that gives her the expected return of the risky lottery with probability 1. For this to be possible, we must abandon the simplifying assumption that $\mathcal{X}$ is some finite, abstract set. Instead, we now assume instead that $\mathcal{X}$ is some interval in $\mathbb{R}$, which we can interpret as the space of wealth levels, X, of the individual. A lottery in this setting is a probability distribution p over $\mathcal{X}$, and we restrict attention to lotteries for which an expected payoff is defined: there exists a real number $\mathrm{E}_p(X)$, such that $\int x \, \mathrm{d}p(x) = \mathrm{E}_p(X)$. For simplicity of notation, we identify a wealth level x with the *degenerate* lottery that gives that prize with probability 1.

A risk neutral decision-maker is one who is always indifferent between the expected payoff of a lottery, *for sure*, to the lottery itself. The individual is risk averse if she finds the sure expected payoff at least as good to the lottery, strictly so if she strictly prefers the expected payoff when the lottery is non-trivial. Formally, binary relation $\succsim$ is:

1. *risk neutral* if for any lottery p, $\mathrm{E}_p(X) \sim p$;

2. *risk averse* if for any lottery p, $\mathrm{E}_p(X) \succsim p$; and

3. *strictly risk averse* if for any lottery p such that $\mathrm{E}_p(X) \neq p$, $\mathrm{E}_p(X) \succ p$.

4. *risk loving* if for any lottery p, $p \succsim \mathrm{E}_p(X)$.

2.1. Expected utility and risk aversion

LET US ASSUME THAT RELATION $\succsim$ has an EU representation, with Bernoulli index $u : \mathcal{X} \to \mathbb{R}$. The following result tells us that the decision-makers attitude toward risk is characterized by the shape of index u.

PROPOSITION 2.1. *EU-representable relation $\succsim$ is:*

1. *risk neutral if, and only if, u is affine.*[2]

[1] We may be tempted to say that if $\succsim$ is convex, in the sense that for all p the set $\{p' : p' \succsim p\}$ is convex, then she dislikes risk. But this would be a mistake: a convex combination of lotteries does not reduce their riskiness..

[2] Function u is affine if for some numbers α and β, we have that $u(x) = \alpha + \beta x$.

2. *risk averse if, and only if, u is concave;*

3. *strictly risk averse if, and only if, u is strongly concave; and*

4. *risk loving if, and only if, u is convex;*

We can easily see that this proposition is true, for a simplified case: lotteries that have only two possible prizes, so that $\mathcal{X} = \{x, x'\}$. In this setting, if lottery p gives probability π to prize x and $1 - \pi$ to x', risk aversion implies that

$$u\left(\pi x + (1-\pi)x'\right) \geq \pi u(x) + (1-\pi)u(x'),$$

so it must be that the Bernoulli utility index u is concave. Under strict risk aversion, if p is not degenerate, the latter inequality must always be strict, so u must be strongly concave. Finally, under risk neutrality we must always have an equality, so u must be both concave and convex, hence affine.

Example 2.1. Consider a decision-maker whose preferences over lotteries admit an EU representation with Bernoulli utility index $u(x) = \sqrt{x}$, for non-negative income levels.

a. Can she be characterized as risk averse, strictly risk averse, or risk neutral?

b. How does your answer change if $u(x) = x$?

c. What if $u(x) = x^2$?

Answer: a. Function $u(x) = \sqrt{x}$ is strictly concave, so the first individual is strictly risk averse.

b. Function $u(x) = x$ is linear, so in this case the agent is risk neutral.

c. Finally, function $u(x) = x^2$ is (strictly) convex, so this individual loves risk!

□

Exercise 2.1. Consider a decision-maker whose preferences over lotteries admit an EU representation with Bernoulli utility index

$$u(x) = \begin{cases} 4x, & \text{if } x < 4; \\ x^2, & \text{otherwise.} \end{cases}$$

Can she be characterized as risk averse, strictly risk averse, risk loving or risk neutral?

2.2. Willingness to pay for insurance

Two measures of how averse to risk a person is are widely used. Given the result above, it is not surprising that these measures are based on the curvature of the Bernoulli utility index. For the remainder of the section, we assume that u is differentiable twice; it is by its second derivative that we will capture the curvature of u. We also assume that $u' > 0$.

Suppose that our decision-maker has an income subject to risk, as it is determined by the lottery p. Let the expectation and the variance of this lottery be, respectively, $\mathrm{E}_p(X) = \bar{X}$ and $\mathrm{V}_p(X) = \mathrm{E}_p[(X - \bar{X})^2] = \Sigma$. How much would she be willing to pay in order to secure an income of $\bar{X}$ instead of running the lottery? Let Γ be this number, so that[3]

$$u(\bar{X} - \Gamma) = U(p) = \mathrm{E}_p[u(X)].$$

Example 2.2. Consider the decision-maker of Example 2.1, and suppose that her income is distributed uniformly over the interval $(0, 1)$. How much would she be willing to pay to insure against this risk?

Answer: When $u(x) = \sqrt{x}$, we want to find Γ such that

$$\sqrt{1/2 - \Gamma} = \int_0^1 \sqrt{x}\,\mathrm{d}x = 2/3,$$

so $\Gamma = 1/18$. For $u(x) = x$, $1/2 - \Gamma = \int_0^1 x\,\mathrm{d}x = 1/2$, so $\Gamma = 0$. Finally, if $u(x) = x^2$,

$$(1/2 - \Gamma)^2 = \int_0^1 x^2\,\mathrm{d}x = 1/3,$$

so in this case $\Gamma = 1/2 - 1/\sqrt{3} < 0$. □

Example 2.3. *A decision-maker has preferences over lotteries that admit an EU representation with cardinal utility index*

$$u(x) = \begin{cases} 2x, & \text{if } x \leq 5; \\ \frac{15}{2} + \frac{1}{2}x, & \text{if } x > 5. \end{cases}$$

a. *What is her attitude toward risk?*

b. *Suppose that her expected income is \$3, but it is subject to a random shock as follows: with probability 1/2, she gains or loses \$1. How much is she willing to pay to insure against this shock?*

c. *Suppose now that her expected income is \$7 while the shock is as before. Is she willing to pay more or less to insure against the shock?*

d. *Suppose, finally, that the person's expected income is \$5, and the shock is still as before. Is she now willing to pay more or less to insure against the shock? How do you make sense of your answers?*

Answer: a. Fig 2.1 shows this person's Bernoulli index. It is a convex function, but not strictly, so she is risk averse, but not strictly. As the function is *not* linear, she is not risk neutral.

[3] After the following expression, if there is no room for confusion we will not write the distribution explicitly, so that E will appear instead of E_p.

b. In this case, her income can be \$2 or \$4 with equal probability. As these values are both on the part of the domain where the function has a constant slope of 2, she behaves as if she were risk neutral, and is willing to pay \$0 for insurance.

c. Now her income can be \$6 or \$8, which lie on the part of the domain where the function has a constant slope of 1/2. Again, she behaves as if she were risk neutral and is willing to pay \$0 for insurance.

d. This case is different, as her low income of \$4 lies in a different part of the domain from her high income of \$6. Here, concavity of the index means that she will be willing to pay a strictly positive premium for insurance. To confirm this, we must solve

$$2 \times (5 - \Gamma) = \tfrac{1}{2}\left[2 \times 4 + \left(\tfrac{15}{2} + \tfrac{1}{2} \times 6\right)\right],$$

which gives $\Gamma = 3/8 > 0$.

The intuition for these seemingly odd answers is that this person is rather odd herself: for small risks when her income is very low or sufficiently high, she acts as if she were risk neutral; but, at an intermediate income (\$5) she behaves as if she were strictly risk averse.

□

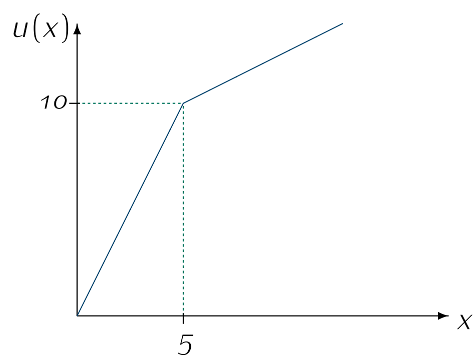

Figure 2.1: Bernoulli index for Example 2.3

Exercise 2.2. Consider the decision-maker of Exercise 2.1.

a. Suppose that this person's wealth is \$1, and she is offered the following gamble: a fair coin is tossed; if it comes up tails, she wins \$3; if it comes up heads, she loses \$1. How much is she willing to pay for this gamble?

b. If her initial wealth was \$5, would she be willing to pay more? Why?

2.3. Absolute risk aversion

FINDING Γ IN GENERAL CAN be complicated, but we can study the last equality by approximating its terms. For the right-hand side, notice that

$$\mathrm{E}\,(u) \approx \mathrm{E}\left[u(\bar{X}) + u'(\bar{X})(X-\bar{X}) + \frac{1}{2}u''(\bar{X})(X-\bar{X})^2\right] = u(\bar{X}) + \frac{1}{2}u''(\bar{X})\Sigma.$$

For the left-hand side, $u(\bar{X}-\Gamma) \approx u(\bar{X}) - u'(\bar{X})\Gamma$. Equating, we get that

$$\Gamma \approx -\frac{1}{2}\frac{u''(\bar{X})}{u'(\bar{X})}\Sigma.$$

It follows that the *coefficient of absolute risk aversion*, defined as

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

is a good measure of the individual's aversion to risk, when her expected wealth is x. To see what type of risk is captured by this coefficient, notice that we could reinterpret things as if the individual had an expected income $\bar{X}$, subject to shocks $Z = X - \bar{X}$, where the shocks have mean $\mathrm{E}\,(Z) = 0$ and variance $\mathrm{V}(Z) = \mathrm{E}\,(Z^2) = \Sigma$. Because of this, one usually understands that $A(\bar{X})$ is a good measure of the individual's attitude to *additive* risk (of mean 0 and variance 2).

EXAMPLE 2.4. Consider the decision-makers of Example 2.1. What are these individuals' coefficients of absolute risk aversion? How do they depend on the individuals' expected income?

Answer: If $u(x) = \sqrt{x}$, then $A(x) = 1/(2x)$, decreasing in income. If $u(x) = x$, then $A(x) = 0$, which is independent of income. Finally, if $u(x) = x^2$, then $A(x) = -1/x$, which is increasing in income. □

EXERCISE 2.3. Consider a decision-maker whose preferences over lotteries admit an EU representation with Bernoulli utility index $u(x) = -\mathrm{e}^{-\alpha x}$.

a. What is the individual's coefficient of absolute risk aversion, and how does it depend on the individual's expected income?

b. What if
$$u(x) = \frac{x^{1-\alpha} - 1}{1-\alpha},$$
where $\rho \neq 1$?

c. What if $u(x) = \ln x$?[4]

EXERCISE 2.4. Consider a decision-maker facing risk. Her preferences over lotteries admit an EU representation with Bernoulli utility index $u(x) = \sqrt{x}$. Suppose that her income is random, and is uniformly distributed over the interval $[0, 4]$. What proportion of her expected income would she be willing to pay as a premium to insure against her income risk?

[4] It's useful to note that this function is the limit of the previous one, as $\alpha \to 1$.

Exercise 2.5 (Harder). Consider a decision-maker whose preferences over lotteries admit an EU representation with Bernoulli utility index $u(x) = -1/x$. She expects to have an income of $w > 0$. Suppose that she discovers that her income is subject to a random shock, so her actual income will be $w + X$, where X is a random variable following the uniform distribution over the interval $[-\bar{x}, \bar{x}]$, with $0 < \bar{x} < w$.

a. What is the decision maker's expected utility in the absence of any shocks to her income? What is her expected utility in the presence of the shock to her income? What is her expected income? Does this make sense? Show that $\lim_{\bar{x}\to 0} \mathrm{E}\,[u(X)] = u(w)$.

b. How much would she be willing to pay to insure against the random shock? What is her coefficient of absolute risk aversion evaluated at her expected income? Let these values be $\Gamma(w, \bar{x})$ and $A(w, \bar{x})$, respectively.

c. Show that $\lim_{\bar{x}\to 0} \Gamma(w, \bar{x}) = 0$ but, still, $\lim_{\bar{x}\to 0} A(w, \bar{x}) = 2/w \neq 0$. Does this make sense?

d. An insurance company offers this person partial insurance as follows: for a premium of Γ, she can reduce her risk to half, so that her income is $w+X/2$. Write the equation that is necessary to solve to determine how much would she be willing to pay for this contract? What is your intuition for the sign of this number?

2.4. Relative risk aversion

Suppose now that, in the same situation as before, we ask what proportion of her expected income the decision-maker would be willing to spend to secure her income? Letting γ be that proportion, we have that

$$u(\bar{X} - \gamma\bar{X}) = U(p) = \mathrm{E}_p[u(X)].$$

Now,

$$u(\bar{X} - \gamma\bar{X}) \approx u(\bar{X}) - u'(\bar{X})\gamma\bar{X},$$

and, hence,

$$\gamma \approx -\frac{1}{2}\frac{u''(\bar{X})}{u'(\bar{X})\bar{X}}\Sigma = -\frac{1}{2}\frac{u''(\bar{X})\bar{X}}{u'(\bar{X})}\sigma,$$

where

$$\sigma = \frac{\Sigma}{\bar{X}^2} = \mathrm{E}\left(\frac{X-\bar{X}}{\bar{X}}\right)^2 = \mathrm{V}\left(\frac{X-\bar{X}}{\bar{X}}\right).$$

We thus get that the *coefficient of relative risk aversion*,

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

is a second measure of the individual's attitude to risk, when her expected wealth is x. The difference is that this measure is designed for *multiplicative risk* of mean 0 and variance 2: suppose that the individual's expected income is $\bar{X}$, but it is subject to proportional shocks $\zeta\bar{X}$, with the random variable

$$\zeta = \frac{X-\bar{X}}{\bar{X}}.$$

Example 2.5. Consider the decision-makers of Example 2.1. What are these individuals' coefficients of relative risk aversion? How do they depend on the individuals' expected incomes?

Answer: If $u(x) = \sqrt{x}$, then $R(\bar{X}) = 1/2$. If $u(x) = x$, then $R(\bar{X}) = 0$. Finally, if $u(x) = x^2$, then $A(\bar{X}) = -1$. In all cases, the coefficients are independent of income. □

Exercise 2.6. Consider the decision-makers of Exercise 2.3. What are these individuals' coefficients of relative risk aversion? How do they depend on the individuals' expected incomes?

Reading. The following is an excerpt from an article in *The Economist*. Discuss it critically.

> *"Prospect theory" is an important contribution to the study of economics. It challenges some of the fundamental assumptions that economists have made concerning human behaviour...*
>
> *To understand prospect theory, you need to know what it disagrees with. The villain in this drama is "expected utility theory", a series of assumptions about human behaviour. Expected utility theory says that people are good at assessing probabilities — people know that their aeroplane (almost definitely) will not crash, so do not feel nervous about the flight. The theory also states that people experience good things or bad things equally — it is as pleasurable to find £5 on the street, as it is painful to lose £5.*
>
> *Prospect theory disagrees with these two assumptions (among others). Its followers have used a barrage of psychological tests—the new vogue in economics—to reach their conclusions. The theory shows that people are in fact terrible at assessing probabilities (they are poor at "probability weighting"). People feel nervous on planes, and no amount of statistical reasoning will rid them of their anxiety.*
>
> *The theory also shows that people find bad things relatively worse than they find good things good. People tend to find losing £5 agonizing, yet are only mildly happy to find £5 on the floor.*

From: Free exchange, *The Economist*, Economics (August 5th, 2013) "Future prostects: Prospect theory and economics".

Note 3

Precautionary savings demand

Consider the case of an expected utility maximizer who has to decide how much to save. Suppose that there are two periods, and the individual's utility over consumption in the two periods, c_1 and C_2, is given by $c_1 + \mathrm{E}\left[u(C_2)\right]$, where we assume that c_1 is known with certainty, while C_2 is a random variable.[1] As before, $u(c_2)$ represents the utility that the individual would derive from knowing that her second-period consumption will be c_2 with certainty.

Now, suppose that the individual has a fixed income $\bar{x}_1$ in the first period, while in the second period she faces uncertainty and her income is the random variable X_2. In the first period, she can invest an amount κ, in exchange for an interest rate ι in the future. That is, if she invests κ, her date-1 consumption is $c_1 = \bar{x}_1 - \kappa$, while it is $C_2 = X_2 + (1 + \iota)\kappa$ in the second period.

3.1. Savings under certainty

Let us remove all uncertainty from the individual's problem, by assuming that her second-period income is fixed at $\bar{X}_2$. Then, her decision problem is

$$\max_{\kappa} \left\{ \bar{x}_1 - \kappa + u(\bar{X}_2 + (1 + \iota)\kappa) \right\}.$$

Letting $\bar{k}$ be the optimal investment, the first-order condition it satisfies is that

$$1 = (1 + \iota) \cdot u'(\bar{X}_2 + (1 + \iota)\bar{k}).$$

Without specifying a functional form for u, this equation characterizes the individual's optimal choice under certainty.

3.2. Precautionary savings

Let us now assume that the individual faces uncertainty about her future income, so that X_2 is a random variable with expectation $\mathrm{E}\left(X_2\right) = \bar{X}_2$ and variance σ^2. If the individual is risk averse, will she necessarily save more than under certainty? In other words, if k^* is the solution to decision problem

$$\max_{\kappa} \{\bar{x}_1 - \kappa + \mathrm{E}\ \left[u\left(X_2 + (1 + \iota)\kappa\right)\right]\},$$

[1] We are simplifying the analysis slightly by assuming that the individual's intertemporal utility is linear in her first-period consumption. Very little changes if her preferences are of the form $v(c_1) + \mathrm{E}\left[u(C_2)\right]$.

and the Bernoulli utility index is concave, is it necessarily the case that $k^* > \bar{k}$?

If the answer to this question is Yes, then it must be that the derivative of the latter maximand is strictly positive when $\kappa = \bar{k}$, namely that

$$-1 + (1+\iota)\mathrm{E}\left[u'(X_2 + (1+\iota)\bar{k})\right] > 0. \tag{3.1}$$

As before, we cannot solve this equation without specifying a functional form for u. But we can again take a second-order approximation of u' around $\bar{X}_2 + (1+\iota)\bar{k}$, to get that $u'(X_2 + (1+\iota)\bar{k})$ is approximately equal to

$$u'(\bar{X}_2 + (1+\iota)\bar{k}) + u''(\bar{X}_2 + (1+\iota)\bar{k})(X_2 - \bar{X}_2) + \frac{1}{2}u'''(\bar{X}_2 + (1+\iota)\bar{k})(X_2 - \bar{X}_2)^2.$$

Taking expectations,

$$\mathrm{E}\left[u'(X_2 + (1+\iota)\bar{k})\right] \approx u'(\bar{X}_2 + (1+\iota)\bar{k}) + \frac{1}{2}\sigma^2 u'''(\bar{X}_2 + (1+\iota)\bar{k}),$$

so that Eq. (3.1) becomes

$$-1 + (1+\iota)u'(\bar{X}_2 + (1+\iota)\bar{k}) + \frac{1+\iota}{2}\sigma^2 u'''(\bar{X}_2 + (1+\iota)\bar{k}) > 0.$$

Now, the first two terms in this inequality cancel each other, so the required condition becomes

$$\frac{1+\iota}{2}\sigma^2 u'''(\bar{X}_2 + (1+\iota)\bar{k}) > 0.$$

Note that concavity of u does not suffice for this inequality, for it only is guaranteed to hold when $u''' > 0$. This condition is known as *prudence*. It does not mean that the individual is risk averse, but that her coefficient of absolute risk aversion is decreasing.

The intuition for this result is simple. Note that the first-order conditions of their respective problems require that

$$u'(\bar{X}_2 + (1+\iota)\bar{k}) = \frac{1}{1+\iota} \tag{3.2}$$

and

$$\mathrm{E}\left[u'(X_2 + (1+\iota)k^*)\right] = \frac{1}{1+\iota}. \tag{3.3}$$

When $u''' > 0$, the marginal utility function u' is convex, as in Fig. 3.1. Risk around $\bar{X}_2$ means, then, that

$$\mathrm{E}\left[u'(X_2 + (1+\iota)\bar{k})\right] > u'(\bar{X}_2 + (1+\iota)\bar{k}).$$

This means that $\bar{k}$ is not optimal under risk, and the individual will choose $k^* > \bar{k}$, in order to lower $\mathrm{E}[u']$ to $1/(1+\iota)$ again. Notice, by the same reason, that if the marginal utility function is linear no adjustment is necessary, while if it is concave the adjustment required is to decrease the level of savings.

Example 3.1. Suppose that $u(x) = \ln(x)$. Compare $\bar{k}$ and k^*. How do you make sense of this result?

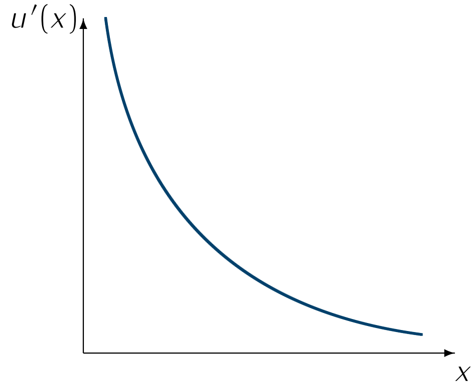

Figure 3.1: Prudence: u', when $u''' > 0$.

Answer: With $u'(x) = 1/x$, and Eq (3.2) becomes

$$1 = \frac{1+\iota}{\bar{X}_2 + (1+\iota)\bar{k}}, \tag{$*$}$$

while Eq. (3.3) is

$$1 = \mathrm{E}\left[\frac{1+\iota}{X_2 + (1+\iota)k^*}\right] \tag{$**$}$$

in this particular case. The graph of $u'(x)$ is a parabola like the one that appears in Fig. 3.1. It follows from that graph that, for any k,

$$u'(E[X_2] + (1+\iota)k) < \mathrm{E}\left[u'(X_2 + (1+\iota)k)\right]$$

From comparing Eqs. ($*$) and ($**$), it then follows that $\bar{k} < k^*$. This is not surprising, as this individual's marginal utility function is convex, so the agent is prudent. In the presence of risk, this agent does demand strictly positive precautionary savings. □

Exercise 3.1. Suppose that $u(x) = -(S-x)^2$, for a (very large) positive constant S.[2] Compare $\bar{k}$ and k^*. How do you make sense of this result?

Exercise 3.2. Suppose that $u(x) = \alpha + \beta x$, for a positive constant β.[3] Compare $\bar{k}$ and k^*. How do you make sense of this result?

Exercise 3.3. In the context of a two-period savings problem, suppose that the individual has a fixed income $\bar{x}$ in both periods. In the first period, she can invest an amount k, in exchange for an interest rate I in the future, where I is a random variable with expectation $\bar{I} > 0$ and variance $\Sigma > 0$. The individual's utility over consumption in the two periods, c_1 and C_2, is given by $c_1 + \mathrm{E}\left[u(C_2)\right]$, where u is a strictly increasing and strictly concave cardinal utility index.

[2] Suppose that S is large enough that the inequality $C_2 < S$ can be taken for granted. This consumption level constitutes a "satiation point."

[3] The sign of α does not matter.

a. Assuming that the interest rate is fixed at $\bar{l}$, find a condition that characterizes her optimal level of savings, $\bar{k}$.

b. Now, allowing for uncertainty on the interest rate, find a condition that characterizes her optimal level of savings, k^*.

c. Assuming now that (over the relevant domain) the individual's marginal utility is (positive, decreasing and) linear, compare $\bar{k}$ and k^*. How do you make sense of this result?

READING. The following excerpt is from an article in *The Economist*. Discuss it critically.

> *Most people save to insure themselves against income and life shocks, for their children's education, and for their retirement. How much you actually save for these uses depends on your income, level of risk aversion, and how risky you perceive your environment to be. Were the Chinese less risk averse than Americans, they might still save more based on perceptions of the relative likelihood of income and other financial shocks. The Great Moderation in America produced a long period of low consumption volatility. Americans may have come to believe that such placidity would persist in the future, leading even the most risk averse to reduce saving.*
>
> *Given this, the most effective fiscal policy for China might be the construction of a better welfare state. But would that cause the Chinese to spend like Americans? An individual's level of risk aversion should stay relatively constant, but how that person calibrates risk may change. The question is, are the Chinese more risk averse or do they simply face greater uncertainty?*
>
> *... Arguably, Asian and white Americans face the same economy-wide risks. The evidence then suggests that Asian Americans have greater innate levels of risk aversion. One other possibility — the survey cited above does not indicate how long a particular family has lived in America. The Asian-American families in the survey could then be calibrating risk based on perceptions ingrained by immigrant relatives or by surrounding immigrant communities.*

From: Free exchange, *The Economist*, Economics (December 19th, 2008) "Are the Chinese nervous or thrifty?"

Note 4

Production decisions under uncertainty

Consider now the case of a producer who needs to decide an output level. Her firm faces uncertainty about the price she will receive. Formally, suppose that the cost of producing q units is $c(q)$, so that the profits she makes if she produces that much are the random variable $\Pi = P \cdot q - c(q)$, where p denotes the price the producer receives. As is usual in economics, we assume that the marginal cost is positive and increasing, so that $c'(q) > 0$ and $c''(q) > 0$ at all output levels.

Now, the price is not chosen by the producer (who is then said to be competitive) and is a random variable with expectation $\bar{P}$ and variance σ^2. The question we want to study first is how this risk affects the profit-maximizing output decisions of the individual.

As before, we assume that the producer has EU preferences over profits, with Bernoulli utility index u, which is assumed to be increasing. Also, we will assume that the individual is risk-averse, so we take u to be concave.

4.1. Output under certainty

Suppose that all uncertainty is removed, so that the price is fixed at $\bar{P}$. Then, the optimal output level is given simply by $\bar{q}$ such that $c'(\bar{q}) = \bar{P}$.

4.2. Output under uncertainty

Now, under uncertainty, the optimal output level q^* solves the problem

$$\max_q \mathrm{E}\left[u\left(P \cdot q - c(q)\right)\right].$$

The question is whether the assumption that the individual is risk averse implies that $q^* \leq \bar{q}$.

The first-order condition that characterizes the optimal output level is that

$$\mathrm{E}\left\{u'\left(P \cdot q^* - c(q^*)\right)\left[P - c'(q^*)\right]\right\} = 0,$$

which we can re-write as

$$\mathrm{E}\left[u'(\Pi^*)(P - \bar{P})\right] = \mathrm{E}\left[u'(\Pi^*)\right][c'(q^*) - \bar{P}]. \tag{4.1}$$

where $\Pi^* = P \cdot q^* - c(q^*)$. Assuming that $q^* > 0$, note that, since u is concave,

$$P > \bar{P} \Rightarrow \Pi^* > \mathrm{E}(\Pi^*) \Rightarrow u'(\Pi^*) \leq u'(\mathrm{E}(\Pi^*)),$$

while

$$P < \bar{P} \Rightarrow \Pi^* < \mathrm{E}(\Pi^*) \Rightarrow u'(\Pi^*) \geq u'(\mathrm{E}(\Pi^*)).$$

We can express these implications as

$$(P - \bar{P})[u'(\Pi^*) - u'(\mathrm{E}(\Pi^*))] \leq 0,$$

and then, taking expectations,

$$\mathrm{E}\left\{(P - \bar{P})\left[u'(\Pi^*) - u'(\mathrm{E}(\Pi^*))\right]\right\} \leq 0.$$

But the latter inequality implies that

$$\mathrm{E}\left[(P - \bar{P})u'(\Pi^*)\right] \leq \mathrm{E}\left[(P - \bar{P})u'(\mathrm{E}(\Pi^*))\right] = 0,$$

so, from Eq. (4.1), $\mathrm{E}[u'(\Pi^*)][c'(q^*) - \bar{P}] \leq 0$. Now, since index u is increasing, it must be that $c'(q^*) \leq \bar{P} = c'(\bar{q})$, and hence that $q^* \leq \bar{q}$, as we anticipated.

Example 4.1. Suppose that the total cost of producing q units is $c(q) = \frac{1}{2}q^2$, and that the producer is an EU maximizer, with Bernoulli utility index $u(x) = \log(\alpha + x)$, for some constant $\alpha > 0$.

a. Suppose initially that the price is known to be $\bar{P} > 0$ per unit. Find the optimal output level, $\bar{q}$, and resulting profit.

b. Suppose now that the price the individual face, P, is a random variable and can be 0 or $2\bar{P}$ with probability 0.5. Write the first-order condition that characterizes the optimal output level, q^*. Do you expect q^* to be smaller than $\bar{q}$? Why, or why not?

c. Suppose, in particular, that $\alpha = 5/8$ and $\bar{P} = 1$. Argue that $q^* = 1/2$, by showing that this value satisfies the first-order condition found in the previous part. Find $\bar{q}$, and compare it to q^*. Is your intuition of the previous part confirmed?

d. Maintaining the assumption that $\alpha = 5/8$ and $\bar{P} = 1$, find the distribution of the maximum profits, Π^*, that the individual will receive. Suppose that someone offers her a future contract, that locks her price at $\bar{P}$. Write an equation that characterizes the maximum value that she would be willing to pay for this future.

Answer: a. Under certainty, the producer will simply maximize her profit, by choosing $\bar{q}$ that solves program

$$\max_q \left\{\bar{P}q - \tfrac{1}{2}q^2\right\}.$$

The first-order conditions of this program give us immediately that $\bar{q} = \bar{P}$.

b. With this distribution of prices, she will choose q^* that solves

$$\max_q \left\{ \tfrac{1}{2}\log(\alpha - \tfrac{1}{2}q^2) + \tfrac{1}{2}\log(\alpha + 2\bar{P}q - \tfrac{1}{2}q^2) \right\}.$$

The first-order condition that characterizes Q^* is that

$$\frac{-q^*}{\alpha - \frac{1}{2}(q^*)^2} + \frac{2\bar{P} - q^*}{\alpha + 2\bar{P}q^* - \frac{1}{2}(q^*)^2} = 0. \qquad (*)$$

(Admittedly, this equation is pretty hard to solve.) Since this producer is risk averse, and since $\bar{P} = \mathrm{E}[P]$, the theory tells us that $q^* \leq \bar{q}$.

c. Under these values, Condition $(*)$ becomes

$$\frac{-q^*}{\frac{5}{8} - \frac{1}{2}(q^*)^2} + \frac{2 - q^*}{\frac{5}{8} + 2q^* - \frac{1}{2}(q^*)^2} = 0. \qquad (**)$$

If we substitute $q^* = 1/2$, the left-hand side of $(**)$ becomes

$$\frac{-\frac{1}{2}}{\frac{5}{8} - \frac{1}{2}\cdot\frac{1}{4}} + \frac{2 - \frac{1}{2}}{\frac{5}{8} + 1 - \frac{1}{2}\cdot\frac{1}{4}} = -\frac{\frac{1}{2}}{\frac{1}{2}} + \frac{\frac{3}{2}}{\frac{3}{2}} = 0.$$

Under certainty, in this case we would have $\bar{q} = 1$, so $\bar{q} > q^*$, as expected.

d. By direct substitution, maximized profits are $\pi^* = -1/8$ or $\pi^* = 7/8$ with equal probabilities. If she took the future contract, her net profit would be

$$\bar{P}q - \tfrac{1}{2}q^2 - \Gamma,$$

with probability 1. The optimal level of production would then be the same as in part 1, since subtracting a constant doesn't change the optimal level of output. From part 1, then, we have that gross profits will be $\bar{\pi} = 1/2$ with probability 1. The maximum she would be willing to pay for the future contract would be Γ such that

$$\log(\alpha + \tfrac{1}{2} - \Gamma) = \tfrac{1}{2}\log(\alpha - \tfrac{1}{8}) + \tfrac{1}{2}\log(\alpha + \tfrac{7}{8}).$$

We should expect $\Gamma > 0$.[1] □

EXERCISE 4.1. Suppose that the total cost of producing q units is $c(q) = \frac{1}{2}q^2$, and that the producer is an expected-utility maximizer, with cardinal utility index $u(x) = \sqrt{A + x}$, where $A > 0$ is very large.

a. Suppose initially that the price is known to be $\bar{P} > 0$ per unit. Find the optimal output level, $\bar{q}$, and resulting profit.

b. Suppose now that the price the individual face, P, is a random variable and can be 0 or $2\bar{P}$ with probability 0.5. Write the first-order condition that characterizes the optimal output level, q^*. Do you expect q^* to be smaller than $\bar{q}$? Why, or why not?

c. Suppose that someone offers her a future contract that locks her price at $\bar{P}$. Write an equation that characterizes the maximum value that she would be willing to pay for this future.

Exercise 4.2. Consider the problem of deciding optimal production levels under *cost* uncertainty. The price is $p > 0$ per unit, but the total cost of producing q units is $\frac{1}{2}Cq^2$, where C is a random variable with expectation $\bar{C} > 0$ and variance Σ. The producer is an EU maximizer with Bernoulli utility index $u(x) = \log(\alpha + x)$, for some constant $\alpha > 0$.

a. Suppose initially that the cost is known to be $\frac{1}{2}\bar{C}q^2 > 0$ per unit. Find the optimal output level, $\bar{q}$, and resulting profit.

b. Suppose now that C can be 0 or $2p$ with equal probability, 0.5. Write the first-order condition that characterizes the optimal output level, q^*. Do you expect q^* to be smaller than $\bar{q}$? Why, or why not?

c. Suppose, in particular, that $\alpha = 4/5$ and $p = 1$. Argue that $q^* = 4/5$, by showing that this value satisfies the first-order condition found in the previous part. Find $\bar{q}$, and compare it to q^*. Is your intuition of the previous part confirmed?

4.3. The effects of sunk costs

Suppose now that the cost function of the firm includes a sunk cost, $\bar{c}$, that we now isolate from its variable costs. The problem is, thus, re-written as $\max_Q [u\,(P \cdot Q - C(Q) - \bar{c})]$. It is easy to see that, in the absence of uncertainty, the optimal level of production is independent of $\bar{c}$: $\mathrm{d}\bar{Q}/\mathrm{d}\bar{c} = 0$. We now want to show that this independence fails when there is price risk. In particular, we want to observe that if the firm exhibits decreasing coefficient of absolute risk aversion, then $\mathrm{d}Q^*/\mathrm{d}\bar{c} < 0$.

Note that, for a given $\bar{c}$, the first-order condition for optimality of Q^* is

$$\mathrm{E}\{u'(\pi^*)[P - C'(Q^*)]\} = 0, \tag{**}$$

where $\pi^* = P \cdot Q^* - C(Q^*) - \bar{c}$. If we differentiate this expression totally, we can solve for

$$\frac{\mathrm{d}Q^*}{\mathrm{d}\bar{c}} = \frac{\mathrm{E}\{u''(\pi^*)[P - C'(Q^*)]\}}{\mathrm{E}\{u''(\pi^*)[P - C'(Q^*)]^2\} - \mathrm{E}\,[u'(\pi^*)]C''(Q^*)}.$$

Since the Bernoulli utility index is increasing and concave, and the cost function is strictly convex, the denominator of this expression is a negative number. For our claim, then, it suffices that we show that the numerator is a positive number.

Now, define $\tilde{P} = C'(Q^*)$, $\tilde{\pi} = \tilde{P} \cdot Q^* - C(Q^*) - \bar{c}$, and $\tilde{A} = A(\tilde{\pi})$. Then, assuming that $Q^* > 0$,

$$P > \tilde{P} \Rightarrow \pi^* > \tilde{\pi} \Rightarrow A(\pi^*) < \tilde{A},$$

while

$$P < \tilde{P} \Rightarrow \pi^* < \tilde{\pi} \Rightarrow A(\pi^*) > \tilde{A}.$$

We can express these implications as $(P - \tilde{P})[A(\pi^*) - \tilde{A}] < 0$, which means that

$$-[P - C'(Q^*)][A(\pi^*) - \tilde{A}]u'(\pi^*) > 0,$$

or, equivalently,

$$-A(\pi^*)u'(\pi^*)[P - C'(Q^*)] > -\tilde{A} \cdot u'(\pi^*)[P - C'(Q^*)].$$

If we take expectations of the last expression,

$$\mathrm{E}\{u''(\pi^*)[P - C'(Q^*)]\} = -\mathrm{E}\{A(\pi^*)u'(\pi^*)[P - C'(Q^*)]\} > -\tilde{A} \cdot \mathrm{E}\{u'(\pi^*)[P - C'(Q^*)]\} = 0,$$

using $(**)$, as we needed.

Note 5

Stochastic dominance

WE NOW WANT TO STUDY "MONOTONICITY" PROPERTIES FOR LOTTERIES. We need a new framework for this, as it would be a mistake to pretend that we can order lotteries using the standard 'greater-than' relation. For simplicity, let us consider the case of lotteries that pay in nonnegative units of some numéraire (money), so that we represent them by the probabilities they assign to any nonnegative number x: a lottery will be a c.d.f. $F : \mathbb{R}_+ \to [0, 1]$.[1]

5.1. First order

WE SAY THAT LOTTERY F IS *as large as* lottery $\tilde{F}$ *in the first-order stochastic sense* if $F(x) \leq \tilde{F}(x)$ for every every possible payoff level x. F is said to *first-order stochastically dominate* $\tilde{F}$ if it is as large, and the above inequality is strict at some payoff level. For simplicity, we will write that $F \geq_1 \tilde{F}$ if F is as large as $\tilde{F}$ in the sense of first-order stochastic dominance, and that $F >_1 \tilde{F}$ if F dominates $\tilde{F}$ in the same sense.

This is an intuitive concept: by definition, both F and $\tilde{F}$ are nondecreasing, so saying that $F \geq_1 \tilde{F}$ means that, at any point, F leaves at least as much probability to be allocated to higher payoffs as $\tilde{F}$.

EXAMPLE 5.1. Consider an EU maximizer, who is presented with the following two lotteries:

1. in lottery F_1, she gets an amount x_1 with probability $1/3$, $2x_1$ with probability $1/3$ or $3x_1$ with probability $1/3$;
2. in lottery F_2, she gets an amount x_2 with probability $1/3$, $2x_2$ with probability $1/3$ or $3x_2$ with probability $1/3$.

Argue that if $x_2 > x_1$, then $F_2 >_1 F_1$. What lottery does this decision-maker prefer? Does this answer depend on the person's specific Bernoulli index?

Answer: The cdf's of the two lotteries are as in Fig. 5.1. From there, it is immediate that $L_2 >_1 L_1$. As long as the individual prefers more wealth to less, namely that $u' > 0$, she prefers L_2. This is independent of u, as long as we restrict attention to increasing indices. □

[1] If a lottery F has a density function, we will denote this function by $f : \mathbb{R}_+ \to \mathbb{R}_+$.

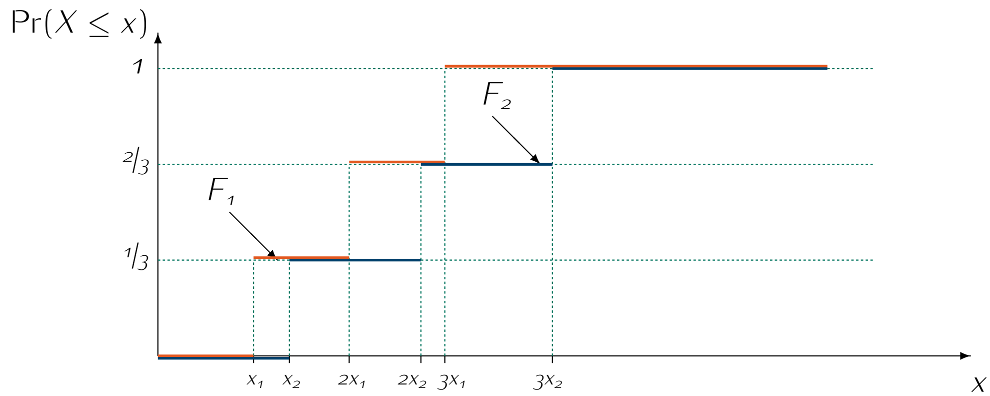

Figure 5.1: The lotteries of Example 5.1.

For any decision maker with EU preferences, if she prefers more wealth to less, then she also ranks lotteries according to first-order stochastic dominance. The following proposition formalizes this observation.

Proposition 5.1. *Given two lotteries, F and $\tilde{F}$:*

1. *if $F >_1 \tilde{F}$, then $\mathrm{E}_F(u) > \mathrm{E}_{\tilde{F}}(u)$ for any increasing utility index u; and*
2. *if $F \neq \tilde{F}$ and it is not true that $F >_1 \tilde{F}$, then for some increasing utility index u one has that $\mathrm{E}_F(u) < \mathrm{E}_{\tilde{F}}(u)$.*

Proof. For simplicity of presentation, we consider here only the case of a discrete random variable giving positive probability only to some integer numbers, and defer the more general case to an appendix. That is, let us assume that F and $\tilde{F}$ give positive probability only to payoffs in the set $\{0, 1, \dots, \bar{x}\}$, for some positive integer $\bar{x}$. For the first statement, by definition,

$$\mathrm{E}_F(u) - \mathrm{E}_{\tilde{F}}(u) = \sum_{x=0}^{\bar{x}} u(x)[F(x) - F(x-1)] - \sum_{x=0}^{\bar{x}} u(x)[\tilde{F}(x) - \tilde{F}(x-1)],$$

with $F(-1) = \tilde{F}(-1) = 0$. Rearranging terms,[2] the right-hand side of this expression is

$$u(\bar{x})[F(\bar{x}) - \tilde{F}(\bar{x})] - u(0)[F(-1) - \tilde{F}(-1)] + \sum_{x=1}^{\bar{x}} [u(x) - u(x-1)][\tilde{F}(x-1) - F(x-1)].$$

Since $F(\bar{x}) = \tilde{F}(\bar{x})$ and $F(-1) = \tilde{F}(-1)$, this expression becomes, simply,

$$\sum_{x=1}^{\bar{x}} [u(x) - u(x-1)][\tilde{F}(x-1) - F(x-1)].$$

[2] Note that for any functions $\upsilon : \{1, \dots, I\} \to \mathbb{R}$ and $\Phi : \{0, 1, \dots I\} \to \mathbb{R}$, if one defines $\Delta\Phi(i) = \Phi(i) - \Phi(i-1)$ for each $i = 1, \dots, I$, then

$$\sum_{i=1}^{I} \upsilon(i)\Delta\Phi(i) = \upsilon(I)\Phi(I) - \upsilon(1)\Phi(0) - \sum_{i=2}^{I} [\upsilon(i) - \upsilon(i-1)]\Phi(i-1).$$

This is the analogous expression, for sums, of integration by parts.

And since u is increasing and $F >_1 \tilde{F}$, this expression is a strictly positive number.

For the second statement, since $F \neq \tilde{F}$ and it is not true that $F >_1 \tilde{F}$, it must be that for some income level x^*, it is true that $F(x^*) > \tilde{F}(x^*)$. Fix one such x^*, and construct the following (nondecreasing) function: $v(x) = 0$ for any $x \leq x^*$, and $v(x) = 1$ for all $x > x^*$. Then, for any cummulative distribution function, $\widehat{F}$, by definition, $\mathrm{E}_{\widehat{F}}(v) = 1 - \widehat{F}(x^*)$, so it follows that $\mathrm{E}_{\tilde{F}}(v) > \mathrm{E}_F(v)$. Since this inequality is strict, we can modify v to construct an increasing index u for which the inequality holds too. □

The following examples provide useful properties of the first-order stochastic dominance relations.

Example 5.2. Argue that relations $\geq_1$ and $>_1$ are transitive.

Answer: Fix lotteries F, $\tilde{F}$ and $\widehat{F}$ such that $F \geq_1 \tilde{F}$ and $\tilde{F} \geq_1 \widehat{F}$. Then, for any value of x, $F(x) \leq \tilde{F}(x)$ and $\tilde{F}(x) \leq \widehat{F}(x)$. This implies that $F(x) \leq \widehat{F}(x)$ for all values of x, and hence that $F \geq_1 \widehat{F}$. For $>_1$, the argument is similar, observing that if $F >_1 \tilde{F}$ abd $\tilde{F} >_1 \widehat{F}$, then $F \neq \widehat{F}$. □

Abusing notation slightly, we will say that a random variable dominates another, in the senses described here, if the distribution of the former dominates the distribution of the latter in the given sense. With this language, we can more easily express the following properties.

Example 5.3 (Difficult!). *Suppose that X, Y and Z are independent random variables. Argue that*

a. if $Y \geq_1 Z$, then $\alpha X + (1-\alpha)Y \geq_1 \alpha X + (1-\alpha)Z$ for any $\alpha \in [0,1]$; and

b. if $Y >_1 Z$, then $\alpha X + (1-\alpha)Y >_1 \alpha X + (1-\alpha)Z$ for any $\alpha \in [0,1)$.

Answer: For the first property, let F, G and H be, respectively, the distributions of X, Y and Z. Note that

$$\begin{aligned}\Pr[\alpha X + (1-\alpha)Y \leq v] &= \int_{-\infty}^{\infty} \Pr\left(Y \leq \frac{v-\alpha x}{1-\alpha} \,\middle|\, X = x\right) \mathrm{d}F(x) \\ &= \int_{-\infty}^{\infty} \Pr\left(Y \leq \frac{v-\alpha x}{1-\alpha}\right) \mathrm{d}F(x) \\ &= \int_{-\infty}^{\infty} G\left(\frac{v-\alpha x}{1-\alpha}\right) \mathrm{d}F(x),\end{aligned}$$

where we used the assumption that X and Y are independent in the second equality. By an identical argument,

$$\Pr[\alpha X + (1-\alpha)Z \leq v] = \int_{-\infty}^{\infty} H\left(\frac{v-\alpha x}{1-\alpha}\right) \mathrm{d}F(x).$$

Since $Y \geq_1 Z$, we know that $H(x) \geq G(x)$ at all x, so

$$\int_{-\infty}^{\infty} G\left(\frac{v-\alpha x}{1-\alpha}\right) \mathrm{d}F(x) \leq \int_{-\infty}^{\infty} H\left(\frac{v-\alpha x}{1-\alpha}\right) \mathrm{d}F(x),$$

which means that

$$\Pr[\alpha X + (1-\alpha)Y \leq v] \leq \Pr[\alpha X + (1-\alpha)Z \leq v].$$

Since this is true for all v, the result follows. The second property is argued in a similar way. □

The following result is very natural:

EXERCISE 5.1. Consider two lotteries F and $\tilde{F}$. Argue that if $F >_1 \tilde{F}$, then $\mathrm{E}_F(X) > \mathrm{E}_{\tilde{F}}(X)$.

5.2. Second order

FIRST-ORDER STOCHASTIC DOMINANCE can sometimes be too strong as a concept of dominance for lotteries. A second, weaker concept is given by the following definition. A lottery F is *as large as* lottery $\tilde{F}$ *in the second-order stochastic sense* if

$$\int_0^x F(s)\,\mathrm{d}s \leq \int_0^x \tilde{F}(s)\,\mathrm{d}s$$

for every every possible payoff level x. F is said to *second-order stochastically dominate* $\tilde{F}$ if it is as large, and the above inequality is strict at some payoff level. As before, we will use $F \geq_2 \tilde{F}$ and $F >_2 \tilde{F}$ to denote stochastic dominance in the second-order sense.

It is immediate that first-order stochastic dominance implies second-order stochastic dominance, but the converse is not true. What the second concept captures is the difference in the "speeds" at which different lotteries accrue probability over 'low' payoffs. The following proposition illustrates the importance of this concept; the proposition is stated without some technical assumptions, which are deferred to the proof given in the appendix.

PROPOSITION 5.2. *Given two lotteries F and $\tilde{F}$:*

1. *if $F >_2 \tilde{F}$, then $\mathrm{E}_F(u) > \mathrm{E}_{\tilde{F}}(u)$ for any increasing and strictly concave utility index u; and*
2. *if $F \neq \tilde{F}$ and it is not true that $F >_2 \tilde{F}$, then for some increasing and strictly concave utility index u one has that $\mathrm{E}_F(u) < \mathrm{E}_{\tilde{F}}(u)$.*

Proof. As before, we can illustrate this result in the discrete case considered in Proposition 5.1. In this case, since the domain of u is not convex, we replace the assumption of concavity of the index by the condition that $[u(x) - u(x-1)] - [u(x-1) - u(x-2)] < 0$ for every $x = 2, \ldots, \bar{x}$. This is the discrete analogous of the requirement that the second derivative of the function be negative. In this setting, suppose that $F >_2 \tilde{F}$, and recall from the proof of Proposition 5.1 that

$$\mathrm{E}_F(u) - \mathrm{E}_{\tilde{F}}(u) = \sum_{x=1}^{\bar{x}+1} [u(x) - u(x-1)][\tilde{F}(x-1) - F(x-1)].$$

Rewriting this expression, as before, we get that its right-hand side equals[3]

$$-\sum_{x=2}^{\bar{x}+1} \left\{ [u(x) - 2u(x-1) + u(x-2)] \sum_{s=0}^{x-2} [\tilde{F}(s) - F(s)] \right\}.$$

[3] The term $[u(\bar{x}+1) - u(\bar{x})][\tilde{F}(\bar{x}) - F(\bar{x})] - [u(1) - u(0)][\tilde{F}(0) - F(0)]$, which also appears in the expression, disappears since $\tilde{F}(\bar{x}) = F(\bar{x})$ and $\tilde{F}(0) = F(0)$.

Since, by assumption, each $u(x) - 2u(x-1) + u(x-2) < 0$ and each $\sum_{s=0}^{x-2}(\tilde{F}(s) - F(s)) \leq 0$, with strict inequality somewhere, it follows that $\mathrm{E}_F(u) > \mathrm{E}_{\tilde{F}}(u)$. $\square$

Exercise 5.2. An oil company, *Pritish Betroleum*, henceforth PB, has been found guilty of causing a major environmental disaster, and the court has decided that it has to pay an amount X in punitive damages. The exact amount of X will be determined after an assessment of the damages, which will be undertaken by one of two environmental laboratories, a or b. Under laboratory i, the probability distribution for X is F_i.[4]

Assume that PB is an EU maximizer and that its Bernoulli utility index is increasing in the company's wealth.

a. Suppose that $F_a \geq_1 F_b$ and that the court allows PB to choose which laboratory to use for the assessment of damages. Which one should they choose?

b. Suppose now that F_a is the uniform distribution over some interval $[x_*, x^*]$, while F_b is the uniform distribution over $[x_* - 1, x^* + 1]$. Argue that $F_a >_2 F_b$.

c. Under the same assumptions as in part 2, suppose that the court allows PB to choose which laboratory to use for the assessment of damages. Which one should they choose?

Exercise 5.3. Let $U[a,b]$ denote the uniform distribution over the interval $[a,b]$. In the context of lotteries that pay in non-negative amounts of wealth, let $L_1 = U[1,2]$, $L_2 = U[1,3]$, $L_3 = U[2,4]$, $L_4 = U[2.5, 3.5]$, $L_5 = U[3,3]$, and $L_6 = U[3,4]$.

a. Draw the distributions of the lotteries, and use their graphs to compare them, whenever possible, according to first- and second-order stochastic dominance.

b. Suppose that $\succsim^A$ are the preferences of a strictly risk-averse individual. Rank the lotteries according to this individual's preferences.

c. Suppose that $\succsim^N$ are the preferences of a risk-neutral individual. Rank the lotteries according to this individual's preferences.

d. Suppose that $\succsim^L$ are the preferences of a risk-loving individual. Rank the lotteries according to this individual's preferences.

The nice properties of first-order stochastic dominance also hold in this case:

Exercise 5.4. Argue that relations $\geq_2$ and $>_2$ are transitive.

Exercise 5.5. Suppose that X, Y and Z are independent random variables. Argue that:

a. if $Y \geq_2 Z$, then $\alpha X + (1-\alpha)Y \geq_2 \alpha X + (1-\alpha)Z$ for any $\alpha \in [0,1]$; and

b. if $Y >_2 Z$, then $\alpha X + (1-\alpha)Y >_2 \alpha X + (1-\alpha)Z$ for any $\alpha \in [0,1)$.

In order to understand what second-order dominance captures, the following result is useful. Intuitively, under the premises of the proposition, lottery $\tilde{F}$ takes probability mass from the "center" of the distribution (i.e. near the mean) and allocates it to both of its extremes; for a risk-averse decision-maker, this makes the lottery worse. Under those premises, $\tilde{F}$ is said to be a *mean-preserving spread* of F.

[4] You can assume that F_a and F_b are continuous.

Example 5.4. Consider two continuous lotteries F and $\tilde{F}$, with densities f and $\tilde{f}$, that assign all the probability mass over the interval $[0, \bar{x}]$, and suppose that $\mathrm{E}_F(X) = \mathrm{E}_{\tilde{F}}(X)$. Argue that if $F >_2 \tilde{F}$, then $\mathrm{V}_F(X) < \mathrm{V}_{\tilde{F}}(X)$.

Answer: Integrating by parts,

$$\mathrm{V}_F[X] - \mathrm{V}_{\tilde{F}}[X] = x^2[F(x) - \tilde{F}(x)]\Big|_0^{\bar{x}} - \int_0^{\bar{x}} 2x[F(x) - \tilde{F}(x)]\,\mathrm{d}x.$$

The first summand on the right-hand side of the expression is zero, since $\mathrm{E}_F[X] = \mathrm{E}_{\tilde{F}}[X]$,[5] Integration by parts of the second term gives

$$-2\left(x\int_0^x [F(s) - \tilde{F}(s)]\,\mathrm{d}s\right)\Big|_0^{\bar{x}} + 2\int_0^{\bar{x}}\int_0^x [F(s) - \tilde{F}(s)]\,\mathrm{d}s\,\mathrm{d}x.$$

The first term in this latter expression is simply $-2\bar{x}\int_0^{\bar{x}}[F(s) - \tilde{F}(s)]\,\mathrm{d}s$ which, again, is null since both lotteries have the same mean. The second term is is negative, since $F >_2 \tilde{F}$. □

Appendix

We now give a more formal presentation of the results in stochastic dominance. We start by giving a version of Proposition 5.1 for continuous lotteries, and its proof.

Proposition. *Consider two lotteries, F and $\tilde{F}$, with densities f and $\tilde{f}$.*

1. *If $F >_1 \tilde{F}$, then, for any increasing and bounded utility index $u \in \mathbf{C}^1$, $\mathrm{E}_F[u(X)] > \mathrm{E}_{\tilde{F}}[u(X)]$.*
2. *Conversely, if $F \neq \tilde{F}$ and it is not true that $F >_1 \tilde{F}$, then for some increasing and bounded utility index $u \in \mathbf{C}^1$ one has that $\mathrm{E}_F[u(X)] < \mathrm{E}_{\tilde{F}}[u(X)]$.*

Proof. For the first statement, integrating by parts and since u is continuously differentiable,

$$\mathrm{E}_F[u(X)] - \mathrm{E}_{\tilde{F}}[u(x)] = \int_0^\infty u(x)[f(x) - \tilde{f}(x)]\,\mathrm{d}x = u(x)[F(x) - \tilde{F}(x)]\Big|_0^\infty - \int_0^\infty u'(x)[F(x) - \tilde{F}(x)]\,\mathrm{d}x.$$

Since $F(0) = \tilde{F}(0) = 0$ (remember that these c.d.f. have density) and $\lim_{x\to\infty} F(x) = \lim_{x\to\infty} \tilde{F}(x) = 1$, and since u is bounded, it follows that $u(x)[F(x) - \tilde{F}(x)]|_0^\infty = 0$. Since $u' > 0$ and $F >_1 \tilde{F}$, we have that[6] $\int_0^\infty u'(x)[F(x) - \tilde{F}(x)]\,\mathrm{d}x < 0$, so $E_F[u(X)] - \mathrm{E}\,\tilde{F}[u(x)] > 0$.

For the second statement, we shall consider two c.d.f. that "cross," so that none of them dominates the other: fix x^* such that $F(x) \geq \tilde{F}(x)$ for all $x \leq x^*$, with strict inequality somewhere, and $F(x) \leq \tilde{F}(x)$ for all $x \geq x^*$. Define the index υ_p, for each positive real number p, by

$$\upsilon_p(x) = \begin{cases} p\exp(\frac{x-x^*}{p}), & \text{if } x \leq x^*; \\ p + \frac{1}{p}\{1 - \exp[-p(x - x^*)]\}, & \text{if } x > x^*. \end{cases}$$

[6] Remember that any c.d.f. is right-continuous.

By construction, this function is differentiable and monotone, with

$$\upsilon_p'(x) = \begin{cases} \exp(\frac{x-x^*}{p}), & \text{if } x \leq x^*; \\ \exp(-\frac{x-x^*}{p}), & \text{if } x \geq x^*. \end{cases}$$

The function is also bounded, with $\lim_{x\to\infty} \upsilon_p(x) = p$. Now, recalling the equation above, we have that $\mathrm{E}_F[u(X)] - \mathrm{E}\,\tilde{F}[u(x)] = -\int_0^\infty u'(x)[F(x) - \tilde{F}(x)]\,\mathrm{d}x$, and the right-hand side of this expression is, by direct substitution,

$$-\int_0^{x^*} \exp\left(\frac{x-x^*}{p}\right)[F(x)-\tilde{F}(x)]\,\mathrm{d}x - \int_{x^*}^\infty \exp\left(-\frac{x-x^*}{p}\right)[F(x)-\tilde{F}(x)]\,\mathrm{d}x,$$

an expression that is ambiguous, in general, by our assumptions. However, since for and $x \leq x^*$ the term $\exp[(x-x^*)/p]$ is increasing in p, while for any $x \leq x^*$ the term $\exp[-(x-x^*)/p]$ is decreasing in p, it follows that for p large enough the first term dominates and the whole expression is negative. □

The result for second-order dominance requires some technical assumptions too:

Proposition. *Let F and $\tilde{F}$ be two continuous lotteries, with densities f and $\tilde{f}$. Suppose that both lotteries have finite mean.*

1. *If $F >_2 \tilde{F}$, then for any increasing, bounded and strictly concave utility index $u \in \mathbf{C}^1$, we have that* $\mathrm{E}_F[u(X)] > \mathrm{E}_{\tilde{F}}[u(X)]$.
2. *Conversely, if $F \neq \tilde{F}$ and it is not true that $F >_2 \tilde{F}$, then for some increasing, bounded and strictly concave utility index $u \in \mathbf{C}^1$ one has that* $\mathrm{E}_F[u(X)] < \mathrm{E}_{\tilde{F}}[u(X)]$.

Proof. For the first statement, let us recall again that $\mathrm{E}_F[u(X)] - \mathrm{E}_{\tilde{F}}[u(x)] = -\int_0^\infty u'(x)[F(x)-\tilde{F}(x)]\,\mathrm{d}x$. By integration by parts again, the right-hand side of the expression is

$$-\left(u'(x)\int_0^x [F(s)-\tilde{F}(s)]\,\mathrm{d}s\right)\Big|_0^\infty + \int_0^\infty \left(u''(x)\int_0^x [F(s)-\tilde{F}(s)]\,\mathrm{d}s\right)\mathrm{d}x.$$

By concavity and boundedness, $\lim_{x\to\infty} u'(x) = 0$, while $\int_0^\infty [F(s)-\tilde{F}(s)]\,\mathrm{d}s \in \mathbb{R}$, since both lotteries have mean, so the first term on this expression vanishes. The second term is positive, since $u'' < 0$ and $F >_2 \tilde{F}$.

The proof of the second statement is similar to its analogous in the extension of Proposition 5.1 above, using the expression we just obtained, and considering the utility indices

$$\upsilon_p(x) = -\frac{1}{p^2}\exp[-p(x-x^*)],$$

for $p > 0$. Details are omitted. □

Note 6

Portfolio theory

Consider a strictly risk-averse individual who derives utility from future consumption. She has a present wealth w that she is to allocate between two assets: a riskless asset, with return $\rho > 0$; and an asset whose return is $R + \rho$, for a random variable R that follows a distribution F, with expectation $\bar{R} > 0$ and density f.

Denoting by k the amount the individual invests in the risky asset, her problem is

$$\max_{k \in [0,w]} \mathrm{E}\left[u\left((w-k)\rho + k(R+\rho)\right)\right] = \max_{k \in [0,w]} \mathrm{E}\left[u(w\rho + kR)\right].$$

The first-order condition for an interior solution k^* is that $\mathrm{E}\left[u'(w\rho + k^*R)R\right] = 0$, which is necessary and sufficient given that the individual is strictly risk averse.[1]

The purpose of this section is to study the response of the optimal level of investment in the risky asset to changes in the distribution of its return. In general, one can capture perturbations to this distribution that affect it in the sense of first- or second-order stochastic dominance. Here, for simplicity, we will concentrate in some specific types of perturbations.

6.1. First-order, simple perturbations

Suppose now that the return of the risky asset is $R + \rho + \vartheta$, where $\vartheta > -\mathrm{E}(R)$ is a constant, and R is distributed as before.[2] Note that higher values of ϑ amount to first-order stochastic improvements in the return of this asset.

The portfolio problem now is

$$\max_{k \in [0,w]} \mathrm{E}\left[u\left(w\rho + k(R+\vartheta)\right)\right].$$

We denote the optimal investment in the risky asset by $k(\vartheta)$, and restrict attention to values of ϑ (and other parameters) for which this solution is interior.

Proposition 6.1. *If the investor has a decreasing coefficient of absolute risk aversion, the optimal investment in the risky asset is increasing in ϑ.*

[1] The fact that $\mathrm{E}(R) > 0$ guarantees that $k^* > 0$. We assume also that $k^* < w$, for simplicity.

[2] We also assume that $F(-\vartheta) > 0$, for interiority of the solution.

In order to argue the proposition, we need to show that $k'(\vartheta) > 0$. By the first-order condition,

$$\mathrm{E}\left[u'\left(w\rho + k(\vartheta)(R+\vartheta)\right)(R+\vartheta)\right] = 0.$$

Differentiating with respect to ϑ gives[3]

$$k'(\vartheta) = -\frac{\mathrm{E}\left[u''(\cdot)k(\vartheta)(R+\vartheta) + u'(\cdot)\right]}{\mathrm{E}\left[u''(\cdot)(R+\vartheta)^2\right]}.$$

Since the denominator of this expression is negative, it suffices for our purposes to prove that its numerator is positive.

Using the definition of the coefficient of absolute risk aversion, A, we can rewrite the numerator as

$$\mathrm{E}\left[u'(\cdot)\right] - \mathrm{E}\left[u'(\cdot)A(\cdot)k(\vartheta)(R+\vartheta)\right].$$

The first of these terms is positive and $k(\vartheta) > 0$, so it suffices to show that

$$\mathrm{E}\left[u'(\cdot)A(\cdot)(R+\vartheta)\right] \le 0.$$

Note that we can write this expectation as[4]

$$\int_{-\infty}^{-\vartheta} u'(\cdot)A(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r + \int_{-\vartheta}^{\infty} u'(\cdot)A(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r.$$

Since A is positive and decreasing, it is immediate that

$$\int_{-\infty}^{-\vartheta} u'(\cdot)A(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r \le A(w\rho)\int_{-\infty}^{-\vartheta} u'(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r.$$

while

$$\int_{-\vartheta}^{\infty} u'(\cdot)A(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r \le A(w\rho)\int_{-\vartheta}^{\infty} u'(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r$$

The sum of these two inequalities yields

$$\begin{aligned}\mathrm{E}\left[u'(\cdot)A(\cdot)(R+\vartheta)\right] &\le A(w\rho)\int_{-\infty}^{\infty} u'(\cdot)(r+\vartheta)f(r)\,\mathrm{d}r \\ &= A(w\rho)\mathrm{E}\left[u'(\cdot)(R+\vartheta)\right] \\ &= 0,\end{aligned}$$

where the last equality comes from the first-order condition of the utility maximization problem.

Example 6.1. In the context of this section, suppose that $u(x) = \ln(x)$, $w > 0$, $\rho = 1$, and that R can be either -1 or 2, with equal probability. Restricting attention to $\vartheta \in (-1/2, 1)$, find $k(\vartheta)$ and argue that $k'(0) > 0$.

[3] In the following expressions, $\cdot$ is used to denote the argument of the different functions, which is the random variable $w\rho + k(\vartheta)(R+\vartheta)$.

[4] In the following, integrals, the argument of the functions u' and A is $w\rho + k(\vartheta)(r+\vartheta)$, but we substitute it for $\cdot$, for simplicity of the expressions.

Answer: The portfolio problem is

$$\max_k \left\{ \tfrac{1}{2}\ln[w+(\vartheta+2)k] + \tfrac{1}{2}\ln[w+(\vartheta-1)k] \right\}.$$

Its first-order condition for the optimum is that

$$\frac{\vartheta+2}{w+(\vartheta+2)k} + \frac{\vartheta-1}{w+(\vartheta-1)k} = 0.$$

By direct computation,

$$k(\vartheta) = \frac{(1+2\vartheta)w}{2(2+\vartheta)(1-\vartheta)}.$$

This function is increasing, as its derivative is proportional to

$$5+2\vartheta+2\vartheta^2 > 0.$$

The result is intuitive, as a higher premium implies a first-order stochastic improvement in the project and the agent has decreasing absolute risk aversion. □

The following exercise shows that the we cannot dispense with the assumption that the coefficient of absolute risk aversion of the investor is decreasing.

Exercise 6.1. In the context of this section, suppose that the investor has quadratic preferences, so that her marginal utility is a linear function $u'(x) = S + \mu x$, where $\mu < 0$ and $S > -\mu$ are constant. In this case, a first order improvement in the return of the risky asset need not increase the investor's optimal level of investment in it: function k need not be increasing in ϑ.

Argue the following steps, which together yield a proof of this claim. For simplicity, suppose that $\rho = w = 1$.

a. The optimal demand is

$$k(\vartheta) = -\frac{(S+\mu)(\mathrm{E}[R]+\vartheta)}{\mu \mathrm{E}[(R+\vartheta)^2]}.$$

b. $k'(\vartheta)$ equals

$$\frac{\mu(S+\mu)\{2(\mathrm{E}[R]+\vartheta)^2 - E[(R+\vartheta)^2]\}}{\mu^2 E[(R+\vartheta)^2]^2}.$$

c. $k'(0) > 0$ only if $\mathrm{V}(R) > \mathrm{E}[R]^2$.

6.2. Second-order perturbations: mean-preserving spreads

Maintaining the notation and assumptions introduced so far, let us now assume that the return of the risky asset is $\rho + \alpha R + (1-\alpha)\mathrm{E}(R)$, for a constant number $\alpha \in [0,1]$. As before, we denote the investor's optimal level of investment in the risky asset as

$$k(\alpha) = \operatorname*{argmax}_{k\in[0,w]} \mathrm{E}\left\{u\left(w\rho + k[\alpha R + (1-\alpha)\mathrm{E}(R)]\right)\right\},$$

which we assume to be interior.

Note that a higher value of α represents a mean-preserving spread in the return of the risky asset, so that we can use a decrease in α to model a (simple) second-order stochastic improvement in this return.

EXERCISE 6.2. Suppose that the investor has quadratic preferences, so that her u'' is a constant, $\mu < 0$. In this case, a second-order improvement in the return of the risky asset increases the investor's optimal level of investment in it: function k is decreasing in α.

Argue each of the following steps, which together yield a proof of the statement.

a. The first-order condition of the investor is that

$$\mathrm{E}\{u'(w\rho + k(\alpha)[\alpha R + (1-\alpha)\mathrm{E}(R)])[\alpha R + (1-\alpha)\mathrm{E}(R)]\} = 0.$$

b. $k'(\alpha)$ equals

$$-\frac{\mathrm{E}\{u''(\cdot)k(\alpha)[R - \mathrm{E}(R)][\alpha R + (1-\alpha)\mathrm{E}(R)] + u'(\cdot)[R - \mathrm{E}(R)]\}}{\mathrm{E}\{u''(\cdot)[\alpha R + (1-\alpha)\mathrm{E}(R)]^2\}}.$$

c. The denominator of the previous expression is negative.

d. The numerator of the expression can be rewritten as

$$\alpha k(\alpha)\mu \mathrm{V}(R) + \mathrm{Cov}[u'(\cdot), R]. \qquad (*)$$

e. The first summand of this last expression is negative, and the second non-positive.

f. Expression $(*)$ is negative and $k'(\cdot) < 0$.

6.3. Diversification and stochastic dominance

CONSIDER NOW THE PROBLEM of an investor who has to choose between risky assets. Let X_1 and X_2 be i.i.d. non-degenerate random variables.[5] Let the common distribution of these variables be F.

Suppose first that the investor has access to only these two assets, and has to choose the proportion α that she invests in X_1, while the remainder $(1-\alpha)$ is invested in X_2. Our first result is that for *any* risk averse individual, any degree of diversification is strictly better than no diversification at all.

PROPOSITION 6.2 (Samuelson). *For any $\alpha \in (0, 1)$, portfolio $\alpha X_1 + (1-\alpha)X_2$ second-order stochastically dominates assets X_1 and X_2.*

[5] That they are independent means that, for any (x_1, x_2), $\Pr(X_1 \leq x_1, X_2 \leq x_2) = \Pr(X_1 \leq x_1) \times \Pr(X_2 \leq x_2)$. They are identically distributed in the sense that for all x, $\Pr(X_1 \leq x) = \Pr(X_2 \leq x)$. They are non degenerate in the sense that for some x, $0 < \Pr(X_1 \leq x) < 1$.

To see why this proposition is true, just consider a strictly risk-averse individual whose Bernoulli index is u. Note that

$$\begin{aligned}\mathrm{E}\,[u\,(\alpha X_1 + (1-\alpha)X_2)] &= \int\!\!\int u[\alpha x_1 + (1-\alpha)x_2]\,\mathrm{d}F(x_1)\,\mathrm{d}F(x_2)\\ &> \int\!\!\int [\alpha u(x_1) + (1-\alpha)u(x_2)]\,\mathrm{d}F(x_1)\,\mathrm{d}F(x_2),\end{aligned}$$

where we used the fact that u is strictly concave and the random variables are independent and non-degenerate. Now,

$$\begin{aligned}\int\!\!\int [\alpha u(x_1) + (1-\alpha)u(x_2)]\,\mathrm{d}F(x_1)\,\mathrm{d}F(x_2) &= \alpha \int u(x_1)\,\mathrm{d}F(x_1) + (1-\alpha)\int u(x_2)\,\mathrm{d}F(x_2)\\ &= \int u(x_1)\,\mathrm{d}F(x_1)\\ &= \mathrm{E}\,[u(X_1)].\end{aligned}$$

Since, then, the inequality

$$\mathrm{E}\,[u\,(\alpha X_1 + (1-\alpha)X_2)] > \mathrm{E}\,[u(X_1)]$$

holds true for all strictly concave indices u, the claim follows from Proposition 5.2.

A more surprising result is the following. Let Y be another asset that is independent from X_1 and X_2 but second-order stochastically dominates X_1. Suppose that perfect diversification between X_1 and X_2 second order stochastically dominates Y. The next result shows that perfect diversification between X_1 and Y is better for *any* risk-averse individual than investing in Y (or in X_1) alone.

Proposition 6.3 (Hadar and Russell). *If $Y >_2 X_1$ but $\frac{1}{2}X_1 + \frac{1}{2}X_2 >_2 Y$, then*

$$\tfrac{1}{2}X_1 + \tfrac{1}{2}Y >_2 Y.$$

This follows immediately from the transitivity and linearity properties of second-order stochastic dominance that we saw in Exercises 5.4 and 5.5:

$$\tfrac{1}{2}X_1 + \tfrac{1}{2}Y >_2 \tfrac{1}{2}X_1 + \tfrac{1}{2}X_2 >_2 Y.$$

Exercise 6.3. Argue that, if X_1, X_2, X_3 and X_4 are i.i.d. assets, then a portfolio containing equal amounts of the four assets second-order stochastically dominates each of them:

$$\tfrac{1}{4}(X_1 + X_2 + X_3 + X_4) \geq_2 X_1.$$

Exercise 6.4. A betting house in London is offering bets on the winners of two tennis matches. Match 1 is between M (for Maria) and S (for Serena); match 2 is between N (for Novak) and R (for Rafa). It is clear that in both cases the players are evenly matched, and that the outcome of one match is independent of the outcome of the other.

The betting house offers three bets. *B1* pays \$2 if M wins match 1, and \$0 if S wins it; *B2* pays \$2 if N wins match 2, and \$0 is R does. Each of these bets has a cost of \$1. Additionally, the house offers a “Combo” bet, C, which consists of half a ticket of each of the two other bets, also at a cost of \$1.

a. Construct and draw the c.d.f. of the net payoff of $B1$ and of Combo bet C.

b. Suppose that you are hired to advise on the betting decision of an individual (who enjoys wealth). She has four options: to buy one ticket for $B1$, one ticket for $B2$, one ticket for C, or to do nothing. What would you advise her to do? Explicitly write how you would rank the bets in your advice.

READING. The following two excerpts are from articles in *The Economist*. Discuss them critically. First,

> *There can be fashions in investing as well as in the arts. Over the past 25 years many university endowments have moved over to the "Yale model", an investment strategy adopted by the Connecticut-based university in the 1980s. Under the leadership of David Swensen, Yale has invested across a wide range of "alternative assets", from private equity and hedge funds to timber...*
>
> *One idea behind the Yale strategy was that endowments have the luxury of time, since their liabilities (paying for new buildings, academic salaries and so on) stretch far into the future. They can thus afford to invest in illiquid asset classes. Such asset classes may offer a better return simply because other investors (mutual funds, for example) are unwilling or unable to deal with illiquidity.*
>
> *In addition, the traditional dominance of domestic equities within institutional portfolios put lots of eggs in one basket. There were other sources of return—the management skills of private-equity houses, the market for distressed bonds—where the trade-off between risk and reward might be better than could be obtained from the S&P 500 index...*
>
> *But the model came into question in 2008 and 2009, when the financial crisis hit. In the year to June 30th 2009, the Yale endowment fell by 24.6% and Harvard's portfolio fell by 27.3%, losing the latter a whopping $10 billion. (there has been a modest recovery since then).*

On the other hand,

> *The financial crisis decimated the wealth of the average American. [...] To a large extent you can blame the housing bubble. For many middle- and lower-middle-income Americans their home made up nearly all of their wealth.*
>
> *They were largely invested in a single asset that did very poorly in this period. Richer Americans, who held other assets, did not see such large declines. This demonstrates how important it is to diversify your wealth, even when you don't have much of it. Why did these Americans have so much housing in their portfolio?*

Obviously there was the housing bubble and sub-prime mortgages, which made levering up to buy a house too easy. But I think it goes deeper than that. Anecdotally there seemed to be a mentality that you should buy as much housing as you can afford, not as much as you need. That may be because of the pervasive view during this period that housing is always a good investment. That idea coupled with the favoured tax treatment of home ownership created an incentive for Americans to put nearly all of their wealth into a single asset. Unless we see a strong rebound in house prices, many Americans may pay the price for that with a grim retirement.

From: Buttonwood, *The Economist* (March 10th, 2011) "Yale may not have the key: why diversification doesn't work"; and Free exchange, *The Economist*, Economics (June 18th, 2012) "Housing and wealth: the perils of not diversifying".

Reading. The following excerpt is from an article in *The New York Times*. Discuss it critically.

It's a classic moment in sports history. With less than 20 seconds left in Game 6 of the 1998 N.B.A. finals and the Chicago Bulls down by one, Michael Jordan goes one-on-one with Bryon Russell of the Utah Jazz. He pushes off (clearly!), Russell stumbles and the ball hits nothing but net. Game over. Bulls win.

Now let's imagine that something different happened. Jordan misses the shot in Game 6, and Game 7 comes down to the same spot: fewer than 20 seconds left with the Bulls down by one. If you're Phil Jackson, the head coach, do you set up the last play for Jordan, or does the ball go to someone else?

Of course the right strategy is to put the ball in Jordan's hands. Just because he missed the shot before doesn't mean it was the wrong strategy to have Jordan shooting the ball in the final seconds. The odds are incredibly high that he will make the shot even though he missed it the night before.

I bring this up because it perfectly captures the investing adage that never seems to die: diversification is "broken." It seems as if this story pops up every year, but it's not really about anything new. Both Joshua M. Brown at The Reformed Broker and Barry Ritholtz at The Big Picture have written blog posts about it recently. Mr. Brown quoted an adviser who said: "Why bother diversifying at all? It's just a drag on performance. What's the point of owning any bonds or international stocks?"

From: Richards, C.,*The New York Times* (August 12th, 2013) "Diversification isn't broken, it just takes a while".

Note 7

Foundations of expected utility theory

IT IS TIME NOW TO EXPLORE HOW GOOD THE EXPECTED UTILITY ASSUMPTION IS. Let us go back to the basic setting where there is a finite set of possible outcomes, and where the set of lotteries over those outcomes is Δ. Henceforth, we assume *only* that some preference relation $\succsim$ is rational, and define $\succ$ and $\sim$ as before. Our goal is to understand what it means to assume that $\succsim$ has an EU representation.

To begin, note that while it may seem natural that the individual's preferences, $\succsim$, are a binary relation over Δ, in doing this we are imposing the condition that the individual cares about the risk (randomness) she faces, and not about the process that ultimately determines that risk; this condition is known as "consequentialism," and there are reasons, both psychological and philosophical, to think that this assumption is restrictive. We won't go into those considerations, focusing more on the types of behavior that are consistent with the decision maker being an EU individual.

7.1. Further properties of preferences

WE SAY THAT RELATION $\succsim$ satisfies *monotonicity* if given two lotteries, p and $\tilde{p}$ such that $p \succ \tilde{p}$, the following statement is true: $\alpha p + (1-\alpha)\tilde{p} \succ \beta p + (1-\beta)\tilde{p}$ if, and only if, $\alpha > \beta$. In words, a decision-maker with monotonic preferences prefers more of a better lottery to more of a worse lottery.

EXAMPLE 7.1. Consider the decision-maker of Exercise 1.2. Argue that these preferences satisfy monotonicity.

Answer: It's best to proof the two properties that define monotonicity independently:

1. That, given $p \succ \tilde{p}$, if it is true that

$$\alpha p + (1-\alpha)\tilde{p} \succ \beta p + (1-\beta)\tilde{p},$$

then it must be that $\alpha > \beta$.

2. that if it is true that

$$\alpha p + (1-\alpha)\hat{p} \succsim \alpha\tilde{p} + (1-\alpha)\hat{p},$$

for some number $0 < \alpha < 1$ and some lottery $\hat{p}$, then it must be true that $p \succsim \tilde{p}$.

For the first claim, note that since $p \succ \tilde{p}$, it must be true that $p_1 + p_2 > \tilde{p}_1 + \tilde{p}_2$. If, moreover, $\alpha p + (1-\alpha)\tilde{p} \succ \beta p + (1-\beta)\tilde{p}$, then one must also have that

$$[\alpha p_1 + (1-\alpha)\tilde{p}_1] + [\alpha p_2 + (1-\alpha)\tilde{p}_2] > [\beta p_1 + (1-\beta)\tilde{p}_1] + [\beta p_2 + (1-\beta)\tilde{p}_2].$$

But one can rewrite this expression as

$$\alpha(p_1 + p_2) + (1-\alpha)(\tilde{p}_1 + \tilde{p}_2) > \beta(p_1 + p_2) + (1-\beta)(\tilde{p}_1 + \tilde{p}_2),$$

which implies the result: it must be true that $\alpha > \beta$.

For the second claim, if $\alpha p + (1-\alpha)\hat{p} \succsim \alpha\tilde{p} + (1-\alpha)\hat{p}$, it must be that

$$[\alpha p_1 + (1-\alpha)\hat{p}_1] + [\alpha p_2 + (1-\alpha)\hat{p}_2] > [\alpha\tilde{p}_1 + (1-\alpha)\hat{p}_1] + [\alpha p_2' + (1-\alpha)p_2''].$$

Rewriting this as

$$\alpha(p_1 + p_2) + (1-\alpha)(\hat{p}_1 + \hat{p}_2) > \alpha(\tilde{p}_1 + \tilde{p}_2) + (1-\alpha)(\hat{p}_1 + \hat{p}_2),$$

one concludes that $p_1 + p_2 \geq \tilde{p}_1 + \tilde{p}_2$, which gives the result: it follows that $p \succsim \tilde{p}$. □

Example 7.2. Consider a decision-maker who faces uncertainty over a finite set of possible outcomes, $\mathcal{X} = \{1, 2, 3, 4\}$, and suppose that her preferences are represented by the following function:

$$U(p) = \begin{cases} 1, & \text{if } p_x = 1/4 \text{ for all } x; \\ 0, & \text{otherwise.} \end{cases}$$

Argue that the individual's preferences *do not* satisfy the following property: for all p and $\tilde{p}$ and all $\alpha \in [0, 1]$, if $p \sim \tilde{p}$, then $\alpha p + (1-\alpha)\tilde{p} \sim \tilde{p}$.

Answer: Consider the lotteries $p = (1/2, 1/2, 0, 0)$ and $p = (0, 0, 1/2, 1/2)$. Since $U(p) = U(\tilde{p}) = 0$, so $p \sim \tilde{p}$). Now, it $\alpha = 1/2$, we have $\hat{p} = \alpha p + (1-\alpha)\tilde{p} = (1/4, 1/4, 1/4, 1/4)$, so $U(\hat{p}) = 1$ and $\hat{p} \succ p$. □

Example 7.3. Consider a decision-maker who faces uncertainty over a finite set of possible outcomes, $\mathcal{X} = \{1, 2, 3\}$, and suppose that her preferences are represented by the following function:

$$U(p) = \min\{p_1, p_2\} + a(1 - p_1 - p_2)$$

Argue that he individual's preferences satisfy the following *convexity* condition: given $p \succ \tilde{p}$, if $0 \leq \alpha \leq 1$, then $\alpha p + (1-\alpha)\tilde{p} \succsim \tilde{p}$.

Answer: If $p \succsim p'$, it must be true that $U(p) \geq U(p')$. Also, note that

$$\min\{\alpha p_1 + (1-\alpha)p'_1, \alpha p_2 + (1-\alpha)p'_2\} \geq \alpha \min\{p_1, p_2\} + (1-\alpha)\min\{p'_1, p'_2\}.$$

From this, it follows that

$$\begin{aligned} U(\alpha p + (1-\alpha)p') &= \min\{\alpha p_1 + (1-\alpha)p'_1, \alpha p_2 + (1-\alpha)p'_2\} \\ &\quad + [\alpha(1-p_1-p_2) + (1-\alpha)(1-p'_1-p'_2)] \\ &\geq \alpha[\min\{p_1, p_2\} + (1-p_1-p_2)] \\ &\quad + (1-\alpha)[\min\{p'_1, p'_2\} + (1-p'_1-p'_2)] \\ &\geq \min\{p'_1, p'_2\} + (1-p'_1-p'_2) \\ &= U(p'). \end{aligned}$$

□

Exercise 7.1. Consider the decision-maker of Exercise 1.3. Argue that this individual's preferences *do not* satisfy the convexity condition defined in Example 7.3.

A very important condition imposes that the decision-maker *values* the outcomes of the lotteries for themselves and then, separately, the randomness induced over them by the lottery:

Definition. *Preference relation* $\succsim$ *satisfies* independence *if, given two lotteries* p *and* $\tilde{p}$*, the following statements are true:*

1. *if* $p \succsim \tilde{p}$*, then for any number* $0 \leq \alpha \leq 1$ *and any lottery* $\hat{p}$ *we have that*

$$\alpha p + (1-\alpha)\hat{p} \succsim \alpha\tilde{p} + (1-\alpha)\hat{p};$$

2. *if for some number* $0 < \alpha < 1$ *and some lottery* $\hat{p}$ *we have that*

$$\alpha p + (1-\alpha)\hat{p} \succsim \alpha\tilde{p} + (1-\alpha)\hat{p},$$

then $p \succsim \tilde{p}$.

The latter property is controversial, and we will come back to it later. The following exercises relate these two properties.

Exercise 7.2. Argue that independence of $\succsim$ implies the following property: for any pair of lotteries p and $\tilde{p}$:

a. if $p \sim \tilde{p}$, then for any $0 \leq \alpha \leq 1$ and any lottery $\hat{p}$, $\alpha p + (1-\alpha)\hat{p} \sim \alpha\tilde{p} + (1-\alpha)\hat{p}$;

b. if for some $0 < \alpha \leq 1$ and some lottery $\hat{p}$ we have that $\alpha p + (1-\alpha)\hat{p} \sim \alpha\tilde{p} + (1-\alpha)\hat{p}$, then $p \sim \tilde{p}$.

c. if $p \succ \tilde{p}$, then for any $0 < \alpha < 1$ and any lottery $\hat{p}$, $\alpha p + (1-\alpha)\hat{p} \succ \alpha\tilde{p} + (1-\alpha)\hat{p}$;

d. if for some $0 < \alpha < 1$ and some lottery $\hat{p}$ we have that $\alpha p + (1-\alpha)\hat{p} \succ \alpha\tilde{p} + (1-\alpha)\hat{p}$, then $p \succ \tilde{p}$; and

e. if $p \succ \tilde{p}$ and $0 < \alpha < 1$, then $p \succ \alpha p + (1-\alpha)\tilde{p}$ and $\alpha p + (1-\alpha)\tilde{p} \succ \tilde{p}$.

Exercise 7.3 (Harder). Argue that independence of $\succsim$ implies its monotonicity.[1]

7.2. The von Neumann–Morgenstern theorem

We now ask the question of when $\succsim$ has an EU representation. A first result is that independence is a necessary condition for this:

Exercise 7.4. Argue that if $\succsim$ has an EU representation then it satisfies independence.

A seminal result is decision theory was obtained by John von-Neumann (Hungary, 1903-1957) and Oskar Morgenstern (Germany 1902-1977).[2] It implies that, under a (technical) assumption of "continuity," independence is also a sufficient condition. Together with Exercise 7.4, it implies that independence is equivalent to the existence of an EU representation for an agent's preferences over lotteries. The full proof of the von Neumann-Morgenstern theorem is beyond these lectures, but a simplified argument is presented next.

For simplicity, we concentrate only on a small subclass of lotteries, rather than on the whole space Δ. We say that a lottery is *simple* if it gives positive probability to at most two outcomes in $\mathcal{X}$.[3] For simplicity, then, we can denote a simple lottery as a triple consisting of a number and two outcomes, $L = (p, x, x')$, with $0 \leq p \leq 1$ and $x, x' \in X$, and with the interpretation that the lottery gives outcome x with probability p, and outcome x' with probability $1-p$. Let $\mathcal{L}_1$ be the space of simple lotteries, $\mathcal{L}_1 = [0,1] \times \mathcal{X} \times \mathcal{X}$.

A *compound lottery* is a device that gives other lottery or lotteries as prizes. We will concentrate on compound lotteries that give positive probability to at most two simple lotteries,[4] and denote them by (p, L, L'), a number and two simple lotteries, $L, L' \in \mathcal{L}_1$. Let $\mathcal{L}_2$ be the space of compound lotteries, $\mathcal{L}_2 = [0,1] \times \mathcal{L}_1 \times \mathcal{L}_1$.

For our argument, we consider only degenerate, simple and our simplified definition of compound lotteries, so we take $\succsim$ as defined over $\mathcal{L} = \mathcal{X} \cup \mathcal{L}_1 \cup \mathcal{L}_2$. In order to keep consistency with the analysis above, we need consider an individual who cares about outcomes, and not about how these outcomes are presented, so we impose the following "consequentialist" assumptions on $\succsim$: for all $p, \tilde{p} \in [0,1]$ and for all $x, x' \in X$, $(p, x, x') \sim (1-p, x', x)$, $(1, x, x') \sim x$ and $[p, (\tilde{p}, x, x'), x'] \sim (p\tilde{p}, x, x')$. For simplicity, suppose also that we can find $x_*, x^* \in X$ such that for every outcome $x \in X$ we have that $x \succsim x_*$ and $x^* \succsim x$.

We will say that $\succsim$ satisfies *continuity* if, for any $x, x', x'' \in \mathcal{X}$ such that $x \succsim x' \succsim x''$, we can find a number $0 \leq p \leq 1$ such that $(p, x, x'') \sim x'$.

Theorem (The von Neumann–Morgenstern Theorem). *If $\succsim$ satisfies independence and continuity, then it has an EU representation (with continuous index u).*

[1] This exercise is a tiny bit more complicated than the others, so here go two hints. First, suppose that you want to write $\beta p + (1-\beta)\tilde{p}$ as $\gamma p + (1-\gamma)(\alpha p + (1-\alpha)\tilde{p})$, given that $\beta > \alpha$; what value must γ have? Then, notice the last property of Exercise 7.2.

[2] Hence the name that we've been using for the EU representations of preferences!

[3] The term "simple" is normally used for lotteries that pay in outcomes and not in other lotteries; here, I am using it is that sense, but making it stronger to require that they pay in only one or two outcomes.

[4] As before, the term "compound" is normally used for lotteries that pay in other lotteries; here, I am using it is that sense, but making it stronger to require that they pay in only one or two lotteries.

Proof. Now, if $\succsim$ satisfies monotonicity, it is relatively easy to construct a utility function representing it over the space of simple lotteries: by continuity, for any lottery in $\mathcal{L}$, we can find $p \in [0, 1]$ such that $L \sim (p, x^*, x_*)$; by monotonicity, such $p \in [0, 1]$ has to be unique; then, just let $U : \mathcal{L} \to \mathbb{R}$ be defined by letting $U(L)$ be the unique number $p \in [0, 1]$ such that $L \sim (p, x^*, x_*)$.

Since $\mathcal{L}$ includes degenerate lotteries, we can define $u : \mathcal{X} \to \mathbb{R}$ by letting $u(x) = U((1, x, x))$. Now, we just want to show that the expected utility property is satisfied in the following sense: for every simple lottery (p, x, x'), $U((p, x, x')) = pu(x) + (1 - p)u(x')$.

Notice that, by construction, $(p, x, x') \sim [U(p, x, x'), x^*, x_*]$, whereas, by independence,

$$(p, x, x') \sim [p, (u(x), x^*, x_*) , (u(x'), x^*, x_*)].$$

By direct computation, it follows that

$$(p, x, x') \sim [pu(x) + (1 - p)u(x'), x^*, x_*],$$

which implies, by monotonicity, that $U(p, x, x') = pu(x) + (1 - p)u(x')$. □

EXERCISE 7.5. Argue that if $\succsim$ satisfies independence, then:

a. $x \succ x'$ and $0 \leq \tilde{p} < p \leq 1$ imply that $(p, x, x') \succ (\tilde{p}, x, x')$;

b. $L \succ L'$ and $0 \leq \tilde{p} < p \leq 1$ imply that $(p, L, L') \succ (\tilde{p}, L, L')$;

c. if $x \succsim x'$, then for any x'' and any $0 \leq p \leq 1$ it is true that $(p, x, x'') \succsim (p, x', x'')$;

d. if for some x'' and some $0 \leq p \leq 1$ it is true that $(p, x, x'') \succsim (p, x', x'')$, then $x \succsim x'$;

e. if $L \succsim L'$, then for any L'' and any $0 \leq p \leq 1$ it is true that $(p, L, L'') \succsim (p, L', L'')$;

f. if for some L'' and some $0 \leq p \leq 1$ it is true that $(p, L, L'') \succsim (p, L', L'')$, then $L \succsim L'$.

7.3. Is independence a good assumption?

INDEPENDENCE IS A STRONG assumption: it implies that the decision-maker is able to evaluate outcomes without worrying about the randomness between them (which is fine, as degenerate lotteries allow for that), and then evaluates the randomness in a way which is perfectly consistent with the "deterministic" evaluation. The second step is controversial: introspection can give you cases in which the winning and losing outcomes are so «related» that the randomness cannot be assessed independently.

7.4. Subjective expected utility

THE FRAMEWORK THAT WE have been considering is *not* the most general setting in which one can study decision making under uncertainty: the decision maker is assumed to choose the probabilities of the outcomes that she faces, and it is assumed that these probabilities are objectively presented to her. While a more general model is beyond the level of our course, we can at least introduce its main difference.

The EU model of Note 1 assumes that the set of contingencies that describe the world and the set of contingencies that the decision maker cares about are the same. Instead, let us continue to denote by $\mathcal{X} \neq \emptyset$ be a set of *outcomes*, namely that the full description of what determines the individual's well-being. Independently, a *state of the world* is a comprehensive description of the status of all the contingencies that can affect the decisions of the individual.[5] Let $\mathcal{S}$ be the set of states of the world.

In this richer setting, an *act* is a mapping of states of the world into outcomes, namely a function $a : \mathcal{S} \to \mathcal{X}$. The Subjective EU Model does not assume that individuals rank (or choose) lotteries over outcomes. Instead, they choose acts: plans of what outcome will come about in each state of the world. The main difference, is that the choice of an act *does not* pin down the probability of each outcome, as the state of the world, which is what will determine the outcome, is out of the control of the individual.

As an alternative to the von-Neumann–Morgenstern analysis, Leonard Savage (USA, 1917-1971) suggested the following model, which is more general: the decision maker has preferences $\succsim$ over the space of acts, which is denoted by $\mathcal{A}$; these preferences have a subjective expected utility (SEU) representation, if there exists a probability distribution π, defined over the set $\mathcal{S}$ of states of the world, and a utility index $u : \mathcal{X} \to \mathbb{R}$ such that

$$a \succsim \tilde{a} \Leftrightarrow \mathrm{E}_{\pi}[u(a)] \geq \mathrm{E}_{\pi}[u(\tilde{a})],$$

where

$$\mathrm{E}_{\pi}[u(a)] = \sum_{s} \pi_s u(a(s)).$$

The main gain of this construction is that the probability distribution π is part of the individual's thought process, and need not coincide with any external distribution (which, incidentally, need not even exist).

[5] In the words of K. Arrow (*Essays on the Theory of Risk Bearing*, 1971, p. 45), it is "a description of the world so complete that, if true and known, the consequences of every action would be known."

Note 8

Non-EU theories

THERE ARE ALTERNATIVE MODELS, THOUGH. As mentioned before, it is not difficult to imagine situations in which a decision maker finds it difficult to assess the probability distribution of some risks that she faces, or where the possible outcomes she can obtain make her preferences violate the independence axiom. Even simple introspection can give us instances of those situations, and it follows from before that the model we have been using so far is not suitable for those situations.

8.1. Some experimental challenges

NEXT, WE PRESENT FOUR experiments, designed by psychologists, where *the majority* of people display behavior that would be hard to reconcile with the EU (or even the SEU) paradigm. These examples have motivated the development of alternative theories, which we then present, briefly.

8.1.1. Allais's paradox

A CANONICAL OBSERVATION IS the following: consider a space of monetary outcomes, and suppose that the following lotteries are available:

$p^1 =$

x	\$0	\$1M	\$5M
p^1_x	0	1	0

$p^2 =$

x	\$0	\$1M	\$5M
p^2_x	0.01	0.89	0.1

$p^3 =$

x	\$0	\$1M	\$5M
p^3_x	0.89	0.11	0

$p^4 =$

x	\$0	\$1M	\$5M
p^4_x	0.9	0	0.1

It has been observed that when asked to compare these lotteries, many people respond that $p^1 \succ p^2$ and $p^4 \succ p^3$. (Are these your preferences too?) Well, it turns out that these preferences cannot have an EU representation! To see this, notice that a representation would require numbers $u(0)$, $u(1)$ and $u(5)$ such that $U(p) = \mathrm{E}_p(u)$. Then, the revealed choices say that u satisfies the following two inequalities:

$$\mathrm{E}_{p^1}(u) = u(1) > 0.01u(0) + 0.89u(1) + 0.1u(5) = \mathrm{E}_{p^2}(u),$$

while

$$\mathrm{E}_{p^3}(u) = 0.89u(0) + 0.11u(1) < 0.9u(0) + 0.1u(5) = \mathrm{E}_{p^4}(u).$$

But, then, from the first equation $0.11u(1) > 0.01u(0) + 0.1u(5)$, while from the second equation $0.11u(1) < 0.01u(0) + 0.1u(5)$.

This observation, proposed by Maurice Allais (France, 1911–2010), suggests that many people don't conform to the assumptions required for their preferences to have an EU representation.

8.1.2. Gains vs. losses

A SECOND SIMPLE EXPERIMENT captures part of what may explain Allais's paradox in a more direct way. Consider the following choice situation:

> You are given \$1K, and are asked to choose between the following two options:
>
> (a) You get an additional \$500 for sure; or
>
> (b) You get an additional \$1K with probability $1/2$ (and otherwise get nothing more).

Now consider this alternative situation:

> You are given \$2K, and are asked to choose between the following two options:
>
> (c) You lose \$500 for sure; or
>
> (d) You lose \$1K with probability $1/2$ (and otherwise lose nothing).

Experimentally, many people prefer (a) in the first situation and (d) in the second one. (Do you?) The problem is that from the point of view of income payoffs the two situations are equivalent, so an EU agent should choose the same option in both.

8.1.3. Ellsberg's one-urn paradox

THERE IS ANOTHER OBSERVATION of usual individual choices that are inconsistent with the von Neumann–Morgestern representation of preferences, and which makes a "problem" very apparent. It is due to Daniel Ellsberg.[1]

Suppose that there is an urn that contains 90 balls. Thirty of them are known to be red, while each of the remaining 60 is known to be blue or green, but the proportion of each of these two colours is not known. Four lotteries are presented:

1. a \$1M prize if a red ball is drawn;
2. a \$1M prize if a blue ball is drawn;
3. a \$1M prize if a red or a green ball are drawn; or
4. a \$1M prize if a blue or a green ball are drawn.

[1] A pretty remarkable guy: he leaked classified information about how the U.S. government had been, let's say, less than candid about the Vietnam war.

It is usually stated by people that they prefer the first lottery to the second lottery, and the fourth lottery to the third lottery. (One more time, are these your preferences too?) But, again, these preferences are not consistent with an EU representation. To see this, suppose that the preferences can be represented with a utility index u. For simplicity of notation let the utility of not getting the prize is 0, and the utility of getting it is 1.[2] Then, the choices say that

$$\Pr(\text{red}) > \Pr(\text{blue}) \text{ and } \Pr(\text{red or green}) < \Pr(\text{blue or green}).$$

But we can rewrite these inequalities as

$$\Pr(\text{red}) > \Pr(\text{blue}) \text{ and } 1 - \Pr(\text{blue}) < 1 - \Pr(\text{red}),$$

which is obviously inconsistent.

8.1.4. Ellsberg's two-urn paradox

Alternatively, suppose that there are two urns: Urn 1 contains 50 red balls and 50 blue balls. Urn 2 contains 100 balls, each of which is known to be red or blue, but the proportion of each of these two colours is not known. Now, four lotteries are offered:

1. a \$1M prize if a red ball is drawn from urn 1;
2. a \$1M prize if a blue ball is drawn from urn 1;
3. a \$1M prize if a red ball is drawn from urn 2; or
4. a \$1M prize if a blue ball is drawn from urn 2.

The observation is that many people are indifferent between the first two lotteries, indifferent between the third and fourth lotteries, and strictly prefer the first lottery to the third lottery. Since the prize is the same, the first indifference requires that in the first urn both colours have probability $1/2$, while the second indifference implies that the same is true for the second urn. But this contradicts the third observation, for if the previous probabilities are correct and the individual's preferences have an EU representation, then the individual should be indifferent between the first and third lotteries.

8.2. Non-EU theories

We now present, briefly, some possible explanations, or models, for the observational paradoxes we just introduced. Some of these models are directly motivated by the experiments.

8.2.1. Prospect theory I: perceived probabilities

For a person with some mathematical training, the concept of probability seems very simple, perhaps. But the fact that the concept requires some understanding suggests that it is not completely natural for human decision makers.

[2] Note that this assumption implies no loss of generality.

Amos Tversky (Israel, 1937–1996) and Daniel Kahneman (Israel, 1934) have proposed a theory of how the human brain deals with uncertainty that distinguishes between what the objective probability of an event is and what a human brain perceives it to be.[3]

Let us imagine that when the actual probability of an outcome is p_x, a decision maker's brain perceives it to be $\pi(p_x)$. Obviously, $0 \leq \pi(p_x) \leq 1$ for all p, and presumably π is an increasing function. In this case, even if we don't question any of the premises of the EU model, we should modify our equations so that, the representation of the individual's preferences should be that, when presented with a lottery p over some outcome space $\mathcal{X}$,

$$U(p) = \sum_x \pi(p_x) u(x),$$

since what matter for the individual's well-being is her perception of risk.

Tversky and Kahneman postulated that our brains tend to under-estimate very low probabilities and to over-estimate very high ones, so that the π function of a typical individual looks like the function of Fig. 8.1. There, the graph of the function is below the diagonal when p is close to 0, and above it when p is near 1.

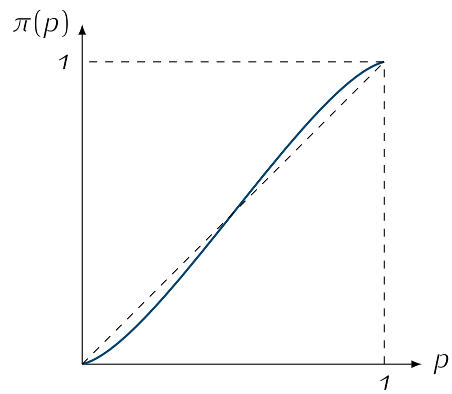

Figure 8.1: Actual and perceived probabilities.

One complication with this model, however, is that $\pi(p)$ need not be a probability distribution even when p is. In that sense, the expression defining function U cannot be called an expected utility. Some extensions of the theory have been proposed to deal with this problem.

8.2.2. Loss aversion

An alternative explanation for Allais's paradox is that when people assess events that have not occurred, they do so by comparing them to some reference point. If the reference point changes, an individual may feel differently about some event, even when that event is objectively the same. In the paradox, when comparing p^1 and p^2 a person may take as a reference point the income of \$1M that she earns for sure with p^1, and the possibility of losing that money may make lottery p^2 very unappealing. If, on the other hand, when comparing p^3 and p^4 the reference point is the

[3] If, for instance, our brain perceives 0.11 and 0.11 as being essentially the same number, it's not surprising that many of us find p^4 better than p^3 in the Allais experiment, regardless of how we felt about p^1 and p^4.

\$0 income that is most likely in both lotteries, how she feels about gambling to get \$1M or \$5M may induce the assessment that seemed inconsistent above.

Again, Tversky and Kahneman suggested an explanation for this phenomenon: people get attached to what they own, and they value their possessions more than what they would be willing to pay for them if they didn't already have them. One way to introduce this idea to a model of uncertainty is to introduce two functions, g and ℓ, that represent how the agent assess gains and losses relative to the reference point, respectively. Assume that both functions are increasing and satisfy that $g(0) = \ell(0) = 0$. When the agent has an income level $\bar{x}$ as her reference point and assesses an alternative income level x, her utility is given by

$$v(x, \bar{x}) = u(x) + g(x - \bar{x}), \quad \text{when } x \geq \bar{x},$$

and by

$$v(x, \bar{x}) = u(x) - \ell(\bar{x} - x), \quad \text{when } x < \bar{x}.$$

This means that, instead of the simple Bernoulli index, the utility over sure outcomes is assessed according to a function like the one in Fig. 8.2. There, the graph presents a kink at the reference income. To the left of the kink, the marginal utility is given by $u'(x) + \ell'(\bar{x} - x)$; to the right, it is $u'(x) + g'(x - \bar{x})$. The kink in the graph is given under the assumption that $\ell'(0) > g'(0)$; this means that the marginal utility is smaller for (infinitesimal) gains than the marginal dis-utility of (infinitesimal) loses. This phenomenon is known as *loss aversion*.

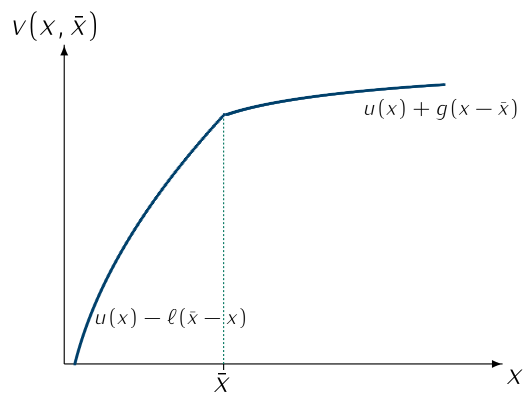

Figure 8.2: Sure income and utility.

8.2.3. Prospect theory 2: gain-loss asymmetries

The observation of the experiment in Subsection 8.1.2 can be explained through a further assumption on the shape of the gain and loss functions we just introduced. Note that if, in the presence of these functions, the individual were "always risk averse," then, in addition to u, the function g should be concave and the function ℓ should be convex.[4] But note that also that the choice of (a) in the first situation of the experiment corresponds to risk-averse behavior, while the choice of (b) in the second situation seems to display love for risk.

[4] So that $-\ell$ is concave.

Tversky and Kahneman, again, suggest an explanation for this: that the loss function need not be convex. Suppose for instance that, given the reference point $\bar{x}$, the utility of sure income levels is as in Fig. 8.3. In this case, the agent displays a very strong aversion to small losses, but this feeling vanishes as the size of the loss increases, which can explain why a sure loss of $500 when the person's reference income is $2K can seem more costly than a loss of $1K with probability $1/2$.

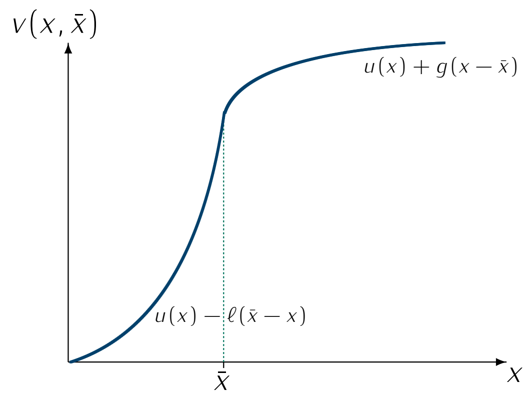

Figure 8.3: Sure income and utility.

8.2.4. Ambiguity

The observations of Ellsberg are particularly illustrative of the problem that the von Neumann–Morgestern theory may be displaying in these cases. Consider the one-urn experiment. Note that the probability of getting the prize in the first lottery is known to be $1/3$, but the probability of getting it in the second lottery is *not known*. With the information given, it could be any number in $\{0, 1/90, 1/45, 1/30 \ldots, 2/3\}$. On the other hand, in the fourth lottery the probability of winning is known to be $2/3$, whereas in the third one it is not, and could be any number in $\{1/3, 31/90, 16/45, 11/30, \ldots, 1\}$. In both choice situations, individuals seem to prefer the lottery where the probability of winning the price is known to that where it is not.

This feature of the problem is known as *ambiguity*.[5] Note that if the individual, in her mind, resolved the ambiguity by picking a probability for blue balls, and applied it consistently in both situations, then the preferences would be different. Say that the individual thinks that the probability of drawing a blue ball is $1/6$, and therefore the probability of a green ball is $1/2$. Then, she should indeed prefer the first lottery to the second one, but should also prefer the third lottery to the fourth one: the probability of getting a red or a green ball is now $1/3 + 1/2 = 5/6$, whereas the probability of drawing a blue or a green ball is $1/6 + 1/2 = 2/3 < 5/6$.[6]

[5] Historically, the feature of not knowing the probability distribution was called *uncertainty*, whereas the fact that the decision problem implied some randomness was called *risk*. But in many cases in which the distribution of random events was known by the decision-maker, the terms uncertainty and risk were used as synonyms, which may give rise to confusion.

[6] Alternatively, if the probability chosen is such that the fourth lottery is preferred to the third one, then the second lottery should be preferred to the first one.

The solutions that have been proposed formalize the idea that people prefer to have more precise information, and in this case this implies that the choice of probabilities is not consistent between the two choice situations. For instance, suppose that an individual is pessimistic in the presence of imprecise information: in each of the two choices she is presented with, she considers the worst case for her. Then, when considering lottery two, she thinks that the probability of winning is null, for she cannot rule out the possibility that all the 60 balls are green. But when she is thinking of the third lottery, she considers the worst possible scenario, which is now different: she cannot rule out that all the 60 balls are blue either!

In this case, her choices would be as observed in Ellsberg's paradox. But, obviously, her preferences are not representable as the expected value of some utility index for given probabilities. Instead, under ambiguity the decision-maker has a set of probability distributions that she cannot rule out, say $\Delta_o \subseteq \Delta$, and evaluates her choices using a worst-case scenario rule:

$$\min_{p \in \Delta_o} \mathrm{E}_p[u].$$

This rule is known as *maximin preferences*, or as *Gilboa-Schmeidler representation*, recognizing the two economists who proposed it.[7]

8.2.5. Capacities

The two-urn paradox suggests another solution. In the first urn, the probabilities are known to be 1/2, while in the second one there is no information about what they are. In the absence of information, a decision maker could just think that both colours of ball are equally likely in the second urn, but since she is basing this conclusion only in the lack of a reason to think otherwise, she may "penalize" gambling on this urn. This would explain the pattern of choices presented by Ellsberg, and can be formalized by weakening the concept of probability, something that we will not do here.

Reading. The following excerpt is from an article in *The Financial Times*. Discuss it critically.

> *It is true that humans have an amazing capacity to analyse complex problems. Unfortunately, the calculating part of our brain depends on intuitions to stir it into action. Before [the work of Nobel Laureate Daniel] Kahneman, we tended to view intuitions as revealing a profound, if ill-defined, understanding of the world. We bowed down before the judgment of experts.*
>
> *Our minds are not naturally good at probabilities; they dislike ambiguity and doubt; an ingrained desire to construct coherent narratives leads us to seek confirming evidence, while disregarding information that refutes our prior view. As a result, we are often more confident than the circumstances warrant. We also depend too much on recent experience when forming judgments—what Mr Kahneman and his long-time collaborator Amos Tversky call the "availability heuristic". Emotions constantly inform our judgments.*

[7] Other models consider agents who do not concentrate all their attention in the worst possible outcome.

What Mr Kahneman calls "cognitive illusions" are a constant menace to investors. We repeatedly fail to anticipate financial disasters—such as the dotcom collapse of 2000 or the implosion of the US housing market a few years later—because at the time they were unfamiliar events.

Conversely, once we start to consider rare outcomes we are prone to exaggerate their likelihood. In today's investment world, every swan is deemed black. Experiments show that people react favourably to words that are repeated to them. "A reliable way to make people believe in falsehoods," writes Mr Kahneman, "is frequent repetition, because familiarity is not easily distinguished from truth. Authoritarian institutions and marketers have always known this fact."

Stockbrokers, who bombard clients with often dubious investment research, also seem aware of this human frailty. Experiments suggest that when people have money on their mind they become more selfish. Perhaps this explains why Wall Street fails on occasion to exhibit the highest ethical standards.

Like everyone else, investors are prone to overconfidence. Confidence usually derives from the coherence of the information at hand and the ease of processing it, rather than its validity. Investors are constantly deluded by compelling stories. In the financial world, skill is often confused with luck. "Professional investors ... ," writes Mr Kahneman damningly, " ... fail a basic test of skill, persistent achievement." Yet in finance we also handle failure badly. We tend to revise the history of our beliefs in the light of events. Agents are blamed for not anticipating events, however unpredictable.

From: Chancellor, E., *The Financial Times* (January 8th, 2012) "Humans are naturally bad investors".

Part II

Asymmetric Information

Note 9

Insurance markets

ONE OF THE CANONICAL EXAMPLES OF ASYMMETRIC INFORMATION IS INSURANCE MARKETS. Implicit in the trade of financial instruments, there is the possibility that informational asymmetries can perturb the functioning of theirl markets. We will consider the simplest possible setting, in which there are two periods and two (idiosyncratic) states of the world that can occur in the second period. We will allow for the existence of two types of agents, who differ in the probabilities with which they face each of the two future states. Importantly, we will assume that the type of each individual is observed *only* by herself.

9.1. Perfect information

TO ESTABLISH A BENCHMARK, let us first assume that there is only one type of agent in the market. In state of the world 1, her income is w; in state 2, she faces a loss and her income is $w-\ell$. The probability of incurring this loss is p. We assume that she is a strictly risk-averse, EU maximizer, so her ex-ante preferences over consumption plans $x=(x_1,x_2)\in\mathbb{R}^2$ are

$$U(x)=(1-p)u(x_1)+pu(x_2).$$

It is easy to see that this individual has downward sloping indifference curves, with marginal rate of substitution

$$\left.\frac{\mathrm{d}x_2}{\mathrm{d}x_1}\right|_{\mathrm{d}U=0}=-\frac{(1-p)u'(x_1)}{pu'(x_2)}.$$

For future reference, it is useful to note that over the full insurance line, the marginal rate of substitution is

$$\left.\frac{\mathrm{d}x_2}{\mathrm{d}x_1}\right|_{\mathrm{d}U=0}=-\frac{1-p}{p}.$$

Since $u''<0$, the marginal rate of substitution is, then, less than this value for plans below the full insurance line.[1]

[1] To see this, note that, along an indifference curve,

$$\frac{\mathrm{d}{x_2}^2}{\mathrm{d}^2x_1}=-\frac{(1-p)\left[u''(x_1)u'(x_2)-u'(x_1)u''(x_2)\frac{\mathrm{d}x_2}{\mathrm{d}x_1}\right]}{pu''(x_2)^2}>0.$$

Now, let Γ be the premium of an insurance contract against the income loss, and let α be this contract's net coverage.[2] As a function of these two variables, the individual preferences are

$$v(\alpha, \Gamma) = (1-p)u(w-\Gamma) + pu(w-\ell+\alpha).$$

For the insurance company, which we assume to be risk neutral, profits are

$$\pi(\alpha, \Gamma) = (1-p)\Gamma - p\alpha.$$

Let us assume that competition between insurance companies drives their profits to 0. Then, it is immediate that for perturbations $\mathrm{d}\alpha$ and $\mathrm{d}\Gamma$ to the contract that the insurance company will offer, $\mathrm{d}\pi = 0$, which means that the *isoprofit* line has slope

$$\left.\frac{\mathrm{d}\alpha}{\mathrm{d}\Gamma}\right|_{\pi=0} = \frac{1-p}{p}.$$

At equilibrium, the isoprofit line and the indifference curve are tangent, so

$$\frac{1-p}{p} = \left.\frac{\mathrm{d}\alpha}{\mathrm{d}\Gamma}\right|_{\pi=0} = \left.\frac{\mathrm{d}\alpha}{\mathrm{d}\Gamma}\right|_{\mathrm{d}v=0} = -\left.\frac{\mathrm{d}x_2}{\mathrm{d}x_1}\right|_{\mathrm{d}U=0} = \frac{(1-p)u'(w-\Gamma)}{pu'(w-\ell+\alpha)}.$$

Since the agent is strictly risk averse, the latter implies that $\Gamma = \ell - \alpha$, which means her gross coverage is $\alpha + \Gamma = \ell$, so she insures fully. Then, from the zero-profit condition,

$$0 = (1-p)\Gamma - p(\ell - \Gamma) = \Gamma - p\ell,$$

which means that the premium is *actuarially fair*. This solution is depicted in Fig. 9.1

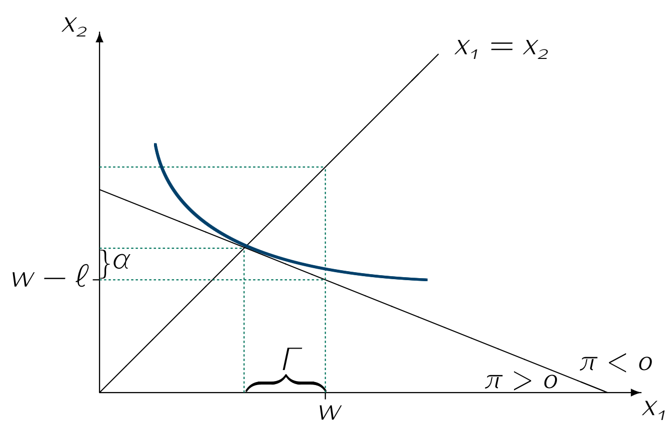

Figure 9.1: Efficient insurance under perfect information.

9.2. Imperfect information

For the case that is of interest, suppose now that there are two types of agents, H and L, and that the agents of type H face a higher risk of experiencing the loss, so that $p_H > p_L$. Graphically,

[2] The gross coverage would be $\alpha + \Gamma$.

this means that at any point in the space of consumption plans, type-H individuals have flatter indifference curves than those of type L.[3] We assume that the proportion of high-risk individuals is $\lambda \in (0, 1)$, and write $\bar{p} = \lambda p^H + (1 - \lambda) p^L$.

In this setting, the definition of (Nash) equilibrium is, intuitively, as follows: an equilibrium is a set of insurance contracts such that, when all consumers choose their optimal contract, no contract makes losses for the insurance companies, and there exists no alternative contract such that, if offered, it would give an insurance company a positive profit, given the consumers that would optimally demand it.

Definition 1 (Pooling equilibrium). *A single contract* (α, Γ) *constitutes a* pooling equilibrium *if:*

1. *low-risk individuals take it:* $V^L(0, 0) \leq V^L(\alpha, \Gamma)$;

2. *high-risk individuals take it:* $V^H(0, 0) \leq V^H(\alpha, \Gamma)$;

3. *insurance profits are zero:* $\bar{\pi} = (1 - \bar{p})\Gamma - \bar{p}\alpha = 0$; *and*

4. *no insurance company finds a profitable deviation: there exists no* (α', Γ') *such that*

$$V^L(\alpha, \Gamma) \leq V^L(\alpha', \Gamma'),\ V^H(\alpha, \Gamma) > V^H(\alpha', \Gamma'),\ \ \textit{and}\ \ (1 - p^L)\Gamma' - p^L\alpha' > 0.$$

In particular, the last condition implies that no insurance company can deviate by offering an alternative contract that attracts the low-risk individuals only, and with which, given that, it would obtain positive profits. An alternative type of equilibrium is the following.

Definition 2 (Separating equilibrium). *A pair of contracts* $\{(\alpha^L, \Gamma^L), (\alpha^H, \Gamma^H)\}$ *is a* separating equilibrium *if:*

1. *low-risk individuals take one contract:*

$$\max\left\{V^L(0, 0), V^L(\alpha^H, \Gamma^H)\right\} \leq V^L(\alpha^L, \Gamma^L);$$

2. *high-risk individuals take the other contract:*

$$\max\left\{V^H(0, 0), V^H(\alpha^L, \Gamma^L)\right\} \leq V^H(\alpha^H, \Gamma^H);$$

3. *insurance profits are zero in both contracts:*

$$\pi^L = (1 - p^L)\Gamma^L - p^L\alpha^L = 0\ \ \textit{and}\ \ \pi^H = (1 - p^H)\Gamma^H - p^H\alpha^H = 0;$$

4. *no insurance company finds a profitable deviation: there exists no* (α', Γ') *such that*

$$V^L(\alpha^L, \Gamma^L) \leq V^L(\alpha', \Gamma'),\ V^H(\alpha^H, \Gamma^H) \leq V^H(\alpha', \Gamma'),\ \ \textit{and}\ \ (1 - \bar{p})\Gamma' - \bar{p}\alpha' > 0.$$

[3] Importantly, note that we are assuming that the two types differ *only* in their probabilities of realizing the loss, and not in any other parameters, nor in their Bernoulli indexes.

9.3. Impossibility of pooling equilibrium

In the case of a pooling equilibrium, the iso-profit line has a slope $(1-\bar{p})/\bar{p}$. By construction, $p^H > \bar{p} > p^L$, which implies that

$$\frac{1-p^H}{p^H} < \frac{1-\bar{p}}{\bar{p}} < \frac{1-p^L}{p^L}.$$

Recalling our observation of Section 9.1, this means that over the full insurance line, the indifference curve of a low-risk individual is flatter than the isoprofit line, which is itself flatter than the indifference curve of a high-risk consumer.[4]

Fig. 9.2 depicts a potential pooling equilibrium, at the contract that yields point A. That insurance firms make null profits follows from the fact that the contract lies on the isoprofit line, and both individuals prefer point A to their endowments.[5]

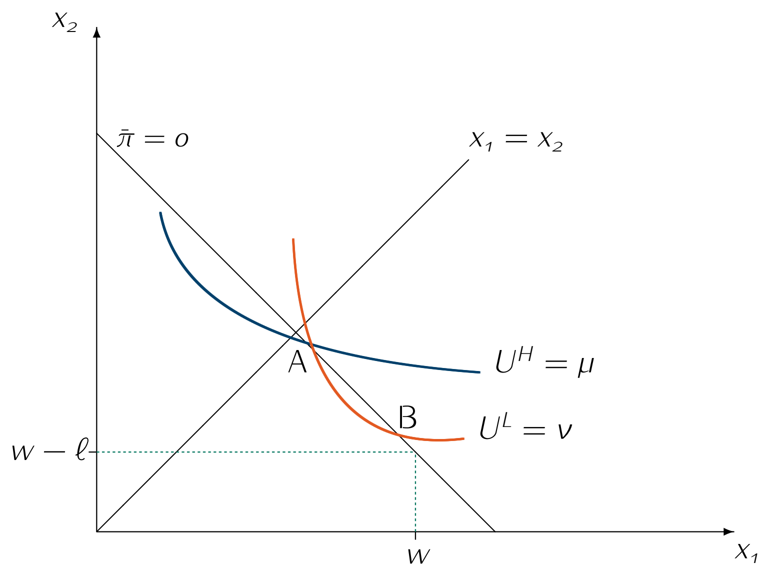

Figure 9.2: A candidate to pooling equilibrium.

Proposition 9.1 (Impossibility of pooling). *There can be no pooling equilibrium in this economy.*

Proof. Consider now Fig. 9.3, where we have added the isoprofit line for

$$\pi' = (1-p^L)\Gamma' - p^L\alpha' = 0.$$

This line is steeper, since $\bar{p} > p^L$. Now consider a contract strictly within the area formed by the lens ADC. It is immediate that every low-risk individual would prefer this, while all high-risk individuals would prefer to remain at A. This implies that the isoprofit line for π' would be the relevant one for this contract, and the firms can choose the contract so as to make strictly positive profits. □

[4] Since we are assuming that the consumers are risk averse, the latter observation will hold true everywhere

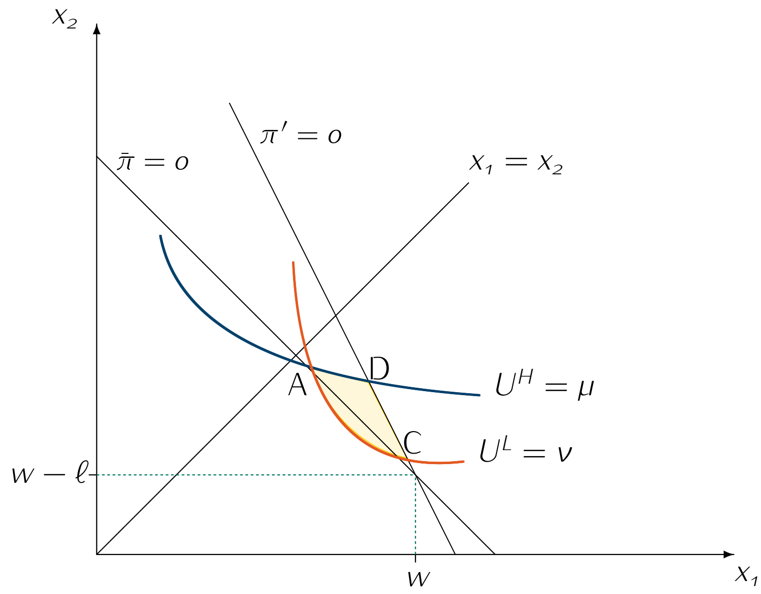

Figure 9.3: There is no pooling equilibrium.

9.4. Does a separating equilibrium exist?

If an equilibrium exists, it must then be one that separates the two types of consumers. Fig. 9.4 depicts a potential separating equilibrium, where point A corresponds to the contract that high-risk individuals take, and point B to the one for low-risk consumers.

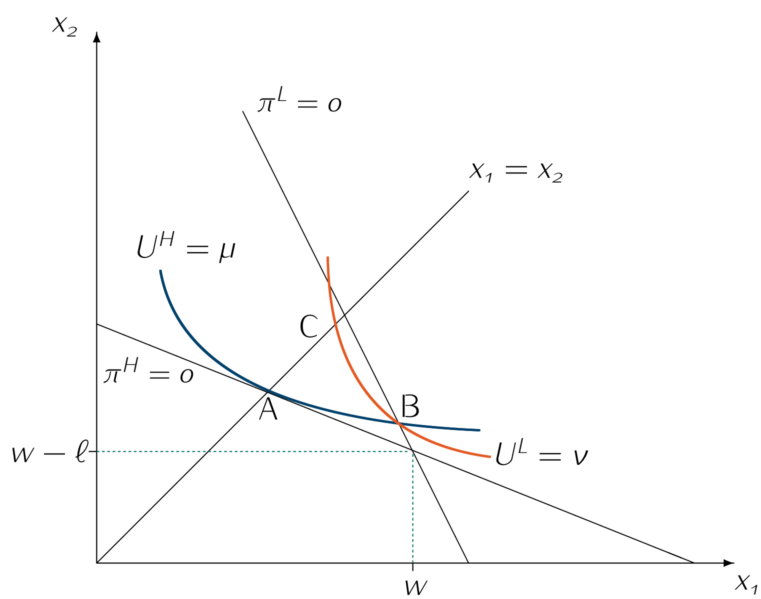

Figure 9.4: A candidate to separating equilibrium

below the full insurance line: the indifference curve gets even flatter there. The former observation, for the same reason, need not be true everywhere, but it certainly holds for plans sufficiently close to the full insurance line.

[5] It is useful to note that point B *cannot* be an equilibrium, as insurance companies could offer an alternative contract that all consumers would prefer to take and, still, would give them strictly positive profits.

If each individual takes the contract that corresponds to her type, it is immediate that both contracts break even. It is also immediate that type-L individuals do take their contract, as they are better off at B than at their endowments and at A. The analogous analysis is true for high-risk individuals, and they too take their contract, at A. Importantly, though, these two contracts result in qualitatively different insurance profiles: while type-H individuals are fully insured at A, those of low risk only get partial insurance. The asymmetry is explained, precisely, by the fact that the insurance company cannot identify the type of each individual, so it needs to use the design of the contracts to elicit this information. If an insurance company offered full insurance to type-L agents, so as at point C, *all* individuals would take this contract, and the company would make losses. The same is not true for the full-insurance contract of risky individuals, as no agent of type L prefers A to B.

Now, for this pair of contracts to constitute an equilibrium, it must also be true that no insurance company can deviate and make higher profits offering an alternative contract. In particular, no contract that would attract both types of consumers can give an insurance company strictly positive profits. Fig. 9.5 depicts the isoprofit consistent with

$$\bar{\pi} = (1 - \bar{p})\Gamma - \bar{p}\alpha = 0.$$

This line lies between the other two isoprofits, and its position depends on the value of λ. In the case of Fig. 9.5, no contract that attracts high risk individuals and gives the insurance company non-negative profits would be taken by low risk individuals, so the pair of contracts is, indeed, a separating equilibrium.

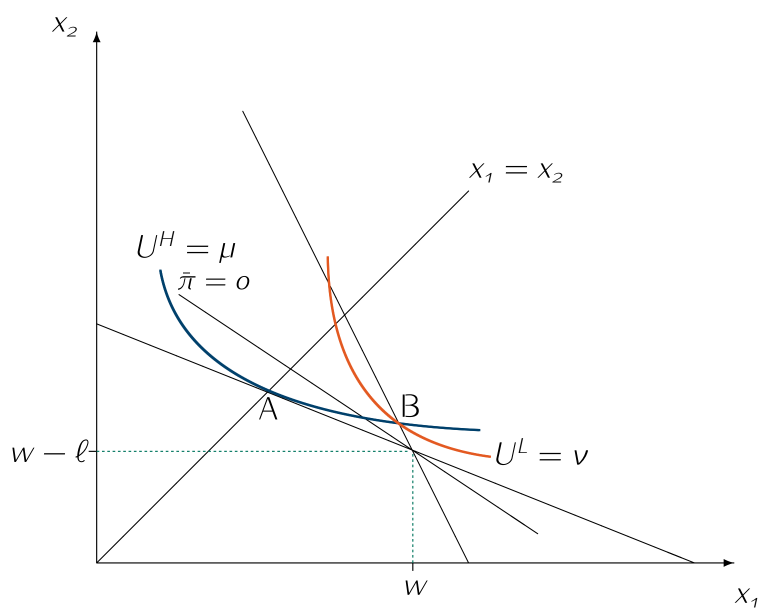

Figure 9.5: A separating equilibrium

Importantly, the situation of Fig. 9.5, does not always occur, and imperfect information may induce non-existence of separating equilibria too.

Proposition 9.2. *For λ close enough to 0, there exists no separating equilibrium.*

Proof. For lower values of λ, the isoprofit for $\bar{\pi} = 0$ must be rotated clockwise. It is clear from Fig. 9.5, that if we rotate this line enough, we can find contracts that would be attractive to both types of individuals and would make zero profits, as in Fig. 9.6. A contract close enough to it would still be preferred by all consumers, and can be chosen to give strictly positive profits to the insurance company. □

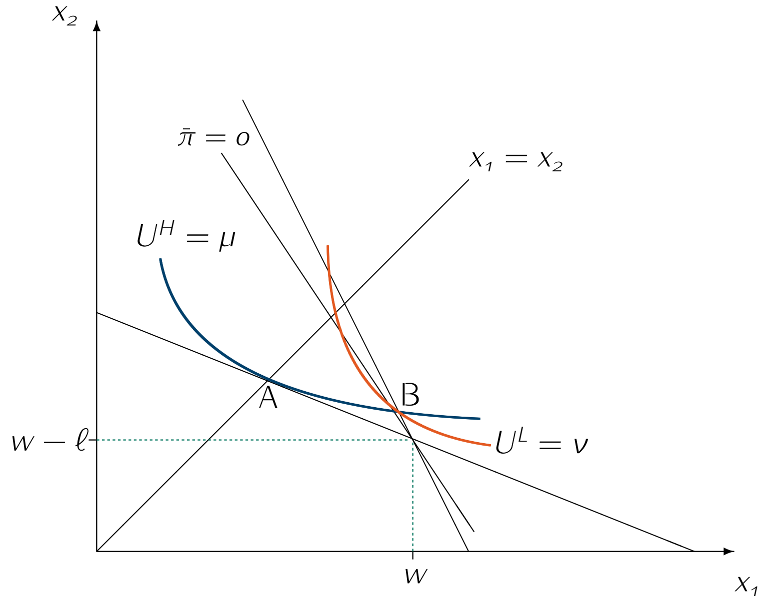

Figure 9.6: No separating equilibrium

Example 9.1. In this canonical insurance problem, suppose that ex-ante preferences over consumption plans are

$$U^\tau(x) = (1 - p_\tau)\ln x_1 + p_\tau \ln x_2,$$

for agents of type $\tau = H, L$.

a. Argue that if the type of each individual was observable, then two insurance contracts would be offered, (Γ_H, α_H) and (Γ_L, α_H), with premium $\Gamma_\tau = p_\tau \ell$ and coverage $\alpha_\tau = (1 - p_\tau)\ell$.

b. Now, suppose that the type of each individual is unobservable, but the market reaches a separating equilibrium. With respect to this equilibrium,

   (i) Argue that the contract for individuals of type H is still the one described in part 1.

   (ii) Argue that the contract for individuals of type L satisfies

$$(1 - p_L)\Gamma_L - p_L \alpha_L = 0$$

   and

$$\ln(w - p_H \ell) = (1 - p_H)\ln(w - \Gamma_L) + p_H \ln(w - \ell + \alpha_L).$$

Answer: We saw the argument in class:

a. Under perfect observation and competition, at equilibrium there is perfect insurance and firms make zero profits. This implies that $(\Gamma_\tau, \alpha_\tau)$ is the solution to the equations:

$$w - \Gamma_\tau = w - \ell + \alpha_\tau$$

and

$$(1 - p_\tau)\Gamma_\tau - p_\tau \alpha_\tau = 0.$$

b. Again, we saw the argument in class: under asymmetric information, H-types get their competitive contract, while L-types have their coverage limited by the incentive compatibility constraints of the H individuals.

If an H individual gets the contract designed for her, her utility is

$$(1 - p_H)\ln(w - \Gamma_H) + p_H \ln(w - \ell + \alpha_H)$$

which is simply $\ln(w - \ell + \alpha_H)$, as this contract offer full insurance. If she takes the contract designed for L-types her utility is

$$(1 - p_H)\ln(w - \Gamma_L) + p_H \ln(w - \ell + \alpha_L).$$

The incentive compatibility constraint is that the latter number cannot be larger than the former. As this constraint binds, Eq. $(**)$ must hold true at equilibrium. Eq. $(*)$ is guaranteed by competition between insurance companies, as each of the equilibrium contracts must make null profits. □

Exercise 9.1. In this canonical insurance problem, suppose that there are three types of agents, H, M and L, distributed in proportions λ_H, λ_M and λ_L. Agents of type H have the highest probability of suffering the loss, $p_H > p_M$, while those of type L have the lowest one, $p_L < p_M$. Ex-ante preferences over consumption plans are

$$U^\tau(x) = (1 - p_\tau)u(x_1) + p_\tau u(x_2),$$

for agents of type τ, where u is an increasing Bernoulli index displaying strict risk aversion.

a. For the case of a pooling contract, write conditions under which individuals of all types take this contract and the insurance company breaks even.

b. For the case of separating contracts, where there is a different insurance contract for each type of costumer, write conditions under which individuals of type M take the contract designed for them, and the insurance company breaks even in this contract.

c. Consider now the case of "partially separating" contracts, where there is one contract intended for individuals of type L, and a different contract for those whose type is either M or H. Write conditions under which individuals of the latter types choose the contract intended for them, and the insurance company breaks even in this contract.

EXERCISE 9.2. A king is worried about the welfare of the elderly in his kingdom and wants to introduce a policy that guarantees some income levels when people retire. There are three types of individuals: A who are the aristocrats; B, who are the bourgeoisie; and C, the countrymen. Their proportions in the kingdom are λ_A, λ_B and λ_C. For each person of type τ, her income when old, in the absence of a pension plan, is the random variable W_τ.

A pension plan consists of a contribution made when the person is young, Γ, and the intended guarantee that the person's income when old will be a given random variable α. In order to do this, the king's policy offers a pension plan for each type, $(\Gamma_\tau, \alpha_\tau)$, for $\tau = A, B, C$. Under this plan, the person contributes Γ_τ when young, and receives the shortfall $\alpha_\tau - W_\tau$ when old.

All individuals have the same preferences over present and future income: if they have x_o when young, and a random income X when old, their utility is $x_o + \mathrm{E}[u(X)]$, where u is an increasing and strictly concave cardinal utility index.

a. If the king does not want to force any of his subjects to take a pension plan, what condition has to be satisfied so that each individual in the kingdom voluntary participates of the policy by signing up to *some* plan.

b. Suppose now that, moreover, the king intends to let everyone choose freely which plan they take, but really wants that people of type C get the income A_C when old. What conditions must be satisfied for that to be the case? If the king also wanted to make sure that nobody else takes the plan intended for these people, what other conditions must be obeyed?

c. Suppose finally that, furthermore, in order to avoid a revolution, the pension plan intended for individuals of types A and B must be identical and that the king also wants to make sure that in expectations the sum of these two pension plans breaks even. What further condition must be satisfied?

READING. The following excerpt is from an article in *The Economist*. Discuss it critically.

> *The idea behind adverse selection is that because [the] people being insured have more information than their insurers about the likelihood of an adverse event, you can't get a market going; only those who expect to lose more than the premium will insure themselves. Which means the insurance company will lose money. Which means it will raise the premiums. Which means the people who expect to lose less than the new, higher premium will drop their coverage. Which means the average loss per insured person will go up. Which means the insurance company will lose money. Which means it will raise the premiums [sic]...*
>
> *But these things are true of any insurance market. You know much better than State Farm how often your wife forgets to lock the door, whether the guy in the next townhouse likes to have a cigarette in bed with his nightly Nembutal-and-Bombay-Sapphire toddy, and how close*

your rottweiler comes to taking off the postman's leg. Since you have a better shot of estimating the probability of events that will require them to pay out on your homeowner's insurance, in theory, this market should not exist. In fact, it does, because people are very risk averse, and also, not so good at calculating actuarial risk. As long as America's public schools continue their appalling record on math education, adverse selection shouldn't be a huge problem.

The real problem is not that people have some sort of excellent secret knowlege about their health that will produce adverse selection; the problem is that some people can't afford to pay the cost of medical care for diseases that have already occurred. This is no more nor less of an issue than the fact that some people cannot afford to replace the contents of their home after it burns down.

That problem is exacerbated by the lunatic structure of America's insurance market, in which most people get their insurance through their employers; that means that people are often thrown out on the insurance market against their will. But it is not a "market failure"; it's hard to think of any market failing worse than one in which an insurance company would write you a policy for something that had already happened.

The problem, then, is whether the government should pay the costs of those who have these sorts of health problems. The answer of those who argue adverse selection is "Yes".

From: Free exchange, *The Economist*, Economics (February 7th, 2007) "It's adverse, but is it selection?"

References: This material is based on Rothschild, M. and J. Stiglitz, "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information." *The Quarterly Journal of Economics* (1976): 629-649.

Note 10

Labor markets

THE SECOND CANONICAL EXAMPLE IS THE EMPLOYEE–EMPLOYER RELATION. We now consider the extent to which informational asymmetries can distort the efficient allocation of resources in a market. Consider, for instance, the case of a bank that is hiring an investment analyst. Suppose that, due to her skills or training, the candidate's marginal productivity is $\vartheta \in [0, \vartheta^*]$, and that as a function of this productivity, her reservation salary is $r(\vartheta)$, for a function with $r' \geq 0$.[1]

10.1. Symmetric information

TO ESTABLISH A BENCHMARK, suppose that the candidate's productivity, ϑ, is observable by the bank. Letting w^* denote the salary offered by the bank, we get that for it to be optimal for the bank, $w^*(\vartheta) = \vartheta$. The candidate will accept an offer if, and only if, $w^*(\vartheta) \geq r(\vartheta)$.

It follows that the bank will hire the candidate if, and only if, $\vartheta \geq r(\vartheta)$, as in Fig. 10.1. This solution is Pareto efficient.

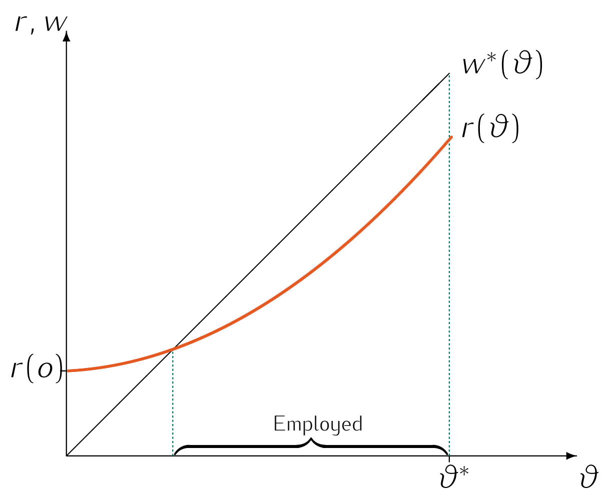

Figure 10.1: Efficient employment under perfect information.

[1] Assume that the candidate's ϑ is drawn from a population such that the distribution of Θ is continuous and has as support the whole set $[0, \vartheta^*]$.

10.2. Asymmetric information: adverse selection

Suppose now that the bank cannot observe the candidate's productivity,[2] while the candidate herself can. Information is now asymmetric, and each agent can condition her actions only on the information they have. While the candidate's decision whether to accept an offer or not will be contingent on her productivity, via $r(\vartheta)$, the offer made by the bank has to be a constant wage, w.

Of course, if the bank is sophisticated, which we assume, the bank will learn something from the decision of the candidate whether to accept an offer or not. An unsophisticated bank's expectation of a candidate's productivity is $\mathrm{E}(\Theta)$. For a sophisticated bank, this is the expectation only at the beginning of the candidate's interview. This type of bank will want to make an offer that they will not regret *if it is accepted by the candidate*. Anticipating that the candidate will accept an offer w if, and only if, $w \geq r(\vartheta)$, the bank's expectation of the candidate's productivity *if she accepts the offer* is $\mathrm{E}[\Theta \mid r(\Theta) \leq w]$. Assuming that the bank is risk neutral, the offer made to the candidate is the solution to

$$\mathrm{E}[\Theta \mid r(\Theta) \leq w] = w.$$

Before continuing to determine the equilibrium of this market, it is useful to understand the behavior of the conditional expectation as a function of w. Assuming that $r' > 0$, function r is invertible and we can write

$$\mathrm{E}[\Theta \mid r(\Theta) \leq w] = \mathrm{E}[\Theta \mid \Theta \leq r^{-1}(w)]$$

for any $w \geq w_* = r(0)$. Since r is increasing, so is its inverse and, under simple assumptions, this function too will be increasing. The exact shape of the function, however, is not easy to pin down, so we will later have to consider different cases. For illustration, consider Fig. 10.2, where a possible function is drawn. Importantly, this function is bounded above at $\mathrm{E}(\vartheta)$, and this bound is reached precisely at a salary $r(\vartheta^*)$.

In this setting, whether the market attains an efficient allocation is undetermined, as this depends on the shape of the bank's conditional expectation function. Fig. 10.3 exhibits a situation where an equilibrium is determined. The equilibrium wage is w_e, which is optimal for the bank and results in a candidate working if her productivity is at ϑ_e or below, and not working otherwise. Importantly, this result is *not* efficient: unlike in the benchmark case, if the candidate's productivity is high, $\vartheta > \vartheta_e$, she is not employed. This is inefficient, as her productivity is, still, above her reservation salary.

This type of inefficiency is known as *adverse selection*, and is due entirely to the bank's impossibility to observe the candidate's productivity.

Of course, a different shape of the conditional expectation function could lead to different qualitative conclusions, in particular in regards to the severity of the adverse selection effect. Fig. 10.4 presents an extreme case of adverse selection, where, due to the bank's lack of information, only the least productive candidate is ever hired. Fig. 10.5 exhibits the other extreme, where the bank makes an offer that is accepted by candidates of all productivities.[3]

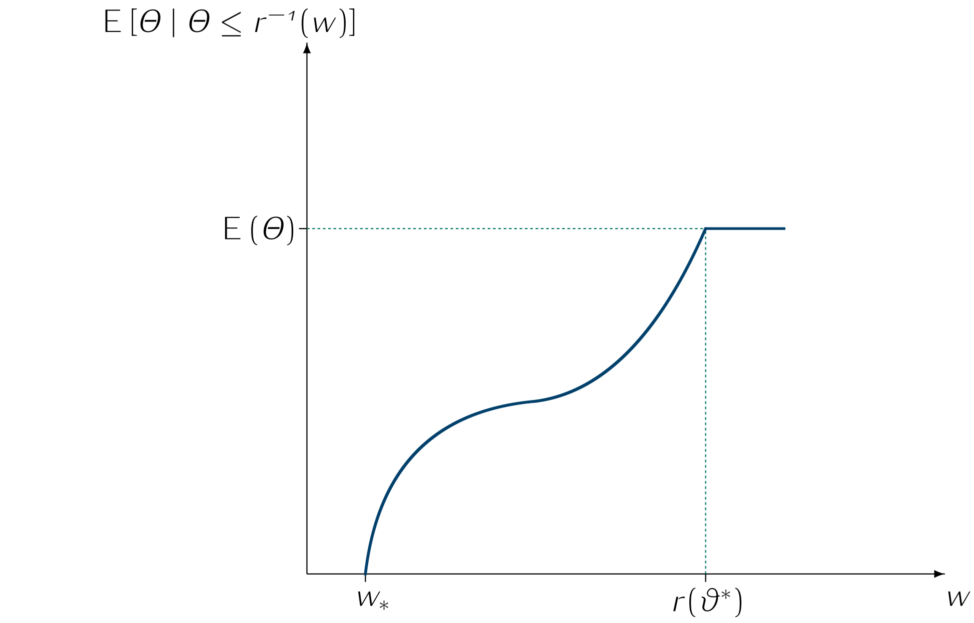

Figure 10.2: Expected productivity conditional on wage acceptance.

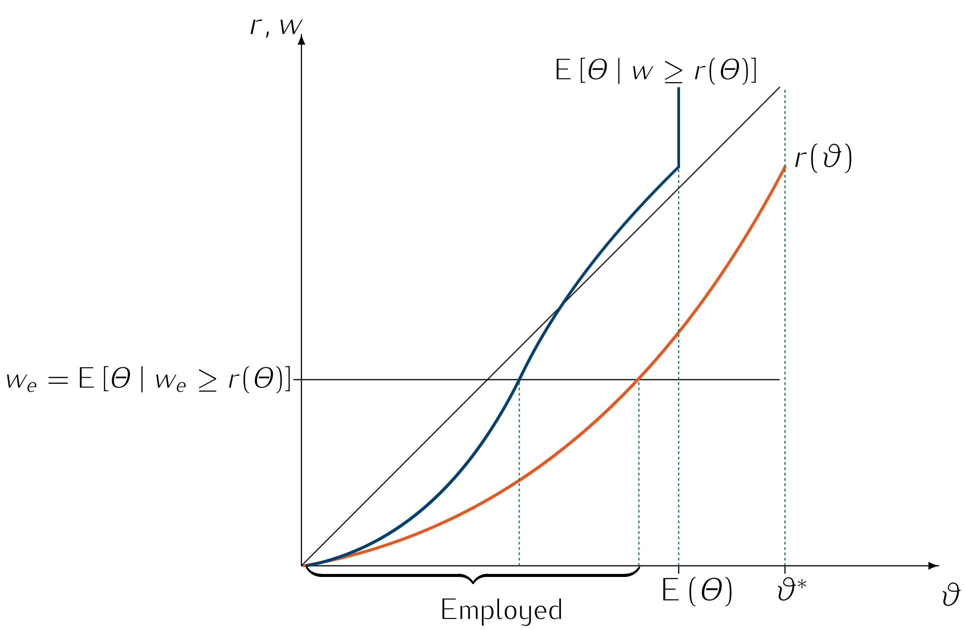

Figure 10.3: Adverse selection: inefficient employment under imperfect information.

Finally, consider Fig. 10.6, where three salary levels are potential equilibria. Since we are assuming that the bank knows the distribution of productivities in the population, an argument is often made that only w_e is an actual equilibrium, since other offers that satisfy the equality $\mathrm{E}\left[\vartheta \mid r(\vartheta) \geq w\right] = w$ need not maximize the bank's expected surplus *globally*.

[2] The bank knows everything else, including the distribution of Θ.

[3] Even in this case, the functioning of markets is not perfect: low productivity candidates are overpaid, and high productivity ones underpaid.

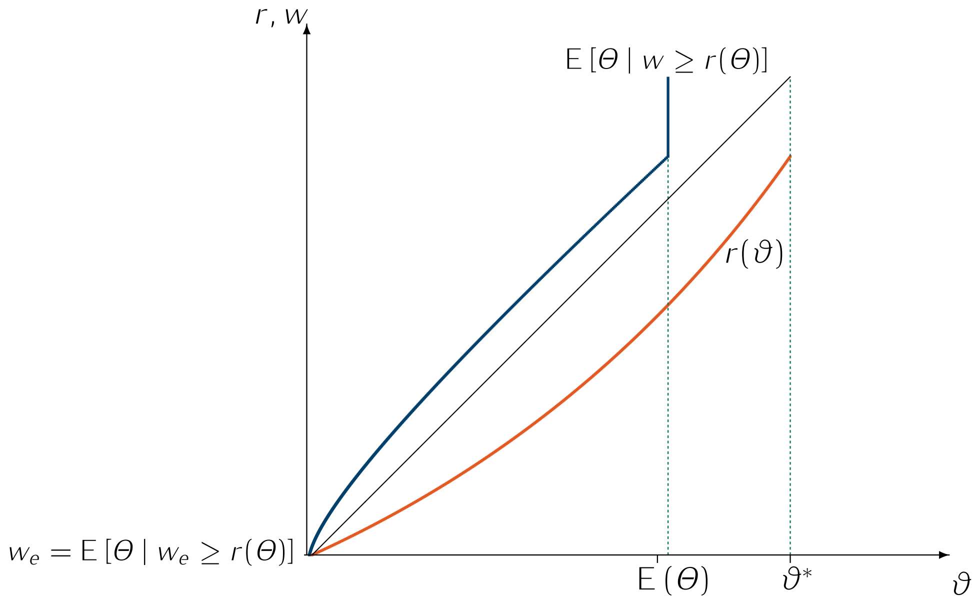

Figure 10.4: Adverse selection: full unemployment under imperfect information.

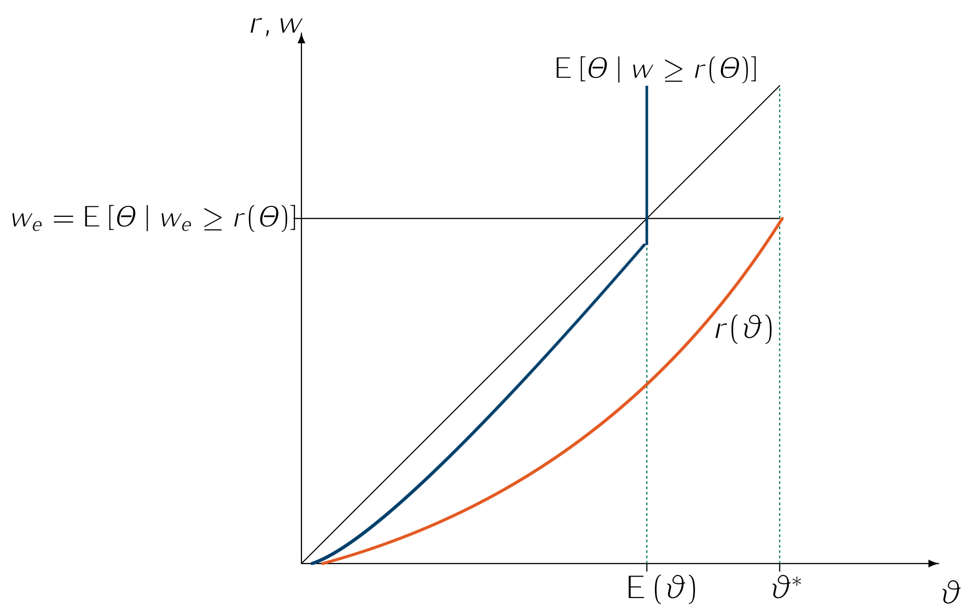

Figure 10.5: Adverse selection: distorted remuneration under imperfect information.

Example 10.1. In the canonical problem of a labor market with asymmetric information, where a firm is trying to hire a new employee, potential employees come from a population where marginal productivity, ϑ, is heterogeneous. In the population, there are equal fractions of potential employees with $\vartheta = 0, 1, 2, 3, 4, 5$. The opportunity cost of a candidate with productivity ϑ is $r(\vartheta)$, with

$$r(0) = 0;\ r(1) = 0.6;\ r(2) = 0.9;\ r(3) = 1.5;\ r(4) = 2.5;\ r(5) = 4.$$

The employee knows her productivity, but cannot make it observable to the firm.

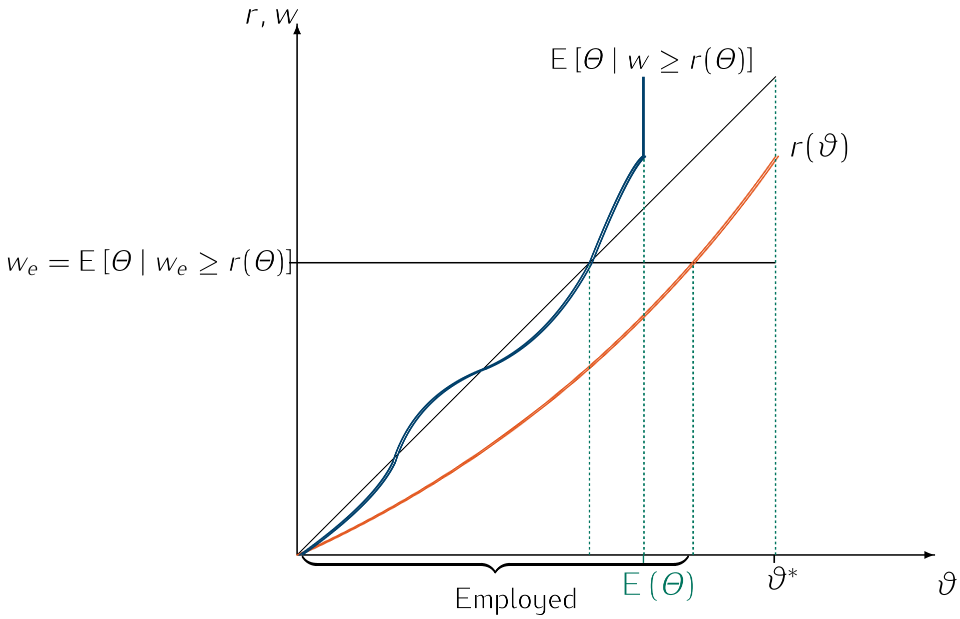

Figure 10.6: Adverse selection: inefficient employment under imperfect information.

a. Denote by $\bar{\vartheta}(w)$ the productivity of the most productive worker who would accept an offer with wage w. Find $\bar{\vartheta}(w)$ for $w = 0, 1/2, 1, 3/2, 2, 3, 4, 5$.

b. For the same values of w, compute $\mathrm{E}\left[\Theta \mid \text{The candidate accepts } w\right]$.

c. Find the equilibrium wage and argue that both inefficient unemployment and distorted remuneration occur at equilibrium.

Answer: a. By individual rationality of the workers, $\bar{\vartheta}(w)$ is the largest value of ϑ for which $r(\vartheta) \leq w$. For example, $r(2) = 0.9 \leq 1 < 1.5 = r(3)$, so $\bar{\vartheta}(1) = 2$. Similarly,

w	0	$1/2$	1	$3/2$	2	3	4	5
$\bar{\vartheta}(w)$	0	0	2	3	3	4	5	5

b. We need to find $\mathrm{E}\left[\Theta \mid \Theta \leq \bar{\vartheta}(w)\right]$, using the distribution of productivity in the population. For instance,

$$\begin{aligned}\mathrm{E}\left[\Theta \mid \Theta \leq \bar{\vartheta}(1)\right] &= \mathrm{E}\left[\Theta \mid \Theta \leq 2\right]\\ &= \Pr(\Theta = 0 \mid \Theta \leq 2) \times 0 + \Pr(\Theta = 1 \mid \Theta \leq 2) \times 1 + \Pr(\Theta = 2 \mid \Theta \leq 2) \times 2\\ &= \tfrac{1}{3} \times 0 + \tfrac{1}{3} \times 1 + \tfrac{1}{3} \times 2\end{aligned}$$

Similarly,

w	0	$1/2$	1	$3/2$	2	3	4	5
$\mathrm{E}\left[\Theta \mid \text{The candidate accepts } w\right]$	0	0	1	$3/2$	$3/2$	2	$5/2$	$5/2$

c. Equilibrium is attained when

$$w = \mathrm{E}\left[\Theta \mid \text{The candidate accepts } w\right].$$

This occurs at $w = 1.5$. At this wage only candidates with $\vartheta \leq 3$ take the job, which leads to inefficient unemployment: people with $\vartheta = 4$ and $\vartheta = 5$, who reject the offer, would be more productive working in the firm than at their alternative. There is also distorted remuneration, as workers with $\vartheta = 0$ and $\vartheta = 1$ are overpaid, while those with $\vartheta = 2$ and $\vartheta = 3$ are underpaid.

A second equilibrium occurs at $w = 1$. □

Exercise 10.1. Consider the case of a firm that is hiring a new manager. The candidate's skills measured by $\sigma \in [\sigma_*, \sigma^*]$. The return of the company, R, is a random variable; under a manager of skill σ, the distribution of R is F_σ. For $\sigma > \sigma'$, $F_\sigma >_1 F_{\sigma'}$.

On the population of candidates, skills are distributed according to L. There is asymmetric information, in that each candidate knows her own σ but the firm only knows L. All agents in this economy are risk neutral.

1. Suppose that the only alternative available to people who don't work for the company is to receive unemployment benefits b. Under what conditions will this firm want to hire a manager? And if such condition holds, will the manager's remuneration reflect her productivity?

2. Suppose now that, instead of unemployment benefits, the candidates have the option of running their own businesses. For an individual of skill σ, the return of her independent business is a random variable ρ, distributed according to G_σ. For $\sigma > \sigma'$, $G_\sigma >_1 G_{\sigma'}$. Write a condition under which, at equilibrium, the firm will hire low-skilled managers, while high-skilled managers choose to remain independent. Argue that this may occur even if the return of the firm under a high-skilled manager is higher than the return she obtains in her independent business.

Exercise 10.2. The problem of adverse selection appears not only on insurance and labor markets, but on any situation where information asymmetries may be salient. The following is an important case: Each firm in an industry is evaluating an indivisible investment project to expand its output. If a firm does not invest, its output is $z > 0$. The investment opportunity is risky: there is a probability that the total output is $x > z$; but the project can also fail, in which case the firm's total output is 0. There are two types of firm, L and H, which differ only in their probabilities of being successful: $p_L < p_H$. The proportion of L firms is $\lambda \in (0, 1)$.

The investment project costs $k > 0$, but it is profitable for both types, in the sense that $p_\tau x - k > z$ for both $\tau = L, H$. Each firm is risk-neutral, and its total available cash is $c < k$. If it chooses to invest in the project, the difference $k - c$ has to be borrowed from (risk-neutral) banks, at an interest rate ι.

1. Find a condition that the interest rate must satisfy for the firm to repay its loan when the output is positive.

2. Assuming perfect competition in financial markets, find the interest rate charged by the banks at a potential pooling equilibrium, and argue that this rate is accepted by all L firms.

3. At that interest rate will type-H firms invest? How does the value of λ affect the conditions under which these firms invest?

READING. The following excerpt is from an article in *The New Yorker*. Discuss it critically.

> *This is the Age of the Incredible Shrinking Everything. Home prices, the stock market, GDP, corporate profits, employment: they're all a fraction of what they once were. Yet amid this carnage there is one thing that, surprisingly, has continued to grow: the paycheck of the average worker. Companies are slashing payrolls: 3.6 million people have lost their jobs since the recession started, with half of those getting laid off in just the past three months. Yet average hourly wages jumped almost four per cent in the past year. It's harder and harder to find and keep a job, but if you've got one you may well be making more than you did twelve months ago.*
>
> *This combination of rising unemployment and higher wages seems improbable. But, as it turns out, it's what history would lead us to expect. Even during the early years of the Great Depression, manufacturing workers actually saw their real wages rise, and wage cuts have been scarce in every recession since. Oil and wheat prices may rise and fall instantaneously to reflect supply and demand, but wages are "sticky": even when the economy goes bad, it takes a lot to make them fall.*
>
> *It's not because businesses are generous that wages are sticky; it's because employers are worried. In part, bosses are afraid of what economists call "adverse selection": if they cut wages, it's the least productive workers who would be the most likely to stay, while the best workers would start looking elsewhere. (Even in a weak economy, businesses still compete for talent.) In a 1997 study of almost two hundred employers, the economists Carl Campbell and Kunal Kamlani found that the threat of losing their best employees was a major reason that bosses didn't cut wages.*

From: Surowiecki, J., "Nice work if you can get it", *The New Yorker*, March 2nd, 2009; Vol. 85 Issue 3, p23-23.

References: This material is based on Akerlof, G.A. "The market for 'lemons': Quality uncertainty and the market mechanism." *The Quarterly Journal of Economics* (1970): 488-500.

Note 11

Screening

THERE ARE WAYS IN WHICH ECONOMIC AGENTS CAN DEAL WITH ADVERSE SELECTION. We now study how, in the context of adverse selection in labor markets, an employer can mitigate the inefficiencies caused by her lack of information. In the same setup as the previous section, suppose for simplicity that there are only two types of workers according to their productivity: ϑ_H and ϑ_L, with $\vartheta_L < \vartheta_H$. Suppose that the low-productivity workers are a proportion λ of the population, that all the agents in the economy are risk neutral, and that the reservation salary of workers of both types are null.

The new dimension that we add now is the assumption that, while a person's productivity cannot be observed, there may exist another variable that can be used to discern it. Suppose, then, that each worker can choose a level of education, $e \geq 0$. To make the analysis interesting, let us assume[1] that a person's education has no effect in her productivity. Instead, suppose that the *cost* of acquiring education varies with the ability of the individual, which is measured by her productivity. Specifically, the marginal cost of an extra unit of education for an individual of type $\tau \in \{H, L\}$ is $c_\tau > 0$, with $c_H < c_L$. The critical assumption here is that highly productive individuals find it less costly to acquire education,[2] so the bank can use the level of education of a potential manager to *screen* for her unobservable productivity.

11.1. Symmetric information

TO ESTABLISH A BENCHMARK, let us first consider the case where the individual's productivity is perfectly observable. In this case, it is easy to see that, since education does not add to a candidate's productivity, the efficient level of education would be $e = 0$ for all individuals, and each worker would receive as salary her (observed) productivity.

11.2. Asymmetric information: screening

THE CASE THAT INSTERESTS US, on the other hand, requires that if a differentiated salary is going to be offered, the separation of workers be made using information extracted endogenously, in

[1] Unrealistically, I hope.

[2] Hmmm...

our case the person's level of education.

As in the problem of insurance contracts, we consider two types of equilibria: one where every worker is paid the same salary, and one where remuneration is differentiated.

11.2.1. Pooling equilibrium

If the bank is to offer only one contract to every candidate, it will find it optimal to pay as salary the expected productivity of the total population,

$$\bar{w} = \bar{\vartheta} = \lambda\vartheta_L + (1-\lambda)\vartheta_H.$$

Since education is costly and not remunerated, both type of workers choose null levels of education, $e_L = e_H = 0$.

To determine whether this situation is an equilibrium, we must study whether the bank would find it profitable to offer a contract that attracts highly productive individuals *only*. In order to do this, a bank can offer a contract $(\tilde{w}, \tilde{e})$, estipulating that the salary $\tilde{w}$ is paid only to an individual with education level at $\tilde{e}$ or above. If this contract is to be successful, it must satisfy that

$$\tilde{w} - c_H\tilde{e} > \bar{w}, \tag{11.1}$$

so that highly productive individuals will accept it, while

$$\tilde{w} - c_L\tilde{e} \leq \bar{w}, \tag{11.2}$$

to guarantee that individuals of low productivity have no incentives to take it.

If we assume that the firm will choose the lowest level of salary that guarantees Eq. (11.2), we have that

$$\tilde{e} = \frac{\tilde{w} - \bar{\vartheta}}{c_L},$$

which we can substitute in Eq. (11.1) to find that the required salary must satisfy that

$$\tilde{w}\left(1 - \frac{c_H}{c_L}\right) > \bar{\vartheta}\left(1 - \frac{c_H}{c_L}\right).$$

Since $c_L > c_H$, the latter amounts, simply to $\tilde{w} > \bar{\vartheta}$.

It follows, then, that no pooling equilibrium can exist: with only a marginally higher salary and an appropriately chosen required level of education, a separating contract would be able to attract highly productive employees only.

11.2.2. Separating equilibrium

Suppose now that the bank offers a pair of contracts that remunerate a person at one of the two possible productivity levels, depending on her level of education. Specifically, the bank chooses a level e_* of education, and offers the candidates the following salary:

$$w = \begin{cases} \vartheta_L, & \text{if } e < e_*, \\ \vartheta_H, & \text{otherwise.} \end{cases}$$

A contract like this is known as a *screening* contract. For it to succeed, it must be that each worker chooses the level of education, and hence the salary, intended for her type. This entails two conditions:

1. That no highly productive individual prefer to choose the low-education salary. If a person were to do this, her optimal level of education would be null, so the condition that is required is simply that
$$\vartheta_H - c_H e_* \geq \vartheta_L,$$
2. That no candidate with low productivity prefer to attain enough education so as to receive as salary the higher level of productivity. This requires that
$$\vartheta_L \geq \vartheta_H - c_L e_*,$$
since the optimal level of education would be null for a person who chooses as salary the lower productivity level.

These two conditions are known as *self-selection constraints*, as they constitute the requirements for the bank to be able to extract from the workers the revelation of their types.

Suppose that the bank fixes the lowest level of education under which these two conditions hold.[3] The two inequalities restrict the level of education so that
$$\frac{\vartheta_H - \vartheta_L}{c_H} \geq e_* \geq \frac{\vartheta_H - \vartheta_L}{c_L}.$$
which is feasible since $\vartheta_H - \vartheta_L > 0$ and $c_H < c_L$. The resulting level of education is, hence,
$$e_* = \frac{\vartheta_H - \vartheta_L}{c_L}.$$

Since education is costly, low-productivity candidates will choose $e_L = 0$, while those with the higher level of productivity will acquire $e_H = e_*$, so that they can access the higher salary at minimum cost.

We still need to determine whether this pair of contracts constitutes an equilibrium, by checking whether it is robust to deviations by the bank to a pooling contract. Before that, it is important to note that, even if this arrangement effectively gets the workers to reveal, through their choices, their private information, adverse selection exists and the outcome is not efficient: while lowly-productive employees will chose no education and will be remunerated at their productivity, those with high productivity have to pay an inefficient cost to access as salary their own productivity. This expenditure is inefficient since education has no value, by assumption of the model.

Now, it only remains to check whether banks would find it profitable to deviate from this pair of contracts, by offering a contract that pools all workers together. If a bank were to do so, it should offer as salary $\bar{\vartheta}$, while requiring no education from the candidate. In this setting, nobody choosing this contract would acquire any education, but for the deviation to be profitable it would need to attract candidates of high productivity, which means that
$$\lambda\vartheta_L + (1-\lambda)\vartheta_H = \bar{\vartheta} > \vartheta_H - c_H e_*.$$
Substituting directly, this is only possible if
$$\lambda < \frac{c_H}{c_L}. \qquad (*)$$

[3] This could be the result, for instance, of competition between banks, *à la* Bertrand.

If Eq. ($*$) fails, the pair of contracts constitutes an equilibrium. If it holds, there can be no equilibrium, since a pooling contract would upset the contracts that would form a separating equilibrium.[4] This occurs when the proportion of low-productivity individuals is small (so that the benefit of separating the two types is not very significant) and/or when the ratio c_H/c_L is large (in which case the separation is very costly).

Example 11.1. In the context of Exercise 10.1, suppose that the company hires a consultant to improve their recruitment results. The consultant tells them the following:

> *The solution here is what economists call "screening". Since more skilled individuals are more productive both working for the company and working independently, we can suggest to them a voluntary application fee, ϕ. Then, we offer them two wages: w^*, to those who contribute the fee; and $w_* < w^*$ to those who don't. With a correctly chosen ϕ, only highly productive candidates will pay it an receive the higher salary. And I will be happy to be paid ϕ for my services.*

Will this screening mechanism work? Why or why not?

Answer: No, this mechanism doesn't work. To be successful, screening requires that the action used to screen the candidates be more costly to take for individuals of lower type. This is the case of education, when education achievement is harder (more costly) to individuals of low ability, as the mechanism may be designed so that only high ability candidates choose to attain the required education. Here, the fee is equally costly for all candidates, regardless of their skill type; if a high-skill individual is willing to pay ϕ, then so will be a low-skill candidate and the mechanism will fail to screen. □

Example 11.2. A bank is trying to hire a financial analyst from a population of economics students. The quantitative ability of a student, Θ, is distributed uniformly over the interval $[0, 1]$. For the bank, the marginal productivity of a hire is increasing in her quantitative ability: the marginal product of an individual with ability ϑ is, in fact, equal to ϑ.

The final exam of a given course is meaningless from the point of view of whether a person is a good financial analyst or not, but a student with higher quantitative abilities finds it easier to get a good grade in that exam. Specifically, for an individual of ability ϑ, the cost of getting a grade g is $(1-\vartheta)g$, for $g \in [0, 1]$. The individual's utility is given by the difference of between the wage she receives and the cost of the grade she attains.

a. If the students' quantitative abilities were observable, what would be the equilibrium wage schedule offered by the bank?

b. Suppose that the students' quantitative ability is unobservable, but the university is willing to reveal to the bank whether a student got a grade of at least γ in the particular exam, or

[4] And we already know that a separating contract *always* upsets a potential pooling equilibrium.

not. (This is known by the students when they enrol in the economics program.) The bank can then offer a (partial) separating contract: for values $w^* > w_*$,[5]

$$w(g) = \begin{cases} w^*, & \text{if } g \geq \gamma; \\ w_*, & \text{otherwise.} \end{cases}$$

Assume that $\gamma > 1/2$. With respect to this contract,

(i) Given w^* and w_*, find the ability level ϑ^* at which a student is indifferent between the two conditions in the wage contract.

(ii) Still taking w^* and w_* as given, argue that students with ability above ϑ^* strictly prefer the salary w^* with grades γ. What grades will students with ability below ϑ^* get?

(iii) Given ϑ^*, what should the salaries w^* and w_* be, if the bank is to break even in the contract?

(iv) To find a contract that could be a separating equilibrium, substitute the salaries of part (c) into the threshold of part (a). Solve for ϑ^*, w^* and w_* as a function of γ only.

(v) Argue that no pooling contract can upset this separating contract: there is no pooling contract that attracts students of all abilities at a salary where the bank breaks even.

Answer: a. If the ability of the student is observed, the bank will offer her marginal productivity, so $w(\vartheta) = \vartheta$ is the wage contract.

b. (i) For the indifferent individual, $w^* - (1 - \vartheta^*)\gamma = w_*$, since she will (optimally) acquire the lowest grades that qualify her for each salary. By direct computation,

$$\vartheta^* = 1 - \frac{w^* - w_*}{\gamma} \tag{$*$}$$

(ii) if $\vartheta > \vartheta^*$, then

$$w^* - (1 - \vartheta)\gamma > w^* - (1 - \vartheta^*)\gamma = w_*,$$

so it follows that any student with ability above ϑ^* will acquire grades γ. Any sudent below ϑ^* will acquire null grades.

(iii) For the bank to break even,

$$w_* = \mathrm{E}\left[\Theta \mid \Theta < \vartheta^*\right] = \frac{\vartheta^*}{2},$$

and

$$w^* = \mathrm{E}\left[\Theta \mid \Theta > \vartheta^*\right] = \frac{1 + \vartheta^*}{2}.$$

(iv) Substituting the two salaries into Eq. ($*$),

$$\vartheta^* = 1 - \frac{1}{2\gamma},$$

and, then,

$$w_* = \frac{1}{2}\left(1 - \frac{1}{2\gamma}\right) \text{ and } w^* = \frac{1}{2}\left(2 - \frac{1}{2\gamma}\right).$$

(v) If it is not to make losses, a pooling contract would offer at most $\bar{w} = \mathrm{E}[\vartheta] = 1/2$. For it to succeed, it has to attract students of all abilities, including $\vartheta = 1$. But this would require that

$$\frac{1}{2} \geq \frac{1}{2}\left(2 - \frac{1}{2\gamma}\right),$$

which is impossible if $\gamma > 1/2$. □

Exercise 11.1. Consider a competitive labor market populated by two types of workers, $\tau \in \{H, L\}$. A worker of type τ produces $\vartheta_\tau \sqrt{e}$, where ϑ_τ denotes her exogenously given productivity level and e the individual's chosen education level. Suppose that the cost of acquiring each unit of education for an individual of type τ is c_τ. As in class, suppose that $\vartheta_H > \vartheta_L$ and $c_H < c_L$.

a. To establish a benchmark, determine the Pareto efficient levels of education of the workers.

b. State the condition, in terms of the parameters of the model, under which a separating equilibrium (screening) cannot be sustained, if workers are intended to acquire the Pareto efficient levels of education.

c. Suppose that the condition just stated holds, so that the efficient separating equilibrium does not exist. Describe a potential, second-best separating equilibrium levels of education.

d. Letting λ be the proportion of type L workers in the market, are there conditions under which the separating contracts just described are indeed an equilibrium?

Exercise 11.2. Consider a competitive labor market populated by two types of workers, $\tau \in \{H, L\}$. A worker of type τ produces $\vartheta_\tau e$, where ϑ_τ denotes her exogenously given productivity level and e the individual's chosen education level. Suppose that the cost of acquiring e units of education for an individual of type τ is $c_\tau e^2$. As in class, suppose that $\vartheta_H > \vartheta_L$ and $c_H < c_L$. Denote by λ be the proportion of type L workers in the market.

a. To establish a benchmark, determine the Pareto efficient levels of education of the workers.

b. State the condition, in terms of the parameters of the model, under which a separating equilibrium (screening) cannot be sustained, if workers are intended to acquire the Pareto efficient levels of education.

c. Suppose that the condition just stated holds, so that the efficient separating equilibrium does not exist. Describe a potential, second-best separating equilibrium contract, including the self-selection constraint that it must satisfy for the L individuals.

Exercise 11.3. Consider a competitive market for student loans. Suppose that each student borrows \$15,000, which are to be repaid, with interest, if, and only if, the student is successful in an economics Ph.D. program. Suppose that that there are two types of students, $\tau \in \{H, L\}$, who differ only in their probability of succeeding, p_τ. As usual, let $p_H > p_L$, suppose that a proportion λ of the students are of type L, and denote by ι the interest rate charged on the loans.

a. Define the equilibrium interest rate in a (pooling) equilibrium, if no screening is possible.

b. Suppose now that, by exerting costly effort, an individual can improve her undergraduate grades, which are observable. For a student of type τ, the cost of getting grades γ is $c_\tau \gamma$, with $c_H < c_L$. Suppose, moreover, that a person's undergrad grades have no impact on the probability of succeeding in the program.

   (i) Derive the values of the two types of loan contracts in a separating equilibrium, including the required grades and the relative values that c_L an c_H must satisfy in order that incentive compatibility requirements are satisfied.

   (ii) Assuming that $p_L = 1/3$, $p_H = 2/3$, $c_L = 1$ and $\lambda = 3/4$, find the critical value that c_H cannot exceed if the equilibrium you just defined is to exist.

Reading. The following excerpt is from an article in *The Economist*. Discuss it critically.

> *Work over the past five years by two economists, Dale Ballou at the University of Massachusetts and Michael Podgursky of the University of Missouri, suggests that the quality of America's teachers has more to do with how they are paid rather than how much. The pay of American public-school teachers is not based on any measure of performance (the NEA opposes merit pay); instead, it is determined by a rigid formula based on experience and years of schooling, factors that Mr Podgursky calls "massively unimportant" in deciding how well students do. The uniform pay scale invites what economists call adverse selection. Since the most talented teachers are also likely to be good at other professions, they have a strong incentive to leave education for jobs in which pay is more closely linked to productivity. For dullards, the incentives are just the opposite.*
>
> *The data are striking: when test scores are used as a proxy for ability, the brightest individuals shun the teaching profession at every juncture. Clever students are the least likely to choose education as a major at university. Among students who do major in education, those with higher test scores are less likely to become teachers. And among individuals who enter teaching, those with the highest test scores are the most likely to leave the profession early.*

From: *The Economist* (August 24th, 2000) "Paying teachers more".

References: This material is based on Spence, A.M. "Job Market Signaling." *The Quarterly Journal of Economics* (1973): 355- 74.

Note 12

Credit markets

OBVIOUSLY, ADVERSE SELECTION CAN AFFECT OTHER IMPORTANT MARKETS. An issue of most interest is how asymmetric information can affect financial markets, and ways in which some market mechanisms may arise that ameliorate this problem. The intuition of the problem is the same as in the classical models studied before, but we now concentrate in the specific context of a market for loans. In particular, we want to determine whether the inability of a bank to observe all the details of an investment project may lead to inefficiency in this market.

For this purpose, suppose that there is a continuum of investment projects, indexed by $\vartheta \in [\vartheta_*, \vartheta^*]$. Parameter ϑ will be the riskiness *type* of the project, in the following sense. Let R denote the return of project ϑ, which is a random variable with distribution F_ϑ. In order to model the type of the project as its measure of riskiness, assume that $\mathrm{E}_{F_\vartheta}(R) = \mathrm{E}_{F_{\vartheta'}}(R)$ for all ϑ and ϑ', but $F_\vartheta >_2 F_{\vartheta'}$ if $\vartheta < \vartheta'$. In words, all projects types have the same expected return, but the lower its type the less risky a project is. Importantly, let us assume that while the entrepreneur can observe the riskiness type of her project, the bank cannot and only knows that it comes from a distribution G over $[\vartheta_*, \vartheta^*]$.

Assume that all projects are of the same size, in the sense that they require a loan of $b > 0$ to finance them, regardless of their type. In order to get this loan, however, the owner of the project has to put down a collateral of c, that the bank will keep in case of default, besides the full return of the funded project. Let us denote by ι the interest rate charged on the loan. Assuming that no further punishment is imposed on a defaulted loan, it follows that the entrepreneur will choose to default when the realization of R, which we denote by r, is such that $c + r < (1 + \iota)b$. Under these assumptions, it follows that the value (net profit) of the entrepreneur is

$$v(r, \iota) = \max\{r - (1 + \iota)b, -c\},$$

while the profits of the bank are

$$\pi(r, \iota) = \min\{r + c, (1 + \iota)b\}.$$

To abstract from complications, let us assume that everybody in the market is risk-neutral. Figs. 12.1 and 12.2 depict these payoff functions. Importantly, note that the possibility of default implies that the entrepreneur *shares with the bank bad outcomes of her project, only.* Technically, notice that although both payoff functions are nondecreasing in r, the entrepreneur has a convex utility function while the bank's is concave.

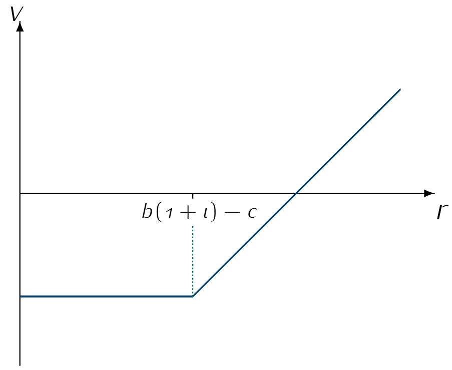

Figure 12.1: net profit of an entrepreneur.

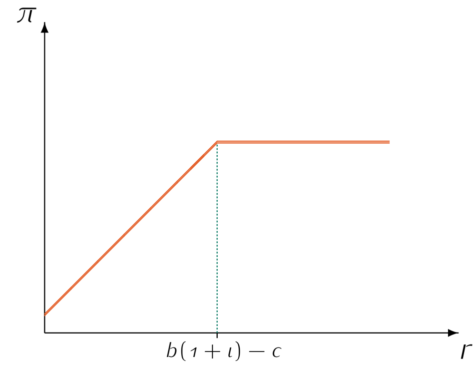

Figure 12.2: profit of the bank.

PROPOSITION 12.1. *Given interest rate ι, there exists a cutoff value of riskiness, $\tilde{\vartheta}$, so that the entrepreneur chooses to borrow and fund her project if, and only of, she is riskier than that, in the sense that her type is $\vartheta \geq \tilde{\vartheta}$.*

Proof. The argument is fairly straightforward. Let $\tilde{\vartheta}$ be such that $\mathrm{E}_{F_{\tilde{\vartheta}}}[v(R,\iota)] = 0$, namely the type of entrepreneur that is indifferent between taking the loan and funding her project, and giving up on the project. Since $\vartheta < \tilde{\vartheta}$ implies that $F_\vartheta >_2 F_{\tilde{\vartheta}}$ and V is convex, it follows that for riskiness below $\tilde{\vartheta}$, $E_{F_\vartheta}[v(R,\iota)] < 0$. □

PROPOSITION 12.2. *The cutoff type, $\tilde{\vartheta}$, is increasing in the interest rate charged by the bank.*

Proof. By definition, for all ι,

$$\begin{aligned} 0 &= \mathrm{E}_{F_{\tilde{\vartheta}}}[v(R,\iota)] \\ &= \mathrm{E}_{F_{\tilde{\vartheta}}}[\max\{R-(1+\iota)b, -c\}] \\ &= -\int_{-\infty}^{(1+\iota)b-c} c\,\mathrm{d}F_{\tilde{\vartheta}}(r) + \int_{(1+\iota)b-c}^{\infty} [r-(1+\iota)b]\,\mathrm{d}F_{\tilde{\vartheta}}(r) \end{aligned}$$

Differentiating with respect to ι,

$$\frac{\partial \mathrm{E}_{F_{\tilde{\vartheta}}}[v(R,\iota)]}{\partial \tilde{\vartheta}}\,\mathrm{d}\tilde{\vartheta} - \left(\int_{(1+\iota)b-c}^{\infty} b\,\mathrm{d}F_{\tilde{\vartheta}}(r)\right)\,\mathrm{d}i = 0,$$

which means that

$$\frac{\mathrm{d}\tilde{\vartheta}}{\mathrm{d}\iota} = b\frac{1-F_{\tilde{\vartheta}}((1+\iota)b-c)}{\partial \mathrm{E}_{F_{\tilde{\vartheta}}}[v(R,\iota)]/\partial\tilde{\vartheta}}.$$

By the argument used in the proof of the previous proposition, the denominator of this expression is positive, which implies that so is the whole expression. □

PROPOSITION 12.3. *The expected profit of the bank, given the type of entrepreneur, is decreasing in the riskiness of the project.*

Proof. This is immediate: since $\vartheta < \vartheta'$ implies that $F_\vartheta >_2 F_{\vartheta'}$, and π is concave. □

PROPOSITION 12.4. *The effect of the interest rate on the expected profit of the bank is ambiguous.*

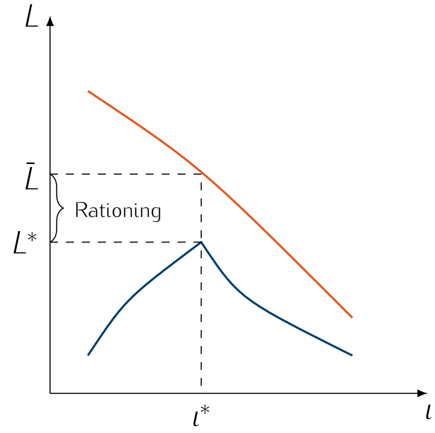

Figure 12.3: Credit rationing under perfect information.

Proof. As the bank cannot discern the type of the project, but knows that only projects above the threshold apply for a loan, its expected profit is

$$\Pi(\iota) = \mathrm{E}_G\left\{\mathrm{E}_{F_\vartheta}[\pi(R,\iota)] \mid \vartheta \geq \tilde{\vartheta}\right\} = \mathrm{E}_G\left\{\mathrm{E}_{F_\vartheta}[\min\{R+c, b(1+\iota)\}] \mid \vartheta \geq \tilde{\vartheta}\right\}.$$

Differentiating,

$$\Pi'(\iota) = b\frac{\partial \mathrm{E}_G\left\{\mathrm{E}_{F_\vartheta}[\pi(R,\iota)] \mid \vartheta \geq \tilde{\vartheta}\right\}}{\partial \iota} + \frac{\partial \mathrm{E}_G\left\{\mathrm{E}_{F_\vartheta}[\pi(R,\iota)] \mid \vartheta \geq \tilde{\vartheta}\right\}}{\partial \tilde{\vartheta}}\frac{\mathrm{d}\tilde{\vartheta}}{\mathrm{d}\iota}.$$

Since $b > 0$, the first summand on the right-hand side is positive. By the previous two propositions, the second summand is the product of a positive and a negative numbers. □

The implication of the last proposition is important. While the demand for loans, L^d, is decreasing in the interest rate, L^s, their supply, need not be increasing. This implies that it is possible for the bank to prefer to ration credit to potential borrowers rather than increase the interest rate, as in Fig. 12.3. The bank will not give loans at a rate above ι^* and will cap its lending at L^*. At this rate, the demand for loans is $\tilde{L} > L^*$, but the bank prefers *not* to serve this quantity, and the standard market mechanism of increasing prices does not eliminate the excess demand, due to its adverse-selection effect on the riskiness of funded projects.

Exercise 12.1. In the context of this problem:

1. Graph the function for the value of the entrepreneur, as a function of the realized return of the project, for given ι and c, and determine how changes in ι and c affect this function.

2. Use your previous answer to argue that the following *is possible*: the bank will offer two different loan contracts (ι_*, c^*) and (ι^*, c_*), where $\iota_* < \iota^*$ and $c_* < c^*$ are chosen such that safer entrepreneurs choose the contract (ι_*, c^*) , while only high-risk entrepreneurs take the (ι^*, c_*) contract .

Exercise 12.2. Assume that (instead of commercial banks) investment banks provide the loan, and participate of the return of the project, in the following sense: the bank stipulates a cut-off value ρ and a profit share $s < 1$; if the return of the project is below the cut-off value, the bank condones (forgives) the loan; otherwise, the bank receives a proportion s of the excess of the return over ρ.

1. Write expressions for the value of the entrepreneur and the profits of the investment bank, as functions of the cut-off value and the realized return of the project: $v(r, \rho, s)$ and $\pi(r, \rho, s)$. Draw the graphs of these two functions, for given ρ and s, as functions of the realized return of the project.

2. Argue that if an entrepreneur of riskiness ϑ is willing to undertake the investment project at a given ρ and s, then any entrepreneur with lower riskiness is willing to do the same.

3. Write a condition to define the "marginal" entrepreneur who is indifferent between investing and not, and argue that this marginal entrepreneur is decreasing in the share of profits charged by the bank, s.

Reading. The following excerpt is from a column in *Harper's*. Discuss it critically.

> *The ever-widening information imbalance between consumers and creditors has only made borrowers easier marks. In a Federal Trade Commission study conducted last year, for instance, nine in ten mortgage customers examining relatively straightforward fixed-rate loan agreements could not figure out the up-front costs on the loan; half could not identify the loan amount. Of all the borrowers who were sold subprime mortgages in the past five years, nearly 60 percent would have qualified for prime mortgages if brokers had offered them; the sub-prime mortgages carried so many rate escalators, prepayment penalties, and other traps that even would-be prime borrowers defaulted.*
>
> *It is time we created the equivalent of a Consumer Product Safety Commission for financial products, an agency whose purpose would be to protect homebuyers and investors from the finance industry's most dangerous offerings. The Financial Product Safety Commission could model itself after the best from the consumer regulatory agencies. For instance, the head of the new agency would be appointed by the president, and its staff of professionals would have civil-service protection and thereby be immune to changing political winds. Although the FPSC would have no hand in setting prices, it would be able to require that companies reveal the true cost of credit. This seemingly small requirement would force into public view essential information about terms and risks that has long been masked and withheld. To achieve this end, the agency could do something as basic as reviewing product disclosures, making sure they were easily comprehensible to the average reader.*

The FPSC *would also "test" products for safety before they had a chance to reach consumers. When the commission found undisclosed fees or bait-and-switch credit modeling, it could allow the offender a period of time to fix the problem, giving lenders the opportunity to minimize government interference. But if a lender failed to act within, say, six months, the agency could impose its own regulations: eliminating confusing paperwork, requiring effective disclosures, and, when necessary, banning outright the most dangerous traps.*

From: Warren, E. and A. Warren Tyagi, *Harper's Magazine* (November, 2008), "Protect financial consumers".

References: This material is based on Stiglitz, J.E. and A. Weiss. "Credit rationing in markets with imperfect information." *The American Economic Review* 71.3 (1981): 393-410.

Note 13

Investment

ADVERSE SELECTION ALSO HAS MACROECONOMIC IMPLICATIONS. We now study how asymmetric information can affect the levels of investment in an economy. We will consider, again, a model with adverse selection, where firms needs to resort to credit markets to fund their investment projects. In this setting, it is natural that information asymmetries make external funding more expensive than internal funding, and can lead to extreme forms of inefficiency in the market.

13.1. A cash-constrained firm

CONSIDER THE CASE OF a firm that is evaluating an indivisible investment project to expand its output. If the firm does not invest, its output is $z > 0$. The investment opportunity is risky: if the firm undertakes it, there is a probability p that the total output is $x > z$; but the project can also fail, in which case the firm's total output is 0. The investment project costs $k > 0$, but it is assumed to be profitable, in the sense that $px - k > z$.

The firm is risk-neutral, and its total available cash is $c < k$. If it chooses to invest in the project, the difference, $k - c$ has to be borrowed in financial markets. Let ι denote the interest rate charged by (risk-neutral) banks. Assuming that the banks cannot enforce collateral requirements, nor any other form of penalty for defaulting on loans, the firm will repay its loan only *only* if the output is positive and

$$x \geq (k - c)(1 + \iota).$$

Under the usual argument of perfect competition in financial markets,[1] the interest rate charged by the bank is

$$\iota = \frac{1}{p} - 1 > 0,$$

which is accepted by the firm, under our assumption that the investment project is profitable.

This outcome is efficient, as profitable investment projects get funded.

[1] And, for the sake of simplicity, ignoring the bank's funding costs.

13.2. Adverse selection

Suppose now that there are two types of firm, L and H, which differ only in their probabilities of being successful in the investment project: $p_L < p_H$. The proportion of L firms is $\lambda \in (0, 1)$. We maintain the assumption that the investment project is profitable for both types, so $p_L x - k > z$.

Since the decision to repay is made *ex-post*, the condition under which repayment occurs is the same for both types: $x \geq (k-c)(1+\iota)$.

13.2.1. Pooling equilibrium

If a contract that is taken by both types of firm, the probability of success considered by the bank is $\bar{p} = \lambda p_L + (1-\lambda)p_H$. As before, equilibrium in the credit market implies that the interest rate is $\bar{\iota} = 1/\bar{p} - 1$.

Firms of type τ will accept this contract only if

$$\begin{aligned} z &\leq p_\tau [x - c - (1+\bar{\iota})(k-c)] + (1-p_\tau)(-c) \\ &= p_\tau x - \frac{p_\tau}{\bar{p}}(k-c) - c \\ &= p_\tau x - k - \left(\frac{p_\tau}{\bar{p}} - 1\right)(k-c), \end{aligned}$$

where the second equality comes from substituting $\bar{\iota}$, and the third one from adding and subtracting k. This is, only if

$$p_\tau x - k - \frac{p_\tau - \bar{p}}{\bar{p}}(k-c) \geq z. \tag{13.1}$$

Since $p_L x - k > z$, this condition is granted for type L firms, as $p_L < \bar{p}$. For H-type firms, the same need not be true, even though $p_H x - k > z$, since $p_H > \bar{p}$.

If Eq. (13.1) holds true for both types of firm, a pooling equilibrium exists. Otherwise, type-H firms are not willing to borrow at the pooling interest rate, and pooling cannot be sustained at equilibrium. The intuition is simple: under pooling, the bank charges a higher interest rate to "good" firms, relative to what it would charge them under perfect information. If this increase is sufficiently high, these firms would choose not to borrow, even though the project would be profitable at the right interest rate. For type-L firms, this effect does not occur, and the pooling interest rate is lower.

13.2.2. Separating equilibrium

Let us now assume that

$$p_H x - k - \frac{p_H - \bar{p}}{\bar{p}}(k-c) < z, \tag{13.2}$$

so that pooling is impossible at equilibrium.

Since we are assuming that collateral requirements are not enforceable, banks have no way of screening firms: if two different interest rates were offered, all firms would take the lowest one. The only possible separating equilibrium is one where only L firms borrow, so that the interest rate is

$$\iota^* = \frac{1}{p_L} - 1 > \bar{\iota}.$$

Since the project is profitable for L firms, they accept this interest rate. For it to be an equilibrium, we must also check that H firms prefer not to invest:[2]

$$p_H x - k - \frac{p_H - p_L}{p_L}(k - c) < z.$$

But, by construction,

$$\frac{p_H - p_L}{p_L} = \frac{p_H}{p_L} - 1 > \frac{p_H}{\bar{p}} - 1 = \frac{p_H - \bar{p}}{\bar{p}},$$

so the latter condition is implied by Eq. (13.2).

13.3. Welfare and efficiency

In this particular setting existence of equilibrium has not been a problem. A pooling equilibrium is possible if the difference between the probabilities of success is not too large, and when this condition fails a separating equilibrium is guaranteed to exist.

This does not mean, in any way, that information asymmetries are not problematic. When the pooling equilibrium holds, adverse selection is not too severe, but some firms are charged a higher interest rate than the one implied by their inherent risk. More importantly, in a separating equilibrium the firms that are most likely to have successful projects choose not to invest.

From a macroeconomic perspective, note that the condition under which the pooling equilibrium fails, Eq. (13.2), depends on the value of the loan that the firm needs to fund investment projects. A policy under which this value is lower would ameliorate the adverse selection problems, as would a policy that facilitates screening of low-probability firms, for instance the introduction of collateral requirements.

References: This material is based on Bernanke, B. and M. Gertler. "Agency costs, net worth, and business fluctuations." *The American Economic Review* (1989): 14-31.

[2] Otherwise, banks would be making strictly positive expected profits, which contradicts the idea that they exhibit free entry and perfect competition.

Part III

Incentives

Note 14

Agency theory

OFTEN, AN ECONOMIC AGENT NEEDS TO LEAVE IMPORTANT DECISIONS IN THE HANDS OF OTHER AGENTS. Our goal now is to study situations in which an decision maker delegates to someone else the undertaking of actions that affect her well-being. Situations like this are ubiquitous, but the main example is the interaction between the owner of a firm and the firm's manager. The interest of the problem arises when the goals of the person who delegates are not perfectly aligned with those of the agent who receives the mandate to act.

14.1. Perfect information

THE CANONICAL AGENCY PROBLEM is the following. A *principal* is to delegate the undertaking of an action to an *agent*. The set of actions that are possible for the agent is A, with representative element a. There is randomness in this setting, represented (for the moment) by a space of states of nature $\mathcal{S}$, with representative element s.

In state of the world s, if the agent's action is a, the principal realizes an output $x(a, s)$. From this output, the principal remunerates the agent. Suppose that $A \subseteq \mathbb{R}$ and $\mathcal{S} \subseteq \mathbb{R}$, and, as a convention, assume that the output is increasing in both a and s. The agent derives utility of her remuneration, φ, and disutility from undertaking higher levels of the action; her preferences are represented by

$$v(\varphi, a), \quad \text{with } \frac{\partial v}{\partial \varphi} > 0 \text{ and } \frac{\partial v}{\partial a} < 0.$$

The principal, on the other hand, only cares about his net output, and his preferences are $u(x-\varphi)$, with $u' > 0$.

For the moment, let us suppose that the agent's remuneration is contingent on the action and the state, and is given by a function $f(a, s)$. Taking this function as given, the agent's problem is

$$\max_{a \in A} \mathrm{E}\{v[f(a, S), a]\},$$

where the random variable over which the expectation is taken is the state of the world, S. The principal, on the other hand, chooses the remuneration *function* so as to solve

$$\max_{f: A \times \mathcal{S} \to \mathbb{R}} \mathrm{E}\{u[x(a, S) - f(a, S)]\}.[1]$$

[1] It may be useful to emphasize that the choice variable in this problem is the whole function $f : A \times \mathcal{S} \to \mathbb{R}$.

In terms of modelling the interaction between the two individuals, the previous problem needs to be better specified: the critical issue here is how the principal believes that the agent will respond to the incentives introduced to her problem by the remuneration function. Even though we will not made this feature explicit, it is useful to think that the remuneration fee is chosen by the principal *first*, taking into account how the agent will optimally respond to it *later*.[2] This means that the principal faces two constraints in his optimization problem:

1. It seems plausible to assume that the agent has alternative activities she can engage in. The principal must, then, guarantee that the fee gives the agent at least the same expected utility as those alternatives, for otherwise she would not participate in the contract. This constraint is known as *individual rationality*.

2. Given the fee schedule, since it is the agent who chooses the action, the principal must take into account how the action that the agent will choose responds to the schedule. That is, when choosing the fee function, the principal must evaluate her objective function at the action that solves the agent's optimization problem, given that function. This constraint is referred to as *incentive compatibility*.

At first glance, it may seem that the incentive compatibility constraint captures all that is interesting about this problem. While this constraint is, to a large extent, the crux of the problem, there is another aspect of it that is important too, and which is less explicit. Note that we are dealing with two agents who are facing exposure to a common risk. In this case, an embedded problem is the one of how these agents will share this risk. To concentrate on this problem, we will first ignore the action of the agent altogether.

14.1.1. Optimal risk sharing

Suppose that the problem involves no action by the agent. Abusing notation in an obvious way, the principal's problem is

$$\max_{f:\mathcal{S}\to\mathbb{R}} \{\mathrm{E}\{u[x(S)-f(S)]\} : \mathrm{E}\{v[f(S)]\} \geq v_*\},$$

where v_* denotes the agent's reservation utility,[3] which has to be guaranteed by the principal through the individual rationality constraint.

Denoting by G the distribution of random variable S, we can rewrite this problem as

$$\max_{f:\mathcal{S}\to\mathbb{R}} \left\{\int u[x(s)-f(s)]\,\mathrm{d}G(s) : \int v[f(s)]\,\mathrm{d}G(s) - v_* = 0\right\},$$

where we have used the monotonicity of function u to write the individual rationality constraint as an equality. The first-order conditions of this problem are that for all state of the world s,

$$u'[x(s)-f(s)] = \lambda v'[f(s)], \tag{14.1}$$

We will next address the conceptual issues raised by this, but there are technical problems that arise here too. In particular, the solution to this optimization problem will require techniques of dynamic optimization, even though the problem is static, where we use the state of the world as index of the problem, instead of time.

[2] From a modelling point of view, this makes things very simple. It rules out, however, the possibility that the fee be negotiated, and assumes that the principal can commit to the schedule.

[3] Namely, the highest expected utility she can attain in an alternative contract.

for some number $\lambda > 0$.[4] Differentiating with respect to s, we obtain that

$$f'(s) = \frac{u''[x(s) - f(s)]x'(s)}{\lambda v''[f(s)] + u''[x(s) - f(s)]}.$$

It is immediate from the latter expression that:

1. If both the principal and the agent are strictly risk averse, then $0 < f'(s) < x'(s)$, which implies that they *share risk*. To find the optimal fee schedule, the latter differential equation has to be solved, using the individual rationality constraint as terminal condition.

2. If the principal is risk neutral, then $f'(s) = 0$ and the optimal schedule is simply $f(s) = v^{-1}(v_*)$.

3. If the agent is risk neutral, then $f'(s) = x'(s)$ and $f(s) = x(s) - \bar{\varphi}$, where $\bar{\varphi}$ is chosen so as to guarantee the individual rationality constraint.

Intuitively, if one of the the two parties in the contract is risk neutral and the other one risk averse, the optimal contract transfers all the risk to the risk-neutral one. If, on the other hand, they are both risk averse, they optimally share the risks implicit in the environment.

14.1.2. Pareto efficiency

We now reintroduce the action of the agent. To make our analysis a little simpler, we assume that

$$\frac{\partial^2 v}{\partial a \, \partial \varphi} = 0,$$

namely that the action and the income of the agent enter her utility function in a *separable* manner.[5]

If the principal chose the action, his problem would be, simply:

$$\max_{a \in A, f : \mathcal{S} \to \mathbb{R}} \{\mathrm{E}\{u[x(a, S) - f(S)]\} : \mathrm{E}\{v[f(S), a]\} \geq v_*\}.$$

We can rewrite this problem as

$$\max_{a \in A, f : \mathcal{S} \to \mathbb{R}} \left\{ \int u[x(a, s) - f(s)] \, \mathrm{d}G(s) : \int v[f(s), a] \, \mathrm{d}G(s) - v_* = 0 \right\},$$

which has as first-order conditions that there exist a positive number λ such that for each state for the world, s,

$$u'[x(a, s) - f(s)] = \lambda \frac{\partial v}{\partial \varphi}[f(s), a],$$

while

$$\int u'[x(a, s) - f(s)] \frac{\partial x}{\partial a}(a, s) \, \mathrm{d}G(s) + \lambda \int \frac{\partial v}{\partial a}[f(s), a] \, \mathrm{d}G(s) = 0,$$

[4] Technical aspects of this solution are found in Appendix A1.

[5] The assumption of risk aversion, which is maintained, is now that

$$\frac{\partial^2 v}{\partial \varphi^2} \leq 0.$$

in addition to the individual rationality constraint. Since we are assuming that the agent's preferences are separable, we can write the latter condition as

$$\mathrm{E}\left\{u'[x(a,S)-f(S)]\frac{\partial x}{\partial a}(a,S)\right\}+\lambda\frac{\partial v}{\partial a}[0,a]=0.$$

Substituting from the first condition, the latter becomes, simply,

$$\mathrm{E}\left\{\frac{\partial v}{\partial \varphi}[f(S),a]\frac{\partial x}{\partial a}(a,S)\right\}=-\frac{\partial v}{\partial a}[0,a], \tag{14.2}$$

which is very intuitive: it equates the marginal disutility of the agent's action with her marginal utility from getting the total marginal income of such action. This intuition suggests that the outcome of this optimization problem is Pareto efficient, which is indeed the case.

Of course, this case is simply a theoretical benchmark, as in reality it is the agent, not the principal, who chooses the action. Moreover, if the principal cannot separately observe the action and the state of the world, he can only make the fee depend on the (observed) output he realizes, so we now write the contract as $f(x)$. All this implies that the action that is chosen is the solution to the problem

$$\max_{a\in A}\mathrm{E}\{v[f(x(a,S)),a]\},$$

which has as first-order conditions that

$$\mathrm{E}\left\{\frac{\partial v}{\partial \varphi}[f(x(a,S)),a]f'[x(a,S)]\frac{\partial x}{\partial a}(a,S)\right\}=-\frac{\partial v}{\partial a}[0,a]. \tag{14.3}$$

Comparison between Eqs. (14.2) and (14.3) yields intuition about the latter: when left to her, the agent only internalizes the part of the marginal product that is transferred to her via the fee, as opposed to the whole benefit of a higher action. It may be tempting to conclude that in order to restore efficiency the principal could offer a schedule where $f'(x)=1$ for all x. This, however, leaves all the risk in the hands of the agent, and need not be efficient if she is not risk neutral, as we saw before.

14.2. Imperfect information: moral hazard

We now study the problem where the principal has to satisfy the incentive compatibility constraint in his design of the optimal fee schedule. For technical reasons, it will be convenient to get rid of the state of the world in our notation. Instead, we will model the uncertainty through the probability distribution of the outcome, and will consider the effects of the agent's action directly over this distribution.

Formally, let X denote the output, which is now a random variable, and suppose that its distribution depends on the action taken by the agent. That is, letting $G:\mathbb{R}\times A\to[0,1]$, we understand that $G(x,a)$ is the probability that $X\leq x$, *when the action of the agent is a*. Keeping with our understanding that the action improves the output, we assume that G is decreasing in a, so that the higher the action of the agent, the better the distribution of the output in the sense of first-order stochastic dominance.

The principal's problem is, with this notation,

$$\max_{a\in A, f:\mathbb{R}\to\mathbb{R}}\left\{\int u[x-f(x)]g(x,a)\,\mathrm{d}x:\int v[f(x),a]g(x,a)\,\mathrm{d}x-v_*=0, a\in A(f)\right\},$$

where

$$A(f) = \operatorname{argmax}_{a' \in A} \int v[f(x), a'] g(x, a')\, dx. \tag{14.4}$$

We have, hence, introduced the incentive compatibility constraint, as given by its definition. Not surprisingly, this way of modelling the incentive compatibility constraint is very difficult to handle, so a technique has been devised by which we substitute this form of the condition by the first-order conditions that characterize the solution of the optimization problem of the agent.[6] From Eq. (14.4), these conditions are that

$$\int \frac{\partial v}{\partial a}[f(x), a] g(x, a)\, dx + \int v[f(x), a] \frac{\partial g}{\partial a}(x, a)\, dx = 0,$$

or, using separability of the agent's preferences (as before), that

$$\int v[f(x), a] \frac{\partial g}{\partial a}(x, a)\, dx = -\frac{\partial v}{\partial a}[0, a].$$

Letting λ be the multiplier of the individual rationality constraint, and μ the one of the incentive compatibility constraint, the first order conditions of the principal's problem are: first, that for all x,

$$\frac{u'[x - f(x)]}{\partial v[f(x), a]/\partial f} = \lambda + \mu \frac{\partial g(x, a)/\partial a}{g(x, a)}; \tag{14.5}$$

while

$$\int u[x - f(x)] \frac{\partial g(x, a)}{\partial a}\, dx = -\mu \frac{d}{da} \left\{ \int v[f(x), a] \frac{\partial g}{\partial a}(x, a)\, dx + \frac{\partial v}{\partial a}[0, a] \right\}. \tag{14.6}$$

To see that we have a complete system to solve, note that Eq. (14.5) gives sufficiently many conditions to solve for $f(x)$, for each x, while for the other three variables, a, μ and λ, we have Eq. (14.6), the individual rationality constraint, and the incentive compatibility constraint. Note also that $\lambda > 0$ is required for optimality (given monotonicity $u' > 0$), but we do not know the sign of μ, as the version of the individual rationality constraint we have used is an equality.

Lemma 1. $\mu > 0$.

Proof. The left-hand side of Eq. (14.6) is positive, since higher a induces a first-order stochastic dominance improvement in $G(\cdot, a)$ and $u' > 0$. On the right-hand side, μ multiplies the second derivative of the agent's objective function with respect to her action. If the first-order conditions characterize the agents incentive compatibility constraint, it must the that the second-order conditions hold, so this second derivative must be negative. □

Theorem 14.1. *The optimal fee schedule fails to display optimal risk sharing.*

Proof. Recall that optimal risk sharing requires, as in Eq. (14.1), that the ratio of $u'[x - f(x)]$ and $\partial v[f(x), a]/\partial f$ be a constant (independent of x). By Eq. (14.5) and Lemma 1, this requires that the ratio of $\partial g(x, a)/\partial a$ and $g(x, a)$ be a constant too. Suppose that this is the case, and let κ denote this constant, so that $\partial g(x, a)/\partial a = \kappa g(x, a)$. Integrating with respect to x,

$$\int \frac{\partial g}{\partial a}(x, a)\, dx = \kappa \int g(x, a)\, dx.$$

[6] See Appendix A2 for more on this approach.

The left-hand side of this equality is null, while $\int g(x,a)\,\mathrm{d}x = 1$, so $\kappa = 0$. But, then, $\partial g(x,a)/\partial a = 0$, which contradicts our assumptions. □

We can, in fact, say more about the distortions to optimal risk sharing, with a simple intuition that is highlighted by the following result.

Corollary 14.1. *Let $f^* : \mathbb{R} \to \mathbb{R}$ be such that, for all x,*

$$u'[x - f^*(x)] = \lambda \frac{\partial v}{\partial f}[a, f^*(x)],$$

so that it shares risks efficiently. Then, $f(x) \geq f^*(x)$ if $\partial g(x,a)/\partial a \geq 0$, while $f(x) < f^*(x)$ if $\partial g(x,a)/\partial a < 0$.

Proof. From Eq. (14.5), if $\partial g(x,a)/\partial a \geq 0$,

$$\frac{u'[x - f(x)]}{\partial v[f(x), a]/\partial f} = \lambda + \mu \frac{\partial g(x,a)/\partial a}{g(x,a)} \geq \lambda = \frac{u'[x - f^*(x)]}{\partial v[a, f^*(x)]/\partial f}.$$

Since $u'' < 0$ and/or $\partial^2 v[a, f^*(x)]/\partial f^2 < 0$, the conclusion follows. □

This last result illustrates that the deviations from optimal risk sharing are intended to give the agent incentives to choose an action that is more convenient. Since a higher action induces, overall, a first-order stochastic dominance improvement in the distribution of output, the outcomes that gain density are "better" than those that lose it. For this outcomes, the principal offer higher remuneration than what would be efficient from the point of view of risk sharing.

Finally, we want to study the shape of the fee function. Before we do that, it is necessary to assume the following property, which is referred to as *the monotone likelihood ratio*:

$$\frac{\partial}{\partial x}\left(\frac{1}{g(x,a)}\frac{\partial g}{\partial a}(x,a)\right) > 0.$$

The intuition of this property is that, since the principal does not observe but infers the action, the property says that the higher the output he observes, the higher the action he infers. This intuition, unfortunately, is not guaranteed under the assumptions we had made before. For more on the intuition, see below.

Theorem 14.2. *The optimal fee schedule is increasing.*

Proof. From Eq. (14.5), differentiating with respect to x, we have that

$$\frac{u''[1 - f'(x)]}{\partial v/\partial \varphi} - \frac{u' f'(x)(\partial^2 v/\partial \varphi^2)}{(\partial v/\partial \varphi)^2} = \mu \frac{\partial}{\partial x}\left(\frac{\partial g/\partial a}{g}\right),$$

where we have ignored the arguments of all functions but f. Immediately,

$$f'(x) = \left(\frac{u''}{\partial v/\partial \varphi} + \frac{u'(\partial^2 v/\partial \varphi^2)}{(\partial v/\partial \varphi)^2}\right)^{-1}\left(\frac{u''}{\partial v/\partial \varphi} - \frac{\partial}{\partial x}\left(\frac{\partial g/\partial a}{g}\right)\right).$$

The result is immediate under the monotone likelihood ratio. □

Exercise 14.1. An entrepreneur has to take a loan of amount L to fund investement on a risky project. Unlike in the cases seen in class, assume now that it is the entrepreneur that designs the repayment contract for the loan. The return of the project, denoted by X, is observable, and the repayment contract can be made contingent on it: $R(x)$ denotes how much the entrepreneur repays when the project has returned x.

Return X is a random variable, whose distribution is influenced by the *effort* exerted by the entrepreneur in running her company: the probability that $X \leq x$, when the effort of the entrepreneur is e, is $F(x, e)$. Assume that $\partial F/\partial e < 0$, so that higher levels of effort induce a first-order stochastic dominance improvement in the return of the project. Effort is *not* observable, and for the entrepreneur the cost of exerting effort e is $\varphi(e)$, which is assumed to be an increasing and convex function.

Suppose that both the entrepreneur and the lender are risk neutral, that the lender only accepts the contract if his expected net return is non-negative, and that, for institutional reasons, the contract has to obey a *limited liability* constraint: that $0 \leq R(x) \leq x$ for all x. Suppose also that the distribution of the return of the project satisfies the *monotone likelihood ratio* property, so that

$$\frac{1}{f(x, e)} \cdot \frac{\partial f}{\partial e}(x, e)$$

is increasing in x, for all e

1. Write the individual rationality constraint of the lender.

2. Ignoring for the moment the limited liability constraint, and assuming that the entrepreneur can commit to a level of effort, white her optimization problem and find its first-order condition.

3. Introducing the limited liability and incentive compatibility constraints, re-write the entrepreneur's problem. Show that the optimal solution is of the form

$$R(x) = \begin{cases} x, & \text{if } x \leq \tilde{x}; \\ 0, & \text{otherwise.} \end{cases}$$

   for some level $\tilde{x}$ of return.

14.3. Validity of the first-order approach

Recall that in our solution, instead of using the original incentive compatibility constraint, Eq. (14.4), we are requiring that

$$\frac{\partial}{\partial a} \int v[f(x), a'] g(x, a') \, dx.$$

In words, this means that, instead of using the agent's optimal action, we are saying that the principal foresees that the agent will choose a critical point of her objective function.

Of course, this could mean that the principal thinks that the agent will choose a minimum, or a saddle point of her objective function, which would be (grossly) wrong. The conditions that we have imposed so far (risk aversion, separability, first-degree stochastic dominance and the monotone likelihood ratio) do not suffice to rule out this possibility. We still need to impose a second-order condition on how the action affects the probability distribution of outcomes. Recalling that $\partial G/\partial a < 0$, it is natural that the required condition is that $\partial^2 G/\partial a^2 > 0$, namely that the improvements on distribution occur at a diminishing rate.

14.4. The monotone likelihood ratio

To SEE THE INTUITION of this condition more clearly, it is useful to consider a discrete case. Suppose that there are only two outcomes, $x^* > x_*$, and two actions $a^* > a_*$. In this case, $g(x, a)$ is the probability that action x occurs when the action is a, so that the *odds* of high output, given action a, are $g(x^*, a)/g(x_*, a)$. These odds measure the likelihood of high output relative to low output.

The monotone likelihood ratio, in this case, is simply that these odds improve with the higher action:

$$\frac{g(x^*, a^*)}{g(x_*, a^*)} > \frac{g(x^*, a_*)}{g(x_*, a_*)}.$$

This expression is equivalent to

$$\frac{g(x^*, a^*)}{g(x^*, a_*)} > \frac{g(x_*, a^*)}{g(x_*, a_*)},$$

which can be approximated, to a first order, by

$$\frac{g(x^*, a_*) + (a^* - a_*)[\partial g(x^*, a_*)/\partial a]}{g(x^*, a_*)} > \frac{g(x_*, a_*) + (a^* - a_*)[\partial g(x_*, a_*)/\partial a]}{g(x_*, a_*)},$$

or, after simplification,

$$\frac{1}{g(x^*, a_*)}\frac{\partial g}{\partial a}(x^*, a_*) > \frac{1}{g(x_*, a_*)}\frac{\partial g}{\partial a}(x_*, a_*).$$

This is analogous to our continuous expression:

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{1}{g(x, a)}\frac{\partial g}{\partial a}(x, a)\right) > 0.$$

READING. The following excerpt is from an article in *Harper's*. Discuss it critically.

> *We must change how financial executives are personally compensated. We should require that stock options be subject to "expensing" (a more transparent accounting that makes clear their full costs). The present stock-option payment structure encourages CEOs to take actions that bloat the short-term reported profits of the firm, thereby inflating the share price, and everyone (except the executives in the know) eventually loses as a result. Their pay must be based on long-term performance, and they should share the losses, not just the gains.*

Certain masterminds of Wall Street exhibited great ingenuity in creating new, highly complex products capable of evading accounting rules and taking full advantage of the housing frenzy. But as they were getting rich off these innovations, they failed to design products that help reduce the risks faced by most people in the housing market. Mortgages that would make it easier for Americans to keep their homes as interest rates rise or the economy spirals downward can be developed. But those in the financial sector have been fixated on their own annual bonuses.

Adam Smith argued that in serving their own interests individuals were led "as if by an invisible hand" to serve the interests of society as a whole. But once again we see that only with the right re- wards can these interests actually be joined.

From: Stiglitz, J., "Realign the interests of Wall Street", *Harper's Magazine*, March 2008.

Appendix: Math

State spaces and probabilities

A *state of the world* is a comprehensive description of the status of *all* the contingencies that may affect an individual, or a group of people, but which are exogenous to them. The *state space* is the universe of all states of nature. An *event* is a collection of states of nature—that is, a subset of the state space.

Let $\mathcal{S}$ be a state space, with states denotes by $s \in \int$. We will focus on two types of state space. If the state space is finite, we write it in the general form $\mathcal{S} = \{1, 2, \ldots, \bar{s}\}$. Alternatively, in some cases the state space will be an interval of the real line, of the form $\mathcal{S} = [0, \bar{s}]$.

If the state space is finite, a *probability distribution* is a vector $p = (p_1, p_2, \ldots, p_{\bar{s}}) \geq 0$ such that

$$\sum_{s=1}^{\bar{s}} p_s = 1.$$

If it is an interval, a probability distribution is a function $p : \mathcal{S} \longrightarrow \mathbb{R}_+$ such that

$$\int_0^{\bar{s}} p(s)\,\mathrm{d}s = 1.$$

The interpretation is simple. If the state space is finite, the probability that an event E realizes is

$$\Pr(E) = \sum_{s \in E} p_s;$$

if the space is an interval, it is

$$\Pr(E) = \int_E p(s)\,\mathrm{d}s,$$

assuming that the integral exists (which it may not, for some $E \subseteq \mathcal{S}$, in which case such E is not considered an event).

Random variables

Given a state space $\mathcal{S}$, a *random variable* is a function $X : \mathcal{S} \longrightarrow \mathbb{R}$. The interpretation is that when state of nature s realizes, the variable of interest takes the value $X(s)$. If the state space is finite, the random variables defined on it are said to be *discrete*; if the space is an interval, any the random variable whose range is also an interval is said to be *continuous*.

Distribution, probability and density function

The *distribution* of random variable X is the function $F : \mathbb{R} \longrightarrow [0, 1]$ defined by

$$F(x) = \Pr(X \leq x) = \Pr(\{s \in \mathcal{S} \mid X(s) \leq x\}).$$

If random variable X is discrete, it can only take finitely many values. Suppose that the range of random variable X is the set $\{x_1, x_2, \ldots, x_N\}$, with $x_1 < x_2 < \ldots < x_N$. Then, the *probability function* of X is the function $f : \mathbb{R} \longrightarrow [0, 1]$ defined by

$$f(x) = \begin{cases} F(x_1), & \text{if } x = x_1; \\ F(x_n) - F(x_{n-1}), & \text{if } x = x_n \text{ for some } n = 2, 3, \ldots N; \\ 0, & \text{elsewhere.} \end{cases}$$

It is true, by construction, that $f(x) = \Pr(\{s \in \mathcal{S} \mid X(s) = x\})$, which is why the function receives its name. In this case, the distribution function is also referred to as *cumulative probability function*, or CPF.

Suppose now that the random variable is continuous, and let $[\underline{x}, \bar{x}]$ be its range. If the distribution of X is differentiable, we call its first derivative the *density function* of X. More precisely, such function is $f : \mathbb{R} \to [0, 1]$ defined by

$$f(x) = \begin{cases} F'(x), & \text{if } \underline{x} < x < \bar{x}; \\ 0, & \text{elsewhere.} \end{cases}$$

By construction

$$\Pr(x_* \leq X \leq x^*) = \int_{x_*}^{x^*} f(x)\,\mathrm{d}x,$$

but $f(x)$ is *not* the probability that X will take the value x (which is null). Instead the density function measures the intensity at which the probability of an event increases if a neighborhood of the corresponding state is included in the event. In the case of continuous random variables, the distribution function is also called *cumulative density function* or CDF.

Moments of a random variable

Let X be a random variable with distribution F, and let $g : \mathbb{R} \to \mathbb{R}$. Obviously, $Y = g(X)$ is a random variable too.

If X is discrete and its range is $\{x_1, x_2, \ldots, x_N\}$, we define the *expectation* of Y as

$$\mathrm{E}[Y] = \sum_{n=1}^{N} g(x_n) f(x_n),$$

where f is the probability function of X. In terms of the fundamental state space,

$$\mathrm{E}[Y] = \sum_{s=1}^{\bar{s}} g(x(s)) p_s.$$

The expectation of X itself, $\mathrm{E}[X] = \sum_{n=1}^{N} x_n f(x_n)$ is also known as its *mean*.

If X is continuous and its range is $[\underline{x}, \bar{x}]$, we define the *expectation* of Y as

$$E[Y] = \int_{\underline{x}}^{\bar{x}} g(x) f(x)\,\mathrm{d}x,$$

where f is the density function of X, assuming that the integral exists.

In any case, the *variance* of random variable X is defined as $\mathrm{Var}[X] = \mathrm{E}\left[(X - \mathrm{E}[X])^2\right]$. This number is always non-negative. It is zero when, and only when, the variable takes one value with probability 1, in which case it is said to be *deterministic* or *degenerate*.

The expectation of a random variable is also known as its *first moment*, while its variance is often called its *second central moment*. Important properties of these numbers are:

THEOREM. *Let X be a random variable with expectation $\mathrm{E}[X] = \bar{X}$ and variance $\mathrm{Var}[X] = \Sigma$. Then, for any constants α and β, $\mathrm{E}[\alpha X + \beta] = \alpha \bar{X} + \beta$ and $\mathrm{Var}[\alpha X + \beta] = \alpha^2 \Sigma$.*

In general, unless $g(x) = \alpha x + \beta$, $\mathrm{E}[g(X)] \neq g(\mathrm{E}[X])$.

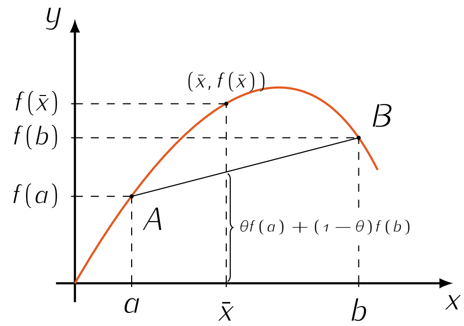

Figure App.1: A concave function

Concave and convex functions

Let I be an interval on the real line. Function f with domain I is said to be *concave* if for all distinct points a and b in I, and for all numbers $0 < \theta < 1$, it is true that

$$f(\theta a + (1-\theta)b) \geq \theta f(a) + (1-\theta) f(b). \qquad (*)$$

It is said to be *convex* if for all distinct points c and d in I, and all numbers $0 < \theta < 1$,

$$f(\theta c + (1-\theta)d) \leq \theta f(c) + (1-\theta) f(d)) \qquad (**)$$

Consider the points $A = (a, f(a))$ and $B = (b, f(b))$ in Fig. App.1. As we let θ range from 0 to 1, we can obtain *any* $x = \theta a + (1-\theta)b$ between a and b. The straight segment connecting A and B consists of the points with coordinates

$$(\theta a + (1-\theta)b, \theta f(a) + (1-\theta) f(b))$$

for $0 \leq \theta \leq 1$. Now, fix any point x between a and b. The RHS of the inequality in $(*)$ is the vertical distance from the horizontal axis to the straight segment. The vertical distance to the curve is the value of the function evaluated at $\bar{x}$, namely the LHS of the inequality. Concavity is the requirement that the latter distance be at least as much as the former, which is to say that the straight segment cannot be above the graph of the function, as in Fig. App.1

Function f is *strictly concave* if Eq. $(*)$ is always satisfied with strict inequality, and *strictly convex* if Eq. $(**)$ is always satisfied with strict inequality.

Note that given a function f, if we let function g be defined by $g(x) = -f(x)$, then f is (strictly) concave if and only if g is (strictly) convex. More importantly, suppose that f is continuous in the interval I and twice differentiable in the interior of I. Then,

$$\begin{aligned} f \text{ is concave on } I &\Leftrightarrow f''(x) \leq 0 \text{ for all } x \text{ in } I \\ f \text{ is convex on } I &\Leftrightarrow f''(x) \geq 0 \text{ for all } x \text{ in } I \\ f \text{ is strictly concave on } I &\Leftarrow f''(x) < 0 \text{ for all } x \text{ in } I \\ f \text{ is strictly convex on } I &\Leftarrow f''(x) < 0 \text{ for all } x \text{ in } I \end{aligned}$$

Partial differentiation

If $z = f(x, y)$, then $\partial z / \partial x$ is the derivative of $f(x, y)$ seen as a function of x, when y is held constant, while $\partial z / \partial y$ is the derivative of $f(x, y)$ seen as a function of y, when x is held constant. Formally,

$$\frac{\partial z}{\partial x} = \frac{\partial f}{\partial x}(x, y) = \lim_{d \to 0} \frac{f(x+d, y) - f(x, y)}{d}$$

and, similarly,

$$\frac{\partial z}{\partial y} = \frac{\partial f}{\partial y}(x, y) = \lim_{d \to o} \frac{f(x, y + d) - f(x, y)}{d}$$

provided that the limits exist.

A very useful result, known as the *chain rule* is that if $z = f(x, y)$ with $x = g(t, s)$ and $y = h(t, s)$, then

$$\frac{\partial z}{\partial t} = \frac{\partial z}{\partial x}\frac{\partial x}{\partial t} + \frac{\partial z}{\partial y}\frac{\partial y}{\partial t}$$

and

$$\frac{\partial z}{\partial s} = \frac{\partial z}{\partial x}\frac{\partial x}{\partial s} + \frac{\partial z}{\partial y}\frac{\partial y}{\partial s}$$

Also given the function $z = f(x, y)$, if we fix a constant c and $\partial z/\partial y \neq o$, then the equation $f(x, y) = c$ *defines implicitly* a function $y(x)$ with

$$y' = -\frac{\partial z/\partial x}{\partial z/\partial y}.$$

This result is known as the *implicit function theorem*.

Calculus of variations

Suppose that we have to solve the variational problem

$$\max_{x:[t_*, t^*] \to \mathbb{R}} \int_{t_*}^{t^*} I[x(t), \dot{x}(t), t] \, \mathrm{d}t,$$

subject to $x(t_*) = x_*$ and $x(t^*) = x^*$. This problem is different from the standard constrained optimization problem is that the control is a function, so the objective is an *functional*, not a function.

The *Euler equation* for this problem is that. at all $t \in [t_*, t^*]$,

$$\frac{\partial I}{\partial x} - \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial I}{\partial \dot{x}}\right) = o.$$

If the integrand is autonomous, in the sense that I does not depend on $\dot{x}$, the Euler equation reduces to $\partial I/\partial x = o$, which is the equation we used in the text.

Index