

Microeconomic Theory 3

Course Notes

by

Lutz-Alexander Busch

Dept. of Economics
University of Waterloo

Revised 2004

©Lutz-Alexander Busch, 1994,1995,1999,2001,2002,2004
Do not quote or redistribute without permission

Contents

1 Preliminaries	1
1.1 About Economists	2
1.2 What is (Micro—)Economics?	2
1.3 Economics and Sex	5
Notation	8
2 Review	9
2.1 Modelling Choice Behaviour ¹	9
2.1.1 Choice Rules	10
2.1.2 Preferences	11
2.1.3 What gives?	13
2.2 Consumer Theory	15
Special Utility Functions	18
2.2.1 Utility Maximization	20
Properties	22
2.2.2 Expenditure Minimization and the Slutsky equation . .	25
2.3 General Equilibrium	27

¹This material is based on Mas-Colell, Whinston, Green, chapter 1

2.3.1	Pure Exchange	27
2.3.2	A simple production economy	33
2.4	Review Problems	35
3	Inter-temporal Economics	39
3.1	The consumer's problem	39
3.1.1	Deriving the budget set	40
	No Storage, No Investment, No Markets	40
	Storage, No Investment, No Markets	40
	Storage, Investment, No Markets	41
	Storage, No Investment, Full Markets	42
	Storage, Investment, Full Markets	43
3.1.2	Utility maximization	44
3.2	Real Interest Rates	47
3.3	Risk-free Assets	47
3.3.1	Bonds	48
3.3.2	More on Rates of Return	49
3.3.3	Resource Depletion	51
3.3.4	A Short Digression into Financial Economics	52
3.4	Review Problems	54
4	Uncertainty	57
4.1	Risk Aversion	61
4.1.1	Comparing degrees of risk aversion	65

4.2	Comparing gambles with respect to risk	67
4.3	A first look at Insurance	69
4.4	The State-Preference Approach	72
4.4.1	Insurance in a State Model	74
4.4.2	Risk Aversion Again	76
4.5	Asset Pricing	77
4.5.1	Diversification	77
4.5.2	Risk spreading	78
4.5.3	Back to Asset Pricing	79
4.5.4	Mean-Variance Utility	81
4.5.5	CAPM	84
4.6	Review Problems	84
5	Information	91
5.1	Search	95
5.2	Adverse Selection	98
5.3	Moral Hazard	103
5.4	The Principal Agent Problem	105
5.4.1	The Abstract P-A Relationship	105
6	Game Theory	113
6.1	Descriptions of Strategic Decision Problems	117
6.1.1	The Extensive Form	117
6.1.2	Strategies and the Strategic Form	121

6.2	Solution Concepts for Strategic Decision Problems	126
6.2.1	Equilibrium Concepts for the Strategic Form	127
6.2.2	Equilibrium Refinements for the Strategic Form	131
6.2.3	Equilibrium Concepts and Refinements for the Extensive Form	134
	Signalling Games	138
6.3	Review Problems	140
7	Review Question Answers	145
7.1	Chapter 2	145
7.2	Chapter 3	155
7.3	Chapter 4	159
7.4	Chapter 6	169

List of Figures

2.1	Consumer Optimum: The tangency condition	22
2.2	Corner Solutions: Lack of tangency	23
2.3	An Edgeworth Box	29
2.4	Pareto Optimal Allocations	30
3.1	Storage without and with markets, no (physical) investment .	44
4.1	Risk Aversion	62
4.2	Risk Neutral and Risk Loving	63
4.3	The certainty equivalent to a gamble	64
4.4	Comparing two gambles with equal expected value	67
4.5	An Insurance Problem in State-Consumption space	75
4.6	Risk aversion in the State Model	76
4.7	Efficient Portfolios and the Market Portfolio	83
5.1	Insurance for Two types	99
5.2	Impossibility of Pooling Equilibria	101
5.3	Separating Contracts could be possible	102
5.4	Separating Contracts definitely impossible	103

6.1	Examples of valid and invalid trees	118
6.2	A Game of Perfect Recall and a Counter-example	120
6.3	From Incomplete to Imperfect Information	121
6.4	A Matrix game — game in strategic form	124
6.5	A simple Bargaining Game	125
6.6	The 4 standard games	126
6.7	Valid and Invalid Subgames	135
6.8	The “Horse”	137
6.9	A Signalling Game	139
6.10	A Minor Perturbation?	140

Chapter 1

Preliminaries

These notes have been written for Econ 401 as taught by me at the University of Waterloo. As such the list of topics reflects the course material for that particular course. It is assumed that the student has mastered the pre-requisites and little or no time is spent on them, aside from a review of standard consumer economics and general equilibrium in chapter 1. I assume that the student has access to a standard undergraduate micro theory text book. Any of the books commonly used will do and will give introductions to the topics covered here, as well as allowing for a review, if necessary, of the material from the pre-requisites.

These notes will not give references. The material covered is by now fairly standard and can be found in one form or another in most micro texts. I wish to acknowledge two books, however, which have served as references: the most excellent book by Mas-Colell, Whinston, and Green, *Microeconomic Theory*, as well as the much more concise Jehle and Reny, *Advanced Microeconomic Theory*. I also would like to acknowledge my teachers Don Ferguson and Glen MacDonald, who have done much to bring microeconomics alive for me.

This preliminary chapter contains extensive quotes which I have found informative, amusing, interesting, and thought provoking. Their sources have been indicated.

1.1 About Economists

By now you may have heard many jokes about economists and noticed that modern economics has a bad reputation in some circles. If you mention your field of study in the bar, you are quite possibly forced to defend yourself against various stereotypical charges (focusing on assumptions, mostly.) The most eloquent quotes about economists that I know of are the following two quotes reproduced from the *Economist*, Sept.4, 1993, p.25:

No real Englishman, in his secret soul, was ever sorry for the death of a political economist, he is much more likely to be sorry for his life. You might as well cry at the death of a cormorant. Indeed how he can die is very odd. You would think a man who could digest all that arid matter; who really preferred 'sawdust without butter'; who liked the tough subsistence of rigid formulae, might defy by intensity of internal constitution all stomachic or lesser diseases. However they do die, and people say that the dryness of the Sahara is caused by a deposit of similar bones.
(Walter Bagehot (1855))

Are economists human? By overwhelming majority vote, the answer would undoubtedly be No. This is a matter of sorrow for them, for there is no body of men whose professional labours are more conscientiously, or consciously, directed to promoting the wealth and welfare of mankind. That they tend to be regarded as blue-nosed kill-joys must be the result of a great misunderstanding. (Geoffrey Crowther (1952))

1.2 What is (Micro—)Economics?

In Introductory Economics the question of what economics *is* has received some attention. Since then, however, this question may have received no further coverage, and so I thought to collect here some material which I to use to start a course. It is meant to provide a background for the field as well as a defense, of sorts, of the way in which micro economics is practiced.

Malinvaud sees economics as follows:

Here we propose the alternative, more explicit definition: *economics is the science which studies how scarce resources are employed for the satisfaction of the needs of men living in society: on the one hand, it is interested in the essential operations of production, distribution and consumption of goods, and on the other hand, in the institutions and activities whose object it is to facilitate these operations.* [...]

The main object of the theory in which we are interested is the analysis of the simultaneous determination of prices and the quantities produced, exchanged and consumed. It is called microeconomics because, in its abstract formulations, it respects the individuality of each good and each agent. This seems a necessary condition *a priori* for logical investigation of the phenomena in question. By contrast, the rest of economic theory is in most cases macroeconomic, reasoning directly on the basis of aggregates of goods and agents.

[E. Malinvaud, *Lectures on Microeconomic Theory*, revised, N-H, 1985, p.1-2.]

This gives us a nice description of what economics is, and in particular what micro theory entails. In following the agenda laid out by Malinvaud a certain amount of theoretical abstraction and rigor have been found necessary, and one key critique heard often is the “attempt at overblown rigor” and the “unrealistic assumptions” which micro theory employs. Takayama and Hildenbrand both address these criticisms in the opening pages of their respective books. First Takayama:

The essential feature of modern economic theory is that it is analytical and mathematical. Mathematics is a language that facilitates the honest presentation of a theory by making the assumptions explicit and by making each step of the logical deduction clear. Thus it provides a basis for further developments and extensions. Moreover, it provides the possibility for more accurate empirical testing. Not only are some assumptions hidden and obscured in the theories of the verbal and “curve-bending” economic schools, but their approaches provide no scope for accurate empirical testing, simply because such testing requires explicit and mathematical representations of the propositions of the theories to be tested.

[...] But yet, economics is a complex subject and involves many things that cannot be expressed readily in terms of mathematics.

Commenting on Max Planck's decision not to study economics, J.M. Keynes remarked that economics involves the "amalgam of logic and intuition and wide knowledge of facts, most of which are not precise." In other words, economics is a combination of poetry and hard-boiled analysis accompanied by institutional facts. This does not imply, contrary to what many poets and institutionalists feel, that hard-boiled analysis is useless. Rather, it is the best way to express oneself honestly without being buried in the millions of institutional facts. [...]

Mathematical economics is a field that is concerned with complete and hard-boiled analysis. The essence here is the *method* of analysis and not the resulting collection of theorems, for actual economies are far too complex to allow the ready application of these theorems. J.M. Keynes once remarked that "the theory of economics does not furnish a body of settled conclusions immediately applicable to policy. It is a method rather than a doctrine, an apparatus of the mind, a technique of thinking, which helps its possessor to draw conclusions."

An immediate corollary of this is that the theorems are useless without explicit recognition of the assumptions and complete understanding of the logic involved. It is important to get an intuitive understanding of the theorems (by means of diagrams and so on, if necessary), but this understanding is useless without a thorough knowledge of the assumptions and proofs.

[Akira Takayama, *Mathematical Economics*, 2nd ed., Cambridge, 1985, p. xv.]

Hildenbrand offers the following:

I cannot refrain from repeating here the quotation from Bertrand Russell cited by F. Hahn in his inaugural lecture in Cambridge: "Many people have a passionate hatred of abstraction, chiefly, I think, because of its intellectual difficulty; but as they do not wish to give this reason they invent all sorts of other that sound grand. They say that all abstraction is falsification, and that as soon as you have left out any aspect of something actual you have exposed yourself to the risk of fallacy in arguing from its remaining aspects alone. Those who argue that way are in fact concerned with matters quite other than those that concern science." (footnote 2, p.2, with reference)

Let me briefly recall the main characteristics of an axiomatic theory of a certain economic phenomenon as formulated by Debreu:

First, the primitive concepts of the economic analysis are selected, and then, each one of these primitive concepts is represented by a mathematical object.

Second, assumptions on the mathematical representations of the primitive concepts are made explicit and are fully specified. Mathematical analysis then establishes the consequences of these assumptions in the form of theorems.

[Werner Hildenbrand, *Twenty Papers of Gerard Debreu*, Econometric Society Monograph 4, Cambridge, 1983, page 4, quoted with omissions.]

1.3 Economics and Sex

I close this chapter with the following thought provoking excerpt from Mark Perlman and Charles R. McCann, Jr., “Varieties of uncertainty,” in *Uncertainty in Economic Thought*, ed. Christian Schmidt, Edward Elgar 1996, p 9-10.

The problem as perceived

As this is an opening paper, let us begin with what was once an established cultural necessity, namely a reference to our religious heritage. What we have in mind is the Biblical story of the Fall of Man, the details of which we shall not bore you with. Rather, we open consideration of this difficult question by asking what was the point of that Book of Genesis story about the inadequacy of Man.

We are told that apparently whatever were God’s expectations, He became disappointed with Man. Mankind and *particularly Womankind*¹ did not live up to His expectations.² In any case, Adam and Eve were informed

¹Much has been made of the failure of women, perhaps that is because men wrote up the history. We should add, in order to avoid deleterious political correctness (and thereby cut off provocation and discussion), that since Eve was the proximate cause of the Fall, and Eve represents sexual attraction or desire, some (particularly St Paul, whose opinion of womankind was problematic) have considered that sexual attraction was in some way even more responsible for the Fall than anything else. **Put crudely, even if economics is not a sexy subject, its origins were sexual.**

²What that says about His omniscience and/or omnipotence is, at the very least, para-

that they had ‘fallen’ from Grace, and all of us have been made to suffer ever since.

From our analytical standpoint there are two crucial questions:

1. What was the sin; and
2. What was the punishment?

The sin seems to have been something combining (1) an inability to follow precise directions; (2) a willingness to be tempted, particularly when one could assert that ‘one was only doing what everyone else (sic) was doing;’³ (3) a greed involving things (something forbidden) and time (instant gratification); (4) an inability to leave well enough alone; and (5) an excessive Faustian curiosity. Naturally, as academic intellectuals, we fancy the fifth reason as best.

But what interests us directly is the second question. It is ‘What was God’s punishment for Adam and Eve’s vicarious sin, for which all mankind suffers?’ Purportedly a distinction has been made between what happened to Man and Woman, but, the one clear answer, *particularly as seen by Aquinas and by most economists ever since*, was that man is condemned to live with the paradigm of scarcity of goods and services and with a schedule of appetites and incentives which are, at best, confusing.

In the more modern terms of William Stanly Jevons, ours is a world of considerable pain and a few costly pleasures. We are driven to produce so that we can consume, and production is done mostly by the ‘sweat of the brow’ and the strength of the back. The study of economics — of the production, distribution and even the consumption of goods and services — it follows, is the result of the Original Sin. When Carlyle called Economics the ‘Dismal Science’, he was, if anything, writing in euphemisms; Economics *per se*, is the Punishment for Sin.

doxical.

³Cf. Genesis, 3:9-12,16,17. [9] But the Lord God called to the man and said to him, ‘Where are you?’ [10] He replied, ‘I heard the sound as you were walking the garden, and I was afraid because I was naked, and I hid myself.’ [11] God answered, ‘Who told you that you were naked? Have you eaten from the tree which I forbade you?’ [12] The man said, ‘The woman you gave me for a companion, she gave me fruit from the tree and I ate.’ [Note: The story, as recalled, suggests that Adam was dependent upon Eve (for what?), and the price of that dependency was to be agreeable to Eve (‘It was really all her fault — I only did what You [God] had laid out for me.’)] ([16] and [17] omitted) [Again, for those civil libertarians amongst us, kindly note that God forced Adam to testify against himself. Who says that the Bill of Rights is an inherent aspect of divine justice? Far from it, in the Last Judgment, pleading the Fifth won’t do at all.]

But, it is another line of analysis, perhaps novel, which we put to you. Scarcity, as the paradigm, may not have been the greatest punishment, because scarcity, as such, can usually be overcome. Scarcity simply means that one has to allocate between one's preferences, and the thinking man ought to be able to handle the situation. We use our reasoning power, surely tied up with Free Will, to allocate priorities and thereby overcome the greater disasters of scarcity. What was the greater punishment, indeed the greatest punishment, is more basic. Insofar as we are aware, it was identified early on by another Aristotelian, one writing shortly before Aquinas, Moses Maimonides. Maimonides suggested that God's real punishment was to push man admittedly beyond the limits of his reasoning power. Maimonides held that prior to the Fall, Adam and Eve (and presumably mankind, generally) knew everything concerning them; after the Fall they only had *opinions*.⁴ Requisite to the wise use of power is understanding and full specification; what was lost was any such claim previously held by man to complete knowledge and the full comprehension of his surroundings. In other words, what truly underlies the misery of scarcity is neither hunger nor thirst, but the lack of knowledge of what one's preference schedule will do to one's happiness. For if one had complete knowledge (including foreknowledge) one could compensate accordingly.

If one pursues Maimonides' line of inquiry, it seems that uncertainty (which is based not on ignorance of what can be known with study of data collection, but also on ignorance tied to the unknowable) is the real punishment.

⁴[omitted]

Notation:

Blackboard versions of these symbols may differ slightly.

I will not distinguish between vectors and scalars by notation. Generally all small variables (x, y, p) are column vectors (even if written in the notes as row vectors to save space.) The context should clarify the usage. Capital letters most often denote sets, as in the consumption set X , or budget set B . Sets of sets are denoted by capital script letters, such as $\mathcal{X} = \{X_1, X_2, \dots, X_k\}$, where $X_i = \{x \in \mathfrak{R}^N \mid \sqrt{\sum_n x_n^2} = i\}$.

I will use the usual partial derivative notation $\partial f / \partial x_1 = f_1(\cdot)$, if no confusion arises, and will often omit the arguments to a function but indicate that there is a function by the previous notation, i.e., $f(\cdot)$ denotes $f(x_1, x_2, \dots, x_n)$, for example. Finally, the Reals \mathfrak{R} will be written as \mathbb{R} on the board, in the more familiar fashion.

\mathfrak{R}	the real numbers, superscript denotes dimensionality
\mathfrak{R}_+	the non-negative reals (include 0); (\mathfrak{R}_{++} are the positive Reals)
\forall	“for all”, for all elements in the set
\exists	“there exists” or “there is at least one element”
\neg	“not”, or negation, the following statement is not true
\cdot	dot product, as in $x \cdot y = x^T y = \sum_n x_i y_i$
\in	“in”, or “element of”
\ni	“such that”. Note: $\forall x \in X \ni A$ is equivalent to $\{x \in X \mid A\}$.
\leq	less or equal to, component wise: $x_n \leq y_n$ for all $n = 1, \dots, N$.
\ll	strictly less than, component wise: $\forall n \in N : x_n < y_n$
\geq	greater or equal (component wise)
\gg	strictly greater (component wise)
\succ	strictly preferred to
\succeq	weakly preferred to
\sim	indifferent to
∂	partial, as in partial derivatives
∇	gradient, the vector of partial derivatives of a function of a vector. $\nabla f(x) = [f_1(\cdot), f_2(\cdot), \dots, f_n(\cdot)]^T$
D_x	derivative operator, the vector (matrix) of partial derivatives of a vector of functions: $D_w x(p, w) = \left[\frac{\partial x_1(\cdot)}{\partial w}, \frac{\partial x_2(\cdot)}{\partial w}, \dots, \frac{\partial x_N(\cdot)}{\partial w} \right]^T$
\iff	“if and only if”, often written “iff”

Chapter 2

Review

Consumer theory is concerned with modelling the choices of consumers and predicting the behaviour of consumers in response to changes in their choice environment (comparative statics.) This raises the question of how one can model choice formally, or more precisely, how one can model decision making. An extensive analysis of this topic would lead us way off course, but the two possible approaches are both outlined in the next section. After that, we will return to the standard way in which consumer decision making in particular is addressed.

2.1 Modelling Choice Behaviour ¹

Any model of choice behaviour at a minimum contains two elements: (1) a description of the available choices/decisions. (2) a description of how/why decisions are made. The first part of this is relatively easy. We specify some set B which is the set of (mutually exclusive) alternatives. For a consumer this will usually be the budget set, i.e., the set of amounts of commodities which can be purchased at current prices and income levels. In principle this set can be anything, however. A simple description of this set will be advantageous, and thus we most often will have $B = \{x \in \mathfrak{R}_+^n \mid p \cdot x \leq m\}$, the familiar budget set for a price taking consumer. Note that much of economics is concerned with how choice changes as B changes. This is the standard question of *comparative statics*, where we ask how choice adjusts as income, m or one of the prices, p_i , $i \in \{1, \dots, n\}$, changes. It does require, however, some additional

¹This material is based on Mas-Colell, Whinston, Green, chapter 1

structure as to what kind of B are possible. So we get a set of sets, \mathcal{B} , which contains all sets B we could consider. For example, for our neo-classical consumer this is $\mathcal{B} = \{\{x \in \mathbb{R}_+^n \mid p \cdot x \leq m\}, p \in S^{n-1}, m \in \mathbb{R}, m > 0\}$. We also need a description of the things which are assembled into these budget sets. In consumer theory this would be a consumption bundle, that is, a vector x indicating the quantities of all the goods. We normally let X be the set of technically feasible choices, or **consumption set**. This set does not take into account what is an economically feasible choice, but only what is in principle possible (for example, it would contain 253 Lear Jets even if those are not affordable to the consumer in question, but it will not allow the consumer to create time from income, say, since that would be physically impossible.)

For the second element, the description of choice itself, there are two fundamental approaches we can take: (1) we can specify the choice for each and every choice set, that is, we can specify *choice rules* directly. For our application to the consumer this basically says that we specify the demand functions of the consumer as the primitive notion in our theory.

Alternatively, (2), we can specify some more primitive notion of “preferences” and derive the choice for each choice set from those. Under this approach a ranking of choices in X is specified, and then choice for a given choice set B is determined by the rule that the highest ranked bundle in the set is taken.

In either case we will need to specify what we mean by “rational”. This is necessary since in the absence of such a restriction we would be able to specify quite arbitrary choices. While that would allow our theory to match any data, it also means that no refutable propositions can be derived, and thus the theory would be useless. Rationality restrictions can basically be viewed as consistency requirements: It should not be possible to get the consumer to engage in behaviour that contradicts previous behaviour. More precisely, if the consumer’s behaviour reveals to us that he seems to like some alternative a better than another, b , by virtue of the consumer selecting a when both are available, this conclusion should not be contradicted in other observations. The way we introduce this restriction into the model will depend on the specific approach taken.

2.1.1 Choice Rules

Since we will not, in the end, use this approach, I will give an abstract description without explicit application to the consumer problem. We are

to specify as the primitive of our theory the actual choice for each possible set of choices. This is done by means of a **choice structure**. A choice structure contains a set of nonempty subsets of X , denoted by \mathcal{B} . For the consumer this would be the collection of all possible budget sets which he might encounter. For each actual budget set $B \in \mathcal{B}$, we then specify a **choice rule** (or correspondence). It assigns a set of chosen elements, denoted $C(B) \subset B$. In other words, for each budget set B we have a “demand function” $C(B)$ specified directly.

The notion of rationality we impose on this is as follows:

Definition 1 *A choice structure satisfies the **weak axiom of revealed preference** if it satisfies the following restriction: If ever the alternatives x and y are available and x is chosen, then it can never be the case that y is chosen and x is not chosen, if both x and y are available. More concisely: If $\exists B \in \mathcal{B} \ni x, y \in B \Rightarrow x \in C(B)$ then $\forall B' \in \mathcal{B} \ni x, y \in B'$ and $y \in C(B') \Rightarrow x \in C(B')$.*

For example, suppose that $X = \{a, b, c, d\}$, and that $\mathcal{B} = \{ \{a, b\}, \{a, b, c\}, \{b, c, d\} \}$. Also suppose that $C(\{a, b\}) = \{a\}$. While the above criterion is silent on the relation of this choice to $C(\{b, c, d\}) = \{b\}$, it does rule out the choice $C(\{a, b, c\}) = \{b\}$. Why? In the first case, a is not an element of both $\{a, b\}$ and $\{b, c, d\}$, so the definition is satisfied vacuously. (Remember, any statement about the empty set is true.) In the second case, the fact that $C(\{a, b\}) = \{a\}$ has revealed a preference of the consumer for a over b , after all both were available but only a was picked. A rational consumer should preserve this ranking in other situations. So if the choice is over $\{a, b, c\}$, then b cannot be picked, only a or c , should c be even better than a .²

You can see how we are naturally lead to think of one choice being better than, or worse than, another even when discussing this abstract notion of direct specification of choice rules. The second approach to modelling decision making, which is by far the most common, does this explicitly.

2.1.2 Preferences

In this approach we specify a more primitive notion, that of the consumer being able to rank any two bundles. We call such a ranking a **preference**

²Just let x be a and y be b in the above definition!

relation which is usually denoted by the symbol \succeq . So, define a binary relation \succeq on X , with the interpretation that for any $x, y \in X$, the statement $x \succeq y$ means that “ x is at least as good as y ”.

From this notion of weak preference we can easily define the notions of strict preference ($x \succ y \Leftrightarrow (x \succeq y \text{ and } \neg(y \succeq x))$), and indifference ($x \sim y \Leftrightarrow (x \succeq y \text{ and } y \succeq x)$).

How is rationality imposed on preferences?

Definition 2 *A preference relation \succeq is rational if it is a complete, reflexive, and transitive binary relation.*

complete: $\forall x, y \in X$, either $x \succeq y$ or $y \succeq x$, or both;

reflexive: $\forall x \in X, x \succeq x$;

transitive $\forall x, y, z \in X, x \succeq y \text{ and } y \succeq z \implies x \succeq z$.

At this point we might want to consider how restrictive this requirement of rational preferences is. Completeness simply says that any two feasible choices can be ranked. This could be a problem, since what is required is the ability to rank things which are very far from actual experience. For example, the consumer is assumed to be able to rank two consumption bundles which differ only in the fact that in the first he is to consume 2 Ducati 996R, 1 Boeing 737, 64 cans of Beluga caviar, and 2 bottles of Dom Perignon, while the other contains 1 Ducati 996R, 1 Aprilia Mille RS, 1 Airbus 320, 57 cans of Beluga caviar, and 2 bottles of Veuve Cliquot. For most consumers these bundles may be difficult to rank, since they may be unsure about the relative qualities of the items. On the other hand, these bundles are very far from the budget set of these consumers. For bundles close to being affordable, it may be much more reasonable to assume that any two can be ranked against each other.

Transitivity could also be a problem. It is crucial to economics since it rules out cycles of preference. This, among other important implications, rules out “dutch books” — a sequence of trades in which the consumer pays at each step in order to change his consumption bundle, only to return to the initial bundle at the end. Yet, it can be shown that it is possible to have real life consumers violate this assumption. One problem is that of “just perceptible differences”. If the changes in the consumption bundle are small enough, the consumer can be led through a sequence of bundles in which he prefers any earlier bundle to a later bundle, but which closes up on the first, so that the last bundle in the sequence is preferred to the first, even though it was worse than all preceding bundles, including the first. The

other common way in which transitivity is violated is through the effects of the “framing problem”. The framing problem refers to the fact that a consumer’s stated preference may depend on the way the question is asked. A famous example of this are surveys about vaccinations. If the benefits are stressed (so many fewer crippled and dead) respondents are in favour, while they are against vaccination programs if the costs are stressed (so many deaths which would not have occurred otherwise.) It is worthwhile to note that what looks like intransitivity can quite often be explained as the outcome of transitive behaviour on more primitive characteristics of the problem. Also note that addiction and habit formation commonly lead to what appears to be intransitivity, but really involves a change in the underlying preferences, or (more commonly) preferences which depend on past consumption bundles as well. In any case, a thorough discussion of this topic is way beyond these notes, and we will assume rational preferences henceforth.

While beautiful in its simplicity, this notion of preferences on sets of choices is also cumbersome to work with for most (but note the beautiful manuscript by Debreu, *Theory of value*). In particular, it would be nice if we could use the abundance of calculus tools which have been invented to deal with optimization problems. Thus the notion of a **utility function** is useful.

Definition 3 *A function $u : X \mapsto \mathfrak{R}$ is a utility function representing the preference relation \succeq if $\forall x, y \in X$, $(x \succeq y) \Leftrightarrow (u(x) \geq u(y))$.*

We will return to utility functions in more detail in the next section.

Choice now is modeled by specifying that the consumer will choose the highest ranked available bundle, or alternatively, that the consumer will maximize his utility function over the set of available alternatives, choosing the one which leads to the highest value of the utility function (for details, see the next section.)

2.1.3 What gives?

If there are two ways to model choice behaviour, which is “better”? How do they compare? Are the two approaches equivalent? Why do economists not just write down demand functions, if that is a valid method?

To answer these questions we need to investigate the relationship between the two approaches. What is really important for us are the following

two questions: If I have a rational preference relation, will it generate a choice structure which satisfies the weak axiom? (The answer will be Yes.) If I have a choice structure that satisfies the weak axiom, does there exist a rational preference relation which is consistent with these choices? (The answer will be Maybe.)

The following statements can be made:

Theorem 1 *Suppose \succeq is rational. Then the choice structure $(\mathcal{B}, C(\cdot)) = (\mathcal{B}, \{x \in B \mid x \succeq y \ \forall y \in B\})$ satisfies the weak axiom.*

In other words, rational preferences always lead to a choice structure which is consistent with the weak axiom (and hence satisfies the notion of rationality for choice structures.)

Can this be reversed? Note first that we only need to deal with $B \in \mathcal{B}$, not all possible subsets of X . Therefore in general more than one \succeq might do the trick, since there could be quite a few combinations of elements for which no choice is specified by the choice structure, and hence, for which we are free to pick rankings. Clearly,³ without the weak axiom there is no hope, since the weak axiom has the flavour of “no contradictions” which transitivity also imposes for \succeq . However, the following example shows that the weak axiom is not enough: $X = \{a, b, c\}$, $\mathcal{B} = \{ \{a, b\}, \{b, c\}, \{c, a\} \}$, $C(\{a, b\}) = \{a\}$, $C(\{b, c\}) = \{b\}$, $C(\{c, a\}) = \{c\}$. This satisfies the weak axiom (vacuously) but implies that a is better than b , b is better than c , and c is better than a , which violates transitivity. Basically what goes wrong here is that the set \mathcal{B} is not rich enough to restrict $C(\cdot)$ much. However, if it were, there would not be a problem as the following theorem asserts:

Theorem 2 *If $(\mathcal{B}, C(\cdot))$ is a choice structure such that (i) the weak axiom is satisfied, and (ii) \mathcal{B} includes all 2 and 3 element subsets of X , then there exists a rational preference ordering \succeq such that $\forall B \in \mathcal{B}$, $C(B) = \{x \in B \mid x \succeq y \ \forall y \in B\}$.*

While this is encouraging, it is of little use to economists. The basic problem is that we rarely if ever have collections of budget sets which satisfy this restriction. Indeed, the typical set of consumer budget sets consists of

³Words such as “clearly”, “therefore”, “thus”, etc. are a signal to the reader that the argument should be thought about and is expected to be known, or at least easily derivable.

nested hyper planes. One way around this problem are the slightly stronger axioms required by Revealed Preference Theory (see Varian's *Advanced Microeconomic Theory* for details.)

Hence, the standard approach to consumer theory employs preferences as the fundamental concept. The next section will review this approach.

2.2 Consumer Theory

Consumers are faced with a **consumption set** $X \subset \mathfrak{R}_+^n$ which is the set of all (non-negative) vectors⁴ $x = (x_1, x_2, \dots, x_n)$ where each coordinate x_i indicates the amount desired/consumed of commodity i .⁵ All commodities are assumed infinitely divisible for simplicity. Often we will have only two commodities to allow simple graphing. A key concept to remember is that a “commodity” is a completely specified good. That means that not only is it a ‘car’, but its precise type, colour, and quality are specified. Furthermore, it is specified when and where this car is consumable, and the circumstances under which it can be consumed. In previous courses these latter dimensions have been suppressed and the focus was on current consumption without uncertainty, but in this course we will focus especially on these latter two. So we will have, for example, today’s consumption versus tomorrow’s, or consumption if there is an accident versus consumption if there is no accident (i.e., consumption contingent on the state of the world.) The consumption set incorporates all physical constraints (no negative consumption, no running time backwards, etc) as well as all institutional constraints. It does not include the economic constraints the consumer faces.

The economic constraints come chiefly in two forms. One is an assumption about what influence, if any, the consumer may have on the prices which are charged in the market. The usual assumption is that of **price taking**, which is to say that the consumer does not have any influence on price, or more precisely, acts under the assumption of not having any influence on price. (Note that there is a slight difference between those two statements!) This assumption can be relaxed, of course, but then we have to specify how price reacts to the consumer’s demands, or at least how the consumer thinks

⁴While I write all vectors as row vectors in order to save space and notation, they are really column vectors. Hence I should really write $x = (x_1, x_2, \dots, x_n)^T$ where the T operator indicates a transpose.

⁵Given the equilibrium concept of market clearing a consumer’s desired consumption will coincide, in equilibrium, with his actual consumption. The difference between the two concepts, while crucial, is therefore seldom made.

that they will react. We will deal with a simple version of that later under the heading of bilateral monopoly, where there are only two consumers who have to bargain over the price. In all but the most trivial settings, situations in which the prices vary with the consumer's actions and the consumer is cognizant of this fact will be modelled using game theory.

More crucially, consumers cannot spend more money than what they have. In other words, their consumption choices are limited to **economically feasible** vectors. The set of economically feasible consumption vectors for a given consumer is termed the consumer's **budget set**: $B = \{x \in X \mid p \cdot x \leq m\}$. Here p is a vector of n prices, and m is a scalar denoting the consumer's monetary wealth. Recall that $p \cdot x$ denotes a 'dot product', so that $p \cdot x = \sum_{i=1}^n p_i x_i$. The left hand side of the inequality therefore gives total expenditure (cost) of the consumption bundle. To add to the potential confusion, we normally do not actually have money in the economy at all. Instead the consumer usually has an **endowment** — an initial bundle of goods. The budget constraint then requires that the value of the final consumption bundle does not exceed the value of this endowment, in other words $B = \{x \in X \mid p \cdot x \leq p \cdot \omega\}$. Here $\omega \in X$ is the endowment. You can treat this as a two stage process during which prices stay constant: First, sell all initially held goods at market prices: this generates income of $p \cdot \omega = m$. Then buy the set of final goods under the usual budget constraint. Note however that for comparative static purposes only true market transactions are important, as evidenced by the change in the Slutsky equation when moving to the endowment model (p.46, for example.)

The behavioural assumption in consumer theory is that **consumers maximize utility**. What we really mean by that is that consumers have the ability to rank all consumption bundles $x \in X$ according to what is termed the consumer's preferences, and will choose that bundle which is the highest ranked (most preferred) among all the available options. We denote preferences by the symbol \succeq , which is a binary relation. That simply means that it can be used to compare two vectors (not one, not three or more.) The expression $x \succeq y$ is read as "consumption bundle x is at least as good as bundle y ", or "x is weakly preferred to y".

These preferences are assumed to be:

- (i) **complete** ($\forall x, y \in X$: either $x \succeq y$ or $y \succeq x$ or both);
- (ii) **reflexive** ($\forall x \in X$: $x \succeq x$);
- (iii) **transitive** ($\forall x, y, z \in X$: $(x \succeq y, y \succeq z) \Rightarrow x \succeq z$);
- (iv) **continuous** ($\{x \mid x \succeq y\}$ and $\{x \mid y \succeq x\}$ are closed sets.)

What this allows us to do is to **represent** the preferences by a function

$u : \mathfrak{R}_+^N \mapsto \mathfrak{R}$, called a **utility function**, which has the property that

$$u(x) > u(y) \iff x \succ y,$$

i.e., that the value of the function evaluated at x is larger than that at y if and only if x is strictly preferred to y . If preferences are also strongly monotonic ($(x \geq y, x \neq y) \Rightarrow x \succ y$) then the function $u(\cdot)$ can also be chosen to be continuous. With some further assumptions it will furthermore be continuously differentiable, and that is what we normally assume.

Note that the utility function $u(\cdot)$ representing some preferences \succeq is **not unique!** Indeed, any other function which is a monotonic increasing transformation of $u(\cdot)$, say $h(u(\cdot))$, $h'(\cdot) > 0$, will represent the same preferences. So, for example, the following functions all represent the same preferences on \mathfrak{R}^2 :

$$x_1^\alpha x_2^\beta; \quad \alpha \ln x_1 + \beta \ln x_2; \quad x_1^a x_2^{(1-a)} - 10^4, \quad a = \frac{\alpha}{\alpha + \beta}; \quad \alpha, \beta > 1.$$

In terms of the diagram of the utility function this means that the function has to be increasing, but that the rate of increase, and changes in the rate of increase, are meaningless. Put differently, the “spacing” of indifference curves, or more precisely, the labels of indifference curves, are arbitrary, as long as they are increasing in the direction of higher quantities. Furthermore, any two functions that have indifference curves of the same shape represent the same preferences. This is in marked contrast to what we will have to do later. Once we discuss uncertainty, the curvature (rate of change in the rate of increase) of the function starts to matter, and we therefore will then be restricted to using positive affine transforms only (things of the form $a + bu(\cdot)$; $b > 0$).

Finally a few terms which will arise frequently enough to warrant definition here:

Definition 4 *The preferences \succeq are said to be **monotone** if*

$$\forall x \in X \quad x \gg y \iff x \succ y.$$

Definition 5 *The preferences \succeq are said to be **strongly monotone** if*

$$\forall x \in X \quad x \geq y \iff x \succ y.$$

Note in particular that Leontief preferences (‘L’-shaped indifference curves) are monotone but not strictly monotone.

Definition 6 *The preferences \succeq are said to be **convex** if $\forall x \in X$ the set $\{y \in X | y \succeq x\}$ is convex.*

In other words, the set of bundles weakly preferred to a given bundle is convex. Applied to the utility function $u(\cdot)$ representing \succeq this means that the upper contour sets of the utility function are convex sets, which, if you recall, is the definition of a **quasi-concave** function.

Definition 7 *A function $u : \mathbb{R}^N \mapsto \mathbb{R}$ is quasi-concave if its upper contour sets are convex. Alternatively, u is quasi-concave if $u(\lambda x + (1 - \lambda)y) \geq \min\{u(x), u(y)\}, \forall \lambda \in [0, 1]$.*

Note also that the lower boundary of an upper contour set is what we refer to as an indifference curve. (Which exists due to our assumption of continuity of \succeq , which is why we had to make that particular assumption.)

Special Utility Functions

By and large utility functions and preferences can take many forms. However, of particular importance in modelling applications and in many theories are those which allow the derivation of the complete indifference map from just one indifference curve. Two kinds of preferences for which this is true exist, homothetic preferences and quasi-linear preferences.

Definition 8 *Preferences \succeq are said to be homothetic if $\forall x, y \in X, \alpha \in \mathbb{R}_{++}, x \sim y \iff \alpha x \sim \alpha y$.*

In other words, for any two bundles x and y which are indifferent to one another, we may scale them by the same factor, and the resulting bundles will also be indifferent. This is particularly clear in \mathbb{R}_+^2 . A bundle x can be viewed as a vector with its foot at the origin and its head at the coordinates (x_1, x_2) . $\alpha x, \alpha > 0$ defines a ray through the origin through that point. Now consider some indifference curve and 2 points on it. These points define two rays from the origin, and homotheticity says that if we scale the distance from the origin by the same factor on these two rays, we will intersect the same indifference curve on both. Put differently, indifference curves are related to one another by “radial blowup”. The utility function which goes along with such preferences is also called homothetic:

Definition 9 A utility function $u : \mathfrak{R}_+^N \mapsto \mathfrak{R}$ is said to be homothetic if $u(\cdot)$ is a positive monotonic transformation of a function which is homogenous of degree 1.

How does this assumption affect indifference curves? Recall that the key about indifference curves is their slope, not the label attached to the level of utility. The slope, or MRS (Marginal Rate of Substitution,) is $u_1(x)/u_2(x)$. Now let $u(x)$ be homothetic, that is $u(x) = h(l(x))$ where $h'(\cdot) > 0$ and $l(\cdot)$ is homogeneous of degree 1. Then we get

$$\frac{u_1(x)}{u_2(x)} = \frac{h'(l(x))l_1(x)}{h'(l(x))l_2(x)} = \frac{l_1(x)}{l_2(x)}$$

by the chain rule. So what you ask? Recall that

Definition 10 A function $h : \mathfrak{R}^N \mapsto \mathfrak{R}$ is homogeneous of degree k if $\forall \lambda > 0$, $h(\lambda x) = \lambda^k h(x)$.

Therefore (in two dimensions, for simplicity) $l(x_1, x_2) = x_1 l(1, x_2/x_1) = x_1 \hat{l}(k)$, where $k = x_2/x_1$. But then $l_1(x) = \hat{l}(k) + \hat{l}'(k)k$, and $l_2(x) = \hat{l}'(k)$, and therefore the marginal rate of substitution of a homothetic utility function is only a function of the ratio of the consumption amounts, not a function of the absolute amounts. But the ratio x_2/x_1 is constant along any ray from the origin. Therefore, the MRS of a homothetic utility function is constant along any ray from the origin! You can verify this property for the Cobb-Douglas utility function, for example.

The other class of preferences for which knowledge of one indifference curve is enough to know all of them is called quasi-linear.

Definition 11 The preference \succeq defined on $\mathfrak{R} \times \mathfrak{R}_+^{N-1}$ is called quasi-linear with respect to the numeraire good 1, if $\forall x, y \in \mathfrak{R} \times \mathfrak{R}_+^{N-1}$, $\alpha > 0$, $x \sim y \iff x + \alpha e_1 \sim y + \alpha e_1$. Here $e_1 = (1, 0, 0, \dots)$ is the base vector of the numeraire dimension.

In terms of the diagram in \mathfrak{R}^2 , we require that indifference between 2 consumption bundles is maintained if an equal amount of the numeraire commodity is added (or subtracted) from both: Indifference curves are related to one another via translation along the numeraire axis!

Definition 12 A utility function $u : \mathfrak{R} \times \mathfrak{R}_+^{N-1} \mapsto \mathfrak{R}$ is called quasi-linear if it can be written as $u(x) = x_1 + \hat{u}(x_{-1})$.

For example the functions $u(x) = x_1 + \sqrt{x_2}$ and $u(x) = x_1 + \ln x_2$ are quasi-linear with respect to commodity 1. Note that the MRS of a quasi-linear function is independent of the numeraire commodity:

$$\frac{u_1(x)}{u_2(x)} = \frac{1}{\hat{u}'(x_2)},$$

so that the indifference curves all have the same slope along a line parallel to the numeraire axis (good 1 in this case.)

2.2.1 Utility Maximization

The problem addressed in consumer theory is to predict consumer behaviour for a price taking consumer with rational preferences. The consumer's behaviour is often summarized by the **Marshallian demand functions**, which give us the desired consumption bundle of the consumer for all prices and income. These demands correspond to the observable demands by consumers, since they depend only on observable variables (prices and income.) This is in contrast to the Hicksian demand functions, which we later derive from the expenditure minimization problem. Marshallian, or ordinary, demand is derived mathematically from a constrained optimization problem:

$$\max_{x \in X} \{u(x) \text{ s.t } p \cdot x \leq m\}.$$

The easiest way to solve this is to realize that (i) normally all commodities are goods, that is, consumer preferences are monotonic, wherefore we can replace the inequality with an equality constraint;⁶ (ii) that normally we have assumed a differentiable utility function so that we can now use a Lagrangian:

$$\begin{aligned} \max_{x \in X} L &= u(x) + \lambda(m - p \cdot x) \\ (FOC_i) \quad L_i &= \frac{\partial}{\partial x_i} u(x) - \lambda p_i = u_i(\cdot) - p_i \lambda = 0 \quad \forall i = 1 \dots n \\ (FOC_\lambda) \quad L_\lambda &= m - p \cdot x = 0 \end{aligned}$$

Of course, there are also second order conditions, but if we know that preferences are convex (the utility function is quasi-concave) and that the budget is (weakly) concave, we won't have to worry about them. If these conditions

⁶In fact, Local Non-satiation is enough to guarantee that the budget is exhausted. Local Non-satiation assumes that for any bundle there exists a strictly preferred bundle arbitrarily close by. This is weaker than monotonicity since no direction is assumed.

are not obviously satisfied in a problem you will have to check the second order conditions!

The above first order conditions are $n + 1$ equations which have to hold as identities at the optimum values of the $n + 1$ choice variables (the n consumption levels of goods, and the multiplier λ .) That means a variety of things, most importantly that the implicit function theorem can be applied if need be (and it is needed if we want to do comparative statics!)

Let us interpret the above first order conditions. The last equation (FOC_λ) states that the optimum is on the budget constraint. Since they are identities, any two of the other first order conditions can be combined in one of the following ways:

$$\left. \begin{array}{l} u_i(\cdot) = \lambda p_i \\ u_j(\cdot) = \lambda p_j \end{array} \right\} \Rightarrow \frac{u_i(\cdot)}{u_j(\cdot)} = \frac{p_i}{p_j} \Rightarrow \frac{u_i(\cdot)}{p_i} = \frac{u_j(\cdot)}{p_j}$$

These have the economic interpretation that the (negative of the) slope of the indifference curve (the MRS) is equal to the ratio of prices. But since the slope of the budget is the negative of the ratio of prices, this implies that a tangency of an indifference curve to the budget must occur at the optimum.⁷

The second form tells us that the marginal utility gained from the last dollar's worth of consumption must be equalized across all goods at the optimum, another useful thing to keep in mind. Also note that $\lambda^* = u_i(\cdot)/p_i$, so that the level of the multiplier gives us the "marginal utility of budget". This can also be seen by considering the fact that $\partial L/\partial m = \lambda$.

The marginal rate of substitution term depends on the amounts of the goods consumed. In other words $u_i(\cdot)/u_j(\cdot)$ is a function of x . Requiring that this ratio take on some specific value thus implicitly defines a relationship between the x_i . This tangency condition thus translates into a locus of consumption bundles where the slope of the indifference curve has the given value. This locus is termed the **Income Expansion path**. This is due to the fact that this locus would be traced out if we were to keep prices constant but increased income, thus shifting the budget out (in a parallel fashion.) Since the budget line provides another locus of consumption bundles, namely those which are affordable at a given income level, solving for the demands then boils down to solving for the intersection of these lines. Now note that the

⁷If you ever get confused about which term "goes on top" remember the following derivation: An indifference curve is defined by $u(x, y) = \bar{u}$. Taking a total differential we obtain $u_x(\cdot)dx + u_y(\cdot)dy = 0$, and hence $dy/dx = -u_x(\cdot)/u_y(\cdot)$. The same works for the budget. As a last resort you can verify that it is the term for the horizontal axis which is "on top".




Figure 2.1: Consumer Optimum: The tangency condition

income expansion path for both homothetic and quasi-linear preferences is linear. Hence solving for demands is particularly simple, since, in geometric terms, we are looking for the intersection of a line and a (hyper-) plane.

In terms of a diagram, we have the usual picture of a set of indifference curves superimposed on a budget set, given in Figure 2.1. The optimum occurs on the highest indifference curve feasible for the budget, i.e., where there is a tangency between the budget and an indifference curve.

Note a couple of things: For one, the usual characterization of “tangency to budget” is only valid if we are not at a corner solution and can compute the relevant slopes. If we are at a corner solution (and this might happen either because our budget ‘stops’ before it reaches the axis or because our indifference curves intersect an axis, as for quasi-linear preferences) then we really need to consider “complementary slackness conditions”. Basically we will have to consider all the non-negativity constraints which we have currently omitted, and their multipliers. We will not bother with that, however. In this course such cases will be fairly obvious and can be solved by inspection. The key thing to remember, which follows from the condition on per dollar marginal utilities, is that if we are at a corner then the indifference curve must lie outside the budget. That means that on the horizontal axis the indifference curve is steeper than the budget, on the vertical axis it is flatter. Finally, if the indifference curve or the budget have a kink, then no tangency exists in strict mathematical terms (since we do not have differentiability), but we will often loosely refer to such a case as a “tangency” anyways.




Figure 2.2: Corner Solutions: Lack of tangency

Properties

As mentioned, the above problem gives rise first of all to the (ordinary) demand functions. These tell us the relationship of demand behaviour to the parameters of the problem, namely the price vector and the income:

$$x(p, m) = \operatorname{argmax}_{x \in X} \{ u(x) + \lambda(m - p \cdot x) \}.$$

The other result of solving the optimization problem is the indirect utility function, which tells us the highest level of utility actually achieved:

$$v(p, m) = \max_{x \in X} \{ u(x) + \lambda(m - p \cdot x) \}.$$

This is useless *per se*, but the function is nevertheless useful for some applications (duality!)

What properties can we establish for these? Two important properties of the (ordinary) demand functions are

Theorem 3 *The (ordinary) demand $x(p, m)$ derived from a continuous utility function representing rational preferences which are monotonic satisfies:*

homogeneous of degree zero: $x(p, m) = x(\alpha p, \alpha m)$, $\forall p, m; \alpha > 0$;

Walras' Law: $p \cdot x(p, m) = m$, $\forall p \gg 0, m > 0$.

Walras' Law has two important implications known as Engel and Cournot Aggregation. Both are derived by simple differentiation of Walras' Law. Engel Aggregation refers to the fact that a change in consumer income must all be spent. Cournot Aggregation refers to the fact that a change in a price may not change total expenditure.

Definition 13 Engel Aggregation:

$$\sum_{i=1}^n p_i \frac{\partial x_i(\cdot)}{\partial m} = 1, \text{ or more compactly } p \cdot D_p x(p, m) = 1.$$

Definition 14 Cournot Aggregation:

$\sum_{i=1}^n p_i \frac{\partial x_i(\cdot)}{\partial p_k} + x_k(p, m) = 0, \forall k = 1, \dots, n$, or $p \cdot D_p x(p, m) + x(p, m)^T = 0^T$.
(Here 0^T is a row vector of zeros.)

Sometimes it is more useful to write these in terms of elasticities (since that is what much of econometrics estimates)⁸.

Define $s_i(p, m) = p_i x_i(p, m)/m$, the expenditure share of commodity i (note how for a Cobb-Douglas utility function this is constant!) Let $\epsilon_{ia}(\cdot)$ denote the elasticity of demand for good i with respect to the variable a . Then the above aggregation laws can be expressed as

$$\sum_{i=1}^n s_i(p, m) \epsilon_{im}(p, m) = 1 \quad \text{and} \quad \sum_{i=1}^n s_i(p, m) \epsilon_{ik}(p, m) i + s_k(p, m) = 0.$$

We will not bother here with the properties of the indirect utility function aside from providing **Roy's Identity**:

$$x(p, m) = \frac{-1}{\partial v(p, m)/\partial m} \cdot \nabla_p v(p, m) = \left[-\frac{\partial v(\cdot)/\partial p_i}{\partial v(\cdot)/\partial m} \right]_{n \times 1}.$$

This is “just” an application of the envelope theorem, or can be derived explicitly by using the first order conditions. The reason this is useful is that it is mathematically much easier to differentiate, compared to solving $L + 1$ nonlinear equations. Hence a problem can be presented by stating an indirect utility function and then differentiating to get demands, a technique often employed in trade theory.⁹

Finally, note that we have **no** statements about demand other than the ones above. In particular, it is **not** necessarily true that (ordinary) demand curves slope downwards ($\partial x(p, m)/\partial p_i \leq 0$.) Similarly, the income derivatives can take any sign, and we call a good **normal** if $\partial x(p, m)/\partial m \geq 0$, while we call it **inferior** if $\partial x(p, m)/\partial m \leq 0$.

⁸Recall that often the logarithms of variables are regressed. A coefficient on the right hand side thus represents $\frac{d \ln x}{d \ln y} = \frac{dx}{dy} \frac{y}{x}$

⁹In graduate courses the properties of the indirect utility function get considerable attention for this reason. It can be shown that any function having the the requisite properties can be the indirect utility function for some consumer.

2.2.2 Expenditure Minimization and the Slutsky equation

Another way to consider the consumer's problem is to think of the least expenditure needed to achieve a given level of utility. Basically, this reverses the process from above. We now fix the indifference curve we want to achieve and minimize the budget subject to that:

$$e(p, u) = \min_{x \in X} \{ p \cdot x + \lambda(\bar{u} - u(x)) \}.$$

The function which tells us the least expenditure for a given price vector and utility level is called the **expenditure function**. The optimal choices for this problem,

$$h(p, u) = \operatorname{argmin}_{x \in X} \{ p \cdot x + \lambda(\bar{u} - u(x)) \},$$

are the **Hicksian demand** functions. These are unobservable (since they depend on the unobservable utility level) but extremely useful. The reason is that the Hicksian demand for good i , $h_i(p, u)$, is (by the envelope theorem) the first partial derivative of the expenditure function with respect to p_i . It follows that the first derivative of the Hicksian demand function with respect to a price is the second partial derivative of the expenditure function. So what, you ask? Consider the following:

Theorem 4 *If $u(x)$ is continuous and represents monotonic, rational preferences and if $p \gg 0$, then $e(p, u)$ is **homogeneous of degree 1 in prices; continuous in u and p ; concave in p ; increasing in u and nondecreasing in p .***

We will not prove these properties, but they are quite intuitive. For concavity in p consider the following: Fix an indifference curve and 2 tangent budget lines to it. A linear combination of prices at the same income is another budget line which goes through the intersection point of the two original budgets and has intermediate axis intercepts. Clearly, it cannot touch or intersect the indifference curve, indicating that higher expenditure (income) would be needed to reach that particular utility level.

So, since the expenditure function is concave in prices, this means that the matrix of second partials, $D_p \nabla_p e(p, u)$, is negative semi-definite. Also, by Young's Theorem, this matrix is symmetric. But it is equal to the matrix of first partials of the Hicksian demands, which therefore is also symmetric and negative semi-definite!

Why is this useful information? It means that Hicksian demand curves do not slope upwards, since negative semi-definiteness requires negative main diagonal entries. This is a statement which we could not make about ordinary demand curves. Also, the cross-price effects of any two goods are equal. This certainly is unexpected, since it requires that the effect of a change in the price of gasoline on bananas is the same as the effect of a change in the price of bananas on gasoline. Finally, Hicksian demands are homogeneous of degree zero in prices.

As mentioned, the Hicksian demand curves are unobservable. So why should we be excited about having strong properties for them, if they can never be tested? One key reason is the **Slutsky equation**, which links the unobservable Hicksian demands to the observable ordinary demands. We proceed as follows: The fact that we are talking about two ways to look at the same optimal point means that

$$x_i(p, e(p, u)) = h_i(p, u),$$

where we recognize the fact that the income and expenditure level must coincide in the two problems. Now taking a derivative of both sides we get

$$\frac{\partial x_i(p, e(p, u))}{\partial p_j} + \frac{\partial x_i(p, e(p, u))}{\partial m} \frac{\partial e(p, u)}{\partial p_j} = \frac{\partial h_i(p, u)}{\partial p_j},$$

and using the envelope theorem on the definition of the expenditure function as well as utilizing the original equality we simplify this to

$$\frac{\partial x_i(p, e(p, u))}{\partial p_j} + \frac{\partial x_i(p, e(p, u))}{\partial m} x_j(p, u) = \frac{\partial h_i(p, u)}{\partial p_j}.$$

The right hand side of this is a typical element of the symmetric and negative semi-definite matrix of price partials of the Hicksian demands. The equality then implies that the matrix of corresponding left hand sides, termed the **Slutsky matrix**, is also symmetric and negative semi-definite.

Reordering any such typical element we get the **Slutsky equation**:

$$\frac{\partial x_i(p, e(p, u))}{\partial p_j} = \frac{\partial h_i(p, u)}{\partial p_j} - \frac{\partial x_i(p, e(p, u))}{\partial m} x_j(p, e(p, u)).$$

This tells us that for any good the total response to a change in price is composed of a substitution effect and an income effect. Therefore ordinary demands can slope upward if there is a large enough negative income effect — which means that we need an inferior good.

It bears repeating that homogeneous demands satisfying Walras' Law with a symmetric, negative semi-definite Slutsky matrix is all that our theory can say about the consumer. Nothing more can be said unless one is willing to assume particular functional forms.

2.3 General Equilibrium

What remains is to model the determination of equilibrium prices. This is not really done in economics, however. Since in the standard model every consumer (and every firm, in a production model) takes prices as given, there is nobody to set prices. Instead a “Walrasian auctioneer” is imagined, who somehow announces prices. The story (and it is nothing more) is that the auctioneer announces prices, and checks if all markets clear at those prices. If not, no trades occur but the prices are adjusted (in the way first suggested in Econ 101: markets with excess demands experience a price increase, those with excess supply a price decrease, the “tâtonnement” process.) Of course, all agents other than the auctioneer must be unaware of this process somehow — otherwise they might start to mis-state demands and supplies in order to affect prices. In any case, there really is no model of equilibrium price determination. This does, of course, not stop us from defining what an equilibrium price is! At its most cynical, equilibrium can be viewed simply as a consistency condition which makes our model logically consistent. In particular, it must be true that, in equilibrium, no consumer finds it impossible to trade the amounts desired at the announced market prices. That is, a key requirement of any equilibrium price vector will be that all economic agents can maximize their objective functions and carry out the resulting plans. Since aggregate demands and supplies are the sum of individual demands and supplies it follows that aggregate demand and supply must balance in equilibrium — all markets must clear!

In what follows we will only review 2 special cases of general equilibrium, the pure exchange economy, and the one consumer, one producer economy.

2.3.1 Pure Exchange

Consider an economy with two goods, $i = 1, 2$, and two consumers, $l = 1, 2$. Each consumer has an initial endowment $\omega^l = (\omega_1^l, \omega_2^l)$ of the two goods, and preferences \succeq_j on \mathbb{R}^2 which can be represented by some utility function $u_j(x) : \mathbb{R}_+^2 \mapsto \mathbb{R}$. The **total endowment** of good i in the economy is the

sum of the consumers' endowments, so there are $\omega_1 = \omega_1^1 + \omega_1^2$ units of good 1, and $\omega_2 = \omega_2^1 + \omega_2^2$ units of good 2.

We can now define the following:

Definition 15 An **allocation** $x = (x^1, x^2)$ is an assignment of a nonnegative consumption vector $x^l \in \mathfrak{R}^2$ to each consumer.

Definition 16 An allocation is **feasible** if $\sum_{l=1}^2 x_i^l = \sum_{l=1}^2 \omega_i^l$ for $i = 1, 2$.

Note that feasibility implies that if consumer 1 consumes (x_1^1, x_2^1) in some feasible allocation, then consumer 2 must consume $(\omega_1 - x_1^1, \omega_2 - x_2^1)$. This fact allows us to represent the set of feasible allocations by a box in \mathfrak{R}^2 . Suppose we measure good 1 along the horizontal axis and good 2 along the vertical axis. Then the feasible allocations lie anywhere in the rectangle formed by the origin and the point (ω_1, ω_2) . Furthermore, if we measure consumer 1's consumption from the usual origin (bottom left), then for any consumption point x we can read off consumer 2's consumption by measuring from the right top 'origin' (the point ω .)

Also note that a budget line for consumer 1, defined by $p \cdot x = p \cdot \omega^1$ passes through both the endowment point and the allocation, and has a slope of p_1/p_2 . But if we consider consumer 2's budget line at these prices we find that it passes through the same two points and has, measured from consumer 2's origin on the top right, the same slope. Hence any price vector defines a line through the endowment point which divides the feasible allocations into the budget sets for the two consumers. These budget sets only have the budget line in common.

As a final piece, we may represent each consumer's preferences by an indifference map in the usual fashion, only that consumer 2's consumption set is measured from the top right origin (at ω in consumer 1's coordinates.) Note that both consumer's indifference curves are *not limited* to lie within the box. Feasibility is a problem for the economy, not the consumer. Indeed, the consumer is assumed to be unaware of the actual amount of the goods in the economy. Hence each consumer's budget and indifference maps are defined in the usual fashion over the respective consumer's consumption set, which here is all of \mathfrak{R}_+^2 , measured from the usual origin for consumer 1 and measured down and to the left from ω for consumer 2! (More bluntly, both a consumer's budget line and her indifference curves can "leave the box".)

Figure 2.3 represents an example of such a pure exchange economy. Here we have assumed that $u_1(x) = \min\{x_1, x_2\}$; $\omega^1 = (10, 5)$, while $u_2(x) = x_1 + x_2$; $\omega^2 = (10, 10)$.




Figure 2.3: An Edgeworth Box

Aside of the notion of feasibility the other key concept in general equilibrium analysis is that of **Pareto efficiency**, or Pareto optimality.

Definition 17 A feasible allocation x is **Pareto Efficient** if there does not exist a feasible allocation y such that $y \succeq_l x$, $\forall j$, and $\exists j \ni y \succ_l x$.

A slightly weaker criterion also exists to deal with preferences that are not strictly monotonic:

Definition 18 A feasible allocation x is **weakly Pareto Efficient** if there does not exist a feasible allocation y such that $y \succ_l x$, $\forall j$.

In words, Pareto efficiency (PO) requires that no feasible alternative exist which leaves all consumers as well off and at least one better off. Weak Pareto efficiency requires that no feasible allocation exist which makes all consumers better off.

Note that these definitions can be tricky to apply since an allocation is defined as efficient if something is **not** true. Moreover, there are infinitely

box to the right of the intersection with the kink line? For such points both consumers cannot be made better off, since there exists no feasible allocation which makes consumer 1 better off. Thus, the allocation q is weakly PO. However, by moving left we can make consumer 2 better off without reducing the utility of consumer 1, so q is not PO in the strict sense.

The set of Pareto optimal allocations in Figure 2.4 therefore is the kink line of consumer 1, that is all allocations

$$\{x \in \mathbb{R}^4 \mid x_1^1 = x_2^1, x_1^1 \leq 15, x_1^2 = 20 - x_1^1, x_2^2 = 15 - x_1^1, x_i^l \geq 0 \forall i, l = 1, 2\}.$$

You can see from the above example that the difference between PO and weak PO can only arise for preferences which have vertical or horizontal sections that could coincide with a boundary.

Theorem 5 *All Pareto Optimal allocations are weakly Pareto optimal. If all preferences are strictly monotonic, then all weakly Pareto optimal allocations are also Pareto optimal.*

Finally note that for convex preferences and differentiable utility functions the condition that the weakly preferred sets be disjoint translates into the familiar requirement that the indifference curves of the consumers must be tangent, that is, an allocation is Pareto optimal if the consumers' marginal rates of substitution equal.

We are now ready to define equilibrium and to see how it relates to the above diagram.

Definition 19 *A Walrasian (competitive) equilibrium for a pure exchange economy is a price vector p^* and allocation x^* such that*

- (1) (**optimality**) for each consumer l , $(x^l)^* \succeq_l x^l \quad \forall x^l \ni p^* x^l = p^* \omega^l$;
- (2) (**feasibility**) for each good i , $\sum_j (x_i^j)^* = \sum_j \omega_i^j$.

This definition is the most basic we can give, and at the same time includes all others. For example, you may remember from introductory economics that a general equilibrium is achieved when “demand equals supply”. But supply in this case is the sum of endowments. Demand means aggregate demand, that is, the sum of all consumers' demands. But a consumer's demand has been defined as a utility maximizing bundle at a given price — hence the notion of demand includes the optimality condition above. We therefore can define equilibrium in the following way:

Definition 20 A **Walrasian (competitive) equilibrium** for a pure exchange economy is a price vector p^* and allocation x^* such that $x^* = x(p^*)$, $\sum_j x_i^l(p^*) = \sum_j \omega_i^l$, for each good i , and for each consumer l , $x^l(p) = \operatorname{argmax}_x \{u_j(x) \text{ s.t. } p \cdot x^l = p \cdot \omega^l\}$.

For an Edgeworth box we can also define equilibrium in yet a third way — as the intersection of consumers' offer curves. An offer curve is the analogue of a demand curve, only drawn in commodity space, not commodity-price space.

Definition 21 The **offer curve** is the locus of all optimal consumption bundles (in X) as price changes:

$\{x \in X \mid \exists p \ni \forall y \text{ with } p \cdot y = p \cdot \omega, x \succeq_j y, \text{ and } p \cdot x = p \cdot \omega\}$ or alternatively $\{x \in X \mid \exists p \ni x = \operatorname{argmax}\{u_j(x) \text{ s.t. } p \cdot x = p \cdot \omega^l\}\}$.

Note how the notion of an offer curve incorporates optimality, so that the intersection of consumers' offer curves, lying on both offer curves, is optimal for both consumers. Since offer curves cannot intersect outside the Edgeworth box, feasibility is also guaranteed.¹⁰

We thus have different ways of finding an equilibrium allocation in a pure exchange economy, and depending on the setting one may be more convenient than another. For example, should both consumers have differentiable utility functions we can normally most easily find demands, and then set demand equal to supply. This is simplified by Walras' Law. Let $z(p)$ denote the aggregate excess demand vector: $z(p) = \sum_j (x^l(p) - \omega^l)$. Then we have.

Definition 22 Walras' Law: $p \cdot z(p) = 0 \quad \forall p \in \mathfrak{R}_+^n$.

Why is this so? It follows from Walras' Law for each consumer, requiring that each consumer l 's demands exhaust the budget, so that the total value of excess demand for a consumer is zero, for all prices, by definition. Hence, summing across all consumers the aggregate value of excess demand must also be zero. In an equilibrium aggregate demand is equal to aggregate supply,¹¹ that is, summing across consumers the excess demand must equal zero for each good, and hence the value must be zero. Hence, if $n - 1$ markets clear

¹⁰This argument assumes as given that offer curves cannot intersect outside the feasible set. Why? You should be able to argue/prove this!

¹¹Yet another definition of a Walrasian equilibrium is $z(p) = 0!$

the value of the excess demand of the n^{th} market must be zero, and this can only be the case if either the price of the good is zero or the actual aggregate excess demand itself. Thus it suffices to work with $n - 1$ markets. The other market will “take care of itself.”

In practice we are also free to choose one of the prices, since we have already shown that (Walrasian) demand is homogeneous of degree zero in prices for each consumer (remember that $m^l = p \cdot \omega^l$ here), and thus aggregate demand must also be homogeneous of degree zero.

So, suppose the economy consists of two Cobb-Douglas consumers with parameters α and β and endowments $((\omega_1^1, \omega_2^1), (\omega_1^2, \omega_2^2))$. Demands are

$$\left(\frac{\alpha p \cdot \omega^1}{p_1}, \frac{(1 - \alpha)p \cdot \omega^1}{p_2} \right), \text{ and } \left(\frac{\beta p \cdot \omega^2}{p_1}, \frac{(1 - \beta)p \cdot \omega^2}{p_2} \right).$$

Hence the equilibrium price ratio p_1/p_2 is

$$p^* = \frac{\alpha\omega_2^1 + \beta\omega_2^2}{(1 - \alpha)\omega_1^1 + (1 - \beta)\omega_1^2}.$$

The allocation then is obtained by substituting the price vector back into the consumer’s demands and simplifying.

On the other hand, if we have non-differentiable preferences it is most often more convenient to use the offer curve approach. For example, a consumer with Leontief preferences will consume on the kink line no matter what the price ratio. If such a consumer were paired with a consumer with perfect substitute preferences, such as in Figure 2.3, we know the price ratio by inspection to be equal to the perfect substitute consumer’s MRS (in Figure 2.3 that is 1.) Why? Since a perfect substitute consumer will normally consume only one of the goods, and only consumes both if the MRS equals the price ratio. The allocation then can be derived as the intersection of consumer 1’s kink line with the indifference curve through the endowment point for consumer 2: in Figure 2.3 the equilibrium allocation thus would be $(x^1, x^2) = ((7.5, 7.5), (12.5, 7.5))$.

2.3.2 A simple production economy

The second simple case of interest is a production economy with one consumer, one firm and 2 goods, leisure and a consumption good. This case is also often used to demonstrate the second welfare theorem in its most simple

setting. Recall that the second welfare theorem states that a Pareto efficient allocation can be “decentralized”, that is, can be achieved as a competitive equilibrium provided that endowments are appropriately redistributed.

Let us therefore approach this problem as is often done in macro economics. Consider the social planner’s problem, which in this case is equivalent to the problem of the consumer who owns the production technology directly. The consumer has preferences over the two goods, leisure and consumption, denoted by l and c respectively, represented by the utility function $u(c, l)$. It has the usual properties (that is, it has positive first partials and is quasi-concave.) The consumer has an endowment of leisure, \bar{L} . Time not consumed as leisure is work, $x = \bar{L} - l$, and the consumption good c can be produced according to a production function $c = f(x)$. Assume that it exhibits non-increasing returns to scale (so $f'(\cdot) > 0$; $f''(\cdot) \leq 0$.) The consumer’s problem then is

$$\max_{c,l} \{u(c, l) \text{ s.t. } c = f(\bar{L} - l) \}$$

or, substituting out for c

$$\max_l u(f(\bar{L} - l), l).$$

The FOC for this problem is $-u_c(\cdot)f'(\cdot) + u_l(\cdot) = 0$, or more usefully, $f'(\cdot) = u_l(\cdot)/u_c(\cdot)$. This tells us that the consumer will equate the marginal product of spending more time at work with the Marginal Rate of Substitution between leisure time and consumption — which measures the cost of leisure time to him.

Now consider a competitive private ownership economy in which the consumer sells work time to a firm and buys the output of the firm. The consumer owns all shares of the firm, so that firm profits are part of consumer income. Suppose the market price for work, x , is denoted by w (for wage), and the price of the output c is denoted by p .

The firm’s problem is $\max_x pf(x) - wx$ with first order condition $pf'(x^*) = w$. The consumer’s problem is $\max_{c,l} \{u(c, l) \text{ s.t. } \pi(p, w) + pc = w(\bar{L} - l) \}$. Here $\pi(p, w)$ denotes firm profits. The FOCs for this lead to the condition $u_l(\cdot)/u_c(\cdot) = w/p$. Hence the same allocation is characterized as in the planner’s problem, as promised by the second welfare theorem. (Review Problem 12 provides a numerical example of the above.)

Note that in this economy there is a unique Pareto optimal point if $f''(\cdot) < 0$. Since a competitive equilibrium must be PO, it suffices to find the PO allocation and you have also found the equilibrium! This “trick” is often exploited in macro growth models, for example.

2.4 Review Problems

It is useful exercise to derive both kinds of demand for the various functions you might commonly encounter. The usual functions are

Cobb-Douglas: $u(x_1, x_2) = x_1^\alpha x_2^{(1-\alpha)}$, $\alpha \in (0, 1)$;

Quasi-Linear: $u(x_1, x_2) = f(x_1) + x_2$, $f' > 0, f'' < 0$;

Perfect Substitute: $u(x_1, x_2) = ax_1 + x_2$, $a > 0$;

Leontief: $u(x_1, x_2) = \min\{ax_1, x_2\}$, $a > 0$;

“Kinked”: $u(x_1, x_2) = \min\{ax_1 + x_2, x_1 + bx_2\}$, $a, b > 0$.

The first thing to do is to determine how the indifference curves for each of these look like. I strongly recommend that you do this. Familiarity with the standard utility functions is assumed in class and in exams.

Budgets are normally easy and are often straight lines. However, they can have kinks. For example, consider the budget for the following problem: A consumer has an endowment bundle of $(10, 10)$ and if he wants to increase his x_2 consumption he can trade 2 x_1 for 1 x_2 , while to increase x_1 consumption he can trade 2 x_2 for 1 x_1 . Such budgets will arise especially in inter temporal problems, where consumers' borrowing and lending rates may not be the same.

You can now combine any kind of budget with any of the above utility functions and solve for the income expansion paths (diagrammatically), and the demands, both ordinary and Hicksian.

You may also want to refresh your knowledge on (price-)offer curves. This is the locus of demands in (x_1, x_2) space which is traced out as one price, say p_1 , is varied while p_2 and m stay constant.

Question 1: Let $X = \{x, y, z, w\}$ and let $(\mathcal{B}, C(\cdot))$ be a choice structure with $\mathcal{B} = \{\{x, y, z\}, \{x, z, w\}, \{w, z, y\}, \{x\}, \{y, w\}, \{z, x\}, \{x, w\}\}$.

- a) Provide a $C(\cdot)$ which satisfies the weak axiom.
- b) Does there exist a preference relation on X which rationalizes your $C(\cdot)$ on \mathcal{B} ?
- c) Is that preference relation rational?
- d) Given this \mathcal{B} and some choice rule which satisfies the weak axiom, are we guaranteed that the choices could be derived from a rational preference relation?
- e) Since you probably ‘cheated’ in part a) by thinking of rational preferences first and then getting the choice rule from that: write down a choice rule $C(\cdot)$ for the above \mathcal{B} which does satisfy the weak axiom but can NOT be rationalized by a preference relation.

Question 2: Consider a consumer endowed with 16 hours of leisure time and \$100 of initial wealth. There are only two goods, leisure and wealth/aggregate consumption, so a consumption bundle can be denoted by (l, c) . Extend the definition of the budget set to be the set of all feasible (l, c) vectors, i.e., $B = \{(l, c) \mid 0 \leq l \leq 16, c \text{ feasible}\}$, where feasibility of consumption c will be determined by total income, consisting of endowed income and work income. In particular note that the boundary of the budget set, \bar{B} , is the set of **maximal** feasible consumption for any given leisure level. Now, in view of this, what is the budget set and its boundary in the following cases (you may want to provide a well labelled diagram, and/or a definition of \bar{B} instead of a set definition for B .) Is the budget set convex?

a) There are two jobs available, each of which can be worked at for no more than 8 hours. An additional restriction is that you have to work at one job full time (8 hours) before you can start the other job. Job 1 pays \$12 per hour for the first 6 hours and \$16 per hour for the next 2 hours. Job 2 pays \$8 per hour for the first two hours and \$14 per hour for the next 6 hours.

b) The same situation as in [a] above, but job 2 now pays \$14 $\frac{2}{3}$ per hour for the next 6 hours (after the first 2).

c) Suppose we drop the restriction that you have to work one job full time before you can work the other. (So you could work one job for 5 hours and the other for 3, for example, which was not possible previously.) Why does it (does it not) matter?

d) What if instead there are four jobs, with job (i) paying \$12 per hour for up to 6 hours, job (ii) paying \$16 per hour for up to 2 hours, (iii) paying \$8 per hour for up to 2 hours and (iv) paying \$14 per hour for up to 6 hours? You may work any combination of jobs, each up to its maximum number of hours.

Question 3: A consumer's preferences over 2 goods can be represented by the utility function $u(x) = x_1^{0.3}x_2^{0.6}$. Solve for the consumer's demand function. If the consumer's endowment is $\omega = (412, 72)$ and the price vector is $p = (3, 1)$, what is the quantity demanded?

Question 4: A consumer's preferences are represented by the utility function

$$U(x_1, x_2) = \min\{x_1 + 5x_2, 2x_1 + 2x_2, 5x_1 + x_2\}.$$

Solve both the utility maximization and the expenditure minimization problem for this consumer. (Note the the Walrasian and Hicksian demands are not functions in this case.) Also draw the price offer curve for p_1 variable, m and p_2 fixed. [Note: Solving the problems does not require differential calculus techniques.]

Question 5: The elasticity of substitution is defined as

$$\epsilon_{1,2} = -\frac{\partial [x_1(p, w)/x_2(p, w)]}{\partial [p_1/p_2]} \frac{p_1/p_2}{x_1(p, w)/x_2(p, w)}.$$

What does this elasticity measure? Why might it be of interest?

Question 6: The utility function $u(x_1, x_2) = (x_1^\rho + x_2^\rho)^{1/\rho}$, where $\rho < 1$ and $\rho \neq 0$, is called the Constant Elasticity of Substitution utility function. The shape of its indifference curves depends on the parameter ρ . Demonstrate that the function looks like a Leontief function $\min\{x_1, x_2\}$ as $\rho \rightarrow -\infty$; like a Cobb-Douglas function x_1x_2 as $\rho \rightarrow 0$; like a perfect substitute function $x_1 + x_2$ as $\rho \rightarrow 1$.

Question 7: A consumer's preferences over 3 goods can be represented by the utility function $u(x) = x_1 + \ln x_2 + 2\ln x_3$. Derive Marshallian (ordinary) demand.

Question 8: Consider a 2 good, 2 consumer economy. Consumer A has preferences represented by $u_A(x) = 4x_1 + 3x_2$. Consumer B has preferences represented by $u_B(x) = 3\ln x_1 + 4\ln x_2$. Suppose consumer A's endowment is $\omega_A = (12, 9)$ while consumer B's endowment is $\omega_B = (8, 11)$. What is the competitive equilibrium of this economy?

Question 9: Consider a 2 good, 2 consumer economy. Consumer A has preferences represented by $u_A(x) = \min\{4x_1 + 3x_2, 3x_1 + 4x_2\}$. Consumer B has preferences represented by $u_B(x) = 3x_1 + 4\ln x_2$. Suppose consumer A's endowment is $\omega_A = (12, 9)$ while consumer B's endowment is $\omega_B = (8, 11)$. What is the competitive equilibrium of this economy?

Question 10*: Suppose $(\mathcal{B}, C(\cdot))$ satisfies the weak axiom. Consider the following two strict preference relations: \succ defined by $x \succ y \iff \exists B \in \mathcal{B} \ni x, y \in B, x \in C(B), y \notin C(B)$; and \succ^* defined by $x \succ^* y \iff \exists B \in \mathcal{B} \ni x, y \in B, x \in C(B)$ and $\nexists B' \in \mathcal{B} \ni x, y \in B', y \in C(B')$. Prove that these two strict preference relations give the same relation on X . Is this still true if the weak axiom is not satisfied by $C(\cdot)$? (Counter-example suffices.)

Question 11*: Consider a 2 good, 2 consumer economy. Consumer A has preferences represented by $u_A(x) = \alpha x_1 + x_2$. Consumer B has preferences represented by $u_B(x) = \beta \ln x_1 + \ln x_2$.

a) Suppose consumer A's endowment is $\omega_A = (10, 10)$ while consumer B's endowment is $\omega_B = (10, 10)$. What is the competitive equilibrium of this economy?

b) Suppose $\alpha = 2, \beta = 1$ and fix the total size of the economy at 20

units of each good as above. As you have seen above, there are two kinds of equilibria, interior, and boundary. For which endowments will interior equilibria obtain, for which ones boundary equilibria?

Question 12*: Consider a private-ownership production economy with one consumption good, one firm and one consumer. Suppose technology is given by $c = f(x) = 4\sqrt{x}$. Let consumer preferences be represented by $u(c, l) = \ln c + \frac{1}{2}\ln l$ and let the consumer have a leisure endowment of $\bar{L} = 16$. **a)** Solve the social planner's problem. **b)** Solve for the competitive equilibrium of the private ownership economy.

Chapter 3

Inter-temporal Economics

In this chapter we will re-label the basic model in order to focus the analysis on the question of inter-temporal choice. This will reveal information about borrowing and savings behaviour and the relationship between interest rates and the time-preferences of consumers. In particular, we will see that interest rates are just another way to express a price ratio between time dated commodities.

3.1 The consumer's problem

In what follows, we will consider the simple case of just two commodities, consumption today — denoted c_1 — and consumption tomorrow — c_2 . These take the place of the two commodities, say apples and oranges, in the standard model of the previous chapter. Since it is somewhat meaningless to speak of income in a two period model, indeed, part of our job here is to find out how to deal with and evaluate income streams, we will employ an endowment model. The consumer is assumed to have an endowment of the consumption good in each period, which we will denote by m_1 and m_2 , respectively. The consumer is assumed to have preferences over consumption in both periods, represented by the utility function $u(c_1, c_2)$.

3.1.1 Deriving the budget set

Our first job will be to determine how the consumer's budget may be expressed in this setting. That, of course, will depend on the technology for storing the good and on what markets exist. Note first of all that it is never possible to consume something before you have it: the consumer will not be able to consume in period 1 any of the endowment in period 2. In contrast, it may be possible to store the good, so that quantities of the period 1 endowment not consumed in period 1 are available for consumption in period 2. Of course, such storage may be subject to losses (depreciation). Ideally, the consumer has some investment technology (such as planting seeds) which allows the consumer to “invest” (denoting that the good is used up, but not in utility generating consumption) period 1 units in order to create period 2 units. Finally, the consumer may be able to access markets, which allow lending and borrowing. We consider these in turn.

No Storage, No Investment, No Markets

If there is no storage and no investment, then anything not consumed today will be lost forever and consumption cannot be postponed to tomorrow. You may want to think of this as 100% depreciation of any stored quantity. No markets means that there is also no way to trade with somebody else in order to either move consumption forward or backward in time. The consumer finds himself therefore in the economically trivial case where he is forced to consume precisely the endowment bundle. The budget set then is just that one point: $B = \{c \in \mathbb{R}_+^2 \mid c_i = m_i, i = 1, 2\}$.

Storage, No Investment, No Markets

This is a slightly more interesting case where consumption can be postponed. Since there are no markets, no borrowing against future endowments is possible. Storage is usually not perfect. Let $\delta \in (0, 1)$ denote the rate of depreciation — for example, our consumption good may spoil, so that the outside layers of the meat will not be edible in the next period, or pests may eat some of the stuff while it is in storage (a serious problem in many countries.) A typical budget set then is

$$B = \{(c_1, c_2) \mid c_2 \leq m_2 + (1 - \delta)(m_1 - c_1), 0 \leq c_1 \leq m_1\}.$$

The quantity $(m_1 - c_1)$ in this expression is the amount of period 1 endowment not consumed in period 1, usually called savings, and denoted S . We can therefore express the consumer's utility maximization problem in three equivalent ways:

$$\begin{aligned} & \max_{c_1, c_2 \in B} u(c) \\ & \max_{c_1 \leq m_1} u(c_1, m_2 + (1 - \delta)(m_1 - c_1)) \\ & \max_{S \leq m_1} u(m_1 - S, m_2 + (1 - \delta)S) \end{aligned}$$

In the second case only the level of consumption in period 1, c_1 , is a choice variable. It implies the consumption in period 2. The third line simply relabels that same maximization and has savings in period 1 — whatever is not consumed — as the choice variable. All of these will give the same answer, of course! The left diagram in Figure 3.1 gives the diagrammatic representation of this optimization problem (with two different (!) preferences indicated by representative indifference curves.)

One thing to be careful about is the requirement that $0 \leq c_1 \leq m_1$. We could employ Kuhn-Tucker conditions to deal with this constraint, but usually it suffices to check after the fact. For example, a consumer with CD preferences $u(c_1, c_2) = c_1 c_2$ faced with a price ratio of unity would like to consume where $c_1 = c_2$ (You ought to verify this!) However, if the endowment happens to be $(m_1, m_2) = (5, 25)$ then this is clearly impossible. We therefore conclude that there is a corner solution and consumption occurs at the endowment point m .

Storage, Investment, No Markets

How is the above case changed if investment is possible? Here I am thinking of physical investment, such as planting, not “market investment”, as in lending. Suppose then that the consumer has the same storage possibility as previously, but also has access to an investment technology. We will model this technology just as if it were a firm, and specify a function that gives the returns for any investment level: $y = f(x)$. This function must either have constant returns to scale or decreasing returns to scale for things to work easily (and for the second order conditions of the utility maximization problem to hold.) Suppose the technology exhibits CRS. In that case the marginal return of investment, $f'(x)$, is constant. It either is larger or smaller than the return of storage, which is $1 - \delta$. Since the consumer will want to maximize the budget set, he will choose whichever has the higher return, so if $f'(x) > (1 - \delta)$ the consumer will invest, and he will store otherwise. The

budget then is

$$B = \begin{cases} \{(c_1, c_2) \mid c_2 \leq m_2 + f(m_1 - c_1), 0 \leq c_1 \leq m_1\} & \text{if } (1 - \delta) < f'() \\ \{(c_1, c_2) \mid c_2 \leq m_2 + (1 - \delta)(m_1 - c_1), 0 \leq c_1 \leq m_1\} & \text{otherwise} \end{cases}$$

What if the investment technology exhibits decreasing returns to scale? Suppose the most interesting case, where $f'(x) > (1 - \delta)$ for low x , but the opposite is true for high x . Suppose further that not only the marginal return falls, but that the case is one where $f(m_1) < (1 - \delta)m_1$, so that the total return will be below that of storage if all first period endowment is invested. What will be the budget set? Clearly (?) any initial amounts not consumed in period 1 should be invested, not stored. But when should the consumer stop investing? The maximal investment \bar{x} is defined by $f'(\bar{x}) = (1 - \delta)$. Any additional unit of foregone period 1 consumption should be stored, since the marginal return of storage now exceeds that of further investment. The budget then is

$$B = \begin{cases} \{(c_1, c_2) \mid c_2 \leq m_2 + f(m_1 - c_1), m_1 - \bar{x} \leq c_1 \leq m_1\} \\ \{(c_1, c_2) \mid c_2 \leq m_2 + f(\bar{x}) + (1 - \delta)(m_1 - \bar{x} - c_1), 0 \leq c_1 \leq m_1 - \bar{x}\} \\ \text{where } \bar{x} \text{ is defined by } (1 - \delta) = f'(\bar{x}) \end{cases}$$

Storage, No Investment, Full Markets

We now allow the consumer to store consumption between periods 1 and 2, with some depreciation, and to trade consumption on markets. Instead of the usual prices, which are an expression of the exchange ratio of, say good 2 for good 1, we normally express things in terms of **interest rates** when we deal with time. Of course, one can always convert between the two without much trouble if some care is taken. The key idea is that a loan of P dollars today will pay back the principal P plus some interest income. At an interest rate r , this is an additional rP dollars. Thus, 1 dollar today “buys” $(1 + r)$ dollars tomorrow. Put differently, the price of 1 of today’s dollars is $(1 + r)$ of tomorrow’s. Similarly, for a payment of F dollars tomorrow, how many dollars would somebody be willing to pay? $F/(1 + r)$, of course, since $F/(1 + r) + rF/(1 + r) = F$. We therefore have an equivalence between P dollars today and F dollars tomorrow, provided that $P(1 + r) = F$, or $P = F/(1 + r)$. By convention the value P is called the **present value** of F , while F is the **future value** of P .

Given these conventions, let us now derive the budget set of a consumer who may borrow and lend, but has no (physical) investment technology.

Storage is not an option in order to move consumption from period 2 to period 1. The consumer may, however, forgo some future consumption for current consumption by purchasing the appropriate borrowing contract. Based on what we have said above, if he is willing to pay, in period 2, an amount of $(m_2 - c_2) > 0$, then in period 1 he can receive at most $(m_2 - c_2)/(1+r)$ units. Thus one constraint in the budget is $m_1 \leq c_1 \leq m_1 + (m_2 - c_2)/(1+r)$. On the other hand, consumption can be postponed in two ways: storage and lending. If the consumer stores the good his constraint on period 2 consumption is as derived previously, $c_2 \leq m_2 + \delta(m_1 - c_1)$. If he uses the market instead, his constraint becomes $c_2 \leq m_2 + (m_1 - c_1)(1+r)$. As long as δ and r are both strictly positive the second of these is a strictly larger set than the first. Since consumption is a good (more is better) the consumer will not use storage, and the effective constraint on future consumption will be $c_2 \leq m_2 + (m_1 - c_1)(1+r)$. This is indicated in the right diagram in Figure 3.1, where there are two “budget lines” for increasing future consumption. The lower of the two is the one corresponding to storage, the higher the one corresponding to a positive interest rate on loans.

Manipulation of the two constraints shows that they are really identical. Indeed, the consumer’s budget can be expressed as any of

$$\begin{aligned}
 B &= \{(c_1, c_2) \mid c_2 \leq m_2 + (1+r)(m_1 - c_1), c_1, c_2 \geq 0\}, \\
 &= \{(c_1, c_2) \mid c_1 \leq m_1 + \frac{(m_2 - c_2)}{1+r}, c_1, c_2 \geq 0\}, \\
 &= \{(c_1, c_2) \mid (m_2 - c_2) + (1+r)(m_1 - c_1) \geq 0, c_1, c_2 \geq 0\}, \\
 &= \{(c_1, c_2) \mid (m_1 - c_1) + \frac{(m_2 - c_2)}{1+r} \geq 0, c_1, c_2 \geq 0\}.
 \end{aligned}$$

The first and third of these are expressed in future values, the second and fourth in present values. It does not matter which you choose as long as you adopt one perspective and stick to it! (If you recall, we are free to pick a numeraire commodity. Here this means we can choose a period and denominate everything in future- or present-value equivalents for this period. This, by the way, is the one key “trick” in doing inter-temporal economics: pick one time period as your “base”, convert everything into this time period, and then stick to it! Finally note that we are interested in the budget line, generally. Replacing inequalities with equalities in the above you note that all are just different ways of writing the equation of a straight line through the endowment point (m_1, m_2) with a slope of $(1+r)$.




Figure 3.1: Storage without and with markets, no (physical) investment

Storage, Investment, Full Markets

What if everything is possible? As above, storage will normally not figure into the problem, so we will ignore it for a moment. How do (physical) investment and market investment interact in determining the budget? As in the case of storage and investment, the first thing to realize is that the (physical) investment technology will be used to the point where its marginal (gross rate of) return $f'(\cdot)$ is equal to that of putting the marginal dollar into the financial market. That is, the optimal investment is determined by $f'(\bar{x}) = (1 + r)$. The second key point is that with perfect capital markets it is possible to borrow money to invest. The budget thus is bounded by a straight line through the point $(m_1 - \bar{x}, m_2 + f(\bar{x}))$ with slope $(1 + r)$.

3.1.2 Utility maximization

After the budget has been derived the consumer's problem is solved in the usual fashion, and all previous equations which characterize equilibrium apply. Suppose the case of markets in which an interest rate of r is charged. Then it is necessary that

$$\frac{u_1(c)}{u_2(c)} = \frac{(1 + r)}{1}.$$

What is the interpretation of this equation? Well, on the right hand side we have the slope of the budget line, which would usually be the price ratio p_1/p_2 . That is, the RHS gives the cost of future consumption in terms of current consumption (recall that in terms of units of goods we have $(1/p_2)/(1/p_1)$.) On the left hand side we have the ratio of marginal utilities, that is, the MRS. It tells us the consumer's willingness to trade off current consumption for future consumption. This will naturally depend on the consumer's time preferences.

Let us take a concrete example. In macro-economics we often use a so-called **time-separable utility function** like $u(c_1, c_2) = \ln c_1 + \beta \ln c_2$; $\beta < 1$. This specification says that consumption is ranked equally within each period (the same sub-utility function applies within each period) but that future utility is not as valuable as today's and hence is discounted. One of the key features of such a separable specification is that the marginal utilities of today's consumption and tomorrow's are independent. That is, if I consider $\partial u(c_1, c_2) / \partial c_i$, I find that it is only a function of c_i and not of c_j . This feature makes life much easier if the goal is to solve some particular model or make some predictions. Note that the **discount factor** β is related to the consumer's **discount rate** ρ : $\beta = 1 / (1 + \rho)$. For this specific function, the equation characterizing the consumer's optimum then becomes

$$\frac{c_2}{\beta c_1} = (1 + r) \quad \Rightarrow \quad \frac{c_2}{c_1} = \beta(1 + r) \quad \Rightarrow \quad \frac{c_2}{c_1} = \frac{1 + r}{1 + \rho}.$$

Some interesting observations follow from this. First of all, if the private discount rate is identical to the market interest rate, $\rho = r$, the consumer would prefer to engage in **perfect consumption smoothing**. The ideal consumption path has the consumer consume the same quantity each period. If the consumer is more impatient than the market, so that $\rho > r$, then the consumer will favour current consumption, while a patient consumer for whom $\rho < r$ will postpone consumption. This is actually true in general with additively separable functions, since $u'(c_1) = u'(c_2)$ iff $c_1 = c_2$.

In order to achieve the preferred consumption path the consumer will have to engage in the market. He will either be a **lender** ($c_1 < m_1$) or a **borrower** ($c_1 > m_1$.) This, of course, depends on the desired consumption mix compared to the consumption mix in the consumer's endowment.

Finally, we can do the usual **comparative statics**. How will a change in relative prices affect the consumer's wellbeing and consumption choices? The usual facts from revealed preference theory apply: If the interest rate increases (i.e., current consumption becomes more expensive relative to future consumption) then

- A lender will remain a lender.
- A borrower will borrow less (assuming normal goods) and may switch to being a lender.
- A lender will become better off.
- A borrower who remains a borrower will become worse off.

- A borrower who switches to become a lender may be worse or better off.

These can be easily verified by drawing the appropriate diagram and observing the restrictions implied by the weak axiom of revealed preference.

We can also see some of these implications by considering the Slutsky equation for this case. In this case the demand function for period 1 consumption is $c_1(p_1, p_2, M)$, where $M = p_1m_1 + p_2m_2$. It follows from the chain rule that

$$\frac{dc_1(\cdot)}{dp_1} = \frac{\partial c_1(\cdot)}{\partial p_1} + \frac{\partial c_1(\cdot)}{\partial M} \frac{\partial M}{\partial p_1},$$

but $\partial M/\partial p_1 = m_1$. The Slutsky equation tells us that

$$\frac{\partial c_1(\cdot)}{\partial p_1} = \frac{\partial h_1(\cdot)}{\partial p_1} - c_1(\cdot) \frac{\partial c_1(\cdot)}{\partial M}$$

and hence we obtain

$$\frac{dc_1(\cdot)}{dp_1} = \frac{\partial h_1(\cdot)}{\partial p_1} - (c_1(\cdot) - m_1) \frac{\partial c_1(\cdot)}{\partial M}.$$

This equation is easily remembered since it is really just the Slutsky equation as usual, where the weighting of the income effect is by **market purchases only**. In the standard model without endowment, all good 1 consumption is purchased, and hence subject to the income effect. With an endowment, only the amount traded on markets is subject to the income effect.

Now to the analysis of this equation. We know that the substitution effect is negative. We also know that for a normal good the income effect is positive. The sign of the whole expression therefore depends on the term in brackets, in other words on the lender/borrower position of the consumer. A borrower has a positive bracketed term. Thus the whole expression is certainly negative and a borrower will consume less if the price of current consumption goes up. A lender will have a negative bracketed term, which cancels the negative sign in the expression, and we know that the total effect is less negative than the substitution effect. In fact, the (positive) income effect could be larger than the (negative) substitution effect and current consumption could go up!

3.2 Real Interest Rates

So far we have used the usual economic notion of prices as exchange rates of goods. In reality, prices are not denoted in terms of some numeraire commodity, but in terms of money (which is **not** a good.) This may lead to the phenomenon that there is a price ratio of goods to money, and that this may change over time, an effect called inflation if money prices of goods go up (and deflation otherwise.) We can modify our model for this case by fixing the price of the numeraire at only one point in time. We then can account for inflation/deflation. Doing so will require a differentiation between nominal and real interest rates, however, because you get paid back the principal and interest on an investment in units of money which has a different value (in terms of real goods) compared to the one you started out with. So, let $p_1 = 1$ be the (money) price of the consumption good in period 1 and let p_2 be the (money) price of the consumption good in period 2. Let the nominal interest rate be r , which means that you will get interest of r units of money per unit. The budget constraint for the two period problem then becomes:

$$B :: \quad p_2 c_2 = p_2 m_2 + (1 + r)(m_1 - c_1).$$

In other words, the total monetary value of second period consumption can at most be the monetary value of second period endowment plus the monetary value of foregone first period consumption.

We can now ask two questions: what is the **rate of inflation** and what is the **real interest rate**.

The rate of inflation is the rate π such that $(1 + \pi)p_1 = p_2$, i.e., $\pi = (p_2 - p_1)/p_1$. The real interest rate can now be derived by rewriting the above budget to have the usual look:

$$c_2 = m_2 + \frac{1 + r}{p_2}(m_1 - c_1) = m_2 + \frac{1 + r}{1 + \pi}(m_1 - c_1) = m_2 + (1 + \hat{r})(m_1 - c_1).$$

Thus, the real interest rate is $\hat{r} = (r - \pi)/(1 + \pi)$, and as a rule of thumb (for small π) it is common to simply use $\hat{r} \sim (r - \pi)$. Note however that this exaggerates the real interest rate.

3.3 Risk-free Assets

Another application of this methodology is for a first look at assets. It is easiest if we first look at financial assets only. A financial asset is really just

a promise to pay (sometimes called an IOU, from the phrase “I owe you”.) If we assume that the promise is known to be kept, that the amount it will pay back is known, and that the value of what it pays back when it does is known, then we have a **risk-free asset**. While there are few such things, a government bond comes fairly close. Note that assets do not have to be financial instruments: they can also be real, such as a dishwasher, car, or plant (of either kind, actually.) By calling those things assets we focus on the fact that they have a future value.

3.3.1 Bonds

A **bond** is a financial instrument issued by governments or large corporations. Bonds are, as far as their issuer is concerned, a loan. A bond has a so called **face value**, which is what it promises to pay the owner of the bond at the **maturity date** T of the the bond. A normal bond also has a a sequence of **coupons**, which used to be literally coupons that were attached to the bond and could be cut off and exchanged for money at indicated dates. If we denote by F the face value of the bond, and by C the value of each coupon, we can now compute the **coupon rate** of the bond. For simplicity, assume a yearly coupon for now (we will see more financial math later which allows conversion of compounding rates into simple rates). In that case the **coupon rate** is simply $\frac{C}{F}100\%$. A **strip bond** is a bond which had all its coupons removed (“stripped”). The coupons themselves then generate a **simple annuity** — a fixed yearly payment for a specified number of years — which could be sold separately. Another special kind of bond is a **consol**, which is a bond which pays its coupon rate forever, but never pays back any face value.

We can now ask what the price of such a bond should be today. Since the bond bestows on its owner the right to receive specified payments on specified future dates, its value is the current value of all these future payments. Future payments are, of course, converted into their present values by using the appropriate interest rate:

$$PV = \frac{1}{(1+r)}C + \frac{1}{(1+r)^2}C + \dots + \frac{1}{(1+r)^T}C + \frac{1}{(1+r)^T}F.$$

Denote the coupon rate by c . Then we can simplify the above equation to yield

$$PV = F \left(\frac{1}{(1+r)^T} + c \sum_{i=1}^T \frac{1}{(1+r)^i} \right).$$

From this equation follows one of the more important facts concerning the price of bonds and the interest rate. Recall that once a bond is issued T , F , and C are fixed. That means that the present price of the bond needs to adjust as the interest rate r varies. Since r appears in the denominator, as the interest rate rises the present value of a bond falls. Bond prices and interest rates are **inversely related**. Is the present price of a bond higher or lower than the maturity value? In the latter case the bond is trading at a discount, in the former at a premium. This will depend on the relationship between the interest rate and the coupon rate. Intuitively, if the coupon rate exceeds the interest rate the bond is more valuable, and thus its price will be higher.

3.3.2 More on Rates of Return

You may happen to own an asset which, as mentioned previously, is the right to a certain stream of income or benefits (services). This asset happens to also have a current market price (we will see later where that might come from and what conditions it will have to satisfy.) The question you now have is, what is this asset's rate of return?

Let us start with the simple most case of no income/services, but an asset that only has a current and a future price, p_0 and p_1 , respectively. Your rate of return then is the money you gain (or loose) as a percentage of the cost of the asset, in other words the **per dollar rate of return** is $\frac{p_1 - p_0}{p_0}$. We can now ask what conditions such a rate of return might have to satisfy in equilibrium. Assume, therefore, a case where there is complete certainty over all assets' returns, i.e., future prices. All consumers would want to buy that asset which has the highest rate of return, since that would move their individual budget line out the most. In equilibrium, if a consumer is willing to hold more than one asset, then it must be true that both assets have the same rate of return. Call this rate r . Then we know that $r = \frac{p_1 - p_0}{p_0} = \frac{p_1}{p_0} - 1$, or $(1 + r) = \frac{p_1}{p_0}$. This condition must hold for all assets which are held, otherwise there would be arbitrage opportunities. It is therefore also known as a **zero arbitrage condition**. Recall that arbitrage refers to the activity of selling high while buying low in order to make a profit. For example, assume that there were an asset for which $p_1/p_0 > 1 + r$. Then what I should do is to borrow money at the interest rate r , say I dollars, and use those funds to buy the good in question, i.e., purchase I/p_0 units. Then wait and sell the goods in the next period. That will yield $I p_1/p_0$ dollars. I also have to pay back my loan, at

$(1+r)I$ dollars, and thus I have a profit of $p_1/p_0 - (1+r)$ per dollar of loan. Note that the optimal loan size would be infinite. However, the resulting large demand would certainly drive up current prices (while also lowering future prices, since everybody expects a flood of the stuff tomorrow), and this serves to reduce the profitability of the exercise. In a zero-arbitrage equilibrium we therefore must have $(1+r) = p_1/p_0$, or, more tellingly, $p_0 = \frac{p_1}{1+r}$. The correct market price for an asset is its discounted future value!

This discussion has an application to the debate about pricing during supply or demand shocks. For example, gasoline prices during the Gulf war, or the alleged price-gouging in the ice-storm areas of Quebec and Ontario: What should the price of an item be which is already in stock? Many people argue that it is unfair to charge a higher price for in-stock items. Only the replacement items, procured at higher cost, should be sold at the higher cost. While this may be “ethical” according to some, it is easily demonstrated to violate the above rule: The price of the good tomorrow will be determined by demand and supply tomorrow, and apparently all are agreed that that price might well be higher due to a large shift out in the demand and/or reduction in supply. Currently I own that good, and have therefore the choice of selling it tomorrow or selling it today. I would want to obtain the appropriate rate of return on the asset, which has to be equal between the two options. Thus I am only willing to part with it now if I am offered a higher price which foreshadows tomorrow's higher price. Should I be forced not to do so I am forced to give money away against my will and better judgment. This would normally be considered unethical by most (just try and force them to give you money.)

Of course, assets are not usually all the same, and we will see this later when we introduce uncertainty. For example, a house worth \$100,000 and \$100,000 cash are not equivalent, since the cash is immediately usable, while the house may take a while to sell — it is less “liquid.” The same is true for thinly traded stocks. Such assets may carry a **liquidity premium** — an illiquidity punishment, really — and will have a higher rate of return in order to compensate for the potential costs and problems in unloading them. This can, of course, be treated in terms of risk, since the realization of the house's value is a random variable, at least in time, if not in the amount. Of course, there are other kinds of risk as well, and in general the future price of the asset is not known. (Note that bonds are an exception to some degree. If you choose to hold the bond all the way to the maturity date you do know the precise stream of payments. If you sell early, you face the uncertain sale price which depends on the interest rate at that point in time.)

Assets may also yield consumption returns while you hold them: a car or house are examples, as are dividend payments of stocks or interest payments of bonds. For one period this is still simple to deal with: The asset will generate benefit (say rent saved, or train tickets saved) of b and we thus compute the rate of return as $\frac{p_1 - p_0 + b}{p_0}$. If the consumer holds multiple assets in equilibrium, then we again require that this be equal to the rate of return on other assets. Complicating things in the real world is the fact that assets often differ in their tax treatment. For example, if the house is a principal residence any capital gains (the tax man's term for $p_1 - p_0$, and to add insult to injury they ignore inflation) are tax free. For another asset, say a painting, this is not true. Equilibrium requires, of course, that the rates of return as perceived by the consumer are equalized, and thus we may have to use an after tax rate for one asset and set it equal to an untaxed rate for another.

3.3.3 Resource Depletion

The simple discounting rules above can also be applied to gain some first insights into resource economics. We can analyse the question of simple resource depletion: at what rate should we use up a non-renewable resource. We can also analyse when a tree (or forest) should to be cut down.

Assume a non-renewable resource currently available at quantity S . For simplicity, first assume a fixed annual demand D . It follows that there are S/D years left, after which we assume that an alternative has to be used which costs C . Thus the price in the last year should be $p_{S/D} = C$. Arbitrage implies that $p_{t+1} = (1+r)p_t$, so that $p_0 = C/(1+r)^{S/D}$. Note that additional discoveries of supplies lower the price since they increase the time to depletion, as do reductions in demand. Lowering the price of the alternative also lowers the current price. Finally, increases in the discount rate lower the price.

This approach has a major flaw, however. It assumes demand and supply to be independent of price. So instead, let us assume some current price p_0 as a starting value and let us focus on supply. When will the owner of the resource be willing to sell? If the market rate of return on other assets is r then the resource, which is just another asset, will also have to generate that rate of return. Therefore $p_1 = (1+r)p_0$, and in general we'd have to expect $p_t = (1+r)^t p_0$. Note that the price of the resource is therefore increasing with time, which, in general equilibrium, means two things: demand will

fall as customers switch to alternatives, and substitutes will become more competitive. Furthermore we might expect more substitutes to be developed. We will ultimately run out of the resource, but it is nearly always wrong to simply use a linear projection of current use patterns. This fact has been established over and over with various natural resources such as oil, tin, copper, titanium, etc.

What about renewable resources? Consider first the ‘European’ model of privately owned land for timber production as an example. Here we have a company who owns an asset — a forest — which it intends to manage in order to maximize the present value of current and future profits. When should it harvest the trees? Each year there is the decision to harvest the tree or not. If it is cut it generates revenue right away. If it continues to grow it will not generate this revenue but instead generate more revenue tomorrow (since it is growing and there will be more timber tomorrow.) It follows that the two rates of return should be equalized, that is, the tree should be cut once its growth rate divided by its current size has slowed to the market interest rate. This fact has a few implications for forestry: Faster growing trees are a better investment, and thus we see mostly fast growing species replanted, instead of, say, oaks, which grow only slowly. (This discussion is *ceteris paribus* — ignoring general equilibrium effects.) Furthermore, what if you don’t own the trees? What if you are the James Bond of forestry, with a (time-limited) license to kill? In that case you will simply cut the trees down either immediately or before the end of your license, depending on the growth rate. Of course, in Canada most licenses are for mature forests, which nearly by definition have slow or no growth — thus the thing to do is to clear cut and get out of there. The Europeans, critical of clear-cutting, forget that they have long ago cut nearly all of their mature forests and are now in a harvesting model with mostly high growth forests.

As a final note, notice that lack of ownership will also impact the re-planting decision. As we will see later in the course, if we treat the logger as an agent of the state, the state has serious incentive problems to overcome within this principal agent framework.

3.3.4 A Short Digression into Financial Economics

I thought it might be useful to provide you with a short refresher or introduction to multi-period present value and compound interest computations. For starters, assume you put \$1 in the bank at 5% interest, computed yearly, and that all interest income is also reinvested at this 5% rate. How much

money will you have in each of the following years? The answer is

$$1.05, 1.05^2, 1.05^3, \dots 1.05^t.$$

The important fact about this is that a simple interest rate and a compounded interest rate are not the same, since with compounding there is interest on interest. For example, if you get a loan at 12%, it matters how often this is compounded. Let us assume it is just simple interest; You then owe \$1.12 for every dollar you borrowed at the end of one year. What you will quickly find out is that banks don't normally do that. They at least compound semi-annually, and normally monthly. Monthly compounding would mean that you will owe $(1 + \frac{.12}{12})^{12} = 1.1268$. On a million dollar loan this would be a difference of \$6825.03. In other words, you are really paying not a 12% interest rate but a 12.6825% simple interest rate. It is therefore very important to be sure to know what interest rate applies and how compounding is applied (semi-annual, monthly, etc.?)

Here is a handy little device used in many circles: **the rule of 72**, sometimes also referred to as the rule of 69. It is used to find out how long it will take to double your money at any given interest rate. The idea is that it will approximately take $72/r$ periods to double your money at an interest rate of r percent. The proof is simple: we want to solve for the t for which

$$\left(1 + \frac{r\%}{100}\right)^t = 2 \quad \Rightarrow \quad t \ln \left(1 + \frac{r\%}{100}\right) = \ln 2.$$

However, for small x we know that $\ln(1+x) \sim x$, thus

$$t \frac{r\%}{100} \sim \ln 2 \quad \Rightarrow \quad t \sim \frac{100 \ln 2}{r\%} = \frac{69.3147}{r\%}$$

but of course 72 has more divisors and is much easier to work with.

The power of compounding also comes into play with **mortgages** or other installment loans. A mortgage is a promise to pay every period for a specified length (typically 25 years, i.e., 300 months) a certain payment p . This is also known as a **simple annuity**. What is the value of such a promise, i.e., its present value? We need to compute the value of the following sum:

$$\delta p + \delta^2 p + \delta^3 p + \dots + \delta^n p.$$

Here $\delta = 1/(1+r)$, where r is the interest rate we use per period. Thus

$$\begin{aligned} \text{PV} &= \delta (p + \delta p + \delta^2 p + \dots + \delta^{n-1} p) \\ &= \delta p (1 + \delta + \delta^2 + \dots + \delta^{n-1}) \end{aligned}$$

$$\begin{aligned}
 &= \delta p \frac{1 - \delta^n}{1 - \delta} \\
 \text{PV} &= \frac{1 - (1 + r)^{-n}}{r} p
 \end{aligned}$$

(Recall in the above derivation that for $\delta < 1$ we have $\sum_{i=1}^{\infty} \delta^i = 1/(1 - \delta)$.)

The above equation relates four variables: the principal amount, the payment amount, the payment periods, and the interest rate. If you fix any three this allows you to derive the fourth after only a little bit of manipulation. A final note: In Canada a mortgage can be at most compounded semi-annually. Thus the effective interest rate per month is derived by solving $(1 + r/2)^2 = (1 + r_m)^{12}$. If you are quoted a 12% interest rate per year the monthly rate is therefore $(1.06)^{1/6} - 1 = 0.975879418\%$. The effective yearly interest rate in turn is $(1.06)^2 - 1 = 12.36\%$, and by law the bank is supposed to tell you about that too. Given the above, and the fact that nearly all mortgages are computed for a 25 year term (but seldom run longer than 5 years, these days), the monthly payment at a 10% yearly interest rate for an additional \$1000 on the mortgage is \$8.95. Before you engage in mortgages it would be a good idea to program your spreadsheet with these formulas and convince yourself how bi-weekly payments reduce the total interest you pay, how important a couple of percentage points off the interest rate are to your monthly budget, etc.

3.4 Review Problems

Question 1: There are three time periods and one consumption good. The consumer's endowments are 4 units in the first period, 20 units in the second, and 1 unit in the third. The money price for the consumption good is known to be $p = 1$ in all periods (no inflation.) Let r_{ij} denote the (simple, nominal) interest rate from period i to j .

- a) State the restrictions on r_{12} , r_{23} and r_{13} implied by zero arbitrage.
- b) Write down the consumer's budget constraint assuming the restriction in (a) holds. Explain why it is useful to have this condition hold (i.e., point out what would cause a potential problem in how you've written the budget if the condition in (a) fails.
- c) Draw a diagrammatic representation of the budget constraint in periods 2 and 3, being careful to note how period 1 consumption influences this diagram.

Question 2: There are two goods, consumption today and tomorrow. Joe

has an initial endowment of $(100, 100)$. There exists a credit market which allows him to borrow or lend against his initial endowment at market interest rates of 0%. A borrowing constraint exists which prevents him from borrowing against more than 60% of his period 2 endowment. Joe also possesses an investment technology which is characterized by a production function $x_2 = 10\sqrt{x_1}$. That is, an investment of x_1 units in period 1 will lead to x_2 units in period 2.

a) What is Joe's budget constraint? A very clearly drawn and well labelled diagram suffices, or you can give it mathematically. Also give a short explanatory paragraph how the set is derived.

b) Suppose that Joe's preferences can be represented by the function $U(c_1, c_2) = \exp(c_1^4 c_2^6)$. (Here $\exp()$ denotes the exponential function.) What is Joe's final consumption bundle, how much does he invest, and what are his transactions in the credit market.

Question 3: Anna has preferences over her consumption levels in two periods which can be represented by the utility function

$$u(c_1, c_2) = \min \left\{ \frac{23}{22} \left(\frac{12}{10} c_1 + c_2 \right), \frac{13}{10} c_1 + c_2 \right\}.$$

a) Draw a carefully labelled representation of her indifference curve map.

b) What is her utility maximizing consumption bundle if her initial endowment is $(9, 8)$ and the interest rate is 25%.

c) What is her utility maximizing consumption bundle if her initial endowment is $(5, 12)$ and the interest rate is 25%.

d) Assume she can lend money at 22% and borrow at 28%. What would her endowment have to be for her to be a lender, a borrower?

e) Assume she can lend money at 18% and borrow at 32%. Would Anna ever trade at all? (Explain.)

Question 4: Alice has preferences over consumption in two periods represented by the utility function $u_A(c_1, c_2) = \ln c_1 + \alpha \ln c_2$, and an endowment of $(12, 6)$. Bob has preferences over consumption in two periods represented by the utility function $u_B(c_1, c_2) = c_1 + \beta c_2$, and an endowment of $(8, 4)$.

a) Draw an appropriately labelled representation of this exchange economy in order to "prime" your intuition. (Indicate the indifference maps and the Contract Curve.)

b) Assuming, of course, that both α and β lie strictly between zero and one, what is the equilibrium interest rate and allocation?

Chapter 4

Uncertainty

So far, it has been assumed that consumers would know precisely what they were buying and getting. In real life, however, it is often the case that an action does not lead to a definite outcome, but instead to one of many possible outcomes. Which of these occurs is outside the control of the decision maker. It is determined by what is referred to as “nature.” These situations are ones of **uncertainty** — it is uncertain what happens. Often, however, the probabilities of the different possibilities are known from past experience, or can be estimated in some other way, or indeed are assumed based on some personal (subjective) judgment. Economists then speak of **risk**.

Note that our “normal” model is already handling such cases if we take it at its most general level: commodities in the model were supposed to be fully specified, and could, in principle, be state contingent. We will develop that interpretation further later on in this chapter. First, however, we will develop a more simple model which is designed to bring the role of probabilities to the fore.

One of the key facts about situations involving risk/uncertainty is that the consumer’s wellbeing does not only depend on the various possible outcomes, and which occurs in the end, but also on how likely each outcome is. The standard model of chapter 2 does not allow an explicit role for such probabilities. They are somehow embedded in the utility function and prices. In order to compare situations which differ only in the probabilities, for example, it would be nice to have probabilities explicitly in the model formulation. A particularly simple model that does this holds the outcomes fixed, they will all be assumed to lie in some set of alternatives X , and focuses on the different probabilities with which they occur. We call such a list

of the probabilities for each outcome a **lottery**.

Definition 1 A **simple lottery** is a list $L = (p_1, p_2, \dots, p_N)$ of probabilities for the N different outcomes in X , with $p_i \geq 0$, $\sum_{i=1}^N p_i = 1$.

If we have a suitably defined continuous space of outcomes, for example \mathfrak{R}_+ for the outcome “wealth”, we can view a probability distribution as a lottery. We will therefore, if it is convenient, use the cumulative distribution function (cdf) $F(\cdot)$ or the probability density function (assuming it exists) $f(x)$ to denote lotteries over (one-dimensional) continuous outcome spaces.¹

Of course, lotteries could have as outcomes other lotteries, and such lotteries are called **compound lotteries**. However, any such compound lottery will lead to a probability distribution over final outcomes which is equivalent to that of some simple lottery. In particular, if some compound lottery leads with probability α_i to some simple lottery i , and if that simple lottery i in turn assigns probability p_n^i to outcome n , then we have the total probability of outcome n given by $p_n = \sum_i \alpha_i p_n^i$.

Assumption 1 Only the probability distribution over “final” outcomes matters to the consumer \implies preferences are over simple lotteries only.

Note that this is a restrictive assumption. For example, it does not allow the consumer to attach any direct value to the fact that he is involved in a lottery, i.e., gambling pleasure (or dislike) in itself is not allowed under this specification. The consumer is also assumed not to care how a given probability distribution over final outcomes arises, only what it is. This is clearly an abstraction which loses quite a bit of richness in consumer behaviour. On the positive side stands the mathematical simplicity of this setting. Probability distributions over some set are a fairly simple framework to work with.

We now have assumed that a consumer cares only about the probability distribution over a finite set of outcomes. Just as before, we will assume that the consumer is able to rank any two such distributions, in the usual way.

¹Recall that a cumulative distribution function $F(\cdot)$ is always defined and that $F(x)$ represents the probability of the underlying random variable taking on a value less or equal to x . If the density exists (and it does not have to!) then we denote it as $f(x)$ and have the identity $F(x) = \int^x f(t)dt$. We can then denote the mean, say, in two different ways: $\mu = \int tf(t)dt = \int t dF(t)$, depending on if we know that the pdf exists, or want to allow that it does not.

Assumption 2 *The preferences over simple lotteries are rational and continuous.*

This latter assumption guarantees us the existence of a utility function representing these preferences. Note that rationality is a strong assumption in this case. In particular, rationality requires transitivity of the preferences, that is, $L \succeq L', L' \succeq \hat{L} \Rightarrow L \succeq \hat{L}$. For different lotteries this may be hard to believe, and there is some evidence that real life consumers violate this assumption. These violations of transitivity are “more common” in this setting compared to the model without risk/uncertainty. Yet, without rationality the model would have no predictive power. We further assume that the preferences satisfy the following assumption:

Assumption 3 *Preferences are Independent of what other outcomes may have been available: $L \succeq L' \Rightarrow \alpha L + (1 - \alpha)\hat{L} \succeq \alpha L' + (1 - \alpha)\hat{L}$.*

This seems sensible on the one hand, since outcomes are mutually exclusive — one and only one of the outcomes will happen in the end — but is restrictive since consumers often express regret: having won \$5000, the evaluation of that win often depends on the fact if this was the top price available or the consolation price, for example. The economic importance of this assumption is that we have only one utility index over outcomes in which preferences over lotteries (the distribution over available outcomes) does not enter. Once we have done this, there is only one more assumption:

Assumption 4 *The utility of a lottery is the expected value of the utilities of its outcomes:*

$$U(L) = \sum_{i=1}^n p_i u(x_i) \quad \left(U(L) = \int u(x) dF(x) \right)$$

This form of a utility function is called a von Neumann-Morgenstern, or vN-M, utility function. Note that this name applies to $U(\cdot)$, not $u(\cdot)$. The latter is sometimes called a Bernoulli utility function. The vN-M utility function is unique only up to a positive affine transformation, that is, the same preferences over lotteries are expressed by $V(\cdot)$ and $U(\cdot)$ if and only if $V(\cdot) = aU(\cdot) + b, a > 0$. We are allowed to scale the utility index and to change its slope, but we are not allowed to change its curvature. The reason for this should be clear. Suppose we compare two lotteries, L and \hat{L} ,

which differ only in that probability is shifted between outcomes k and j and outcomes m and n . Suppose $U(L) > U(\hat{L})$, so:

$$U(L) = \sum_{i=1}^n p_i u(x_i) > U(\hat{L}) = \sum_{i=1}^n \hat{p}_i u(x_i)$$

$$\sum_{i=1}^n p_i u(x_i) - \sum_{i=1}^n \hat{p}_i u(x_i) > 0$$

$$(p_k - \hat{p}_k)u(x_k) + (p_j - \hat{p}_j)u(x_j) + (p_m - \hat{p}_m)u(x_m) + (p_n - \hat{p}_n)u(x_n) > 0$$

$$(p_k - \hat{p}_k)(u(x_k) - u(x_j)) + (p_m - \hat{p}_m)(u(x_m) - u(x_n)) > 0$$

This comparison clearly depends on both, the differences in probabilities as well as the differences in the utility indices of the outcomes. If we multiply $u(\cdot)$ by a constant, it will factor out of the last line above. If, however, we were to transform the function $u(\cdot)$, even by a monotonic transformation, we would change the difference between outcome utilities, and this could change the above comparison. In fact, as we shall see later, the curvature of the Bernoulli utility index $u(\cdot)$ is crucial in determining the consumer's behaviour with respect to risk, and will be used to measure the consumer's risk aversion.

Before we proceed to that, some famous paradoxes relating to uncertainty and our assumptions.

Allais Paradox: The Allais paradox shows that consumers may not satisfy the axioms we had assumed. It considers the following case: Consider a space of outcomes for a lottery given by $C = (25, 5, 0)$ in hundred thousands of dollars. Subjects are then asked which of two lotteries they would prefer,

$$L_A = (0, 1, 0) \quad \text{or} \quad L_B = (.1, .89, .01).$$

Often consumers will indicate a preference for L_A , probably because they foresee that they would regret to have been greedy if they end up with nothing under lottery B. On the other hand, if they are asked to choose between

$$L_C = (0, .11, .89) \quad \text{or} \quad L_D = (.1, 0, .9)$$

the same consumers often indicate a preference for lottery D. Note that there is little regret possible here, you simply get a lot larger winning in exchange for a slightly lower probability of winning under D . These choices, however, violate our assumptions. This is easily checked by assuming the existence of some $u(\cdot)$: The preference for A over B then indicates that

$$u(5) > .1u(25) + .89u(5) + .01u(0)$$

$$.11u(5) > .1u(25) + .01u(0)$$

$$.11u(5) + .89u(0) > .1u(25) + .9u(0)$$

and the last line indicates that lottery C is preferred to D!

Ellsberg Paradox: This paradox shows that consumers may not be consistent in their assessment of uncertainty. Consider an urn with 300 balls in it, of which precisely 100 are known to be red. The other 200 are blue or green in an unknown proportion (note that this is uncertainty: there is no information as to the proportion available.) The consumer is again offered the choice between two pairs of gambles:

$$\begin{aligned} \text{Choice 1} &= \begin{cases} L_A : \$1000 \text{ if a drawn ball is red.} \\ L_B : \$1000 \text{ if a drawn ball is blue.} \end{cases} \\ \text{Choice 2} &= \begin{cases} L_C : \$1000 \text{ if a drawn ball is NOT red.} \\ L_D : \$1000 \text{ if a drawn ball is NOT blue.} \end{cases} \end{aligned}$$

Often consumers faced with these two choices will choose A over B and will choose C over D. However, letting $u(0)$ be zero for simplicity, this means that $p(R)u(1000) > p(B)u(1000) \Rightarrow p(R) > p(B) \Rightarrow (1 - p(R)) < (1 - p(B)) \Rightarrow p(\neg R) < p(\neg B) \Rightarrow p(\neg R)u(1000) < p(\neg B)u(1000)$. Thus the consumer should prefer D to C if choice were consistent.

Other problems with expected utility also exist. One is the intimate relation of risk aversion and time preference which is imposed by these preferences. There consequently is a fairly active literature which attempts to find a superior model for choice under uncertainty. These attempts mostly come at the expense of much higher mathematical requirements, and many still only address one or the other specific problem, so that they too are easily 'refuted' by a properly chosen experiment.

4.1 Risk Aversion

We will now restrict the outcome space X to be one-dimensional. In particular, assume that X is simply the wealth/total consumption of the consumer in each outcome. With this simplification, the basic attitudes of a consumer concerning risk can be obtained by comparing two different lotteries: one that gives an outcome for certain (a degenerate lottery), and another that has the same expected value, but is non-degenerate. So, let L be a lottery on $[0, \bar{X}]$ given by the probability density $f(x)$. It generates an expected value of wealth of $\int_0^{\bar{X}} xf(x)dx = C$. We can now compare the consumer's utility from obtaining C for certain, and that from the lottery L (which has

expected wealth C .) Compare

$$U(L) = \int_0^{\bar{X}} u(x)f(x)dx \quad \text{to} \quad u\left(\int_0^{\bar{X}} xf(x)dx\right).$$

Definition 2 A risk-averse consumer is one for whom the expected utility of any lottery is lower than the utility of the expected value of that lottery:

$$\int_0^{\bar{X}} u(x)f(x)dx < u\left(\int_0^{\bar{X}} xf(x)dx\right).$$




Figure 4.1: Risk Aversion

The astute reader may notice that this is Jensen's inequality, which is one way to define a concave function, in this case $u(\cdot)$ (see Fig. 4.1.) This is also the reason why only affine transformations were allowed for expected utility functions. Any other transformation would affect the curvature of the Bernoulli utility function $u(\cdot)$, and thus would change the risk-aversion of the consumer. Clearly, consumers with different risk aversion do not have the same preferences, however.² Note that a concave $u(\cdot)$ has a diminishing marginal utility of wealth, an assumption which is quite familiar from introductory courses. Risk aversion therefore implies (and is implied by) the fact

²To belabour the point, consider preferences over wealth represented by $u(w) = w$. In the standard framework of chapter 1 positive monotonic transformations are ok, so that the functions w^2 and \sqrt{w} both represent identical preferences. It is easy to verify that these two functions lead to a quite different relationship between the expected utility and the utility of the expected wealth than the initial one, however. Thus they cannot represent the same preferences in a setting of uncertainty/risk.




Figure 4.2: Risk Neutral and Risk Loving

that additional units of wealth provide additional utility, but at a decreasing rate. Of course, consumers do not have to be risk-averse.

Risk neutral and **risk loving** are defined in the obvious way: The first requires that

$$\int u(x)f(x)dx = u\left(\int xf(x)dx\right).$$

while the second requires

$$\int u(x)f(x)dx > u\left(\int xf(x)dx\right).$$

There is a nice diagrammatic representation of these available if we consider only two possible outcomes (Fig. 4.2).

There are two other ways in which we might define risk aversion, and both reveal interesting facts about the consumer's economic behaviour. The first is by using the concept of a **certainty equivalent**. It is the answer to the question "how much wealth, received for certain, is equivalent (in the consumer's eyes, according to preferences) to a given gamble/lottery?" In other words:

Definition 3 The **certainty equivalent** $C(f, u)$ for a lottery with probability distribution $f(\cdot)$ under the (Bernoulli) utility function $u(\cdot)$ is defined by the equation

$$u(C(f, u)) = \int u(x)f(x)dx.$$

Again a diagram for the two-outcome case might help (Fig. 4.3).




Figure 4.3: The certainty equivalent to a gamble

A risk averse consumer is one for whom the certainty equivalent of any gamble is less than the expected value of that gamble. One useful economic interpretation of this fact is that the consumer is willing to pay (give up expected wealth) in order to avoid having to face the gamble. Indeed, the maximum amount which the consumer would pay is the difference $\int w f(w) dw - C(f, u)$. This observation basically underlies the whole insurance industry: risk-averse consumers are willing to pay in order to avoid risk. A well diversified insurance company will be risk neutral, however, and therefore is willing to provide insurance (assume the risk) as long as it guarantees the consumer not more than the expected value of the gamble: Thus there is room to trade, and insurance will be offered. (More on that later.)

Another way to look at risk aversion is to ask the following question: If I were to offer a gamble to the consumer which would lead either to a win of ϵ or a loss of ϵ , how much more than fair odds do I have to offer so that the consumer will take the bet? Note that a fair gamble would have an expected value of zero (i.e., 50/50 odds), and thus would be rejected by the (risk averse) consumer for sure. This idea leads to the concept of a **probability premium**.

Definition 4 The **probability premium** $\pi(u, \epsilon, w)$ is defined by

$$u(w) = (0.5 + \pi(\cdot)) u(w + \epsilon) + (0.5 - \pi(\cdot)) u(w - \epsilon).$$

A risk-averse consumer has a positive probability premium, indicating that the consumer requires more than fair odds in order to accept a gamble.

It can be shown that all three concepts are equivalent, that is, a consumer with preferences that have a positive probability premium will be one for whom the certainty equivalent is less than the expected value of wealth and for whom the expected utility is less than the utility of the expectation. This is reassuring, since the certainty equivalent basically considers a consumer with a property right to a gamble, and asks what it would take for him to trade to a certain wealth level, while the probability premium considers a consumer with a property right to a fixed wealth, and asks what it would take for a gamble to be accepted.

4.1.1 Comparing degrees of risk aversion

One question we can now try to address is to see which consumer is more risk averse. Since risk aversion apparently had to do with the concavity of the (Bernoulli) utility function it would appear logical to attempt to measure its concavity. This is indeed what Arrow and Pratt have done. However, simply using the second derivative of $u(\cdot)$, which after all measures curvature, will not be such a good idea. The reason is that the second derivative will depend on the units in which wealth and utility are measured.³ Arrow and Pratt have proposed two measures which largely avoid this problem:

Definition 5 *The Arrow-Pratt measure of (absolute) risk aversion is*

$$r_A = -\frac{u''(w)}{u'(w)}.$$

The Arrow-Pratt measure of relative risk aversion is

$$r_R = -\frac{u''(w)w}{u'(w)}.$$

Note that the first of these in effect measures risk aversion with respect to a fixed amount of gamble (say, \$1). The latter, in contrast, measures risk aversion for a gamble over a fixed percentage of wealth. These points can be demonstrated as follows:

Consider a consumer with initial wealth w who is faced with a small fair bet, i.e., a gain or loss of some small amount ϵ with equal probability.

³You can easily verify this by thinking of the units attached to the second derivative. If the first derivative measures change in utility for change in wealth, then its units must be u/w , while the second derivative is like a rate of acceleration. Its units are u/w^2 .

How much would the consumer be willing to pay in order to avoid this bet? Denoting this payment by I we need to consider (note that $w - I$ is the certainty equivalent)

$$0.5u(w + \epsilon) + 0.5u(w - \epsilon) = u(w - I).$$

Use a Taylor series expansion in order to approximate both sides:

$$\begin{aligned} 0.5(u(w) + \epsilon u'(w) + 0.5\epsilon^2 u''(w)) + 0.5(u(w) - \epsilon u'(w) + 0.5\epsilon^2 u''(w)) \\ \sim u(w) - Iu'(w) \quad . \end{aligned}$$

Collecting terms and simplifying gives us

$$0.5\epsilon^2 u''(w) \sim -Iu'(w) \quad \Rightarrow \quad I \sim \frac{\epsilon^2}{2} \times \frac{-u''(w)}{u'(w)}.$$

Thus the required payment is proportional to the absolute coefficient of risk aversion (and the dollar amount of the gamble.)

On the other hand,

$$\frac{u''w}{u'} = \frac{du' w}{dw u'} = \frac{du'/u'}{dw/w} \sim \frac{\% \Delta u'}{\% \Delta w}.$$

Thus the relative coefficient of risk-aversion is nothing but the elasticity of marginal utility with respect to wealth. That is, it measures the responsiveness of the marginal utility to wealth changes.

Comparing across consumers, a consumer is said to be more risk averse than another if (either) Arrow-Pratt coefficient of risk aversion is larger. This is equivalent to saying that he has a lower certainty equivalent for any given gamble, or requires a higher probability premium.

We can also compare the risk aversion of a given consumer for different wealth levels. That is, we can compute these measures for the same $u(\cdot)$ but different initial wealth. After all, r_A is a function of w . It is commonly assumed that consumers have (absolute) risk aversion which is decreasing with wealth. Sometimes the stronger assumption of decreasing relative risk aversion is made, however. Note that a constant absolute risk aversion implies increasing relative risk aversion. Finally, note also that the only functional form for $u(\cdot)$ which has constant absolute risk aversion is $u(w) = -e^{-aw}$.

You may wish to verify that a consumer exhibiting decreasing absolute risk aversion will have a decreasing difference between initial wealth and the certainty equivalent (a declining maximum price paid for insurance) on the one hand, and a decreasing probability premium on the other.




Figure 4.4: Comparing two gambles with equal expected value

4.2 Comparing gambles with respect to risk

Another type of comparison of interest is not across consumers or wealth levels, as above, but across different gambles. Faced with two gambles, when do we want to say that one is riskier than the other? We could try to approach this question with purely statistical measures, such as comparisons of the various moments of the two lotteries' distributions. This has the major problem, however, that the consumer may in general be expected to be willing to trade off a higher expected return for higher variance, say. Because of this, a definition based directly on consumer preferences is preferable. Two such measures are commonly employed in economics, first and second order stochastic dominance.

Let us first focus on lotteries with the same expected value. For example, consider the two gambles depicted in Fig. 4.4. The first is a gamble over w_1 and w_2 . The second is a gamble over w_3 and w_4 . Both have an identical expected value of $E(w)$. Nevertheless a risk averse consumer clearly will prefer the second to the first, as inspection of the diagram verifies.

Note that in Fig. 4.4 $E(w) - w_1 > E(w) - w_3$ and $w_2 - E(w) > w_4 - E(w)$. This clearly indicates that the second lottery has a lower variance, and thus that a risk averse consumer prefers to have less variability for a given mean. With multiple possible outcomes the question is not so simple anymore, however. One could construct an example with two lotteries that have the same mean and variance, but which differ in higher moments. What are the "obvious" preferences of a risk averse consumer about skurtosis, say?

This has led to a more general definition for comparing distributions which have the same mean:

Definition 6 Let $F(x)$ and $G(x)$ be two cumulative distribution functions for a one-dimensional random variable (wealth). Let $F(\cdot)$ have the same mean as $G(\cdot)$. $F(\cdot)$ is said to dominate $G(\cdot)$ according to **second order stochastic dominance** if for every non-decreasing concave $u(x)$:

$$\int u(x)dF(x) \geq \int u(x)dG(x)$$

In words, a distribution second order stochastically dominates another if they have the same mean and if the first is preferred by all risk-averse consumers.

This definition has economic appeal in its simplicity, but is one of those definitions that are problematic to work with due to the condition that **for all possible** concave functions something is true. In order to apply this definition easily we need to find other tests.

Lemma 1 Let $F(x)$ and $G(x)$ be two cumulative distribution functions for a one-dimensional random variable (wealth). $F(\cdot)$ dominates $G(\cdot)$ according to **second order stochastic dominance** if

$$\int tg(t)dt = \int tf(t)dt, \text{ and } \int_0^x G(t)dt \geq \int_0^x F(t)dt, \forall x.$$

I.e., if they have the same mean and there is more area under the cdf $G(\cdot)$ than under the cdf $F(\cdot)$ at any point of the distribution.⁴

A concept related to second order stochastic dominance is that of a mean preserving spread. Indeed it can be shown that the two are equivalent.

Definition 7 Let $F(x)$ and $G(x)$ be two cumulative distribution functions for a one-dimensional random variable (wealth). $G(\cdot)$ is a **mean preserving spread** compared to $F(\cdot)$ if x is distributed according to $F(\cdot)$ and $G(\cdot)$ is the distribution of the random variable $x + z$, where z is distributed according to some $H(\cdot)$ with $\int z dH(z) = 0$.

⁴Note that the condition of identical means also implies a restriction on the total areas below the cumulative distributions. After all, $\int_{\underline{x}}^{\bar{x}} tdF(t) = [tF(t)]_{\underline{x}}^{\bar{x}} - \int_{\underline{x}}^{\bar{x}} F(t)dt = \bar{x} - \int_{\underline{x}}^{\bar{x}} F(t)dt$.

The above gives us an easy way to construct a second order stochastically dominated distribution: Simply add a zero mean random variable to the given one.

While it is nice to be able to rank distributions in this manner, the condition of equal means is restrictive. Furthermore, it does not allow us to address the economically interesting question of what the trade off between mean and risk may be. The following concept is frequently employed in economics to deal with such situations.

Definition 8 Let $F(x)$ and $G(x)$ be two cumulative distribution functions for a one-dimensional random variable (wealth). $F(\cdot)$ is said to dominate $G(\cdot)$ according to **first order stochastic dominance** if for every non-decreasing $u(x)$:

$$\int u(x)dF(x) \geq \int u(x)dG(x)$$

This is equivalent to the requirement that $F(x) \leq G(x), \forall x$.

Note that this requires that any consumer, risk averse or not, would prefer F to G . It is often useful to realize two facts: One, a first order stochastically dominating distribution F can be obtained from a distribution G by shifting up outcomes randomly. Two, first order stochastic dominance implies a higher mean, but is stronger than just a requirement on the mean. The other moments of the distribution get involved too. In other words, just because the mean is higher for one distribution than another does not mean that the first dominates the second according to first order stochastic dominance!

4.3 A first look at Insurance

Let us use the above model to investigate a simple model of insurance. To be concrete, assume an individual with current wealth of \$100,000 who faces a 25% probability to lose his \$20,000 car through theft. Assume the individual has vN-M expected utility. The individual's expected utility then is

$$U(\cdot) = 0.75u(100,000) + .25u(80,000).$$

Now assume that the individual has access to an insurance plan. Insurance works as follows: The individual decides on an amount of coverage, C . This

coverage carries a premium of π per dollar. The contract specifies that the amount C will be paid out if the car has been stolen. (Assume that this is all verifiable.) How would our individual choose the amount of coverage? Simple: maximize expected utility. Thus

$$\max_C \{0.75u(100,000 - \pi C) + 0.25u(80,000 - \pi C + C)\}.$$

The first order condition for this problem is

$$(-\pi)0.75u'(100,000 - \pi C) + (1 - \pi)0.25u'(80,000 - \pi C + C) = 0.$$

Before we further investigate this equation let us verify the second order condition. It requires

$$(-\pi)^2 0.75u''(100,000 - \pi C) + (1 - \pi)^2 0.25u''(80,000 - \pi C + C) < 0.$$

Clearly this is only satisfied if $u(\cdot)$ is concave, in other words, if the consumer is risk averse.

So, what does the first order condition tell us? Manipulation yields the condition that

$$\frac{u'(100,000 - \pi C)}{u'(80,000 - \pi C + C)} = \frac{(1 - \pi)}{3\pi}$$

which gives us a familiar looking equation in that the LHS is a ratio of marginal utilities. It follows that total consumption under each circumstance is set so as to set the ratio of marginal utility of wealth equal to some fraction which depends on price and the probabilities. Even without knowing the precise function we can say something about the insurance behaviour, however. To do so, let us compute the **actuarially fair premium**. The expected loss is \$5,000, so that an insurance premium which collects that amount for the \$20,000 insured value would lead to zero expected profits for the insurance firm: $0.75\pi C + 0.25(\pi C - C) = 0 \Rightarrow \pi = 0.25$. An actuarially fair premium simply charges the odds (there is a 1 in 4 chance of a loss, after all.) If we use this fair premium in the above first order condition we obtain

$$\frac{u'(100,000 - \pi C)}{u'(80,000 - \pi C + C)} = 1.$$

Since the utility function is strictly concave it can have the same slope only at the same point, and we conclude that⁵

$$(100,000 - \pi C) = (80,000 - \pi C + C) \Rightarrow C = 20,000.$$

⁵Ok, read that sentence again. Do you understand the usage of the word ‘Since’? I am **not** “cancelling” the u' terms, because those indicate a function. Instead the equation tells us that numerator and denominator must be the same. But for what values of the independent variable wealth does the function $u(\cdot)$ have the same derivative? For none, if $u(\cdot)$ is strictly concave. Therefore the function must be evaluated at the same level of the independent variable.

This is one of the key results in the analysis of insurance: **at actuarially fair premiums a risk averse consumer will fully insure.** Note that the consumer will not bear any risk in this case: wealth will be \$95,000 independent of if the car is stolen, since a \$5,000 premium is due in either case, and if the car is actually stolen it will be replaced. As we have seen before, this will make the consumer much better off than if he is actually bearing the gamble with this same expected wealth level. If you draw the appropriate diagram you can verify that the consumer does not have to pay any of the amount he would be willing to pay (the difference between the expected value and the certainty equivalent.)

If we had a particular utility function we could now also compute the maximal amount the consumer would be willing to pay. We have to be careful, however, how we set up this problem, since simply increasing π will reduce the amount of coverage purchased! So instead, let us approach the question as follows: What fee would the consumer be willing to pay in order to have access to actuarially fair insurance? Let F denote the fee. Then we have the consumer choose between

$$u(95,000 - F) \quad \text{and} \quad 0.25u(80,000) + 0.75u(100,000).$$

(Note that I have skipped a step by assuming full insurance. The left term is the expected utility of a fully insured consumer who pays the fee, the right term is the expected utility of an uninsured consumer. You should verify that the lump sum fee does not stop the consumer from fully insuring at a fair premium.) For example, if $u(\cdot) = \ln(\cdot)$ then simple manipulation yields $F \sim 426$.

It is important to note why we have set up the problem this way. Consider the alternative (based on these numbers and the logarithmic function) and assume that the total payment of \$5,426 which is made in the above case of a fair premium plus fee, were expressed as a premium. Then we get that $\pi = 5426/20000 = 0.2713$. The first order condition for the choice of C then requires that (recall that $\partial \ln(x)/\partial x = 1/x$)

$$\frac{(80,000 + 0.7287C)}{(100,000 - 0.2713C)} = \frac{0.7287}{0.8139} = 0.895318835 \quad \Rightarrow C = 9,810.50.$$

As you can see, if the additional price is understood as a per dollar charge for insured value, the consumer will not insure fully. Of course this is an implication of the previous result — the consumer now faces a premium which is not actuarially fair. Indeed, we could also compute the premium for which the consumer will cease to purchase any insurance. For logarithmic utility like this we would want to compute (remember, we are trying to find

when $C = 0$ is optimal)

$$\frac{80,000}{100,000} = \frac{1 - \pi}{3\pi} \Rightarrow \pi = 0.2941.$$

As indicated before, there is room to trade between insurance providers and risk averse consumers. Indeed, as you can verify in one of the questions at the end of the chapter, there is room for trade between two risk averse consumers if they face different risk or if they differ in their attitudes towards risk (degree of risk aversion.)

4.4 The State-Preference Approach

While the above approach lets us focus quite well on the role of probabilities in consumer choice, it is different in character to the ‘maximize utility subject to a budget constraint’ approach we have so much intuition about. In the first order condition for the insurance problem, for example, we had a ratio of marginal utilities on the one side — but was that the slope of an indifference curve?

As mentioned previously, we can actually treat consumption as involving contingent commodities, and will do so now. Let us start by assuming that the outcomes of any random event can be categorized as something we will refer to as the *states of the world*. That is, there exists a set of mutually exclusive states which are adequate to describe all randomness in the world. In our insurance example above, for example, there were only two states of the world which mattered: Either the car was stolen or it was not. Of course, in more general settings we could think of many more states (such as the car is stolen and not recovered, the car is stolen but recovered as a write off, the car is stolen and recovered with minor damage, etc.) In accordance with this view of the world we now will have to develop the idea of **contingent commodities**. In the case of our concrete example with just two states, a contingent commodity would be delivered only if a particular state (on which the commodity’s delivery is contingent) occurs. So, if there are two states, good and bad, then there could be two commodities, one which promises consumption in the good state, and one which promises consumption in the bad state. Notice that you would have to buy both of these commodities if you wanted to consume in both states. Notice also that nothing requires that the consumer purchase them in equal amounts. They are, after all, different commodities now, even if the underlying good which gets delivered in each state is the same. Finally, note that if one of these commodities were missing

you could not assure consumption in both states (which is why economists make such a fuss about “complete markets” — which essentially means that everything which is relevant can be traded. It does not have to be traded, of course, that is up to people’s choices, but it should be available should someone want to trade.) Of course, after the fact (*ex post* in the lingo) only one of these states does occur, and thus only the set of commodities contingent on that state are actually consumed. Before the fact (before the uncertainty is resolved, called *ex ante*) there are two different commodities available, however.

Once we have this setting we can proceed pretty much as before in our analysis. To be concrete let there be just two states, good and bad. We will now index goods by a subscript b or g to indicate the state in which they are delivered. We will further simplify things by having just one good, consumption (or wealth). Given that there are two states, that means that there are two distinct (contingent) commodities, c_g and c_b . We may now assume that the consumer has our usual vN-M expected utility.⁶ If the individual assessed a probability of π to the good state occurring, then we would obtain an expected utility of consumption of $U(c_g, c_b) = \pi u(c_g) + (1 - \pi)u(c_b)$.

This expression gives us the expected utility of the consumer. The consumers’ objective is to maximize expected utility, as before. It might be useful at this point to assess the properties of this function. As long as the utility index applied to consumption in each state, $u(\cdot)$, is concave, this is a concave function. It will be increasing in each commodity, but at a decreasing rate. We can also ask what the marginal rate of substitution between the commodities will be. This is easily derived by taking the total derivative along an indifference curve and rearranging:

$$\pi u'(c_g)dc_g + (1 - \pi)u'(c_b)dc_b = 0, \quad \frac{dc_b}{dc_g} = -\frac{\pi u'(c_g)}{(1 - \pi)u'(c_b)}.$$

Note the fact that the MRS now depends not only on the marginal utility of wealth but also on the (subjective) probabilities the consumer assesses for each state! Even more importantly, we can consider what is known as the **certainty line**, that is, the locus of points where $c_g = c_b$. Since the marginal utility of consumption then is equal in both states (we have state independent utility here, after all, which means that the same $u(\cdot)$ applies in each state), it follows that the slope of an indifference curve on the certainty

⁶Note that this is somewhat more onerous than before now: imagine the states are indexed by good health and poor health. It is easy to imagine that an individual would evaluate material wealth differently in these two cases.

line only depends on the probability the consumer assesses for each state. In this case, it is $\pi/(1 - \pi)$.

The other ingredient is the budget line, of course. Since we have two commodities, each might be expected to have a price, and we denote these by p_g, p_b respectively. The consumer who has a total initial wealth of W may therefore consume any combination which lies on the budget line $p_g c_g + p_b c_b = W$, while a consumer who has an endowment of consumption given by (w_g, w_b) may consume anything on the budget line $p_g c_g + p_b c_b = p_g w_g + p_b w_b$. Where do these prices come from? As before, they will be determined by general equilibrium conditions. But if contingent markets are well developed and competitive, and there is general agreement on the likelihood of the states, then we might expect that a dollar's worth of consumption in a state will cost its expected value, which is just the dollar times the probability that it needs to be delivered. (I.e., a kind of zero profit condition for state pricing.) Thus we might expect that $p_g = \pi$ and $p_b = (1 - \pi)$.

The budget line also has a slope, of course, which is the rate at which consumption in one state can be transferred into consumption in the other state. Taking total derivatives of the budget we obtain that the budget slope is $dc_b/dc_g = p_g/p_b$. Combining this with our condition on "fair" pricing in the previous paragraph, we obtain that the budget allows transformation of consumption in one state to the other according to the odds.

4.4.1 Insurance in a State Model

So let us reconsider our consumer who was in need of insurance in this framework. In order to make this problem somewhat neater, we will reformulate the insurance premium into what is known as a **net premium**, which is a payment which only accrues in the case there is no loss. Since the normal insurance contract specifies that a premium be paid in either case, we usually have a payment of $premium \times Amount$ in order to obtain a net benefit of $Amount - premium \times Amount$. One dollar of consumption added in the state in which an accident occurs will therefore cost $premium/(1 - premium)$ dollars in the no accident state. Thus, let $p_b = 1$ and let $p_g = P$, the net premium. The consumer will then solve

$$\max_{c_b, c_g} \{ \pi u(c_g) + (1 - \pi)u(c_b) \} \quad \text{s.t.} \quad P c_g + c_b = P(100,000) + 80,000.$$

The two first order conditions for the consumption levels in this problem are

$$\pi u'(c_g) - \lambda P = 0 \quad \text{and} \quad (1 - \pi)u'(c_b) - \lambda = 0.$$




Figure 4.5: An Insurance Problem in State-Consumption space

Combining them in the usual way we obtain

$$\frac{\pi u'(c_g)}{(1 - \pi)u'(c_b)} = P.$$

Now, as we have just seen the LHS of this is the slope of an Indifference curve. The RHS is the slope of the budget, and so this says nothing but the familiar “there must be a tangency”.

We have also derived $P = \pi/(1 - \pi)$ for a fair net premium before. Thus we get that

$$\frac{\pi u'(c_g)}{(1 - \pi)u'(c_b)} = \frac{\pi}{1 - \pi},$$

which requires that

$$\frac{u'(c_g)}{u'(c_b)} = 1 \quad \Rightarrow \quad \frac{c_g}{c_b} = 1.$$

Thus this model shows us, just as the previous one, that a risk averse consumer faced with a fair premium will choose to fully insure, that is, choose to equalize consumption levels across the states. A diagrammatic representation of this can be found in diagram 3.5, which is standard for insurance problems. The consumer has an endowment which is off the certainty (45-degree) line. The fair premium defines a budget line along which the consumer can reallocate consumption from the good (no loss) state to the bad (loss) state. Optimum occurs where there is a tangency, which must occur on the certainty line since then the slopes are equalized. The picture looks perfectly “normal”, that is, just as we are used from introductory economics.

4.4.2 Risk Aversion Again

Given the amount of time spent previously on risk-aversion, it is interesting to see how risk-aversion manifests itself in this setting. Intuitively it might be apparent that a more risk averse consumer will have indifference curves which are more curved, that is, exhibit less substitutability (recall that a straight line indifference curve means that the goods are perfect substitutes, while a kinked Leontief indifference curve means perfect complements.) It therefore stands to reason that we might be interested in the rate at which the MRS is falling. It is, however, much easier to think along the lines of certainty equivalents: Consider two consumers with different risk aversion, that is, curvature of indifference curves. For simplicity, let us consider a point on the certainty line and the two indifference curves for our consumers through that common point (see Fig. 4.6).




Figure 4.6: Risk aversion in the State Model

Assume further that consumer B's indifference curve lies everywhere else above consumer A's. We can now ask how much consumption we have to add for each consumer in order to keep the consumer indifferent between the certain point and a consumption bundle with some given amount less in the bad state. Clearly, consumer B will need more compensation in order to accept the bad state reduction. Looked at it the other way around, this means that consumer B is willing to give up more consumption in the good state in order to increase bad state consumption. Note that both assess the same probabilities on the certainty line, since the slopes of their ICs are the same. How does this relate to certainty equivalents? Well, a budget line at fair odds will have the slope $-\frac{\pi}{1-\pi}$. Consider three such budget lines which are all parallel and go through the certain consumption point and the two

gambles which are equivalent for the consumer to the certain point. Clearly (from the picture) consumer B's budget is furthest out, followed by consumer A's, and furthest in is the budget through the certain point. But we know that parallel budgets differ only in the income/wealth they embody. Thus there is a larger reduction in wealth possible for B without reducing his welfare, compared to A. The wealth reduction embodied in the lower budget is the equivalent of the certainty equivalent idea before. (The expected value of a given gamble on such a budget line is given by the point on the certainty line and that budget, after all.)

4.5 Asset Pricing

Any discussion of models of uncertainty would be incomplete without some coverage of the main area in which all of this is used, which is the pricing of assets. As we have seen before, if there is only time to contend with but returns or future prices are known, then asset pricing reduces to a condition which says that the current price of an asset must relate to the future price through discounting. In the "real world" most assets do not have a future price which is known, or may otherwise have returns which are uncertain — stocks are a good example, where dividends are announced each year and their price certainly seems to fluctuate. Our discussion so far has focused on the avoidance of risk. Of course, even a risk averse consumer will accept some risk in exchange for a higher return, as we will see shortly. First, however, let us define two terms which often occur in the context of investments.

4.5.1 Diversification

Diversification refers to the idea that risk can be reduced by spreading one's investments across multiple assets. Contrary to popular misconceptions it is not necessary that their price movements be negatively correlated (although that certainly helps.) Let us consider these issues via a simple example.

Assume that there exists a project A which requires an investment of \$9,000 and which will either pay back \$12,000 or \$8,000, each with equal probability. The expected value of this project is therefore \$10,000. Now assume that a second project exists which is just like this one, but (and this is important) which is completely independent of the first. How much each pays back in no way depends on the other. Two investors now could each invest \$4,500 in each project. Each investor then has again a total

investment of \$9,000. How much do the projects pay back? Well, each will pay an investor either \$6,000 or \$4,000, each with equal probability. Thus an investor can receive either \$12,000, \$10,000, or \$8,000. \$12,000 or \$8,000 are received one quarter of the time, and half the time it is \$10,000. The total expected return thus is the same. BUT, there is less risk, since we know that for a risk-averse consumer $0.5u(12)+0.5u(8) < 0.25u(12)+0.25u(8)+0.5u(10)$ since $2(0.25u(12) + 0.25u(8)) < 2(0.5u(10))$.

Should the investor have access to investments which have negatively correlated returns (if one is up the other is down) risk may be able to be eliminated completely. All that is needed is to assume that the second project above will pay \$8,000 when the first pays \$12,000, and that it will pay \$12,000 if the first pays \$8,000. In that case an investor who invests half in each will obtain either \$6,000 and \$4,000 or \$4,000 and \$6,000: \$10,000 in either case. The expected return has not increased, but there is no risk at all now, a situation which a risk-averse consumer would clearly prefer.

4.5.2 Risk spreading

Risk spreading refers to the activity which lies at the root of insurance. Assume that there are 1000 individuals with wealth of \$35,000 and a 1% probability of suffering a \$10,000 loss. If the losses are independent of one another then there will be an average of 10 losses per period, for a total \$100,000 loss for all of them. The expected loss of each individual is \$100, so that all individuals have an expected wealth of \$34,900. A mutual insurance would now collect \$100 from each, and everybody would be reimbursed in full for their loss. Thus we can guarantee the consumers their expected wealth for certain.

Note that there is a new risk introduced now: in any given year more (or less) than 10 losses may occur. We can get rid of some of this by charging the \$100 in all years and retaining any money which was not collected in order to cover higher expenses in years in which more than 10 losses occur. However, there may be a string of bad luck which might threaten the solvency of the plan: but help is on the way! We could buy insurance for the insurance company, in effect insuring against the unlikely event that significantly more than the average number of losses occurs. This is called **re-insurance**. Since an insurance company has a well diversified portfolio of (independent) risks, the aggregate risk it faces itself is low and it will thus be able to get fairly cheap insurance.

These kind of considerations are also able to show why there may not be any insurance offered for certain losses. You may recall the lament on the radio about the fact that homeowners in the Red River basin were not able to purchase flood insurance. Similarly, you can't get earth-quake insurance in Vancouver, and certain other natural disasters (and man-made ones, such as wars) are excluded from coverage. Why? The answer lies in the fact that all insured individuals would have either a loss or no loss at the same time. That would mean that our mutual insurance above would either require no money (no losses) or \$10,000,000. But the latter requires each participant to pay \$10,000, in which case you might as well not insure! (A note aside: often the statement that no insurance is available is not literally correct: there may well be insurance available, but only at such high rates that nobody would buy it anyways. Even at low rates many people do not carry insurance, often hoping that the government will bail them out after the fact, a ploy which often works.)

4.5.3 Back to Asset Pricing

Before we look at a more general model of asset pricing, it may be useful to verify that a risk-averse consumer will indeed hold non-negative amounts of risky assets if they offer positive returns. To do so, let us assume the simple most case, that of a consumer with a given wealth w who has access to a risky asset which has a return of r_g or $r_b < 0 < r_g$. Let x denote the amount invested in the risky asset. Wealth then is a random variable and will be either $w_g = (w - x) + x(1 + r_g)$ or $w_b = (w - x) + x(1 + r_b)$. Suppose the good outcome occurs with probability π . What will be the choice of x ?

$$\max_{0 \leq x \leq w} \{ \pi u(w + r_g x) + (1 - \pi) u(w + r_b x) \}.$$

The first and second order conditions are

$$\begin{aligned} r_g \pi u'(w + r_g x) + r_b (1 - \pi) u'(w + r_b x) &= 0 \\ r_g^2 \pi u''(w + r_g x) + r_b^2 (1 - \pi) u''(w + r_b x) &< 0 \end{aligned}$$

The second order condition is satisfied trivially if the consumer is risk averse. To show under what circumstances it is not optimal to have a zero investment consider the FOC at $x = 0$:

$$r_g \pi u'(w) + r_b (1 - \pi) u'(w) \quad ? \quad 0.$$

The LHS is only positive if $\pi r_g + (1 - \pi) r_b > 0$, that is, if expected returns are positive. Notice also that in that case there will be some investment! Of

course this is driven by the fact that not investing guarantees a zero rate of return. Investing is a gamble, which the consumer dislikes, but also increases returns. Even a risk-averse consumer will take some risk for that higher return!

Now let us consider a more general model with many assets. Assume that there is a risk-free asset (one which yields a certain return) and many risky ones. Let the return for the risk-free asset be denoted by R_0 and the returns for the risky assets be denoted by \tilde{R}_i , each of which is a random variable with some distribution. Initial wealth of the consumer is w . Finally, we can let x_i denote the fraction of wealth allocated to asset $i = 0, \dots, n$. In the second period (we will ignore time discounting for simplicity and clarity) wealth will be a random variable the distribution of which depends on how much is invested in each asset. In particular, $\tilde{w} = w_0 \sum_{i=0}^n x_i \tilde{R}_i$, with the budget constraint that $\sum_{i=0}^n x_i = 1$. We can transform this expression as follows:

$$\tilde{w} = w \left[\left(1 - \sum_{i=1}^n x_i\right) R_0 + \sum_{i=1}^n x_i \tilde{R}_i \right] = w \left[R_0 + \sum_{i=1}^n x_i (\tilde{R}_i - R_0) \right].$$

The consumer's goal, of course, is to maximize expected utility from this wealth by choice of the investment fractions. That is,

$$\max_{\{x\}_i} \{ \mathbf{E}u(\tilde{w}) \} = \max_{\{x\}_i} \left\{ \mathbf{E}u \left(w \left[R_0 + \sum_{i=1}^n x_i (\tilde{R}_i - R_0) \right] \right) \right\}.$$

Differentiation yields the first order conditions

$$\mathbf{E}u'(\tilde{w})(\tilde{R}_i - R_0) = 0, \quad \forall i.$$

Now we will do some manipulation of this to make it look more presentable and informative. You may recall that the covariance of two random variables, X, Y , is defined as $\text{COV}(X, Y) = \mathbf{E}XY - \mathbf{E}X\mathbf{E}Y$. It follows that $\mathbf{E}u'(\tilde{w})\tilde{R}_i = \text{COV}(u', \tilde{R}_i) + \mathbf{E}u'(\tilde{w})\mathbf{E}\tilde{R}_i$. Using this fact and distributing the subtraction in the FOC across the equal sign, we obtain for each risky asset i the following equation:

$$\mathbf{E}u'(\tilde{w})R_0 = \mathbf{E}u'(\tilde{w})\mathbf{E}\tilde{R}_i + \text{COV}(u'(\tilde{w}), \tilde{R}_i).$$

From this it follows that in equilibrium the expected return of asset i must satisfy

$$\mathbf{E}\tilde{R}_i = R_0 - \frac{\text{COV}(u'(\tilde{w}), \tilde{R}_i)}{\mathbf{E}u'(\tilde{w})}.$$

This equation has a nice interpretation. The first term is clearly the risk-free rate of return. The second part therefore must be the risk-premium which the asset must garner in order to be held by the consumer in a utility maximizing portfolio. Note that if a return is positively correlated with wealth — that is, if an asset will return much if the consumer is already rich — then it is negatively correlated with the marginal utility of wealth, since that is decreasing in wealth. Thus the expected return of such an asset must exceed the risk free return if it is to be held. Of course, assets which pay off when wealth otherwise would be low can have a lower return than the risk-free rate since they, in a sense, provide insurance.

4.5.4 Mean-Variance Utility

The above pricing model required us to know the covariance and expectation of marginal utility, since, as we have seen before, it is the fact that marginal utility differs across outcomes which in some sense causes risk-aversion. A nice simplification of the model is possible if we specify at the outset that our consumer likes the mean but dislikes the variance of random returns, i.e., the mean is a good, the variance is a bad. We can then specify a utility function directly on those two characteristics of the distribution. (The normal distribution, for example is completely described by these two moments. If distributions differ in higher moments, this formulation would not be able to pick that up, however.)

Recall that for a set of outcomes (w_1, w_2, \dots, w_n) with probabilities $(\pi_1, \pi_2, \dots, \pi_n)$

$$\text{The mean is } \mu_w = \sum_{i=1}^n \pi_i w_i,$$

$$\text{and the variance is } \sigma_w^2 = \sum_{i=1}^n \pi_i (w_i - \mu_w)^2.$$

We now define utility directly on these: $u(\mu_w, \sigma_w)$, although it is standard practice to actually use the standard deviation as I just have done. Risk aversion is now expressed through the fact that we assume that

$$\frac{\partial u(\cdot)}{\partial \mu_w} = u_1(\cdot) > 0 \quad \text{while} \quad \frac{\partial u(\cdot)}{\partial \sigma_w} = u_2(\cdot) < 0.$$

We will now focus on two portfolios only (the validity of this approach will be shown in a while.) The risk free asset has a return of r_f , the risky

asset (the “market portfolio”) has a return of m_s with probability π_s . Let $r_m = \sum \pi_s m_s$ and $\sigma_m = \sqrt{\sum \pi_s (m_s - r_m)^2}$. Assume that a fraction x of wealth is to be invested in the risky asset (the market).

The expected return for a fraction x invested will be

$$r_x = \sum \pi_s (x m_s + (1 - x) r_f) = (1 - x) r_f + x \sum \pi_s m_s.$$

The variance of this portfolio is

$$\sigma_x^2 = \sum \pi_s (x m_s + (1 - x) r_f - r_x)^2 = \sum \pi_s (x m_s - x r_m)^2 = x^2 \sigma_m^2.$$

The investor/consumer will maximize utility by choice of x :

$$\begin{aligned} \max_x \{ & u(x r_m + (1 - x) r_f, x \sigma_m) \} \\ \text{FOC} \quad & u_1(\cdot) [r_m - r_f] + u_2(\cdot) \sigma_m = 0 \\ \text{SOC} \quad & u_{11}(\cdot) [r_m - r_f]^2 + u_{22}(\cdot) \sigma_m^2 + 2 \sigma_m [r_m - r_f] u_{12}(\cdot) \leq 0 \end{aligned}$$

Assuming that the second order condition holds, we note that we will require $[r_m - r_f] > 0$ since $u_2(\cdot)$ is negative by assumption. We may also note that we can rewrite the FOC as

$$\frac{-u_2(\cdot)}{u_1(\cdot)} = \frac{r_m - r_f}{\sigma_m}.$$

The LHS of this expression is the MRS between the mean and the standard deviation, that is, the slope of an indifference curve. The RHS can be seen to be the slope of the budget line since the budget is a mix of two points, $(r_f, 0)$ and (r_m, σ_m) , which implies that the tradeoff of mean for standard deviation is rise over run: $(r_m - r_f)/\sigma_m$.

In a general equilibrium everybody has access to the same market and the same risk free asset. Thus, everybody who does hold any of the market will have the same MRS — a result analogous to the fact that in our usual general equilibria everybody will have the same MRS. Of course, this is just a requirement of Pareto Optimality.

In this discussion we had but one risky asset. In reality there are many. As promised, we derive here the justification for considering only the so-called market portfolio. The basic idea is simple. Assume a set of risky assets. Since we are operating in a two dimensional space of mean versus standard deviation, one can certainly ask what combination of assets (also known as a portfolio) will yield the highest mean for a given standard deviation, or, which is often easier to compute, the lowest standard deviation for a given mean.

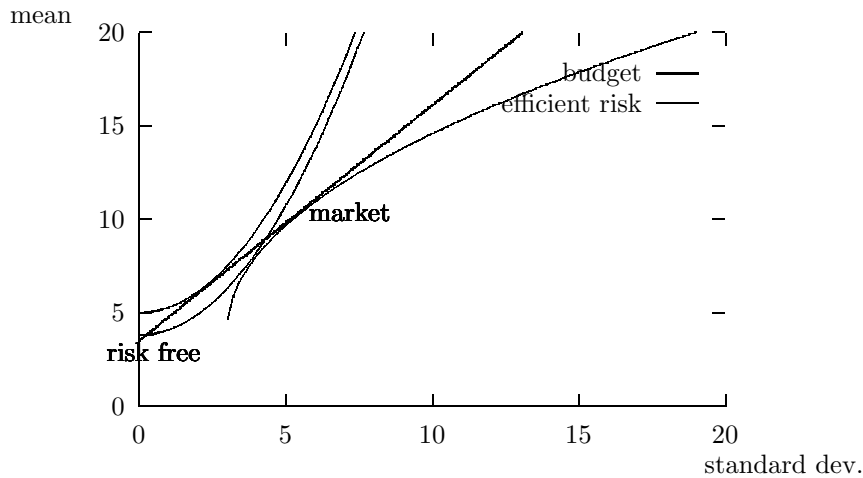


Figure 4.7: Efficient Portfolios and the Market Portfolio

Let x denote the vector of shares in each of the assets so that $\sum_I x = 1$. Define the mean and standard deviation of the portfolio x as $\mu(x)$ and $\sigma(x)$. Then it is possible to solve

$$\begin{aligned} \max_x \{ \mu(x) + \lambda(s - \sigma(x)) \} & \quad \text{or} \\ \min_x \{ \sigma(x) + \lambda(m - \mu(x)) \}. \end{aligned}$$

For each value of the constraint there will be a portfolio (an allocation of wealth across the different risky assets) which achieves the optimum. It turns out (for reasons which I do not want to get into here: take a finance course or do the math) that this will lead to a frontier which is concave to the standard variation axis.

Now, by derivation, any proportion of wealth which is held in risky assets will be held according to one of these portfolios. But which one? Well, the consumer can combine any one of these assets with the risk free asset in order to arrive at the final portfolio. Since this is just a linear combination, the resulting “budget line” will be a straight line and have a positive slope of $(\mu_x - \mu_f)/\sigma_x$, where μ_x, σ_x are drawn from the efficient frontier. A simple diagram suffices to convince us that the highest budget will be that which is just tangent to the efficient frontier. The point of tangency defines the market portfolio we used above!

Note a couple more things from this diagram. First of all it is clear from the indifference curves drawn that the consumer gains from the availability of risky assets on the one hand, and of the risk-free asset on the other. Second, the market portfolio with which the risk free asset will be mixed will depend on the return from the risk free asset. Imagine sliding the risk free return

up in the diagram. The budget line would have to become flatter and this means a tangency to the efficiency locus further to the right. Finally, the precise mix between market and risk free will depend on preferences. Indeed, there might be people who would want to have more than their entire wealth in the market portfolio, that is, the tangency to the budget would occur **to the right** of the market portfolio. This requires a “leveraged” portfolio in which the consumer is allowed to hold negative amounts of certain assets (short-selling.)

4.5.5 CAPM

The above shows that all investors face the same price of risk (in terms of the variance increase for a given increase in the mean.) It does not tell us anything about the risk-return trade-off for any given asset, however. Any risk unrelated to the market risk can be diversified away in this setting, however, so that any unsystematic risk will not attract any excess returns. An asset must, however, earn higher returns to the extent that it contributes to the market risk, since if it did not it would not be held in the market portfolio. Consideration of this problem in more detail (assume a portfolio with a small amount of this asset held and the rest in the market, compute the influence of the asset on the portfolio return and variance, rearrange) will yield the famous CAPM equation involving the asset’s ‘beta’ a number which is published in the financial press:

$$\mu_x = \mu_f + (\mu_m - \mu_f) \frac{\sigma_{X,M}}{\sigma_M^2}.$$

Here x denotes asset x , m the market, f the risk free asset, and $\sigma_{X,M}$ is the covariance between asset x and the market. The ratio $\frac{\sigma_{X,M}}{\sigma_M^2}$ is referred to as the asset’s **beta**.

4.6 Review Problems

Question 1: Demonstrate that all risk-averse consumers would prefer an investment yielding wealth levels 24, 20, 16 with equal probability to one with wealth levels 24 or 16 with equal probability.

Question 2: Compute the certainty equivalent for an expected utility maximizing consumer with (Bernoulli) utility function $u(w) = \sqrt{w}$ facing a gamble over \$3600 with probability α and \$6400 with probability $1 - \alpha$.

Question 3: Determine at what wealth levels a consumer with (Bernoulli) utility function $u(w) = \ln w$ has the same absolute risk aversion as a consumer with (Bernoulli) utility function $u(w) = 2\sqrt{w}$. How do their relative risk aversions compare at that level. What does that mean?

Question 4: Consider an environment with two states — call them rain, R , and shine, S , — with the probability of state R occurring known to be π . Assume that there exist two consumers who are both risk-averse, vN-M expected utility maximizers. Assume further that the endowment of consumer A is $(10, 5)$ — denoting 10 units of the consumption good in the case of state R and 5 units in the case of S — and that the endowment of consumer B is $(5, 10)$. What are the equilibrium allocation and price? (Provide a well labelled diagram and supporting arguments for any assertions you make.)

Question 5: Anna has \$10,000 to invest and wants to invest it all in one or both of the following assets: Asset G is gene-technology stock, while asset B is stock in a bible-printing business. There are only two states of nature to worry about, both of which occur with equal probability. One is that gene-technology is approved and flourishes, in which case the return of asset B is 0% while the return of asset G is 80%. The other is that religious fundamentalism takes hold, and gene-technology is severely restricted. In that case the return of asset B will be 40% but asset G has a return of (-40%). Anna is a risk-averse expected-utility maximizer with preferences over wealth represented by $u(w)$.

a) State Anna's choice problem mathematically.

b) What proportion of the \$10,000 will she invest in the gene-technology stock? (I.e. solve the above maximization problem.)

Question 6*: Assume that consumers can be described by the following preferences: they are risk averse over wealth levels, but they enjoy gambling for its consumption attributes (i.e., while they dislike the risk on wealth which gambling implies, they get some utility out of partaking in the excitement (say, they get utility out of daydreaming about what they could do if they won.)) Let us further assume that consumers only differ with respect to their initial wealth level and their consumption utility from gambling, but that all consumers have the same preferences over wealth. In order to simplify these preferences further, assume that wealth and gambling are separable. We can represent these preferences over wealth, w , and gambling, $g \in \{0, 1\}$ by some $u(w, g) = v(w) + \mu_i g$, where $v(w)$ is strictly concave, and μ_i is an individual parameter for each consumer. Finally, they are assumed to be expected utility maximizers.

a) Assume for now that the gambling is engaged in via a lottery, in

which consumers pay a fixed price p for one lottery ticket, and the ticket is either a “Try Again” (no prize) or “Winner” (Prize won.) Also assume for simplicity that they can either gamble once or not at all (i.e., each consumer can only buy one ticket.)

i) First verify that in such an environment the government can make positive profits. (I.e., verify that some consumers will buy a ticket even if the ticket price exceeds the expected value of the ticket.)

ii) If consumers have identical μ_i , how does the participation of consumers then depend on their initial wealth level if their preferences exhibit {constant| decreasing} {absolute |relative} risk aversion? (The above notation means: consider all (sensible) permutations.)

iii) Assume now that preferences are characterized by decreasing absolute and constant relative risk aversion. Verify that the utility functions $v(w) = \ln w$ and $v(w) = \sqrt{w}$ satisfy this assumption. Also assume that the consumption enjoyment of gambling is decreasing in wealth, that is, consumers with high initial wealth have low μ_i . (They know what pain it is to be rich and don't daydream as much about it.) Who would gamble then?

Question 7*: Assume that a worker only cares about the income he generates from working and the effort level he expends at work. Also assume that the worker is risk averse over income generated and dislikes effort. His preferences can be represented by $u(w, e) = \sqrt{w} - e^2$. The worker generates income by accepting a contract which specifies an effort level e and the associated wage rate $w(e)$. (Since leisure time does not enter in his preferences he will work full time (supply work inelastically) and we can normalize the wage rate to be per period income.) The effort level of the worker is not observed directly by the firm, and thus the worker has an incentive to expend as little effort as possible. The firm, however, cares about the effort expended, since it affects the marginal product it gets from employing the worker. It can conduct random tests of the worker's effort level with some probability α . These tests reveal the true effort level employed by the worker. If the worker is not tested, then it is assumed that he did indeed work at the specified effort level and will receive the contracted wage. If he is tested and found to have shirked (not supplied the correct effort level) then a penalty p is assessed and deducted from the wage of the worker.

a) What relationship has to hold between $w(e)$, e , α and p in order for a worker to provide the correct effort level if p is a constant (i.e., not dependent on either the contracted nor the observed effort level)?

b) If we can make p depend on the actual deviation in effort which we observe, what relationship has to be satisfied then?

c) Is there any economic insight hidden in this? Think about how these problems would change if the probability of detection somehow depended on

the effort level (i.e., the more I deviate from the correct effort level, the more likely I might be caught.)

Question 8: A consumer is risk averse. She is faced with an uncertain consumption in the future, since she faces the possibility of an accident. Accidents occur with probability π . If an accident occurs, it is either really bad, in which case she loses B , or minor, in which case her loss is $M < B$. Bad accidents represent only $1/5$ of all accidents, all other accidents are minor. Without the accident her consumption would be $W > B$.

a) Derive her optimal insurance purchases if the magnitude of the loss is publicly observable and verifiable and insurance companies make zero profits. More explicitly, if the type of accident is verifiable then a contract can be written contingent on the type of loss. The problem thus is equivalent to a problem where there are two types of accident, each of which occurs with a different probability and can be insured for separately at fair premiums.

b) Now consider the case if the amount of loss is private information. In this case only the fact that there was an accident can be verified (and hence contracted on), but the insurer cannot verify if the loss was minor or major, and hence pays only one fixed amount for any kind of accident. Assume zero profits for the insurer, as before, and show that the consumer now over-insures for the minor loss but under-insures for the bad loss, and that her utility thus is lowered. (Note that the informational distortion therefore leads to a welfare loss.)

Question 9: Assume a mean-variance utility model, and let μ denote the expected level of wealth, and σ its variance. Take the boundary of the efficient risky asset portfolios to be given by $\mu = \sqrt{\sigma - 16}$. Assume further that there exists a risk-free asset which has mean zero and standard deviation zero (if this bothers you, you can imagine that this is actually measuring the increase in wealth above current levels.) Let there be two consumers who have mean-variance utility given by $u(\mu, \sigma) = \mu - \frac{\sigma^2}{64}$ and $u(\mu, \sigma) = 3\mu - \frac{\sigma^2}{64}$ respectively. Derive their optimal portfolio choice and contrast their decisions.

Question 10: Fact 1: The asset pricing formula derived in class states that

$$\mathbf{E}\tilde{R}_i = R_0 - \frac{\mathbf{Cov}(U'(\tilde{w}), \tilde{R}_i)}{\mathbf{E}U'(\tilde{w})}.$$

Fact 2: A disability insurance contract can be viewed as an asset which pays some amount of money in the case the insured is unable to generate income from work.

Use Facts 1 and 2 above to explain why disability insurance can have

a negative return, (that is, why the price of the contract may exceed the expected value of the payment) if viewed as an asset in this way.

Question 11: TRUE/FALSE/UNCERTAIN: Provide justification for your answers via proof, counter-example or argument.

1) The optimal amount of insurance a risk-loving consumer purchases is characterized by the First Order Condition of his utility maximization problem, which indicates a tangency between an indifference curve and a budget line.

2) For a consumer who is initially a borrower, the utility level will definitely fall if interest rates increase substantially.

3) The utility functions

$$\frac{(3000 * \ln x_1 + 6000 * \ln x_2)}{12} + 2462 \quad \text{and} \quad \text{Exp}(x_1^{1/3} x_2^{2/3})$$

represent the same consumer preferences.

4) A risk-averse consumer will never pay more than the expected value of a gamble for the right to participate in the gamble. A risk-lover would, on the other hand.

5) The market rate of return is 15%. The stock of Gargleblaster Inc. is known to increase to \$117 next period, and is currently trading for \$90. This market (and the current stock price of Gargleblaster Inc.) is in equilibrium.

6) Under Risk-Variance utility functions, all consumers who actually hold both the risk-free asset and the risky asset will have the same Marginal Rate of Substitution between the mean and the variance, but may not have the same investment allocation.

Question 12*: Suppose workers have identical preferences over wealth only, which can be represented by the utility function $u(w) = 2\sqrt{w}$. Workers are also known to be expected utility maximizers. There are three kinds of jobs in the economy. One is a government desk job paying \$40,000.00 a year. This job has no risk of accidents associated with it. The second is a bus-driver. In this job there is a risk of accidents. The wage is \$44,100.00 and if there is an accident the monetary loss is \$11,700.00. Finally a worker could work on an oil rig. These jobs pay \$122,500.00 and have a 50% accident probability. These are all market wages, that is, all these jobs are actually performed in equilibrium.

a) What is the probability of accidents in the bus driver occupation?

b) What is the loss suffered by an oil rig worker if an accident occurs there?

c) Suppose now that the government institutes a workers' compensation scheme. This is essentially an insurance scheme where each job pays a fair premium for its accident risk. Suppose that workers can buy this insurance

in arbitrary amounts at these premiums. What will the new equilibrium wages for bus-drivers and oil rig workers be? Who gains from the workers' compensation?

d) Now suppose instead that the government decides to charge only one premium for everybody. Suppose that of the workers in risky jobs 40% are oil rig workers and 60% are bus drivers. Suppose that they can buy as much or little insurance as they wish. How much insurance do the two groups buy? Who is better off, who is worse off in this case (at the old wages)? Can we say what the new equilibrium wages would have to be?

Question 13*: Prove that state-independent expected utility is homothetic if the consumer exhibits constant relative risk aversion. (This question arises since indifference curves do have the same slope along the certainty line. So could they have the same slope along any ray from the origin? In that case they would be homothetic.)

Chapter 5

Information

In the standard model of consumer choice discussed in chapter 1, as well as the model of uncertainty developed in chapter 3, it was assumed that the decision maker knows all relevant information. In chapter 1 this meant that the consumer knows the price of all goods, as well as the precise features of each good (all characteristics relevant to the consumer.) In chapter 3 this in particular implied that the consumer has information about the probabilities of states or outcomes. Not only that, this information is symmetric, so that all parties to a transaction have the same information. Hence in chapter 3 the explicit assumption that the insurance provider has the same knowledge of the probabilities as the consumer.

What if these assumptions fail? What if there is no complete and symmetric information? Fundamentally, one of the key problems is asymmetric information — when one party to a transaction knows something relevant to the transaction which the other party does not know. This quite clearly will lead to problems, since Pareto efficiency necessitates that all available information is properly incorporated. Consider, for example, the famous “Lemon’s Problem”:¹ Suppose a seller knows the quality of her used car, which is either high or low. The seller attaches values of \$5000 or \$1000 to the two types of car, respectively. Buyers do not know the quality of a used car and have no way to determine it before purchase. Buyers value good used cars at \$6000 and bad used cars at \$2000. Note that it is Pareto efficient in either case for the car to be sold. Will the market mechanism work in this case? Suppose that it is known by everybody that half of all cars are good, and half are bad. To keep it simple, suppose buyers and sellers are risk

¹This kind of example is due to Akerlof (1970) *Quarterly Journal of Economics*, a paper which has changed economics.

neutral. Buyers would then be willing to pay at most \$4000 for a gamble with even odds in which they either receive a good or a bad used car. Now suppose that this were the market price for used cars. At this price only bad cars are actually offered for sale, since good car buyers rather keep theirs, since their valuation of their own car exceeds the price. It follows that this price cannot be an equilibrium price. It is fairly easy to verify that the only equilibrium would have bad cars trade at a price somewhere between \$1000 and \$2000 while good cars are not traded at all. This is not Pareto efficient.

Of course, it is not necessary that there be asymmetric information. Suppose that there exist many restaurants, each offering slightly different combinations of food and ambiance. Consumers have tastes over the characteristics of the meals (how spicy, what kind of meat, if any meat at all; Italian, French, eastern European, Japanese, Egyptian, etc.) as well as the kind of restaurant (formal, romantic, authentic, etc.) as well as the general quality of the cooking within each category. In a Pareto efficient general equilibrium each consumer must frequent (subject to capacity constraints) the most preferred restaurant, or if that is full the next preferred one. Can we expect this to be the equilibrium?² To see why the general answer may be “**No**” consider a risk averse consumer in an environment where a dining experience is necessary in order to find out all relevant characteristics of a restaurant. Every visit to a new place carries with it the risk of a really unpleasant experience. If the expected benefit of finding a better place than the one currently known does not exceed the expected cost of having a bad experience, the consumer will not try a new restaurant, and hence will not find the best match!

What seems to be important then, are two aspects of information: One, can all relevant information be acquired before the transactions is completed?; Two, is the information symmetric or asymmetric?

Aside from a classification of problems into asymmetric or symmetric information, it is common to distinguish between three classes of goods, based on the kind of informational problems they present: search goods, experience goods, and (less frequently) faith goods. A **search good** is one for which the consumer is lacking some information, be it price or some attribute of the good, which can be fully determined before purchase of the good. Anytime

²This is a question similar to one very important in labour economics: are workers matched to the correct jobs, that is, are the characteristics of workers properly matched to the required characteristics of the job? These kind of problems are analysed in the large **matching** literature. In this literature you will find interesting papers on stable marriages — is everybody married to their most preferred partner, or could a small coalition swap partners and increase their welfare?

the consumer can discover all relevant aspects before purchasing we speak of a search good. Supposing that search is costly, which seems reasonable, we can then model the optimal search behaviour of consumers by considering the (marginal) benefits and (marginal) costs of search. Applying such thinking to labour markets we can study the efficiency effects of unemployment insurance; or we can apply it to advertising, which for such goods focuses on supplying the missing information, and is therefore possibly efficiency enhancing.

For some goods such a determination of all relevant characteristics may not be possible. Nobody can explain to you how something tastes, for example; you will have to consume the good to find out. Similarly for issues of quality. Inasmuch as this refers to how long a durable good lasts, this can only be determined by consuming the good and seeing when (and if) it breaks. Such goods are called **experience goods**. The consumer needs to purchase the good, but the purchase and consumption of the good (or service!) will fully inform the consumer. This situation naturally leads to consumers trying a good, but maybe not necessarily finding the best match for them. Advertising will be designed to make the consumer try the product — free samples could be used.³ Why would the consumer not necessarily find the best product? If there is a cost to unsatisfactory consumption experiences this will naturally arise. As in the restaurant example above. Similar examples can be constructed for hair cuts, and many other services.

What if the consumer never finds out if the product performs its function? This is the natural situation for many medical and religious services. The consumer will not discover the full implications until it is (presumably) too late. Such goods are termed **faith goods** and present tremendous problems to markets. In our current society the spiritual health of consumers is not judged to be important, and so the market failure for religious services is not addressed.⁴ Health, in contrast, is judged important — since consumers value it, for one, and since there are large externalities in a society with a social safety net — and thus there is extensive regulation for health care services in most societies.⁵ Since education also has certain attributes of a

³It used to be legal for cigarette manufacturers to distribute “sample packs” for free, allowing consumers to experience the good without cost. The fact that nicotine is addictive to some is only helpful in this regard, as any local drug pusher knows: they also hand out free sample packs in schools.

⁴A convincing argument can be made that societies which prescribe one state religion do so not in an attempt to maximize consumer welfare but tax revenue and control. The Inquisition, for example, probably had little to do with a concern for the welfare of heretics. Note also that I am speaking especially with respect to religions in which the “afterlife” plays a large role. We will encounter them again in the chapter on game theory.

⁵Interesting issues arise when the provision of health care services is combined with the

faith good — in the sense that it is very costly to unlearn, or to repeat the education — we also see strong regulation of education in most societies.⁶

Note that religion and health care probably differ in another dimension: in health care there is information asymmetry between provider and consumer; it may be argued that in religion both parties are equally uninformed.⁷ Presumably the fact that information is asymmetric makes it easier to exploit the uninformed party on the one hand, and to regulate on the other.⁸ The government attempts to combat this informational asymmetry by certifying the supply.

Aside from all the above, there are additional problems with information. These days everybody speaks of the “information economy”. Clearly information is supposed to have some kind of value and therefore should have a price. However, while that is clear conceptually, it is far from easy to incorporate information into our models. Information is not a good like most others. For one, it is hard to measure. There are, of course, measurements which have been derived in communications theory — but they often measure the amount of meaningful signals versus some noise (as in the transmission of messages.) These measurements measure if messages have been sent, and how many. Economists, in contrast, are concerned with what kind of message actually contains relevant information and what kind may be vacuous. Much care is therefore taken to define the informational context of any given decision problem (we will encounter this again in the game theory part, where we will use the term **information set** to denote all situations in which the same information has been gathered, loosely speaking.)

Aside from the problem of defining and measuring information, it is also a special good since it is **not rivalrous** (a concept you may have encountered in ECON 301): the fact that I possess some information in no way impedes your ability to have the same information. Furthermore, the fact that I “consume” the information somehow (let’s say by acting on it) does not stop you from consuming it or me from using it again later. There are therefore

provision of spiritual services.

⁶This is regulation for economically justifiable reasons. Because education is so costly to undo or repeat it is also frequently meddled with for “societal engineering” reasons.

⁷I am not trying to belittle beliefs here, just pointing to the fact that while a medical doctor may actually know if a prescribed treatment works (while the patient does not), neither the religious official or the follower of a faith know if the actions prescribed by the faith will “work”.

⁸Note the weight loss and aphrodisiac markets, or cosmetics, for example. Little difference between these and the “snake oil cures” of the past seems to exist. Regulation can take the simple form that only “verifiable” claims may be made.

large public good aspects to information which require special consideration.⁹ The long and short of this is that standard demand-supply models often don't work, and that markets will in general misallocate if information is involved, which makes it even more important to have a good working model. A complete study of this topic is outside the scope of these notes, however. In what follows we will only outline some specific issues in further detail.

5.1 Search

One of the key breakthroughs in the economics of information was a simple model of information acquisition. The basic idea is a simple one (as all good ideas are.) A consumer lacks information — say about the price at which a good may be bought. Clearly it is in the consumer's interest not to be “fleeced,” which requires him to have some idea about what the market price is. In general the consumer will not know what the lowest price for the product is, but can go to different stores and find out what their price is. The more stores the consumer visits the better his idea about what the correct price might be — the better his information — but of course the higher his cost, since he has to visit all these stores. An optimizing consumer may be expected to continue to find new (hopefully) lower prices as long as the marginal benefit of doing so exceeds the marginal cost of doing so. Therefore we “just” need to define benefits and costs and can then apply our standard answer that the optimum is achieved if the marginal cost equals the marginal benefit!

The problem with this is the fact that we will have to get into sampling distributions and other such details (most of which will not concern us here) to do this right. The reason for this is that the best way to model this kind of problem is as the consumer purchasing messages (normally: prices) which are drawn from some distribution. For example: let us say that the consumer knows that prices are distributed according to some cumulative distribution function $F(z) = \int_0^z f(p)dp$, where $f(p)$ is the (known) density. If the consumer were to obtain n price samples (draws) from this distribution

⁹This is the problem in the Patent protection fight: Once somebody has invented (created the information for) a new drug the knowledge should be freely available — but this ignores the general equilibrium question of where the information came from. The incentives for its creation will depend on what property rights are enforceable later on. For while two firms both may use the information, a firm may only really profit from it if it is the sole user.

then the probability that any given sample (say p_0) is the lowest will be

$$[1 - F(p_0)]^{n-1} f(p_0).$$

(This formula should make sense to you from Econometrics or Statistics: we are dealing with n independent random variables.) From this it follows that the expected value of the lowest price after having taken a sample of n prices is

$$p_{low}^n = \int_0^{\infty} p[1 - F(p)]^{n-1} f(p) dp.$$

Note that this expected lowest sample decreases as n increases, but at a decreasing rate: the difference between sampling n times and $n - 1$ times is

$$p_{low}^n - p_{low}^{n-1} = - \int_0^{\infty} pF(p)[1 - F(p)]^{n-2} f(p) dp < 0.$$

So additional sampling leads to a benefit (lower expected price) but with a diminishing margin.

What about cost? Even with constant (marginal) cost of sampling we would have a well defined problem and it is easy to see that individuals with higher search costs will have lower sample sizes and thus pay higher expected prices. Also note that the lowest price paid is still a random variable, and hence consumers do not buy at the same price (which is inefficient, in general!) Computing the variance you would observe that dispersion of the lowest price is decreasing in n — that means that the lower the search costs the ‘better’ the market can be expected to work. Indeed, competition policy is concerned with this fact in some places and attempts to generate rules which require that prices be posted (lowering the cost of search).

Any discussion of search would be lacking if we did not point out that search is normally sequential. In the above approach n messages were bought, with n predetermined. This is the equivalent of visiting n stores for a quote irrespective of what the received quotes are. The dynamic problem is the much more interesting one and has been quite well studied. We will attempt to distill the most important point and demonstrate it by way of a fairly simple example. The key insight into this kind of problem is that it is often optimal to employ a **stopping rule**.¹⁰ That is, to continue sampling until a price has been obtained which is below some preset limit, at which point search is abandoned and the transaction occurs (a sort of “good enough” attitude.) The price at which one decides to stop is the **reservation**

¹⁰It does depend on the distributional assumptions we make on the samples — independence makes what follows true.

price — the highest price one is willing to pay! In order to derive this price we will have to specify if one can ‘go back’ to previous prices, or if the trade will have fallen through if one walks away. The latter is the typical setup in the matching literature in labour economics, the former is a bit easier and we will consider it first.

Suppose the cost of another sample (search effort) is c . The outcome of the sample is a new price p , which will only be of benefit if it is lower than the currently known minimum price. Evaluated at the optimal reservation price p_R , the expected gain from an additional sample is therefore the savings $p_R - p$, “expected over” all prices $p < p_R$. If these expected savings are equal to the cost of the additional sample, then the consumer is just indifferent between buying another sample or not, and thus the reservation price is found:

$$p_R \text{ satisfies } \int_0^{p_R} (p_R - p)f(p)dp = c.$$

Next, consider the labour market, where unemployed workers are searching for jobs. This is the slightly more complex case where the consumer cannot return to a past offer. Also, the objective is to find the highest price. First determine the value of search, V , which is composed of the cost, the expected gain above the wage currently on the table, and the fact that a lower wage might arise which would indicate that another search is needed. Thus, assuming a linear utility function for simplicity (no risk aversion), and letting p stand for wages

$$V = -c + \int_{p_R}^{\infty} pf(p)dp + V \int_0^{p_R} f(p)dp \implies V = \frac{-c + \int_{p_R}^{\infty} pf(p)dp}{1 - F(p_R)}.$$

Note that I assume stationarity here and the fact that one p_R will do (i.e., the fact that the reservation wage is independent of how long the search has been going on.) All of these things ought to be shown for a formal model. We now will ask what choice of p_R will maximize the value above (which is the expected utility from searching.) Taking a derivative we get

$$\frac{-p_R f(p_R)(1 - F(p_R)) + f(p_R)(-c + \int_{p_R}^{\infty} pf(p)dp)}{(1 - F(p_R))^2} = 0.$$

Simplifying and rearranging we obtain

$$\int_{p_R}^{\infty} pf(p)dp - p_R = c - p_R F(p_R).$$

The LHS of this expression is the expected increase in the wage, the RHS is the cost of search, which consists of the actual cost of search and the fact

that the current wage is foregone if a lower value is obtained. What now will be the effect of a decrease in the search cost c , for example if the government decides to subsidize non-working (searching) workers? This would lower c , and the LHS would have to fall to compensate, which will occur only if a higher p_R is chosen. Of course, a higher p_R lowers the RHS further. In mathematical terms it is easy to compute that

$$\frac{dp_R}{dc} = \frac{-1}{1 - F(p_R)}.$$

Unemployment insurance will increase the reservation wage (and thus unemployment — note the pun: it ensures unemployment!) The reason is that our workers can be more choosy and search for the “right” job. They become more discriminating in their job search. Note that this is **not** necessarily bad. If the quality of the match (suitability of worker and firm with each other) is reflected in a higher wage, then this leads to better matches (fewer Ph.D. cab drivers). This may well be desirable for the general equilibrium efficiency properties of the model.

Now, this model is quite simplistic. More advanced models might take into account eligibility rules. In those models unemployment insurance can be shown to cause some workers to take bad jobs (because that way they can qualify for more insurance later.) Similar models can also be used to analyse other matching markets. The market for medical interns comes to mind, or the marriage market.

In closing let us note that certain forms of advertising will lower search costs (since consumers now can determine cheaply who has what for sale at which price) and thus are efficiency enhancing (less search, less resources spent on search, and lower price dispersion in the market.) Other forms of advertising (image advertising) do not have this function, however, and will have to be looked at in a different framework. This is where the distinction between search goods and experience goods comes in.

5.2 Adverse Selection

Most problems which will concern us in this course are actually of a different nature than the search for information above. What we are interested in most are the problems which arise because information is **asymmetric**. This means that two parties to a transaction do not have the same information, as is the case if the seller knows more about the quality (or lack thereof) of




Figure 5.1: Insurance for Two types

his good than the buyer, or if the buyer knows more about the value of the good to himself than the seller. In these types of environments we run into two well known problems, that of **adverse selection** (which you may think of as “hidden property”) and **moral hazard** (“hidden action.”) We will deal with the former in this section.

Adverse selection lies at the heart of Akerlof’s Lemons Problem. These kind of markets lead naturally to the question if the informed party can send some sort of **signal** to reveal information. But how could they do so? Simply stating the fact that, say, the car is of good quality will not do, since such a statement is free and would also be made by sellers of bad cars (it is not **incentive compatible**.) Sometimes this problem can be fixed via the provision of a warranty, since a warranty makes the statement that the car is good more costly to sellers of bad cars than of cars which are, in fact, good.

Let us examine these issues in our insurance model. Assume two states, good and bad, and assume two individuals who both have wealth endowment $(w_g, w_b); w_g > w_b$. Suppose that these individuals are indistinguishable to the insurance company, but that one individual is a good risk type who has a probability of the good state occurring of π_H , while the other is a bad risk type with probability of the good state of only $\pi_L < \pi_H$. To be stereotypical and simplify the presentation below, assume that the good risk is female, the bad risk male, so that grammatical pronouns can distinguish types in what follows. As we have seen before, the individuals’ indifference curves will have a slope of $-\pi_i/(1 - \pi_i)$ on the certainty line.

If there were full information both types could buy insurance at fair premiums and would choose to be on the certainty line on their respective budget lines (with the high risk type on a lower budget line and with a lower consumption in each state.) However, with asymmetric information this is not a possible outcome. The bad risk cannot be distinguished from the good risk *a priori* and therefore can buy at the good risk premium. Now, if he were to maximize at this premium we know that he would over insure — and this action would then distinguish him from the good risk. The best he can do without giving himself away is to buy the same insurance coverage that our good risk would buy, in other words to mimic her. Thus both would attempt to buy full insurance for a good type.

We now have to ask if this is a possible equilibrium outcome. The answer is NO, since the insurance company now would make zero (expected) profits on her insurance contract, but would lose money on his insurance contract. Consider the profit function (and recall that $p_i = \pi_i$):

$$(1 - \pi_H)(w_g - w_b) - (1 - \pi_L)(w_g - w_b) < 0.$$

Foreseeing this fact, the insurance company would refuse to sell insurance at these prices. Well then, what is the equilibrium in such a market? There seem to be two options: either both types buy the same insurance (this is called **pooling** behaviour) or they buy different contracts (this is called **separating** behaviour.)

Does there exist a pooling equilibrium in this market? Consider a larger market with many individuals of each of the two types. Let f_H denote the proportion of the good types (π_H) in the market. An insurer would be making zero expected profits from selling a common policy for coverage of I at a premium of p to all types if and only if

$$f_H(\pi_H p I - (1 - \pi_H)(1 - p)I) + (1 - f_H)(\pi_L p I - (1 - \pi_L)(1 - p)I) = 0.$$

This requires a premium

$$p = (1 - \pi_L) - f_H(\pi_H - \pi_L).$$

Note that at this premium the good types subsidize the bad types, since the former pay too high a premium, the latter too low a premium. In Figure 5.2, this premium is indicated by a zero profit locus (identical to the consumers' budget) at an intermediate slope (labelled 'market'.) Any proposed equilibrium with pooling would have to lie on this line and be better than no insurance for both types. Such a point might be point M in the figure. However, this cannot be an equilibrium. In order for it to be an equilibrium

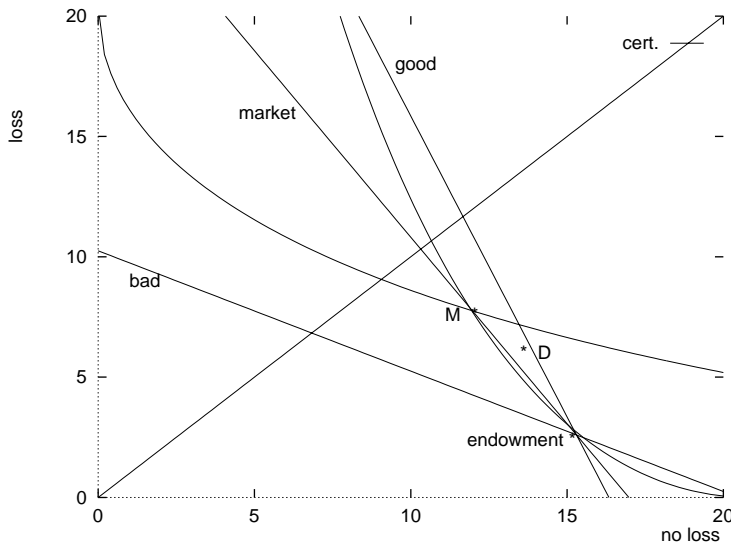


Figure 5.2: Impossibility of Pooling Equilibria

nobody must have any actions available which they could carry out and prefer to what is proposed. Consider, then, any point in the area to the right of the zero profit line, below the bad type's indifference curve, and above the good type's indifference curve: a point such as D . This contract, if proposed by an insurance company, offers less insurance but at a better price. Only the good type would be willing to take it (she would end up on a higher indifference curve). Since it is not all the way over on the good type zero profit line, the insurance company would make strictly positive profits. Of course, all bad types would remain at the old contract M , and since this is above the zero profit line for bad types whoever sold them this insurance would make losses. Notice that the same arguments hold whatever the initial point. It follows that a pooling equilibrium cannot exist.

Well then, does a separating equilibrium exist? We now would need two contracts, one of which is taken by all bad types and the other by all good types. Insurance offerers would have to make zero expected profits. Of course, since each type takes a different contract the insurer will know who is who from their behaviour. This suggests that we look at contracts which insure the bad risks fully at fair prices. For the bad risk types to accept this type of contract the contract offered to the good risk type must be on a lower indifference curve. It must also be on the fair odds line for good types for there to be zero profits. Finally it must be acceptable to the good types. This suggests a point such as A in the Figure 5.3. There now are two potential problems with this. One is that the insurer and insured of good type would like to move to a different point, such as B , after the type




Figure 5.3: Separating Contracts could be possible

is revealed. But if that would indeed happen then it can't, in equilibrium, because the bad types would foresee it and pretend to be good in order then to be mistaken for good. The other problem is that the stability of such separating equilibria depends on the mix of types in the market. If, as in the Figure 5.3, the market has a lot of bad types this kind of equilibrium works. But what if there are only a few bad types? In that case an insurer could deviate from our proposed contract and offer a pooling contract which makes strictly positive profits and is accepted by all in favour over the separating contract. This is point C in Figure 5.4. Of course, while this deviation destroys our proposed equilibrium it is itself not an equilibrium (we already know that no pooling contract can be an equilibrium.) This shows that a small proportion of bad risks who can't be identified can destroy the market completely!

Notice the implications of these findings on life or health insurance markets when there are people with terminal diseases such as AIDS or Hepatitis C, or various forms of cancer. The patient may well know that he has this problem, but the insurance company sure does not. Of course, it could require testing in order to determine the risk it will be exposed to. Our model shows that doing so would make sure that everybody has access to insurance at fair rates and that there is no cross-subsidization. This is clearly efficient. However, it will reduce the welfare of consumers with these diseases. Indeed, given that AIDS, for example, means a very high probability of seriously expensive treatment, the insurance premiums would be very high (justifiably so, by the way, since the risk is high.) What ought society do




Figure 5.4: Separating Contracts definitely impossible

about this? Economics is not able to answer that question, but it is able to show the implications of various policies. For example, banning testing and forcing insurance companies to bear the risk of insuring such customers might make the market collapse, in the worst scenario, or will at least lead to inefficiency for the majority of consumers. It would also most likely tend to lead to “alternative” screening devices — the insurance company might start estimating consumers’ “lifestyle choices” and then discriminate based on that information. Is that an improvement? An alternative would be to let insurance operate at fair rates but to directly subsidize the income of the affected minority. (This may cause “moral” objections by certain segments of the population.)

5.3 Moral Hazard

Another problem which can arise in insurance markets, and indeed in any situation with asymmetric information, is that of hidden actions being taken. In our insurance example it is often possible for the insured to influence the probability of a loss: is the car locked? Are there anti-theft devices installed? Are there theft and fire alarms, sprinklers? Do the tires have enough tread depth, and are they the correct tire for the season (in the summer a dedicated summer tire provides superior breaking and road holding to a M&S tire, while in the winter a proper winter tire is superior.) Some of these actions are observable and will therefore be included in the conditions of the insurance

contract. For example, in Germany your insurance will refuse to pay if you do not have “approved” tires, or if the car was not locked and this can be established. Indeed, insurance for cars without anti-theft devices is now basically unaffordable since theft rates are so high since the ‘iron curtain’ opened.¹¹ If you insure your car as parked in a closed parking garage and then leave it out over night your insurance may similarly refuse to pay! Other actions are not so easily verified, however. How aggressively do you drive? How many risks do you decide to take on a given day? This is often not observable but nevertheless in your control. If the avoidance of accidents is costly to the insured in any way, then he can be expected to pick the optimal level of the (unobservable) action — optimal for himself, not the insurance company, that is.

As a benchmark, let us consider an individual who faces the risk of loss L . The probability of the loss occurring depends on the amount of preventive action, A , taken and is denoted $\pi(A)$. It would appear logical to assume that $\pi'(A) < 0, \pi''(A) > 0$. The activity costs money, and cost is $c(A)$ with $c'(A) > 0, c''(A) > 0$. In the absence of insurance the individual would choose A so as to

$$\max_A \{ \pi(A)u(w - L - c(A)) + (1 - \pi(A))u(w - c(A)) \}.$$

The first order condition for this problem is

$$\begin{aligned} \pi'(A)u(w - L - c(A)) - c'(A)\pi(A)u'(w - L - c(A)) &- \\ \pi'(A)u(w - c(A)) - c'(A)(1 - \pi(A))u'(w - c(A)) &= 0. \end{aligned}$$

Thus the optimal A^* satisfies

$$c'(A^*) = \frac{\pi'(A^*)(u(w - L - c(A^*)) - u(w - c(A^*)))}{\pi(A^*)u'(w - L - c(A^*)) - (1 - \pi(A^*))u'(w - c(A^*))}.$$

Now consider the consumer’s choice when insurance is available. To keep it simple we will assume that the only available contract has a premium of $\pi(A^*)$ and is for the total loss L . Note that this would be a fair premium if the consumer continues to choose the level A^* of abatement activity. However, the maximization problem the consumer now faces becomes (the consumer will assume the premium fixed)

$$\max_A \{ \pi(A)u(w - pL - c(A)) + (1 - \pi(A))u(w - pL - c(A)) \}.$$

The first order condition for this problem is

$$\pi'(A)u(\cdot) - c'(A)\pi(A)u'(\cdot) - \pi'(A)u(\cdot) - c'(A)(1 - \pi(A))u'(\cdot) = 0.$$

¹¹Yes, I am suggesting that most stolen cars end up in former eastern block countries — it is a fact.

Thus the optimal A^* satisfies

$$c'(\hat{A}) = 0.$$

Clearly (and expectedly) the consumer will now not engage in the activity at all. It follows that the probability of a loss is higher and that the insurance company would lose money. If we were to recompute the level of A at a fair premium for $\pi(0)$ we would find that the same is true: no effort is taken. The only way to elicit effort is to expose the consumer to the right incentives: there must be a reward for the hidden action. A deductible will accomplish this to some extent and will lead to at least some effort being taken. This is due to the fact that the consumer continues to face risk (which is costly.) The amount of effort will in general not be the optimal, however.

5.4 The Principal Agent Problem

To conclude the chapter on information here is an outline of the leading paradigm for analysing asymmetric information problems. Most times one (uninformed) party wishes to influence the behavior of another (informed) party, the problem can be considered a principal-agent problem. This kind of problem is so frequent that it deserved its own name and there are books which nearly exclusively deal with it. We will obviously have only a much shorter introduction to the key issues in what follows.

First note that this is a frequent problem in economics, to say the least. Second, note that at the root of the problem lies asymmetric information. We are, generically, concerned with situations in which one party — the principal — has to rely on some other party — the agent — to do something which affects the payoff of the first party. This in itself is not the problem, of course. What makes it interesting is that the principal cannot tell if the agent did what he was asked to do, that is, there is a lack of information, and that the agent has incentives not to do as asked. In short, there is a moral-hazard problem. The principal-agent literature explores this problem in detail.

5.4.1 The Abstract P-A Relationship

We will frame this analysis in its usual setting, which is that of a risk-neutral principal who tries to maximize (expected) profits, and one agent, who will

have control over (unobservable) actions which influence profits. Call the principal the owner and the agent the manager. The agent/manager will be assumed to have choice of a single item, what we shall call his effort level $e \in [\underline{e}, \bar{e}]$. This effort level will be assumed to influence the level of profits before manager compensation, called gross-profits, which you may think of as revenue if there are no other inputs. More general models could have multiple dimensions at this stage. The manager's preferences will be over monetary reward (wages/income) and effort: $U(w, e)$. We assume that the manager is an expected utility maximizer. Also assume that the manager is risk-averse (that is $u_1(\cdot, \cdot) > 0, u_{11}(\cdot, \cdot) < 0$) and that effort is a bad ($u_2(\cdot, \cdot) < 0$.)

The owner can not observe the manager's effort choice. Instead, only the realization of gross profits is observed. Note at this stage that this in itself does not yet cause a moral-hazard problem, since if the relationship between effort and profit is known we can simply invert the profit function to deduce the effort level. Therefore we need to introduce some randomness into the model. Let ρ be a random variable which also influences profits and which is not directly observable by the owner either. It could be known to the manager, but we will have to assume that it will become known only after the fact, so that the manager cannot condition his effort choice on the realization of ρ . Let the relationship between effort, profits and the random variable be denoted by $\Pi(e, \rho)$. Note that $\Pi(e, \rho)$ will be a random variable itself. All expectations in what follows will be taken with respect to the distribution of ρ . $E\Pi(e, \rho)$, for example, will be the expected profits for effort level e .

Since the owner can only observe profits, the most general wage contract he can offer the manager will be a function $w(\Pi)$. This formulation includes a fixed wage as well as all wage plus bonus schemes, or share contracts (where the manager gets a fixed share of profits.)

Let us first, as a benchmark, determine the efficient level of effort, which would be provided under full information. In that case we would have to solve the Pareto problem, that is, solve

$$\max_{e,w} \{E\Pi(e, \rho) - w \quad \text{s.t.} \quad U(w, e) = \bar{u}\}.$$

Here \bar{u} is a level of utility sufficient to make the manager accept the contract. Note also that in formulating this problem like this I have already used the knowledge that a risk-neutral owner should fully insure a risk-averse manager by offering a constant wage.

Assuming no corner solutions, it is easy to see that we would want to set the effort level such that the marginal benefit of effort (in terms of higher expected profits) is related to the marginal cost (in terms of higher disutility

of effort, which will have to be compensated via the wages). In particular, we need that

$$\mathbf{E}\Pi_e(\cdot, \cdot) = \frac{w}{U_w(\cdot, \cdot)} U_e(\cdot, \cdot).$$

What if effort is not observable? In that case we will solve the model “backwards”: given any particular wage contract faced by the manager, we will determine the manager’s choice. Then we will solve for the optimal contract, “foreseeing” those choices. Assume in what follows that the owner wants to support some effort level \hat{e} (which is derived from this process.)

So, our manager is faced with two decisions. One is to determine how much effort to provide given he accepts the contract:

$$\max_e \{\mathbf{E}U(w(\Pi(e, \rho)), e)\}.$$

This leads to FOC

$$\mathbf{E}[U_w(\cdot, \cdot)w'(\cdot)\Pi_e(\cdot, \cdot) + U_e(\cdot, \cdot)] = 0,$$

which determines the optimal e^* . The other is the question if to accept the contract at all, which requires that

$$\max_e \{\mathbf{E}U(w(\Pi(e, \rho)), e)\} = \mathbf{E}U(w(\Pi(e^*, \rho)), e^*) \geq U_0.$$

Here U_0 is the level of utility the manager can obtain if he does not accept the contract but instead engages in the next best alternative activity (in other words, it is the manager’s opportunity cost.)

Both of these have special names and roles in the principal-agent literature. The latter one is called the **participation constraint**, or **individual rationality constraint**. That is, any contract is constrained by the fact that the manager must willingly participate in it. Thus, the contract $w^*(\Pi)$ must satisfy

$$\text{IR : } \max_e \{\mathbf{E}U(w(\Pi(e, \rho)), e)\} \geq U_0.$$

The other constraint is that the manager’s optimal choice should, in equilibrium, correspond to what the owner wanted the manager to do. That is, if the owner wants to elicit effort level \hat{e} , then it should be true that the manager, in equilibrium, actually supplies that level. This is called the **incentive compatibility constraint**. Mathematically it says:

$$\text{IC : } \hat{e} = \operatorname{argmax}_e \{\mathbf{E}U(w(\Pi(e, \rho)), e)\}.$$

Here ‘argmax’ indicates the argument which maximizes. In other words, we want $\hat{e} = e^*$.

Now we can write down the principal's problem. The principal 'simply' wishes to maximize his own payoff subject to both, the participation and incentive compatibility constraints. This problem is easily stated

$$\max_{e,w(\Pi)} \{ \mathbf{E} [\Pi(e, \rho) - w(\Pi(e, \rho))] + \lambda_P (U_0 - U(w(\Pi(e, \rho)), e)) + \lambda_I (U_w(\cdot, \cdot)w'(\cdot)\Pi_e(\cdot, \cdot) + U_e(\cdot, \cdot)) \},$$

but hard to solve in general cases. We will, therefore, only look at three special cases in which some answers are possible.

First, let us simplify by assuming a risk neutral manager. In that case we can have $U(w, e) = w - v(e)$, where $v(e)$ is the disutility of effort. Now note that we would not have to insure the manager in this case since he evaluates expected wealth the same as some fixed payment. Also note that there is no reason to over pay the manager compared to the outside option, that is, he needs to be given only U_0 . Before we charge ahead and solve this brute force, let us think for a moment longer about the situation. We will not be insuring the manager. He will furthermore have the correct incentive to expend effort if he cares as much about profits as the owner. Thus it would stand to reason that if we were to sell him the profits he would be the owner and thus take the optimal action! But this is equivalent to paying him a wage which is equal to the profits less some fixed return for the owners. So, let us propose a wage contract of $w = \Pi - p$, and assume for the moment that the IR constraint can be satisfied (by choice of the correct p .) With such a contract the manager will choose an effort level such that

$$\Pi_e(e, \rho) - v'(e) = 0.$$

Note that this is the efficient effort level under full information. The owner's problem now simply is to choose the largest p such that the manager still accepts the contract!

The other special case one can consider is that of an infinitely risk-averse manager. We can model this as a manager who has lexicographic preferences over his minimum wealth and effort. Independent of the effort levels the manager will prefer a higher minimum wealth, and for equal minimum wealth the manager will prefer the lower effort. For simplicity also assume that the lowest possible profit level is independent of effort (that is, only the probabilities of profits depend on the effort, not the level). In that case a manager will always face the same lowest wage for any wage function and thus will not be able to be enticed to provide more than the minimum level of effort.

As you can see, we might expect anything between efficient outcomes and completely inefficient outcomes, largely depending on the precise situation. Let us return for a moment to the general setting above. We will restrict it somewhat by assuming that the manager's preferences are separable: $U(w, e) = u(w) - v(e)$ with the obvious assumptions of $u'(\cdot) > 0, u''(\cdot) < 0, v'(\cdot) \geq 0, v''(\cdot) > 0, v'(\underline{e}) = 0, v'(\bar{e}) = \infty$. The last two are typical "Inada conditions" designed to ensure an interior solution. We will also assume that the cumulative distribution of profits exists and depends on effort, $F(\Pi, e)$, and that this distribution is over $\Pi \in [\underline{\Pi}, \bar{\Pi}]$, has a density $f(\Pi, e) > 0$, and satisfies first-order stochastic dominance.

In this formulation the manager will solve

$$\max_e \left\{ \int_{\underline{\Pi}}^{\bar{\Pi}} u(w(\Pi)) f(\Pi, e) d\Pi - v(e) \right\}.$$

This has first order condition

$$\int_{\underline{\Pi}}^{\bar{\Pi}} u(w(\Pi)) f_e(\Pi, e) d\Pi - v'(e) = 0.$$

We will ignore the SOC for a while. The participation constraint will be

$$\int_{\underline{\Pi}}^{\bar{\Pi}} u(w(\Pi)) f(\Pi, e) d\Pi - v(e) \geq U_0.$$

The owner will want to find a wage structure and effort level to

$$\max_{w(\cdot), e} \left\{ \int_{\underline{\Pi}}^{\bar{\Pi}} [(\Pi - w(\Pi)) f(\Pi, e) + \lambda_P (u(w(\Pi)) - v(e) - U_0) f(\Pi, e) + \lambda_C (u(w(\Pi)) f_e(\Pi, e) - v'(e) f(\Pi, e))] d\Pi \right\}.$$

This will have to be differentiated, which is quite messy and, as pointed out before, hard to solve. However, consider the differentiation with respect to the wage function:

$$-f(\Pi, e) + \lambda_P f(\Pi, e) u'(w(\Pi)) + \lambda_C f_e(\Pi, e) u'(w(\Pi)) = 0.$$

Rewriting we get something which is informative:

$$\frac{1}{u'(w(\Pi))} = \lambda_P + \lambda_C \frac{f_e(\Pi, e)}{f(\Pi, e)}.$$

The left hand side of this is the inverse of the slope of the wealth utility function. It is therefore increasing in the wage. The right hand side consists of the constraint multipliers (which will both be positive since the constraints are binding) multiplied by some factors. Of particular interest is the last term, the ratio of the derivative of the density to the density. Let us further specialize the problem and suppose that there are only two effort levels, and that the owner wants to get the high effort level. In that case it is easy to see that the above equation will become

$$\frac{1}{u'(w(\Pi))} = \lambda_P + \lambda_C \frac{f_H(\Pi) - f_L(\Pi)}{f_H(\Pi)} = \lambda_P + \lambda_C \left(1 - \frac{f_L(\Pi)}{f_H(\Pi)}\right).$$

The term $\frac{f_L(\Pi)}{f_H(\Pi)}$ is called a *likelihood ratio*. It follows that the higher the relative probability that the effort was high for a given realized profit level, the higher the manager's wage. Indeed, if the likelihood ratio is decreasing then then the wage function must be increasing with realized profits. What is going on here is that higher profits are a correct signal of higher effort, and thus we will make wages increase with the signal.

Finally, assume that there are only two profit levels, with the high level of profits more likely under high effort. In that case it can be shown that

$$0 < \frac{w_H - w_L}{\Pi_H - \Pi_L} < 1.$$

When does this principal-agent relationship come into play? We already mentioned the insurance problem, where the insurer wants to elicit the appropriate amount of care from the insured. It also occurs in politics, where the electorate would like to insure that the elected officials actually carry out the electorate's will (witness the 'recall' efforts in BC, presumably designed to make the wage function steeper, since under the old system any punishment only occurred after 4 years.) In essence a principal-agent relationship with the attendant problems exists any time there is a delegation of responsibilities with less than perfect monitoring. Most frequently in economics this is the case in the production process. Owners of firms delegate to managers, who have superior information about market circumstances, technologies, etc. Managers in turn delegate to subordinates — who often have superior information about the details of the production process. (Do quality problems arise inherently from the production process employed, or are they due to negligent work?) Any time a firm subcontracts another firm there is another principal-agent relationship set up. Indeed, even in our class there are multiple principal-agent relationships. The university wants to ensure that I treat you appropriately, teach the appropriate material, and teach "well".

You, in some sense, are the principal of the university, having similar goals. However, you inherently do not know what it is you “ought” to be taught. You may also have no way to determine if you are taught well — especially since there is a potential conflict between your future self (desiring rigorous instruction) and your current self (desiring a ‘pleasant’ course.) There are various incentive systems in place to attempt to address these problems, and heated debate about how successful they are!

As an aside: the principal-agent problems arising from delegation are often held responsible for limits on firm size (i.e., are taken to cause decreasing returns to scale.)

Chapter 6

Game Theory

Most of the models you have encountered so far had one distinguishing feature: the economic agent, be it firm or consumer, faced a simple decision problem. Aside from the discussion of oligopoly, where the notion that one firm may react to another's actions was mentioned, none of our decision makers had to take into account other's decisions. For the price taking firm or consumer prices are fixed no matter what the consumer decides to do. Even for the monopolist, where prices are not fixed, the (inverse) demand curve is fixed. These models therefore can be treated as maximization problems in the presence of an (exogenous) constraint.

Only when duopoly was introduced did we need to discuss what, if any, effect one firm's actions might have on the other firm's actions. Usually this problem is avoided in the analysis of the problem, however. For example, in the standard Cournot model we suppose a fixed market demand schedule and then determine one firm's optimal (profit maximizing) output under the assumption that the other firm produces some **fixed** output. By doing this we get the optimal output for each possible output level by the other firm (called reaction function) for each firm.¹ It then is argued that each firm must correctly forecast the opponent's output level, that is, that in equilibrium each firm is on its reaction function. This determines the equilibrium output level for both firms (and hence the market price).²

¹So, suppose two fixed marginal cost technology firms, and that inverse market demand is $p = A - BQ$. Each firm i then solves $\max_{q_i} \{(A - B(q_i + q_{-i}))q_i - c_i q_i\}$ which has FOC $A - Bq_{-i} - 2Bq_i - c_i = 0$ and thus the optimal output level is $q_i = (A - c_i)(2B)^{-1} - 0.5q_{-i}$.

²To continue the above example: $q_1 = (A - c_1)(2B)^{-1} - 0.5((A - c_2)(2B)^{-1} - 0.5q_1) \implies q_1 = (A - 2c_1 + c_2)(3B)^{-1}$. Thus $q_2 = (A - 2c_1 + c_2)(3B)^{-1}$ and $p = (A + c_1 + c_2)(3B)^{-1}$.

In this kind of analysis we studiously avoid allowing a firm to consider an opponent's actions as somehow dependent on its own. Yet, we could easily incorporate this kind of thinking into our model. We could not only call it a reaction function, but actually consider it a reaction of some sort. Doing so requires us to think about (or model) the reactions to reactions, etc.. Game theory is the term given to such models. The object of game theory is the analysis of **strategic decision problems** — situations where

- The outcome depends on the decisions of multiple ($n \geq 2$) decision makers, such that the outcome is not determined unilaterally;
- Everybody is aware of the above fact;
- Everybody assumes that everybody else conforms to fact 2;
- Everybody takes all these facts into account when formulating a course of action.

These points especially interesting if there exists a **conflict of interest** or a **coordination problem**. In the first case, any payoff gains to one player imply payoff losses to another. In the second case, both players' payoffs rise and fall together, but they cannot agree beforehand on which action to take. Game theory provides a formal language for addressing such situations.

There are two major branches in game theory — **Cooperative game theory** and **Non-cooperative game theory**. They differ in their approach, assumptions, and solution concepts. Cooperative game theory is the most removed from the actual physical situation/game at hand. The basis of analysis is the set of feasible payoffs, and the payoffs players can obtain by not participating in the first place. Based upon this, and without any knowledge about the underlying rules, certain properties which it is thought the solution ought to satisfy are postulated — so called axioms. Based upon these axioms the set of points which satisfy them are found. This set may be empty — in which case the axioms are not compatible (for example Arrow's Impossibility Theorem) — have one member, or have many members. The search is on for the fewest axioms which lead to a unique solution, and which have a "natural" interpretation, although that is often a more mathematical than economic metric. One of the skills needed for this line of work is a pretty solid foundation in functional analysis — a branch of mathematics concerned with properties of functions. We will not talk much about this type of game theory in this course, although it will pop up once later, when we talk about bargaining (in the Nash Bargaining Solution). As a final note on this subject: this type of game theory is called "cooperative" not necessarily because

players will cooperate, but since it is assumed that players' commitments (threats, promises, agreements) are binding and can be enforced. Concepts such as the Core of an exchange economy³ or the Nash bargaining solution are all cooperative game theory notions.

In this chapter we will deal exclusively with non-cooperative game theory. In this branch the focus is more on the actual rules of the game — it is thus a useful tool in the analysis of how the rules affect the outcome. Indeed, in the mechanism design and the implementation literature researchers basically design games in order to achieve certain outcomes. Non-cooperative game theory can be applied to games in two broad categories differing with respect to the detail employed in modelling the situation at hand (in fact, there are many ways in which one can categorize games; by the number of players, the information possessed by them, or the question if there is room for cooperation or not, among others.) The most detailed branch employs the **extensive form**, while a more abstracted approach employs the **strategic form**.

While we are ultimately concerned with economic problems, much of what we do in the next pages deals with toy examples. Part of the reason for the name game theory is the similarity of many strategic decision problems to games. There are certain essential features which many economic situations have in common with certain games, and a study of the simplified game is thus useful preparation for the study of the economic situation. Some recurring examples of particular games are the following (you will see that many of these games have names attached to them which are used by all researchers to describe situations of that kind.)

Matching Pennies: 2 players simultaneously⁴ announce Head or Tails. If the announcements match, then player 2 pays player 1 one dollar; if they

³The 'Core' refers to the set of all Pareto Optimal allocations for which each player/trader achieves at least the same level of utility as in the original allocation. For a two player exchange economy it is the part of the Contract Curve which lies inside the trading lens. The argument is that any trade, since it is voluntary, must improve each player at least weakly, and that two rational players should not stop trading until all gains from trade are exhausted. The result follows from this. It is not more specific, since we do not know the trading procedures. A celebrated result in economics (due to Edgeworth) is that the core converges to the competitive equilibrium allocation as the economy becomes large (i.e., players are added, so that the number of players grows to infinity.) This demonstrates nicely that in a large economy no one trader has sufficient market power to influence the market price.

⁴This means that no information is or can be transmitted. It does not necessarily mean actually simultaneous actions. While simultaneousness is certainly one way to achieve the goal, the players could be in different rooms without telephones or shared walls.

do not match, then player 1 pays player 2 one dollar. This is an example of a zero-sum game, the focus of much early game theoretic analysis. The distinguishing feature is that what is good for one is necessarily bad for the other player. The game is one of pure conflict. It is also a game of incomplete information.

Battle of the Sexes: Romeo and Juliet rather share an activity than not. However, Romeo likes music better than sports while Juliet likes sports better than music. They have to choose an activity without being able to communicate with each other. This game is a coordination game — the players want to coordinate their activities, since that increases their payoffs. There is some conflict, however, since each player prefers coordination on a different activity.

Prisoners' Dilemma: Probably one of the most famous (and abused) games, it captures the situation in which two players have to choose to cooperate or not (called 'defect') simultaneously. One form is that both have to announce one of two statements, either "give the other player \$3000" or "give me \$1000." The distinguishing feature of this game is that each player is better off by not cooperating — independently of what the other player does — but as a group they would be better off if both cooperated. This is a frequent problem in economics (for example in duopolies, where we will use it later.)

"Punishment Game:" This is not a standard name or game. I will use it to get some ideas across, however, so it might as well get a name. It is a game with sequential moves and perfect information: first the child chooses to either behave or not. Based on the observation of the behaviour, the parents then decide to either punish the child, or not. The child prefers not to behave, but punishment reduces its utility. The parents prefer if the child behaves, but dislike punishing it.

There are many other examples which we will encounter during the remainder of the course. There are **repeated games**, where the same so called stage game is repeated over and over. While the stage game may be a simultaneous move game, such as the Prisoners' Dilemma, after each round the players get to observe either payoffs or actual actions in the preceding stage. This allows them to condition their play on the opponents' play in the past (a key mechanism by which cooperation and suitable behaviour is enforced in most societies.)

6.1 Descriptions of Strategic Decision Problems

6.1.1 The Extensive Form

We will start rather formally. In what follows we will use the correct and formal way to define games and you will therefore not have to reformatize some foggy notions later on — the down side is that this may be a bit unclear on first reading. Much of it is only jargon, however, so don't be put off!

The first task in writing down a game is to give a complete description of the players, their possible actions, the timing of these actions, the information available to the players when they take the actions, and of course the payoffs they receive in the end. This information is summarized in a *game tree*. Now, a game tree has to satisfy certain conditions in order to be sensible: It has to be finite (otherwise, how would you write it down?), and it has to be connected (so that we can get from one part of the tree to another), and it has to be like a tree in that we do not have loops in it, where two branches join up again.⁵ All this is formally said in the following way:

Definition 1 *A game tree Γ (also called a topological tree) is a finite collection of nodes, called **vertices**, connected by lines, called **arcs**, so as to form a figure which is connected (there exists a set of arcs connecting any one vertex to any other) and contains no simple closed curves (there does not exist a set of arcs connecting a vertex to itself.)*

In Figure 6.1 A, B, and C satisfy the definition, D, E, and F do not.

We also need a sense of where the game starts and where it ends up. We will call the start of the game the **distinguished node/vertex** or more commonly the **root**. We can then define the following:

Definition 2 *Let Γ be a tree with root A . Vertex C **follows** vertex B if the sequence of arcs connecting A to C passes through B . C **follows** B **immediately** if C follows B and there is one arc connecting C to B . A vertex is called **terminal** if it has no followers.*

⁵Finite can be dropped. Indeed we will consider infinite games, such as the Cournot duopoly game where each player has infinitely many choices. The formalism to this extension is left to more advanced courses and texts.




Figure 6.1: Examples of valid and invalid trees

We are now ready to define a game formally by defining a bunch of objects and a game tree to which they apply. These objects and the tree will capture all the information we need. What is this information? We need the number of players, the order of play, i.e., which player plays after/together with what other player(s), what each player knows when making a move, what the possible moves are at each stage, what, if any, exogenous moves exist, and what the probability distribution over them is, and of course the payoffs at the end. So, formally, we get the following:

Definition 3 *A n-player game in extensive form comprises:*

1. A **game tree** Γ with root A ;
2. A function, called **payoff function**, associating a vector of length n with each terminal vertex of Γ ;
3. A partition $\{S_0, S_1, \dots, S_n\}$ of the set of nonterminal nodes of Γ (the **player sets**);
4. For each vertex in S_0 a probability distribution over the set of immediate followers;
5. for each $i \in \{1, 2, \dots, n\}$ a partition of S_i into subsets S_i^j (**information sets**), such that $\forall B, C \in S_i^j$, B and C have the same number of immediate followers;
6. for each S_i^j an **index set** I_i^j and a 1-1 map from I_i^j to the set of immediate followers of each vertex in S_i^j .

This is a rather exhaustive list. Notice in particular the following: (1) While the game is for n players, we have $(n + 1)$ player sets. The reason is that “nature” gets a player set too. Nature will be a very useful concept. It allows us to model a non-strategic choice by the environment, most often the outcome of some randomization (as when there either is an accident or not, but none of the players is choosing this.) (2) Nature, since it is non-strategic, does not have a payoff in the game, and it does not have any information sets. (3) Our information sets capture the idea that a player can not distinguish the nodes within them, since every node has the same number of possible moves and they are the same (i.e., their labels are the same.)

Even with all the restrictions already implied, there are still various ways one could draw such a game. Two very important issues deal with assumptions on information — properties of the information sets, in other words. The first is a restriction on information sets which will capture the idea that players do not forget any information they learn during the game.

Definition 4 *A n -person game in extensive form is said to be a game of perfect recall if all players never forget information once known, and if they never forget any of their own previous moves: i.e., if $x, x' \in S_i^j$ then neither x nor x' is a predecessor of the other one, and if \hat{x} is a predecessor of x and the same player moves at x and \hat{x} (i.e., $\hat{x}, x \in S_i$), then there exists some \tilde{x} in the same S_i^j as \hat{x} which is a predecessor of x' , and which has the same index on the arc leading to x' as that from \hat{x} to x .*

This definition bears close reading, but is quite intuitive in practice: if two nodes are in a player’s information set, then one cannot follow the other — else the player in effect forgets that he himself moved previously in order to get to the current situation. Furthermore, there is a restriction on predecessors: it must be true that either both nodes have the same predecessor in a previous information set of this player, or if not, that “the same action” was chosen at the two different predecessor nodes. Otherwise, the player should remember the index he chose previously and thus the current two nodes cannot be indistinguishable.

We will always assume perfect recall. In other words, all our players will recall all their own previous moves, and if they have learned something about their opponents (such as that the opponent took move “Up” the second time he moved) then they will not forget it. In practice, this means that players learn something during the game, in some sense.

Games of perfect recall still allow for the possibility of players not know-




Figure 6.2: A Game of Perfect Recall and a Counter-example

ing something about the past of the game — for example in games where players move simultaneously, a player would not know his opponent's move at that time. Such games are called games of imperfect information. The opposite is a game of perfect information:

Definition 5 *A n -person game in extensive form is said to be a game of perfect information if all information sets are singletons.*

It is important to note a crucial linguistic difference here: Games of **Incomplete Information** are not the same as games of **Imperfect Information**. Incomplete information refers to the case when some feature of the extensive form is not known to one (or more) of the players. For example the player may not know the payoff function, or even some of the possible moves, or their order. In that case, the player could not write down the extensive form at all! In the case of imperfect information the player can write down the extensive form, but it contains non-trivial information sets.

A famous theorem due to Harsanyi shows that it is possible to transform a situation of incomplete information to one of imperfect information if players are Bayesian.⁶ In that case, uncertainty about the number of players,

⁶By this we mean that they use Bayes' formula to update their beliefs. Bayes' formula




Figure 6.3: From Incomplete to Imperfect Information

available moves, outcomes, etc., can be transformed into uncertainty about payoffs only. From that, one can construct a complete information game with imperfect information, and the equilibria of these games will coincide. For example, take the case of a market entrant who does not know if the monopolist enjoys fighting an entrant, or not. The entrant thus does not know the payoffs of the game. However, suppose it is known that there are two types of monopolist, one that enjoys a fight, and one that does not. Nature is assumed to choose which one is actually playing, and the probability of that choice is set to coincide with the entrant’s priors.⁷ After this transformation we have a well specified game and can give the extensive form — as in Figure 6.3. This turns out to be one of the more powerful results, since it makes our tools very useful to lots of situations. If we could not transform incomplete information into imperfect information, then we could not model most interesting situations — which nearly all have some information that players don’t know.

6.1.2 Strategies and the Strategic Form

In order to analyze the situation modelled by the extensive form, we employ the concept of **strategies**. These are complete, contingent plans of behaviour in the game, not just a “move,” which refers to the action taken at any particular information set. You should think of a strategy as a complete game plan, which could be given to a referee or a computer, and they would then play the game for you according to these instructions, while you just

says that the probability of an event A occurring given that an event B has occurred will be the probability of the event A occurring times the probability of B happening when A does, all divided by the probability of B occurring:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(B|A)P(A)}{P(B)}.$$

⁷A **prior** is the ex-ante belief of a player. The ex-post probability is called a **posterior**.

watch what happens and have no possibility to change your mind during the actual play of the game. This is a very important point. The players have to submit a **complete** plan before they start the game, and it has to cover all eventualities — which translates into saying that the plan has to specify moves for all information sets of the player, even those which prior moves of the same player rule out!

Definition 6 A **pure strategy** for player $i \in \{1, 2, \dots, n\}$ is a function σ_i that associates every information set S_i^j with one element of the index set I_i^j , $\sigma_i : S_i^j \mapsto I_i^j$.

Alternatively, we can allow the player to randomize.⁸ This randomization can occur on two levels, at the level of each information set, or at the level of pure strategies:

Definition 7 A **behavioural strategy** for player $i \in \{1, 2, \dots, n\}$ is a function β_i that associates every information set S_i^j with a probability distribution over the elements of the index set I_i^j .

Definition 8 A **mixed strategy** μ_i for player $i \in \{1, 2, \dots, n\}$ is a probability distribution over the pure strategies $\sigma_i \in \Sigma_i$.

Notice that these are not the same concepts, in general. However, under perfect recall one can find a behavioural strategy corresponding to each mixed strategy, and so we will only deal with mixed strategies (which are more properly associated with the strategic form, which we will introduce shortly.) Mixed strategies are also decidedly easier to work with.

We can now consider what players' payoffs from a game are. Consider pure strategies only, for now. Each player has pure strategy σ_i , giving rise to a **strategy vector** $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) = (\sigma_i, \sigma_{-i})$. In general, σ does not determine the outcome fully, however, since there may be moves by nature. We

⁸This is a somewhat controversial issue. Do people flip coins when making decisions? Nevertheless it is pretty much generally accepted. In some circumstances randomization can be seen as a formal equivalent of bluffing: Take Poker for example. Sometimes you fold with a pair, sometimes you stand, sometimes you even raise people. This could be modelled as a coin flip. In other instances the randomizing distribution is explained by saying that while each person plays some definite strategy, a population may not, and the randomization probabilities just correspond to the proportion of people in the population who play that strategy. We will not worry about it, however, and assume randomization as necessary (and sometimes it is, as we will see!)

therefore use von Neumann-Morgenstern expected utility to evaluate things. In general, the payoffs players receive from a strategy combination (vector) are therefore expected payoffs. In the end, players will, of course, arrive at precisely one terminal vertex and receive whatever the payoff vector is at that terminal vertex. Before the game is played, however, the presence of nature or the use of mixed strategies implies a probability distribution over terminal vertices, and the game and strategies are thus evaluated using expected payoffs. Define the following:

Definition 9 *The expected payoff of player i given $\sigma = (\sigma_i, \sigma_{-i})$, is $\pi_i(\sigma)$. The vector of expected payoffs for all players is $\pi(\sigma) = (\pi_1(\sigma), \dots, \pi_n(\sigma))$.*

Definition 10 *The function $\pi(\sigma)$ associated with the n -person game Γ in extensive form is called the **strategic form** associated with Γ . (It is also known as the “normal form,” but that language is coming out of use.)*

We will treat the strategic form in this fashion — as an abbreviated representation of the sometimes cumbersome extensive form. This is the prevalent view nowadays, and this interpretation is stressed by the term “strategic form.” There is a slightly different viewpoint, however, since so-called matrix-games were actually analyzed first. Thus, one can also see the following definition:

Definition 11 *A game G in strategic (normal) form is a 3-tuple (N, S, U) , where N is the player set $\{1, \dots, n\}$, S is the strategy set $S = S_1 \times S_2 \dots \times S_n$, where S_i is player i 's strategy set, and U is the payoff function $U : S \mapsto \mathcal{R}^n$; and a set of rules of the game, which are implicit in the above.*

This is a much more abstract viewpoint, where information is not only suppressed, but not even mentioned, in general. The strategic form can be represented by a matrix (hence the name matrix games.) Player 1 is taken to choose the row, player 2 the column, and a third player would be choosing among matrices. (For more than three players this representation clearly loses some of its appeal.) Figure 6.4 provides an example of a three player game in strategic form.

A related concept, which is even more abstract, is that of a game form. Here, only outcomes are specified, not payoffs. To get a game we need a set of utility functions for the players.

MATRIX A				←	PLAYER 3	→	MATRIX B			
1\2	<i>L</i>	<i>R</i>	<i>C</i>				1\2	<i>L</i>	<i>R</i>	<i>C</i>
<i>U</i>	(1, 1, 1)	(2, 1, 2)	(1, 3, 2)				<i>U</i>	(1, 1, 2)	(2, 1, 3)	(1, 3, 1)
<i>C</i>	(1, 2, 1)	(1, 1, 1)	(2, 3, 3)				<i>C</i>	(1, 2, 2)	(1, 1, 0)	(2, 3, 4)
<i>D</i>	(2, 1, 2)	(1, 1, 3)	(3, 1, 1)				<i>D</i>	(2, 1, 0)	(1, 1, 5)	(3, 1, 2)

Figure 6.4: A Matrix game — game in strategic form

Definition 12 A **game form** is a 3-tuple (N, S, O) where N and S are as defined previously and O is the set of physical outcomes.

You may note a couple of things at this point. For one, different extensive form games can give rise to the same strategic form. The games may not even be closely related for this to occur. In principle, realizing that the indices in the index sets are arbitrary, and that we can relabel everything without loss of generality (does it matter if we call a move “UP” or “OPTION 1”?), any extensive form game with, say, eight strategies for a player will lead to a matrix with eight entries for that player. But we could have one information set with eight moves, or we could have three information sets with two moves each. We could also have two information sets, one with four moves, one with two. The extensive forms would thus be widely different, and the games would be very different indeed. Nevertheless, they could all give rise to the same matrix. Does this matter? We will have more to say about this later, when we talk about solution concepts. The main problem is that one might want all games that give rise to the same strategic form to have the same solution — which often they don’t. What is “natural” in one game may not be “natural” in another.

The second point concerns the fact that the strategic form is not always a more convenient representation. Figure 6.5 gives an example.

This is a simple bargaining game, where the first player announces if he wants 0, 50 or 100 dollars, then the second player does the same. If the announcements add to \$100 or less the players each get what they asked for, if not, they each pay one dollar. While the extensive form is simple, the strategic form of this game is a 3×27 matrix!⁹

⁹Why? Since a strategy for the second player is an announcement after **each** announcement by the first player it is a 3-tuple. For each of the elements there are three possible moves, so that there are 3^3 different vectors that can be constructed.

1\2	(0, 0, 0)	(0, 0, 50)	...	(100, 100, 100)
0	(0, 0)	(0, 0)	...	(0, 100)
50	(50, 0)	(50, 0)	...	(-1, -1)
100	(100, 0)	(-1, -1)	...	(-1, -1)




Figure 6.5: A simple Bargaining Game

Before we go on, Figure 6.6 below and on the next page gives the four games listed in the beginning in both their extensive and strategic forms. Note that I am following the usual convention that in a matrix the first payoff belongs to the row player, while in an extensive form payoffs are listed by index.

Matching Pennies


1\2	<i>H</i>	<i>T</i>
<i>H</i>	(1, -1)	(-1, 1)
<i>T</i>	(-1, 1)	(1, -1)

Battle of the Sexes

1\2	<i>M</i>	<i>S</i>
<i>M</i>	(50, 30)	(5, 5)
<i>S</i>	(1, 1)	(30, 50)

Prisoners' Dilemma

1\2	<i>C</i>	<i>D</i>
<i>C</i>	(3, 3)	(0, 4)
<i>D</i>	(4, 0)	(1, 1)



The “Education Game”




Figure 6.6: The 4 standard games

6.2 Solution Concepts for Strategic Decision Problems

We have developed two descriptions of strategic decision problems (“games”). How do we now make a prediction as to the “likely” outcome of this?¹⁰ We will employ a **solution concept** to “solve” the game. In the same way in which we impose certain conditions in perfect competition (such as, “markets clear”) which in essence say that the equilibrium is a situation where everybody is able to carry out their planned actions (in that case, buy and sell as much as they desire at the equilibrium price), we will impose conditions on the strategies of players (their planned actions in a game). Any combination of strategies which satisfy these conditions will be called an equilibrium. Since there are many different conditions one could impose, the equilibrium is usually qualified by a name, such as “these strategies constitute a Nash equilibrium (Bayes Nash equilibrium, perfect equilibrium, subgame perfect equilibrium, the Cho-Kreps criterion,...) The equilibrium outcome is determined by the equilibrium strategies (and moves by nature.) In general, you will have to get used to the notion that there are many equilibrium outcomes for a game. Indeed, in general there are many equilibria for one game. This is part of the reason for the many equilibrium concepts, which try to “refine away” (lingo for “discard”) outcomes which do not appear to be sensible. There are about 280 different solution concepts — so we will only deal with a select few which have gained wide acceptance and are easy to work with (some of the others are difficult to apply to any given game.)

¹⁰It is sometimes not quite clear what we are trying to do: tell players how they should play, or determine how they will play. There are some very interesting philosophical issues at stake here, for a discussion of which we have neither the time nor the inclination! However, let it be noted here that the view taken in this manual is that we are interested in prediction only, and do not care one iota if players actually determine their actions in the way we have modeled.

6.2.1 Equilibrium Concepts for the Strategic Form

We will start with equilibrium concepts for the strategic form. The first is a very persuasive idea, which is quite old: Why don't we eliminate a strategy of a player which is strictly worse than all his other strategies no matter what his opponents do? This is known as **Elimination of (Strictly) Dominated Strategies**. We will not formally define this since there are various variants which use this idea (iterated elimination or not, of weakly or strictly dominated strategies), but the general principle should be clear from the above. What we will do, is to define what we mean by a dominated strategy.

Definition 13 *Strategy a **strictly dominates** strategy b if the payoff to the player is larger under a , independent of the opponents' strategies:*

$$a \text{ strictly dominates } b \text{ if } \pi_i(a, s_{-i}) > \pi_i(b, s_{-i}) \quad \forall s_{-i} \in S_{-i}.$$

A similar definition can be made for **weakly dominates** if the strict inequality is replaced by a weak inequality. Other authors use the notion of a dominated strategy instead:

Definition 14 *Strategy a is **weakly (strictly) dominated** if there exists a mixed strategy α such that $\pi_i(\alpha, s_{-i}) \geq (>) \pi_i(a, s_{-i})$, $\forall s_{-i} \in \Sigma_i$ and $\pi_i(\alpha, s_{-i}) > \pi_i(a, s_{-i})$ for some $s_{-i} \in \Sigma_i$.*

If we have a 2×2 game, then elimination of dominated strategies may narrow down our outcomes to one point. Consider the "Prisoners' Dilemma" game, for instance. 'Defect' strictly dominates 'Cooperate' for both players, so we would expect both to defect. On the other hand, in "Battle of the Sexes" there is no dominated (dominating) strategy, and we would still not know what to predict. If a player has more than two strategies, we also do not narrow down the field much, even if there are dominated strategies. In that case, we can use **Successive Elimination of Dominated Strategies**, where we start with one player, then go to the other player, back to the first, and so on, until we can't eliminate anything. For example, in the following game

$1 \setminus 2$	(l, l)	(r, r)	(l, r)	(r, l)
L	$(2, 0)$	$(2, -1)$	$(2, 0)$	$(2, -1)$
R	$(1, 0)$	$(3, 1)$	$(3, 1)$	$(1, 0)$

player 1 does not have a dominated strategy. Player 2 does, however, since (r, l) is strictly dominated by (l, r) . If we also eliminate weakly dominated strategies, we can throw out (l, l) and (r, r) too, and then player 1 has a dominated strategy in L . So we would predict, after successive elimination of weakly dominated strategies, that the outcome of this game is $(R, (l, r))$.

There are some criticisms about this equilibrium concept, apart from the fact that it may not allow any predictions. These are particularly strong if one eliminates weakly dominated strategies, for which the argument that a player should never choose those appears weak. For example you might know that the opponent will play that strategy for which you are indifferent between two strategies. Why then would you eliminate one of these strategies just because somewhere else in the game (where you will not be) one is worse than the other?

Next, we will discuss the probably most widely used equilibrium concept ever, **Nash equilibrium**.¹¹ This is the most universally accepted concept, but it is also quite weak. All other concepts we will see are refinements of Nash, imposing additional constraints to those imposed by Nash equilibrium.

Definition 15 *A Nash equilibrium in pure strategies is a set of strategies, one for each player, such that each player's strategy maximizes that player's payoff, taking the other players' strategies as given:*

$$\sigma^* \text{ is Nash iff } \forall i, \forall \sigma_i \in \Sigma_i, \quad \pi_i(\sigma_i^*, \sigma_{-i}^*) \geq \pi_i(\sigma_i, \sigma_{-i}^*).$$

Note that the crucial feature of this equilibrium concept: each player takes the others' actions as given and plays a best response to them. This is the mutual best response property we first saw in the Cournot equilibrium, which we can now recognize as a Nash equilibrium.¹² Put differently, we only check against deviations by **one** player at a time. We **do not** consider mutual deviations! So in the Prisoners' Dilemma game we see that one player alone cannot gain from a deviation from the Nash equilibrium strategies

¹¹Nash received the Nobel price for economics in 1994 for this contribution. He extended the idea of mutual best responses proposed by von Neumann and Morgenstern to n players. He did this in his Ph.D. thesis. von Neumann and Morgenstern had thought this problem too hard when they proposed it in their book *Games and Economic Behaviour*.

¹²Formally we now have 2 players. Their strategies are $q_i \in [0, P^{-1}(0)]$. Restricting attention to pure strategies, their payoff functions are $\pi_i(q_1, q_2)$, so the strategic form is $(\pi_1(q), \pi_2(q))$. Denote by $b_i(q_{-i})$ the best response function we derived in footnote 1 of this chapter. The Nash equilibrium for this game is the strategy vector $(q^*_1, q^*_2) = (b_1(q^*_2), b_2(q^*_1))$. This, of course, is just the computation performed in footnote 2.

(Defect,Defect). We do not allow or consider agreements by both players to defect to (Cooperate,Cooperate), which would be better!

A Nash equilibrium in pure strategies may not exist, however. Consider, for example, the “Matching Pennies” game: If player 2 plays ‘H’ player 1 wants to play ‘H’, but given that, player 2 would like ‘T’, but given that 1 would like ‘T’, ... We may need mixed strategies to be able to have a Nash equilibrium. The definition for a mixed strategy Nash equilibrium is analogous to the one above and will not be repeated. All that changes is the definition of the strategy space. Since an equilibrium concept which may fail to give an answer is not that useful (hence the general disregard for elimination of dominated strategies) we will consider the question of existence next.

Theorem 1 *A Nash equilibrium in pure strategies exists for perfect information games.*

Theorem 2 *For finite games a Nash equilibrium exists (possibly in mixed strategies.)*

Theorem 3 *For (N, S, U) with $S \in \mathcal{R}^n$ compact and convex and $U_i : S \mapsto \mathcal{R}$ continuous and strictly quasi concave in s_i , a Nash equilibrium exists.*

Remarks:

1. Nash equilibrium is a form of rational expectations equilibrium (actually, a rational expectations equilibrium is a Nash equilibrium, formally.) As in a rational expectations equilibrium, the players can be seen to “expect” their opponent(s) to play certain strategies, and in equilibrium the opponents actually do, so that the expectation was justified.
2. There is an apparent contradiction between the first existence theorem and the fact that Nash equilibrium is defined on the strategic form. However, you may want to think about the way in which assuming perfect information restricts the strategic form so that matrices like the one for matching pennies can not occur.
3. If a player is to mix over some set of pure strategies $\{\sigma_i^1, \sigma_i^2, \dots, \sigma_i^k\}$ in Nash equilibrium, then all the pure strategies in the set must lead

to the same expected payoff (else the player could increase his payoff from the mixed strategy by changing the distribution.) This in turn implies that the fact that a player is to mix in equilibrium will impose a restriction on the other players' strategies! For example, consider the matching pennies game:

1\2	<i>H</i>	<i>T</i>
<i>H</i>	(1, -1)	(-1, 1)
<i>T</i>	(-1, 1)	(1, -1)

For player 1 to mix we will need that $\pi_1(H, \mu_2) = \pi_1(T, \mu_2)$. If β denotes the probability of player 2 playing *H*, then we need that $\beta - (1 - \beta) = -\beta + (1 - \beta)$, or $2\beta - 1 = 1 - 2\beta$, in other words, $\beta = 1/2$. For player 1 to mix, player 2 must mix at a ratio of $1/2 : 1/2$. Otherwise, player 1 will play a pure strategy. But now player 2 must mix. For him to mix (the game is symmetric) we need that player 1 mixes also at a ratio of $1/2 : 1/2$. We have, by the way, just found the unique Nash equilibrium of this game. There is no pure strategy Nash, and if there is to be a mixed strategy Nash, then it must be this. (Notice that we know there is a mixed strategy Nash, since this is a finite game!)

The next equilibrium concept we mention is **Bayesian Nash Equilibrium (BNE)**. This will be for completeness sake only, since we will in practice be able to use Nash Equilibrium. BNE concerns games of incomplete information, which, as we have seen already, can be modelled as games of imperfect information. The way this is done is by introducing "types" of one (or more) player(s). The type of a player summarizes all information which is not public (common) knowledge. It is assumed that each type actually knows which type he is. It is common knowledge what distribution the types are drawn from. In other words, the player in question knows who he is and what his payoffs are, but opponents only know the distribution over the various types which are possible, and do not observe the actual type of their opponents (that is, do not know the actual payoff vectors, but only their own payoffs.) Nature is assumed to choose types. In such a game, players' expected payoffs will be contingent on the actual types who play the game, i.e., we need to consider $\pi(\sigma_i, \sigma_{-i} | t_i, t_{-i})$, where t is the vector of type realizations (potentially one for each player.) This implies that each player type will have a strategy, so that player i of type t_i will have strategy $\sigma_i(t_i)$. We then get the following:

Definition 16 *A Bayesian Nash Equilibrium is a set of type contingent strategies $\sigma^*(t) = (\sigma_1^*(t_1), \dots, \sigma_n^*(t_n))$ such that each player maximizes his expected utility contingent on his type, taking other players' strategies as given, and using the priors in computing the expectation:*

$$\pi_i(\sigma_i^*(t_i), \sigma_{-i}^*|t_i) \geq \pi_i(\sigma_i(t_i), \sigma_{-i}^*|t_i), \quad \forall \sigma_i(t_i) \neq \sigma_i^*(t_i), \quad \forall i, \quad \forall t_i \in T_i.$$

What is the difference to Nash Equilibrium? The strategies in a Nash equilibrium are not conditional on type: each player formulates a plan of action before he knows his own type. In the Bayesian equilibrium, in contrast, each player knows his type when choosing a strategy. Luckily the following is true:

Theorem 4 *Let G be an incomplete information game and let G^* be the complete information game of imperfect information that is Bayes equivalent: Then σ^* is a Bayes Nash equilibrium of the normal form of G if and only if it is a Nash equilibrium of the normal form of G^* .*

The reason for this result is straight forward: If I am to optimize the expected value of something given the probability distribution over my types and I can condition on my types, then I must be choosing the same as if I wait for my type to be realized and maximize then. After all, the expected value is just a weighted sum (hence linear) of the conditional on type payoffs, which I maximize in the second case.

6.2.2 Equilibrium Refinements for the Strategic Form

So how does Nash equilibrium do in giving predictions? The good news is that, as we have seen, the existence of a Nash equilibrium is assured for a wide variety of games.¹³ The bad news is that we may get too many equilibria, and that some of the strategies or outcomes make little sense from a “common sense” perspective. We will deal with the first issue first. Consider the following game, which is a variant of the Battle of the Sexes game:

¹³One important game for which there is no Nash equilibrium is Bertrand competition between 2 firms with different marginal costs. The payoff function for firm 1, say, is

$$\pi_1(p_1, p_2) = \begin{cases} (p_1 - c_1)Q(p_1) & \text{if } p_1 < p_2 \\ \alpha(p_1 - c_1)Q(p_1) & \text{if } p_1 = p_2 \\ 0 & \text{otherwise} \end{cases}$$

which is not continuous in p_2 and hence Theorem 3 does not apply.

$1 \setminus 2$	M	S
M	(6, 2)	(0, 0)
S	(0, 0)	(2, 6)

This game has three Nash equilibria. Two are in pure strategies — (M, M) and (S, S) — and one is a mixed strategy equilibrium where $\mu_1(S) = 1/4$ and $\mu_2(S) = 3/4$. So what will happen? (Notice another interesting point about mixed strategies here: The expected payoff vector in the mixed strategy equilibrium is $(3/2, 3/2)$, but any of the four possible outcomes can occur in the end, and the actual payoff vector can be any of the three vectors in the game.)

The problem of too many equilibria gave rise to **refinements**, which basically refers to additional conditions which will be imposed on top of standard Nash. Most of these refinements are actually applied to the extensive form (since one can then impose restrictions on how information must be consistent, and so on.) However, there is one common refinement on the strategic form which is sometimes useful.

Definition 17 *A dominant strategy equilibrium is a Nash equilibrium in which each player's strategy choice (weakly) dominates any other strategy of that player.*

You may notice a small problem with this: It may not exist! For example, in the game above there are no dominating strategies, so that the set of dominant strategy equilibria is empty. If such an equilibrium does exist, it may be quite compelling, however.

There is another commonly used concept, that of **normal form perfect equilibrium**. We will not use this much, since a similar perfection criterion on the extensive form is more useful for what we want to do later. However, it is included here for completeness. Basically, normal form perfect will refine away some equilibria which are “knife edge cases.” The problem with Nash is that one takes strategies of the opponents as given, and can then be indifferent between one's own strategies. Normal form perfect eliminates this by forcing one to consider completely mixed strategies, and only allowing pure strategies that survive after the limit of these completely mixed strategies is taken. This eliminates many of the equilibria which are only brought about by indifference. We first define an “approximate” equilibrium for completely mixed games, then take the limit:

Definition 18 A completely mixed strategy for player i is one that attaches positive probability to every pure strategy of player i : $\mu_i(s_i) > 0 \forall s_i \in S_i$.

Definition 19 A n -tuple $\mu(\epsilon) = (\mu_1, \dots, \mu_n)$ is an ϵ -perfect equilibrium of the normal form game G if μ_i is completely mixed for all $i \in \{1, \dots, n\}$, and if

$$\mu_i(s_j) \leq \epsilon \text{ if } \pi_i(s_j, \mu_{-i}) \leq \pi_i(s_k, \mu_{-i}), s_k \neq s_j, \epsilon > 0.$$

Notice that this restriction implies that any strategies which are a poor choice, in the sense of having lower payoffs than other strategies, must be used very seldom. We can then take the limit as “seldom” becomes “never:”

Definition 20 A Perfect Equilibrium is the limit point of an ϵ -perfect equilibrium as $\epsilon \rightarrow 0$.

To see how this works, consider the following game:

1\2	T	B
t	(100, 0)	(-50, -50)
b	(100, 0)	(100, 0)


The pure strategy Nash equilibria of this game are (t, T) , (b, B) , and (b, T) . The unique normal form perfect equilibrium is (b, T) . This can easily be seen from the following considerations. Let α denote the probability with which player 1 plays t , and let β denote the probability with which player 2 plays T . 2’s payoff from T is zero independent of α . 2’s payoff from B is -50α , which is less than zero as long as $\alpha > 0$. So, in the ϵ -perfect equilibrium we have to set $(1 - \beta) < \epsilon$, that is $\beta > 1 - \epsilon$ in any ϵ -perfect equilibrium. Now consider player 1. His payoff from t will be $150\beta - 50$, while his payoff from b is 100. His payoff from t is therefore less than from b for all β , and we require that $\alpha < \epsilon$. As $\epsilon \rightarrow 0$, both α and $(1 - \beta)$ thus approach zero, and we have (b, T) as the unique perfect equilibrium.

While the payoffs are the same in the perfect equilibrium and all the Nash equilibria, the perfect equilibrium is in some sense more stable. Notice in particular that a very small probability of making mistakes in announcing or carrying out strategies will not affect the nPE, but it would lead to a potentially very bad payoff in the other two Nash equilibria.¹⁴

¹⁴Note that an nPE is Nash, but not vice versa.

6.2.3 Equilibrium Concepts and Refinements for the Extensive Form

Next, we will discuss equilibrium concepts and refinements for the extensive form of a game. First of all, it should be clear that a Nash equilibrium of the strategic form corresponds one-to-one with a Nash equilibrium of the extensive form. The definition we gave applies to both, indeed. Since our extensive form game, as we have defined it so far, is a finite game, we are also assured existence of a Nash equilibrium as before. Consider the following game, for example, here given in both its extensive and strategic forms:



This game has 3 pure strategy Nash equilibria: $(D, (d, d))$, $(U, (u, u))$, and $(U, (u, d))$.¹⁵ What is wrong with this? Consider the equilibrium $(D, (d, d))$. Player 1 moves first, and his move is observed by player 2. Would player 1 really believe that player 2 will play d if player 1 were to choose U , given that player 2's payoff from going to u instead is higher? Probably not. This is called an **incredible threat**. By threatening to play 'down' following an 'Up', player 2 makes his preferred outcome, D followed by d , possible, and obtains his highest possible payoff. Player 1, even though he moves first, ends up with one of his worst payoffs.¹⁶ However, player 2, if asked to follow his strategy, would rather not, and play u instead of d if he finds himself after a move of U . The move d in this information set is only part of a best reply because under the proposed strategy for 1, which is taken as given in a Nash equilibrium, this information set is never reached, and thus it does not matter (to player 2's payoff) which action is specified. This is a type of behaviour which we may want to rule out. This is done most easily by requiring all moves to be best replies for their part of the game, a concept we will now make more formal. (See also Figure 6.7)

¹⁵There are also some mixed strategy equilibria, namely $(U, (P_2^1(u) = 1, P_2^2(u) = \alpha))$ for any $\alpha \in [0, 1]$, and $(D, (P_2^1(u) \leq 1/4, P_2^2(u) = 0))$.

¹⁶It is sometimes not clear if the term 'incredible threat' should be used only if there is some actual threat, as for example in the education game when the parents threaten to punish. The more general idea is that of an action that is not a best reply at an information set. In this sense the action is not credible at that point in the game.




Figure 6.7: Valid and Invalid Subgames

Definition 21 Let V be a non-terminal node in Γ , and let Γ_V be the game tree comprising V as root and all its followers. If all information sets in Γ are either completely contained in Γ_V or disjoint from Γ_V , then Γ_V is called a **subgame**.

We can now define a subgame perfect equilibrium, which tries to exclude incredible threats by assuring that all strategies are best replies in all proper subgames, not only along the equilibrium path.¹⁷

Definition 22 A strategy combination is a **subgame perfect equilibrium (SPE)** if its restriction to every proper subgame is a subgame perfect equilibrium.

In the example above, only $(U, (u, d))$ is a SPE. There are three proper subgames, one starting at player 2's first information set, one starting at his second information set, and one which is the whole game tree. Only u is a best reply in the first, only d in the second, and thus only U in the last.


Remarks:

1. Subgame Perfect Equilibria exist and are a strict subset of Nash Equilibria.
2. Subgame Perfect equilibrium goes hand in hand with the famous “backward induction” procedure for finding equilibria. Start at the end of the game, with the last information sets before the terminal nodes,

¹⁷The equilibrium path is, basically, the sequence of actions implied by the equilibrium strategies, in other words the implied path through the game tree (along some set of arcs.)


and determine the optimal action there. Then back up one level in the tree, and consider the information sets leading up to these last decisions. Since the optimal action in the last moves is now known, they can be replaced by the resulting payoffs, and the second last level can be determined in a similar fashion. This procedure is repeated until the root node is reached. The resulting strategies are Subgame Perfect.

3. Incredible Threats are only eliminated if all information sets are singletons, in other words, in games of perfect information. As a counter-example consider the following game:



In this game there is no subgame starting with player 2's information set after 1 chose B or C , and therefore the equilibrium concept reverts to Nash, and we get that $(A, (c, c))$ is a SPE, even though c is strictly dominated by a in the non-trivial information set.

4. Notwithstanding the above, Subgame Perfection is a useful concept in repeated games, where a simultaneous move game is repeated over and over. In that setting a proper subgame starts in every period, and thus at least incredible threats with regard to future retaliations are eliminated.
5. Subgame Perfection and normal Form Perfect lead to different equilibria. Consider the game we used before when we analyzed nPE:



As we had seen before, the nPE is (b, T) , but since there are no subgames, the SPE are all the Nash equilibria, i.e., (b, T) , (t, T) and (b, B) .




Figure 6.8: The “Horse”

As we can see, SPE does nothing to prevent incredible threats if there are no proper subgames. In order to deal with this aspect, the following equilibrium concept has been developed. For games of imperfect information we cannot use the idea of best replies in all subgames, since a player may not know at which node in an information set he is. We would like to ensure that the actions taken at an information set are best responses nevertheless. In order to do so, we have to introduce what the player believes about his situation at that information set. By introducing a belief system — which specifies a probability distribution over all the nodes in each of the player’s information sets — we can then require all actions to be best responses given the belief system. Consider the example in Figure 6.8, which is commonly called “The Horse.” This game has two pure strategy Nash equilibria, (A, a, r) and (D, a, l) . Both are subgame perfect since there are no proper subgames at all. The second one, (D, a, l) , is “stupid” however, since player 2 could, if he is actually asked to move, play d , which would improve his payoff from 1 to 4. In other words, a is not a best reply for player 2 if he actually gets to move.

Definition 23 A system of beliefs ϕ is a vector of beliefs for each player, ϕ_i , where ϕ_i is a vector of probability distributions, ϕ_i^j , over the nodes in each of player i ’s information sets S_i^j :

$$\phi_i^j : x_k^j \mapsto [0, 1], \quad \sum_{k=1}^K \phi_i^j(x_k^j) = 1; \quad \forall x_k^j \in S_i^j.$$

Definition 24 An assessment is a system of beliefs and a set of strategies, (σ^*, ϕ^*) .

Definition 25 An assessment (σ, ϕ) is **sequentially rational** if

$$E^{\phi^*} [\pi_i(\sigma_i^*, \sigma_{-i}^* | S_i^j)] \geq E^{\phi^*} [\pi_i(\sigma_i, \sigma_{-i}^* | S_i^j)], \forall i, \forall \sigma_i \in \Sigma_i, \forall S_i^j.$$

Definition 26 An assessment (σ^*, ϕ^*) is **consistent** if

$(\sigma^*, \phi^*) = \lim_{n \rightarrow \infty} (\sigma_n, \phi_n)$, where σ_n is a sequence of completely mixed behavioural strategies and ϕ_n are beliefs consistent with σ_n being played (i.e., obtained by Bayesian updating.)

Definition 27 An assessment (σ^*, ϕ^*) is a **sequential equilibrium** if it is consistent and sequentially rational.

As you can see, some work will be required in using this concept! Reconsider the horse in Figure 6.8. The strategy combination (A, a, r) is a sequential equilibrium with beliefs $\alpha = 0$, where α denotes player 3's probability assessment of being at the left node. You can see this by considering the following sequence of strategies: 1 plays D with $(1/n)^2$, which converges to zero, as required. 2 plays d with $(1/n)$, also converging to zero. The consistent belief for three thus is given by (from Bayes' Rule)

$$\alpha(n) = \frac{(1/n)^2}{(1/n)^2 + (1 - (1/n)^2)(1/n)} = \frac{n}{n^2 + n - 1} = \frac{1}{n + 1 - 1/n},$$

which converges to zero as n goes to infinity. As usual, the tougher part is to show that there is no sequence which can be constructed that will lead to (D, a, l) . Here is a short outline of what is necessary: In order for 3 to play l we need beliefs which put at least a probability of $1/3$ on being at the left node. We thus need that 1 plays down almost surely, since player 2 will play d any time 3 plays l with more than a $1/4$ probability. But as weight shifts to l for 3, and 2 plays d , sequential rationality for 1 requires him to play A ($4 > 3$). This destroys our proposed setup.

Signalling Games

This is a type of game used in the analysis of quality choice, advertising, warranties, or education and hiring. The general setup is that an informed party tries to convey information to an uninformed party. For example, the fact that I spend money advertising should convey the information that my product is of high quality to consumers who are not informed about the quality of my product. There are other sellers of genuinely low quality,

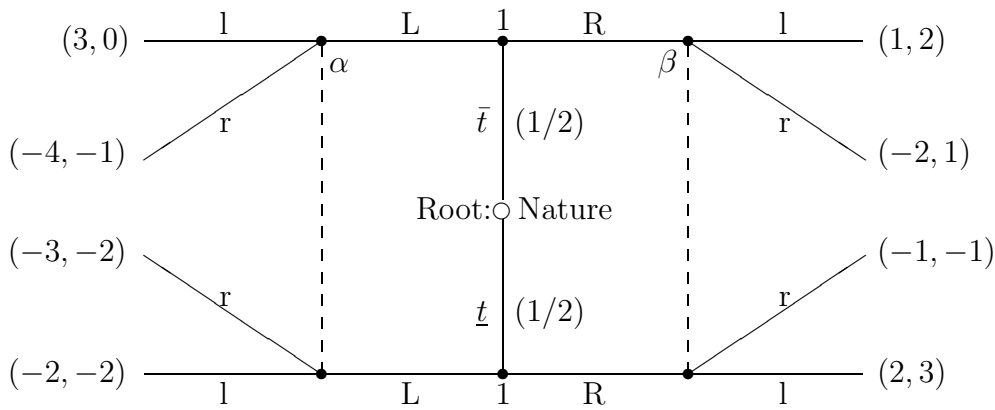


Figure 6.9: A Signalling Game

however, and they will try to mimic my actions. In order to be credible my action should therefore be hard (costly) to mimic. The equilibrium concept used for this type of game will be sequential equilibrium, since we have to model the beliefs of the uninformed party.

Consider the game in Figure 6.9. Nature determines if player 1 is a high type or a low type. Player 1 moves left or right, knowing his type. Player 2 then, without knowing 1's true type, moves left or right also. The payoffs are as indicated. This type of game is called a **game of asymmetric information**, since one of the parties is completely informed while the other is not. The usual question is if the informed party can convey its **private information** or not. In the above game, the Nash equilibria are the following: Let player 1's strategy vector (S_1, S_2) indicate 1's action if he is the low and high type, respectively, while player 2's strategy vector (s_1, s_2) indicates 2's response if he observes L and R , respectively. Then we get that the pure strategy Nash equilibria are $((R, R), (r, r))$, $((R, L), (r, r))$, and $((L, R), (l, l))$. Now introduce player 2's beliefs. Let 2's belief of facing a low type player 1 be denoted by α if 2 observes L , and by β if 2 observes R . We then can get two sequential equilibria: $((L, R), (l, r), (\alpha = 1, \beta = 0))$ and $((R, R), (r, r), (\alpha = 0, \beta = 0.5))$ (It goes without saying that you should try to verify this claim!). The first of these is a **separating equilibrium**. The action of player 1 completely conveys the private information of player 1. Only low types move Left, only high types move Right, and the move thus reveals the type of player. The second equilibrium, in contrast, is a **pooling equilibrium**. Both types take the same move in equilibrium, and no information is transmitted. Notice, however, that the belief that $\alpha = 0$ is somewhat stupid. If you find yourself, as player 2, inadvertently in your first information set, what should you believe? Your beliefs here say that

you think it must have been the high type who made the mistake. This is “stupid”, since the high type’s move L is strictly dominated by R , while the low type’s move L is not dominated by R . It would be more reasonable to assume, therefore, that if anything it was the low type who was trying to “tell you something” by deviating (you are not on the equilibrium path if you observe L , remember!)

There are refinements that impose restrictions like this last argument on the beliefs out of the equilibrium path, but we will not go into them here. Look up the Cho-Kreps criterion in any good game theory book if you want to know the details. The basic idea is simple: of the equilibrium path you should only put weight on types for whom the continuation equilibria off the equilibrium path are actually better than if they had followed the proposed equilibrium. The details are, of course, messy.

Finally, notice a last problem with sequential equilibrium. Minor perturbations of the extensive form change the equilibrium set. In particular, the two games in Figure 6.10 have different sequential equilibria, even though the games would appear to be quite closely related.




Figure 6.10: A Minor Perturbation?

6.3 Review Problems

Question 1: Provide the definition of a 3-player game in extensive form. Then draw a well labelled example of such a game in which you indicate all the elements of the definition.

Question 2: Define “perfect recall” and provide two examples of games

which violate perfect recall for **different** reasons.

Question 3: Assume that you are faced with some finite game. Will this game have a Nash equilibrium? Will it have a Subgame Perfect Equilibrium? Why can you come to these conclusions?

Question 4: Consider the following 3 player game in strategic form:

		<i>Left</i>		Player 3			<i>Right</i>	
1\2	<i>L</i>	<i>R</i>	<i>C</i>		1\2	<i>L</i>	<i>R</i>	<i>C</i>
<i>U</i>	(1, 1, 1)	(2, 1, 2)	(1, 3, 2)		<i>U</i>	(2, 2, 2)	(4, 2, 4)	(2, 6, 4)
<i>C</i>	(1, 2, 1)	(1, 1, 1)	(2, 3, 3)		<i>C</i>	(5, 0, 1)	(1, 1, 1)	(0, 1, 1)
<i>D</i>	(2, 1, 2)	(1, 1, 3)	(3, 1, 1)		<i>D</i>	(3, 2, 3)	(2, 2, 4)	(4, 2, 2)

Would elimination of weakly dominated strategies lead to a good prediction for this game? What are the pure strategy Nash equilibria of this game? Describe **in words** how you might find the mixed strategy Nash equilibria. Be clear and concise and **do not** actually attempt to solve for the mixed strategies.

Question 5: Find the mixed strategy Nash equilibrium of this game:

1\2	<i>L</i>	<i>R</i>	<i>C</i>
<i>U</i>	(1, 4)	(2, 1)	(4, 2)
<i>C</i>	(3, 2)	(1, 1)	(2, 3)

Question 6: Consider the following situation and construct an extensive form game to capture it.

A railway line passes through a town. Occasionally, accidents will happen on this railway line and cause damage and impose costs on the town. The frequency of these accidents depends on the effort and care taken by the railway — but these are unobservable by the town. The town may, if an accident has occurred, sue the railway for damages, but will only be successful in obtaining damages if it is found that the railway did not use a high level of care. For simplicity, assume that there are only two levels of effort/care (high and low) and that the courts can determine with certainty which level was in fact used. Also assume that going to court costs the railway and the town money, that effort is costly for the railway (high effort reduces profits), that accidents cost the railway and the town money and that this cost is independent of the effort level (i.e., there is a “standard accident”). Finally, assume that if the railway is “guilty” it has to pay the town’s damages and court costs.

Question 7: Consider the following variation of the standard Battle of the Sexes game: with probability α Juliet gets informed which action Romeo has taken before she needs to choose (with probability $(1 - \alpha)$ the game is as usual.)

- a) What is the subgame perfect equilibrium of this game?
- b) In order for Romeo to be able to insist on his preferred outcome, what would α have to be?

Question 8: (Cournot Duopoly) Assume that inverse market demand is given by $P(Q) = (Q - 10)^2$, where Q refers to market output by all firms. Assume further that there are n firms in the market and that they all have zero marginal cost of production. Finally, assume that all firms are Cournot competitors. This means that they take the other firms' outputs as given and consider their own inverse demand to be $p_i(q_i) = (\sum_{j \neq i} q_j + q_i - 10)^2$. Derive the Nash equilibrium output and price. (That is, derive the multilateral best response strategies for the output choices: Given every other firm's output, a given firm's output is profit maximizing for that firm. This holds for all firms.) Show that market output converges to the competitive output level as n gets large. (HINT: Firms are symmetric. It is then enough for now to focus on symmetric equilibria. One can solve for the so-called reaction function of one firm (it's best reply function) which gives the profit maximizing level of output for a given level of joint output by all others. Symmetry then suggests that each firm faces the same joint output by its $(n - 1)$ competitors and produces the same output in equilibrium. So we can substitute out and solve.)

Question 9*: Assume that a seller of an object knows its quality, which we will take to be the probability with which the object breaks during use. For simplicity assume that there are only two quality levels, high and low, with breakdown probabilities of 0.1 and 0.4, respectively. The buyer does not know the type of seller, and can only determine if a good breaks, but not its quality. The buyer knows that 1/2 of the sellers are of high quality, and 1/2 of low quality. Assume that the seller receives a utility of 10 from a working product and 0 from a non-working product, and that his utility is linear in money (so that the price of the good is deducted from the utility received from the good.) If the seller does not buy the object he is assumed to get a utility level of 0. The cost of the object to the sellers is assumed to be 2 for the low quality seller and 3 for the high quality seller. We want to investigate if signalling equilibria exist. We also want to train our understanding of sequential equilibrium, so use that as the equilibrium concept in what follows.

- a) Assume that sellers can only differ in the price they charge. Show that no separating equilibrium exist.

b) Now assume that the sellers can offer a warranty which will replace the good once if it is found to be defective. Does a separating equilibrium exist? Does a pooling equilibrium exist?

Question 10*: Assume a uniform distribution of buyers over the range of possible valuations for a good, $[0, 2]$.

a) Derive the market demand curve.

b) There are 2 firms with cost functions $C_1(q_1) = q_1/10$ and $C_2(q_2) = q_2^2$. Find the Cournot Equilibrium and calculate equilibrium profits.

c) Assume that firm 1 is a Stackelberg leader and compute the Stackelberg equilibrium. (This means that firm 1 moves first and firm 2 gets to observe firm 1's output choice. The Stackelberg equilibrium is the SPE for this game.)

d) What is the joint profit maximizing price and output level for each firm? Why could this not be attained in a Nash equilibrium?

Chapter 7

Review Question Answers

7.1 Chapter 2

Question 1:

a) There are multiple $C(\cdot)$ which satisfy the Weak Axiom. Note, however, that you have to check back and forth to make sure that the WA is indeed satisfied. (I.e., $C(\{x, y, z\}) = \{x\}$, $C(\{x, y\}) = \{x, y\}$ does not satisfy the axiom since while the check for x seems to be ok, you also have to check for y , and there it fails.) One choice structure that does work is $C(\{x, y, z\}) = \{x\}$, $C(\{x, z, w\}) = \{x\}$, $C(\{y, w, z\}) = \{w\}$, $C(\{y, w\}) = \{w\}$, $C(\{x, z\}) = \{x\}$, $C(\{x, w\}) = \{x\}$, $C(\{x\}) = \{x\}$.

b) Yes (I thought of that first, actually, in deriving the above) it is $x \succeq w \succeq y \succeq z$.

c) Yes, it is transitive.

d) I was aiming for an application of our Theorem: our set of budget sets \mathcal{B} does not contain all 2 and 3 element subsets of X . Missing are $\{x, y, w\}$, $\{x, y\}$, $\{y, z\}$, $\{w, z\}$.

e) The best way to go about this one is to determine where we can possibly get this to work. Examination of the sets B shows that the two choices y, x only appear in one of the sets and thus must be our key if we want to satisfy the WA without having rational preferences. Some fiddling reveals that the following works: $C(\{x, y, z\}) = \{x\}$, $C(\{x, z, w\}) = \{w\}$, $C(\{y, w, z\}) = \{y\}$, $C(\{y, w\}) = \{y\}$, $C(\{x, z\}) = \{x\}$, $C(\{x, w\}) = \{w\}$, $C(\{x\}) = \{x\}$. The problem is intransitivity, since the above implies that $y \succeq w \succeq x \succeq z$ but we also have $x \succeq y$!

Question 2: Here you have to make sure to maximize income for any given

work hrs	Part a			Partb		
	1 then 2	2 then 1	Max	1 then 2	2 then 1	Max
1	112	108	112	112	108	112
2	124	116	124	124	116	124
3	136	130	136	136	$130\frac{2}{3}$	136
4	148	144	148	148	$145\frac{1}{3}$	148
5	160	158	160	160	160	160
6	172	172	172	172	$174\frac{2}{3}$	$174\frac{2}{3}$
7	188	186	188	188	$189\frac{1}{3}$	$189\frac{1}{3}$
8	204	200	204	204	204	204
9	212	212	212	212	216	216
10	220	224	224	220	228	228
11	234	236	236	$234\frac{2}{3}$	240	240
12	248	248	248	$249\frac{1}{3}$	252	252
13	262	260	262	264	264	264
14	276	272	276	$278\frac{2}{3}$	276	$278\frac{2}{3}$
15	290	288	290	$293\frac{1}{3}$	292	$293\frac{1}{3}$
16	304	304	304	308	308	308

Table 7.1: Table 1: Computing maximal income

amount of work. In parts (a) and (b) you have to choose to work either job 1 then job 2 (after 8 hours in job 1) or job 2 then job 1. Simply plotting the two and then taking the outer hull (i.e., the highest frontier) for each leisure level gives you the frontier. In (a) they only cross twice (at 9 and 12 hours of work) while in part (b) they cross 4 times. You can best see this effect by considering a table in which you tabulate total hours worked against total income, computed by doing job 1 first, and by doing job 2 first. This is shown in Table 1. In neither part a) nor in part b) is the budget set convex.

c) This is a possibly quite involved problem. The intuitive answer is that it will not matter since marginal and average pay is (weakly) increasing in both jobs. Here is a more general treatment of these questions:

We really are faced with an maximization problem, to max income given the constraints, for any given total amount worked. Let h_1 and h_2 denote hours worked in jobs 1 and 2, respectively. Then the objective function is $I(h_1, h_2) = h_1w_1(h_1) + h_2w_2(h_2)$, where $w_i(h_i)$ are the wage schedules.

The wage schedules have the general form $w_1(h_1) = \begin{cases} \underline{w}_1 & \text{if } h_1 \leq C_1 \\ \overline{w}_1 & \text{if } h_1 \geq C_1 \end{cases}$ and

$w_2(h_2) = \begin{cases} \underline{w}_2 & \text{if } h_2 \leq C_2 \\ \overline{w}_2 & \text{if } h_2 \geq C_2 \end{cases}$, where $\underline{w}_i < \overline{w}_i$. I ignore here that no hours

above 8 are possible for either job, choosing to put that information into the constraints later.

Consider now the iso-income curves in (h_1, h_2) space which result. We will have four regions to consider, namely $A = \{(h_1, h_2) | h_1 \leq C_1, h_2 \leq C_2\}$, $B = \{(h_1, h_2) | h_1 \leq C_1, h_2 \geq C_2\}$, $C = \{(h_1, h_2) | h_1 \geq C_1, h_2 \leq C_2\}$, $D = \{(h_1, h_2) | h_1 \geq C_1, h_2 \geq C_2\}$. The slope of the iso-income curves for the regions is easily seen to be the negative of the ratio of wages, so we have $S(A) = -\underline{w}_1/\underline{w}_2$, $S(B) = -\underline{w}_1/\bar{w}_2$, $S(C) = -\bar{w}_1/\underline{w}_2$, $S(D) = -\bar{w}_1/\bar{w}_2$. It is obvious that $S(C) < S(D)$ and $S(A) < S(B)$, as well as that $S(C) < S(A)$ and $S(D) < S(B)$. This implies, of course, that $S(C) < S(A) < S(B)$ as well as that $S(C) < S(D) < S(B)$, with the comparison of $S(A)$ to $S(D)$ indeterminate. (But luckily not needed in any case.) The important fact which follows from all of this is that the iso-income curves are all concave to the origin and piece-wise linear.

Now superimpose the choice sets onto this. Note that without any restrictions $H = h_1 + h_2$, that is, for any given number of hours H the hours in each job are “perfect substitutes”. These iso-hour curves are all straight lines with a slope of -1 . (For our parameters all of $S(A), S(C), S(D)$ are less than -1 , while $S(B) > -1$, but this is not important.) For parts (a) and (b) the feasible set consists of the boundaries of the 8×8 square of feasible hours, where either $h_i = 0, h_j < 8$, or where $h_i = 8, 0 \leq h_j \leq 8$. The choice set is thus given by the intersection of the iso-hour lines with the feasible set (the box boundary). In part (c) this restriction is removed and the whole interior of the box is feasible. Due to the concavity to the origin of the iso-income lines this is of no relevance, however. Note how I have used our usual techniques of iso-objective curves and constraint sets to approach this problem. Works pretty well, doesn't it!

d) Now we “clearly” take up jobs in decreasing order of pay, starting with the highest paid and progressing to the lower paid ones in order. The resulting budget set will be convex.

Question 3: The consumer will

$$\max_x \{x_1^{0.3} x_2^{0.6} + \lambda(m - x_1 p_1 - x_2 p_2)\}$$

which leads to the first order conditions

$$0.3x_1^{-0.7} x_2^{0.6} = \lambda p_1, \quad 0.6x_1^{0.3} x_2^{-0.4} = \lambda p_2, \quad x_1 p_1 + x_2 p_2 = m.$$

The utility function is quasi-concave (actually, strictly concave in this case) and the budget set convex, so the second order conditions will be satisfied. Combining the first two first order conditions we get

$$\frac{0.3x_2}{0.6x_1} = \frac{p_1}{p_2} \implies x_2 = \frac{2p_1}{p_2} x_1.$$

Substitute into the budget constraint and simplify:

$$x_1 p_1 + \frac{2p_1}{p_2} x_1 p_2 = m \implies x_1 = \frac{m}{3p_1}.$$


Now use this to solve for x_2 : $x_2 = \frac{2m}{3p_2}$. So $(x_1(p, m), x_2(p, m)) = \left(\frac{m}{3p_1}, \frac{2m}{3p_2}\right)$.

To find the particular quantity demanded, simply plug in the numbers and simplify:

$$x_1 = \frac{3 \times 412 + 1 \times 72}{3 \times 3} = \frac{412 + 24}{3} = \frac{436}{3};$$

$$x_2 = \frac{2(3 \times 412 + 1 \times 72)}{3 \times 1} = \frac{2(412 + 24)}{1} = 872.$$

Question 4: The key is to realize that this utility function is piece-wise linear with line segments at slopes -5 , -1 , $-1/5$, from left to right. The segments join at rays from the origin with slopes 3 and $1/3$. Properly speaking, neither the Hicksian nor the Marshallian demands are functions. The function has either a perfect substitute or Leontief character. In the former the substitution effects approach infinity, in the latter they are zero. Demands are easiest derived from the price offer curve, which is a nice zigzag line. It starts at the intercept of the budget with the vertical axis (point A). It follows the indifference curve segment with -5 slope to the ray with slope 3. Call this point B. From there it follows the ray with slope 3 until that ray intersects a budget drawn from A with a slope of 1 (point C). It then continues on this budget and the coinciding indifference curve segment to the ray with slope $1/3$ (point D). Up along that ray to an intersection with a budget from A with slope $1/5$ (point E), along that budget to the intercept with the horizontal axis (point F), and then along the horizontal axis off to infinity.



Now we can solve for the demands along the different pieces of the offer curve and get the Marshallian demand. Note that demand is either a whole range, or a “proper demand”. The ranges can be computed from the endpoints (i.e., A to B, C to D, E to F.) Along the rays demand is solved as for

a Leontief consumer: we know the ratio of consumption, we know the budget. So for example on the first ray segment (B to C) we know that $x_2 = 3x_1$. Also, $p_1x_1 + p_2x_2 = w$. Hence $x_1(p_1 + 3p_2) = w$, and $x_1 = w/(p_1 + 3p_2)$. (For the Hicksian demand we simply need to fix one indifference curve and compute the points along it. We then get either a segment (like A to B above), or we stay at a kink for a range of prices.) The demands for good 1 therefore are

$$x_1(p, w) = \begin{cases} 0, & \text{if } p_1/p_2 > 5; \\ [0, 5w/(8p_1)], & \text{if } 5 = p_1/p_2; \\ w/(3p_2 + p_1), & \text{if } 5 > p_1/p_2 > 1; \\ [w/(4p_1), 3w/(4p_1)], & \text{if } p_1/p_2 = 1; \\ 3w/(3p_1 + p_2), & \text{if } 1 > p_1/p_2 > 1/5; \\ [3w/(8p_1), w/p_1], & \text{if } p_1/p_2 = 1/5; \\ w/p_1, & \text{if } 1/5 > p_1/p_2. \end{cases}$$

$$h_1(p, u) = \begin{cases} 0, & \text{if } p_1/p_2 > 5; \\ [0, u/8], & \text{if } p_1/p_2 = 5; \\ u/8, & \text{if } 5 > p_1/p_2 > 1; \\ [u/8, 3u/16], & \text{if } p_1/p_2 = 1; \\ 3u/16, & \text{if } 1 > p_1/p_2 > 1/5; \\ [3u/16, u], & \text{if } p_1/p_2 = 1/5; \\ u, & \text{if } 1/5 > p_1/p_2. \end{cases}$$

The demands for good 2 are similar and left as exercise. The income expansion paths and Engel curves can be whole regions at price ratios 1,5,1/5, otherwise the income expansion paths are the axes or rays, and the Engel curves are straight increasing lines.

Question 5: The elasticity of substitution measures by how much the consumption ratio changes as the price ratio changes (both measured in percentages.) In other words, as the price ratio changes the slope of the budget changes and we know this will cause a change in the ratio of the quantity demanded of the goods. But by how much? The higher the value of the elasticity, the larger the response in demands.

Question 6: First we need to realize that the utility index which each function assigns to a given consumption point does not have to be the same. Instead, as long as the MRS is identical at every point, two utility functions represent the same preferences. So instead of taking the limit of the utility function directly, we will take the limits of the MRS and compare those to the MRSs of the other functions.

$$MRS = \frac{u_1}{u_2} = \frac{(1/\rho)(x_1^\rho + x_2^\rho)^{(1-\rho)/\rho}(\rho x_1^{\rho-1})}{(1/\rho)(x_1^\rho + x_2^\rho)^{(1-\rho)/\rho}(\rho x_2^{\rho-1})} = \frac{x_1^{\rho-1}}{x_2^{\rho-1}} = \frac{x_2^{1-\rho}}{x_1^{1-\rho}}.$$

The MRSs for the other functions are

$$\text{CD: } \frac{x_2}{x_1}; \quad \text{Perfect Sub: } 1; \quad \text{Leon: } 0, \text{ or } \infty.$$

So, consider the Leontief function $\min\{x_1, x_2\}$. Its MRS is 0 or ∞ . But as $\rho \rightarrow -\infty$ we see that $(x_2/x_1)^{1-\rho} \rightarrow (x_2/x_1)^\infty$. But if $x_2 > x_1$ the fraction is greater than 1 and an infinite power goes to infinity. If $x_2 < x_1$ the fraction is less than one and the power goes to zero. The Cobb-Douglas function x_1x_2 has MRS x_2/x_1 . But as $\rho \rightarrow 0$ the MRS of our function is just that. The perfect substitute function $x_1 + x_2$ has a constant MRS of 1. But as $\rho \rightarrow 1$ the MRS of our function is $(x_2/x_1)^0 = 1$. Therefore the CES function “looks like” those three functions for those choices of ρ . The parameter ρ essentially controls the curvature of the IC’s.

Question 7: Set up the consumer’s optimization problem:

$$\max_{x_1, x_2, x_3} \{x_1 + \ln x_2 + 2\ln x_3 + \lambda(m - p_1x_1 - p_2x_2 - p_3x_3)\}.$$


The FOCs are

$$1 - \lambda p_1 = 0; \quad \frac{1}{x_2} - \lambda p_2 = 0; \quad \frac{2}{x_3} - \lambda p_3 = 0$$

and the budget. The first of these allows us to solve for $\lambda = 1/p_1$. Therefore the second and third give us $x_2 = p_1/p_2$ and $x_3 = 2p_1/p_3$. Combining this with the budget we get $x_1 = m/p_1 - 3$. Of course, this is only sensible if $m > 3p_1$. If it is not we must be at a corner solution. In that case $x_1 = 0$ and all money is spent on x_2 and x_3 . The second and third FOC above tell us that $x_3/x_2 = 2p_2/p_3$. Hence (remember $x_1 = 0$ now) $m = p_2x_2 + 2p_2x_2$ and $x_2 = m/(3p_2)$ while $x_3 = 2m/(3p_3)$. So we get

$$x(p, m) = \begin{cases} \left(\frac{m}{p_1} - 3, \frac{p_1}{p_2}, \frac{2p_1}{p_3} \right) & \text{if } m > 3p_1 \\ \left(0, \frac{m}{3p_2}, \frac{2m}{3p_3} \right) & \text{if } m \leq 3p_1 \end{cases}$$

Question 8: Here we have a pure exchange economy with 2 goods and 2 consumers. We can best represent this in an Edgeworth box.




Suppose x_1 is on the horizontal axis and x_2 is on the vertical, and let consumer B have the lower left hand corner as origin, consumer A the upper right hand corner. (I made this choice because I like to have the “harder” consumer oriented the usual way.) The dimensions of the box are 20 by 20 units. The first thing to do is to find the Pareto Set (the contract curve), since we know that any equilibrium has to be Pareto efficient. The MRS for person A is $4/3$, the MRS for person B is $3x_2/(4x_1)$. Therefore the Pareto Set is defined by $x_2/x_1 = 16/9$ (in person B 's coordinates.) This is a straight ray from B 's origin with a slope greater than 1, and therefore above the main diagonal. The Pareto set is this ray and the portion of the upper boundary of the box from the ray's intersection point to the origin of A . There now are 2 possibilities for the equilibrium. Either it is on the ray, and therefore must have a price ratio of $4/3$. Or it is on the upper boundary of the box, in which case the price ratio must be below $4/3$, but we know that B 's consumption level for good 2 is 20. In the first case we have 2 equations defining equilibrium. The ray, $x_2 = 16x_1/9$, and the budget line $(x_2 - 11) = 4(8 - x_1)/3$. From this we get $16x_1/9 - 11 = 32/3 - 12x_1/9$ and from that $28x_1/9 = 65/3$ and thus $x_1 = 195/28 < 20$. It follows that $x_2 = (16 \times 195)/(9 \times 28) = 780/63 = 260/21 < 20$. Since both of B 's consumption points are strictly within the interior of the box, we are done. All that remains is to compute A 's allocation. The equilibrium is therefore

$$(p^*, (x^A), (x^B)) = \left(\frac{4}{3}, \left(20 - \frac{195}{28}, 20 - \frac{780}{63} \right), \left(\frac{195}{28}, \frac{780}{63} \right) \right).$$

Question 9: Again we have a square Edgeworth box, 20×20 . Again I choose to put consumer B on the bottom left origin. B 's preferences are quasi-linear with respect to x_1 , A 's are piece-wise linear with slopes $4/3$ and $3/4$ which

meet at the kink line $x_2 = x_1$ (in A's coordinates!) which coincides with the main diagonal. The MRS for B's preferences is $3x_2/4$. We have Pareto optimality when $3x_2/4 = 4/3 \rightarrow x_2 = 16/9$ and when $3x_2/4 = 3/4 \rightarrow x_2 = 1$. So, the Pareto Set is the vertical axis from B's origin to $x_2^B = 1$, the horizontal line $x_2^B = 1$ to the main diagonal (the point $(1, 1)$ in other words), up the main diagonal to the point $(x_1^B, x_2^B) = (16/9, 16/9)$, from there along the horizontal line $x_2^B = 16/9$ to the right hand edge of the Box, and then up that border to A's origin.



By inspection, the most likely candidate for an equilibrium is a price ratio of $4/3$ with an allocation on the second horizontal line segment. Let us attempt to solve for it. First, the budget equation (in B's coordinates) is $3(x_2 - 11) = 4(8 - x_1)$. Second, we are presuming that $x_2 = 16/9$. So we get $16/3 - 33 = 32 - 4x_1$, or $x_1 = 14\frac{11}{12}$. Since this is less than 20 we have found an interior point and are done. The equilibrium is

$$(p^*, (x^A), (x^B)) = \left(\frac{4}{3}, \left(5\frac{1}{12}, \frac{164}{9} \right), \left(14\frac{11}{12}, \frac{16}{9} \right) \right).$$

Question 10: To prove this we need to show the implication in both directions: (\Leftarrow) : Suppose $x \succ y$. Then $\exists B, x, y \in B$ with the property that $x \in C(B), y \notin C(B)$. Consider all other $B' \in \mathcal{B}$ with the property that $x, y \in B'$. By the Weak Axiom $\nexists B'$ with $y \in C(B')$ since otherwise the set B would violate the weak axiom (applied to the choice y with the initial set B' .) Therefore $x \succ^* y$.

(\Rightarrow) : Let $x \succ^* y$. The first part of the definition requires $\exists B, x, y \in B, x \in C(B)$. By the weak axiom there are two possibilities: either all $B' \in \mathcal{B}$ with $x, y \in B'$ have $\{x, y\} \in C(B')$ or none have $y \in C(B')$. The second part of

the definition requires us to be in the second case, but then $y \notin C(B)$, and so $x \succ y$.

If the WA fails a counter example suffices: Let $X = \{x, y, z\}$, $\mathcal{B} = \{\{x, y\}, \{x, y, z\}\}$, $C(\{x, y\}) = \{x\}$, $C(\{x, y, z\}) = \{x, y\}$. This violates the WA. $C(\{x, y\}) = \{x\}$ demonstrates that $x \succ y$ by definition. On the other hand it is not true that $x \succ^* y$ (let $B = \{x, y\}$ and $B' = \{x, y, z\}$ in the definition of \succ^*).

Question 11:

a) This is another 20 by 20 box, with the endowment in the centre. Suppose B's origin on the bottom left, A's the top right. As in question 8, A's indifference curves have a constant MRS of α and are perfect substitute type. B's ICs have a MRS of $\beta x_2/x_1$ and are Cobb-Douglas. The contract curve in the interior must have the MRSs equated, so it occurs where (in B's coordinates) $x_2/x_1 = \alpha/\beta$. This is a straight ray from B's origin and depending on the values of α and β it lies above or below the main diagonal. Since these cases are (sort of) symmetric we pick one, and assume that $\alpha/\beta > 1$. The contract curve is this ray and then the part of the upper edge of the box to A's origin.

As in question 8 there are two cases for the competitive equilibrium. Either it occurs on the part of the Contract curve interior to the box, or it occurs on the boundary of the box. In the first case the slope of the budget and hence the equilibrium price must be α , since both MRSs have that slope along the ray and in equilibrium the price must equal the MRS. Note that the budget now coincides with A's indifference curve through the endowment point. The equilibrium allocation is determined by the intersection of the contract curve and this budget/IC. So we have two equations in two unknowns:

$$\alpha = \frac{x_2 - 10}{10 - x_1} \quad \text{and} \quad x_2 = \frac{\alpha}{\beta}x_1.$$

Hence $\alpha(10 - x_1) = \alpha x_1/\beta$ or $\alpha\beta 10 = x_1(\alpha + \alpha\beta)$ and thus $x_1 = \beta 10/(1 + \beta)$ and $x_2 = \alpha 10/(1 + \beta)$. These are the consumption levels for B. A gets the rest. The equilibrium thus would be $(p^*, (x_1^A, x_2^B), (x_1^B, x_2^B)) =$

$$\left(\alpha, \left(10 \frac{2 + \beta}{1 + \beta}, 10 \frac{2(1 + \beta) - \alpha}{1 + \beta} \right), \left(\frac{\beta 10}{1 + \beta}, \frac{\alpha 10}{1 + \beta} \right) \right)$$

which only makes sense if the allocation indeed is interior, that is, as long as $10\alpha/(1 + \beta) < 20$, or $(\alpha - \beta) < (2 + \beta)$.

If that is not true we find ourselves in the other case. In that case we know that we are looking for an equilibrium on the upper boundary of the box and thus know that $x_2^B = 20$ while $x_2^A = 0$. It remains to determine p and the allocations for good 1. At the equilibrium point the budget must be flatter

than A's IC (so that A chooses to only consume good 1). The allocation must also be the optimal choice for B and hence the budget must be tangent to B's IC, since for B this is an interior consumption bundle (interior to B's consumption set, that is.) So we again have to solve two equations in 2 unknowns:

$$\frac{\beta c_2}{c_1} = p \quad \text{and} \quad p = \frac{c_2 - 10}{10 - c_1} \quad \text{while} \quad c_2 = 20.$$

It follows that $\beta 20(10 - c_1) = 10c_1$ and therefore $c_1 = 20\beta/(1 + 2\beta)$. A gets the rest. The equilibrium is therefore

$$(p^*, (x_1^A, x_2^A), (x_1^B, x_2^B)) = \left(1 + 2\beta, \left(20 \frac{1 + \beta}{1 + 2\beta}, 0 \right), \left(\frac{\beta 20}{1 + 2\beta}, 20 \right) \right).$$

b) All endowments above and to the right of the line $x_2 = 40 - 2x_1$ in B's coordinates will lead to a boundary equilibrium. All those on this line and below will lead to an interior equilibrium with $p = 2$.

Question 12:

a) The social planner's problem is

$$\max_l \left\{ \ln(4\sqrt{16-l}) + \frac{1}{2}\ln(l) \right\}$$

which has first order condition

$$-\frac{1}{4\sqrt{16-l}} \frac{2}{\sqrt{16-l}} + \frac{1}{2l} = 0.$$

Hence $16 - l = l$ and so $l^* = 8$, $x^* = 8$, $c^* = 8\sqrt{2}$.

b) Since the consumer's problem requires profits, we solve for the firm first. $\max_x \{p4\sqrt{x} - wx\}$ has FOC $2p/\sqrt{x} = w$ and leads to firm labour demand of $x(p, w) = 4p^2/w^2$, consumption good supply of $c(p, w) = 8p/w$, and profits of $\pi(p, w) = 4p^2/w$.

The consumer will

$$\max_{c,l} \left\{ \ln c + \frac{1}{2}\ln l + \lambda(16w + \pi(p, w) - pc - wl) \right\}$$

which has first order conditions $1/c - \lambda p = 0$; $1/(2l) - \lambda w = 0$; $16w + \pi(p, w) = pc + wl$. The first 2 imply that $pc = 2lw$. Substituting into the third and using the profits computed above yields demand of $c(p, w) = 32w/(3p) + 8p/(3w)$ and leisure demand of $l(p, w) = 16/3 + 4p^2/(3w^2)$.

We can now solve for the equilibrium price ratio. Take any one market and set demand equal to supply. For the goods market this implies $32w/(3p) +$

$8p/(3w) = 8p/w$, and hence $p^2/w^2 = 2$. Substituting into the demands and supplies this gives $l^* = 8$, and hence all values are the same as in the social planner's problem in part a). You may want to verify that you could have solved for the price ratio from the labour market.

The complete statement of the general equilibrium is: The equilibrium price ratio is $p/w = \sqrt{2}$, the consumer's allocation is $(c, l) = (8\sqrt{2}, 8)$, and the firm produces $8\sqrt{2}$ units consumption good from 8 units input. Note that we cannot state profits without fixing one of the prices. So let $w = 1$ (so that we use labour as numeraire), then $p = \sqrt{2}$ and profits are 8.

7.2 Chapter 3

Question 1:

a) Zero arbitrage means that whichever way I move between periods, I get the same final answer. In particular I could lend in period 1 to collect in period three, or I could lend in period 1 to period 2, and then lend the proceeds to period 3. Hence the condition is

$$(1 + r_{12})(1 + r_{23}) = (1 + r_{13}).$$


Note that if we were to treat r_{13} not as a simple interest rate but as a compounding one, we'd get $(1 + r_{12})(1 + r_{23}) = (1 + r_{13})^2$ instead.

b) You have to adopt one period as your viewpoint and then put all other values in terms of that period (by discounting or applying interest). With period 3 as the viewpoint I use period 3 future values for everything:

$$B = \{(c_1, c_2, c_3) | (1 + r_{13})c_1 + (1 + r_{23})c_2 + c_3 = (1 + r_{13})m_1 + (1 + r_{23})m_2 + m_3\}$$

Note that any other viewpoint is equally valid. The restriction in (a) means that it does not matter which interest rate I use to compute the forward value of c_1 , say. Indeed, without that restriction I would get an infinite budget if it is possible to borrow infinite amounts. With some borrowing constraints in place I would have to compute the highest possible arbitrage profits for the various periods and compute the resulting budget.

c) This is a standard downward sloping budget line in (c_2, c_3) space with a slope of $-(1 + r_{23})$. It does not necessarily have to go through the endowment point (m_2, m_3) , however. It will be below that point if $c_1 > m_1$ and above that point if $c_1 < m_1$.

**Question 2:**

a) The easy way to get this is to first ignore the technology. The market budget is a straight line with a slope of -1 through the point $(100, 100)$, which is truncated at the point $(160, 40)$, where the budget becomes vertical. Note that the gross rate of return is 1 since the interest rate is 0. Now consider the technology and the implications of zero arbitrage: Joe can move consumption from period 1 to period 2 in two ways, via the financial market, or via “planting”. Both must yield the same gross rate of return at the optimum (why? we know that at the optimum of a maximization problem the last dollar allocated to each option must yield the same marginal benefit.) The gross rate of return at the margin is nothing but the marginal product of the technology, however. So, compute the MP ($5/\sqrt{x_1}$) and find the investment level at which the MP is 1.

$$\frac{5}{\sqrt{x_1}} = 1 \quad \longrightarrow \quad 5 = \sqrt{x_1} \quad \longrightarrow \quad x_1 = 25.$$

At optimal use at an interior optimum Joe invests 25 units (and collects 50 in the next period.) This means that from any point on the financial market budget Joe can move left 25 and up 50. So that gives a straight line with slope -1 which starts at $(0, 225)$ and goes to $(135, 90)$. After this point there is a corner solution in technology choice: Joe cannot use the market any more. The technology therefore may give a higher return than the market. So the budget follows the (flipped over to the left) technology frontier down to the point $(160, 40)$, and down to $(160, 0)$ from there.

b) First simplify the preferences (this step is not necessary!). Applying a natural logarithm gives the function $\hat{U}(c_1, c_2) = c_1^4 c_2^6$ which represents the same preferences. Applying the 10th root gives $\tilde{U}(c_1, c_2) = c_1^4 c_2^6$ which also represents the same preferences and is recognized as a Cobb-Douglas. Now you can either compute the MRS ($2c_2/3c_1$) and set that equal to 1 (since most of the budget has a slope of -1 and we know that $MRS = Slope$ at the optimum.) That gives you two equations in two unknowns, and we can solve:

$$c_2 = 225 - c_1, \quad c_2 = 3c_1/2 \quad \longrightarrow \quad 450 = 5c_1 \quad \longrightarrow \quad c_1 = 90, \quad c_2 = 135.$$

We then double check that the assumption that we are on the -1 sloped portion of the budget was correct, which it is (by inspection.) Or you could use the demand function for C-D, so you know

$$(c_1, c_2) = \left(\frac{.4M}{p_1}, \frac{.6M}{p_2} \right) = \left(\frac{.4 \times 225}{1+0}, \frac{.6 \times 225}{1} \right) = (90, 135).$$

Now this is his final consumption bundle. In order to get there he invested 25 units, so on the “market budget” line he must have started at $(115, 85)$, and that required him to borrow 15 units.

In summary, he borrows 15, giving him 115, of which he invests 25, so he has 90 left to consume. In the next period he gets 100 from his endowment, 50 from the investment, for a total of 150, of which he has to use 15 to pay back the loan, so he can consume 135!

Question 3: I will not draw the diagram but describe it. You should refer to a rough diagram while reading these solutions to make sense of them.

a) The indifference curves have two segments with a slope of -1.3 and -1.2 respectively. The switch (kink) occurs where

$$23 \left(\frac{12}{10}c_1 + c_2 \right) = 22 \left(\frac{13}{10}c_1 + c_2 \right) \quad \rightarrow \quad c_2 = (22 \times 13 - 23 \times 12)c_1/10 = c_1.$$

b) Note that the budget has a slope of 1.25 which is less than 1.3 and more than 1.2, so she consumes at the kink. Thus she is on the kink line and the budget:

$$c_1 = c_2 \text{ and } -1.25 = \frac{c_2 - 8}{c_1 - 9} \quad \rightarrow \quad c_1 = c_2 = 77/9.$$

c) Again she consumes at the kink, so

$$c_1 = c_2 \text{ and } -1.25 = \frac{c_2 - 12}{c_1 - 5} \quad \rightarrow \quad c_1 = c_2 = 73/9.$$

d) Here we need to work back. Note that at the slopes implied by the interest rates she continues to consume at her kink line. The reason is that both 1.25 and 1.28 are bigger than 1.2, the slope of her lower segment, but less than 1.3, the slope of the steep segment. Hence optimal consumption is at the kink and she borrows if she has less period 1 endowment than period 2 endowment. She lends money if she has larger period 1 endowment than period 2 endowment. So for all endowments above the main diagonal she is

a borrower, for all endowments below a lender.

e) Now she never trades. To lend money the budget slope is 1.18 which is less than either of her IC segment slopes. She would not want to lend ever at this rate no matter what her endowment. On the other hand, suppose she were to borrow. The budget slope is 1.32 which is steeper than even her steepest IC segment. She would not borrow. Thus she remains at the kink in her budget (the endowment point) no matter where it is.

Question 4: Again I will not draw the diagram but describe it. You should refer to a rough diagram while reading these solutions to make sense of them.

a) This is a 20 by 10 box. Suppose A's origin on the bottom left, B's the top right. A's indifference curves have a MRS of $c_2/(\alpha c_1)$ and are nice C-D type curves. B's ICs have a MRS of $1/\beta$ and are straight lines. The contract curve in the interior must have the MRSs equated (from Econ 301: for differentiable utility functions an interior Pareto optimum has a tangency), so it occurs where $c_2/c_1 = \alpha/\beta$. This is a straight ray from A's origin and depending on the values of α and β it lies above or below the main diagonal. Since these cases are (sort of) symmetric we pick one, and assume that $\alpha/\beta > 1/2$. The contract curve is this ray and then the part of the upper edge of the box to B's origin.

b) There are two cases, either the Contract curve ray is shallow enough that the equilibrium occurs on it, or it is so steep that the equilibrium occurs on the top boundary of the box. In the first case the slope of the budget and hence the equilibrium price must be $1/\beta$, since both MRSs have that slope along the ray and in equilibrium the price must equal the MRS. So the equilibrium interest rate is $(1 - \beta)/\beta$. Note that the budget now coincides with player B's indifference curve through the endowment point. Hence the ray of the contract curve must intersect that, and it does so only if it intersects the top boundary to the right of the intersection of B's IC with the boundary. The latter occurs at $(4(3 - \beta), 10)$. The former occurs at $(10\beta/\alpha, 10)$. So the interior solution obtains if $10\beta/\alpha > 4(3 - \beta)$, or if $10\beta > 12\alpha - 4\alpha\beta$. In that case the equilibrium allocations are derived by solving the intersection of the budget and the ray:

$$c_2 = \alpha c_2/\beta \text{ and } -\frac{1}{\beta} = \frac{c_2 - 6}{c_1 - 12} \quad \rightarrow \quad c_1^A = \frac{6(2 + \beta)}{1 + \alpha}, c_2^A = \frac{\alpha 6(2 + \beta)}{\beta(1 + \alpha)}.$$

B gets the remainder.

In the other case, when the ray fails to intersect B's IC, we know that we are looking for an equilibrium on the upper boundary of the box (so $c_2^A = 10$ and $c_2^B = 0$.) At this point we must have a budget flatter than B's IC (so that B chooses to only consume good 1). It must also be tangent to

A's IC, since for player A this is an interior consumption bundle (interior to his consumption set, that is.) So we require $1 + r = 10/(\alpha c_1)$ to have the tangency, and we require $1 + r = (10 - 6)/(12 - c_1)$ in order to be on the budget line. These are two equations in two unknowns again, so we solve: $c_1^A = 60/(5 + 2\alpha)$ and $r = (5 - 4\alpha)/(6\alpha)$. B gets the rest of good 1, of course.

7.3 Chapter 4

Question 1: We wish to show that for any concave $u(x)$

$$\frac{1}{3}u(24) + \frac{1}{3}u(20) + \frac{1}{3}u(16) \geq \frac{1}{2}u(24) + \frac{1}{2}u(16).$$

We can do the following: first bring the $u(24)$ and $u(16)$ to the RHS:

$$\frac{2}{6}u(20) \geq \frac{1}{6}u(24) + \frac{1}{6}u(16).$$

Then multiply both sides by 3:

$$u(20) \geq \frac{1}{2}u(24) + \frac{1}{2}u(16).$$

The LHS of this represents a certain outcome of 20, the RHS a lottery with 2 equally likely outcomes. Now note that

$$\frac{1}{2}24 + \frac{1}{2}16 = 20.$$

That is, the expected value of the lottery on the RHS of the last inequality above is equal to the expected value of the degenerate lottery on the LHS. Therefore this penultimate inequality must be true, since it coincides with the definition of a risk averse consumer. (utility of expectation greater than expectation of utility.)

Question 2: The certainty equivalent is defined by

$$U(CE) = \sum p_i u(x_i) = \int u(x) dF(x).$$

Using the particular function we are given:

$$\begin{aligned} \sqrt{CE} &= \alpha\sqrt{3600} + (1 - \alpha)\sqrt{6400} \\ CE &= (\alpha 60 + (1 - \alpha)80)^2 = (80 - 20\alpha)^2. \end{aligned}$$

Note that the expected value of the gamble is $E(w) = \alpha 3600 + (1 - \alpha)6400 = 6400 - \alpha 2800$ and thus the maximal fee this consumer would pay for access to fair insurance would be the difference $E(w) - CE = 400\alpha(1 - \alpha)$.

Question 3: The coefficient of absolute risk aversion is defined as $r_A = -u''(w)/u'(w)$. Computing this for both functions we get

$$u(w) = \ln w \longrightarrow r_A = \frac{1}{w}; \quad u(w) = 2\sqrt{w} \longrightarrow r_A = \frac{1}{2w}.$$

Therefore the two consumers exhibit equal risk aversion if the second consumer has half the wealth of the first. Their relative risk aversion coefficients (defined as $-u''(w)w/u'(w)$) are 1 and 1/2, respectively. That means that while, if the logarithm consumer has twice the wealth as the root consumer, he will have the same attitude towards a fixed dollar amount gamble, he will be more risk averse with respect to a gamble over a given proportion of wealth. (Note that the two statements don't contradict one another: a \$1 gamble represents half the percentage of wealth for a consumer with twice the wealth!)

Question 4: Here we need an Edgeworth Box diagram, which is a square, 15 units a side. Suppose we have consumer A on the bottom left origin (B then goes top right). Suppose also that we put state R on the horizontal. Note that the certainty line is the main diagonal of the box! This observation is crucial, since it means that there is no aggregate risk!

General equilibrium requires that demand is equal to supply for each good, but we can't find those here (not knowing the consumers' tastes), so it is not useful information. But we also know that in general equilibrium the price ratio must equal each consumer's MRS (since GE is Pareto optimal and that requires MRSs to be equalized, at least for interior allocations.) Note that the two MRSs here are

$$MRS_A = \frac{\pi u'_A(c_R^A)}{(1 - \pi)u'_A(c_S^A)} \quad MRS_B = \frac{\pi u'_B(c_R^B)}{(1 - \pi)u'_B(c_S^B)}$$

On the certainty line (the main diagonal) $c_S^A = c_R^A$ and $c_S^B = c_R^B$, so $MRS_A = MRS_B = \pi/(1 - \pi)$. In other words, the certainty line for each consumer coincides and together they are the set of Pareto optimal points.

Hence the equilibrium price ratio must be $p^* = \pi/(1 - \pi)$.

The allocation is now easily computed: we know the price ratio and the endowment, hence the budget line for the consumers. We also know that

consumption is equal in both states. So

$$\frac{\pi}{1 - \pi} = \frac{c_S^A - 5}{10 - c_R^A} = \frac{c^A - 5}{10 - c^A} \rightarrow c_S^A = c_R^A = 5(1 + \pi)$$

and since $c_i^B = 15 - c_i^A$ we get $c_S^B = c_R^B = 5(2 - \pi)$.

Question 5: **a)** $\max_x \{0.5u(10000(1 + 0.8x)) + 0.5u(10000(1.4 - 0.8x))\}$

b) The FOC for this is

$$0.5 \times 0.8u'(10000(1 + 0.8x)) - 0.5 \times 0.8u'(10000(1.4 - 0.8x)) = 0$$

$$\text{implies :} \quad u'(10000(1 + 0.8x)) = u'(10000(1.4 - 0.8x))$$

$$\text{implies :} \quad 10000(1 + 0.8x) = 10000(1.4 - 0.8x)$$

since she is risk averse. It follows that $1 + .8x = 1.4 - .8x$, and therefore that $1.6x = 0.4$, so that $x = 0.25$. One quarter, or 25% are invested in gene technology.

Question 6:

i) Denote the probability with which a ticket wins by π and the prize by P . A fair price for this lottery ticket would have to be a fraction p per dollar of prize such that $\pi(P - pP) - (1 - \pi)pP = 0$, or $p = \pi$. Let us start with this as a benchmark case (we know that normally such a lottery would not be accepted.) Utility maximization requires that for a gambling consumer $v(w_0) \leq \pi v(w_0 + (1 - p)P) + (1 - \pi)v(w_0 - pP) + \mu_i$. Thus all consumers for whom $\mu_i \geq v(w_0) - \pi v(w_0 + (1 - p)P) - (1 - \pi)v(w_0 - pP)$ purchase a ticket. At a fair gamble this is

$$\begin{aligned} \mu_i &\geq v(w_0) - \pi v(w_0 + (1 - \pi)P) - (1 - \pi)v(w_0 - \pi P) \\ &> v(w_0) - v(\pi(w_0 + (1 - \pi)P) + (1 - \pi)(w_0 - \pi P)) \\ &= v(w_0) - v(w_0) \end{aligned}$$

(the second strict inequality follows from the definition of risk aversion). Clearly a strictly positive μ is required. Can the government make money on this? Well, assume that the price p above is fair ($p = \pi$) and let there be an additional charge of q . Now all consumers gamble for whom $\mu_i \geq v(w_0) - \pi v(w_0 + (1 - \pi)P - q) - (1 - \pi)v(w_0 - \pi P - q)$. While such a μ_i is larger than before, it exists (for small q in any case) as long as things are sufficiently smooth and the μ_i go that high. Note that those who gamble have a high utility for it (a high taste parameter μ_i) in this setting. Note that this implies that even though they lose money on average they have a higher

welfare. (The anti-gambling arguments in public policy debates therefore come in two flavours: (i) your gambling is against my (religious) beliefs, and thus it ought to be banned, (ii) there are externalities: your lost money is really not yours but should have bought a lunch for your child/spouse/dog. Since your child/spouse/dog can't make you stop, we will on their behalf.)

ii) Now μ is fixed. Of course, the decision to gamble will still depend on the same inequality, namely

$$\mu > v(w_0) - \pi v(w_0 + (1 - \pi)P - q) - (1 - \pi)v(w_0 - \pi P - q).$$

We thus can translate this question into the question of how the right hand side depends on w_0 and how this dependency relates to the different behaviours of risk aversion with wealth. So, is the right hand side increasing or decreasing with wealth, and is this a monotonic relationship? The right hand side is related, of course, to the utility loss from going to the expected utility from the expected value (ignoring q for a minute.) Intuitively, we would expect the difference to be declining in wealth for constant absolute risk aversion: Constant absolute risk aversion implies a constant difference between the expected value and the certainty equivalent.¹ Let this difference be the base of a right triangle. Orthogonal to that we have the side which is the required distance between the two utilities. The third side must have a declining slope as wealth increases since it is related to the marginal utility of wealth at the certainty equivalent, which is declining in wealth by assumption. There you go, I'd expect the utility difference must fall with wealth.

More formally, consider the original inequality again and approximate the RHS by its second order Taylor series expansion (that way we get first and second derivatives, which we want in order to form r_A :

$$\begin{aligned} v(w_0) - \pi v(w_0^\oplus) - (1 - \pi)v(w_0 - \pi P - q) &\approx \\ v(w_0) - \pi(v(w_0) + \oplus v'(w_0) + & \\ \oplus^2 v''(w_0)/2) - (1 - \pi)(v(w_0) - \ominus v'(w_0) + \ominus^2 v''(w_0)/2) & \\ = -\pi \oplus v'(w_0) - \pi \oplus^2 v''(w_0)/2(1 - \pi)(\ominus v'(w_0) - \ominus^2 v''(w_0)/2) & \\ = v'(w_0) [(1 - \pi) \ominus (1 - \ominus r_A/2) - \pi \oplus (1 - \oplus r_A/2)] & \end{aligned}$$

This looks more like it! Now note that we use \ominus and \oplus as positive quantities (which are not equal: \ominus is larger!) Furthermore we know that (a) this quantity must be positive and (b) that π is probably a very small number. Now, if r_A is constant then the term in brackets is constant, but of course

¹Is there a general proof for that? Note that constant r_A has for example the functional form $u(w) = -e^{-aw}$, for which the above is certainly true.

$v'(w)$ falls with w and thus the right hand side of our initial inequality (way above) falls. Any given μ is therefore more likely to be larger than it. Thus rich consumers participate, poor consumers don't if we have constant absolute risk aversion. If we have decreasing absolute risk aversion this effect is strengthened. Now, since relative risk aversion is just $r_A w$, it follows that constant relative risk aversion requires a decreasing absolute risk aversion, and that decreasing relative risk aversion requires an even more decreasing absolute risk aversion. Thus in all cases the rich gamble and the poor don't. (Note here that they are initially rich. Since they lose money on average they will become poor and stop gambling.)

iii) If $v(w) = \ln w$ then $v'(w) = 1/w$ and $v''(w) = -1/w^2$. Therefore $r_A = 1/w$, with $\partial r_A / \partial w < 0$, and $r_R = 1$. If $v(w) = \sqrt{w}$ then $v'(w) = 1/2\sqrt{w}$ and $v''(w) = -w^{3/2}/4$. Therefore $r_A = 1/(2w)$ and $r_R = 1/2$. We now know two pieces of information: the consumers' risk aversion to a given size gamble is declining with wealth. This would, *ceteris paribus* make them more likely to purchase the gamble for a constant μ (see above). But μ now is also declining with wealth. The final outcome therefore depends on what declines faster, and we can't make a definite statement. (As an aside note the following. Suppose we are talking stock market participation here. Then it might be reasonable to assume that the utility of participating in it is increasing in wealth, on average, and so we get higher participation by wealthier people. Now, if the stock market on average is a bad bet we get mean reversion in wealth, while if the stock market is on average more profitable than savings accounts etc we get the rich getting richer. If you now run a voting model where the mean voter wins, you get the desire to redistribute (i.e., tax the investing and profiting rich and give the cash to those who have a too high marginal utility of wealth to invest themselves.) Note also that progressive taxes reduce the returns of a given investment proportional to wealth, counteracting the above effect of more participation by wealthy individuals. ...)

See how much fun you can have with these simple models and a willingness to extrapolate wildly?)

Question 7: This question forms part of a typical incomplete information contracting environment. Here we focus only on the consumer's behaviour.

a) Assume that the worker has a contractual obligation to provide an effort level of E . Once he has signed the contract, however, he knows that his actual effort is not observable and thus would try to shirk. Expected utility is maximized for

$$e^* = \operatorname{argmax}\{\alpha\sqrt{w(E) - p} + (1 - \alpha)\sqrt{w(E) - e^2}\}.$$

The first order condition for this problem is $-2e = 0$ **if** $e \neq E$. I.e., given the worker shirks he will go all the way (after all, the punishment does not

depend on the severity of the crime in any way.) Thus we need to ensure that the worker will not shirk at all, which is the case if $\sqrt{w(E)} - E^2 \geq \alpha\sqrt{w(E) - p} + (1 - \alpha)\sqrt{w(E)}$, or $E^2 \leq \alpha(\sqrt{w(E)} - \sqrt{w(E) - p})$. If the wage function satisfies this inequality for all E , it will elicit the correct effort levels in all cases.

b) Now we have a potentially variable punishment. Given some job with contractual obligation E , the worker now will maximize expected utility and set

$$e^* = \operatorname{argmax}\{\alpha\sqrt{w(E) - p(E - e)} + (1 - \alpha)\sqrt{w(E)} - e^2\}.$$

The FOC for this problem is $\alpha p'(\cdot)(2\sqrt{w(E) - p(E - e)})^{-1} - 2e = 0$. (There are also second order conditions which need to hold!) This implies that the worker will play off the cost of shirking against the gains from doing so. We need to make sure that this equation is only satisfied for $e^* = E$, in which case he “voluntarily” chooses the contracted level. This clearly requires a positive $p'(\cdot)$. In particular, $\alpha p'(0)(2\sqrt{w(E)})^{-1} - 2E = 0$. Note: We could also vary the detection/supervision probability and make α depend on E . Then we get $e^* = \operatorname{argmax}\{\alpha(E)\sqrt{w(E) - p} + (1 - \alpha(E))\sqrt{w(E)} - e^2\}$. As in (a), if the worker deviates he will go all the way here. So the problem is similar to (a), only the wage schedule is now different since $\alpha(E)$ can also vary now. What this shows us is that we tend to want a punishment and a detection probability which both depend on the deviation from the correct level. (This is going to be a question about the technology available: some technologies may be able to detect flagrant shirking more readily than slight shirking.)

c) What this seems to indicate is that we would like to make punishments fit the crime. (So for example, if the punishment for a hold-up with a weapon is as severe as if somebody actually gets shot during it, then I might as well shoot people when I’m at it and I think that helps (and if it does not increase the effort the police put into finding me.)) Furthermore, if detection is a function of the actual effort level (the more you fudge the books the more likely will you be detected) then we need lower punishments, *ceteris paribus*, since the increasing risk will provide some disincentive to cheat anyways.

Question 8:

a) Let C_B denote the coverage purchased for bad losses, and C_M the coverage for minor losses. Zero profits imply that the premiums p_B and p_M for bad and minor losses, respectively, are $p_B = \pi/5$ and $p_M = 4\pi/5$. Hence the consumer’s expected utility maximization problem becomes

$$\max_{C_B, C_M} \left\{ (1 - \pi)u(W - p_M C_M - p_B C_B) + \pi \left(\frac{1}{5}u(W - p_M C_M - p_B C_B + C_B - B) + \right. \right.$$

$$\left. \frac{4}{5}u(W - p_M C_M - p_B C_B + C_M - M) \right\}$$

The first order conditions for this problem are

$$\begin{aligned} -p_M(1 - \pi)u'(n) - p_M \frac{\pi}{5}u'(b) + (1 - p_M) \frac{4\pi}{5}u'(m) &= 0 \\ -p_B(1 - \pi)u'(n) + (1 - p_B) \frac{\pi}{5}u'(b) - p_B \frac{4\pi}{5}u'(m) &= 0 \end{aligned}$$

Using the fair premiums this simplifies to

$$\begin{aligned} -(1 - \pi)u'(n) - \frac{\pi}{5}u'(b) + \left(1 - \frac{4\pi}{5}\right)u'(m) &= 0 \\ -(1 - \pi)u'(n) + \left(1 - \frac{\pi}{5}\right)u'(b) - \frac{4\pi}{5}u'(m) &= 0 \end{aligned}$$

Hence

$$\left(1 - \frac{4\pi}{5}\right)u'(m) - \frac{\pi}{5}u'(b) = \left(1 - \frac{\pi}{5}\right)u'(b) - \frac{4\pi}{5}u'(m)$$

and thus $u'(b) = u'(m)$, which finally implies that $u'(n) = u'(b) = u'(m)$ and therefore that

$$0 = C_B - B = C_M - M.$$

As expected, the consumer buys full insurance for each accident type separately.

b) Now only one coverage can be purchased, denote it by C , and will be paid in case of either accident. Zero profits imply that the premium p is $p = \pi$. Hence the consumer's expected utility maximization problem becomes

$$\max_C \left\{ (1 - \pi)u(W - pC) + \pi \left(\frac{1}{5}u(W - pC + C - B) + \frac{4}{5}u(W - pC + C - M) \right) \right\}$$

The first order condition for this problem is

$$-p(1 - \pi)u'(n) + (1 - p) \frac{\pi}{5}u'(b) + (1 - p) \frac{4\pi}{5}u'(m) = 0$$

Using the fair premium this simplifies to

$$u'(n) = \frac{1}{5}u'(b) + \frac{4}{5}u'(m)$$

and hence either

$$u'(b) > u'(n) > u'(m) \quad \text{or} \quad u'(b) < u'(n) < u'(m).$$

Thus either $W - B + (1 - \pi)C < W - \pi C < W - M + (1 - \pi)C$ or $W - B + (1 - \pi)C > W - \pi C > W - M + (1 - \pi)C$, but this implies either $-B + 1C < 0 < -M + 1C$ or $-B + 1C > 0 > -M + 1C$. Since $B > M$ by definition we obtain that $B > C > M$, the consumer over insures against minor losses, and is under insured against big losses.

Question 9: From Figure 3.7 in the text, we can take the consumers' budget line to be the line from the risk free asset point (the origin in this case) to a tangency with the efficient portfolio frontier. Now this tangency occurs where the margin is equal to the average, so that $\sqrt{\sigma - 16}/\sigma = (2\sqrt{\sigma - 16})^{-1}$. That means that the market portfolio has $2(\sigma - 16) = \sigma$ or $\sigma = 32$. Therefore $\mu = 4$. The slope of the portfolio line thus is $4/32$. For an optimal solution the consumer's MRS must equal the slope of the portfolio line. For the two consumers given the MRS is $\sigma/32$ and $\sigma/96$. Thus the optima are $\sigma = 4$, $\mu = 1/2$ and $\sigma = 12$, $\mu = 3/2$. As expected, the consumer with the higher marginal utility for the mean will have a higher mean at the same prices (and given that both have the same disutility from variance.)

Question 10: The asset pricing formula implies that the expected return of the insurance equals the expected risk-free return less a covariance term. If insurance has a lower expected return than the risk-free asset, this covariance term must be positive. In the denominator we have the expected marginal utility, guaranteed to be positive. Thus the numerator must be positive. This means that $Cov(u'(w), R_i) > 0$. But since $u''(w) < 0$ this implies that the covariance between w and R_i is negative, that is, if wealth is low the return to the policy is high, if wealth is high, the return to the policy is low. That of course is precisely the feature of disability insurance which replaces income from work if and only if the consumer is unable to work.

Question 11:

1) False. The second order condition would indicate a minimum as demonstrated here: $\max_C \{ \pi u(w - L - pC + C) + (1 - \pi)u(w - pC) \}$ has FOC $\pi(1 - p)u'(w - L + (1 - p)C) - p(1 - \pi)u'(w - pC) = 0$. The second order condition for a maximum is $\pi(1 - p)^2 u''(w - L + (1 - p)C) + p^2(1 - \pi)u''(w - pC) \leq 0$. Note that $1 \geq \pi, p \geq 0$, so that the SOC requires $u''(\cdot)$ to be negative for at least one of the terms. A risk-lover has, by definition, $u''(\cdot) > 0$.

2) Uncertain. We can draw 2 diagrams to demonstrate. In both we have two intersecting budget lines, one steeper, one flatter. The flatter one corresponds to the initial situation. They intersect at the consumer's endowment. Since the consumer is a borrower, the initial consumption point is below and

to the right of the endowment on the initial budget. The indifference curve through this point is tangent to this budget. It may, however, cut the new budget (so that the IC tangent to the new budget represents a higher level of utility) or lie everywhere above it (in which case utility falls.)

3) True. Apply the following positive monotonic transformations to the first function: -2462 , $\times 12$, collect terms in one logarithm, take exponential, take the 9000th root. What you get is the second function.

4) True. A risk averse consumer is defined as having $u(\int xg(x)dx) > \int u(x)g(x)dx$. Let the consumer have initial wealth w and suppose he could participate in a lottery which leads to a change in his initial wealth by x , distributed as $f(x)$. Suppose the payment for this lottery is p . If this payment is equal to the expected value of the lottery then the consumer will not have a change in expected wealth, but will face risk. Thus by definition he would not buy this lottery. If the payment is less, then the expected value of wealth from participating in the lottery exceeds the initial wealth. Depending on by how much, the consumer may purchase. A risk loving consumer, of course, would already buy at when the expected net gain is zero. (This argument could be made more precise, and you should try to put it into equations!)

5) False. The market rate of return is 15%. Gargleblaster stock has a rate of return of $(117 - 90)/90 = 30\%$. This violates zero arbitrage.

6) True. All consumers face the same budget line in mean-variance space. At an interior optimum (and assuming their MRS is defined) they all consume on this line where the tangency to their indifference curve occurs. This may be anywhere along the line, depending on tastes, but the slope is dictated by the market price for risk.

Question 12:

a) Since workers work as bus driver and at a desk job we require

$$2\sqrt{40000} = \alpha 2\sqrt{44100 - 11700} + (1 - \alpha)2\sqrt{44100}$$

Therefore

$$\alpha = \frac{\sqrt{44100} - \sqrt{40000}}{\sqrt{44100} - \sqrt{32400}} = \frac{210 - 200}{210 - 180} = \frac{1}{3}.$$

b) Since workers work on oil rigs and at a desk job we require

$$2\sqrt{40000} = 0.5 \times 2\sqrt{122500 - Loss} + 0.5 \times 2\sqrt{122500}.$$

Thus $400 = \sqrt{122500 - Loss} + 350$ and hence $50 = \sqrt{122500 - Loss}$ or $Loss = 120000$.

c) At fair premiums the workers will fully insure. That is, they suffer their expected loss for certain. For a bus driver the expected loss is $11700/3 = 3900$. Thus the bus driver wage must satisfy $2\sqrt{40000} = 2\sqrt{w - 3900}$ and

hence it is \$43900. For the oil rig worker the expected loss is \$60000, and their wages will fall to \$100000 under workers compensation. Note the the condition that workers take all jobs together with a fixed desk job wage fixes the utility level in equilibrium for workers. However, the wage premium for risky jobs will not have to be paid: the wages of the risky occupations fall which benefits the firms in those industries by lowering their wage costs. (This is why industries are in favour of workers' compensation.)

d) The average probability of an accident now is $0.4 \times 0.5 + 0.6 \times 1/3 = 0.4$. If we were to use this as a fair premium (but see below!) this premium is too high for bus drivers, who will under insure, and too low for oil rig workers, who will over insure. Indeed, the bus drivers will choose to buy insurance C_b such that $6\sqrt{44100 - 0.4C_b} = 8\sqrt{32400 + .6C_b}$ (Take the first order condition for $\max_{C_b} \{(1/3)\sqrt{32400 + 0.6C_b} + (2/3)\sqrt{44100 - 0.4C_b}\}$, bring the \sqrt from the denominator into the numerators and loose the 1/30 on both sides.) Thus we require $9(44100 - 0.4C_b) = 16(32400 + 0.6C_b)$, or $C_b = 10(9 \times 44100 - 16 \times 32400)/(6 \times 16 + 4 \times 9) = -9204$. What does this mean? It means that the bus drivers would like to bet on themselves having an accident buying negative amounts of insurance! (The ultimate in under insurance!) Note that the governments expected profit from bus drivers is $-0.4 \times 9204/3 + 1.2 \times 9204/3 = 2454.40 > 0$.

The oil rig workers would need to solve $\max_{C_o} \{\sqrt{2500 + 0.6C_o} + \sqrt{122500 - 0.4C_o}\}$, which leads to $3\sqrt{122500 - 0.4C_o} = 2\sqrt{2500 + 0.6C_o}$ and thus $C_o = 182083.33$. Note that the govt loses money on them, since $(0.5 \times 0.4 - 0.5 \times 0.6) \times 182083.33 = -18208.30$.

Overall then the govt makes losses of $5810.68N$, where N is the number of workers in risky occupations. At the old wages both groups are better off (and thus there would be an influx of desk workers and a reallocation towards oil rigs.) In order to break even the insurance rates would have to be changed, in particular raised. It also seems that the govt would ban the purchase of negative insurance amounts. In which case the bus drivers would find it optimal to buy no insurance, and then premiums would have to be 0.5 for the govt to break even. This would be deemed unjust by all involved, and so in practice the govt forces all workers to buy a fixed amount of insurance!

In principle we could compute equilibrium wages if we treat the insurance purchase as a function of the wage. So, for example we know from the above that $C_b(w) = 10(9 \times w - 16 \times (w - 11700))/(6 \times 16 + 4 \times 9)$. We then solve for $s\sqrt{40000} = (2/3)\sqrt{w + (1 - 0.4)C_b(w)} + (4/3)\sqrt{w - 11700 - 0.4C_b(w)}$. The details are left to the reader.

The important point here is that it is important to charge the correct premiums. If that is not done things will work out funny. That in turn leads to real life plans which do not allow a choice — workers have to insure, the amount is dictated (often capped, that is, the insured amount is a function of the wage up to a maximum.) You can see that such plans can be quite complicated and that it can be quite complicated to figure out who would want to do what, what the distributional implications are, etc.

Question 13: Let us translate the question into notation: We are to show that $\frac{u'(c_1)}{u'(c_2)} = k$ if $\frac{c_2}{c_1} = \lambda$ if the function $u(\cdot)$ satisfies $\frac{u''(w)w}{u'(w)} = a, \forall w$.

$$\frac{u'(c_1)}{u'(c_2)} = k \text{ and } \frac{c_2}{c_1} = \lambda \implies u'(c_1) = k(\lambda)u'(\lambda c_1).$$

If the left and right hand side of that last expression are identical functions, then their derivatives must equal: $u''(c_1) = k(\lambda)\lambda u''(\lambda c_1)$, but we know that $k(\lambda) = u'(c_1)/u'(\lambda c_1)$, so that

$$u''(c_1) = \frac{u'(c_1)}{u'(\lambda c_1)} \lambda u''(\lambda c_1) \implies \frac{u''(c_1)}{u'(c_1)} = \lambda \frac{u''(\lambda c_1)}{u'(\lambda c_1)}$$

Thus the MRS is constant for any consumption ratio λ if

$$\frac{u''(c_1)c_1}{u'(c_1)} = \frac{u''(\lambda c_1)\lambda c_1}{u'(\lambda c_1)} \quad \forall \lambda,$$

which is constant relative risk aversion.

7.4 Chapter 6

Question 1: A 3-player game in extensive form comprises a game tree, Γ , a payoff vector of length three for each terminal node, a partition of the set of non-terminal nodes into player sets S_0, S_1, S_2, S_3 , a partition of the player sets S_1, S_2, S_3 into information sets. Further, a probability distribution for each node in S_0 over the set of immediate followers and for each S_i^j an index set I_i^j and a 1-1 mapping from I_i^j to the set of immediate followers of the nodes in S_i^j . Any carefully labelled game tree diagram will do. It does not even have to have nature (i.e., S_0 could be empty.)

Question 2: Perfect recall is when each player never forgets any of his own previous moves (so that for any two nodes within an information set one

may not be a predecessor of the other and any two nodes may not have a common predecessor in another information set of that player such that the arc leading to the nodes differs) and never forgets information once known (so that any two nodes in a player's information set may not have predecessors in distinct previous information sets of this player.) Counter examples as in the text, or any game which violates these requirements.

Question 3: Yes, any finite game has a Nash equilibrium, possibly in mixed strategies. This follows from the Theorem we have in the text. The game therefore will also have a SPE (they are a subset of the Nash equilibria, but keep in mind that the condition of subgame perfection may have no 'bite', in which case we revert to Nash.)

Question 4: Player 1 has no weakly dominated strategies since $u_1(D, L, Left) > u_1(U, L, Left)$ but $u_1(D, R, Left) < u_1(U, R, Left)$, while $u_1(C, L, Right) > u_1(U, L, Right)$. Player 3 does also not have a weakly dominated strategy. Depending on the opponents' moves he gets a higher payoff sometimes in the left and sometimes in the right matrix. Player 2 does have weakly dominated strategies: Both L and R are weakly dominated by C .

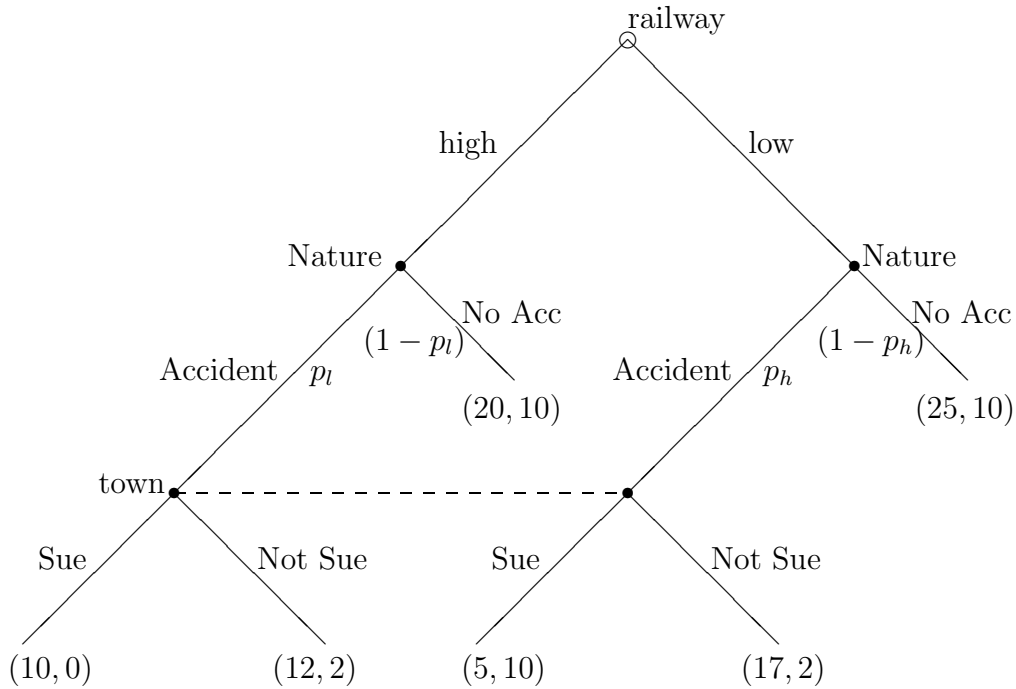
This does not leave us with a good prediction yet, aside from the fact that 2 can be argued to play C . However, if we now consider repeated elimination we can narrow down the answer to what is also the unique Nash equilibrium in pure strategies in this case, $(D, C, Right)$.

To find mixed strategy Nash we assign probabilities to the strategies for players, so let $\mu_1 = Pr(U), \mu_2 = Pr(C), \gamma_1 = Pr(L), \gamma_2 = Pr(R)$, and $\alpha = Pr(Left)$. We can then compute the payoffs for players for each of their pure strategies. So for example $u_1(U, \gamma, \alpha) = \alpha(\gamma_1 + 2\gamma_2 + (1 - \gamma_1 - \gamma_2) + (1 - \alpha)(2\gamma_1 + 4\gamma_2 + 2(1 - \gamma_1 - \gamma_2))$. We then can ask, when is player 1, say, actually willing to mix? Only if the payoff from the pure strategies in the support of the mixed strategy are equal, so that the player does not care.

Question 5: This one is made easier by the fact that strategy R is (strictly) dominated, so that it will never be used in any mixed strategy equilibrium (or indeed any equilibrium.) Hence this is really just a 2×2 matrix we need to consider. Let $\alpha = Pr(U)$ and $\beta = Pr(L)$, so that $Pr(C) = 1 - \alpha$ and $Pr(C) = 1 - \beta$. Then for player 1 to mix we require $\beta + 4(1 - \beta) = 3\beta + 2(1 - \beta)$, hence $2 = 4\beta$, and hence $\beta = 0.5$. So if player 2 mixes with this probability then player 1 is indifferent between his two strategies. Now look at player 2: For 2 to be indifferent between the two strategies L and C we require $4\alpha + 2(1 - \alpha) = 2\alpha + 3(1 - \alpha)$. Hence $3\alpha = 1$ and thus $\alpha = 1/3$. Thus the

mixed strategy Nash equilibrium is $((1/3, 2/3), (1/2, 1/2))$. Note that the game has no pure strategy Nash equilibria.

Question 6: This is a two player game (the court is not a strategic player and does not receive any payoffs.) The most natural extensive form for such a situation is probably as in the game tree on the next page.



Here it is important to note that $p_h > p_l$, reflecting the fact that if low care is taken the accident probability is higher. I have arbitrarily assigned payoffs which satisfy the description. High level of care costs the railway 5, accidents impose a cost of 8 on both parties, legal costs are 2 for each party.

Let us try and find the Nash equilibrium of this game. As an exercise let us first find the strategic form:


$R \backslash T$	<i>Sue</i>	<i>NotSue</i>
<i>high</i>	$(20 - 10p_l, 10 - 10p_l)$	$(20 - 8p_l, 10 - 8p_l)$
<i>low</i>	$(25 - 20p_h, 10)$	$(25 - 8p_h, 10 - 8p_h)$

Note that *low* strictly dominates *high* for the railway if $0.5 > (2p_h - p_l)$ while *high* strictly dominates *low* if $p_h - p_l > 5/8$. In those cases the

Nash equilibria are (low, Sue) and $(high, NotSue)$, respectively. Otherwise there will be a mixed strategy equilibrium. Let α be the probability with which the railway uses the high effort level. The town is indifferent iff its expected payoffs from the two strategies are the same, that is, if $10 - 10\alpha p_l = 10 - 8p_h + 8\alpha(p_h - p_l)$. This is the case if $\alpha = 4p_h / (4p_h + p_l)$. For lower α it prefers to *Sue*, for higher α it prefers to *NotSue*. Letting β denote the probability with which the town sues, the railway expects to receive $20 - 8p_l - 2\beta p_l$ from *high* and $25 - 8p_h - 12\beta p_h$ from *low*. It is indifferent if $\beta = (5 - 8(p_h - p_l)) / (12p_h - 2p_l)$. So the mixed strategy Nash equilibrium is

$$(\alpha, \beta) = \left(\frac{4p_h}{4p_h + p_l}, \frac{5 - 8(p_h - p_l)}{12p_h - 2p_l} \right) \quad \text{if } p_h - \frac{5}{8} > p_l > 2p_h - \frac{1}{2}.$$

Question 7: There are two ways to draw this game. We can have nature move first and then Romeo (who does not observe nature's move.) Or we can have Romeo move first, and then nature determines if the move is seen. The game tree for the first case is as drawn below.



A strategy vector in this game is $(s_R, (s_J^1, s_J^2, s_J^3))$. Subgames start at information sets J^1 and J^2 , the only other subgame is the whole tree. In the subgame perfect equilibrium Juliet therefore is restricted to (S, M, \cdot) . Let α denote Romeo's probability of moving *S*, and β Juliet's (in J^3 .) Romeo's (expected) payoff from *S* is $30p + (1-p)(30\beta + 1 - \beta)$ and his payoff from *M* is $50p + (1-p)(5\beta + 50(1 - \beta))$. The β for which he is indifferent is $(49 - 29p) / (74(1 - p))$. Note that this is increasing in p and that $\beta = 1$ if $p = 5/9$! Juliet has payoffs of $50\alpha + 5(1 - \alpha)$ and $\alpha + 30(1 - \alpha)$ from moving *S*

and M , respectively, in J^3 . Hence she is indifferent if $\alpha = 25/74$. Of course, pure strategy equilibria may also exist, and we get the SPE equilibria to be

$$\left(\frac{25}{74}, \left(S, M, \frac{49 - 29p}{74(1 - p)} \right) \right), (S, (S, M, S)), (M, (S, M, M)) \text{ if } p < \frac{5}{9}.$$

Note that $(S, (S, M, S))$ requires that $30 > 50p + 5(1 - p)$, or $p < 5/9$ also. $(M, (S, M, M))$ requires that $50 > 30p + (1 - p)$, or $p < 49/29$, which is always true. What if $p > 5/9$? In that case the equilibrium in which the outcome is coordination on S (preferred by Juliet) does not exist, and neither does the mixed strategy equilibrium. Hence the unique equilibrium if $p \geq 5/9$ is $(M, (S, M, M))$. Romeo can effectively insist on his preferred outcome.

Question 8: Each firm will

$$\max_{q_i} \left\{ \left(\sum_{j \neq i} q_j + q_i - 10 \right)^2 q_i - 0q_i \right\}.$$

The FOC for this problem is

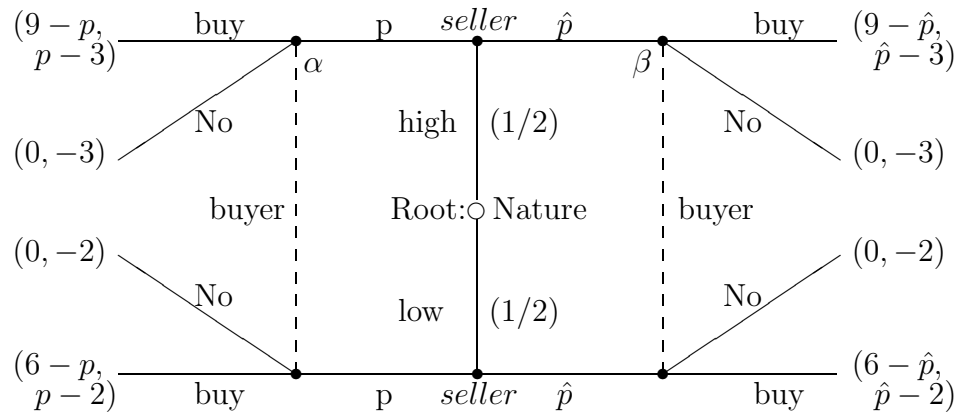
$$2 \left(\sum_{j \neq i} q_j + q_i - 10 \right) q_i + \left(\sum_{j \neq i} q_j + q_i - 10 \right)^2 = 0$$

Hence if $\sum_{j \neq i} q_j + q_i - 10 \neq 0$ we require

$$2q_i + \left(\sum_{j \neq i} q_j + q_i - 10 \right) = 0$$

and get the reaction function $q_i = (10 - \sum_{j \neq i} q_j)/3$. With identical firms we then know that in equilibrium $q_j = q_i$, so that $\sum_{j \neq i} q_j = (n - 1)q_i$. Hence $3q_i = 10 - (n - 1)q_i$ and we get that $q_i^* = 10/(n + 2) \forall i$. Total market output then is $10n/(n + 2)$. Note that total market output approaches 10 from below as n gets large. Market price for a given n is $400/(n + 2)^2$, which approaches zero as n gets large. (Note that the marginal cost is zero and hence the perfectly competitive price is zero!)

Question 9: In the first instance the sellers can only vary price. To clarify ideas, let us focus on two prices only (as would be needed for a separating equilibrium.) The game then is as depicted below. We are to show that no separating equilibrium exists. If it did, it would have to be the two prices as indicated, where one firm charges one price (presumably the high quality firm charging the higher price) the other another. But given that, the consumer knows (in equilibrium) which firm produced the product. It is easy to see that the low quality firm would deviate to the higher price (being then mistaken



for the high quality firm so that the consumer buys) since costs are unaffected by such a move, but a higher price is received.

At this point the remainder is non-trivial and left for summer study! The key is that the consumer now has an information set for each price-warranty pair, and that there are two nodes in it, one for each type of firm.

Question 10: What was not stated in the question was the fact that each consumer buys either one or no units. Each buyer purchases a unit of the good if and only if the price is below the valuation of the buyer. Hence total market demand is given by the number of buyers with a valuation above p , or $1 - F(p)$. $F(v)$ is the cumulative distribution for the uniform distribution on $[0, 2]$. Since the pdf for the uniform distribution on $[0, 2]$ is 0.5, we have $F(v) = \int_0^v 0.5 dt = 0.5v$. Hence market demand is $1 - 0.5p$ and inverse market demand is $2(1 - Q)$.

A Cournot equilibrium is nothing but a Nash equilibrium in the game in which firms simultaneously choose output levels. Hence firm 1 solves

$$\max_{q_1} \{2(1 - q_1 - q_2)q_1 - q_1/10\}$$

which leads to FOC $2(1 - q_1 - q_2) - 2q_1 - 1/10 = 0$ and the reaction function $q_1(q_2) = 19/40 - q_2/2$.

Firm 2 solves

$$\max_{q_2} \{2(1 - q_1 - q_2)q_2 - q_2^2\}$$

which leads to FOC $2(1 - q_1 - q_2) - 2q_2 - 2q_2 = 0$ and the reaction function $q_2(q_1) = 1/3 - q_1/3$.

The Nash equilibrium then is $(q_1, q_2) = (37/100, 21/100)$. Market price is $42/50$. Profits for the two firms are $74 \times 37/10000$ for firm 1 and $(82 \times 21 - 21^2)/10000$ for firm 2, so that joint profit is $(74 \times 37 + 82 \times 21 - 21^2)/10000$.

In the Stackelberg leader case we consider the SPE of the game in which firm 1 chooses output first and firm 2, after observing firm 1's output choice, picks its output level. Firm 1, the Stackelberg leader, therefore takes firm 2's reaction function as given. Thus firm 1 solves

$$\max_{q_1} \left\{ 2 \left(1 - q_1 - \left(\frac{1}{3} - \frac{q_1}{3} \right) \right) q_1 - q_1/10 \right\}$$

The FOC for this is $4(1 - 2q_1)/3 - 1/10 = 0$ and hence $q_1 = 37/80$. Thus $q_2 = 43/240$. Market price is $86/240$. Profits for the two firms are $(62 \times 37)/(240 \times 80)$ and $43 \times 43/240^2$. Joint profit thus is $(186 \times 37 + 43 \times 43)/240^2$.

Joint profit maximization would require that the firms solve

$$\max_{q_1, q_2} \left\{ 2(1 - q_1 - q_2)(q_1 + q_2) - q_1/10 - q_2^2 \right\}.$$

This has FOCs

$$\begin{aligned} 2(1 - q_1 - q_2) - 2(q_1 + q_2) - 1/10 &= 0 \\ 2(1 - q_1 - q_2) - 2(q_1 + q_2) - 2q_2 &= 0 \end{aligned}$$

so that we know that $2q_2 = 1/10$ or $q_2 = 1/20$. Hence $2(1 - q_1 - 1/20) - 2(q_1 + 1/20) - 1/10 = 0$ and $2 - 4q_1 - 3/10 = 0$ and $q_1 = 7/40$. Market price then is $2(31/40)$. Joint profits are $2(31/40)(9/40) - 8/400$. This cannot be attained as a Nash equilibrium because neither output level is on the firm's reaction function, and only output levels on the reaction function are, by design, a best response.