


Mathematical Methods for Economic Analysis*

Paul Schweinzer

School of Economics, Statistics and Mathematics
Birkbeck College, University of London
7-15 Gresse Street, London W1T 1LL, UK
Email: P.Schweinzer@econ.bbk.ac.uk
Tel: 020-7631.6445, Fax: 020-7631.6416



* This version (9th March 2004) is preliminary and incomplete; I am grateful for corrections or suggestions.

Contents

I	Static analysis	9
1	Mathematical programming	11
1.1	Lagrange theory	12
1.1.1	Non-technical Lagrangian	12
1.1.2	General Lagrangian	15
1.1.3	Jargon & remarks	17
1.2	Motivation for the rest of the course	18
2	Fundamentals	21
2.1	Sets and mappings	22
2.1.1	Sets	22
2.2	Algebraic structures	24
2.2.1	Relations	25
2.2.2	Extrema and bounds	28
2.2.3	Mappings	29
2.3	Elementary combinatorics	31
2.4	Spaces	32
2.4.1	Geometric properties of spaces	32
2.4.2	Topological properties of spaces	35
2.5	Properties of functions	41
2.5.1	Continuity	41
2.5.2	Differentiability	42
2.5.3	Integrability	43
2.5.4	Convexity, concavity	44
2.5.5	Other properties	45
2.6	Linear functions on \mathbb{R}^n	48
2.6.1	Linear dependence	48
2.6.2	Determinants	49
2.6.3	Eigenvectors and eigenvalues	50
2.6.4	Quadratic forms	52
3	First applications of fundamentals	55
3.1	Separating hyperplanes	55
3.2	The irrationality of playing strictly dominated strategies	57
3.3	The Envelope Theorem	60
3.3.1	A first statement	60

3.3.2	A general statement	63
3.4	Applications of the Envelope Theorem	65
3.4.1	Cost functions	65
3.4.2	Parameterised maximisation problems	67
3.4.3	Expenditure minimisation and Shephard's Lemma	68
3.4.4	The Hicks-Slutsky equation	69
3.4.5	The Indirect utility function and Roy's Identity	69
3.4.6	Profit functions and Hotelling's Lemma	69
4	Kuhn-Tucker theory	73
4.1	Theory	73
4.1.1	The Lagrangian	73
4.1.2	The extension proposed by Kuhn and Tucker	76
4.2	A cookbook approach	78
4.2.1	The cookbook version of the Lagrange method	79
4.2.2	The cookbook version of the Kuhn-Tucker method	79
4.2.3	A first cookbook example	80
4.2.4	Another cookbook example	81
4.2.5	A last cookbook example	82
4.3	Duality of linear programs	85
5	Measure, probability, and expected utility	89
5.1	Measure	89
5.1.1	Measurable sets	89
5.1.2	Integrals and measurable functions	95
5.2	Probability	98
5.3	Expected utility	102
6	Machine-supported mathematics	113
6.1	The programs	113
6.2	Constrained optimisation problems	118
7	Fixed points	121
7.1	Motivation	121
7.1.1	Some topological ideas	121
7.1.2	Some more details [†]	126
7.1.3	Some important definitions	127
7.2	Existence Theorems	129
7.2.1	Brouwer's Theorem	129
7.2.2	Kakutani's Theorem	131
7.2.3	Application: Existence of Nash equilibria	131
7.2.4	Tarski's Theorem	133
7.2.5	Supermodularity	136

II	Dynamic analysis	143
8	Introduction to dynamic systems	145
8.1	Elements of the theory of ordinary differential equations	146
8.1.1	Existence of the solution	148
8.1.2	First-order linear ODEs	149
8.1.3	First-order non-linear ODEs	151
8.1.4	Stability and Phase diagrams	152
8.1.5	Higher-order ODEs	156
8.2	Elements of the theory of ordinary difference equations	157
8.2.1	First-order (linear)	157
8.2.2	Second-order (linear) ODEs	159
8.2.3	Stability	162
9	Introduction to the calculus of variation	165
9.1	Discounting	166
9.2	Depreciation	166
9.3	Calculus of variation: Derivation of the Euler equation	167
9.4	Solving the Euler equation	169
9.5	Transversality condition	171
9.6	Infinite horizon problems	174
10	Introduction to discrete Dynamic Programming	177
10.1	Assumptions	177
10.2	Definitions	178
10.3	The problem statement	178
10.4	The Bellman equation	179
11	Deterministic optimal control in continuous time	181
11.1	Theory I	181
11.2	Theory II	188
11.3	Example I: Deterministic optimal control in the Ramsey model	189
11.4	Example II: Extending the first Ramsey example [†]	193
11.5	Example III: Centralised / decentralised equivalence results [‡]	194
11.5.1	The command optimum	196
11.5.2	The decentralised optimum	196
12	Stochastic optimal control in continuous time	203
12.1	Theory III	203
12.2	Stochastic Calculus	205
12.3	Theory IV	208
12.4	Example of stochastic optimal control	212
12.5	Theory V	215
12.6	Conclusion	216
A	Appendix	217
B	Some useful results	227

C Notational conventions

235

Overview

We will start with a refresher on linear programming, particularly Lagrange Theory. The problems we will encounter should provide the motivation for the rest of the first part of this course where we will be concerned mainly with the mathematical foundations of optimisation theory. This includes a revision of basic set theory, a look at functions, their continuity and their maximisation in n -dimensional vector space (we will only occasionally glimpse beyond finite spaces). The main results are conceptual, that is, not illustrated with numerical computation but composed of ideas that should be helpful to understand a variety of key concepts in modern Microeconomics and Game Theory. We will look at two such results in detail—both illustrating concepts from Game Theory: (i) that it is not rational to play a strictly dominated strategy and (ii) Nash’s equilibrium existence theorem. We will not do many proofs throughout the course but those we will do, we will do thoroughly and you will be asked to proof similar results in the exercises.

The course should provide you with the mathematical tools you will need to follow a master’s level course in economic theory. Familiarity with the material presented in a ‘September course’ on the level of Chiang (1984) or Simon and Blume (1994) is assumed and is sufficient to follow the exposition. The justification for developing the theory in a rigorous way is to get used to the precise mathematical notation that is used in both the journal literature and modern textbooks. We will seek to illustrate the abstract concepts we introduce with economic examples but this will not always be possible as definitions are necessarily abstract. More readily applicable material will follow in later sessions. Some sections are flagged with daggers[†] indicating that they can be skipped on first reading.

The main textbook we will use for the Autumn term is (Sundaram 1996). It is more technical and to an extent more difficult than the course itself. We will cover about a third of the book. If you are interested in formal analysis or are planning to further pursue economic research, I strongly encourage you to work through this text. If you find yourself struggling, consult a suitable text from the reference section.

The second part of the course (starting in December) will be devoted to the main optimisation tool used in dynamic settings as in most modern Macroeconomics: Dynamic Control Theory. We will focus on the Bellman approach and develop the Hamiltonian in both a deterministic and stochastic setting. In addition we will derive a cookbook-style recipe of how to solve the optimisation problems you will face in the Macro-part of your economic theory lectures. To get a firm grasp of this you will need most of the fundamentals we introduced in the Autumn term sessions. The main text we will use in the Spring term is (Obstfeld 1992); it forms the basis of section (11) and almost all of (12). You should supplement your reading by reference to (Kamien and Schwartz 1991) which, although a very good book, is not exactly cheap. Since we will not touch upon more than a quarter of the text you should only buy it if you lost your heart to dynamic optimisation.

Like in every mathematics course: Unless you already know the material covered quite well, there is no way you can understand what is going on without doing at least some of the exercises indicated at the end of each section. Let me close with a word of warning: This is only the third time these notes are used for teaching. This means that numerous mistakes, typos and ambiguities have been removed by the students using the notes in previous years. I am most grateful to them—but I assure you that there are enough remaining. I apologise for these and would be happy about comments and suggestions. I hope we will have an interesting and stimulating course.

*Paul Schweinzer, Summer 2002.*¹

¹ I am grateful to John Hillas, University of Auckland, for allowing me to use part of his lecture notes for the introductory section on the Lagrangian. I owe a similar debt to Maurice Obstfeld, University of California at Berkeley, for allowing me to incorporate his paper into the spring-term part of the notes. Pedro Bacao—endowed with endless energy and patience—read the whole draft and provided very helpful comments and suggestions and an anonymous referee contributed most detailed notes and corrections.

“The good Christian should beware of mathematics and all those who make empty prophecies. The danger already exists that mathematicians have made a covenant with the devil to darken the spirit and to confine man in the bonds of hell.”

Saint Augustine (4th C.)

Part I

Static analysis

Chapter 1

Mathematical programming

Reading: Your micro-textbooks (Varian 1992b) and (Mas-Colell, Whinston, and Green 1995) can be relied on for further illustrations and examples. An undergraduate-level but yet formal introduction is contained in (Nicholson 1990, chapter II.4).

We will come back to all of the topics glanced over in this introductory chapter more systematically—and on the basis of (Sundaram 1996)—as soon as we have worked through chapter 2. The main purpose of this chapter is to review concepts you already know from undergraduate micro courses and to develop questions which we cannot easily answer without getting more formal.

A typical *linear program* (i.e. an optimisation problem for which both objective and constraints are linear in the variables \mathbf{x}) in ‘standard form’ consists of:

- (a) a set of linear equations or inequalities: $A\mathbf{x} = \mathbf{b}$,
- (b) sign constraints on (some of) the unknowns: $\mathbf{x} \geq 0$,
- (c) a linear form to be minimised (maximised): $\min \mathbf{c}^T \mathbf{x}$.

Solutions to (a) & (b) are called *feasible* and solutions to (a), (b) & (c) are called *optimal*. Every *primal* linear program has a *dual* statement. The dual program for the above example: $A\mathbf{x} = \mathbf{b}$, $\mathbf{x} \geq 0$, $\mathbf{c}^T \mathbf{x} = \min$ is $\mathbf{y}^T A \leq \mathbf{c}^T$, $\mathbf{y}^T \mathbf{b} = \max$. For now we just state this as a fact but after we have discussed Lagrange theory in some detail we will return to the topic of duality in section 4.3.

Exl: (Franklin 1980, p4) Portfolio choice problem. Suppose a bank has £100m and wants to put part of its money into loans (L), another part into securities (S) and keeps also some cash. Loans earn high interest (say, 10%), securities do not (say, 5%); securities, however, have a liquidity advantage. Hence we formalise:

$0.1L + 0.05S$	objective (max)
$L \geq 0; S \geq 0; L + S \leq 100$	initial constraints
$L - 3S \leq 0$	liquidity constraint: keep 25% of funds liquid
$L \geq 30$	loan balance constraint: keep some cash for best customers

We best solve this problem graphically as in figure 1.1. ◁




Figure 1.1: A linear program. The dashed lines are level sets, the shaded area is the feasible set.

1.1 Lagrange theory

1.1.1 Non-technical Lagrangian

Consider the problem of a consumer who seeks to distribute his income across the purchase of the two goods he consumes, subject to the constraint that he spends no more than his total income. Let us denote the amount of the first good that he buys x_1 and the amount of the second good x_2 , the prices of the two goods p_1 and p_2 , and the consumer's income y . The utility that the consumer obtains from consuming x_1 units of good 1 and x_2 of good two is denoted $u(x_1, x_2)$. Thus the consumer's problem is to maximise $u(x_1, x_2)$ subject to the constraint that $p_1x_1 + p_2x_2 \leq y$. (We shall soon write $p_1x_1 + p_2x_2 = y$, i.e., we shall assume that the consumer must (wants to) spend all of his income.) Before discussing the solution of this problem let us write it in a more 'mathematical' way:

$$\begin{aligned} \max_{x_1, x_2} u(x_1, x_2) \\ \text{s.t.: } p_1x_1 + p_2x_2 = y \end{aligned} \quad (1.1.1)$$

We read this as "Choose x_1 and x_2 to maximise $u(x_1, x_2)$ subject to the constraint that $p_1x_1 + p_2x_2 = y$." Let us assume, as usual, that the indifference curves (i.e., the sets of points (x_1, x_2) for which $u(x_1, x_2)$ is a constant) are convex to the origin. Let us also assume that the indifference curves are nice and smooth. Then the point (x_1^*, x_2^*) that solves the maximisation problem (1.1.1)

is the point at which the indifference curve is tangent to the budget line as given in figure 1.2. One thing we can say about the solution is that at the point (x_1^*, x_2^*) it must be true that the




Figure 1.2: A non-linear maximisation program.

marginal utility with respect to good 1 divided by the price of good 1 must equal the marginal utility with respect to good 2 divided by the price of good 2. For if this were not true then the consumer could, by decreasing the consumption of the good for which this ratio was lower and increasing the consumption of the other good, increase his utility. This is shown in figure 1.3. Marginal utilities are, of course, just the partial derivatives of the utility function. Thus we have:




Figure 1.3: A point where marginal utilities differ cannot be optimal.

$$\frac{\frac{\partial u}{\partial x_1}(x_1^*, x_2^*)}{p_1} = \frac{\frac{\partial u}{\partial x_2}(x_1^*, x_2^*)}{p_2}. \quad (1.1.2)$$

The argument we have just made seems very ‘economic.’ But it is easy to give an alternative argument that does not explicitly refer to the economic intuition. Let x_2^u be the function that defines the indifference curve through the point (x_1^*, x_2^*) , i.e.,

$$u(x_1, x_2^u(x_1)) \equiv \bar{u} \equiv u(x_1^*, x_2^*).$$

Now, totally differentiating this identity gives:

$$\frac{\partial u}{\partial x_1}(x_1, x_2^u(x_1)) + \frac{\partial u}{\partial x_2}(x_1, x_2^u(x_1)) \frac{dx_2^u}{dx_1}(x_1) = 0.$$

That is,

$$\frac{dx_2^u}{dx_1}(x_1) = -\frac{\frac{\partial u}{\partial x_1}(x_1, x_2^u(x_1))}{\frac{\partial u}{\partial x_2}(x_1, x_2^u(x_1))}.$$

Now $x_2^u(x_1^*) = x_2^*$. Thus the slope of the indifference curve at the point (x_1^*, x_2^*) is

$$\frac{dx_2^u}{dx_1}(x_1^*) = -\frac{\frac{\partial u}{\partial x_1}(x_1^*, x_2^*)}{\frac{\partial u}{\partial x_2}(x_1^*, x_2^*)}. \quad (1.1.3)$$

We know that the budget line is defined as $x_1 p_1 + x_2 p_2 = y$. Setting $x_2 = 0$, we get $x_1 = \frac{y}{p_1}$ and similarly $x_2 = \frac{y}{p_2}$ for $x_1 = 0$. Hence we know that the slope of the budget line in figure 1.2, is given by

$$-\frac{x_1}{x_2} = -\frac{y/p_1}{y/p_2} = -\frac{p_2 y}{p_1 y} = -\frac{p_2}{p_1}.$$

Combining this with (1.1.3) gives result (1.1.2). Since we also know that (x_1^*, x_2^*) must satisfy

$$p_1 x_1^* + p_2 x_2^* = y \quad (1.1.4)$$

we have two equations in two unknowns and we—given that we know the utility function, p_1, p_2 and y —go happily away and solve the problem (not bothering for the moment with the inequality constraint that we skillfully ignored above). What we shall develop is a systematic and useful way to obtain the conditions (1.1.2) and (1.1.4). Let us first denote the common value of the ratios in (1.1.2) by λ . That is,

$$\frac{\frac{\partial u}{\partial x_1}(x_1^*, x_2^*)}{p_1} = \lambda = \frac{\frac{\partial u}{\partial x_2}(x_1^*, x_2^*)}{p_2}.$$

and we can rewrite this and (1.1.4) as:

$$\begin{aligned} \frac{\partial u}{\partial x_1}(x_1^*, x_2^*) - \lambda p_1 &= 0 \\ \frac{\partial u}{\partial x_2}(x_1^*, x_2^*) - \lambda p_2 &= 0 \\ y - p_1 x_1^* - p_2 x_2^* &= 0. \end{aligned} \quad (1.1.5)$$

Now we have three equations in x_1^*, x_2^* , and the new artificial (or auxiliary) variable λ . Again we can, perhaps, solve these equations for x_1^*, x_2^* , and λ . Consider the following function known as

the *Lagrangian*:

$$\mathcal{L}(x_1, x_2, \lambda) = u(x_1, x_2) + \lambda(y - p_1x_1 - p_2x_2).$$

Now, if we calculate $\frac{\partial \mathcal{L}}{\partial x_1}$, $\frac{\partial \mathcal{L}}{\partial x_2}$, and $\frac{\partial \mathcal{L}}{\partial \lambda}$, and set the results equal to zero we obtain exactly the equations given in (1.1.5). We shall now describe this technique in a somewhat more general way.

1.1.2 General Lagrangian

Suppose that we have the following more general maximisation problem:

$$\begin{aligned} \max_{x_1, \dots, x_n} f(x_1, \dots, x_n) \\ \text{s.t.: } g(x_1, \dots, x_n) = c. \end{aligned} \tag{1.1.6}$$

We construct the Lagrangian as above as:

$$\mathcal{L}(x_1, \dots, x_n, \lambda) = f(x_1, \dots, x_n) + \lambda(c - g(x_1, \dots, x_n))$$

then (x_1^*, \dots, x_n^*) solves (1.1.6) and there is a value of λ , say λ^* such that (for $i = 1, \dots, n$):

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x_i}(x_1^*, \dots, x_n^*, \lambda^*) &= 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda}(x_1^*, \dots, x_n^*, \lambda^*) &= 0. \end{aligned} \tag{1.1.7}$$

Notice that the conditions (1.1.7) are precisely the first order conditions for choosing x_1, \dots, x_n to maximise \mathcal{L} , once λ^* has been chosen. This provides an intuition into this method of solving the constrained maximisation problem. In the constrained problem we have told the decision maker that he must satisfy $g(x_1, \dots, x_n) = c$ and that he should choose among all points that satisfy this constraint the point at which $f(x_1, \dots, x_n)$ is greatest. We arrive at the same answer if we tell the decision maker to choose any point he wishes but that for each unit by which he violates the constraint $g(x_1, \dots, x_n) = c$ we shall take away λ units from his payoff. Of course we must be careful to choose λ to be the correct value. If we choose λ too small the decision maker may choose to violate his constraint—e.g., if we made the penalty for spending more than the consumer's income very small the consumer would choose to consume more goods than he could afford and to pay the penalty in utility terms. On the other hand, if we choose λ too large the decision maker may violate his constraint in the other direction, e.g., the consumer would choose not to spend any of his income and just receive λ units of utility for each unit of his income.

It is possible to give a more general statement of this technique, allowing for multiple constraints. (We should always have fewer linearly independent constraints than we have variables.) Suppose we have more than one constraint. Consider the problem with n unknowns and m constraints:

$$\begin{aligned} \max_{x_1, \dots, x_n} f(x_1, \dots, x_n) \\ \text{s.t.: } \begin{aligned} g_1(x_1, \dots, x_n) &= c_1 \\ &\vdots \\ g_m(x_1, \dots, x_n) &= c_m. \end{aligned} \end{aligned} \tag{1.1.8}$$

Again we construct the Lagrangian:

$$\begin{aligned}\mathcal{L}(x_1, \dots, x_n, \lambda_1, \dots, \lambda_m) = \\ f(x_1, \dots, x_n) + \lambda_1(c_1 - g_1(x_1, \dots, x_n)) + \dots + \lambda_m(c_m - g_m(x_1, \dots, x_n))\end{aligned}$$

and again if (x_1^*, \dots, x_n^*) solves (1.1.8) there are values of $\boldsymbol{\lambda}$, say $\lambda_1^*, \dots, \lambda_m^*$ such that (for $i = 1, \dots, n, j = 1, \dots, m$):

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x_i}(x_1^*, \dots, x_n^*, \lambda_1^*, \dots, \lambda_m^*) &= 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda_j}(x_1^*, \dots, x_n^*, \lambda_1^*, \dots, \lambda_m^*) &= 0.\end{aligned}$$

In order for our procedure to work in general we also have to assume that a **constraint qualification** holds.¹ But since we have not yet developed the tools to explain this condition, so we just mention it here in passing and explain it in detail later—after all, this chapter is supposed to be motivational. (For a full treatment, see section 4, especially the discussion of (4.1.4).)

Exl: (Sundaram 1996, 127) We want to

$$\begin{aligned}\max_{x_1, x_2} f(x_1, x_2) &= x_1^2 - x_2^2 \\ \text{s.t.} \quad g(x_1, x_2) &= 1 - x_1^2 - x_2^2 = 0.\end{aligned}$$

This means the constraint set \mathcal{D} is defined as

$$\mathcal{D} = \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1^2 + x_2^2 = 1\}$$

which is just the unit circle in \mathbb{R}^2 . For future reference we state that $f(\cdot)$ is continuous and \mathcal{D} is compact. Therefore, by the Weierstrass Theorem, we know that there exists a solution to this problem. Without going into details, we convince ourselves that the constraint qualification holds

$$\nabla g((x_1^*, x_2^*)) = (2x_1, 2x_2).$$

Since $x_1^2 + x_2^2 = 1$ implies that x_1 and x_2 cannot be both zero on \mathcal{D} , we must have $\text{rank}(\nabla g((x_1^*, x_2^*))) = 1$ at all $(x_1, x_2) \in \mathcal{D}$. Therefore the constraint qualification holds.

We set up the Lagrangian as

$$\mathcal{L}(x_1, x_2, \lambda) = x_1^2 - x_2^2 + \lambda(1 - x_1^2 - x_2^2)$$

and find the critical points of \mathcal{L} to be

$$\begin{aligned}2x_1 - 2\lambda x_1 &= 0 \\ -2x_2 - 2\lambda x_2 &= 0 \\ x_1^2 + x_2^2 &= 1.\end{aligned}\tag{1.1.9}$$

¹ For *linear* $f(\cdot), g(\cdot)$ we are doing more than is required here because for linear programming problems the linearity of constraints is sufficient for the constraint qualification to hold. But since we develop the theory for non-linear programming it is really good practise to *always* check this condition.

From the first two we get

$$\begin{aligned} 2x_1(1 - \lambda) &= 0 \\ 2x_2(1 + \lambda) &= 0 \end{aligned}$$

and conclude that for $\lambda \neq \pm 1$ these can only hold for $x_1 = x_2 = 0$. This, however, violates the third equation in (1.1.9). Therefore we know that $\lambda = \pm 1$ and we find the four possible solutions to be

$$(x_1, x_2, \lambda) = \begin{cases} (+1, 0, +1) \\ (-1, 0, +1) \\ (0, +1, -1) \\ (0, -1, -1) \end{cases} \quad (1.1.10)$$

We evaluate f at these possible solutions and find that $f(1, 0, 1) = f(-1, 0, 1) = 1$ beats the other two. Since the critical points of \mathcal{L} contain the global maxima of f on \mathcal{D} , the first two points must be global maximisers of f on \mathcal{D} . \triangleleft

Exl: We want to maximise $f(x) = 3x + 5$. What produces $\frac{\partial f}{\partial x} = 0$? Is this plausible? What goes wrong? \triangleleft

1.1.3 Jargon & remarks

Rem: We call an inequality constraint $x \geq y$ **binding** (i.e. satisfied with equality) if $x = y$ and **slack** (i.e., not satisfied with equality) if $x > y$.

Rem: As a convention, in a constrained optimisation problem we put a *negative* multiplier sign on a minimisation problem and a positive sign on a maximisation problem.

Rem: We very casually transformed the inequality constraints in this section into equality constraints. This is **not** always possible. Especially ignoring non-negativity constraints can be fatal: Consider

$$\begin{aligned} \max_{x_1, x_2} f(x_1, x_2) &= x_1 + x_2 \\ \text{s.t.: } (x_1, x_2) &\in \mathcal{B}(p, I) = \{I - p_1x_1 - p_2x_2 \geq 0\}. \end{aligned}$$

Although the objective is continuous and the budget set $\mathcal{B}(p, I)$ is compact we cannot set up the problem implying that $x_i > 0$ since we get ‘corner’ solutions (i.e. solutions where $x_i = 0, x_{j \neq i} \neq 0$)

$$(x_1^*, x_2^*) = \begin{cases} (\frac{I}{p_1}, 0) & \text{if } p_1 < p_2 \\ (x_1 \in [0, \frac{I}{p_1}], x_2 = \frac{I - p_1x_1}{p_2}) & \text{if } p_1 = p_2 \\ (0, \frac{I}{p_2}) & \text{if } p_2 < p_1 \end{cases}$$

that imply that $x_i = 0$ if $p_i > p_{j \neq i}$. Thus the non-negativity constraints ‘bite’.

If we had ignored these considerations and set-up the Lagrangian mechanically we would have got the following system of critical points for \mathcal{L}

$$\begin{aligned} 1 - \lambda p_1 &= 0 \\ 1 - \lambda p_2 &= 0 \\ p_1x_1 + p_2x_2 &= I. \end{aligned}$$

This implies $p_1 = p_2$ which is the only case when the non-negativity constraints do not bite—which is the clearly very specific case where the budget line and the objective coincide. The Lagrangian method fails to pick up a solution in any other case.

1.2 Motivation for the rest of the course

Notice that we have been referring to the set of conditions which a solution to the maximisation problem must satisfy. (We call such conditions necessary conditions.) So far we have not even claimed that there necessarily is a solution to the maximisation problem. There are many examples of maximisation problems which have no solution. One example of an unconstrained problem with no solution is

$$\max_x 2x$$

which maximises the choice of x for the function $2x$ over an unbounded choice set. Clearly the greater we make x the greater is $2x$, and so, since there is no upper bound on x there is no maximum. Thus we might want to restrict maximisation problems to those in which we choose x from some bounded set.

Again, this is not enough. Consider the problem

$$\max_{0 \leq x \leq 1} \frac{1}{x}.$$

This function has a graph that looks like the one in figure 1.4. The smaller we make x the greater




Figure 1.4: The graph of $\frac{1}{x}$ in a small neighbourhood of $x=0$.

is $\frac{1}{x}$ and yet at zero $\frac{1}{x}$ is not even defined. We could define the function to take on some value at zero, say 7. But then the function would not be continuous. Or we could leave zero out of the feasible set for x , say $0 < x \leq 1$. Then the set of feasible x is not closed. Since there would

obviously still be no solution to the maximisation problem in these cases we shall want to restrict maximisation problems to those in which we choose x to maximise some continuous function from some closed (and because of the previous example) bounded set. (We call a set of numbers, or more generally a set of vectors, that is both closed and bounded a compact set.) Is there anything else that could go wrong? No! An important result known as the *Weierstrass Theorem* says that if the function to be maximised is continuous and the set over which we are choosing is both closed and bounded (i.e., is compact), then there is a solution to the maximisation problem. We will look at the Weierstrass Theorem later on in the course in more detail when we have introduced the necessary concepts more rigorously.

Notice that in defining compact sets we typically use inequalities, such as $x \geq 0$. However, in Section 1.1 we did not consider such constraints, but rather considered only equality constraints. Even in the example of utility maximisation at the beginning of that section, there were implicitly constraints on x_1 and x_2 of the form $x_1 \geq 0, x_2 \geq 0$.

A truly satisfactory treatment would make such constraints explicit. It is possible to explicitly treat the maximisation problem with inequality constraints, at the price of a little additional complexity. We shall return to this question later in the course when we discuss *Kuhn-Tucker theory*.

Also, notice that had we wished to solve a minimisation problem we could have transformed the problem into a maximisation problem by simply multiplying the objective function by -1 . That is, if we wish to minimise $f(x)$ we could do so by maximising $-f(x)$. The maxima and minima have exactly the same structure—thus we cannot tell whether (x_1^*, \dots, x_n^*) identifies a minimum or a maximum. Again, this is an illustration of the importance of looking at the higher order derivatives—as we will do later on.

Exercises

Exc 1.1: (Sundaram 1996, 142:5.1) (i) Find the maxima and the minima of $f(x, y) = x^2 - y^2$ on the unit circle $x^2 + y^2 = 1$ using the Lagrange multipliers method. (ii) Using the substitution $y^2 = 1 - x^2$, solve the same problem as a single variable unconstrained problem. Do you get the same results? Why or why not?

Exc 1.2: (Sundaram 1996, 142:5.2) (i) Show that the problem of maximising $f(x, y) = x^3 + y^3$ on the constraint set $\mathcal{D} = \{(x, y) \mid x + y = 1\}$ has no solution. (ii) Show also that if the Lagrangian method were used on this problem, the critical points of the Lagrangian have a unique solution. (iii) Is the point identified by this solution either a local maximum or a (local or global) minimum?

Exc 1.3: (Sundaram 1996, 142:5.3) Find the maxima of the following functions subject to the specified constraints:

- | | |
|--|---|
| 1. $f(x, y) = xy$
s.t.: $x^2 + y^2 = 2a^2$; | 2. $f(x, y) = \frac{1}{x} + \frac{1}{y}$
s.t.: $(\frac{1}{x})^2 + (\frac{1}{y})^2 = (\frac{1}{a})^2$; |
| 3. $f(x, y, z) = x + y + z$
s.t.: $\frac{1}{x} + \frac{1}{y} + \frac{1}{z} = 1$; | 4. $f(x, y, z) = xyz$
s.t.: $\frac{1}{x} + \frac{1}{y} + \frac{1}{z} = 1$; |
| 5. $f(x, y) = x + y$
s.t.: $xy = 16$; | 6. $f(x, y, z) = x^2 + 2y - z^2$
s.t.: $2x - y = 0$ and $x + z = 6$. |

Exc 1.4: (Sundaram 1996, 142:5.5) Consider the problem:

$$\begin{aligned} & \max_{x,y} x^2 + y^2 \\ & \text{s.t.: } (x - 1)^3 - y^2 = 0. \end{aligned} \tag{1.2.1}$$

1. Solve the problem geometrically.
2. Show that the method of Lagrange multipliers does not work in this case. Can you explain why?

Exc 1.5: (Sundaram 1996, 142:5.11) Consider the problem of maximising the utility function

$$\begin{aligned} & \max_{x,y} u(x, y) = x^{\frac{1}{2}} + y^{\frac{1}{2}} \\ & \text{s.t.: } \{(x, y) \in \mathbb{R}_+^2 \mid px + y = 1\}. \end{aligned}$$

Show that if the non-negativity constraints $x, y \geq 0$ are ignored, and the problem is written as an equality constrained one, the resulting Lagrangian has a unique critical point. Does this critical point identify a solution to the problem? Why or why not?

Chapter 2

Fundamentals

Reading: Further discussion and proofs of most topics can be found in (Sundaram 1996). More advanced but perfectly readable treatments are (Rudin 1975) or (Royden 1988). These references apply to the whole chapter.

This chapter is where the fun really starts. If this were a proper mathematics course, we would only have ‘fundamentals’ sections like this one but since it is not, we will treat one particular theme of (mathematical) economics in each of the chapters to follow. In many respects it would be more sensible to bring *all* the basic results needed in these later sections together in a couple of proper fundamentals chapters but for practical reasons this does not seem feasible in an MSc economics course. Therefore we will see a lot of definitions and basic results in the opening pages of the later chapters as well—needless to say this will necessitate some repetition.

We begin with relations and orders that are directly useful in the construction of preference orderings on sets of goods bundles. We clarify the basic distinction between countably infinite and uncountably infinite sets and expand on the issue of the equivalence of sets.

We then discuss basic properties of spaces and in particular vector spaces. We focus on the Euclidean versions of norms, distances, inner products and so forth but also discuss such notions in a more general fashion. We then take a look at boundary and extreme points of sets, in particular the infimum and supremum. Our discussion of sequences then leads on to the topic of convexity (of sets and functions).

The section on the properties of functions includes a discussion of the important concepts of continuity and differentiability. We make a distinction between single- and set-valued functions. After that we have all the required elements in place for the principal result we are after in this chapter: The Weierstrass Theorem. We then discuss some properties of functions that are useful in microeconomics: Homogeneity, homotheticity, and Euler’s Theorem. We close with a refresher on quadratic forms.

What follows will be dry. The reason for this is that before we can discuss useful applications (which are the subject of the following chapter), we have to introduce some basic concepts in a more formal way than you may be used to. All this might seem a bit cumbersome at the beginning but an important reason for doing the definitions properly is to get used to the rigorous notation

that is used in the profession today. There is simply no way around that.

2.1 Sets and mappings

Reading: For a much more comprehensive introduction refer to the first chapter of (Kolmogorov and Fomin 1970). As emphasised there, ‘sets’ and ‘elements’ are concepts that are difficult to define in a way that does not simply replace the word ‘set’ by something equivalent like ‘class’, ‘family’, ‘collection’ etc and the word ‘element’ by something like ‘member.’ So we follow the usual ‘naïve’ approach to set theory in mathematics and regard the notions of a set and its elements as well-understood and primitive.¹

2.1.1 Sets

We denote set membership by ‘ \in ’; writing $x \in X$ means that the element x is contained in the set X . If it is not contained in X , we write $x \notin X$. Elements of sets are enclosed in curly brackets $\{\cdot\}$. To define a set X of real numbers whose members have a certain property P we write $X = \{x \in \mathbb{R} : P(x)\}$. Set inclusion is denoted by ‘ \subset ’; $X \subset Y$ means that all elements of X are also contained in Y .² The empty set ‘ \emptyset ’ denotes a set without any elements. It is a subset of every set.

Let X and Y be any two sets. We denote the **union** of these two sets—i.e. the set of all elements belonging to at least one of the two sets—by $X \cup Y$. Similarly their **intersection**—i.e. the set of all elements belonging to both sets—by $X \cap Y$ (see figure 2.1). Both concepts can be




Figure 2.1: In the left panel the shaded area is the union $X \cup Y$ (left) and in the right panel the intersection $X \cap Y$ (right) of two sets X and Y .

extended to operate on an arbitrary number of sets X_i, Y_i indexed by $i \in I$ (I is called an index set). The corresponding notation is

$$X = \bigcup_{i \in I} X_i; \quad Y = \bigcap_{i \in I} Y_i.$$

¹ The attempts to provide a rigorous axiomatic, i.e. non-naïve, basis for set theory have a fascinating history. One of the most famous contributions is Gödel’s Theorem which indicates that both an affirmative and a negative answer to the same question (of a certain type) may be consistent with any particular axiomatisation.

² We adopt the convention that $X \subset X$, i.e. a set is always a subset of itself. If we want to specify a proper subset we write $X \subset Y, X \neq Y$.

If sets do not have any elements in common, i.e. if $X \cap Y = \emptyset$, they are said to be **disjoint**. Union and intersection have the following properties. \cap and \cup are **commutative**:

$$X \cup Y = Y \cup X, \quad X \cap Y = Y \cap X,$$

associative:

$$(X \cup Y) \cap Z = A \cup (Y \cap Z), \quad (X \cap Y) \cup Z = A \cap (Y \cup Z),$$

and **distributive**:

$$(X \cup Y) \cap Z = (X \cap Z) \cup (Y \cap Z), \quad (X \cap Y) \cup Z = (X \cup Z) \cap (Y \cup Z).$$

The **difference** of two sets $X \setminus Y$ (the order matters!) means all elements of X that are not contained in Y . The **symmetric difference** $X \Delta Y$ is defined as $(X \setminus Y) \cup (Y \setminus X)$ —both concepts are illustrated in figure 2.2.




Figure 2.2: The difference $X \setminus Y$ (left) and symmetric difference $X \Delta Y$ (right) of two sets X and Y .

We denote the **complement** of a set X as X^C . Obviously this requires some underlying basic set that X is part of. We can, e.g., view some set X as a subset of the real numbers \mathbb{R} . Then $X^C = \mathbb{R} \setminus X$. In the following we use the symbol R to denote this universal set.

An important property of union and intersection we can state now is the **duality principle**

$$\begin{aligned} R \setminus \bigcup_i X_i &= \bigcap_i (R \setminus X_i), \\ R \setminus \bigcap_i X_i &= \bigcup_i (R \setminus X_i). \end{aligned}$$

In words the principle (also called *de Morgan's law*) says that the complement of a union of an arbitrary number of sets equals the intersection of the complements of these sets.³

Subsets of the real line \mathbb{R} specified by their end points are called **intervals**. The kind of bracket used to denote the (beginning or) end of the interval specifies if that (beginning or) end point is included in the interval or not. We use a square bracket to specify inclusion and a round bracket to denote exclusion. Therefore, for $a \leq b \in \mathbb{R}$, the interval $(a, b) = \{x \in \mathbb{R} : a < x < b\}$ is called **open**, $[a, b] = \{x \in \mathbb{R} : a \leq x \leq b\}$ is called **closed**, and $[a, b) = \{x \in \mathbb{R} : a \leq x < b\}$, $(a, b] = \{x \in \mathbb{R} : a < x \leq b\}$ are neither open nor closed. More general notions of open- and closedness will be discussed in the following section.

³ If we use the formulation 'arbitrary number of sets' in this way we express the idea that the stated property holds for more than just a finite number of sets.

2.2 Algebraic structures

Reading: A comprehensive reference is (Humphreys and Prest 1989) and a fun introduction with many examples is *Frequently Asked Questions in Mathematics* which can be downloaded (for free) from <http://db.uwaterloo.ca/~alopez-o/math-faq/math-faq.html>.

Def: Let X be a non-empty set. A **binary operation** \diamond on X is a rule for combining two elements $x, x' \in X$ to produce some result $x \diamond x'$ specified by the rule \diamond .

Exl: The prime examples for binary operations \diamond and \odot are the operations of addition and multiplication. \triangleleft

Rem: Consider a set X and a binary operation \diamond . If for all $x, y \in X$ it is true that also $x \diamond y \in X$, the set is frequently referred to as being **closed under** the operation of \diamond .

Def: A tuple $\langle X, \diamond \rangle$ consisting of a X and a binary operation \diamond is called a **group** if \diamond satisfies

1. associativity,
2. there exists an identity element wrt \diamond , and
3. every element has an inverse wrt \diamond .

If \diamond only satisfies (i), it is called a **semigroup** and if it satisfies only (i) and (ii), it is called a **monoid**. If \diamond satisfies all of (i-iii) and in addition is commutative, then $\langle X, \diamond \rangle$ is called a commutative or **Abelian group**.

Def: A tuple $\langle X, \diamond, \odot \rangle$ consisting of a set X and two binary operations \diamond and \odot is called a **ring** if $\langle X, \diamond \rangle$ is an Abelian group and \odot satisfies

1. associativity, and
2. is distributive over \diamond .

If \odot is—in addition—commutative, $\langle X, \diamond, \odot \rangle$ is called a **commutative ring**.

Def: A subset \mathcal{J} of the ring $R = \langle X, \diamond, \odot \rangle$ which forms an additive group and has the property that, whenever $x \in R$ and $y \in \mathcal{J}$, then $x \diamond y$ and $y \diamond x$ belong to \mathcal{J} is called an **ideal**.

Def: Let $\langle X, \diamond, \odot \rangle$ be a commutative ring and $a \neq 0 \in X$. If there is another $b \neq 0 \in X$ such that $a \odot b = 0$, then a is called a **zero divisor**. A commutative ring is called an **integral domain** if it has no zero divisors.

Def: Let $\langle X, \diamond, \odot \rangle$ be a ring. If $\langle X \setminus \{0\}, \odot \rangle$ is an Abelian group, $\langle X, \diamond, \odot \rangle$ is called a **field**.

Def: A ring $\langle X, \diamond, \odot \rangle$ is called **ordered** if for $a, b \in S \subset X$, both $x \diamond y$ and $x \odot y$ are again in S .

Def: A ring $\langle X, \diamond, \odot \rangle$ is called **complete** if every non-empty subset which possess an upper bound has a least upper bound in X .⁴

Rem: It is well known that there are polynomials (or more generally, functions) which do not possess solutions in \mathbb{R} . An example is $x^2 + 1$ which can only be solved in \mathbb{C} . Therefore **function fields** are generally fields over \mathbb{C} .

Def: A (function) field $\langle X, \diamond, \odot \rangle$ is called **algebraically closed** if every polynomial splits into linear factors.⁵

⁴ Please confer to page 28 for precise definitions of the concepts of an upper bound and a least upper bound (sup).

⁵ Linear factors are factors not containing $x \in X$ to any power of two or higher, i.e. they are of the form $ax + b$.

Exl:

- $\langle \mathbb{Z}, + \rangle$ is a group.
- $\langle \mathbb{Z}, +, \cdot \rangle$ is a commutative ring (with an identity (called a *unit* wrt \cdot)).
- $\langle \mathbb{Z}, +, \cdot \rangle$ is an integral domain.
- $\langle \mathbb{Q}, +, \cdot \rangle$ is a ring.
- $\langle \mathbb{R}_+, +, \cdot \rangle$ is an ordered ring ($\langle \mathbb{R}_-, +, \cdot \rangle$ is not).
- $\langle \mathbb{R}, +, \cdot \rangle$ is a complete ring.
- $\langle \mathbb{R}, +, \cdot \rangle$ is a field.
- $\langle \mathbb{C}, +, \cdot \rangle$ is an algebraically closed field.
- The set of even integers is an ideal in the ring of integers.

Further examples can be found in the exercises section 2.4.2. \triangleleft

Exl: $x^2 + x - 6$ can be factored as $(x + 3)(x - 2)$ so is contained in the field $\langle \mathbb{R}, +, \cdot \rangle$. But we cannot factor $x^2 + 1$ in the field of reals while we can do so in the field of complex numbers $\langle \mathbb{C}, +, \cdot \rangle$ where $x^2 + 1 = (x - i)(x + i)$. You may recall that $i^2 = -1$. \triangleleft

2.2.1 Relations

We want to find a criterion for a decomposition (or partition) of a given set into classes. This criterion will be something that allows us to assign some elements of the set to one class and other elements to another. We will call this criterion a relation. Not every relation, however, can partition a set into classes. Consider for instance the criterion that $x, y \in \mathbb{R}$ are members of the same class iff $y > x$: If $y > x$, y is a member of the class but then x cannot be an element of the same class because $x < y$. Moreover, since $y > y$ is not true, y cannot even be assigned to the class containing itself! Let M be a set, $x, y \in M$, and let certain ordered pairs (x, y) be called ‘labeled.’ If (x, y) are labeled, we say that x is related to y by the binary relation R , symbolically xRy . For instance, if M is the set of circles, xRy may define classes with equal diameter—in such cases dealing with some notion of equality of elements, R is called an equivalence relation. This, by the way, is a general result: The relation ‘ R ’ acts as a criterion for assigning two elements of a set M to the same class iff R is an equivalence relation on M .

The purpose of the enterprise is to find suitable relations that will allow us to describe preferences of individuals over sets of consumption bundles. This introduces some kind of order over the set of all consumption bundles. Later we will try to express the same order by means of functions because they give us much more flexibility in solving economic problems (e.g. there is no derivative for relations).

Def: A **relation** R from M into M is any subset R of $M \times M$. We interpret the notation ‘ xRy ’ as $(x, y) \in R$. R is called a (binary) **equivalence relation** on M if it satisfies:

1. reflexivity $xRx, \forall x \in M$;
2. symmetry $xRy \implies yRx, \forall x, y \in M$; and
3. transitivity $xRy \wedge yRz \implies xRz, \forall x, y, z \in M$.

Exl: Take the modulo- b relation as an example for an equivalence relation R . xRy holds if the integer division $\frac{x}{b}$ leaves the same remainder $m \in \mathbb{N}/\emptyset$ on M as some $y \neq x \in M$. Take, for example, the modulo-2 relation ($b = 2$) for $x = 5, y = 7$. Since $\frac{5}{2}$ leaves the same remainder as $\frac{7}{2}$, namely 1, x and y belong to the same ‘equivalence class.’ Let’s check the above properties: (1) reflexivity holds because $\frac{x}{2}$ always leaves the same remainder, (2) symmetry holds because $x = 5, y = 7$ always leave the same remainder over $b = 2$, and (3) transitivity holds because if for some z the remainder for $\frac{z}{2}$ is 1, then all three remainders are the same. A familiar application should be adding hours (modulo 12, 24) or minutes (modulo 60). Can you find the algebraic structure the modulo relation gives rise to? \triangleleft

Def: A relation R on a nonempty set M is called a **partial ordering** (\preceq) on M if it satisfies:

1. reflexivity $xRx, \forall x \in M$;
2. antisymmetry $xRy \wedge yRx \implies y = x, \forall x, y \in M$; and
3. transitivity $xRy \wedge yRz \implies xRz, \forall x, y, z \in M$.

The set M is then called *partially ordered*, or *poset* for short.

Exl: Suppose $M \subset \mathbb{R}$ and xRy means $x \leq y$, then R is a partial ordering and M is a poset. Moreover, since for *any* $x, y \in \mathbb{R}$ we have either $x \leq y$ or $y \leq x$, i.e. all elements of the set are comparable under ‘ \leq ’, the set \mathbb{R} is not merely partially but *totally* ordered. Such a set is called a **chain**. \triangleleft

Def: A **linear extension** of a partially ordered set P is an ordering without repetition (a permutation) of the elements p_1, p_2, \dots of P such that $i < j$ implies $p_i < p_j$.

Exl: Consider the poset $P = \{\{1, 2\}, \{3, 4\}\}$. The linear extensions of P are 1234, 1324, 1342, 3124, 3142, and 3412, all of which have 1 before 2 and 3 before 4. \triangleleft

Def: Let $\mathbb{N}_n = \{1, 2, \dots, n\} \subset \mathbb{N}$. We call a set A **finite** if it is empty, or there exists a one-to-one mapping of A onto \mathbb{N}_n for some natural number n . A set A is called **countably infinite** if all of \mathbb{N} is required. Sets that are either finite or countably infinite are called **countable**.

Exl: The set of all integers \mathbb{Z} is countable. As is the set of all rational numbers \mathbb{Q} . Cantor showed that \mathbb{R} is *not* countable (this is exercise 2.4.2.) \triangleleft

Theorem 1. *The union of a finite or countable number of countable sets A_1, A_2, \dots is countable.*

Proof. Wlg we assume that the sets A_1, A_2, \dots are disjoint (we could always construct disjoint sets by making some of the sets smaller—and thus preserving their countability property). Suppose

we can write the elements of A_1, A_2, \dots in the form of the infinite table below

$$\begin{array}{l} A_1 = \{ a_{11} \quad a_{12} \quad a_{13} \quad a_{14} \quad \dots \quad a_{1i} \quad \dots \} \\ A_2 = \{ a_{21} \quad a_{22} \quad a_{23} \quad a_{24} \quad \dots \quad a_{2i} \quad \dots \} \\ A_3 = \{ a_{31} \quad a_{32} \quad a_{33} \quad a_{34} \quad \dots \quad a_{3i} \quad \dots \} \\ \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \\ A_n = \{ a_{n1} \quad a_{n2} \quad a_{n3} \quad a_{n4} \quad \dots \quad a_{ni} \quad \dots \} \\ \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \end{array}$$

Now we count all elements in the above table in the ‘diagonal’ way sketched below

$$\begin{array}{cccccccc} a_{11} & \rightarrow & a_{12} & & a_{13} & \rightarrow & a_{14} & \dots & a_{1i} & \dots \\ & \swarrow & & \nearrow & & \swarrow & & & & \\ a_{21} & & a_{22} & & a_{23} & \dots & a_{24} & \dots & a_{2i} & \dots \\ \downarrow & \nearrow & & \swarrow & & & & & & \\ a_{31} & & a_{32} & & a_{33} & \dots & a_{34} & \dots & a_{3i} & \dots \\ & & \swarrow & & & & & & & \\ a_{n1} & & a_{n2} & & a_{n3} & \dots & a_{n4} & \dots & a_{ni} & \dots \\ \vdots & & \vdots & & \vdots & & \vdots & & \vdots & \end{array}$$

this is supposed to mean that we start counting at a_{11} , and then count a_{12}, a_{21}, \dots following the arrows. This procedure will associate a unique number with any a_{ij} in the above matrix and will eventually cover all elements of all sets. Clearly this establishes a one-to-one correspondence with the set \mathbb{Z}_+ . \square

Theorem 2. *Every infinite set has a countably infinite subset.*⁶

Proof. Let M be an infinite set. Choose any element $a_1 \in M$, then choose any $a_2 (\neq a_1) \in M$, then $a_3 (\neq a_1, a_2) \in M$, and go on picking elements from M in this fashion forever. We will never run out of elements to pick because M is infinite. Hence we get a (by construction) countable subset $N \subset M$ with $N = \{a_1, a_2, a_3, \dots\}$. \square

Def: Two (infinite) sets M and N are **equivalent**, symbolically $M \sim N$, if there is a one-to-one relationship between their elements.

Exl: The sets of points in any two closed intervals $[a, b]$ and $[c, d]$ are equivalent (see figure 2.3). The set of points x in $(0, 1)$ is equivalent to the set of all points y on the whole real line; the formula

$$y = \frac{1}{\pi} \arctan x + \frac{1}{2}$$

is an example of the required one-to-one relationship. \triangleleft

Rem: Every infinite set is equivalent to one of its proper (i.e. not equal) subsets.

⁶ Hence countably infinite sets are the smallest infinite sets.




Figure 2.3: A one-to-one correspondence between $p \in [a, b]$ and $q \in [c, d]$.

Rem: The set of real numbers in $[0, 1]$ is uncountable. Any open interval (a, b) is an uncountable set.

Def: A **partition** Z of a set M is a collection of disjoint subsets of M , the union of which is M .

Rem: Any equivalence relation R on M creates a partition Z on M . One example is the modulo- b operation we encountered previously.

Def: The **power set** of a set X is the set of all subsets of X . It is denoted by $\wp(X)$. If X is infinite, we say it has the ‘power of infinity.’

Def: The **cardinality** of a set X is the number of the elements of X . It is denoted by $|X|$. A set with cardinality 1 is called a **singleton**.

Exl: A set with n elements has 2^n subsets. E.g. the set $M = \{1, 2, 3\}$ can be decomposed into $\wp(M) = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{2, 3\}, \{1, 3\}, \{1, 2, 3\}\}$ with $|\wp(M)| = 2^3 = 8$. \triangleleft

2.2.2 Extrema and bounds

Def: An **upper bound** of $X \subset \mathbb{R}$ is an element of the set

$$U(X) = \{u \in \mathbb{R} \mid u \geq x\}$$

for all $x \in X$. The set of **lower bounds** $L(X)$ is the obvious opposite for $l \leq x$. A set X is said to be **bounded** if both a $u \in U(X)$ and a $l \in L(X)$ exist.

Exl: $(0, 1) \subset \mathbb{R}$ is bounded. \mathbb{R} is not bounded. \triangleleft

Def: An **infimum** is the largest lower bound of the set X , it is written $\inf X$. A **supremum** is the smallest upper bound and written $\sup X$.

Def: If the infimum (supremum) of a set X is contained in X it is called a **minimum** (**maximum**) of X .

Exl: Let $X = (0, 1)$, the open unit interval. $U(X) = \{x \in \mathbb{R} \mid x \geq 1\}$, $L(X) = \{x \in \mathbb{R} \mid x \leq 0\}$. So $\sup X = 1$ and $\inf X = 0$. We cannot, however, find a minimum or maximum of X . \triangleleft

Exl: In the multi-dimensional case, the infimum (supremum) can be quite different from the minimum (maximum). To see this, we first need to agree on an interpretation of what $x \leq y$

means for $x, y \in \mathbb{R}^n$ for $n > 1$, i.e., we need to agree on a metric. Let us settle on $x \leq y$ if $x_i \leq y_i$ for all $i = 1, \dots, n$, but $x_j < y_j$ for some j . Consider now the set X of points on the unit circle in \mathbb{R}^2 with its centre at $(1,1)$.⁷ Sketching possible $\min X$ ($\max X$) and $\inf X$ ($\sup X$) results in figure 2.4. Notice that neither a minimum nor a maximum exists for X in this setting. \triangleleft




Figure 2.4: The difference between infimum (supremum) and minimum (maximum).

Rem: If the set $U(X)$ is empty, we set $\sup X = \infty$. Conversely, if the set $L(X)$ is empty, we set $\inf X = -\infty$. While this is only a convention, it is an important one because it ensures that the infimum and supremum of a set always exist while the minimum and maximum of a set may well be non-existent. Indeed, it can be shown that if the maximum (minimum) exists, it must be equal to the supremum (infimum).

Rem: Any poset P can have at most one element b satisfying bRy for any $y \in P$. To see this interpret R as ' \leq ' and suppose that both b_1, b_2 have the above bounding property. Then, if it is true that $b_1 \leq b_2$ and $b_2 \leq b_1$, we get from antisymmetry $b_1 = b_2$. Such an element b is called the **least** element of P . By the same argument, both $\inf P$ and $\sup P$ are unique (if they exist—which we ensure adhering to the convention stated in the previous remark).

2.2.3 Mappings

Def: Given two sets $X, Y \subset \mathbb{R}^n$, a **function** $f : X \mapsto Y$ maps elements from the **domain** X to the **range** Y .⁸ Alternatively, we write for $\mathbf{x} \in X, \mathbf{y} \in Y, \mathbf{y} = f(\mathbf{x})$. $f : X \mapsto \mathbb{R}$ are called real(-valued) functions.⁹

Rem: We (informally) choose to call a mapping a function only if it possesses a unique value of $\mathbf{y} = f(\mathbf{x})$ for each $\mathbf{x} \in X$.

Def: A function is **into** if $f(X) \subset Y$ and **onto** if $f(X) = Y$.¹⁰

Def: A function $f : X \mapsto Y$ is called **injective** if, for every $y \in Y$, there exists *at most* one $x \in X$ such that $f(x) = y$.

⁷ We have not yet introduced the notation but—trust me—we refer to the set $X = \{x \in \mathbb{R}^2 : \|x - (1,1)\| = 1\}$.

⁸ Sometimes, $\text{im}(f) \subset Y$ is called the *image* of f and $\ker(f) = \{x \in X \mid f(x) = 0\}$ the *kernel* (or null space) of f .

⁹ This should not be confused with the term 'functional' which is used to designate real functions on function or linear spaces.

¹⁰ Thus every onto mapping is into but not conversely.

Def: A function $f : X \mapsto Y$ is called **surjective** if, for every $y \in Y$, there exists *at least one* $x \in X$ such that $f(x) = y$.

Def: A function $f : X \mapsto Y$ is called **bijective** if, for every $y \in Y$, there exists *exactly one* $x \in X$ such that $f(x) = y$. Hence a bijective mapping is one-to-one and onto.

Figure 2.5 shows examples for such mappings.

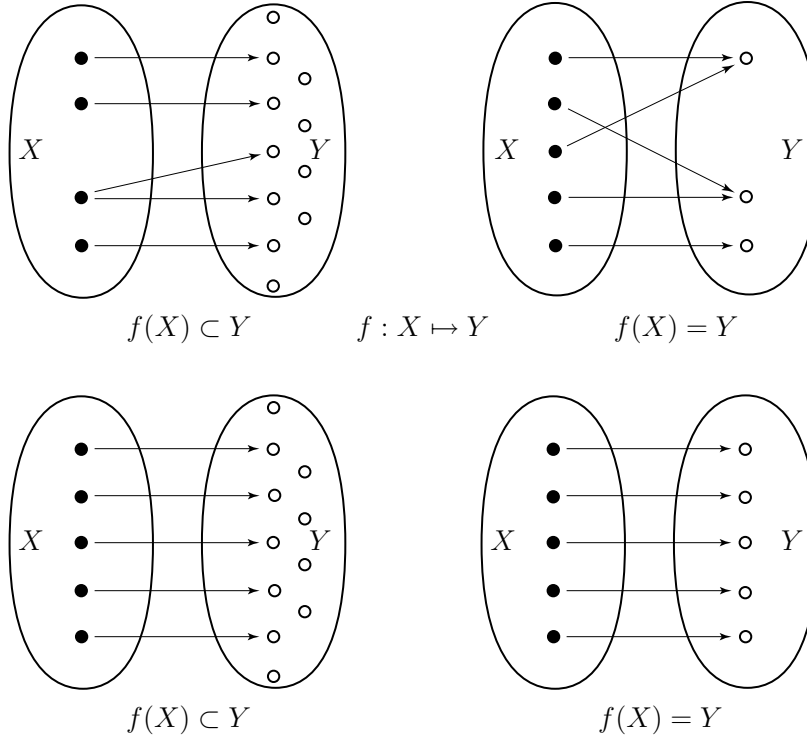


Figure 2.5: An injective and into (left-top), surjective and onto (right-top), one-to-one and into (left-bottom) and a one-to-one and onto (i.e. bijective) mapping (right-bottom).

Def: Let M and M' be any two partially ordered sets, let f be a bijection of M onto M' and let $x, y \in M$, $f(x), f(y) \in M'$. Then f is said to be **order-preserving** (or *isotone*) if $x \leq y \implies f(x) \leq f(y)$. (An order preserving bijection f such that $f(x) \leq f(y) \iff x \leq y$ is called an **isomorphism**).

Rem: If a function is onto and one-to-one (i.e. bijective), its **inverse** $x = f^{-1}(y)$ exists.

Def: A function $f : \mathbb{R}^n \mapsto \mathbb{R}^m$ is called a **linear** function if it satisfies:

1. $f(x + y) = f(x) + f(y)$, $\forall x, y \in \mathbb{R}^n$, and
2. $f(\alpha x) = \alpha f(x)$, $\forall x \in \mathbb{R}^n$, $\forall \alpha \in \mathbb{R}$.

Rem: A function $f(x) = \alpha x + \beta$, ($\alpha, \beta \in \mathbb{R}$) does not satisfy the above conditions for $b \neq 0$. Functions of the form $f(x) = \alpha x + \beta$ are called **affine** functions; these functions are generally taken to be the sum of a linear function for $x \in \mathbb{R}^n$ and a constant $b \in \mathbb{R}^m$.

Def: A function $f : \mathbb{R}^n \mapsto \mathbb{R}^m$ is called a **bilinear** function if it satisfies:

$$f(x, \alpha y + \beta z) = \alpha f(x, y) + \beta f(x, z)$$

for $\alpha, \beta \in \mathbb{R}$ and $x, y, z \in \mathbb{R}^n$.

Def: A function $f : X \mapsto Y$, $x_1, x_2 \in X$, is called **nondecreasing** if $x_1 \leq x_2 \implies f(x_1) \leq f(x_2)$ and **nonincreasing** if $x_1 \leq x_2 \implies f(x_1) \geq f(x_2)$. A **monotonic** function is either nonincreasing or nondecreasing.

2.3 Elementary combinatorics

Proposition 3. *Let A, B be finite sets with $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_p\}$. Then there are m^n distinct bijective mappings of the form $f : A \mapsto B$.*

Proof. For any $i \in [1, n]$, there are p possible choices of $f(a_i)$. As every choice is independent from the other ones, the total number of possibilities is n times p or p^n . \square

Def: The **factorial** (symbolically “!”) is defined for $n \in \mathbb{N}$ as

$$n! \equiv \begin{cases} n(n-1)(n-2)\dots 2 \cdot 1 & \text{for } n > 1 \\ 1 & \text{otherwise.} \end{cases}$$

Proposition 4. *Let A, B be finite sets with $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_p\}$. Then the set of all the injections $f : A \mapsto B$ is finite and its number is A_n^p .*

Proof. For $f(a_1)$ we have n choices, for $f(a_2)$, we have $n-1$ choices, and so on. Each choice reduces the number of further possible choices of elements in B by one. Hence the total number of choices is $n(n-1)(n-2)\dots(n-p+1) \equiv A_n^p$. \square

Def: A **permutation** is an arrangement (ordering) of n out of p distinct objects without repetitions. Hence the number of permutations is

$$A_n^p = \frac{n!}{(n-p)!}.$$

In particular, the number of ways to order n objects is $A_n^n = \frac{n!}{(n-n)!} = n!$.

Def: **Combinations** are k -element subsets of n distinct elements. The number of combinations is

$$\binom{n}{p} \equiv \frac{n!}{p!(n-p)!}.$$

In particular, we have

$$\binom{n}{0} = 1, \quad \text{and} \quad \binom{n}{1} = n.$$

Def: **Variations** are arrangements of k out of n distinct objects with repetitions allowed. The number of variations is n^k .

2.4 Spaces

2.4.1 Geometric properties of spaces

Def: Let $X \subset \mathbb{R}^n$, then the set $\mathcal{S}[X]$ is the set of all linear combinations of vectors from X . $\mathcal{S}[X]$ is called the **span** of X and we say that X spans $\mathcal{S}[X]$.

Def: The **Cartesian plane** \mathbb{R}^2 is spanned by the **Cartesian product** (or **cross product**) of $\mathbb{R} \times \mathbb{R} = \{(x, y) : x \in \mathbb{R}, y \in \mathbb{R}\}$.

Def: A n -dimensional **vector** \mathbf{x} is a finite, ordered collection of elements x_1, \dots, x_n of some space X . Hence $\mathbf{x} \in X^n$. A vector is defined as a column vector. Its **transpose** is a row vector \mathbf{x}^T usually written down as (x_1, x_2, \dots, x_n) .

Def: A (real) **linear space** (or **vector space**) X is a set of vectors $X \subset \mathbb{R}^n$ and the set \mathbb{R} , together with an addition and multiplication operation, such that for any $\mathbf{x}, \mathbf{y}, \mathbf{z} \in X$ and $r, s \in \mathbb{R}$, the following properties hold:

1. commutativity: $\mathbf{x} + \mathbf{y} = \mathbf{y} + \mathbf{x}$;
2. associativity: $(\mathbf{x} + \mathbf{y}) + \mathbf{z} = \mathbf{x} + (\mathbf{y} + \mathbf{z})$ and $r(s\mathbf{x}) = (rs)\mathbf{x}$;
3. distributivity: $r(\mathbf{x} + \mathbf{y}) = r\mathbf{x} + r\mathbf{y}$ and $(r + s)\mathbf{x} = r\mathbf{x} + s\mathbf{x}$;
4. identity: there exists $0 \in X$, $1 \in \mathbb{R}$ such that $\mathbf{x} = \mathbf{x} + 0 = \mathbf{x}$ and $1\mathbf{x} = \mathbf{x}1 = \mathbf{x}$.

Def: Let X be a linear space. A function $\|\cdot\| : X \mapsto [0, \infty]$ is called a **norm** iff for all $\mathbf{x}, \mathbf{y} \in X$ and all $\alpha \in \mathbb{R}$:

1. positivity: $\|\mathbf{x}\| = 0$ iff $\mathbf{x} = 0$, otherwise $\|\mathbf{x}\| > 0$,
2. homogeneity: $\|\alpha\mathbf{x}\| = |\alpha| \|\mathbf{x}\|$, and
3. triangle inequality: $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$.

Def: The **Euclidean norm** is defined as

$$\|x\| = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2}.$$

Rem: Notice that the norm of $x \in \mathbb{R}$ is the absolute of x , denoted by $|x|$. Intuitively, the absolute is the distance from the origin. This intuition is unchanged for higher dimensional spaces.

Exl: The **\mathbf{L}_1 -norm** $\|\cdot\|_1$ is a vector norm defined for the vector $\mathbf{x} \in \mathbb{C}^n$ as

$$\|\mathbf{x}\|_1 = \sum_r^n |x_r|.$$

<

Exl: The **ℓ_2 -norm** $\|\cdot\|_2$ is also a vector norm defined for the complex vector \mathbf{x} as

$$\|\mathbf{x}\|_2 = \sqrt{\sum_{r=1}^n |x_r|^2}.$$

As you can see, for the discrete case this is just the Euclidean Norm—the difference becomes only

clear when it is applied to a function $\phi(x)$. Then this norm is called **L₂-norm** and is defined as

$$\|\phi\|^2 \equiv \phi \cdot \phi \equiv \langle \phi | \phi \rangle \equiv \int |\phi(x)|^2 dx.$$

◁

Exl: The **sup-norm** measures the highest absolute value x takes on the entire set X : $\|x\| = \sup_{x \in X} |x|$. This is often used when the elements of X are functions: $\|f\| = \sup_{x \in X} |f(x)|$. It is also used when the elements are infinite sequences: $x = (x_1, x_2, \dots)$; then $\|x\| = \sup_k |x_k|$. ◁

Rem: The above norms are also used to define normed linear spaces (see below for definition) which then come under names such as L_1 - or L_2 -spaces. The space of bounded, infinite sequences is usually denoted ℓ_∞ and the normed linear space which consists of infinite sequences for which the norm

$$\|x\| = \left(\sum (x_k)^p \right)^{\frac{1}{p}}$$

is finite, is denoted ℓ_p .

Def: Let X be a set. A **distance (metric)** is a function $d : X \times X \mapsto [0, \infty]$ satisfying:

1. positivity: $d(\mathbf{x}, \mathbf{y}) = 0$ iff $\mathbf{x} = \mathbf{y}$, otherwise $d(\mathbf{x}, \mathbf{y}) > 0$,
2. symmetry: $d(\mathbf{x}, \mathbf{y}) = d(\mathbf{y}, \mathbf{x})$, and
3. triangle inequality: $d(\mathbf{x}, \mathbf{y}) \leq d(\mathbf{x}, \mathbf{z}) + d(\mathbf{z}, \mathbf{y})$ for any $\mathbf{z} \in X$.

Def: The **Euclidean distance (Euclidean metric)** is defined as

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{(x_1 - y_1)^2 + \dots + (x_n - y_n)^2}.$$

Def: The **inner product (or scalar product or dot product)** of two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ is defined through the following properties:

1. positivity: $\mathbf{x} \cdot \mathbf{x} = 0$ iff $\mathbf{x} = 0$, otherwise $\mathbf{x} \cdot \mathbf{x} > 0$,
2. symmetry: $\mathbf{x} \cdot \mathbf{y} = \mathbf{y} \cdot \mathbf{x}$, and
3. bi-linearity: $(a\mathbf{x} + b\mathbf{y}) \cdot \mathbf{z} = a\mathbf{x} \cdot \mathbf{z} + b\mathbf{y} \cdot \mathbf{z}$, and $\mathbf{x} \cdot (a\mathbf{y} + b\mathbf{z}) = \mathbf{x} \cdot a\mathbf{y} + \mathbf{x} \cdot b\mathbf{z}$ for any $\mathbf{z} \in \mathbb{R}^n, a, b \in \mathbb{R}$.

(Sometimes, the inner product $\mathbf{x} \cdot \mathbf{y}$ is denoted by angular brackets as $\langle \mathbf{x}, \mathbf{y} \rangle$.)

Def: The **Euclidean inner product** is defined as

$$\mathbf{x} \cdot \mathbf{y} = x_1 y_1 + x_2 y_2 + \dots + x_n y_n.$$

Rem: The inner product of two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ $\mathbf{x} \cdot \mathbf{y} = \sum x_i y_i$ can also be defined as

$$\mathbf{x} \cdot \mathbf{y} \equiv |\mathbf{x}| |\mathbf{y}| \cos(\theta)$$

where θ is the angle between the vectors.¹¹ This has the immediate consequence that

$$\mathbf{x} \cdot \mathbf{y} = 0 \quad \text{iff} \quad \mathbf{x} \text{ and } \mathbf{y} \text{ are perpendicular.}$$

¹¹ Recall that $\sin(0^\circ) = 0$, $\cos(0^\circ) = 1$ and $\sin(90^\circ) = 1$, $\cos(90^\circ) = 0$.

Hence the inner product has the geometric interpretation of being proportional to the length of the orthogonal projection of \mathbf{z} onto the unit vector $\hat{\mathbf{y}}$ when the two vectors are placed so that their tails coincide. This is illustrated in figure 2.6 below. The unit vector $\hat{\mathbf{y}}$ is generated from \mathbf{y} through rotating and scaling of both vectors; \mathbf{z} is generated from \mathbf{x} .




Figure 2.6: Geometric interpretation of the inner product as proportional to the length of the projection of one vector onto the unit vector.

Def: The (n -dimensional) **Euclidean space** (or linear real space) \mathbb{R}^n is the cross product of a finite number n of real lines ($\mathbb{R} \times \dots \times \mathbb{R}$) equipped with the Euclidean inner product

$$\mathbf{x} \cdot \mathbf{y} = \sum_{i=1}^n x_i y_i$$

for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.

Rem: As a consequence of this, the following identities hold in Euclidean space:

1. $\|x\| \equiv \sqrt{x \cdot x}$, and
2. $d(x, y) \equiv \|x - y\|$.

Rem: The subspace $\mathbb{R}_+^n \equiv \{\mathbf{x} \in \mathbb{R}^n | \mathbf{x} \geq 0\}$ is sometimes called the **positive orthant**. Likewise, $\mathbb{R}_-^n \equiv \{\mathbf{x} \in \mathbb{R}^n | \mathbf{x} \leq 0\}$ is called the **negative orthant**.

Def: The **Cauchy-Schwartz inequality** relates the inner product to the norm as

$$|x \cdot y| \leq \|x\| \|y\|.$$

Exl: Let $x = (3, 2)$ and $y = (4, 5)$. Then the inner product $|x \cdot y| = 22 = 3 \cdot 4 + 2 \cdot 5$ is smaller than the product of the norms $\|x\| \|y\| = 23.09 \approx \sqrt{533} = \sqrt{13}\sqrt{41}$. Thus, we have confirmed the Cauchy-Schwartz inequality in \mathbb{R}^2 . \triangleleft

Each of the following spaces takes the definition of the linear space as a basis and extends it with only either the norm, the metric, or the inner product. Only in Euclidean real space where the Cauchy-Schwartz inequality ensures that all three properties are fulfilled and interlinked, we can be sure to be able to count on all three properties.

Def: A **normed linear space** $(X, \|\cdot\|)$ consists of the vector space X and the norm $\|\cdot\|$.

Def: A **metric space** is a pair (X, d) where X is a set and d a metric.¹²

Def: An **inner product space** is a . . .

Def: A **Banach space** is a complete linear space, that is, a linear space in which every Cauchy sequence converges to an element of the space.

Def: A **Hilbert space** is a Banach space with a symmetric bilinear inner product $\langle x, y \rangle$ defined such that the inner product of a vector with itself is the square of its norm

$$\langle \mathbf{x}, \mathbf{y} \rangle = \|\mathbf{x}\|^2.$$

An attempt to replace the notion of norm, metric, or inner product by a purely set-theoretical concept leads to topological spaces. The precise definition is more difficult than the above ones but the intuition is easy: Think of an object that we want to measure, say 1×1 square in a set $X \subset \mathbb{R}^2$. Then we can obviously ‘cover’ the square with the whole set X but for certain subsets of X this may also be possible.¹³ Using these subsets (called *open sets*) we can, in a sense, measure the extent of the square in relation to the structure of the open sets. The coarser the open sets the coarser is this measure and the finer the open sets of X the finer gets the measure. This is formalised below.

Def: A **topological space** is a pair (X, τ) where X is a set and τ is a collection of subsets of X satisfying:

1. both X and \emptyset belong to τ , and
2. the intersection of a finite number and the union of an arbitrary number of sets in τ belong to τ .

Rem: The sets in τ are called **open sets** and τ is called the **topology**. The intuition behind topological spaces is easy to grasp. (Munkres 2000, 77) puts it like this: “We can think of a topological space as being something like a truckload full of gravel—the pebbles and all unions of collections of pebbles being the open sets. If now we smash the pebbles into smaller ones, the collection of open sets has been enlarged, and the topology, like the gravel, is said to have been made finer by the operation.

Def: Suppose that for each pair of distinct points x and y in a topological space T , there are neighbourhoods $N_\varepsilon(x)$ and $N_{\varepsilon'}(y)$ such that $N_\varepsilon(x) \cap N_{\varepsilon'}(y) = \emptyset$. Then T is called a **Hausdorff space** (or *T_2 space*) and said to satisfy the Hausdorff axiom of separation.

Rem: Each pair of disjoint points in a Hausdorff space has a pair of disjoint neighbourhoods.

2.4.2 Topological properties of spaces

Def: A **sequence** $\{x_n\}$ is a function whose domain is the set of positive integers \mathbb{N}_+ (e.g. $1, 2, 3 \dots$).

¹² Notice that the metric is part of the definition of a metric space while in the case of a linear space—which was defined on the sum and product of its elements—the metric was defined independently.

¹³ The precise definition of a cover can be found on page 89, but the intuition is easy: It is a collection of open sets whose union includes the set to be covered.

Def: Let $k : \mathbb{N}_+ \mapsto \mathbb{N}_+$ be some increasing rule assigning a number $k(n)$ to each index n of a sequence $\{x_n\}$. A **subsequence** $\{x_{k(n)}\}$ of the sequence $\{x_n\}$ is a new sequence whose n -th element is the $k(n)$ -th element of the original sequence $\{x_n\}$.

Rem: Sequences have infinitely many elements. Subsequences are order preserving infinite subsets of sequences.

Def: A sequence $\{x_k\}$ is said to be **convergent** to x , if the distance $d(x_k, x)$ tends towards zero as k goes to infinity, symbolically $\{x_k\} \xrightarrow[k \rightarrow \infty]{} x$. x is called the **limit** of $\{x_k\}$.

Rem: A sequence can have at the most one limit.

Def: Alternatively, given a real sequence $\{x_k\}$, $k, l \in \mathbb{N}$, the finite or infinite number

$$\bar{x}_0 = \inf \left\{ \sup_{l \geq k} x_l \right\}$$

is called the **upper limit** (or $\lim \sup$) of $\{x_k\}$. The **lower limit** (or $\lim \inf$) of $\{x_k\}$ is defined as

$$\underline{x}_0 = \sup \left\{ \inf_{l \geq k} x_k \right\}.$$

The sequence $\{x_k\}$ converges if $x_0 = \bar{x}_0 = \underline{x}_0$. Then x_0 is called the limit of $\{x_k\}$.

Def: A sequence $\{x_n\}$ is said to be a **Cauchy sequence** if for each $\varepsilon > 0$ there exists an N such that $\|x_n - x_m\| < \varepsilon$ for all $n, m > N$.¹⁴

Exl: (Sundaram 1996, 8) $\{x_n\} = \frac{1}{k}$ is a convergent sequence for all $k \in \mathbb{N}$ and it converges to the limit $x = 0$ as k goes to infinity: $x_k \rightarrow 0$ for $k \rightarrow \infty$. $\{x_n\} = k$, $k \in \mathbb{N}$, on the other hand, is not convergent and does not possess a limit. \triangleleft

Def: A sequence $\{\mathbf{x}_k\} \in \mathbb{R}^n$ is called **bounded** if there exists a $M \in \mathbb{R}$ such that $\|\mathbf{x}_k\| \leq M$ for all k .

Def: A space X with the property that each Cauchy sequence is convergent to a $x \in X$ is called **complete**.

Theorem 5. *Every convergent sequence in \mathbb{R}^n is bounded.*

Proof. Suppose $x_k \implies x$. Let $\varepsilon = 1$ as the acceptable neighbourhood in the definition of convergence. Then there is a $k(\varepsilon = 1)$ such that for all $k > k(\varepsilon = 1)$, $d(x_k, x) < 1$. Since $d(x_k, x) = \|x_k - x\|$, an application of the triangle inequality yields for any $k \geq k(\varepsilon = 1)$

$$\begin{aligned} \|x_k\| &= \|(x_k - x) + x\| \\ &\leq \|x_k - x\| + \|x\| \\ &< 1 + \|x\|. \end{aligned}$$

Defining $M = \max \{\|x_1\|, \|x_2\|, \dots, \|x_{k(1)-1}\|, 1 + \|x\|\}$ establishes $M \geq \|x_k\|$ for all k . \square

Def: An open ε -**neighbourhood** of a point in $\mathbf{x} \in \mathbb{R}^n$ is given by the set

$$N_\varepsilon(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^n : \|(\mathbf{x} - \mathbf{y})\| < \varepsilon\}.$$

¹⁴ Notice that we do not require an explicit statement of the limit x in this case.

Def: A set $X \subset \mathbb{R}^n$ is **open** if for every $\mathbf{x} \in X$, there exists an ε such that $N_\varepsilon(\mathbf{x}) \subset X$.

Def: A set $X \subset \mathbb{R}^n$ is **bounded** if for every $\mathbf{x} \in X$, there exists an $\varepsilon < \infty$ such that $X \subset N_\varepsilon(\mathbf{x})$.

Def: A point $\mathbf{x} \in \mathbb{R}^n$ is called a **contact point** of the set $X \subset \mathbb{R}^n$ if every $N_\varepsilon(\mathbf{x})$ of \mathbf{x} contains at least one point of X . The set of all contact points of X (which by definition contains the set X) is called the **closure** of X . It is denoted by $[X]$ or \overline{X} .

Def: A set X is called **closed** if it coincides with its closure $X = [X]$.

Def: A set $X \subset \mathbb{R}^n$ is called **compact** if for all sequences of points $\{x_k\}$ such that $x_k \in X$ for each k , there exists a subsequence $\{x_{m(k)}\}$ of $\{x_k\}$ and a point $x \in X$ such that $\{x_{m(k)}\} \rightarrow x$.

Theorem 6. (Heine-Borel) *A set $X \subset \mathbb{R}^n$ is compact iff it is closed and bounded.*¹⁵

Proof. See (Weibull 2001, 49). □

Corollary 7. *Every closed and bounded subset B of \mathbb{R}^n is compact.*

Def: A subset $X \subset Y \subset \mathbb{R}^n$ is **dense** in Y if $Y \subset [X]$.

Intuitively, a subset $X \subset Y$ is dense in Y if the points in X are ‘everywhere near’ to the points in Y , i.e. if each neighbourhood of every point in Y contains some point from X (i.e. is a contact point of X).

Exl: $\mathbb{N} \subset \mathbb{R}$ is not dense in \mathbb{R} but $\mathbb{Q} \subset \mathbb{R}$ is. Likewise, $(0, 1) \subset \mathbb{R}$ is a dense subset of $[0, 1] \subset \mathbb{R}$. \triangleleft

Def: A point $\mathbf{x} \in \mathbb{R}^n$ is called a **limit point** of the set $X \subset \mathbb{R}^n$ if every $N_\varepsilon(\mathbf{x})$ of \mathbf{x} contains infinitely many points of X . The limit point itself may or may not be an element of X .

Theorem 8. (Bolzano-Weierstrass) *Every infinite and bounded set in \mathbb{R}^n has at least one limit point.*

Proof. See (Royden 1988). □

Def: A point $\mathbf{x} \in X \subset \mathbb{R}^n$ is called an **isolated point** of X if there exists a ‘sufficiently small’ $\varepsilon > 0$ such that $N_\varepsilon(\mathbf{x})$ contains no points of X other than \mathbf{x} itself.

Def: A point $\mathbf{x} \in X \subset \mathbb{R}^n$ is called an **interior point** of X if there exists an $\varepsilon > 0$ such that $N_\varepsilon(\mathbf{x}) \subset X$. The set X is called **open** if all its points are interior.

Exl: Any set consisting of a finite number of points is closed. Consider the set $X = \{1, 2, 3, 4\}$. We cannot choose any ‘small’ neighbourhood $\varepsilon > 0$ around, say 3, that includes another element from X . Hence there are no contact points and the closure of X equals X which is another way of saying that X is closed. \triangleleft

Theorem 9. *A subset X of \mathbb{R} is open iff its complement $\mathbb{R} \setminus X$ is closed.*

Proof. If X is open, then for every point $\mathbf{x} \in X$ there exists a $\varepsilon > 0$ such that $N_\varepsilon(\mathbf{x}) \subset X$. Hence there are no contact points of $\mathbb{R} \setminus X$ in X , therefore all contact points of $\mathbb{R} \setminus X$ must be in $\mathbb{R} \setminus X$ itself and hence $\mathbb{R} \setminus X$ must be closed.

Conversely, if $\mathbb{R} \setminus X$ is closed, then for any point $\mathbf{x} \in X$ there must exist an $N_\varepsilon(\mathbf{x}) \subset X$. Otherwise

¹⁵ Beware, compactness for general sets X is different—the theorem can be found on page 126.

every $N_\varepsilon(\mathbf{x})$ would also contain points of $R \setminus X$, meaning that \mathbf{x} would be a contact point of $R \setminus X$ not contained in $R \setminus X$. Hence X is open. \square

The above theorem leads naturally to the topological definition of open- and closedness. We will go further than this in our discussion of fixed points later in the course but for now let us just state these definitions to complement the above definitions of open- and closedness.

Def: A set $O \subset \mathbb{R}$ is called **open** if it is the union of open intervals¹⁶

$$O = \bigcup_{i \in I} A_i.$$

A set is called **closed** if its complement is open.

Convexity

Def: A set $X \subset \mathbb{R}^n$ is **convex** if for all $\mathbf{x}, \mathbf{y} \in X$, $\lambda \in [0, 1]$

$$\lambda \mathbf{x} + (1 - \lambda) \mathbf{y} \in X.$$

(The empty set is by definition convex.) X is a **strictly convex** if $x \neq y$ and all points $\lambda \mathbf{x} + (1 - \lambda) \mathbf{y}$ for $\lambda \in (0, 1)$ are interior points.

Rem: All hyperplanes are convex sets.

Def: The **convex hull** $co(S)$ is the smallest convex set containing S .¹⁷

Rem: Some sets:

1. There are sets that are neither open nor closed such as e.g. $[0, 1)$ which is bounded but not closed.
2. The empty set \emptyset and all of the real line \mathbb{R} are *both* open *and* closed. This follows from theorem 9.
3. Open intervals, however, may well be convex as e.g. $(0, 1)$ is.
4. The set of non-negative real numbers $[0, \infty)$ is closed by convention since each (sub)sequence converges.¹⁸
5. $[k, \infty)$, $k \in \mathbb{N}$ is closed but not bounded; consequently it is non-compact.
6. There are further examples in figure 2.7.

As additional and very important examples, you are asked in the exercises to think about the open-/closed-/boundedness of (subsets of) \mathbb{R} .

¹⁶ We are more general than required here because it can be shown that a countable number of open intervals will suffice to cover \mathbb{R} . These sets can even be chosen to be pairwise disjoint.

¹⁷ Notice that rather contrary to the common usage of the word hull, the convex hull contains the interior of the set.

¹⁸ Sometimes, the space of real numbers \mathbb{R} is called *quasi-closed* and only the 'extended real line', the space $\overline{\mathbb{R}} = \mathbb{R} \cup \pm\infty$ is taken to be closed. This generalises to higher dimensions.

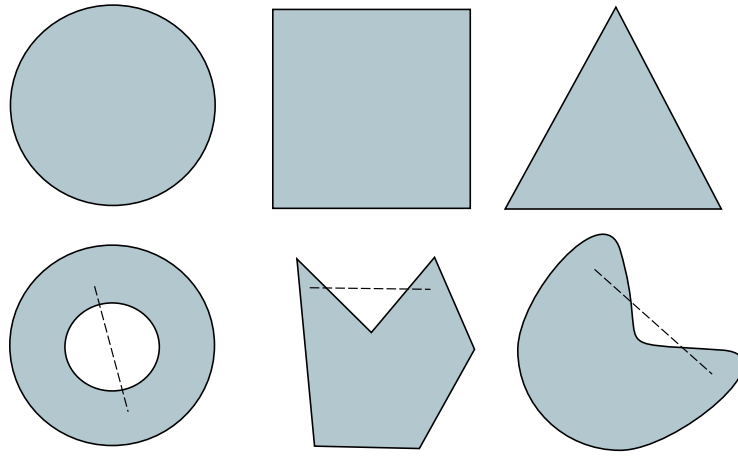


Figure 2.7: (Top) Some convex sets and some non-convex sets (bottom).

Exercises 1

Exc 2.1: Give an example of a binary relation which is

- reflexive and symmetric, but not transitive;
- reflexive, but neither symmetric nor transitive;
- symmetric, but neither reflexive nor transitive;
- transitive, but neither reflexive nor symmetric.

Exc 2.2: For $a, b \in \mathbb{R}$, find

$$(a) \bigcup_{n=1}^{\infty} [a + \frac{1}{n}, b - \frac{1}{n}]; \quad (b) \bigcap_{n=1}^{\infty} (a + \frac{1}{n}, b - \frac{1}{n}).$$

Exc 2.3: Show that the set of real numbers \mathbb{R} is uncountable. (A look at the proof of Theorem 1 should help.)

Exc 2.4: Reconsider the example for the infimum and supremum from page 28. Redraw figure 2.4 using the Euclidean norm as the basis for the interpretation of the relation ' \leq '. Indicate $\max X$, $\min X$, $\sup X$, and $\inf X$ for the set $X = \{x \in \mathbb{R}^2 : \|x, (1, 1)\| = 1\}$.

Exc 2.5: (Sundaram 1996, 67:1.4) Let $\{x_k\}$ be a sequence in \mathbb{R}^n . Let $m(\cdot)$ be a rule assigning a value $m(k) \in \mathbb{N}$ to each $k \in \mathbb{N}$. and let $m(\cdot)$ be strictly increasing (i.e. $m(k+1) > m(k)$). Given $\{x_k\}$, we can define a **subsequence** $\{x_{m(k)}\}$ whose k^{th} element is the $m(k)^{\text{th}}$ element of the original sequence. I.e. a subsequence is any infinite subset of the original sequence that preserves the ordering of terms. The limits of these subsequences are called limit points.

Give an example of a sequence $\{x_k\}$ which has exactly n limit points, where $n \in \{1, 2, \dots\}$.

Exc 2.6: (Sundaram 1996, 67:1.6) Let $x, y \in \mathbb{R}^n$. Show that for any sequences $\{x_k\}, \{y_k\}$ such that $x_k \rightarrow x$ and $y_k \rightarrow y$, we have $\lim_{k \rightarrow \infty} d(x_k, y_k) = d(x, y)$, where d is the Euclidean metric on \mathbb{R}^n .

Exc 2.7: (Sundaram 1996, 68:1.17) Prove that the 'closed interval' $[0, 1]$ is, in fact, a closed set, and that the 'open interval' $(0, 1)$ is, in fact, an open set. Prove also that $[0, 1)$ and $(0, 1]$ are neither open nor closed.

Exc 2.8: (Sundaram 1996, 68:1.18) Consider the set $\mathbb{Z}_+ = \{0, 1, 2, \dots\}$ of non-negative integers viewed as a subset of \mathbb{R} . Is it closed, open or neither?

Exc 2.9: (Sundaram 1996, 68:1.19) Is \mathbb{R}^n viewed as a subset of itself an open set? Is it a closed set? Explain your answer.

Exc 2.10: (Sundaram 1996, 68:1.20) Give an example of an open set in \mathbb{R} that is not convex.

Exc 2.11: (Sundaram 1996, 68:1.21) Give an example of a compact set in \mathbb{R} that is not convex.

2.5 Properties of functions

Def: Let $f : X \mapsto \mathbb{R}$ be a real valued function. Then the closure of the set $\{x : f(x) \neq 0\}$ is called the **support** of f .

Def: A function $f : \mathbb{R} \mapsto \mathbb{R}$ is said to be **subadditive** if it satisfies the inequality $f(x + y) \leq f(x) + f(y)$ whenever $x, y \in \mathbb{R}$. It is called **superadditive** if $f(x + y) \geq f(x) + f(y)$.

2.5.1 Continuity

Def: A function $f : X \mapsto Y$ from $X \subset \mathbb{R}^n, Y \subset \mathbb{R}^m$ is said to be **continuous** at $x \in X$, if for all $\varepsilon > 0$, there exists a $\delta > 0$ such that $y \in X$ and

$$d(x, y) < \delta \implies d(f(x), f(y)) < \varepsilon.$$

f is continuous on X if it is continuous at every $x \in X$.




Figure 2.8: A continuous function (left) and a discontinuous function (right).

Exl: Let $f : \mathbb{R} \mapsto \mathbb{R}, f(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ 1 & \text{for } x > 0 \end{cases}$ $f(\cdot)$ is discontinuous only at $x = 0$. \triangleleft

Exl: Let $f(x) = \begin{cases} 1.2 & \text{for } x = 1 \\ x & \text{for } x \neq 1. \end{cases}$ $f(\cdot)$ is discontinuous only at $x = 1$. \triangleleft

Because continuity is such a fundamental concept, we give two alternative definitions.

Def: A real function $f : X \mapsto f(X)$ is **continuous** if $f^{-1}(Y)$ is open for each open set $Y \subset f(X)$.

Def: A function $f : X \mapsto Y$ from $X \subset \mathbb{R}^n, Y \subset \mathbb{R}^n$ is called (locally) **Lipschitz continuous** if for every compact subset $C \subset X$ there exists some real number λ such that for all $x, y \in C$

$$\|f(x) - f(y)\| \leq \lambda \|x - y\|. \quad (2.5.1)$$

As suggested by the following example, Lipschitz continuity is more stringent than continuity: Indeed every Lipschitz continuous function is continuous but the reverse is not true.

Exl: Evidently $f(x) = x^2$ is Lipschitz continuous, and so is the continuous but not (everywhere) differentiable function $f(x) = |x|$. In contrast, the continuous but not (everywhere) differentiable

function $f(x) = \begin{cases} 0 & \text{for } x < 0 \\ \sqrt{x} & \text{for } x \geq 0 \end{cases}$ is not Lipschitz continuous because its *rhs* slope at $x = 0$ is $+\infty$, and there exists no $\lambda \in \mathbb{R}$ that meets (2.5.1) for all y near $x = 0$. \triangleleft

Theorem 10. (Weierstrass) *If $f : S \mapsto \mathbb{R}$ is continuous on a closed & bounded set $S \subset \mathbb{R}^n$, then there exists a maximum and a minimum point for f in S .*

Proof. See (Sundaram 1996, 96). □

2.5.2 Differentiability

Def: A function $f : X \mapsto Y$ from $X \subset \mathbb{R}^n, Y \subset \mathbb{R}^m$ is said to be **differentiable** at $x \in X$ if there exists a unique $m \times n$ matrix A such that for all $\varepsilon > 0$, there exists a $\delta > 0$ such that $y \in X$ and

$$\|x - y\| < \delta \implies \frac{\|f(x) - (f(y) - A(x - y))\|}{\|x - y\|} < \varepsilon. \quad (2.5.2)$$

The matrix A is called the **derivative** of f (see figure 2.9). f is said to be differentiable on X if it is differentiable for all $x \in X$.




Figure 2.9: Illustration of the derivative of $y = f(x)$.

Rem: Intuitively, the derivative of f at x is the best affine approximation $g(y) = Ay + b$ to $f(\cdot)$ at x . ‘Best’ means that

$$\frac{\|f(y) - g(y)\|}{\|y - x\|} \xrightarrow{y \rightarrow x} 0.$$

To see this, we set $g(y) = Ay + b = f(x)$ and obtain $b = f(x) - Ax$. Thus $g(y) = Ay + f(x) - Ax = A(y - x) + f(x)$. Therefore $\|f(y) - g(y)\| = \|f(y) - A(y - x) + f(x)\|$ which is what we used in the above definition (2.5.2).

Exl: Let $f : \mathbb{R} \mapsto \mathbb{R}$, $f(x) = \begin{cases} -x & \text{for } x \leq 0 \\ x & \text{for } x > 0 \end{cases}$ $f(\cdot)$ is not differentiable at $x = 0$ (note that we require the defining property for all ε) and differentiable everywhere else. f is continuous everywhere. \triangleleft

Def: Let $f : X \mapsto Y$, $X \subset \mathbb{R}^n, Y \subset \mathbb{R}^m$, be differentiable. Then the derivative of f , denoted Df (or $f'(\cdot)$ or $\nabla f(\cdot)$), itself forms a function $Df : X \mapsto X \times Y$. If this function Df is continuous everywhere, f is called **continuously differentiable**; we write f is C^1 . If f is twice continuously differentiable, we write f is C^2 and so on. A function that is sufficiently continuously differentiable (at the extreme C^∞) is called **smooth**.

Rem: If two functions $f(\cdot)$ and $g(\cdot)$ are differentiable, their sum $f(\cdot) + g(\cdot)$ is differentiable as well. Moreover, their product $f(\cdot) \circ g(\cdot)$ is also differentiable (chain rule).

Rem: A function must be continuous in order to be differentiable. The reverse, however, needs not to hold. For example the trajectories created by a Brownian motion (i.e. a random walk in continuous time) is—with probability 1—continuous everywhere but nowhere differentiable.

Rem: The derivative of $f(x_1, \dots, x_n)$ with respect to x_i while holding all other $x_{j \neq i}$ constant is called the **partial derivative** of f wrt x_i , symbolically $\frac{\partial f}{\partial x_i}$. It looks at variations only along the i^{th} dimension of \mathbf{x} .

Rem: The derivative of $f(x_1, \dots, x_n)$ with respect to x_i while allowing all other $x_{j \neq i}$ to vary is called the **total derivative** of f wrt x_i , symbolically

$$\frac{df}{dx_i} = \sum_{j=1}^n \frac{\partial f}{\partial x_j}(x_1, \dots, x_n) \frac{dx_j}{dx_i}$$

where the variation $\frac{dx_i}{dx_i} = 1$. The total derivative looks at variations along all dimensions of \mathbf{x} simultaneously (see figure 2.10).

Def: The **gradient** of a scalar function $f(\mathbf{x})$; $\mathbf{x} \in \mathbb{R}^n$ is defined as

$$\nabla f(\mathbf{x}) = \left(\frac{\partial f(\mathbf{x})}{\partial x_1}, \dots, \frac{\partial f(\mathbf{x})}{\partial x_n} \right).$$

2.5.3 Integrability

Def: Consider a function $f : [a, b] \mapsto \mathbb{R}$, $a < b$. If $a = x_0 < x_1 < x_2 < \dots < x_n = b$, then $\mathcal{Z} = \{x_0, x_1, \dots, x_n\}$ is called a **partition** of the interval $[a, b]$ with **gap** $\|\mathcal{Z}\| = \max \{x_k - x_{k-1} \mid 1 \leq k \leq n\}$. If the gaps $\xi_k \in [x_k - x_{k-1}]$ for all $1 \leq k \leq n$, then the ξ are called intermediary points of \mathcal{Z} .

The sum

$$S(f, \xi, \mathcal{Z}) = \sum_{k=1}^n f(\xi_k)(x_k - x_{k-1})$$

is called the **Riemann sum**.

Def: If ξ is such that $f(\xi_k) = \inf f([x_{k-1}, x_k])$ for all $1 \leq k \leq n$, then we call $L(f, \mathcal{Z}) \equiv S(f, \xi, \mathcal{Z})$ the **lower Darboux sum**. If ξ is such that $f(\xi_k) = \sup f([x_{k-1}, x_k])$ for all $1 \leq k \leq n$, then we call $U(f, \mathcal{Z}) \equiv S(f, \xi, \mathcal{Z})$ the **upper Darboux sum**.




Figure 2.10: Illustration of the difference between a partial and a total derivative.

Def: A function $f : [a, b] \mapsto \mathbb{R}$, $a < b$, is called **Riemann integrable** if $\lim_{n \rightarrow \infty} S(f, \xi_n, \mathcal{Z}_n)$ exists for any sequence of partitions $\{\mathcal{Z}_n\}$ and intermediate points $\{\xi_n\}$ with $\lim_{n \rightarrow \infty} \{\|\mathcal{Z}_n\|\} \rightarrow 0$.

Def: If a function $f : [a, b] \mapsto \mathbb{R}$, $a < b$, is Riemann integrable, the limit

$$\lim_{n \rightarrow \infty} S(f, \xi_n, \mathcal{Z}_n) \text{ is called } \int_a^b f(x) dx.$$

Unfinished. Some rules, a statement and proof of the fundamental theorem of calculus, and a discussion of improper integrals remain to be written. But I guess I'll keep the Lebesgue story with the measure theory ...

2.5.4 Convexity, concavity

Def: Let $S \subset \mathbb{R}^n$ be convex, $\mathbf{x}, \mathbf{y} \in S$. A function $f : S \mapsto \mathbb{R}$ is **concave** on S if for all $\lambda \in [0, 1]$

$$f(\lambda \mathbf{x} + (1 - \lambda) \mathbf{y}) \geq \lambda f(\mathbf{x}) + (1 - \lambda) f(\mathbf{y}).$$

Again the function is **strictly concave** if the permissible unit interval for λ is open. $f(\mathbf{x})$ is (**strictly**) **convex** if $-f(\mathbf{x})$ is (strictly) concave.

Def: Let $f : \mathcal{D} \mapsto \mathbb{R}$ and $a \in \mathbb{R}$. Then the **upper level set** (or upper contour set, or 'better-than' set) of f is defined as the set

$$U_f(a) = \{x \in \mathcal{D} \mid f(x) \geq a\}.$$

Def: Let $S \subset \mathbb{R}^n$ be convex & $\mathbf{x} \in S$. A function $f(\mathbf{x})$ is **quasiconcave** if its upper level (contour) set $P = \{\mathbf{x} : f(\mathbf{x}) \geq \text{const}\}$ is convex (see figure 2.12). It is **strictly quasiconcave** if the function




Figure 2.11: (Left) a (strictly) concave function $f(\cdot)$ and (right) a (strictly) convex function $u(\cdot)$.

has no straight stretches (see figure 2.13).¹⁹




Figure 2.12: A quasiconcave function $f(\cdot)$ of two variables x_1, x_2 .

Theorem 11. *A quasiconcave (quasiconvex) function attains a maximum (minimum) on a compact subset of \mathbb{R}^n . The maximum (minimum) is unique if the function is strictly quasiconcave (strictly quasiconvex).*

Proof. See (Sundaram 1996, 214). □

2.5.5 Other properties

Def: A function $f(x_1, \dots, x_n)$ is **homogeneous** of degree r in $D \subset \mathbb{R}^n$ if for all $t \in \mathbb{R}_+$ and all $\mathbf{x} \in D$ ²⁰

$$f(tx_1, \dots, tx_n) \equiv t^r f(x_1, \dots, x_n). \tag{2.5.3}$$

¹⁹ Quasiconcavity (quasiconvexity) is not the same as pseudoconcavity (pseudoconvexity).

²⁰ If a function is homogeneous of degree 1, we often just call it homogeneous.




Figure 2.13: A strictly quasiconcave function $f(\cdot)$ over a convex set.

Def: D is called an (open) **cone** (as illustrated in figure 2.14) if $tx \in D$ whenever $x \in D$ and $t > 0$.




Figure 2.14: A function which is homogeneous of degree 1.

Theorem 12. (Euler) Let $f(\mathbf{x}) = f(x_1, \dots, x_n)$ be C^1 . $f(\mathbf{x})$ is homogeneous of degree r in the open cone D iff for all $\mathbf{x} \in D \subset \mathbb{R}^n$

$$r f(\mathbf{x}) = \sum_{i=1}^n x_i f'_i(\mathbf{x}) = \mathbf{x} \nabla f(\mathbf{x}).$$

Proof. The defining identity for homogeneity (2.5.3) of degree r says $f(t\mathbf{x}) \equiv t^r f(\mathbf{x})$. We differentiate wrt t and obtain

$$\mathbf{x}\nabla f(t\mathbf{x}) = rt^{r-1}f(\mathbf{x}).$$

Which is for $t = 1$

$$\mathbf{x}\nabla f(\mathbf{x}) = rf(\mathbf{x}).$$

□

Exl: Let $g(x) = x^2$ which is homogeneous of degree 2. $g'(x) = 2x$. Let's look at $x = 4$. Euler's Theorem tells us that

$$xg'(x) = 4 \cdot 2 \cdot 4 = 2 \cdot 16 = rg(x)$$

which can be very useful if we know the derivative of a function but not the function itself. ((Mas-Colell, Whinston, and Green 1995, 928) contains more on homogeneous functions and the Euler Theorem.) ◁

Rem: If $f(\mathbf{x})$ is homogeneous of degree r , then $\frac{\partial f}{\partial x_i}$ is homogeneous of degree $r - 1$.

Def: $f(\mathbf{x}) = f(x_1, \dots, x_n)$ is **homothetic** in the cone D if for all $\mathbf{x}, \mathbf{y} \in D$ and all $r > 0$

$$f(\mathbf{x}) = f(\mathbf{y}) \implies f(r\mathbf{x}) = f(r\mathbf{y}).$$




Figure 2.15: A homothetic function $f(\cdot)$: If \mathbf{x} and \mathbf{y} are on the same level set, then so are $r\mathbf{x}$ and $r\mathbf{y}$.

Rem: Let $f(\mathbf{x})$ be a continuous, homothetic function (see figure 2.15) defined on the cone D . Assume f is strictly increasing along each ray in D , i.e. for each $\mathbf{x}_0 \in D$, $f(t\mathbf{x}_0)$ is a strictly increasing function of t . Then there exist a homogeneous function g and a strictly increasing function F such that $f(\mathbf{x}) = F(g(\mathbf{x}))$ for all $\mathbf{x} \in D$.

Def: A **correspondence** is a set-valued function $\mathbf{y} = f(x) \subset Y$.

2.6 Linear functions on \mathbb{R}^n

Def: The **zero vector** $\theta = \{0, \dots, 0\} \in \mathbb{R}^n$ is an n -dimensional vector consisting of n zeros.

Def: The **unit vector** $e_i \in \mathbb{R}^n$ is an n -dimensional vector consisting solely of zeros except for position i , with $1 \leq i \leq n$, where it is 1.

Def: A $m \times n$ **matrix** A is a table of real numbers with m rows and n columns. The number in row i and column j is denoted a_{ij} and one writes $A = ((a_{ij})_{i=1}^n)_{j=1}^m$.

We justify the title of this section with the following

Proposition 13. *A function $f : \mathbb{R}^n \mapsto \mathbb{R}^m$ is linear iff there exists a $m \times n$ matrix A such that $f(\mathbf{x}) = A\mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^n$.*

Proof. Suppose f is linear and let $x = (x_i e_i)_{i=1}^n = \sum_i x_i e_i$. For any such x we have

$$f(x) = f\left(\sum_i x_i e_i\right) = \sum_i f(x_i e_i) = \sum_i x_i f(e_i) = \sum_i x_i a_i,$$

where $a_i = f(e_i) \in \mathbb{R}^m$. Letting A be an $m \times n$ matrix with columns a_1, \dots, a_n , we thus obtained $f(x) = Ax$. Conversely, suppose A is an $m \times n$ matrix such that $f(\mathbf{x}) = A\mathbf{x}$ for all $x \in \mathbb{R}^n$. Then it is easy to see that f is linear. \square

2.6.1 Linear dependence

Def: A **linear subspace** of \mathbb{R}^n is the subset L that contains all linear combinations of all of its points.

Def: A set of n vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^n$ is called **linearly dependent** iff there exists a set of scalars k_1, \dots, k_n , not all of which are zero, such that $\sum \mathbf{v}_k = \theta$. It is called **linearly independent** if this sum is zero only for all $k_i = 0$.

Thus, the collection of vectors $\mathbf{v} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\} \in \mathbb{R}^n$ is linearly independent iff

$$\sum_{i=1}^n k_i \mathbf{v}_i = \theta \quad \Rightarrow \quad k_i = 0 \quad \forall i.$$

Exl: Let $v_1 = (1, 1)$ and $v_2 = (2, 2)$. Then $\mathbf{k} = (2, -1)$ results in

$$\sum \mathbf{v}_k = \begin{pmatrix} 1 \\ 1 \end{pmatrix} (2) + \begin{pmatrix} 2 \\ 2 \end{pmatrix} (-1) = 0.$$

Clearly, v_1 and v_2 are linearly dependent. \triangleleft

Def: The set of all vectors that can be written as linear combinations of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ is called the **span** of those vectors. The set A of n linearly independent n -vectors spans the \mathbb{R}^n (while a set of n linearly dependent vectors spans only a linear subspace $X \subset \mathbb{R}^n$). The set A forms a **basis** for \mathbb{R}^n .

Def: The vectors $\mathbf{b}_1, \dots, \mathbf{b}_d$ form a **basis** of the linear subspace L if they:

1. are linearly independent and
2. span L such that $L = \{\mathbf{y} : \mathbf{y} = c_1 \mathbf{b}_1 + \dots + c_d \mathbf{b}_d\}$.

Exl: The vectors $v_1 = (1, 0)$ and $v_2 = (0, 1)$ form a basis of \mathbb{R}^2 . \triangleleft

Rem: Any two bases for a given subspace X have the same number of elements.

Def: The **dimension** of the linear subspace L is the maximum number of independent vectors spanning L .

Def: The **rank** of a $m \times n$ matrix A whose determinant $|A|$ is non-zero is the $\min(\text{column dimension, row dimension})$ and denoted $\text{rank}(A)$. If $|A| = 0$ we set $\text{rank}(A) = 0$.

Rem: If A is an $n \times n$ matrix, then the equation system $A\mathbf{x} = \mathbf{b}$ has a unique solution \mathbf{x} for every $\mathbf{b} \in \mathbb{R}^n$ iff $\text{rank}(A) = n$ (i.e. A has independent columns).

Def: If the inverse of a square matrix A^{-1} does not exist, A is called **singular**.

Def: The **trace** of a square $n \times n$ matrix A is the sum of its diagonal entries

$$\text{tr}(A) = \sum_{i=1}^n a_{ii}.$$

2.6.2 Determinants

Let us focus on the special case of linear functions $f : \mathbb{R}^n \mapsto \mathbb{R}^n$.

Def: The **determinant** of a 2×2 matrix

$$\begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix}$$

is defined as $a_{11}a_{22} - a_{21}a_{12}$. If $A = (a_{ij})_{n \times n}$ is an $n \times n$ matrix, the determinant of A is a number

$$|A| = a_{i1}A_{i1} + \dots + a_{in}A_{in} = \sum_{j=1}^n a_{ij}A_{ij}$$

where A_{ij} , the cofactor of the element a_{ij} , is²¹

$$A_{ij} = (-1)^{i+j} \begin{vmatrix} a_{11} & \cdots & a_{1j} & \cdots & a_{1n} \\ \vdots & & \vdots & & \vdots \\ a_{i1} & \cdots & a_{ij} & \cdots & a_{in} \\ \vdots & & \vdots & & \vdots \\ a_{n1} & \cdots & a_{nj} & \cdots & a_{nn} \end{vmatrix}.$$

Rem: As illustrated in figure 2.16, we can interpret the determinant as the absolute value of the volume of the hypercube spanned by the generating vectors.

Rem: The inverse of a $n \times n$ matrix A exists iff $|A| \neq 0$ (that is, A has full rank²² n).

²¹ I do not know how to cross out the row vector $(a_{i1}, \dots, a_{ij}, \dots, a_{in})$ and the column vector $(a_{1j}, \dots, a_{ij}, \dots, a_{nj})^T$. However, in the following equation, they should be crossed out—all I've managed so far are the lines between columns/rows 2 and 3.

²² The rank of a matrix is defined in definition 2.6.1.




Figure 2.16: The geometrical interpretation of the determinant.

Rem: Given the equation system $\bar{x} = A^{-1}d$, the solutions for $j = 1, \dots, n$ are given using **Cramer's rule** by

$$\bar{x}_j = \frac{|A_j|}{|A|} = \frac{1}{|A|} \begin{vmatrix} a_{11} & \dots & a_{1j-1} & d_1 & a_{1j+1} & \dots & a_{1n} \\ a_{21} & \dots & a_{2j-1} & d_2 & a_{2j+1} & \dots & a_{2n} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ a_{n1} & \dots & a_{nj-1} & d_n & a_{nj+1} & \dots & a_{nn} \end{vmatrix}$$

2.6.3 Eigenvectors and eigenvalues

Let $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ be some linear function and $x = f(x)$ of f (which always exists because of continuity of f).²³ Clearly, the zero vector θ is always a ‘trivial’ fixed point under any such linear mapping—eigenvectors define the collection of all other fixed points of f .

Def: $\mathbf{x} \in \mathbb{R}^n$ is called an **eigenvector** of f if $\mathbf{x} \neq \theta$ and $f(\mathbf{x}) = k\mathbf{x}$ for some scalar k . The associated k is called an **eigenvalue** of f .

Rem: For any scalar k , let $V(k)$ denote the nonempty subspace of \mathbb{R}^n for elements $\mathbf{x} \in V(k)$ of which $f(\mathbf{x}) = k\mathbf{x}$. If θ is the only element of $V(k)$, then k is *not* an eigenvalue of f , otherwise $V(k)$ contains an eigenvector associated with k .

Def: If k is an eigenvalue of f , then $V(k)$ is called the **eigenspace** of k .

Theorem 14. Suppose $\{k_1, \dots, k_n\}$ are distinct eigenvalues of f and let $\theta \neq \mathbf{x}_i \in V(k_i)$ for $i = 1, \dots, n$. Then the collection of eigenvectors $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is linearly independent.

Proof. See (Weibull 2001, 102f). □

Corollary 15. A linear function $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ can have at most n distinct eigenvalues.

Theorem 16. (Perron-Frobenius) Every positive square matrix has at least one positive eigenvalue and an accompanying non-negative eigenvector.

Proof. See (Weibull 2001, 103f). □

²³ If you need details, please consider section 7.

After having established that there are eigenvalues out there, let us now try to find them. Consider any $n \times n$ matrix A . \mathbf{x} is an eigenvector and k an eigenvalue iff $(\mathbf{x}, k) \in \mathbb{R}^{n+1}$ solves the system of n equations $A\mathbf{x} = k\mathbf{x}$, or, equivalently, the system

$$(kI - A)\mathbf{x} = \theta.$$

Notice that we have one more unknown than equations in this system. For fixed k , the matrix $B = kI - A$ represents the linear function $g : \mathbb{R}^n \mapsto \mathbb{R}^n$ defined by $g(x) = (kI - A)\mathbf{x}$. The equation $g(x) = \theta$ has a solution $x \neq \theta$ iff g is *not* a bijection, which is the case iff

$$\det[kI - A] = 0. \tag{2.6.1}$$

(2.6.1) is called the **characteristic equation** associated with the square matrix A . It is one equation in one unknown, k .

Proposition 17. *A scalar k is an eigenvalue iff it satisfies the characteristic equation.*

Proof. See (Weibull 2001, 104). □

The left-hand side of (2.6.1) is a polynomial of degree n in k

$$\det[kI - A] = p(k) = k^n + c_{n-1}k^{n-1} + \dots + c_1k + c_0.$$

Allowing for both real and complex roots, any such polynomial p of degree n can be factorised on the form

$$p(k) = (k - k_1)(k - k_2) \dots (k - k_n)$$

for not necessarily distinct (real or complex) numbers k_1, \dots, k_n . These are called the **algebraic roots** of the characteristic equation. Hence every square matrix A has exactly n eigenvalues.

Exl: Consider the following four examples of a function mapping \mathbb{R}^2 into itself:

1. $e(\mathbf{x}) = \theta \forall \mathbf{x} \in \mathbb{R}^2$,
2. $f(\mathbf{x}) = (3x_1, 2x_2) \forall \mathbf{x} \in \mathbb{R}^2$,
3. $g(\mathbf{x}) = (-x_2, x_1) \forall \mathbf{x} \in \mathbb{R}^2$,
4. $h(\mathbf{x}) = (x_1, 0) \forall \mathbf{x} \in \mathbb{R}^2$.

In matrix representation, these linear functions are

$$E = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad F = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}, \quad G = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad H = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

Then the characteristic equations are

$$\begin{aligned} k^2 &= 0, & \text{roots: } k_1 &= k_2 = 0, \\ (k - 3)(k - 2) &= 0, & \text{roots: } k_1 &= 3, k_2 = 2, \\ k^2 + 1 &= 0, & \text{roots: } k_{1,2} &= \pm i, \\ (k - 1)k &= 0, & \text{roots: } k_1 &= 1, k_2 = 0. \end{aligned}$$

Notice that $e(\cdot)$ ‘collapses,’ $f(\cdot)$ ‘dilutes,’ $g(\cdot)$ ‘rotates,’ and $h(\cdot)$ ‘projects’ their respective ranges into their domains. \triangleleft

Proposition 18.

$$\prod_{i=1}^n k_i = \det(A) \quad \text{and} \quad \sum_{i=1}^n k_i = \text{tr}(A).$$

2.6.4 Quadratic forms

Def: A **polynomial** of degree n in k is an expression

$$p(k) = k^n + c_{n-1}k^{n-1} + \dots + c_1k + c_0.$$

Def: A **form** is a polynomial expression in which each component term has a uniform degree.

Def: A **quadratic form** on \mathbb{R}^n is a function $g_A : \mathbb{R}^n \mapsto \mathbb{R}$ of the form

$$g_A(\mathbf{x}) = \mathbf{x}'A\mathbf{x} = \sum_{i=1}^n \sum_{j=1}^n a_{ij}x_i x_j,$$

$\mathbf{x} \in \mathbb{R}^n$, $A = (a_{ij})$ being a symmetric $n \times n$ matrix. Since A fully describes the quadratic form, sometimes A itself is referred to as quadratic form.

Exl: The polynomial $f(x,y)=x^2 + y^2$ is a quadratic form with $A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. \triangleleft

Rem: If $f : \mathbb{R}^n \mapsto \mathbb{R}$ is C^2 and $\mathbf{x} \in \mathbb{R}^n$, then the matrix of 2^{nd} partials $D^2f(\mathbf{x})$ (called the **Hessian**) defines a quadratic form.

Theorem 19. (Schwartz’s Theorem) Given a C^2 function $f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^n$, the order of computing the second partial derivatives is immaterial. Formally

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \frac{\partial^2 f}{\partial x_j \partial x_i} \quad \forall i, j = 1, \dots, n \quad (2.6.2)$$

Rem: Hence the Hessian is symmetric.

Def: A quadratic form is said to be²⁴

1. **positive definite** if $\mathbf{x}'A\mathbf{x} > \mathbf{0}$ for all $\mathbf{x} \neq \mathbf{0} \in \mathbb{R}^n$,
2. **positive semi-definite** if $\mathbf{x}'A\mathbf{x} \geq \mathbf{0}$ for all $\mathbf{x} \neq \mathbf{0} \in \mathbb{R}^n$,
3. **negative definite** if $\mathbf{x}'A\mathbf{x} < \mathbf{0}$ for all $\mathbf{x} \neq \mathbf{0} \in \mathbb{R}^n$,
4. **negative semi-definite** if $\mathbf{x}'A\mathbf{x} \leq \mathbf{0}$ for all $\mathbf{x} \neq \mathbf{0} \in \mathbb{R}^n$, and
5. **indefinite** if none of the above applies.

Rem: If

1. $\mathbf{x}'A\mathbf{x}$ is positive definite, then $a_{ii} > \mathbf{0}$ for $i = 1, \dots, n$.
2. $\mathbf{x}'A\mathbf{x}$ is positive semi-definite, then $a_{ii} \geq \mathbf{0}$ for $i = 1, \dots, n$.

²⁴ Recall that $\mathbf{x} \neq \mathbf{0}$ means that at least one component of the vector must be different from zero.

3. $\mathbf{x}'A\mathbf{x}$ is negative definite, then $a_{ii} < \mathbf{0}$ for $i = 1, \dots, n$.
4. $\mathbf{x}'A\mathbf{x}$ is negative semi-definite, then $a_{ii} \leq \mathbf{0}$ for $i = 1, \dots, n$.

Rem: If

1. $\mathbf{x}'A\mathbf{x}$ is positive definite, then all eigenvalues of A $\lambda_i > \mathbf{0}$ for $i = 1, \dots, n$.
2. $\mathbf{x}'A\mathbf{x}$ is positive semi-definite, then all eigenvalues of A $\lambda_i \geq \mathbf{0}$ for $i = 1, \dots, n$.
3. $\mathbf{x}'A\mathbf{x}$ is negative definite, then all eigenvalues of A $\lambda_i < \mathbf{0}$ for $i = 1, \dots, n$.
4. $\mathbf{x}'A\mathbf{x}$ is negative semi-definite, then all eigenvalues of A $\lambda_i \leq \mathbf{0}$ for $i = 1, \dots, n$.
5. $\mathbf{x}'A\mathbf{x}$ is indefinite, then there is at least one positive and one negative eigenvalue of A .

Rem: Assume the function f is C^2 . Then $f : A \mapsto \mathbb{R}$ is (for every $\mathbf{x} \in A$)

1. strictly convex *iff* the Hessian $D^2f(\mathbf{x})$ is positive semi-definite,
2. convex *iff* the Hessian $D^2f(\mathbf{x})$ is positive definite,
3. concave *iff* the Hessian $D^2f(\mathbf{x})$ is negative definite, and
4. strictly concave *iff* the Hessian $D^2f(\mathbf{x})$ is negative semi-definite.

Exl: $A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ for $\mathbf{x} \in \mathbb{R}^2$ represents $\begin{bmatrix} x_1 & x_2 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$ which is the quadratic form $\mathbf{x}'A\mathbf{x} = x_1^2 + x_2^2$. Since $2x_1^2 + 2x_2^2 > 0$ whenever $\mathbf{x} \neq 0$, it is positive definite. \triangleleft

Exl: $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ for $\mathbf{x} \in \mathbb{R}^2$ represents the positive semi-definite quadratic form $\mathbf{x}'A\mathbf{x} = x_1^2$ since $x_1^2 = 0$ for e.g. $\mathbf{x} = (0, 1) \neq 0$. \triangleleft

Exl: $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ is indefinite. \triangleleft

Rem: The reason why we are so concerned with quadratic forms is that the Hessian that we create by taking second derivatives of our objective (of general functional form) at a critical point describes a quadratic form. If we find that this Hessian is locally (i.e. around the critical point) negative (semi-)definite, our necessary conditions are sufficient for a local maximum as well. This is much easier to check than to examine all critical points of a general function employing the (general) approaches we used in the previous chapters.

Exercises 2

Exc 2.12: We say that a function $f(x_1, \dots, x_n)$ is *non-decreasing* if $x'_i \geq x_i$ for each i implies that $f(x'_1, \dots, x'_n) \geq f(x_1, \dots, x_n)$, is *increasing* if $x'_i > x_i$ for each i implies $f(x'_1, \dots, x'_n) > f(x_1, \dots, x_n)$ and is *strictly increasing* if $x'_i \geq x_i$ and $x'_j > x_j$ for at least one j implies that $f(x'_1, \dots, x'_n) > f(x_1, \dots, x_n)$. Show that if f is non-decreasing and strictly concave then it must be strictly increasing.

Exc 2.13: Prove the remark on page 47: If $f(\mathbf{x})$ is homogeneous of degree r , then $\frac{\partial f}{\partial x_i}$ is homogeneous of degree $r - 1$.

Exc 2.14: (Sundaram 1996, 68:1.25) Let $B \subset \mathbb{R}^2$ be as follows

$$B = \{(x, y) \in \mathbb{R}^2 : y = \sin \frac{1}{x}, x > 0 \cup \{(0, 0)\}\}.$$

Is B closed? open? bounded? compact?

Exc 2.15: (Sundaram 1996, 69:1.26) Let $A \subset \mathbb{R}^2$ be defined as

$$B = \{(x, y) \in \mathbb{R}^2 : 1 < x < 2 \text{ and } y = x\}.$$

Is A open? bounded? compact?

Exc 2.16: (Sundaram 1996, 70:1.35) Show that the set $S \subset \mathbb{R}$ is convex *iff* it is an interval, i.e. it is of the form $[a, b]$, $[a, b)$, $(a, b]$, or (a, b) . (We do not preclude the possibility that $a = -\infty$ and/or $b = \infty$.)

Exc 2.17: (Sundaram 1996, 70:1.36) Show that if $S \subset \mathbb{R}^n$ is convex (i.e. the convex combination of any two points in S is also in S), then the convex combination of any finite collection of points from S is also in S .

Exc 2.18: (Sundaram 1996, 70:1.37) Give an example of a bounded and convex set $S \subset \mathbb{R}$ such that $\sup S \in S$ but $\inf S \notin S$.

Exc 2.19: (Sundaram 1996, 71:1.46) Suppose $f : \mathbb{R}^n \mapsto \mathbb{R}$ is a continuous function. Show that the set

$$\{\mathbf{x} \in \mathbb{R}^n \mid f(\mathbf{x}) = 0\}$$

is a closed set.

Exc 2.20: (Sundaram 1996, 71:1.47) Let $f : \mathbb{R}^n \mapsto \mathbb{R}$ be defined by

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } 0 \leq x \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

Find an open set O such that $f^{-1}(O)$ is not open and find a closed set C such that $f^{-1}(C)$ is not closed.

Exc 2.21: (Sundaram 1996, 71:1.48) Give an example of a function $f : \mathbb{R} \mapsto \mathbb{R}$ which is continuous at exactly two points (say, at 0 and 1), or show that no such function can exist.

Chapter 3

First applications of fundamentals

Reading: Alternative derivations and further examples are contained in (Varian 1992b), (Mas-Colell, Whinston, and Green 1995), (Kreps 1998) and, on a more introductory level, (Nicholson 1990).

As promised in the last chapter, we will do a set of first *economic* applications in this chapter: We start with some Game Theory and then we will employ the Envelope Theorem to derive a number of results that prove to be very useful in microeconomics. Again we will have to start with a host of definitions to introduce the game theoretic ideas that we will use. But we will do this in a quite relaxed fashion—relying on the assumption that each of you has already had at least some exposure to game theory before. If this is not the case, both (Gibbons 1992) and (Fudenberg and Tirole 1991) give much more detail than you will need for our little example.

In our first application we shall prove that it is not rational to play strictly dominated strategies. This may seem obvious at first sight, but at the very least it will give you a clear understanding of what makes a proof. Or, more importantly maybe, what does not constitute a proper proof.

The second application, the Envelope Theorem, will haunt you throughout the program and pop up in Macro as well as in Industrial Organisation, Game Theory and Micro. So it really can be argued that it is a good idea to understand it now and get it behind you. Another argument for this strategy is that no-one else will have the time to go as slowly over the theorem as we will—so it will only become more difficult.

3.1 Separating hyperplanes

Def: Let $\mathbf{p} \neq 0$ be a point in \mathbb{R}^n and a a scalar. Then the set $H(\mathbf{p}, a) = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{p} \cdot \mathbf{x} = a\}$ is called a **hyperplane** in \mathbb{R}^n .

Exl: Consider our usual micro budget sets. For income m , two goods \mathbf{x} , and prices \mathbf{p} , the budget set \mathcal{B} is usually defined as

$$\mathcal{B} = \{\mathbf{x} \in \mathbb{R}^2 \mid p_1x_1 + p_2x_2 \leq m\}.$$

(As usual, we replace the inequality with an equality if we refer to the frontier only.) According

to the above definition, this defines the budget frontier as a hyperplane: $\mathbf{p}\mathbf{x} = m$. We show an example in figure 3.1. \triangleleft




Figure 3.1: A price vector and the budget hyperplane it defines.

Def: A set X is bounded by a hyperplane if the whole of X lies (weakly) on one side of the plane. I.e. $\mathbf{p} \cdot \mathbf{x} \leq a$ for all $\mathbf{x} \in X$. In this case, $H(\mathbf{p}, a)$ is called a **supporting hyperplane** for X (see figure 3.2).

Rem: So a hyperplane in \mathbb{R}^3 is a plane and a hyperplane in \mathbb{R}^2 is simply a line. In general a hyperplane in \mathbb{R}^n is a set in \mathbb{R}^{n-1} .

Exl: Let $X = \{\mathbf{x} = (x_1, x_2) \in \mathbb{R}_+^2 \mid x_1x_2 \geq 1\}$, $\mathbf{p} = (1, 1)$, and $a = 2$. The set X is bounded by the rectangular hyperbola $x_1x_2 = 1$. Then the hyperplane $H(\mathbf{p}, a) = \{\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2 \mid x_1+x_2 = 2\}$ shown in figure 3.2 bounds X since the $H((1, 1), 2)$ and X share the point $\mathbf{p} = (1, 1)$ while X lies to the northeast of $H((1, 1), 2)$ everywhere else. \triangleleft




Figure 3.2: A supporting hyperplane $H((1, 1), 2)$.

Def: Two sets $X, Y \subset \mathbb{R}^n$ are said to be **separated** by a hyperplane if the whole of X lies (weakly) on one side of the plane while the whole of Y lies (weakly) on the other side. I.e. for all

$\mathbf{x} \in X, \mathbf{y} \in Y$ we have $\mathbf{p} \cdot \mathbf{x} \leq a$ while $\mathbf{p} \cdot \mathbf{y} \geq a$.

Theorem 20. (Minkowski) Suppose that (i) $X, Y \subset \mathbb{R}^n$, (ii) X, Y are convex, and (iii) the interior of the sets does not intersect: $\text{int}(X) \cap \text{int}(Y) = \emptyset$. Then there is some non-zero $\boldsymbol{\alpha} \in \mathbb{R}^n$ and some $k \in \mathbb{R}$ such that

$$\begin{aligned} \mathbf{x} \in X &\implies \boldsymbol{\alpha} \mathbf{x} \leq k \\ \mathbf{x} \in Y &\implies \boldsymbol{\alpha} \mathbf{x} \geq k \end{aligned}$$

3.2 The irrationality of playing strictly dominated strategies

In this section we shall put the instruments we developed to some use. First, however, we require some additional economics.

Def: S_i is the non-empty and finite set of pure strategies for player $i \in I = \{1, \dots, n\}$.

Def: Player i 's **pure strategy** is $s_i \in S_i$; a **pure strategy profile** for all players is denoted by $\mathbf{s} \in S_1 \times \dots \times S_n$.

Def: Player i 's **mixed strategy** is a probability distribution over her set of pure strategies $\sigma_i \in \Delta(S_i)$. A mixed strategy profile is denoted by $\boldsymbol{\sigma} \in \Delta(S_1) \times \dots \times \Delta(S_n)$.¹

Rem: It is worthwhile to reflect a moment on the difference between $\Delta(S_1 \times \dots \times S_n)$ and $\Delta(S_1) \times \dots \times \Delta(S_n)$. The first can represent correlated beliefs while the second is necessarily a product belief.

Def: A profile of all of i 's **opponents' strategies** (excluding i 's strategy) is denoted by $\mathbf{s}_{-i} \in S_{-i}$ or $\boldsymbol{\sigma}_{-i} \in \Delta(S_1) \times \dots \times \Delta(S_{i-1}) \times \Delta(S_{i+1}) \times \dots \times \Delta(S_n)$.

Def: Player i 's **payoff** (more precisely, her **expected utility**) $u : S_1 \times \dots \times S_n \mapsto \mathbb{R}$ from playing s_i while the opponents play \mathbf{s}_{-i} is denoted by $u_i(s_i, \mathbf{s}_{-i})$.

Def: Player i 's **belief** is a probability distribution over i 's opponents' strategies $\mu_i \in \Delta(S_{-i})$.

Exl: Let $S_1 = \{T, B\}$ and $S_2 = \{L, R\}$. A belief of player 1 might be that player 2 mixes at 50:50 between her 2 strategies. This belief would be represented by the mixed strategy $\sigma_1 = \frac{1}{2}L + \frac{1}{2}R$.

◁

Def: Strategy s_i is called **Bayesian rational** if there is some belief $\mu_i \in \Delta(S_{-i})$ such that $u_i(s_i, \mu_i) \geq u_i(s'_i, \mu_i)$ for all $s'_i \in S_i$.

Exl: Consider the payoff matrix for a 2-player game given in figure 3.2. (This representation is called normal form and is denoted by $G = \{2, S_1 \times S_2, u\}$.) By convention, player I is the row-player, II is the column player. The strategy sets are $S_1 = \{T, M_1, M_2, B\}$, $S_2 = \{L, R\}$. As we can see from the right panel, $\{T, M_1, M_2\}$ are Bayesian rational, but B is not. ◁

Def: A mixed strategy $\sigma_i \in \Delta(S_i)$ **strictly dominates** $s_i \in S_i$ if $u_i(\sigma_i, \mathbf{s}_{-i}) > u_i(s_i, \mathbf{s}_{-i})$ for all $\mathbf{s}_{-i} \in S_{-i}$.

Ass: We will confine ourselves to looking only at (i) a finite number of (ii) pure strategies. Nothing essential changes if these assumptions are relaxed.

¹ Nothing is said about why players should wish to play mixed strategies. In particular, we do *not* assume that they actually do this. But this is an entirely different story ...

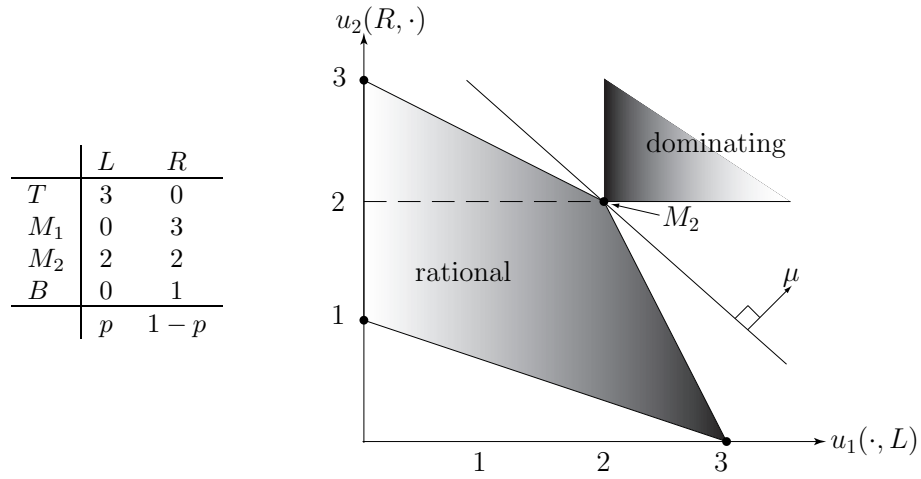


Figure 3.4: Playing a strictly dominated strategy cannot be rational.

1. Both X and Y are finite because we are in a 2 player finite game.
2. X is convex because it is generated by convex combinations (simplex). Y is convex because it is the strictly preferred set to M_2 —which is simply the positive orthant above M_2 .
3. $\text{int}(X) \cap \text{int}(Y) = \emptyset$ because M_2 is at the boundary of X and Y is the strictly dominating set of M_2 .

Thus we can apply Minkowski's Theorem and conclude that for some non-zero $\alpha \in \mathbb{R}^n$ and some $k \in \mathbb{R}$, we have:

$$\begin{aligned} \mathbf{x} \in X &\implies \alpha \mathbf{x} \leq k \\ \mathbf{x} \in Y &\implies \alpha \mathbf{x} \geq k. \end{aligned}$$

Let us apply this result to our setting:

$$\alpha(u_i(s_i^*, s_{-i}))_{s_{-i} \in S_{-i}} = k$$

because this vector is in both X and Y . Now for every $s_i \in S_i$ we have

$$\alpha(u_i(s_i, s_{-i}))_{s_{-i} \in S_{-i}} \leq k.$$

Therefore by playing α , s_i^* maximises expected utility.

Now we normalise the general vector α to a probability that sums to 1 by multiplying it with some constant non-zero c :

$$\mu_i = \alpha_i c.$$

Hence s_i^* maximises expected utility with beliefs μ_i . Consequently μ_i is Bayesian rational and we are done. □

3.3 The Envelope Theorem

Reading: (Milgrom 1999), (Mas-Colell, Whinston, and Green 1995, 964ff).

3.3.1 A first statement

In this section we shall examine a theorem that is particularly useful in the study of consumer and producer theory, as well as the analysis of incentive constraints in contract and game theory. It also has numerous applications in macroeconomics. It examines a setting where we have to maximise a function that depends on its variables (i.e. the argument(s) x that we variate in order to find the maximum of the function) and other parameters (i.e. the argument(s) a that we take as given). To be a bit more precise, the theorem describes conditions under which the value of a parameterised optimisation problem (which can be unconstrained or constrained) is a differentiable function of the parameters and, if this optimal value exists, it provides us with a formula for the derivative. Let us examine the canonical parameterised unconstrained optimisation problem

$$\max_{x,a} f(x(a), a). \quad (3.3.1)$$

Usually there will also be constraints of the form $g(x(a), a)$ present but for the first example we will only look at the unconstrained case. As you will see presently, the proof of the Envelope Theorem is very simple—more importantly, the theorem has a very clear intuition that we try to convey in figure 3.5 for the case $x, a \in \mathbb{R}$: Suppose our problem is to investigate the behaviour




Figure 3.5: The (unconstrained) Envelope Theorem: $f(x^*(\hat{r} + \varepsilon), \hat{r} + \varepsilon) \simeq f(x^*(\hat{r}), \hat{r} + \varepsilon)$.

of the maximum on the upper semi-circle shifted to the right by its radius. The variable of this problem is x and the parameter the radius r . For each optimisation step, the radius is kept constant. Let's start with a particular radius \hat{r} ; the maximum $x^*(\hat{r}) = f(x^*(\hat{r}), \hat{r}) = \hat{r}$. Now we change the parameter by a very small amount $\varepsilon > 0$ to $\hat{r} + \varepsilon$. The new maximum is

$x^*(\hat{r} + \varepsilon) = f(x^*(\hat{r} + \varepsilon), \hat{r} + \varepsilon) = \hat{r} + \varepsilon$. However at a maximum the function is flat (the first derivative is zero). Thus when we want to know by how much the maximised value of $f(x(r), r)$ has changed it does not matter (very much) whether or not we take account of how the maximiser $x(r)$ changes or not: $f(x(\hat{r}), \hat{r})$ and $f(x(\hat{r}), \hat{r} + \varepsilon)$ give almost the same values.

Exl: Maximise the unconstrained parameterised objective of

$$\boxed{f(x, a) = ax - x^2, \quad a \in \mathbb{R}_+} \quad (3.3.2)$$

where a is a constant parameter and x is the variable. The *foc* is just

$$f_x(x, a) = 0 \Leftrightarrow a - 2x = 0 \Leftrightarrow x^* = \frac{a}{2} \quad (3.3.3)$$

A family of functions $f(x, a)$ for different parameters a is shown in figure 3.6. Evidently, the optimal value of x depends on a , therefore it is justified to write it as $x(a)$. The optimal loci $x(a)^*$ are shown as dots.




Figure 3.6: The Envelope Theorem in the unconstrained setting.

Let's now define a *value* function $v(a)$ that parameterises solutions of (3.3.2)

$$v(a) \equiv f(x^*(a), a) = ax(a) - x(a)^2.$$

Let's examine the behaviour of the $v(a)$ under variation of a by taking total derivatives

$$\frac{dv(a)}{da} \equiv \frac{df(x^*(a), a)}{da} = f_a(x^*(a), a) = x(a) + ax'(a) - 2x(a)x'(a)$$

which equals

$$v'(a) = x(a) + x'(a)[a - 2x(a)] = x(a) + x'(a)[0] = x(a).$$

This employs the fact (3.3.3) that $f_x(x^*(a), a) = a - 2x(a) = 0$ because we are only looking at maximised values of x . Hence

$$v'(a) = \frac{\partial f(x^*(a), a)}{\partial a} :$$

the value function only depends on the partial $f_a(\cdot)$.

Integrating back (ignoring the constant), we obtain something like e.g.

$$v(a) = \frac{1}{2}x(a)^2$$

which looks indeed just like the function connecting the *foc*-dots in figure 3.6. \triangleleft

We have already seen a number of other applications of this setting. For instance when we optimised in the previous sections using the Lagrange method, we treated the prices p as parameters of the (Marshallian) demand functions $x^*(p, m)$ we derived. Actually, we treated prices as constant during the maximisation. We will now analyse how these derived functions react to small changes of the underlying parameters. Given their interpretation as shadow prices, we should expect the Lagrange multipliers λ to play some role in this.

Def: We denote the maximum value of $f(\mathbf{x}^*(\mathbf{a}), \mathbf{a})$, obtained as solution to problem (3.3.1), by $v(\mathbf{a})$ and call this function the **value** function.²

For all possible values of the parameter-vector \mathbf{a} , we can construct the value function $v(\mathbf{a}) = f^*(\mathbf{a})$ as an (upper) envelope of all the different $f(\mathbf{x}(\mathbf{a}), \mathbf{a})$ —which is where the name for the theorem comes from. For one variable $x(a)$ and one parameter a , the simplest statement of the Envelope Theorem is

$$\frac{dv(a)}{da} = \frac{\partial f(x^*(a), a)}{\partial a}.$$

The equivalent statement for constrained problems can be written in terms of the Lagrangian as

$$\frac{dv(a)}{da} = \frac{\partial \mathcal{L}(x^*(a), a, \lambda)}{\partial a}.$$

According to this remarkable result, the total effect of a small change in a on the value of $v(a)$ is found by simply differentiating the Lagrangian $\mathcal{L}(x(a), a, \lambda)$ *partially* wrt a and thus treating x as well as λ as constants: the indirect effect of changing a on x and λ will be zero because we look only at optimal values of $f(\cdot)$. This makes life a lot easier. And it is extremely easy to see

² In the second term, we will be almost exclusively concerned with optimisation of value functions of the form $v(t) = f(x(t), t)$ where t is a time dependent parameter.

why this is true: the chain rule gives

$$\frac{dv(a)}{da} = \frac{\partial f(x(a), a)}{\partial a} + \frac{\partial f(x(a), a)}{\partial x} \frac{dx(a)}{da}. \quad (3.3.4)$$

We immediately see that the second half of the *rhs* in the above equation is zero because it contains the *foc* of problem (3.3.1) which is³

$$\frac{\partial f(x(a), a)}{\partial x} = 0. \quad (3.3.5)$$

This *must* hold since $v(a) = f^*(a)$ represents the *optimal* values of $f(\cdot)$. The intuition of this is equally easy to grasp: In figure 3.7 we represent $f(x, \cdot)$ for different values of $x = x_1, x_2, x_3$. Since for all a we know that $v(a) = \max_x f(x, a)$, the value function must give the optimal value of $f(x, a)$; it must form the upper envelope. Let us fix a at \hat{a} . Evidently we have $v(\hat{a}) = f(x(\hat{a}), \hat{a})$ and $v(\hat{a}) \geq f(x(\hat{a}), \hat{a})$, that is, $f(x(\hat{a}), \hat{a})$ lies weakly below $v(\hat{a})$ and touches the envelope only at $\hat{a} = a$. This intuition is confirmed by (3.3.4).




Figure 3.7: An upper envelope of a family of functions.

3.3.2 A general statement

Let us reconsider the maximisation problem we met in our first discussion of the Envelope Theorem (3.3.1) but now we include n variables, k parameters, and m constraints

$$\begin{aligned} & \max_{x_1, \dots, x_n} f(x_1, \dots, x_n, a_1, \dots, a_k) \\ & \text{s.t.} \quad \begin{aligned} & g_1(x_1, \dots, x_n, a_1, \dots, a_k) = c_1 \\ & \vdots \\ & g_m(x_1, \dots, x_n, a_1, \dots, a_k) = c_m. \end{aligned} \end{aligned} \quad (3.3.6)$$

³ For completeness the additional *foc* for the constrained case is $\frac{\partial \mathcal{L}}{\partial \lambda}(x_1^*, \dots, x_n^*, \lambda^*) = 0$.

Again we construct the Lagrangian \mathcal{L} as:

$$\begin{aligned} & \mathcal{L}(x_1, \dots, x_n, \lambda_1, \dots, \lambda_m; a_1, \dots, a_k) \\ &= f(x_1, \dots, x_n, a_1, \dots, a_k) + \sum_{j=1}^m \lambda_j (c_j - g_j(x_1, \dots, x_n, a_1, \dots, a_k)). \end{aligned}$$

Let $(x_1^*(a_1, \dots, a_k), \dots, x_n^*(a_1, \dots, a_k))$ and $(\lambda_1(a_1, \dots, a_k), \dots, \lambda_m(a_1, \dots, a_k))$ be the values of \mathbf{x} and $\boldsymbol{\lambda}$ that solve this problem. Now let

$$v(a_1, \dots, a_k) = f(x_1^*(a_1, \dots, a_k), \dots, x_n^*(a_1, \dots, a_k), a_1, \dots, a_k) \quad (3.3.7)$$

That is, $v(a_1, \dots, a_k)$ is the maximised value of the function f when the parameters are (a_1, \dots, a_k) . The Envelope Theorem says that the derivative of $v(\cdot)$ when varying parameter a is equal to the derivative of $\mathcal{L}(\cdot)$ at the maximising values of \mathbf{x} and $\boldsymbol{\lambda}$. Let us state this more precisely:⁴

Theorem 22 (Envelope Theorem). *Consider the value function $v(\mathbf{a}) = f(\mathbf{x}^*(\mathbf{a}), \mathbf{a})$ for the constrained, parameterised problem (3.3.6). Assume $v(\cdot)$ is differentiable at \mathbf{a} and let $\boldsymbol{\lambda}$ be the vector of Lagrange multipliers associated with the solution $\mathbf{x}^*(\mathbf{a})$ at \mathbf{a} . Then for all $h = 1, \dots, k$*

$$\begin{aligned} \frac{\partial v}{\partial a_h}(\mathbf{a}) &= \frac{\partial \mathcal{L}}{\partial a_h}(x_1^*(\mathbf{a}), \dots, x_n^*(\mathbf{a}), \lambda_1(\mathbf{a}), \dots, \lambda_m(\mathbf{a}), \mathbf{a}) \\ &= \frac{\partial f}{\partial a_h}(x_1^*(\mathbf{a}), \dots, x_n^*(\mathbf{a}), \mathbf{a}) - \sum_{j=1}^m \lambda_j(\mathbf{a}) \frac{\partial g_j}{\partial a_h}(x_1^*(\mathbf{a}), \dots, x_n^*(\mathbf{a}), \mathbf{a}). \end{aligned} \quad (3.3.8)$$

Proof. We use the function $v(\cdot)$ of maximised values of $f(\cdot)$ as defined in (3.3.7). Therefore

$$\frac{\partial v}{\partial a_h}(\mathbf{a}) = \frac{\partial f}{\partial a_h}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \frac{\partial x_i^*}{\partial a_h}(\mathbf{a}). \quad (3.3.9)$$

From the *foc* (3.3.5) we know that

$$\frac{\partial f}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) - \sum_{j=1}^m \lambda_j(\mathbf{a}) \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = 0$$

which we rewrite to

$$\frac{\partial f}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = \sum_{j=1}^m \lambda_j(\mathbf{a}) \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}). \quad (3.3.10)$$

Since the optimising arguments $\mathbf{x}^*(\mathbf{a})$ satisfy the constraints $g(\cdot)$, we also have for each j

$$g_j(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \equiv c_j.$$

⁴ Unfortunately, the lines are not wide enough to expand all the vectors of variables and parameters. But by now you surely got the idea—we shall proceed using vector notation $\mathbf{a} = a_1, \dots, a_k$.

Since the above is an identity, we may differentiate both sides wrt a_h

$$\sum_{i=1}^n \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \frac{\partial x_i^*}{\partial a_h}(\mathbf{a}) + \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = 0$$

or

$$\sum_{i=1}^n \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \frac{\partial x_i^*}{\partial a_h}(\mathbf{a}) = -\frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}). \quad (3.3.11)$$

Substituting (3.3.10) into (3.3.9) gives

$$\frac{\partial v}{\partial a_h}(\mathbf{a}) = \frac{\partial f}{\partial a_h}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) + \sum_{i=1}^n \left[\sum_{j=1}^m \lambda_j(\mathbf{a}) \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \right] \frac{\partial x_i^*}{\partial a_h}(\mathbf{a})$$

which, by changing the order of summation, turns to

$$\frac{\partial v}{\partial a_h}(\mathbf{a}) = \frac{\partial f}{\partial a_h}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) + \sum_{j=1}^m \lambda_j(\mathbf{a}) \left[\sum_{i=1}^n \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \frac{\partial x_i^*}{\partial a_h}(\mathbf{a}) \right]. \quad (3.3.12)$$

Now substituting (3.3.11) into (3.3.12) gives

$$\frac{\partial v}{\partial a_h}(\mathbf{a}) = \frac{\partial f}{\partial a_h}(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) - \sum_{j=1}^m \lambda_j(\mathbf{a}) \frac{\partial g_j}{\partial a_h}(\mathbf{x}^*(\mathbf{a}), \mathbf{a})$$

which is the desired result.⁵

□

3.4 Applications of the Envelope Theorem

To motivate our discussion of the Envelope Theorem we will now give a whole series of examples.

3.4.1 Cost functions

Our first particular case is the examination of the relation between short and long run average cost curves. You may remember from micro that in general we assume that the average cost of producing some good is a function of the amount of the good to be produced. The short run average cost function is defined to be the function which for any quantity, Q , gives the average cost of producing that quantity, taking as given the scale of operation, i.e., the size and number of plants and other fixed capital which we assume cannot be changed in the short run (whatever that is). On the other hand, the long run average cost function gives, as a function of Q , the average cost of producing Q units of the good, with the scale of operation selected to be the optimal scale for that level of production.

That is, if we let the scale of operation be measured by a single variable k , say, and we let the

⁵More elegant proofs can be found in (Sydsæter and Hammond 1995, 678ff) or (Mas-Colell, Whinston, and Green 1995, 966). Notice that the proof also applies in the presence of inequality constraints because the set of binding constraints should be unaltered in the neighbourhood of the optimal a .

short run average cost of producing Q units when the scale is k be given by $SRAC(Q, k)$ and the long run average cost of producing Q units by $LRAC(Q)$ then we have:

$$LRAC(Q) = \min_k SRAC(Q, k)$$

Let us denote, for a given value Q , the optimal level of k by $k(Q)$. That is, $k(Q)$ is the value of k that minimises the right hand side of the above equation.

Graphically, for any fixed level of k the short run average cost function can be represented by a curve (normally assumed to be U-shaped) drawn in two dimensions with quantity on the horizontal axis and cost on the vertical axis. Now think about drawing one short run average cost curve for each of the (infinite) possible values of k . One way of thinking about the long run average cost curve is as the ‘bottom’ or envelope of these short run average cost curves. Suppose that we consider a point on this long run or envelope curve. What can be said about the slope of the long run average cost curve at this point. A little thought should convince you that it should be the same as the slope of the short run curve through the same point. (If it were not then that short run curve would come below the long run curve, a contradiction.) That is,

$$\frac{dLRAC(Q)}{dQ} = \frac{\partial SRAC(Q, k)}{\partial Q}$$




Figure 3.8: An application of the Envelope Theorem to long- & short-term cost functions.

3.4.2 Parameterised maximisation problems

Consider the problem

$$\begin{aligned} & \max_x f(x) \\ \text{s.t.: } & g(x, a) = 0. \end{aligned}$$

We set up the Lagrangian as

$$\mathcal{L} = f(x) - \lambda g(x, a) \tag{3.4.1}$$

and obtain the foc as

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= f_x - \lambda g_x(x, a) = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= g(x, a) = 0. \end{aligned} \tag{3.4.2}$$

Solving these gives a set of solution candidate $x^*(a)$ and a $\lambda^*(a)$ that satisfy the *foc*. We can rewrite (3.4.1) as

$$\mathcal{L}(a) = f(x^*(a)) - \lambda^*(a)g(x^*(a), a)$$

and find the maximum of $\mathcal{L}(a)$ by taking total derivatives

$$\frac{d\mathcal{L}(a)}{da} = [f_x(x^*(a)) - \lambda^*(a)g_x(x^*(a), a)] \frac{dx^*(a)}{da} - g(x^*(a), a) \frac{d\lambda^*(a)}{da} - \lambda^*(a)g_a(x, a).$$

Since we know from (3.4.2) that $f_x(x^*(a)) - \lambda^*(a)g_x(x^*(a), a) = 0$ and $g_x(x, a) = 0$ we conclude that

$$\frac{d\mathcal{L}(a)}{da} = -\lambda^*(a)g_a(x, a).$$

Only the direct effect remains.

We now apply this general notation to the **utility maximisation** problem:

$$\begin{aligned} & \max_{x_1, x_2} u(x_1, x_2) \\ \text{s.t.: } & p_1x_1 + p_2x_2 - y = 0. \end{aligned}$$

Let us call the maximised utility at given prices p and income y the *indirect* utility function and denote it by $v(p_1, p_2, y)$. Now let us look at the effect of a change in y with p_1, p_2 held constant. By the Envelope Theorem

$$\frac{\partial v}{\partial y} = \frac{\partial}{\partial y} \{u(x_1, x_2) + \lambda(y - p_1x_1 + p_2x_2)\} = 0 + \lambda 1 = \lambda.$$

This is the reason why λ is sometimes referred to as a *shadow price*: Because it gives precisely the marginal utility of income or, in other words, the increase in utility one obtains by raising ones income. As we see above, the effect of the changes in income on the optimal goods bundle x^* are ignored. In effect the vector x^* is treated as constant.

3.4.3 Expenditure minimisation and Shephard's Lemma

Let us consider the dual problem to the utility maximisation problem we looked at above. In this case the consumer minimises expenditures px subject to attaining a given level of utility \bar{u}

$$\begin{aligned} \min_{x_1, \dots, x_n} \sum_{j=1}^n p_j x_j \\ \text{s.t.: } u(x_1, \dots, x_n) - \bar{u} = 0. \end{aligned} \quad (3.4.3)$$

We set up the Lagrangian in the obvious way as

$$\mathcal{L}(x, \lambda; p) = \sum_{j=1}^n p_j x_j + \lambda (\bar{u} - u(x_1, \dots, x_n))$$

and denote the *minimised* value of the expenditure function by $e(p_1, \dots, p_n, \bar{u})$. The Envelope Theorem states that

$$\frac{\partial e(p, \bar{u})}{\partial p_j} = \frac{\partial \mathcal{L}(x^*, \lambda; p)}{\partial p_j}$$

and hence we obtain at the point x^* which solves the minimisation problem

$$\frac{\partial e(p, \bar{u})}{\partial p_j} = \frac{\partial}{\partial p_j} \left\{ \sum_{j=1}^n p_j x_j^* + \lambda (\bar{u} - u(x_1^*, \dots, x_n^*)) \right\} = x_j^* - \lambda = x_j^* \equiv h_j(p, \bar{u}).$$

This is called the *Hicksian* (or compensated) demand for good j . We denote it by $h_j(p_1, \dots, p_n, \bar{u})$ to indicate that it is a function of prices and utility. This differs from the *Marshallian* (or uncompensated) demand $x_j(p_1, \dots, p_n, \bar{y})$ for good j which is a function of prices and income. The above result is known as *Shephard's Lemma*. Again you see that the optimal consumption vector x^* is treated as constant—substitution effects are ignored. That this is only an approximate result can be seen from figure 3.9.




Figure 3.9: The approximation involved in Shephard's Lemma: Area ABC is ignored although it is part of the compensation required to keep the consumer equally well off after a change in prices δ .

3.4.4 The Hicks-Slutsky equation

It can be shown (by the application of the Duality Theorem (26)) that the compensated demand at utility \bar{u} , i.e., $h_i(p_1, \dots, p_n, \bar{u})$ is identically equal to the uncompensated demand at income $e(p_1, \dots, p_n, \bar{u})$, i.e., $x_i(p_1, \dots, p_n, e(p_1, \dots, p_n, \bar{u}))$. Thus totally differentiating this identity

$$x_i(p_1, \dots, p_n, e(p_1, \dots, p_n, \bar{u})) \equiv h_i(p_1, \dots, p_n, \bar{u})$$

with respect to p_k we obtain

$$\frac{\partial x_i}{\partial p_k} + \frac{\partial x_i}{\partial y} \frac{\partial e}{\partial p_k} = \frac{\partial h_i}{\partial p_k}$$

which by Shephard's Lemma gives

$$\frac{\partial x_i}{\partial p_k} + \frac{\partial x_i}{\partial y} h_k = \frac{\partial h_i}{\partial p_k}.$$

So for all $i, k = 1, \dots, n$

$$\frac{\partial x_i}{\partial p_k} = \frac{\partial h_i}{\partial p_k} - \frac{\partial x_i}{\partial y} h_k \quad (3.4.4)$$

which is known as the *Slutsky equation*.

3.4.5 The Indirect utility function and Roy's Identity

Again let $v(p_1, \dots, p_n, y)$ be the indirect utility function from the utility maximisation example, that is, the maximised value of the direct utility function $u(x_1, \dots, x_n)$. Then by the Envelope Theorem

$$\frac{\partial v}{\partial p_j} = \frac{\partial u}{\partial p_j} - \lambda x_j(p_1, \dots, p_n, y) = -\lambda x_j(p_1, \dots, p_n, y)$$

since $\frac{\partial u}{\partial p_j} = 0$. Now, since we have already shown in the above utility maximisation example that $\lambda = \frac{\partial v}{\partial y}$ we have

$$x_j(p_1, \dots, p_n, y) = -\frac{\partial v / \partial p_j}{\partial v / \partial y}.$$

This is known as *Roy's Identity*.

3.4.6 Profit functions and Hotelling's Lemma

Consider the problem of a firm that maximises profits subject to technology constraints. Let $\mathbf{x} = (x_1, \dots, x_n)$ be a vector of net outputs ('netputs'), i.e., x_i is positive if the firm is a net supplier of good i , negative if the firm is a net user of that good. Let us assume that we can write the technology constraints as $F(\mathbf{x}) = 0$. Thus the firm's problem is

$$\begin{aligned} & \max_{x_1, \dots, x_n} \sum_{i=1}^n p_i x_i \\ & \text{s.t.: } F(x_1, \dots, x_n) = 0. \end{aligned}$$

Let $\varphi_i(\mathbf{p})$ be the value of x_i that solves this problem, i.e. the net supply of commodity i at price vector \mathbf{p} . We call this maximised value the profit function which is given by

$$\Pi(\mathbf{p}) = \sum_{i=1}^n p_i \varphi_i(\mathbf{p}).$$

Hence by the Envelope Theorem

$$\frac{\partial \Pi}{\partial p_i}(\mathbf{p}) = \varphi_i(\mathbf{p})$$

which is known as *Hotelling's Lemma*.

Let us now pull together the above results and draw a chart of how these results interact in the context of consumer theory. The (primary) utility maximisation problem is shown left and the (dual) expenditure minimisation problem right in figure 3.10.

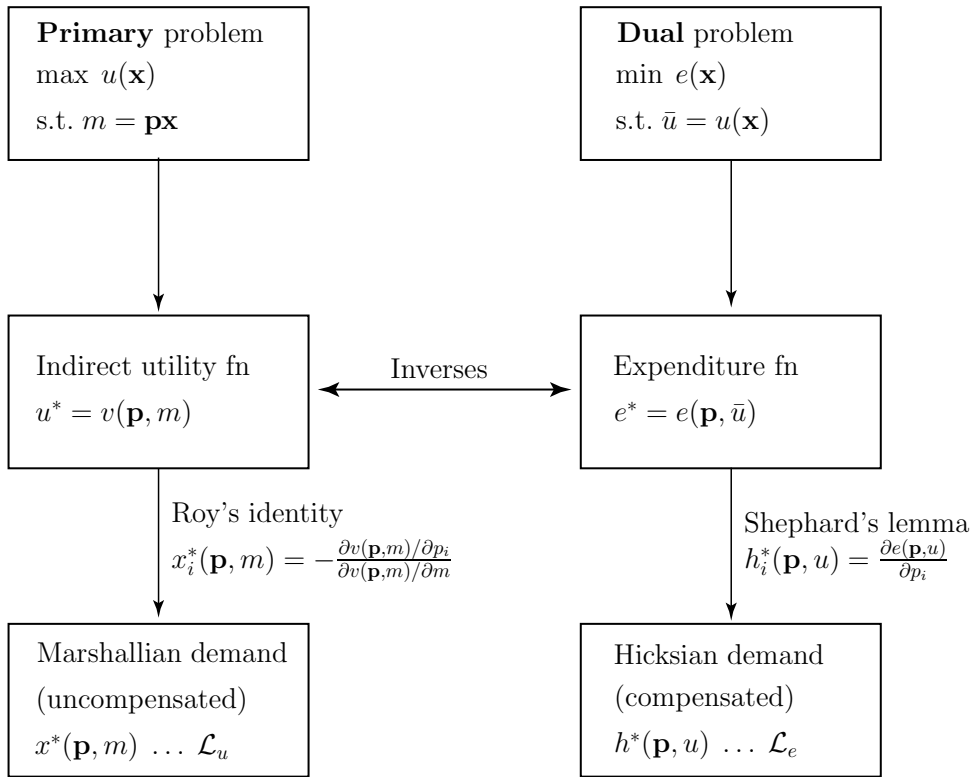


Figure 3.10: Demand concepts in consumer theory.

Exercises

Exc 3.1: What changes in our proof of theorem 21 if we (i) do not allow for mixed strategies?, or if we (ii) do allow for correlated beliefs, i.e. beliefs that are not independent (product) beliefs? Draw a picture in the spirit of figure 3.4 and concentrate on the set $X = \{(u_i(\sigma_i, s_{-i})_{s_{-i} \in S_{-i}} \mid \sigma_i \in \Delta(S_i))\}$.

Exc 3.2: This exercise is an attempt to convince you that the envelope theorem is useful in

real applications. You do not have to understand the jargon—just setting up the maximisation problem and differentiate away will lead you to victory.

Consider a direct regulation mechanism based on transfers $T(P(\theta))$ and price regulation as proposed by (Baron and Myerson 1982).⁶ A monopolist announces a price $P(\theta)$ based on its cost type $\theta \in [\theta_L, \theta_H]$ and the regulator responds by setting the optimal transfer $T(P(\theta))$ (positive or negative) in order for the monopolist to produce a socially useful quantity $Q(P(\theta))$. (We are not concerned with the regulators' response here.)

1. Set up the profit maximisation problem of the firm.
2. Derive the firms' optimal price announcement.
3. Now derive the firm's optimal type announcement.
4. Are profits increasing or decreasing in the quantity produced?
5. What is the relation to the Envelope Theorem?

Exc 3.3: Show that the utility function

$$u(x_1, x_2) = k(x_1 - a)^\alpha(x_2 - b)^\beta$$

where $k, a, b, \alpha, \beta > 0$, represents the same preferences as the utility function

$$u(x_1, x_2) = \delta \ln(x_1 - a) + (1 - \delta) \ln(x_2 - b) \quad (3.4.5)$$

where

$$\delta = \frac{\alpha}{\alpha + \beta}.$$

Let prices be (p_1, p_2) and let income be m . Assume $m > p_1a + p_2b$.

Exc 3.4: (Optional) Assume that a consumer's preferences over (x_1, x_2) are represented by (3.4.5). Let prices be (p_1, p_2) and income be m . Assume $m > p_1a + p_2b$.⁷

1. find the direct demand functions $x_i(\mathbf{p}, m)$ by maximising utility s.t. $m \geq p_1x_1 + p_2x_2$;
2. show that the indirect utility function is given by

$$v(p, m) = \ln(m - ap_1 - bp_2) + \delta \ln \delta + (1 - \delta) \ln(1 - \delta) - \delta \ln p_1 - (1 - \delta) \ln p_2;$$

3. find the Hicksian demands $h_i(\mathbf{p}, \bar{u})$ by minimising expenditure s.t. $u(\mathbf{x}) \geq \bar{u}$;
4. show that the expenditure function is given by

$$e(\mathbf{p}, \bar{u}) = ap_1 + bp_2 + e^{\bar{u}} \delta^{-\delta} (1 - \delta)^{-(1-\delta)} p_1^\delta p_2^{(1-\delta)}.$$

5. Interpret a and b . Why do we need to assume $m > p_1a + p_2b$?

⁶ The reference is given here just to satisfy your curiosity—it is not necessary to read the article in order to solve the exercise.

⁷ This problem is not harder than the rest but it is definitely longer (apart from that it really belongs into a micro course). Only do it if you like this kind of thing or need the exercise.

Chapter 4

Kuhn-Tucker theory

Reading: (Sundaram 1996).

In section 1.1 we discussed optimisation problems in the presence of *equality* constraints in a somewhat informal setting. By now we have developed the tools to make this discussion more precise. We will then look at an extension to the Lagrange technique called the Kuhn-Tucker Theorem, which gives (first-order) necessary conditions for finding local maxima for optimisation problems of $f(\cdot)$ with *inequality* constraints $h_i : \mathbb{R}^n \mapsto \mathbb{R}$, $i = 1, \dots, l$ forming some constraint set \mathcal{D} . As we will see, the combination of the Lagrange Theorem with the Kuhn-Tucker Theorem to solve mixed optimisation problems (in the presence of both equality and inequality constraints) is straightforward.

4.1 Theory

4.1.1 The Lagrangian

We will first recapitulate what we know about the Lagrangian and develop our intuition for the equality constrained case. Remember that our problem was to maximise a general function $f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^n$ subject to some (non-linear) constraint $g(\mathbf{x}) = c$. In order to allow us to draw a picture, let $\mathbf{x} \in \mathbb{R}^2 = (x, y)$. Our problem is therefore to

$$\begin{aligned} & \max_{x,y} f(x, y) \\ & \text{s.t.} : g(x, y) = c. \end{aligned} \tag{4.1.1}$$

As shown in figure 4.1, the graph of $f(\cdot)$ can be (locally) approximated by a quadratic form. We recall from section 2.6.4 that in the case depicted above, we look at a negative semi-definite quadratic form that identifies the underlying objective as (locally) quasiconcave. The ‘normal’ constraint $g(\cdot)$ is the projection of the constraint curve G into the Cartesian plane. The unconstrained maximum of $f(\cdot)$ is at the point A but since we look at a constrained problem, the best we can do is B —the highest point on G (this is illustrated in the right panel of figure 4.1). Understanding why this procedure works is easy. We denote the Lagrangian of the above problem

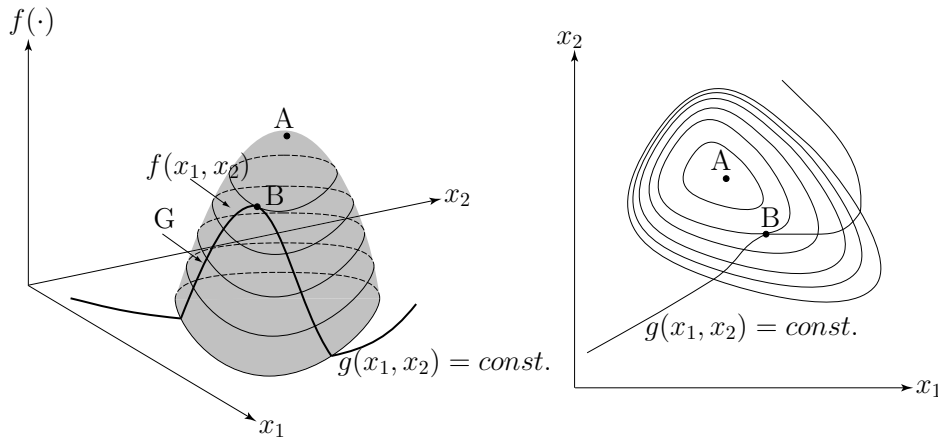


Figure 4.1: Constrained optimisation—the Lagrangian method.

by

$$\mathcal{L}(\mathbf{x}, \lambda) = f(\mathbf{x}) - \lambda(g(\mathbf{x}) - c)$$

which leads to the *foc*

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x_i}(x_1^*, \dots, x_n^*, \lambda^*) = 0 &\implies \nabla f(\mathbf{x}^*) = \lambda \nabla g(\mathbf{x}^*) \\ \frac{\partial \mathcal{L}}{\partial \lambda}(x_1^*, \dots, x_n^*, \lambda^*) = 0 &\implies g(\mathbf{x}^*) = c. \end{aligned}$$

Exl: Consider

$$\begin{aligned} \max_{x,y} f(x,y) &= x^2 + y^2 \\ \text{s.t.} \quad g(x,y) &= x + y = 2. \end{aligned}$$

Then we obtain the *foc* as

$$\begin{aligned} \nabla f(\mathbf{x}^*) = \lambda &\implies (2x, 2y) = \lambda(1, 1) \\ g(\mathbf{x}^*) = c &\implies x + y = 2. \end{aligned}$$


Which is at $\mathbf{x}^* = (1, 1)$

$$\begin{aligned} (2, 2) &= 2(1, 1) \\ 1 + 1 &= 2. \end{aligned}$$

By the way: In this case, the Lagrangian does *not* give the right answer. \triangleleft

The first condition means that the gradient vectors of $f(\cdot)$ and $g(\cdot)$ are proportional at \mathbf{x}^* . In the case of a single linear constraint $g(x, y)$ in the Cartesian plane, this gives rise to the interpretation in figure 4.2. You should keep in mind, however, that the price vector is only orthogonal to the budget hyperplane everywhere in the case of a linear constraint.

Theorem 23. (Lagrange) Suppose $f(x, y)$ and $g(x, y)$ have continuous partial derivatives in a domain A of the Cartesian plane, and that (x^*, y^*) is both an interior point of A and a local extreme point for $f(x, y)$ s.t.: $g(x, y) = c$. Suppose further that $g'_1(x^*, y^*)$ and $f'_1(x^*, y^*)$ are not both 0. Then there exists a unique number λ such that the Lagrangian function $\mathcal{L}(x, y, \lambda) =$

Figure 4.2: The interpretation of the Lagrangian *foc*.

$f(x, y) - \lambda(g(x, y) - c)$ has a stationary point at (x^*, y^*) .

Proof. See (Sundaram 1996, 135). □

Global sufficiency for the existence of the solution (x^*, y^*) to a maximisation (minimisation) problem is given by concavity (convexity) of $\mathcal{L}(x, y, \lambda)$. We can check this by finding the bordered Hessian $\overline{H}(x^*, y^*) > 0$ in the maximisation case or $\overline{H}(x^*, y^*) < 0$ in the minimisation case. These conditions are precisely the same in the Lagrange and Kuhn-Tucker cases. In order to formulate them, we need some additional notation. Remember that we look at a constrained optimisation problem of $f(\mathbf{x})$ with n variables x_i and m constraints g_j . We first define

$$\mathbf{y} \equiv \{y_1, \dots, y_m, y_{m+1}, \dots, y_{m+n}\} = \{\lambda_1, \dots, \lambda_m, x_1, \dots, x_n\}$$

in order to have a single (vector) variable in our function of critical values $\mathcal{L}(\cdot)$.

Def: The determinant of the matrix of second derivatives (the so called Hessian) of the Lagrange function $\mathcal{L}(\mathbf{y})$ is called the **bordered Hessian** and denoted $\overline{H}(\mathbf{y}^*)$. More precisely, for the C^2 function $\mathcal{L}(\mathbf{y})$, $\mathbf{y} \in \mathbb{R}^{n+m}$, $r = 1, \dots, n$, the bordered Hessian is defined as¹

$$\overline{H}_r(\mathbf{y}^*) = \begin{vmatrix} 0 & \cdots & 0 & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_1 \partial y_{m+1}} & \cdots & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_1 \partial y_{m+n}} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_m \partial y_{m+1}} & \cdots & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_m \partial y_{m+n}} \\ \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+1} \partial y_1} & \cdots & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+1} \partial y_m} & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+1} \partial y_{m+1}} & \cdots & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+1} \partial y_{m+n}} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+n} \partial y_1} & \cdots & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+n} \partial y_m} & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+n} \partial y_{m+1}} & \cdots & \frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_{m+n} \partial y_{m+n}} \end{vmatrix}. \quad (4.1.2)$$

¹ We will show how to construct the bordered Hessian for constrained problems in the example given below in (4.2.6). For a more detailed definition (involving more notation) cf. (Sundaram 1996, 120).

Rem: We know from section 2.6.4 that $\frac{\partial^2 \mathcal{L}(\mathbf{y})}{\partial y_i^2}$ is a quadratic form.

Proposition 24. *From what we know about quadratic forms (section 2.6.4) we summarise that*

1. *a necessary condition for $f(\mathbf{x})$ to be quasiconcave is that $(-1)^r \overline{H}_r(\mathbf{x}^*) \geq 0$ for $r = 1, \dots, n$, and all $\mathbf{x} \in \mathcal{D}$;*
2. *a sufficient condition for $f(\mathbf{x})$ to be quasiconcave is that $(-1)^r \overline{H}_r(\mathbf{x}^*) > 0$ for $r = 1, \dots, n$, and all $\mathbf{x} \in \mathcal{D}$.*

Rem: Therefore we conclude that a sufficient condition for \mathbf{x}^* to be a local extremum of $f(\mathbf{x})$ on \mathcal{D} is that

1. \mathbf{x}^* satisfies the first order conditions, and
2. the bordered Hessian $\overline{H}(\mathbf{x}^*) > 0$ in the maximisation case and $\overline{H}(\mathbf{x}^*) < 0$ in the minimisation case.

4.1.2 The extension proposed by Kuhn and Tucker

We return to our problem (4.1.1) but this time, we use inequality constraints

$$\begin{aligned} \max_{x,y} f(x,y) \\ \text{s.t.: } g(x,y) \leq c. \end{aligned} \tag{4.1.3}$$

This gives rise to the same Lagrangian function, but we are going to interpret it differently. Our problem (4.1.3) can be geometrically interpreted as in the following figure (which shows level sets of the function represented to the left of figure 4.1).




Figure 4.3: Constrained optimisation—the Kuhn-Tucker method.

In the right panel of figure 4.3, the constraint $g(\cdot)$ is binding at the locally optimal solution C . In the left panel, however, the constraint is not binding at the global maximum A which is an interior point of the feasible set \mathcal{D} . Therefore, our true optimum is not achieved at the frontier of the budget set but where our constraint is slack. Therefore the multiplier λ associated with the problem (4.1.3) is only non-zero in the right panel. Shadow prices are zero on the left.

The situation described in the left panel cannot be represented as the solution to a standard-Lagrangian problem because there we are forced to consider only solution candidates *on* the constraint.

This idea generalises naturally to the case of multiple constraint functions $h_i(\cdot)$. Recall from section 2 that $\text{rank}(X)$ denotes the rank of a matrix X and the gradient $\nabla f(\mathbf{x})$ of a function $f(\mathbf{x}) : \mathbb{R}^n \mapsto \mathbb{R}$ is the (column) vector of dimension n whose i^{th} entry is $\frac{\partial f}{\partial x_i}(\mathbf{x})$, $i = 1, \dots, n$.

The Kuhn-Tucker extension of the Lagrangian theory allows precisely for situations like that depicted at the left of figure 4.3. Loosely speaking it forces the multiplier to be zero if the constraint does not bind, i.e. $g(\cdot) > 0$. In any other case (and in good Lagrangian fashion), the two gradients of the objective and the constraint are forced into line.

Theorem 25. (Kuhn-Tucker) *Let $f : \mathbb{R}^n \mapsto \mathbb{R}$ and $g_i : \mathbb{R}^n \mapsto \mathbb{R}$, $i = 1, \dots, l$ be C^1 . Suppose \mathbf{x}^* is a local maximum of $f(\cdot)$ on $\mathcal{D} = U \cap \{\mathbf{x} \in \mathbb{R}^n \mid g_i(\mathbf{x}) \geq 0, i = 1, \dots, l\}$ with $U \subset \mathbb{R}^n$ open. Let E be the set of effective constraints at the candidate solution under consideration and let $g_E = (g_i)_{i \in E}$ and suppose*

$$\text{rank}(\nabla g_E(\mathbf{x}^*)) = |E|. \quad (4.1.4)$$

Then, there exists a vector $\boldsymbol{\lambda}^ = (\lambda_1^*, \dots, \lambda_l^*) \in \mathbb{R}^l$ such that the following conditions are met²*

$$\begin{aligned} (KT-1): \lambda_i^* \geq 0 \text{ and } \lambda_i^* g_i(\mathbf{x}^*) = 0 \text{ for } i = 1, \dots, l, \text{ and} \\ (KT-2): \nabla f(\mathbf{x}^*) + \sum_{i=1}^l \lambda_i^* \nabla g_i(\mathbf{x}^*) = 0. \end{aligned} \quad (4.1.5)$$

Proof. See (Sundaram 1996, 165). □

Rem: The terminology for (KT-1) as ‘complementary slackness’ condition arises from the observation that by the feasibility of \mathbf{x}^* , we must have $g_i(\mathbf{x}^*) \leq 0$ for each i ; therefore, for $\lambda_i^* g_i(\mathbf{x}^*) = 0$ to hold alongside with $\lambda_i^* \geq 0$, we must have $\lambda_i^* = 0$ if $g_i(\mathbf{x}^*) < 0$, and $g_i(\mathbf{x}^*) = 0$ if $\lambda_i^* > 0$. That is, if the inequality for $g_i(\cdot)$ is slack (i.e. not binding), then the other one for λ_i^* is binding (and vice versa).

The complementary slackness condition is the major addition introduced by Kuhn-Tucker that enables the Lagrange approach to cope with inequality constrained optimisation problems.

Condition: The awkward business of defining $f(\cdot)$ only locally on the intersection of \mathcal{D} with an open set U guarantees the existence of a local solution. In general we will not be able to formulate our problems this conveniently. Therefore we will have to check the existence of a solution in each case separately by checking the conditions on the Weierstrass Theorem.

Def: The above full rank condition (4.1.4) is called the **constraint qualification**. It says that the number of effective constraints $|E|$ at the candidate solution \mathbf{x}^* must equal the minimum dimension of the Jacobian of the effective constraints at \mathbf{x}^* .

² While (KT-1) is usually called the ‘complementary slackness’ condition(s), (KT-2) is referred to as the ‘Kuhn-Tucker condition(s)’.

Rem: The constraint qualification ensures that the constraints are independent at the local maximum \mathbf{x}^* . More specifically, it requires the rank of the gradient matrix of the constraint set at the solution candidate to be equal to the number of constraints. This is to ensure that every direction of displacement $z \in \mathbb{R}^n$ having no first-order effect on the constraints (i.e. satisfying $\nabla g_i(\mathbf{x}^*) \cdot z = 0$ for every i) must also have no first-order effect on the objective function (i.e. satisfying $\nabla f(\mathbf{x}^*) \cdot z = 0$).




Figure 4.4: Indispensability of the constraint qualification.

As shown in figure 4.4, at a local maximiser $\bar{x} = \mathbf{x}^*$, the gradient of the objective is a linear combination of the gradients of the constraint functions. The figure shows a situation where we maximise $f(x_1, x_2)$ on $D = \{(x_1, x_2) \in \mathbb{R}^2 : g_i(x_1, x_2) = \bar{b}_i\}$ for $i = 1, 2$. While the point \bar{x} is a global constrained maximum (it is the only vector in the constrained set!), we see that $\nabla f(\bar{x})$ is *not* spanned by the vectors $\nabla g_1(\bar{x})$ and $\nabla g_2(\bar{x})$. Since $\nabla g_1(\bar{x}) = -\nabla g_2(\bar{x})$, the constraint qualification is violated.

Condition: The constraint qualification is the second condition (in addition to existence) that we need to satisfy on top of the above two (necessary) conditions (KT-1) and (KT-2).

Notation: From now on, we will say that a pair $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ meets the first-order necessary conditions (or Kuhn-Tucker conditions) for a maximum, if it satisfies the above two conditions (KT-1) and (KT-2) in (4.1.5).

Rem: Because we have now constraints of the form $g_i(\cdot) \geq c$, we cannot use the substitution method any more.

4.2 A cookbook approach

We will now develop a cookbook approach of how to apply the Kuhn-Tucker Theorem to inequality constrained optimisation problems. Our approach closely resembles the methodology that we developed in section 1.1 for the solution of equality constrained problems.³

³ As in the case of the Lagrangian, we will apply our approach to both linear and non-linear optimisation problems.

4.2.1 The cookbook version of the Lagrange method

1. Form the Lagrangian $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^l \mapsto \mathbb{R}$ for the n variables x_j and l constraints $g_i(\mathbf{x}) = c_i$ associated with the multipliers λ_i

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) + \sum_{i=1}^l \lambda_i (g_i(\mathbf{x}) - c_i).$$

2. Compute the critical points (*foc*)

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x_j}(\mathbf{x}, \boldsymbol{\lambda}) &= 0, \quad j = 1, \dots, n. \\ \frac{\partial \mathcal{L}}{\partial \lambda_i}(\mathbf{x}, \boldsymbol{\lambda}) &= 0 \Leftrightarrow g_i(\mathbf{x}) = c_i, \quad i = 1, \dots, l. \end{aligned} \tag{4.2.1}$$

Only with equality constraints will setting the partial derivative $\frac{\partial \mathcal{L}}{\partial \lambda_i} = 0$ recover the constraint $g_i(\mathbf{x}) = c_i$ and therefore result in a valid procedure to derive the *foc*. If we have inequality constraints in place, we have to use the more involved conditions (4.1.5).

3. We solve the $n + l$ equations in (4.2.1) and, given a solution to our problem exists at all, we are done.

Condition: A global maximum \mathbf{x}^* exists to the given equality-constrained problem.

Condition: The constraint qualification is met at \mathbf{x}^* .

Exl: See (Sundaram 1996, 124) or the exercises from section 1.1. ◁

4.2.2 The cookbook version of the Kuhn-Tucker method

The generalisation to inequality constraints is just a little more difficult. The problem is now to

$$\begin{aligned} \max_{\mathbf{x}} f(\mathbf{x}) \\ \text{s.t.: } \mathbf{x} \in \mathcal{D} = \{g(\mathbf{x}) \geq 0\} \end{aligned}$$

1. Form the Lagrangian $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^l \mapsto \mathbb{R}$

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) + \sum_{i=1}^l \lambda_i g_i(\mathbf{x}).$$

2. Compute the critical points of \mathcal{L} by solving the following 2 systems of equations

$$\frac{\partial \mathcal{L}}{\partial x_j}(\mathbf{x}, \boldsymbol{\lambda}) = 0, \quad j = 1, \dots, n.$$

The derivation of the equations for the multipliers (the *complementary slackness* conditions) makes the key difference to the pure Lagrangian approach. As usual, we take derivatives wrt our multipliers $i = 1, \dots, l$ and obtain our original constraints

$$\frac{\partial \mathcal{L}}{\partial \lambda_i}(\mathbf{x}, \boldsymbol{\lambda}) \geq 0 \Leftrightarrow g_i(\mathbf{x}) \geq 0.$$

But in addition we impose the complementary slackness condition which means that simultaneously

$$\lambda_i \geq 0, \lambda_i \frac{\partial \mathcal{L}}{\partial \lambda_i}(\mathbf{x}, \boldsymbol{\lambda}) = 0.$$

This condition is rather tricky. It requires λ_i to be non-negative, and either $\lambda_i = 0$ if $g_i(\mathbf{x}) > 0$ or $\lambda_i > 0$, if $g_i(\mathbf{x}) = 0$. We can alternatively formulate this as

$$\lambda_i \geq 0, \lambda_i [g(\mathbf{x})] = 0. \quad (4.2.2)$$

It is useful to be systematic about this: In order to fulfil (4.2.2) we can have either

- (i) the constraint binding $g_i(\mathbf{x}) = 0$ and the multiplier slack $\lambda_i > 0$, or
- (ii) the constraint slack $g_i(\mathbf{x}) > 0$ and the multiplier binding $\lambda_i = 0$.

Just as before, the multipliers can be interpreted as ‘shadow prices’ associated with increasing the right-hand side of the resource constraint $g(\mathbf{x}) \geq 0$ by one unit. With this interpretation, shadow prices are non-negative, and if the resource constraint is not binding because $g(\mathbf{x}) > 0$ at the optimum, this means that the price associated with increasing $g(\mathbf{x})$ by one unit is zero.

Now we denote the set of critical points $\mathbf{x} \in U$ by M .

3. Evaluate the objective $f(\cdot)$ at each x in the set

$$\{\mathbf{x} \mid \text{there is a } \boldsymbol{\lambda} \text{ such that } (\mathbf{x}, \boldsymbol{\lambda}) \in M\}.$$

Usually, the value of x that maximises f over this set is also a solution to the original maximisation problem. It is *not* guaranteed, though, that the above procedure really picks up a maximum. There are two conditions attached for the existence of a $\boldsymbol{\lambda}^*$ such that $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ is a critical point of \mathcal{L} :

Condition: A global maximum \mathbf{x}^* exists to the given inequality-constrained problem.

Condition: The constraint qualification is met at \mathbf{x}^* .

4.2.3 A first cookbook example

Consider the problem

$$\begin{aligned} & \max_{x,y} x^2 + y^2 + y - 1 \\ \text{s.t.: } & (x, y) \in \mathcal{D} = \{(x, y) \in \mathbb{R}^2 \mid g(x, y) = x^2 + y^2 \leq 1\}. \end{aligned}$$

1. Step: Set up the Lagrangian as

$$\mathcal{L}(x, y, \lambda) = x^2 + y^2 + y - 1 - \lambda(x^2 + y^2 - 1).$$

2. Step: The critical points are the solutions (x, y, λ) to

$$\begin{aligned} 2x - 2\lambda x &= 0 \\ 2y + 1 - 2\lambda y &= 0 \\ \lambda &\geq 0 \quad (= 0 \text{ if } x^2 + y^2 < 1). \end{aligned}$$

The first condition tells us that either $\lambda = 1$ or $x = 0$. If $\lambda = 1$, then the second condition yields $1 = 0$, a contradiction. Thus $x = 0$ and $y = \pm 1$ because in the first case we take λ to be nonzero. We conclude that the candidates for optimality in this case are $f(0, 1)$ and $f(0, -1)$.

Now we check the case where $x = 0$ and $\lambda = 0$ and therefore $x^2 + y^2 < 1 \Leftrightarrow -1 < y < 1$. Then the third condition implies $\lambda = 0$ and therefore the second condition says $y = -\frac{1}{2}$. Therefore $f(0, -\frac{1}{2})$ is our third candidate.

3. Step: Evaluate the candidates

$$\begin{aligned} f(0, 1) &= 1 \\ f(0, -1) &= -1 \\ f(0, -\frac{1}{2}) &= -\frac{5}{4}. \end{aligned} \tag{4.2.3}$$

Let us now check whether the two conditions we attached to the application of the Kuhn-Tucker Theorem hold.

Condition: (Existence) Since we maximise a continuous function $f(\cdot)$ over a closed and bounded set \mathcal{D} , we know that there is a solution to the problem. We conclude that $f(0, 1) = 1$ actually solves the problem.

Condition: (Constraint qualification) The constraint gradient is given by $\nabla g(x, y) = (2x, 2y)$ which, for all the possible candidates in (4.2.3), always has rank 1. Therefore our constraint qualification holds. ◁

4.2.4 Another cookbook example

This example is taken from (Sundaram 1996, 155). Let $g(x, y) = 1 - x^2 - y^2$. Consider the problem of maximising the objective $f(x, y) = x^2 - y$

$$\begin{aligned} &\max_{x, y} x^2 - y \\ \text{s.t.: } &(x, y) \in \mathcal{D} = \{(x, y) \in \mathbb{R}^2 \mid g(x, y) = 1 - x^2 - y^2 \geq 0\}. \end{aligned}$$

1. Step: Set up the Lagrangian as

$$\mathcal{L}(x, y, \lambda) = x^2 - y + \lambda(1 - x^2 - y^2).$$

2. Step: The critical points are the solutions (x, y, λ) to

$$\begin{aligned} 2x - 2\lambda x &= 0 \\ -1 - 2\lambda y &= 0 \\ \lambda \geq 0, (1 - x^2 - y^2) &\geq 0, \lambda(1 - x^2 - y^2) = 0. \end{aligned}$$

For the first of these to hold we must have either $x = 0$ or $\lambda = 1$. If $\lambda = 1$, the second equation gives $y = -\frac{1}{2}$ and $x^2 + y^2 = 1$ follows from the third. Therefore we get a set of these two points (x, y, λ) and a third for the case of $x = 0$ and $\lambda > 0$

$$M = \left\{ \left(+\frac{\sqrt{3}}{2}, -\frac{1}{2}, 1 \right), \left(-\frac{\sqrt{3}}{2}, -\frac{1}{2}, 1 \right), \left(0, -1, \frac{1}{2} \right) \right\}$$

with $f(x, y)_{1,2} = \frac{3}{4} + \frac{1}{2} = \frac{5}{4}$ for the first two and $f(x, y)_3 = 1$ for the third.

3. Step: We know that a global maximum for f on \mathcal{D} must arise on one of the three points in M . Evidently the first two are equally good and better than the third. Hence we have two solutions to the given optimisation problem, namely the two points $(x, y) \in \left\{ \left(+\frac{\sqrt{3}}{2}, -\frac{1}{2} \right), \left(-\frac{\sqrt{3}}{2}, -\frac{1}{2} \right) \right\}$.

Again we check whether the two conditions we attached to the application of the KT Theorem hold.

Condition: (Existence) Our feasible set \mathcal{D} is the closed unit disk—this set is compact. Since the objective is continuous (all polynomials are continuous), we know by the Weierstrass Theorem that a maximum exists. Thus the first condition above is met.

Condition: (Constraint qualification) The constraint qualification is also met: At a point (x, y) where the constraint binds (i.e., $(x, y) \in \mathcal{D}$), we have $x^2 + y^2 = 1$ and we know that either $x \neq 0$ or $y \neq 0$. Since $\nabla g(x, y) = (-2x, -2y)$ at all $(x, y) \in \mathcal{D}$ it follows that when g is effective, we must have $\text{rank}(\nabla g(x, y)) = 1$. Thus the constraint qualification holds for if the optimum occurs at a (x, y) where $g(x, y) = 0$. If the optimum occurs at a (x, y) where $g(x, y) > 0$, no constraints are binding, the set \mathcal{D} is empty and the constraint qualification is met trivially.◁

4.2.5 A last cookbook example

This example is taken from (Sundaram 1996, 158). Consider

$$\begin{aligned} &\max_{x_1, x_2} x_1 + x_2 \\ \text{s.t.: } &(x_1, x_2) \in \mathcal{D} = \{(x_1, x_2) \in \mathbb{R}^2 \mid I - \mathbf{p}\mathbf{x} \geq 0, x_1 \geq 0, x_2 \geq 0\}. \end{aligned}$$

Implicitly, this defines the three constraints

$$\begin{aligned} g_1(\mathbf{x}) &= I - \mathbf{p}\mathbf{x} \geq 0 \\ g_2(\mathbf{x}) &= x_1 \geq 0 \\ g_3(\mathbf{x}) &= x_2 \geq 0. \end{aligned}$$

1. Step: Set up the Lagrangian as

$$\mathcal{L}(x_1, x_2, \lambda_1, \lambda_2, \lambda_3) = x_1 + x_2 + \lambda_1 x_1 + \lambda_2 x_2 + \lambda_3(I - p_1 x_1 - p_2 x_2).$$

2. Step: The critical points of $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ are the solutions to

$$\begin{aligned} 1 + \lambda_1 - \lambda_3 p_1 &= 0 \\ 1 + \lambda_2 - \lambda_3 p_2 &= 0 \\ \lambda_1 \geq 0, x_1 \geq 0, \lambda_1 x_1 &= 0 \\ \lambda_2 \geq 0, x_2 \geq 0, \lambda_2 x_2 &= 0 \\ \lambda_3 \geq 0, I - p_1 x_1 - p_2 x_2 \geq 0, \lambda_3(I - p_1 x_1 - p_2 x_2) &= 0 \end{aligned} \tag{4.2.4}$$

We have to check all possible combinations of the three inequality constraints: We can immediately rule out $g_1 = g_2 = 0$ since this implies $g_3 > 0$. Since the utility function is strictly increasing, all available income will be spent at the optimal point (no interior solution). Therefore, we know that $g_3 = 0$. Therefore we have only 3 possible values for the set of effective constraints at the optimum g_E . These are

$$\begin{aligned} g_E &= \{g_1, g_3\} \\ g_E &= \{g_2, g_3\} \\ g_E &= \{g_3\}. \end{aligned} \tag{4.2.5}$$

(a) $g_E = \{g_3\}$ means that only g_3 holds with equality. Hence $x_1 > 0$ and $x_2 > 0$. By KT-1, we know that $\lambda_1 = \lambda_2 = 0$. Substituting this into the above (4.2.4) yields

$$\lambda_3 p_1 = \lambda_3 p_2 = 1.$$

This implies $p \equiv p_1 = p_2$ and therefore $\lambda_3 = \frac{1}{p}$. But this makes any point a critical point provided that $x_1 \in (0, \frac{I}{p})$, and $x_2 = \frac{I - p x_1}{p}$; all these points give $u(\cdot) = \frac{I}{p}$.

(b) $g_E = \{g_2, g_3\}$ means that $x_2 = 0$ and $I - \mathbf{p}\mathbf{x} = 0$. Thus $x_1 = \frac{I}{p} > 0$ implying $\lambda_1 > 0$. Substituting this into the above (4.2.4) results in

$$\lambda_3 p_1 = 1 \leq 1 + \lambda_2 = \lambda_3 p_2.$$

This implies that a critical point can only have $p_1 \leq p_2$. This gives a unique critical point at which $g_1 = g_3 = 0$ and

$$(x_1, x_2, \lambda_1, \lambda_2, \lambda_3) = \left(\frac{I}{p_1}, 0, 0, \frac{p_2}{p_1} - 1, \frac{1}{p_1}\right); u(\cdot) = \frac{I}{p_1}.$$

(c) $g_E = \{g_1, g_3\}$ results in similar considerations as (ii) and the unique critical point

$$(x_1, x_2, \lambda_1, \lambda_2, \lambda_3) = \left(0, \frac{I}{p_2}, \frac{p_1}{p_2} - 1, 0, \frac{1}{p_2}\right); u(\cdot) = \frac{I}{p_2}.$$

3. Step: Evaluating our three candidates gives our solution as

$$\mathbf{x}^* = \begin{cases} (0, \frac{I}{p_2}) & \text{if } p_1 > p_2 \\ (\frac{I}{p_1}, 0) & \text{if } p_1 < p_2 \\ (x_1 \in (0, \frac{I}{p}), \frac{I - px_1}{p}) & \text{if } p \equiv p_1 = p_2. \end{cases}$$

Again we check whether the two conditions that we attached to the application of the Kuhn-Tucker Theorem hold.

Condition: (Existence) Our usual appeal to the Weierstrass Theorem establishes that maximum exists.

We also check the bordered Hessians

$$\overline{H}_2(\mathbf{x}^*) = \begin{vmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & -p_1 & -p_2 \\ 1 & 0 & -p_1 & 0 & 0 \\ 0 & 1 & -p_2 & 0 & 0 \end{vmatrix} = 0 \quad (4.2.6)$$

and

$$\overline{H}_1(\mathbf{x}^*) = \begin{vmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -p_1 \\ 1 & 0 & -p_1 & 0 \end{vmatrix} = 0.$$

This says that our objective is quasiconcave and quasiconvex at the same time. The only function that can do this is a (hyper)plane. Therefore we know that our necessary conditions are also sufficient.

Condition: (Constraint qualification) Again we have to check all possible combinations of the three inequality constraints. For the same reasons as in (4.2.5) we arrive at the 3 possibilities

$$\begin{aligned} g_E &= \{g_1, g_3\} \\ g_E &= \{g_2, g_3\} \\ g_E &= \{g_3\}. \end{aligned}$$

We have to check the constraint qualification for each of these cases.

1. $g_E = \{g_1, g_3\}$ results in

$$\nabla g_E(\mathbf{x}) = \begin{bmatrix} 1 & 0 \\ -p_1 & -p_2 \end{bmatrix}$$

at any (x_1, x_2) . Since prices are strictly positive, this matrix has full rank, so the constraint qualification holds.

2. $g_E = \{g_2, g_3\}$ results in

$$\nabla g_E(\mathbf{x}) = \begin{bmatrix} 0 & 1 \\ -p_1 & -p_2 \end{bmatrix}$$

so the constraint qualification holds.

3. $g_E = \{g_3\}$ results in

$$\nabla g_E(\mathbf{x}) = [-p_1, -p_2]$$

which implies that this constraint qualification holds as well.◁

4.3 Duality of linear programs

As promised in section (1), we will discuss the concept of duality in a little more detail now. As we mentioned in the introduction, a most interesting fact about every linear programming problem is that there is another ‘dual’ problem associated with it: If the primal program is a maximisation problem (with, say, N variables x and K constraints λ), then the dual program is a minimisation problem (with K variables and N constraints) and vice versa. This is a general feature of duality: We get a variable in the dual program for each constraint in the primary program and we get a constraint in the dual program for each variable in the primary program. We begin with an example.

Exl: (Mas-Colell, Whinston, and Green 1995, 967) We look at a (primary) maximisation problem with $N = K = 2$

$$\begin{aligned} \max_{x_1, x_2} \quad & x_1 + x_2 \\ \text{s.t.} \quad & 2x_1 + x_2 \leq 4 \\ & x_1 + 3x_2 \leq 7 \end{aligned}$$

The associated (dual) minimisation problem is represented by

$$\begin{aligned} \min_{\lambda_1, \lambda_2} \quad & 4\lambda_1 + 7\lambda_2 \\ \text{s.t.} \quad & 2\lambda_1 + \lambda_2 \geq 1 \\ & \lambda_1 + 3\lambda_2 \geq 1 \end{aligned}$$

Figure 4.5 below represents these two problems graphically. ◁

Solving the primary problem we obtain $\mathbf{x}^* = (1, 2)$, yielding an objective value of $1 + 2 = 3$ while the dual gives us $\boldsymbol{\lambda}^* = (\frac{2}{5}, \frac{1}{5})$ yielding an objective value of $4(\frac{2}{5}) + 7(\frac{1}{5}) = 3$. As somewhat expected they give the same value. But why? A little playing around (in more general form for the coefficient matrix \mathbf{A} on the variables and constant vector \mathbf{c} on the constraints) leads us to find that the solution value to the primal problem cannot be larger than the solution value to the dual problem

$$\mathbf{f} \cdot \mathbf{x}^* \leq (\mathbf{A}^T \boldsymbol{\lambda}^*) \cdot \mathbf{x}^* = \boldsymbol{\lambda}^* \cdot (\mathbf{A} \mathbf{x}^*) \leq \boldsymbol{\lambda}^* \cdot \mathbf{c} = \mathbf{c} \cdot \boldsymbol{\lambda}^*.$$

The *duality theorem of linear programming* now bluntly says that $\mathbf{f} \cdot \mathbf{x}^* = \mathbf{c} \cdot \boldsymbol{\lambda}^*$.




Figure 4.5: (Left) A primal linear program and its dual to the (right).

Theorem 26. Duality Theorem of Linear Programming⁴ Suppose that the primal problem

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}} \quad & \mathbf{f} \cdot \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} \leq \mathbf{c} \end{aligned}$$

attains a maximum at $v \in \mathbb{R}$. Then v is also the minimum value attained by the dual problem

$$\begin{aligned} \min_{\boldsymbol{\lambda} \in \mathbb{R}} \quad & \mathbf{c} \cdot \boldsymbol{\lambda} \\ \text{s.t.} \quad & \mathbf{A}^T \boldsymbol{\lambda} \geq \mathbf{f}. \end{aligned}$$

Proof. See (Mas-Colell, Whinston, and Green 1995, 968). □

There is much more to say about duality than this—it really is the key to understanding the the difference between direct (Marshallian) demand and compensated (Hicksian) demand and the associated concepts like indirect utility and expenditure functions (the same is true for production theory). We have, however, covered the basic ideas and the interested reader must be referred elsewhere: Any good text on (linear) programming such as (Franklin 1980, chapter 1) or (Luenberger 1984) will treat duality; (Cornes 1992) and (Sundaram 1996, chapter 2) provide a host of examples. If you are looking for an illustration of the concept of duality to consumer theory (i.e. the duality of the utility maximisation and the expenditure minimisation approaches), please take a look at the discussion of the envelope theorem and figure 3.10, in particular.

⁴ We talked about the *MiniMax Theorem* previously and said that it implies that for any constant sum game, there exists a saddle point where the payoffs of both players are the same. The Duality Theorem is a more general version of the same result: Maximising player one's payoff function and minimising player two's loss function yields the same value because they are dual problems. Hence the Duality Theorem implies the MiniMax Theorem.

Exercises

Exc 4.1: (Sundaram 1996, 124) Let f and g be functions on \mathbb{R}^2 defined by $f(x, y) = 2x^3 - 3x^2$, and $g(x, y) = (3 - x)^3 - y^2$, respectively. Consider the problem of maximising f over the set $\mathcal{D} = \{(x, y) \in \mathbb{R}^2 \mid g(x, y) = 0\}$.

Exc 4.2: (Sundaram 1996, 148) Let $f : \mathbb{R}^2 \mapsto \mathbb{R}$ and $h : \mathbb{R}^2 \mapsto \mathbb{R}$ be given by $f(x, y) = -(x^2 + y^2)$, and $h(x, y) = (x - 1)^3 - y^2$, respectively. Consider the problem of maximising f over the set $\mathcal{D} = \{(x, y) \in \mathbb{R}^2 \mid h(x, y) \geq 0\}$.

Exc 4.3: (Sundaram 1996, 168:6.1) Solve the following maximisation problem

$$\begin{aligned} & \max_{x_1, x_2} \ln x_1 + \ln x_2 \\ \text{s.t.:} & \text{ (i) } x_1^2 + x_2^2 = 1, \text{ (ii) } x_1 \geq 0, \text{ (iii) } x_2 \geq 0. \end{aligned}$$

Exc 4.4: (Sundaram 1996, 168f:6.2) A firm produces two outputs y and z using a single input x . The set of attainable output levels $H(x)$ from an input use of x , is given by

$$H(x) = \{(x, y) \in \mathbb{R}_+^2 \mid y^2 + z^2 \leq x\}.$$

The firm has available to it a maximum of one unit of input x . Letting (p_1, p_2) denote output prices, determine the firm's optimal output mix.

Exc 4.5: (Sundaram 1996, 170:6.9) An agent who consumes three commodities has a utility function given by

$$u(\mathbf{x}) = x_1^{\frac{1}{3}} + \min\{x_2, x_3\}.$$

Given an income of I and prices $\mathbf{p} = (p_1, p_2, p_3)$, describe the consumer's utility maximisation problem. Can the Weierstrass and/or Kuhn-Tucker Theorems be used to obtain and characterise a solution? Why or why not?

Chapter 5

Measure, probability, and expected utility

Reading: A very readable introduction to measure theory is (Capiński and Kopp 1999)—(Wheeden and Zygmund 1977) is more complete but a harder read.

This section attempts to form a minimal but rigorous introduction to measure theory. We will look at the theory of *Lebesgue* measure that we will need in order to discuss some aspects of the notion of a random draw and infinite sets. The purpose of this chapter is to introduce the modern notion of probability theory as used in all areas of economics. In particular, we will lay the foundations for the discussion of probability and expected utility theory in the next chapter.

5.1 Measure

5.1.1 Measurable sets

Let us informally define the **length** of an interval $I = [a, b]$ to be $l(I) = |b - a|$. Notice that this length operator is invariant to the different types of interval limits $\{[(\}a, b\{)]\}$ we can use. Now consider a singleton $\{c\}$ with interval length $l(c) = |c - c| = 0$. What we are after is to show that a set that can be constructed as the union of a finite number of null-sets still has length zero. Any interval $[a, b]$ with $a \neq b$, however, will have non-zero length.

Def: A collection of sets $\mathcal{I} = \{I_n\}_{n \geq 1}$ is said to **cover** the set A if

$$A \subset \bigcup_{n \geq 1} I_n.$$

If \mathcal{I} covers A and there exists a subset \mathcal{J} of \mathcal{I} which also covers A , then \mathcal{J} is called a **subcover** of A .

Def: A **null-set** $A \subset \mathbb{R}$ is a set that may be covered by a sequence of intervals of arbitrarily small

total length, i.e. given any $\varepsilon > 0$ we can find a sequence $\{I_n : n \geq 1\}$ of intervals such that

$$A \subset \bigcup_{n=1}^{\infty} I_n \quad \text{and} \quad \sum_{n=1}^{\infty} l(I_n) < \varepsilon.$$

Rem: Clearly, any singleton $\{x\}$ is a null-set. More generally, any countable set A is also a null-set and any finite union of countable sets is a null-set as well. Not obvious at all is, however, that even uncountable sets can be null. The prime example for such a set is the **Cantor set** C .

Before we discuss this set, let us state an observation which you perhaps remember from school. Every $x \in [0, 1] \subset \mathbb{R}$ can be approximated arbitrarily closely by an infinite rational expansion of the form

$$x = \frac{a_1}{n} + \frac{a_2}{n} + \frac{a_3}{n^2} + \cdots + \frac{a_k}{n^k} + \cdots \quad (5.1.1)$$

where the numbers a_1, a_2, \dots can only take the integer values $0, 1, \dots, n$.

Exl: For the construction of the Cantor set we start with the unit interval $C_0 = [0, 1] \subset \mathbb{R}$ and take out the ‘open middle third’ $(\frac{1}{3}, \frac{2}{3})$ such that we are left with $C_1 = [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$. We proceed by again removing the open middle thirds of each of the remaining two subintervals and are left with $C_2 = [0, \frac{1}{9}] \cup [\frac{2}{9}, \frac{3}{9}] \cup [\frac{6}{9}, \frac{7}{9}] \cup [\frac{8}{9}, 1]$. We go on deleting the middle third of each of these subintervals until we reach, after n steps, the set $C_n = [0, \frac{1}{3^n}] \cup \cdots \cup [1 - \frac{1}{3^n}, 1]$. Then the Cantor set C is defined to be

$$C = \bigcap_{i=0}^{\infty} C_i.$$

Since C consists of 2^n intervals of length $\frac{1}{3^n}$ each, we know that the total length of the set deleted from the Cantor set $l([0, 1] \setminus C)$ is

$$\frac{1}{3} + \frac{2}{9} + \frac{4}{27} + \cdots = \sum_{i=1}^{\infty} \frac{2^i}{3^i} = \sum_{i=1}^{\infty} \left(\frac{2}{3}\right)^i \stackrel{i \rightarrow \infty}{=} 1.$$

Hence, in the limit, we delete *all* of the unit interval. Yet given any $\varepsilon > 0$, we can choose an n large enough to ensure $l(C) = \left(\frac{2}{3}\right)^n < \varepsilon$. Since $C \subset C_n$ by construction and C_n consists of a finite sequence of intervals that can be covered with a set of length smaller than ε , we know that C is indeed a null-set.

Now similarly to our argument in the proof of theorem 1, we can express each $x \in [0, 1]$ in ‘ternary’ form, i.e. (5.1.1) with $n = 3$ as

$$x = \frac{a_1}{3} + \frac{a_2}{3^2} + \cdots + \frac{a_k}{3^k} + \cdots = \sum_{k=1}^{\infty} \frac{a_k}{3^k} \quad (5.1.2)$$

with each $a_k \in \{0, 1, 2\}$. To believe that this sequence can indeed approximate every $x \in [0, 1]$, it helps to remember a basic fact about infinite power series $\sum_{i \geq 0} ct^i$, which, for $|t| < 1$, sum to

$$c + ct + ct^2 + \cdots = \sum_{i=0}^{\infty} ct^i = \frac{c}{1-t}.$$

This corresponds to a ternary expansion with, e.g., all $a_k = 1$, of

$$1 + \frac{1}{3} + \frac{1}{9} + \frac{1}{27} + \dots = \sum_{i=0}^{\infty} \frac{1}{3^i} = \sum_{i=0}^{\infty} \left(\frac{1}{3}\right)^i = \frac{1}{1 - \frac{1}{3}} = \frac{3}{2}.$$

Subtracting the leading term which is not in our sequence (5.1.2), we find that

$$\frac{1}{3} + \frac{1}{9} + \frac{1}{27} + \dots = \frac{1}{2}.$$

Obviously the sequence of all $a_k = 0$ sums to 0 and the sequence of all $a_k = 2$ sums to 1. So we can indeed generate all $x \in [0, 1]$. But what we really want to find is a representation for $y \in C$. It can be easily seen, however, that we have $y \in C$ iff $a_k \neq 1, \forall k$ in the above expansion (5.1.2) (can you see why?).

To see that C is indeed uncountable, we associate a point $z \in [0, 1]$ of the form

$$z = \sum_{k=1}^{\infty} \frac{b_k}{2^k}$$

with each point $x \in C$ expanded in (5.1.2) but instead of the a_k used there, we use

$$b_n = \begin{cases} 1 & \text{if } a_n = 2, \\ 0 & \text{otherwise.} \end{cases}$$

In this fashion we set up a one-to-one correspondence of the elements of C with $[0, 1]$ which we know to be uncountable. Hence we proved that the Cantor set is uncountable. \triangleleft

Our discussion of null-sets is important because their contribution is what we want to ignore in our general concept of length l . One way of incorporating this idea in a length is Lebesgue's outer measure.

Def: The (Lebesgue) **outer measure** of any set $A \subset \mathbb{R}$ is the non-negative real number

$$m^*(A) = \inf(Z_A)$$

where¹

$$Z_A = \left\{ \sum_{n=1}^{\infty} l(I_n) \mid I_n \text{ are intervals, } A \subset \bigcup_{n=1}^{\infty} I_n \right\}.$$

It can be shown that the outer measure $m^*(\cdot)$ has all the nice properties we would like it to have:

1. $m^*(A) = 0$ for any null-set A ,
2. $m^*([a, b]) = |b - a|$ for an interval $[a, b]$,
3. $m^*(\cdot)$ is monotone (the bigger the set, the greater its measure),
4. $m^*(A) = m^*(A + t)$, i.e. the outer measure is invariant to translations t , and

¹ Since the above series may not converge, let us further adhere to the following conventions: $m^*(\cdot) + \infty = \infty$, $\infty + \infty = \infty$, and $0 \times \infty = 0$.

5. $m^*(\bigcup_n E_n) \leq \sum_n m^*(E_n)$, i.e. the outer measure is countably subadditive for any sequence of sets $\{E_n\}_{n \geq 1}$.

We will, however, only prove the first two properties and confine ourselves with a discussion of the others. Proofs of the remaining properties can be found in (Kolmogorov and Fomin 1970), (Royden 1988) or any other good textbook on real analysis.

Theorem 27. $A \subset \mathbb{R}$ is a null-set iff $m^*(A) = 0$.

Proof. From the definition of a null-set we obtain a sequence of intervals $\{I_n\}$ covering A such that $\sum l(I_n) < \varepsilon$ for any $\varepsilon > 0$. Hence the infimum of such a collection must be zero.

Conversely, if $m^*(A) = 0$, then by the definition of inf there exists a $z \in Z_A$ for any $\varepsilon > 0$ such that $z < \varepsilon$. But a $z \in Z_A$ is a cover of A and therefore A must be a null-set. \square

Theorem 28. The outer measure $m^*(I)$ of an interval I equals its length $l(I)$.

Proof. Clearly, $m^*(I) = l(I) = \infty$ if I is unbounded. Hence we can restrict attention to bounded intervals.

(i) To show that $m^*(I) \leq l(I)$, let us cover I using a sequence of intervals $\{I_n\}_{n \geq 1}$ of the form $I_1 = I$, $I_2 = [0, 0]$, $I_3 = [0, 0], \dots$. The length of this sequence is equal to $l(I)$ and therefore we know that $l(I) \in Z_I$. This is sufficient since the infimum of a set of lengths cannot exceed that of any of its elements.

(ii) To show that $m^*(I) \geq l(I)$, let $I = [a, b]$ and $a < b$ wlg.² For an arbitrary $\varepsilon > 0$ we choose a sequence of intervals $\{I_n\}_{n \geq 1}$ covering $[a, b]$ such that

$$\sum_{n=1}^{\infty} l(I_n) \leq m^*([a, b]) + \frac{\varepsilon}{2}. \quad (5.1.3)$$

Denote the endpoints of the above sequence by $I_n = \{[a_n, b_n]\}_{n \geq 1}$ and extend the sequence of closed intervals I_n to the sequence of open intervals J_n

$$\left\{ \left(a_n - \frac{\varepsilon}{2^{n+2}}, b_n + \frac{\varepsilon}{2^{n+2}} \right) \right\}_{n \geq 1}.$$

Clearly,

$$l(I_n) = l(J_n) - \frac{\varepsilon}{2^{n+1}} \quad \text{or} \quad \sum_{n=1}^{\infty} l(I_n) = \sum_{n=1}^{\infty} l(J_n) - \frac{\varepsilon}{2}.$$

We insert this into (5.1.3) and obtain

$$\sum_{n=1}^{\infty} l(J_n) \leq m^*([a, b]) + \varepsilon \quad (5.1.4)$$

which apparently covers $[a, b]$.

From the Heine-Borel Theorem (6) we know that a finite number of intervals n is sufficient to

² We leave the other 2 cases of $I = (a, b)$ and $I = (a, b]$ or $I = [a, b)$ as an exercise.

cover $[a, b]$ because $[a, b]$ is compact in \mathbb{R} . So there exists a finite m such that³

$$[a, b] \subset \bigcup_{n=1}^m J_n.$$

Let, again, be the endpoints of the open intervals denoted $J_n = \{(c_n, d_n)\}$. Let

$$c = \inf_{c \geq 1} \{c_n\} \quad \text{and} \quad d = \sup_{c \geq 1} \{c_n\}$$

with $c < a$ and $d > b$ and therefore $l([a, b]) < d - c$. Since the intervals $\{J_n\}_{n=1}^m$ must overlap by construction, we have

$$l([a, b]) < d - c < \sum_{n=1}^m l(J_n).$$

Since with non-negative terms J_n , the finite sum above is less than the infinite sum in (5.1.3), (5.1.3) follows from the above and (5.1.4). \square

Until now we have discussed the properties of Lebesgue outer measure only for sets along the real line. We will drop this simplification now and use a discussion of the subadditivity of the outer measure to generalise our idea and define the general notion of a Lebesgue measure.

Above, we claimed without proof that the outer measure is countably subadditive for general sets E_n . Let us restate this property here:

$$m^*\left(\bigcup_n E_n\right) \leq \sum_n m^*(E_n). \quad (5.1.5)$$

As it turns out, this is the key to understanding the limitations of the outer measure. We would really like to obtain 5.1.5 with equality, i.e. we would like to show *general* additivity if the sets $\{E_n\}$ are pairwise disjoint. That, alas, is out of reach. What we can do, however, is to stop here, register that our problem is precisely to claim general additivity, and postulate the desired properties for our theory. Which is what we will do. Hence for pairwise disjoint $\{E_n\}_{n \geq 1}$ we claim countable additivity

$$m^*\left(\bigcup_n E_n\right) = \sum_n m^*(E_n). \quad (5.1.6)$$


Similarly we require our concept of ‘length’ for sets to be additive. In particular, we shall want our sets E to be able to split every other set A in an additive manner. Formally we define

Def: A set $E \subset \mathbb{R}$ is (Lebesgue-) **measurable** if for every set $A \subset \mathbb{R}$ we have

$$m^*(A) = m^*(A \cap E) + m^*(A \cap E^c)$$

where $E^c \equiv \mathbb{R} \setminus E$. We denote the whole class of (Lebesgue-)measurable sets by \mathcal{M} and write $E \in \mathcal{M}$.

³ For ease of notation we keep using the (infinite) sequence $\{J_n\}$ but we keep in mind that we only use the first finite m terms.

Figure 5.1: A measurable set $A \in \mathcal{M}$.

Rem: Once again, we know that by defining \mathcal{M} in this way we impose the implicit restriction that the sets we use to additively split the measurable set must be disjoint (as E and E^C obviously are). This was indeed the reason why we previously laid so much emphasis on the fact that open sets can be covered by a (finite) number of *disjoint* intervals. If we do not impose this restriction, for any $A, E \subset \mathbb{R}$, we are only able to claim

$$E \in \mathcal{M} \Leftrightarrow m^*(A) \leq m^*(A \cap E) + m^*(A \cap E^c) \quad (5.1.7)$$

as a criterion for checking measurability.

Theorem 29. *Any null-set N is measurable.*

Proof. We know from Theorem 27 that $m^*(N) = 0$ if N is a null-set. So for any $A \subset \mathbb{R}$ we have

$$m^*(A \cap N) \leq m^*(N)$$

since $A \cap N \subset N$. And

$$m^*(A \cap N^C) \leq m^*(A)$$

since $A \cap N^C \subset A$. Adding the two together implies (5.1.7). \square

Theorem 30. *Any interval is measurable.*⁴

Proof. Suppose $I = [a, b]$ and take any $\varepsilon > 0$ and $A \subset \mathbb{R}$. Then we choose a cover for A such that

$$m^*(A) \leq \sum_{n=1}^{\infty} l(I_n) \leq m^*(A) + \varepsilon.$$

Now let $I'_n = I_n \cap [a, b]$, $\forall n$. Clearly these cover $A \cap [a, b]$ such that we also have

$$m^*(A \cap [a, b]) \leq \sum_{n=1}^{\infty} l(I'_n).$$

⁴ We again only bother to prove the case of the closed interval. The theorem, however, holds in general.

Similarly let $I_n'' = I_n \cap [-\infty, b]$ and $I_n''' = I_n \cap [a, \infty]$, $\forall n$ cover $[a, b]^C$ such that we have

$$m^*(A \cap [a, b]^C) \leq \sum_{n=1}^{\infty} l(I_n'') + \sum_{n=1}^{\infty} l(I_n''').$$

Putting the three inequalities together again implies (5.1.7). \square

Theorem 31. *We conclude that*

1. $\mathbb{R} \in \mathcal{M}$;
2. if $E \in \mathcal{M}$, then if $E^C \in \mathcal{M}$;
3. if $\{E_n\}_{n \geq 1} \in \mathcal{M}$, then

$$\bigcup_{n=1}^{\infty} E_n \in \mathcal{M};$$

4. if $\{E_n\}_{n=1}^{\infty} \in \mathcal{M}$, then

$$m^*\left(\bigcup_{n=1}^{\infty} E_n\right) = \sum_{n=1}^{\infty} m^*(E_n).$$

Proof. See (Capiński and Kopp 1999, 29ff). \square

Def: For any $E \in \mathcal{M}$ we write $m(E)$ instead of $m^*(E)$ and call $m(E)$ the **Lebesgue measure** of the set E .

Rem: Hence the Lebesgue measure $m = \mathcal{M} \mapsto [0, \infty]$ is a countably additive set function defined on a σ -algebra \mathcal{M} of measurable sets (properties 1-3 above). The Lebesgue measure of an interval is its length, and the Lebesgue measure of a null-set is zero. We will return to the discussion of the Lebesgue measure during the discussion of probability.

We can use the tools we developed above for Lebesgue measures on arbitrary sets: Any set Ω carrying an outer measure can be equipped with a measure μ defined on a σ -algebra \mathcal{F} of subsets of Ω . The resulting triple $(\Omega, \mathcal{F}, \mu)$ is then called a **measure space**. We will confine ourselves, however, with only one particular example: that of probability theory.

Def: Let $\Omega = (\Omega, \mathcal{F}, \mu)$ be any measure space. A set $E \in \mathcal{F}$ is called an **atom** of Ω such that $\mu(E) > 0$ and, whenever $E \supset F \in \mathcal{F}$, either

$$\mu(F) = 0, \quad \text{or} \quad \mu(E \setminus F) = 0.$$

If Ω has no atoms, it is called **atomless**.

5.1.2 Integrals and measurable functions

With measurable sets we measured the Lebesgue-measure of sets. With measurable functions we do not really measure the functions but the areas under them: We integrate. Therefore our intention is to motivate the idea of the *Lebesgue-integral* by showing some of the shortcomings of the usually employed *Riemann-integral* you know from school (or the September course). Let's

refresh our memories. Let $f : [a, b] \mapsto \mathbb{R}$ be a bounded function with $a, b \in \mathbb{R}$ and P be a partition of $[a, b]$. Then, for each $i \leq n$, we define the **upper (Riemann) sum** as

$$U(P, f) = \sum_{i=1}^n \sup_{a_{i-1} \leq x \leq a_i} f(x)(a_i - a_{i-1}),$$

and the **lower (Riemann) sum** as

$$L(P, f) = \sum_{i=1}^n \inf_{a_{i-1} \leq x \leq a_i} f(x)(a_i - a_{i-1}).$$

The **upper integral** is the infimum of the set of upper sums (that become smaller and smaller) for successively finer partitions and the **lower integral** is the supremum of the set of lower sums (that become larger and larger) for successively finer partitions of the interval we integrate over. A function is said to be **Riemann-integrable** on $[a, b]$ if the lower and the upper integral coincide. Their common value is called the **Riemann-integral**.

A useful criterion for checking the integrability is the

Theorem 32. (Riemann Criterion) *A function $f : [a, b] \mapsto \mathbb{R}$ is Riemann-integrable iff for every $\varepsilon > 0$ there exists a partition P_ε such that $U(P_\varepsilon, f) - L(P_\varepsilon, f) < \varepsilon$.*

The important result that links the theory of integration to that of differentiation is

Theorem 33. (Fundamental Theorem of Calculus) *Let $F : [a, b] \mapsto \mathbb{R}$ be equipped with the continuous derivative $f : [a, b] \mapsto \mathbb{R}$, then*

$$F(b) - F(a) = \int_a^b f(x) dx. \quad (5.1.8)$$

To see the shortcomings associated with this notion of integrability let us look for an answer to the following questions:

1. How do we integrate over null-sets? Apparently the ‘area’ above or below the function we want to integrate depends on the length of the interval $[a, b]$ we integrate over. If $a = b$ we expect both areas to be zero. But then, by (1), any countable set must also have zero ‘area’!
2. How can we integrate over unbounded sets or functions without resorting to ‘improper’ integrals?
3. How can we integrate functions whose values are ‘awkwardly’ distributed? What we mean by ‘awkward’ is exemplified by the **indicator** function of \mathbb{Q} over $[0, 1]$

$$\mathbf{1}_{\mathbb{Q}}(x) \equiv \begin{cases} 1 & \text{if } x \in \mathbb{Q} \\ 0 & \text{otherwise.} \end{cases} \quad (5.1.9)$$

If we additionally require that each sub-interval of $[0, 1]$ which we allow to be picked for integration contains both elements of \mathbb{Q} and \mathbb{R} , then we will always get the ‘upper sum’ as 1 and each ‘lower sum’ as 0. Consequently the Riemann-integral does not exist because the function is discontinuous at *each* point.

4. How can we integrate discontinuous functions?

What we have achieved until now is the following: We have developed a concept of some null-set N that allows us to define certain properties on e.g. $\mathbb{R} \cup N$ as if we would define them on just \mathbb{R} . This really gives already an answer to most of the questions above, formally

Def: We say that a function f has a property P **almost everywhere** (a.e.) (or **almost surely**) if f has this property at all points of its domain except (possibly) on some null-set.

Exl: The function

$$f(x) = \begin{cases} 0 & \text{for } x = 0, \\ 1 & \text{otherwise.} \end{cases}$$

is a.e. continuous because $\{0\}$ is a null-set and f is continuous on $\mathbb{R} \setminus \{0\}$. \triangleleft

Given the Riemann approach to the problem of finding the area under general functions, the Lebesgue integral is really very easy to understand. Consider $f : X \mapsto Y$. We simply partition not the domain X but the range Y of the function into disjoint (short) intervals $\{I_i\} = \{[a_i, a_{i+1}]\}_{i \geq 1}$ of the same length $l([a_i, a_{i+1}])$ and look at the pre-images of these intervals in the domain $f^{-1}([a_i, a_{i+1}])$. In general these pre-images will not be intervals but more complicated sets. If these sets are measurable, we will call the function f measurable, too.

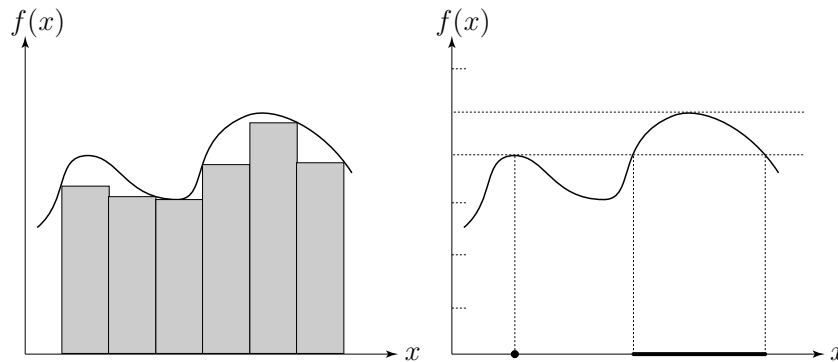


Figure 5.2: Riemann (left) and Lebesgue integrals (right).

More formally, for each interval $[a_i, a_{i+1}]$ we define a number c_i corresponding to the ‘height’ of the interval I_i . For instance we can use $c_i = \inf I_i$. Then, given all $\{I_i\}_{i \geq 1} \in \mathcal{M}$, (the simplest form of) the Lebesgue integral is defined as

$$\sum_{i=1}^N c_i m(f^{-1}(I_i)).$$

Def: Suppose E is a measurable set. We say that a function $f : E \mapsto \mathbb{R}$ is (Lebesgue-)measurable if for any interval $I \subset \mathbb{R}$ ⁵

$$f^{-1}(I) = \{x \in \mathbb{R} : f(x) \in I\} \in \mathcal{M}.$$

Exl: Any continuous function is measurable. Moreover, every monotone function is measurable.

⁵ A function is sometimes called Borel-measurable if the sets $f^{-1}(I)$ are Borel sets but we will not make any distinction.

Even the indicator function (5.1.9) defined on page 96 is measurable. \triangleleft

Theorem 34. *The set of real-valued measurable functions defined on $E \in \mathcal{M}$ is a vector space closed under addition and multiplication.*⁶

Proof. See (Capiński and Kopp 1999, p58). □

Theorem 35. *If $\{f_n\}$ is a sequence of measurable functions defined on $E \in \mathcal{M}$, then*

$$\max_{n \leq k} f_n, \quad \min_{n \leq k} f_n, \quad \sup_{n \in \mathbb{N}} f_n, \quad \inf_{n \in \mathbb{N}} f_n, \quad \limsup_{n \rightarrow \infty} f_n, \quad \liminf_{n \rightarrow \infty} f_n.$$

are measurable as well.

Proof. See (Capiński and Kopp 1999, p62f). □

Theorem 36. *If $f : E \mapsto \mathbb{R}$ is measurable and $E \in \mathcal{M}$, and we consider an arbitrary function $g : E \mapsto \mathbb{R}$ for which $\{x : f(x) = g(x)\}$ is a null-set, then $g(\cdot)$ is measurable as well.*

Proof. See (Capiński and Kopp 1999, p63). □

Corollary 37. *If $\{f_n\}$ is a sequence of measurable functions defined on $E \in \mathcal{M}$ and $f_n(x) \rightarrow f(x)$ almost everywhere, then $f(\cdot)$ is measurable.*

Proof. See (Capiński and Kopp 1999, p63). □

5.2 Probability

When we talk about a *random variable* we mean some quantity that depends on the outcome of some random experiment such as the toss of a die, draw from an urn, a laboratory measurement, etc. Such a random variable takes values on a *sample space*. The sample space is the collection of possible outcomes of an experiment and is usually denoted by Ω . The possible individual outcomes themselves are denoted $\omega \in \Omega$.

Def: A **random variable** X is a function $X : \Omega \mapsto \mathbb{S}$ taking the sample space into the state space.

An *event* $E \subset \Omega$ is a set of possible outcomes for which the event occurs. The collection of events we consider is usually denoted by \mathcal{F} . The biggest possible such collection would be the power set $\wp(\Omega)$ but it turns out that that is ‘too big’ for most problems with continuous distributions. For technical reasons relating to probabilities (i.e. that we want them to be non-negative and sum to 1) we require \mathcal{F} to have certain closure properties. The mathematical construct that meets these requirements is called a σ -algebra. To clarify the origin of the term, we begin by recalling the concepts of a field, a ring and an algebra from the introductory chapter 2.

Def: A **field** $\mathcal{F} = \langle X, +, \cdot \rangle$ is a system of elements X which satisfies for all $a, b, c \in X$ the **field axioms** for both addition and multiplication

⁶ I.e. if f, g are measurable functions, then so is $f + g$ and fg .

1. commutativity: $a + b = b + a$ and $ab = ba$;
2. associativity: $a + (b + c) = (a + b) + c$ and $a(bc) = (ab)c$;
3. distributivity: $a(b + c) = ab + ac$ and $(a + b)c = ac + bc$;
4. identity: $a + 0 = a = 0 + a$ and $a \cdot 1 = a = 1 \cdot a$;
5. inverse: $a + (-a) = 0$ and $aa^{-1} = 1 = a^{-1}a$ for $a \neq 0$ (a^{-1} is called the multiplicative inverse);

Rem: A **ring** $\mathcal{R} = \langle X, +, \cdot \rangle$ is an Abelian group under addition and a semigroup under multiplication.

Examples for fields include the complex numbers \mathbb{C} , the rational numbers \mathbb{Q} , and real numbers \mathbb{R} , but not the integers \mathbb{Z} , which only form a ring.

Def: An set E is called the **unit** of a system of sets \mathcal{S} , if $E \in \mathcal{S}$ and $A \cap E = A$ for all $A \in \mathcal{S}$.

Def: An **algebra** is a ring containing a unit.

Def: A **σ -algebra** (or **Borel algebra**) is the collection \mathcal{F} of all measurable subsets of R fulfilling

1. $R \in \mathcal{F}$, $\emptyset \in \mathcal{F}$,
2. if $E \in \mathcal{F}$ then $E^c \in \mathcal{F}$,
3. if $E_n \in \mathcal{F}$ for all $n = 1, 2, \dots$ and $E_j \cap E_k = \emptyset$ for $j \neq k$, then

$$m^*\left(\bigcup_{n=1}^{\infty} E_n\right) = \sum_{n=1}^{\infty} m^*(E_n).$$

Exl: Using the above notation, both (i) $\mathcal{F} = \{\emptyset, R\}$, and (ii) $\mathcal{F} = \emptyset(R)$ are σ -algebras of R . \triangleleft

Def: Given any nonempty system of sets R , the unique σ -algebra containing R and contained in every algebra contained in R is called the **minimal** σ -algebra generated by R .

Def: The **Borel sets** are the subsets of the real line belonging to the minimal σ -algebra generated by the set of all open (or closed, or half-open) intervals (a, b) , $a < b \in \mathbb{R}$. This particular (minimal) σ -algebra is typically denoted \mathcal{B} .

Exl: By the above definition, all intervals are Borel sets and thus belong to \mathcal{B} . Since \mathcal{B} is a σ -algebra, all open sets must belong to \mathcal{B} , too, because any open set is a countable union of intervals. Since the complements of elements are also contained in a σ -algebra, \mathcal{B} also contains all closed intervals. (You are asked in the exercises to show this.) \triangleleft

Def: A **probability assignment** (or **probability measure**) is a rule assigning a number $0 \leq \Pr(A) \leq 1$ to each event $A \in \mathcal{F}$:

1. $0 \leq \Pr(A) \leq 1$ for any $A \in \mathcal{F}$;
2. $\Pr(\Omega) = 1$;
3. For any sequence of disjoint events A_1, A_2, \dots (i.e., $A_i \cap A_j = \emptyset$ for $i \neq j$),

$$\Pr\left(\bigcup_{i \geq 1} A_i\right) = \sum_{i \geq 1} \Pr(A_i).$$

Def: A **probability (measure) space** is a triple $(\Omega, \mathcal{F}, \Pr)$ where Ω is an arbitrary set, \mathcal{F} is a σ -field of subsets of Ω , and \Pr is a measurable function on \mathcal{F} such that $\Pr(\emptyset) = 0$ and $\Pr(\Omega) = 1$.

Def: Suppose that events $X, E \in \mathcal{F}$ and $\Pr(E) > 0$, then the number

$$\Pr(X|E) = \frac{\Pr(X \cap E)}{\Pr(E)} \quad (5.2.1)$$

is called the **conditional probability** of X given E . (5.2.1) is called **Bayes' Rule**.

Rem: Bayes' Rule computes the conditional probability of X, E when the events X and E are not independent. The interpretation is that there is a prior probability for X that is revised to the posterior $\Pr(X|E)$ after the event E happened.

The probability of $\Pr(X \cap E)$ —i.e. that both events happen together—is calculated as

$$\Pr(X \cap E) = \Pr(E) \Pr(X|E),$$

where $\Pr(X|E)$ is the conditional probability of X given E . Turning this around leads to

$$\Pr(X|E) = \frac{\Pr(X \cap E)}{\Pr(E)}.$$

Since $\Pr(X|E) \Pr(E) = \Pr(X \cap E) = \Pr(E|X) \Pr(X)$, **Bayes' Rule** follows as

$$\Pr(X|E) = \frac{\Pr(E|X) \Pr(X)}{\Pr(E)}.$$

Exl: Assume an agent who wants to submit a bid for an object is uncertain whether the value of the object v is 0 or 3. He puts prior probability $\Pr(v = 3) = 1/2$ on each of the two possible states. Further assume that the agent receives some signal $s \in \{0, 1\}$ before he bids which gives him additional information. The signal's conditional probabilities $\Pr(s|v)$ are given by:

$\Pr(s v)$	$v = 0$	$v = 3$
$s = 0$	$1/2$	$1/2$
$s = 1$	$1/4$	$3/4$

Given the signal $s = 1$, he will revise his prior $\Pr(v = 3)$ as follows by applying Bayes' Rule

$$\Pr(v = 3|s = 1) = \frac{\Pr(s = 1|v = 3) \Pr(v = 3)}{\Pr(s = 1)}.$$

Clearly, $\Pr(s = 1|v = 3) = 3/4$, the prior $\Pr(v = 3) = 1/2$, and $\Pr(s = 1) = (1/2)1/4 + (1/2)3/4 = 4/8$. So the posterior $\Pr(v = 3|s = 1)$ is

$$\Pr(v = 3|s = 1) = \frac{3/4 \cdot 1/2}{1/2} = \frac{3}{4}.$$

Given the signal $s = 0$, the agent learns nothing, i.e. his posterior equals his prior. \triangleleft

A natural way of saying that the event X is probabilistically independent of another event E is that if, whenever E takes place, this has no influence on the probability of X , i.e. $\Pr(X|E) = \Pr(X)$. By the above definition this implies

Def: The events $X, E \in \mathcal{F}$ are **independent** if

$$\Pr(X \cap E) = \Pr(X) \Pr(E).$$

The distribution of a random variable X is the assignment of probabilities to the event $X \in A$ for each set $A \subset S$, i.e., the specification of probabilities to sets of outcomes of X . For real-valued random variables we specify the probabilities for events of the form $[X \leq x_0]$ (i.e., $X \in A_{x_0}$ for sets $A_{x_0} = (-\infty, x_0]$). This function is called the cumulative distribution function of X .

Def: A **cumulative distribution function** (cdf) gives the probability of random variable X being smaller or equal to x_0

$$F(X) \equiv \int_{-\infty}^{x_0} f(x) dx.$$

If this cdf is a continuous and differentiable function, its derivative $f(x_0) \equiv F'(x_0)$ is called the *probability density function* (pdf) of X , and X is said to have a continuous distribution. If instead it is piecewise constant, changing only by jumps of size p_k at the possible values x_k of X , then the function assigning $p_k = f(x_k)$ is called the *probability function* for X and X is called discrete.

Def: A **probability density function** (pdf) gives the probability density of a random variable X taking a particular value x_0

$$f(X) \equiv \Pr(x_0).$$

Rem: Let X be a random variable with pdf $f(X)$ and C^1 cdf $F(X)$. Then we have the following useful properties

1. $F'(X) = f(X)$,
2. $\int_{-\infty}^{+\infty} f(X) dx = 1$.

Obviously, the probability assignment $A \mapsto \Pr(X \in A)$ is the probability distribution of X . It is easy to show that this assignment satisfies the rules required in the definition above. Hence it is a probability measure on \mathcal{F} , the subsets of R . It is often denoted μ_X and gives the probability $\mu_X(A) = \Pr(X \in A)$ that X will lie in each subset $A \subset S$ of possible values of X . For continuous and discrete random variables, respectively, the distributions are given by

$$\mu_X(A) = \int_A f(x) dx \quad \text{or} \quad \mu_X(A) = \sum_{x_i \in A} p_i.$$

The ‘same’ random variable X can be constructed in lots of different ways, with lots of different sample-spaces and even different state spaces—but the distribution μ_X is unique, whatever Ω we use. In fact, for real-valued random variables we can *always* use the **canonical probability space** $(\mathbb{R}, \mathcal{B}, \mu_X)$, thinking of the ‘random experiment’ as simply observing X , and the ‘possible outcomes’ as all real numbers $x \in \mathbb{R}$. Then the probability assignment is just the distribution

$$\mu_X(A) = \Pr(X \in A).$$

Exl: The above example of Bayes' Rule is an application of Bayes' Rule to random variables having discrete distributions. This remark concerns itself with the obvious extension to a continuously, independently and identically distributed distributed random variable X . Suppose there is a prior belief on some parameter μ which is given by a density $g(\mu)$. Let the conditional density of X given μ be $f(X|\mu)$. We construct the posterior $g(\mu|X)$ in the same way as in the previous example and obtain Bayes' Rule for the continuously distributed random variable X as

$$g(\mu|X) = \frac{g(\mu)f(X|\mu)}{\int f(X|\mu)g(\mu)d\mu}.$$

If $g(\mu)$ and $f(X|\mu)$ are both normal, O'Hara (1995, p87) derives the posterior distribution as normal, too. This is a very nice feature: If the prior is normal and the random variable is normally distributed, too, then the posterior remains normal. \triangleleft

5.3 Expected utility

Reading: A systematic treatment of expected utility theory is (Mas-Colell, Whinston, and Green 1995, chapter 6).

Expected utility theory is the workhorse of modern economics under uncertainty. The theory has been developed by John von Neumann and Oskar Morgenstern in the late 20's of last century and was at the centre of their later 'Theorie der Gesellschaftsspiele.' In this chapter we will explain the well-known problems the theory faces and try to convince you that in general the theory can overcome some but not all of these problems. We will indicate in which direction the theory can be modified to accommodate all known 'paradoxes.' We will introduce and discuss measures of risk attitude of decision makers. We will end with the discussion of several measures of the riskiness of generalised lotteries that will allow us to apply the integral results we derived in the previous chapter.

Def: The **expected value** of a random variable x drawn from the continuous distribution $F(x)$, symbolically $x \sim F(x)$ is

$$\mathbb{E}[x] = \int_{-\infty}^{+\infty} x dF(x).$$

Def: A **lottery** $v = (p_1, \dots, p_n)$, $p_i \geq 0$, $\sum p_i = 1$ is a collection of probabilities for the sure outcomes x_1, \dots, x_n . We denote the space of lotteries by L .

Def: In the discrete case, a **von Neumann-Morgenstern expected utility function** $U(\cdot)$ represents preferences among lotteries $v = (p_1, \dots, p_n) \in L$ such that⁷

$$U(v) = \sum_{i=1}^n p_i u(x_i).$$

Def: If we represent the probabilities in a lottery by a continuous random variable, the **expected**

utility of this random variable $x \sim F(x)$ is⁸

$$U = \int_{-\infty}^{+\infty} u(x)dF(x).$$

Rem: Since the single events in a cdf $F(\cdot)$ sum to one, $F(\cdot)$ can be interpreted as a lottery over mutually exclusive outcomes.

Rem: Bernoulli utility functions of risk averse agents in $[x, u(x)]$ space are concave. The more concave they are, the more risk-averse is the agent (as depicted in figure 5.3). Conversely, the expected utility functions of risk loving agents are strictly convex. Notice that this result follows directly from **Jensen's inequality**

$$\int_{-\infty}^{+\infty} u(x)dF(x) \leq u\left(\int_{-\infty}^{+\infty} x dF(x)\right). \quad (5.3.1)$$




Figure 5.3: Expected utility functions of a risk averse (left) and a risk loving (right) decision-maker.

Rem: The level sets of expected utility functions for n outcomes can be represented in an $n - 1$ dimensional simplex. In its three dimensional case, a simplex $\Delta = \{\mathbf{p} \in \mathbb{R}^n : \sum p_i = 1\}$ can be graphically represented by an equilateral triangle with altitude 1. Each perpendicular then can be interpreted as the probability of the outcome at the opposing vertex. Thus every point in the triangle represents a lottery. This is shown in figure 5.4 for the three sure outcomes £3, £2, £1 placed at the vertices. Notice that these vertices represent degenerate lotteries (i.e. $p_i = 1, p_{j \neq i} = 0$). The Lottery $Q = \{1 : p_1, 2 : p_2, 3 : p_3\}$, with all probabilities equal to one third, is drawn below.

Def: Given a Bernoulli utility function $u(\cdot)$, the **certainty equivalent** of $F(\cdot)$, denoted $c(F, u)$ is the amount of money for which the individual is indifferent between the gamble $F(\cdot)$ and the certain amount $c(F, u)$; that is,

$$u(c(F, u)) = \int_{-\infty}^{+\infty} u(x)dF(x).$$

⁸ Notice that $U(\cdot)$ is linear in $F(\cdot)$.

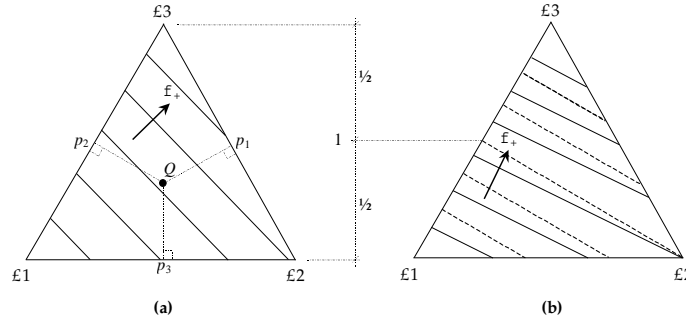


Figure 5.4: Level set of an expected utility function.

Def: Assume $u(\cdot)$ is twice differentiable and increasing. Then the Arrow-Pratt **coefficient of absolute risk aversion** $r(\cdot)$ of Bernoulli utility functions (CARA) defined over outcomes x_i is

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

Rem: Motivation for CARA. We know that linearity of $u(\cdot)$ —that is $u''(\cdot) = 0$ —implies risk neutrality. This suggests that we should look at the function’s curvature to get an idea of the risk behaviour of the agent. The following figure should make this idea clearer: at x , both $u_1(\cdot)$ and $u_2(\cdot)$ give the same utility $u(x)$, but the lotteries $\{x - \varepsilon, x + \varepsilon\}$ are evaluated differently by the two agents. Hence the shape of the utility function gives an indication of the degree of risk aversion.




Figure 5.5: Differing degrees of risk aversion.

Since the certainty equivalent $c(F, u)$ is bigger for the utility function with the higher curvature, we would expect that risk aversion increases with larger curvature. This is indeed the case. The proof, however, is left to you.

Rem: Indifference curves that show a CARA higher than unity have a steeper slope than depictions of the mathematical expectation that are shown as dashed lines in figure 5.4. Steeper indifference curves cross the (dashed) iso-expectation lines in the above figure in direction of

increasing expectation. Thus the agent has to be made indifferent for accepting higher risk by being given higher expected payoffs—she is risk averse (left of figures 5.3 & 5.4). The reverse holds if the agent accepts higher risks with lower expectations—she is risk loving (right of figures 5.3 & 5.4).

Exl: Consider the Bernoulli utility functions $u_1(\mathbf{x}) = x^{\frac{1}{2}}$ and $u_2(\mathbf{x}) = x^2$. Compare the CARA of the two utility functions. We get the following tables

agent 1	$x = 1$	$x = 2$	$x = 3$	$x = 4$
$u(x) = x^2$	1	4	9	16
$u'(x) = 2x$	2	4	6	8
$u''(x) = 2$	2	2	2	2
agent 2	$x = 1$	$x = 2$	$x = 3$	$x = 4$
$u(x) = x^{\frac{1}{2}}$	1	$\sqrt{2}$	$\sqrt{3}$	2
$u'(x) = \frac{1}{2}x^{-\frac{1}{2}}$	$\frac{1}{2}$	$\frac{1}{2\sqrt{2}}$	$\frac{1}{2\sqrt{3}}$	$\frac{1}{2\sqrt{4}}$
$u''(x) = -\frac{1}{4}x^{-\frac{3}{2}}$	$-\frac{1}{4}$	$-\frac{1}{4\sqrt{8}}$	$-\frac{1}{4\sqrt{27}}$	$-\frac{1}{4\sqrt{64}}$
agent 1	$x = 1$	$x = 2$	$x = 3$	$x = 4$
$r(x) = -\frac{u''(x)}{u'(x)}$	1	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{1}{4}$
agent 2	$x = 1$	$x = 2$	$x = 3$	$x = 4$
$r(x) = -\frac{u''(x)}{u'(x)}$	$\frac{1}{2}$	$\frac{\sqrt{2}}{2\sqrt{8}}$	$\frac{\sqrt{3}}{2\sqrt{27}}$	$\frac{\sqrt{4}}{2\sqrt{64}}$

The corresponding functions can be seen in figures 5.6 and 5.7. ◁




Figure 5.6: The derivatives and CARA of x^2 .

Def: The probability distribution of F **first-order stochastically dominates** (fosd) that of G




Figure 5.7: The derivatives and CARA of $x^{\frac{1}{2}}$.

if for non-decreasing $u(\cdot)$ ⁹

$$\int_{-\infty}^{+\infty} u(x)dF(x) \geq \int_{-\infty}^{+\infty} u(x)dG(x).$$

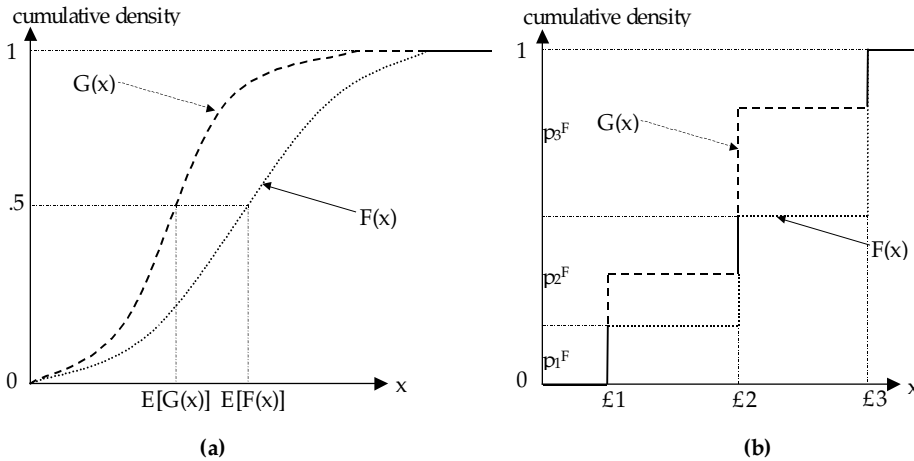


Figure 5.8: Stochastic dominance in the discrete case (left) and the continuous case (right). Distribution F first-order stochastically dominates G if distribution F is always at or below the distribution G .

Rem: Fost is a straightforward extension of the relation ‘ \succ ’ to the stochastic case: More is better. ‘ \succ ’ describes a weak order over the continuous probability distributions F and G rather than over the simple probability measures u and v . So essentially, the discrete probabilities in a lottery are replaced by a continuous distribution function. Fost implies that the mean of F is larger than that of G ; it does not imply, however, that each probability value of $F(x_i)$ is larger than that of $G(x_i)$. Actually, the discrete version of the definition of fost is $F(x_i) \leq G(x_i)$ for every x_i . Stochastic dominance of distribution F over G means that F gives unambiguously higher returns

⁹ Notice that it follows that the mean of F is bigger than that of G . We’ll skip the limits from now on if there is no risk of confusion.

than G .

Def: Let two probability distributions F, G share the same mean (i.e. $\int x dF(x) = \int x dG(x)$). Then distribution F **second-order stochastically dominates** (sosd) distribution G if for concave $u(\cdot)$

$$\int u(x) dF(x) \geq \int u(x) dG(x). \quad (5.3.2)$$

Rem: Sosd is about distributions with the same mean (expected value) but reflecting different risks such as a mean preserving spread of a given distribution. Consequently distributions exhibiting sosd have the same expectations, but the dominated distribution poses a higher risk. Hence, if F sosd's G , for all x_0 we have

$$\int_{-\infty}^{x_0} G(t) dt \geq \int_{-\infty}^{x_0} F(t) dt.$$

Rem: The idea of Sosd is captured by a **mean-preserving spread** of a given distribution $F(\cdot)$. This is best illustrated in an example.

Exl: Let $F(\cdot)$ be an even probability distribution between 2 and 3 pounds (i.e. giving probability $\frac{1}{2}$ to both outcomes). We now 'spread' the 2 dollars outcome to an even probability between 1 and 2 pounds, and the 3 pounds outcome to an even probability between 3 and 4 pounds (i.e. giving probability $\frac{1}{4}$ to all outcomes). We call the resulting distribution $G(\cdot)$. Notice that our probability manipulations do not change the mean of the distributions $E(F(\cdot)) = E(G(\cdot))$. This example is shown in figure 5.9 along another example for the continuous case. \triangleleft




Figure 5.9: A mean-preserving spread in the discrete case (left) and the continuous case (right).

Exl: Let x be uniformly distributed on the unit interval: $x \sim U_{[0,1]}$. The expectation of x on $U_{[0,1]}$ is $\frac{1}{2}$. (The general expectation of x on $U_{[a,b]}$ is $\frac{b+a}{2}$.) The probability density $f(x)$ is given by $\frac{F(1)-1}{1} = 1$. Integrating over the left half unit interval for $x \leq \frac{1}{2}$ gives $\int_0^{\frac{1}{2}} x dx = \frac{1}{8}$. \triangleleft

Exl: The expectation of the minimum draw $E[\min(X)]$ from n random draws of the random variable $A \sim U_{[a,b]}$ is $\frac{b-a}{1+n}$. What is the expectation of the 2^{nd} smallest draw? What's the relevance of this for (second price) auctions? \triangleleft

Exl: (Application of integrating by parts.) Consider the following (artificial) problem

$$\max_{p(\theta)} \int_a^b [\alpha p(\theta) + (1 - \alpha)\pi(\theta)] f(\theta) d\theta \quad (5.3.3)$$

Where $\theta \sim F_{[a,b]}$ and $\pi'(\theta)$, the first derivative of $\pi(\theta)$, is known as $\pi'(\theta) = p(\theta) - 1$. Exogenously, we find that $\pi(a) = 1$ and $\pi(b) = 0$. We know that $F(a) = 0$ and $F(b) = 1$.

We do not really know how to do this since the function $\pi(\theta)$ does not (directly) depend on the maximand $p(\theta)$. But fortunately we (should) know how to integrate by parts, which gives

$$\int_a^b f(\theta)g'(\theta) d\theta = f(\theta)g(\theta)\Big|_a^b - \int_a^b f'(\theta)g(\theta) d\theta. \quad (5.3.4)$$

Now we rewrite (5.3.3) using the following substitutions

$$\begin{aligned} f' &\equiv f(\theta); \text{ therefore } f = F(\theta); \\ g &\equiv \pi(\theta); \text{ therefore } g' = \pi'(\theta). \end{aligned}$$

Thus integrating by parts gives the result of (5.3.4) as

$$\int_a^b f(\theta)\pi(\theta) d\theta = \pi(\theta)F(\theta)\Big|_a^b - \int_a^b \pi'(\theta)F(\theta) d\theta.$$

This simplifies to

$$\begin{aligned} \int_a^b f(\theta)\pi(\theta) d\theta &= (\pi(b)F(b) - \pi(a)F(a)) - \int_a^b \pi'(\theta)F(\theta) d\theta \\ &= 0 - \int_a^b \pi'(\theta)F(\theta) d\theta \\ &= - \int_a^b (p(\theta) - 1)F(\theta) d\theta. \end{aligned}$$

This could be simplified further but without ‘real’ application, it is difficult to see why this would be helpful.¹⁰ It is, however, presently helpful to reconsider our optimisation problem (5.3.3) with our above result substituted

$$\max_{p(\theta)} \int_a^b \{\alpha p(\theta)f(\theta) - (1 - \alpha)[p(\theta) - 1]F(\theta)\} d\theta.$$

We have entirely removed the dependency on $\pi(\theta)$ and can happily maximise by taking derivatives wrt $p(\theta)$. \triangleleft

¹⁰ The model by (Baron and Myerson 1982) we looked at in exercise 3.4.6 is such a ‘real’ application.

Exercises

Exc 5.1: Prove the following assertion: If instead of the family of all intervals we take the family of all open intervals as the basis of the definition of the Borel sets (page 99), then the minimal σ -algebra \mathcal{B} remains unchanged.

Exc 5.2: Consider the problem

$$\begin{aligned} & \max_{x,y} x^2 + 2y^2 - x \\ \text{s.t.:} & \quad g(x, y) = x^2 + y^2 \leq 1. \end{aligned} \tag{5.3.5}$$

(i) Write down the Lagrangian and find the critical points.

(ii) What is the complementary slackness condition?

(iii) There are 5 candidates for (x, y) satisfying the necessary conditions. Find the solution to (5.3.5).

Exc 5.3: Solve the problem

$$\begin{aligned} & \max_{x,y} y - x^2 \\ \text{s.t.:} & \quad (i) \quad g_1(x, y) = x + e^{-x} \leq y \\ & \quad (ii) \quad g_2(x) = x \geq 0. \end{aligned}$$

Exc 5.4: Solve the problem

$$\begin{aligned} & \max_{x,y} xe^{y-x} - 2ey \\ \text{s.t.:} & \quad (i) \quad g_1(x, y) = 0 \leq y \leq 1 + \frac{x}{2} \\ & \quad (ii) \quad g_2(x) = x \geq 0. \end{aligned}$$

Exc 5.5: Consider the problem

$$\begin{aligned} & \max_{x,y} \ln x_1(x_2 + x_3) \\ \text{s.t.:} & \quad (i) \quad g_1(x) = x_1 + x_2 + x_3 \leq 1 \\ & \quad (ii) \quad g_2(x) = x_1 \geq 1 \\ & \quad (iii) \quad g_3(x) = x_1^2 + x_2^2 \leq 2. \end{aligned} \tag{5.3.6}$$

(i) Write down the necessary conditions.

(ii) Find *all* points satisfying the necessary conditions. (Hint: There are many.)

(iii) Find the solution to (5.3.6).

Exc 5.6: Consider the problem

$$\begin{aligned} & \max_x (x - 1)^3 \\ \text{s.t.:} & \quad (i) \quad g_1(x) = -x \leq 0, \\ & \quad (ii) \quad g_2(x) = x \leq 2. \end{aligned}$$

(i) Find the solution graphically and (ii) algebraically.

Exc 5.7: Consider the problem

$$\begin{aligned} & \max_{x,y} 3x + y \\ \text{s.t.:} & \quad (i) \quad g_1(x, y) = y \leq (1 - x)^3, \\ & \quad (ii) \quad g_2(x) = x \geq 0, \\ & \quad (iii) \quad g_3(y) = y \geq 0. \end{aligned}$$

(i) Find the solution graphically and show that the constraint qualification is violated.

(ii) Solve the problem algebraically using the Kuhn-Tucker approach. The unique solution is $(x, y) = (0, 1)$ with associated multipliers $\lambda = (1, 0, 0)$. Is this an optimum?

Exc 5.8: Show that *fosd* (as defined in (5.3.2)) implies $F(x) \leq G(x)$ for all discrete x (as claimed in our remark on page 106).

Exc 5.9: (Mas-Colell, Whinston, and Green 1995, 213:6.D.1) The purpose of this exercise is to prove Proposition 6.D.1. (page 195) in a two-dimensional probability simplex. This proposition reads as:

Proposition 6.D.1: The distribution of monetary payoffs $F(\cdot)$ first-order stochastically dominates the distribution $G(\cdot)$ iff $F(x) \leq G(x)$ for every x .

Suppose that there are three monetary outcomes: 1 pound, 2 pounds, and 3 pounds. Consider the probability simplex of figure 6.B.1 (b) (page 169—we have done this in ETA).

1. For a given lottery L over these outcomes, determine the region of the probability simplex in which lie the lotteries whose distributions first-order stochastically dominate the distribution of L .
2. Given a lottery L , determine the region of the probability simplex in which lie the lotteries L' such that $F(x) \leq G(x)$ for every x , where $F(\cdot)$ is the distribution of L' and $G(\cdot)$ is the distribution of L . (Hint: We get the same region.)

Exc 5.10: Optional (Mas-Colell, Whinston, and Green 1995, 213:6.D.4) The purpose of this exercise is to verify the equivalence of the three statements of Proposition 6.D.2 (page 199) in a two-dimensional probability simplex. This proposition reads as:

Proposition 6.D.2: Consider two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean. Then the following statements are equivalent:

(i) $F(\cdot)$ second-order stochastically dominates $G(\cdot)$.

(ii) $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$.

(iii) $\int_0^\infty G(t) dt \geq \int_0^\infty F(t) dt$ for all x .

Suppose there are three monetary outcomes: 1, 2, and 3 pounds. Consider the probability simplex of figure 6.B.1 (b).

1. If two lotteries have the same mean, what are their positions relative to each other in the probability simplex?
2. Given a lottery L , determine the region of the probability simplex in which lie the lotteries

- L' whose distributions are second-order stochastically dominated by the distribution of L .
3. Given a lottery L , determine the region of the probability simplex in which lie the lotteries L' whose distributions are mean preserving spreads of L .
 4. Given a lottery L , determine the region of the probability simplex in which lie the lotteries L' for which condition (iii) above holds. (Hint: 2-4 give the same region.)

Chapter 6

Machine-supported mathematics

Reading: Mathematica 4.1: (Wolfram 1999),¹ Maple 7: Springer’s maple book,² MatLab 6: Whatever MatLab book. An excellent source of applications of Mathematica to economics is (Varian 1992a).

The purpose of this chapter is to familiarise you with the computer as a tool to solve mathematical problems applied to economics (or, preferably, an MSc in Economics). While the first section is little more than an overview of selected relevant commands, the second section consists of an in-depth illustration of how to use the software to solve Lagrangian constrained optimisation problems. This includes an introduction Maple programming. While nothing in this chapter will help you at your exams, it may be expected that it is a vital supplement for ‘real-life’ quantitative analysis.

6.1 The programs

General principles

- The current version is Mathematica 4.1. There is a version for students. Mathematica’s main competitors are *Maple VII* and *MatLab*—all packages have comparable features and can be used interchangeably for most purposes. The two aforementioned competitors have built-in tutorials that you can run to accustom yourself with their use.
- All naming in the program is case specific. You have to type exactly as described.
- Equations are given in the form $lhs == rhs$.
- All function-arguments are enclosed in `[]`.
- Simultaneous equations can be combined either in a list or with `&&`.
- A single variable or a list of variables can be specified; a list is enclosed in `{}`.
- Pressing the keys *shift + enter* simultaneously evaluates the current line.
- *FullSimplify[.]* simplifies expressions.
- A blank (space) is required between symbols. A blank means multiply.

¹ A free online version is available at <http://documents.wolfram.com/v4/>.

² A 30-days trial version is available from <http://www.maplesoft.com/>.

Individual commands

- $Det[M = \{\{2, 3, 4, 6\}, \{7, 8, 5, 2\}, \{2, 3, 1, 1\}, \{3, 5, 4, -1\}\}]$ gives the determinant of the matrix

$$\begin{bmatrix} 2 & 3 & 4 & 6 \\ 7 & 8 & 5 & 2 \\ 2 & 3 & 1 & 1 \\ 3 & 5 & 4 & -1 \end{bmatrix}.$$

- $Series[eqn, y, x]$ solves a differential equation for the function y , with independent variable x .

Solving general equations

- $Solve[eqns, vars]$ attempts to solve an equation or set of equations for the variables $vars$.
- $Solve[eqns, vars, elims]$ attempts to solve the equations for $vars$, eliminating the variables $elims$.
- $Solve[eqns]$ tries to solve for all variables in $eqns$.
- $Solve$ gives solutions in terms of rules of the form $x \mapsto sol$. When there are several solutions, $Solve$ gives a list of them.
- When there are several variables, the solution is given in terms of lists of rules: $ax \mapsto a, y \mapsto a, a$.
- When a particular root has multiplicity greater than one, $Solve$ gives several copies of the corresponding solution.
- $Solve$ gives generic solutions only. It discards solutions that are valid only when the parameters satisfy special conditions. $Reduce$ gives the complete set of solutions.
- $Solve$ will not always be able to get explicit solutions to equations. It will give the explicit solutions it can, then give a symbolic representation of the remaining solutions in terms of Root objects. If there are sufficiently few symbolic parameters, you can then use N to get numerical approximations to the solutions.
- $Solve$ gives $\{\}$ if there are no possible solutions to the equations.
- $Solve[eqns, , Mod \mapsto Modular]$ solves equations with equality required only modulo an integer. You can specify a particular modulus to use by including the equation $Modulus == p$. If you do not include such an equation, $Solve$ will attempt to solve for the possible moduli.

Exl: $Solve[3x+9==0, x]$ produces $\{\{x \mapsto -3\}\}$. ◁

Exl: $Solve[x^2 + 2bx + c == 0, x]$ produces $\{\{x \mapsto -b - \sqrt{b^2 - c}\}, \{x \mapsto -b + \sqrt{b^2 - c}\}\}$. ◁

See (Wolfram 1999, Sections 1.5.7, 3.4.4).

Functions and their plots

1. $Plot[f, \{x, xmin, xmax\}]$ generates a plot of f as a function of x from $xmin$ to $xmax$.
 - $Plot[\{f_1, f_2, \dots\}, \{x, xmin, xmax\}]$ plots several functions f_i .
 - Plot evaluates its arguments in a non-standard way (see (Wolfram 1999, Section A.4.2)). You should use $Evaluate$ to evaluate the function to be plotted if this can safely be done before specific numerical values are supplied.

- *Plot* has the same options as *Graphics*.
- *Plot* uses the default setting *Axes* \rightarrow *True*.
- *Plot* initially evaluates f at a number of equally spaced sample points specified by *PlotPoints*. Then it uses an adaptive algorithm to choose additional sample points, attempting to produce a curve in which the bend between successive segments is less than *MaxBend*. It subdivides a given interval by a factor of at most *PlotDivision*.
- You should realise that with the finite number of sample points used, it is possible for *Plot* to miss features in your function. To check your results, you should increase the setting for *PlotPoints*.
- *Plot* returns a *Graphics* object.

Exl: `Plot[{Sin[x], Sin[2x], Sin[3x]}, {x, 0, 2 Pi}]`; produces \triangleleft




Figure 6.1: Output of a `Plot[.]` example.

Exl: `Plot[{x^2, 2x, 2}, {x, 0, 4}]`; produces the output of (part of) figure 5.6. \triangleleft




Figure 6.2: Output of `Plot[{x^2, 2x, 2}, {x, 0, 4}]`.

See (Wolfram 1999, Section 1.9.1).

2. `ContourPlot[f[x], {x, xmin, xmax}]` generates a plot of the level sets of f in the stated range of x from $xmin$ to $xmax$.
 - `ContourPlot[f[x, y], {x, xmin, xmax}, {y, ymin, ymax}]` generates a plot of the level sets of f in the stated range of x and y .

Exl: `ContourPlot[x^3 + y^3, {x, -5, 5}, {y, -5, 5}]` produces \triangleleft




Figure 6.3: The level sets of $f(x, y) = x^3 + y^3$.

Linear programming

- `LinearProgramming[]`
- `ConstrainedMax[]`, `ConstrainedMin[]` are implementations of the `LinearProgramming[]` function.

Exl: `ConstrainedMax[x + y, {x < 1, y < 2}, {x, y}]` gives `{3, {x - > 1, y - > 2}}`. ◁

See (Wolfram 1999, Section 3.9.9).

Derivatives

- $D[f, x]$ gives the partial derivative $\frac{\partial f}{\partial x}$.
- $D[f, \{x, n\}]$ gives the multiple derivative $\frac{\partial^n f}{\partial x^n}$.
- $D[f, x_1, x_2, \dots]$ gives $\frac{\partial f}{\partial x_1} \frac{\partial f}{\partial x_2} \dots$
- $D[f, x]$ can be input as $\partial_x f$. The character ∂ can be substituted by `PartialD[]`. The variable x is entered as a subscript.
- All quantities that do not explicitly depend on the x_i are taken to have zero partial derivative.
- $D[f, x_1, \dots, Non - Constants \rightarrow \{v_1, \dots\}]$ specifies that the v_i implicitly depend on the x_i , so that they do not have zero partial derivatives.
- The derivatives of built-in mathematical functions are evaluated when possible in terms of other built-in mathematical functions.
- Numerical approximations to derivatives can be found using `N`.
- D uses the chain rule to simplify derivatives of unknown functions.
- $D[f, x, y]$ can be input as `a`. The character $\partial_{x,y} f$.
- Total derivatives can be obtained with `Dt`.

Exl: $\partial_x \text{Sin}[x]$ produces `Cos[x]`. ◁

Exl: The chain rule looks like this: $\partial_x f[g[x]]$ and produces $f'[g[x]]g'[x]$. ◁

See (Wolfram 1999, Sections 1.5.2, 3.5.1).

Integrals

- *Integrate*[f, x] gives the indefinite integral $\int f(x) dx$.
- *Integrate*[$f, \{x, xmin, xmax\}$] gives the definite integral $\int_{x_{\min}}^{x_{\max}} f(x) dx$.
- *Integrate*[$f, \{x, xmin, xmax\}, \{y, ymin, ymax\}$] gives the multiple integral $\int_{x_{\min}}^{x_{\max}} dx \int_{y_{\min}}^{y_{\max}} f(x, y) dy$.
- *Integrate*[f, x] can be entered as $\int f dx$.
- Multiple integrals use a variant of the standard iterator notation. The first variable given corresponds to the outermost integral, and is done last.
- *Integrate* can evaluate integrals of rational functions. It can also evaluate integrals that involve exponential, logarithmic, trigonometric and inverse trigonometric functions, so long as the result comes out in terms of the same set of functions.
- *Integrate* can give results in terms of many special functions.
- *Integrate* carries out some simplification on integrals it cannot explicitly do.
- You can get a numerical result by applying N to a definite integral.
- You can assign values to patterns involving *Integrate* to give results for new classes of integrals.
- The integration variable can be any expression. However, *Integrate* uses only its literal form. The object $d(x^n)$, for example, is not converted to $nx^{n-1}dx$.
- For indefinite integrals, *Integrate* tries to find results that are correct for almost all values of parameters.
- For definite integrals, the following options can be given:
 - *PrincipalValue* \rightarrow *False* ... whether to find Cauchy principal values
 - *GenerateConditions* \rightarrow *True* ... whether to generate answers that involve conditions on parameters
 - *Assumptions* \rightarrow {...} ... assumptions to make about relations between parameters
- *Integrate* can evaluate essentially all indefinite integrals and most definite integrals listed in standard books of tables.
- In *StandardForm*, *Integrate*[f, x] is output as a $\int f dx$.

Exl: $\int x^n dx$ produces $\frac{x^{1+n}}{1+n}$. \triangleleft

See (Wolfram 1999, Sections 1.4.4, 1.5.3, 3.5.6).

Differential equations

- *DSolve*[eqn, y, x] solves a differential equation for the function y , with independent variable x .
- *DSolve*[{ eqn_1, eqn_2, \dots }, { y_1, y_2, \dots }, x] solves a list of differential equations.
- *DSolve*[$eqn, y, \{x_1, x_2, \dots\}$] solves a partial differential equation.
- *DSolve*[$eqn, y[x], x$] gives solutions for $y[x]$ rather than for the function y itself.
- Differential equations must be stated in terms of derivatives such as $y'[x]$, obtained with D , not total derivatives obtained with Dt .

- *DSolve* generates constants of integration indexed by successive integers. The option *DSolveConstants* specifies the function to apply to each index. The default is *DSolveConstants*→*C*, which yields constants of integration $C[1], C[2], \dots$
- *DSolveConstants*→(*Module*{*C*}, *C*)& guarantees that the constants of integration are unique, even across different invocations of *DSolve*.
- For partial differential equations, *DSolve* generates arbitrary functions $C[n][\dots]$.
- Boundary conditions can be specified by giving equations such as $y'[0] == b$.
- Solutions given by *DSolve* sometimes include integrals that cannot be carried out explicitly by *Integrate*. Dummy variables with local names are used in such integrals.
- *DSolve* sometimes gives implicit solutions in terms of *Solve*.
- *DSolve* can solve linear ordinary differential equations of any order with constant coefficients. It can solve also many linear equations up to second order with non-constant coefficients.
- *DSolve* can find general solutions for linear and weakly non-linear partial differential equations. Truly non-linear partial differential equations usually admit no general solutions.

Exl: *DSolve*[$y'[x] == 2 a x, y[x], x$] produces $\{\{y[x] \mapsto a x^2 + C[1]\}\}$. ◁

Exl: *DSolve*[$y'[x] == a y[x], y[x], x$] produces $\{\{y[x] \mapsto e^{ax} C[1]\}\}$. ◁

See (Wolfram 1999, Sections 1.5.8, 3.5.10).

6.2 Constrained optimisation problems

The first example of a linear program on page 12 is very easy, too.

Let us start with an example.

Exl: Find the extrema of $f(x, y) = x^2 + y^2$ subject to the constraint $x^4 + y^4 = 1$ equivalent to $g(x, y) = x^4 + y^4 - 1$. We solve the equations setting

$$\nabla f = \lambda \nabla g, \quad \text{or} \quad (2x, 2y) = \lambda(4x^3, 4y^3)$$

subject to $x^4 + y^4 = 1$. That is, we must solve a system of 3 nonlinear equations in 3 unknowns:

$$2x = 4\lambda x^3, \quad 2y = 4\lambda y^3, \quad x^4 + y^4 = 1.$$

In Maple we go about doing this like this:

```
> restart;
> f := (x,y) -> x^2+y^2;
> g := (x,y) -> x^4+y^4-1;
> eq1 := diff(f(x,y),x)=lambda*diff(g(x,y),x);
> eq2 := diff(f(x,y),y)=lambda*diff(g(x,y),y);
> eq3 := g(x,y)=0;
> solution := solve({eq1,eq2,eq3},{x,y,lambda});
```

```
> f(1,0);  
> f(-1,0);  
> allvalues(solution[3]);
```

Discussion: The first few lines of code define f and g as well as the equations to be solved. In this example we are using only one constraint for demonstration purposes, but our code would be easily modified to using two or more constraints. Then we ask Maple to solve our system; this is where the magic occurs. While this problem might have looked innocuous at the beginning, it has 8 solutions—in typing $f(1,0)$; and $f(-1,0)$; we are checking the values of f at the first two solutions to determine what the extrema of f are. When we get to the third solution we find that unknowns are expressed as solutions to equations; by typing `allvalues(solution[3]);` we express all the unknowns of the third solution numerically so we can evaluate f at the third solution. Since we find that one of the unknowns is imaginary (Maple uses I to denote $-1^{\frac{1}{2}}$), we do not need to evaluate f at the third solution. The 5 remaining solutions are found accordingly.

◁

Exercises

Solve some of your previous assignments for yourself in Maple.

Chapter 7

Fixed points

Reference: I found (Franklin 1980) to be a very readable and insightful source of ideas particularly on fixed points. On a more fundamental level, (Munkres 2000) is an excellent and complete treatment of general topology and fixed point theory. There are a few pages in the mathematical appendix of (Mas-Colell, Whinston, and Green 1995) that give a very brief summary of the most commonly used fixed point theorems. A most valuable and concise introduction to fixed point theory from an economists' point of view is Border (1985). Ferreira (1996) is a short but fruitful introduction on the subject of fixed points.

7.1 Motivation

7.1.1 Some topological ideas

Informally speaking, topology is the study of those properties that an object retains under deformation—specifically, bending, stretching and squeezing, but not breaking or tearing. Thus, a triangle is topologically equivalent to a circle but not to a straight line segment because as soon as you have collapsed the triangle into a line, you cannot extract the triangle back without using a non-bijective mapping.¹ Similarly, a solid cube made of modeling clay could be deformed into a ball by kneading. It could not, however, be molded into a solid torus (ring) unless a hole were bored through it or two surfaces were joined together. A solid cube is therefore not topologically equivalent to a ring.

More precisely, if there are given two geometric objects or sets of points (and a metric) and if some two-way transformation (or mapping) takes each point p of either set into one and only one point p' of the other and if the transformation is continuous in the sense (that is made more precise in Definition (7.1.2)) that points close to p become points close to p' and vice versa, then the transformation is called a **homeomorphism** and the two sets are said to be **topologically equivalent**. Topology is, then, the study of properties that remain invariant under homeomorphisms. Below is a simple example of such a homeomorphism $h(\cdot)$ of the straight line x into the curve y . Before we go on let us define this in a less chatty way.

¹ Let us refer to *both* functions and correspondences as mappings.




Figure 7.1: A homeomorphism $h(\cdot)$ of the straight line x into a curved set y .

Def: Let (X, d_X) and (Y, d_Y) be two metric spaces and let $f : X \mapsto Y$ be a mapping of the first into the second space so that $y = f(x) \in Y$ for each $x \in X$. The function $f(\cdot)$ is said to be **continuous** at x_0 if, given any $\varepsilon > 0$, there exists a $\delta > 0$ such that $d_Y(f(x), f(x_0)) < \varepsilon$ whenever $d_X(x, x_0) < \delta$. $f(\cdot)$ is said to be continuous if it is continuous at every $x \in X$.

Def: Given the above two metric spaces (X, d_X) and (Y, d_Y) , let $f(\cdot)$ be a one-to-one mapping of X onto Y and suppose that both f and its inverse $f^{-1} : Y \mapsto X$ exist and are continuous. Then f is called a **homeomorphism** (between X and Y).

Exl: The function

$$y = f(x) = \frac{1}{\pi} \arctan x + \frac{1}{2}$$

establishes a homeomorphism between all of \mathbb{R} , i.e. $(-\infty, \infty)$, and $(-1, 1)$. \triangleleft

You may agree that it is an interesting question to ask whether there are points in the transformed objects or sets that are invariant (i.e. do not move) under homeomorphisms. We call such points fixed points. Spaces like x and y of figure 7.1 are maybe not terribly interesting. But fortunately, our theorems about fixed points generalise to more important applications. It is true, for instance, that at any instant there is always at least one place on the Earth's surface where the wind is not blowing. This fact is a direct consequence of the fixed-point theorems that we will look at in the following section, which state that every homeomorphism of a round spherical surface onto itself either leaves at least one point fixed or sends some point into its diametric opposite.

Def: Let $A \subset \mathbb{R}^n$ and $f : A \mapsto A$ be a mapping of the A into itself. Then a vector $x \in A$ is called a **fixed point** of $f(\cdot)$ if $x = f(x)$ (if $f(\cdot)$ is a function), or $x \in f(x)$ (if $f(\cdot)$ is a correspondence).

Before we talk about fixed point existence theorems, we should pause for a minute and contemplate what we are doing. Look at a fixed point equation in the definition above: The result of the mapping is the original point—in other words, the mapping leaves the original point unchanged. For a mapping $f(\cdot)$ of the unit disk onto itself, this looks like in figure 7.2 below. Now suppose the mapping leaves (at least) one point in the same place: then we have $x = y = f(x)$, a fixed point. You see: *every* equation can be a fixed point equation. If this does not convince you,




Figure 7.2: Two mappings of the unit disk; the right panel shows a fixed point.

maybe this will: consider the equation $g(x) = 0$. We can turn this equation into a fixed point equation by simply writing

$$x = x + g(x)$$

or, if you prefer

$$x = x - 79g(x).$$

Since we can do this to any equation, our discussion is extremely general: we are indeed talking about all possible equations. The drawback in our discussion is that none of the results we will look at below is constructive—an exception is the *Contraction Mapping Theorem* which we will look at now.²

Def: Let (X, d) be a metric space. A function $f : X \mapsto X$ is called a **contraction** if there is some $0 \leq \theta < 1$ such that for every pair $x, y \in X$

$$d[f(x), f(y)] \leq \theta d(x, y). \quad (7.1.1)$$

Rem: Every contraction mapping is continuous because (7.1.1) says that for every sequence $\{x_n\}_{n=1}^{\infty}$, $f(x_n) \mapsto f(x)$ whenever $x_n \mapsto x$.

Theorem 38. (*Contraction Mapping Theorem (Banach)*) Every contraction mapping f defined on a complete metric space X has a unique fixed point.

Proof. Pick any point $x_0 \in X$ and construct the Cauchy sequence³

$$x_1 = f(x_0); x_2 = f(x_1) = f(f(x_0)), \dots; x_n = f(x_{n-1}) = f^n(x_0).$$

² If you forgot all the stuff about metric spaces, Cauchy sequences and completeness it is a good idea to look these up (section 2) once more now.

³ We use the notation $f^n(x) = f(f(\dots(x)\dots))$ to denote a function applied n times to its argument. $d(x, y)$ is the distance between x, y .

By the convergence criterion of Cauchy sequences we have

$$\begin{aligned} d(x_n, x_{n'}) &= d(f^n(x_0), f^{n'}(x_0)) \leq \theta^n d(x_0, x_{n-n'}) \\ &\leq \theta^n [d(x_0, x_1) + d(x_1, x_2) + \cdots + d(x_{n'-n-1}, x_{n'-n})] \\ &\leq \theta^n d(x_0, x_1) [1 + \theta + \theta^2 + \cdots + \theta^{n'-n-1}] \\ &< \theta^n d(x_0, x_1) \frac{1}{1-\theta}. \end{aligned}$$

This can be made arbitrarily small for sufficiently large n , since $\theta < 1$. Since X is complete and the sequence $\{x_n\}$ is a Cauchy sequence, we know that it has a limit

$$x = \lim_{n \rightarrow \infty} x_n.$$

Then, by continuity of $f(\cdot)$,

$$f(x) = f(\lim_{n \rightarrow \infty} x_n) = \lim_{n \rightarrow \infty} f(x_n) = \lim_{n \rightarrow \infty} x_{n+1} = x.$$

Hence the existence of the fixed point $f(x) = x$ is ensured. Concerning the uniqueness of x , notice that since both $f(x) = x$, and $f(y) = y$, (7.1.1) becomes

$$d(x, y) = \theta d(x, y).$$

Since $\theta < 1$, $d(x, y) = 0$ and hence $x = y$. □

Exl: Consider a function $f : X \mapsto X$, $X = [a, b]$ satisfying a ‘Lipschitz’ condition⁴

$$|f(x_1) - f(x_2)| \leq \theta |x_1 - x_2|, \quad (7.1.2)$$

with constant $\theta < 1$. Then $f(\cdot)$ is a contraction mapping and, by the above theorem,

$$x_0, \quad x_1 = f(x_0), \quad x_2 = f(x_1), \quad \dots$$

converges to a unique fixed point which is the root of the equation $f(x) = x$. In particular, the contraction condition (7.1.2) holds if $f(\cdot)$ is C^1 on X such that

$$f'(x) \leq \theta < 1.$$

This is illustrated in figure 7.3 for the cases of a positively and negatively sloped function $f(\cdot)$. The alert reader will undoubtedly recognise the connection to the heroic but futile attempts of the two hedgehogs on the cover page. ◁

Rem: The above theorem can be used to prove existence and uniqueness of the solution to a system of linear equations $Ax = x$. Additional applications, many of which are very interesting, are obtained if the underlying metric space is a function space: then existence and uniqueness

⁴ If you are curious about this condition take a look at the definition of Lipschitz continuity, p.41.




Figure 7.3: ‘Successive approximation’ for the cases of $0 < f'(x) < 1$ (left) and $-1 < f'(x) < 0$ (right).

theorems for differential and integral equations such as the Picard-Lindelöf Theorem (which we will meet in a later chapter) can be proved using this technique.

Contraction mappings are well behaved and typically easy to study. Unfortunately, (7.1.1) is often too strong a requirement to be applicable. Therefore we take the chance of introducing an alternative concept which is less stringent in its requirements as a motivation for the following discussion of fixed point theorems.

Def: Let X be a metric space. A mapping $T : X \mapsto X$ is **nonexpansive** if, for every $x, y \in X$,

$$d(Tx, Ty) \leq d(x, y).$$

If X is a normed space, then T is nonexpansive if

$$\|Tx - Ty\| \leq \|x - y\|.$$

If T is linear, this reduces to

$$\|Tx\| \leq \|x\|$$

for all $x \in X$. Thus, a linear operator $T : X \mapsto X$ is nonexpansive if its norm satisfies $\|T\| \leq 1$.

Since any of the above properties of T implies continuity of T , one could be easily led to believe that any nonexpansive mapping $T : X \mapsto X$ has a fixed point. But unfortunately this is not the case. For example, the ‘translation’ $f : f + g$, with $g \neq 0$ in a Banach space is nonexpansive and does not have any fixed points. If $g = 0$, however, this changes dramatically: the identity mapping is nonexpansive and each point in its domain is a fixed point.

This raises the problem of defining and characterising classes of nonexpansive mappings having fixed points. There is also the related and important uniqueness problem. We will restrict attention to three results which are frequently employed in economic applications. Before we look at these theorems, however, we first introduce a number of ideas that help to generalise our approach.

7.1.2 Some more details[†]

The following two subsections are optional. They can be skipped without losing much in terms of understanding the fixed point results in the next section. If you want to get a feeling for the generality with which these theorems apply, however, you should work through the present subsections as well.

Def: A **topological space** is a pair (X, τ) where X is a set and τ is a collection of subsets of X satisfying:

1. both X and \emptyset belong to τ , and
2. the intersection of a finite number and the union of an arbitrary number of sets in τ belong to τ .

Rem: The sets in τ are called **open sets** and τ is called the **topology**.

Exl: Points in the n -dimensional Euclidean space are n -tuples of numbers (x_1, x_2, \dots, x_n) . Basis elements are the interiors of balls in which the distance between $\mathbf{p} = (x_1, x_2, \dots, x_n)$ and $\mathbf{q} = (y_1, y_2, \dots, y_n)$ is the square root of $d(\mathbf{p}, \mathbf{q}) = (y_1 - x_1)^2 + (y_2 - x_2)^2 + \dots + (y_n - x_n)^2$. \triangleleft

Rem: For a collection of objects that is to be regarded as a topological space with its elements as points, it is sufficient to describe the subsets of the collection that are to be called open sets. One possibility is to define all subsets as open, even the one-point sets. The resulting topology is called the **discrete topology**.

Def: Let (X, τ) and (X', τ') be topological spaces. Let $f : X \mapsto Y$ be a bijection. If both the function f and the inverse function $f^{-1} : Y \mapsto X$ exist and are continuous, then f is called a **homeomorphism**.

Def: Let (X, τ) and (X', τ') be topological spaces. The function $f : X \mapsto X$ is said to be **continuous** if for all sets $A \in \tau'$ the set $f^{-1}(A) = \{x \in X \mid f(x) \in A\} \in \tau$. We express this by saying the inverse of open sets are open.

Def: A collection of open sets $\{O_\alpha\}_{\alpha \in A}$ is called an **open cover** of the set M if⁵

$$M \subset \bigcup_{\alpha \in A} O_\alpha.$$

Theorem 39. (Heine-Borel) A set M is **compact** iff every open cover of M has a finite subcover, i.e., a finite subcollection of $\{O_\alpha\}_{\alpha \in A}$, say $\{O_1, \dots, O_k\}$, such that $M \subset \bigcup_{i=1}^k O_i$.

Proof. See (Weibull 2001, 49). □

Exl: Let M be the closed and unbounded interval $[0, \infty)$. Let $k \in (A = \mathbb{N})$ and define $O_k = (k - 2, k + 2)$. Then $\{O_k\}_{k \in \mathbb{N}}$ is an open cover of M but there exists no finite subcover. \triangleleft

Def: A topological space (X, τ) is called a **Hausdorff space** if for any two points $\mathbf{x}, \mathbf{y} \in X$ with $\mathbf{x} \neq \mathbf{y}$ there are open neighbourhoods of \mathbf{x} and \mathbf{y} say $U(x)$ and $U(y)$ such that $U(x) \cap U(y) = \emptyset$.

⁵ This is actually an alternative definition of an open set. Its advantage is that it readily extends to the infinite case.

Def: The **open ε -ball** about the point \mathbf{x} in the metric space (X, d) denoted $B_\varepsilon(\mathbf{x})$ is the set $\{\mathbf{y} \in X \mid d(\mathbf{x}, \mathbf{y}) < \varepsilon\}$. A set is **open** if for any $\mathbf{x} \in X$ there is some $\varepsilon > 0$ such that $B_\varepsilon(\mathbf{x}) \subset X$. A set is **closed** if its complement is open.

Rem: A sequence of points $(\mathbf{x}_i) \in X$ is said to converge to $\mathbf{x} \in X$ if for any $\varepsilon > 0$ there is some integer N such that for all $n > N$, $x_n \in B_\varepsilon(\mathbf{x})$. \mathbf{x} is then called the limit of the sequence and the sequence itself is called convergent. This allows an alternative definition of a closed set: A set is **closed** if every convergent sequence in the set converges to a point in the set.

Def: A sequence (\mathbf{x}_i) in a metric space (X, d) is called **Cauchy sequence** if

$$\lim_{p, q \rightarrow \infty} d(x_p, x_q) = 0.$$

Def: A metric space (X, d) is called **complete metric space** if every Cauchy sequence in (X, d) converges to a limit in X .

7.1.3 Some important definitions

We start by weakening the concept of continuity for functions, then define set-valued functions (or correspondences) and develop similar ideas there. Then we shall have all the tools in place which are necessary for the discussion of fixed point existence theorems.

Def: A real function $f : X \subset \mathbb{R}^n \mapsto \mathbb{R}$ is called **upper semi-continuous** at $x_0 \in X$ if, given any $\varepsilon > 0$, there is a $x \in N_\varepsilon(x_0)$ such that $f(x) < f(x_0) + \varepsilon$. Similarly, f is called **lower semi-continuous (lsc)** at $x_0 \in X$ if, given any $\varepsilon > 0$, there is a $x \in N_\varepsilon(x_0)$ such that $f(x) > f(x_0) - \varepsilon$. $f(\cdot)$ is called usc (lsc) if it is usc (lsc) everywhere.

Exl: Take a continuous function f and ‘replace’ the value the function takes at x_0 with a value larger than the original $f(x_0)$. The resulting function is upper semi-continuous. Conversely, if we decrease the original $f(x_0)$, the function becomes lower semi-continuous. Figure 7.4 shows such a function that is usc at x_0 and lsc at x_1 . (That is, f is neither usc nor lsc.) \triangleleft




Figure 7.4: A function f that is usc at x_0 and lsc at x_1 .

Def: Given $X \subset \mathbb{R}^n$, a **correspondence** F between the sets X and $Y \subset \mathbb{R}^k$, written $F : X \rightrightarrows Y$, is a mapping that assigns a set $F(x)$ to every $x \in X$.

Def: The **graph** of a correspondence $F : X \rightrightarrows Y$ is the set $\Gamma(F) = \{(x, y) \in X \times Y \mid y \in F(x)\}$.

Rem: The graph $\Gamma(F)$ is a set which can be checked for open or closedness as any other set. The same is true for the graph $\Gamma(F(x_0))$ at a certain x_0 . (Formally $\Gamma(F(x_0)) = \{y \in Y \mid y \in F(x_0)\}$.)

Def: Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^k$ be closed. Then the correspondence $F : X \rightrightarrows Y$ is **upper hemicontinuous** (uhc) if it has a closed graph and the images of compact sets are bounded, that is, for every compact set $B \subset X$, the set $F(B) = \{y \in Y : y \in F(x) \text{ for some } x \in B\}$ is bounded.

Since this is the core definition of this section, we give an alternative definition.

Def: The correspondence $F : X \rightrightarrows Y$ is called **uhc** at x_0 , if for any open set $V \supset F(x_0)$, there exists a neighbourhood $N_\varepsilon(x_0)$ such that $F(x) \subset V$ if $x \in N_\varepsilon(x_0)$ for $\varepsilon > 0$. $F(\cdot)$ is called uhc if it is uhc for all $x_0 \in X$.

Def: Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^k$ be compact. Then the correspondence $F : X \rightrightarrows Y$ is **lower hemicontinuous** (lhc) if, for any $(x, y) \in \Gamma(F)$ and any sequence $\{x_m\} \in X$ such that $\{x_m\} \rightarrow x$, we can find a sequence $\{y_m\} \in Y$ such that $\{y_m\} \rightarrow y$ and $y_m \in F(x_m)$ for each x_m .

Exl: Consider the example in the left panel of figure 7.5 for a correspondence $F : X \rightrightarrows Y$. Roughly speaking, uhc is only compatible with discontinuities that make the image set $F(x)$ ‘implode’ as to the right of x_0 (e.g. at $(x, F(x))$) in the right panel, while it does allow for it to ‘explode’. Notice that to the right of x_1 the closed graph requirement fails while the correspondence is still uhc (the closed graph requirement is stronger than uhc).

Let us now check for the complementary idea of lhc at x_0 in the right hand panel of figure 7.5. As one would expect, lhc is only compatible with ‘explosions’ of the image set but disallows for ‘implosions.’ So let us check for the point (x_0, y_0) , $y_0 \in F(x_0)$. We construct a sequence $\{x_m\} \rightarrow x_0$ and try to find a sequence of points $\{y_m\}$ in Y for which the $y_m \in \Gamma(F(x_m))$. We cannot find any: all sequences we can find converge to y' but none to y_0 . Hence F is not lhc at x_0 . The middle panel shows a correspondence which is neither uhc nor lhc (and is thus discontinuous) at x_0 . ◁




Figure 7.5: A correspondence F that is lhc but not uhc (left), one that is uhc but not lhc (right) and one that is neither (in between).

Rem: Notice that the definition of uhc and lhc reduce to the usual definition of continuity for single-valued correspondences (i.e. a single-valued uhc/lhc correspondence is continuous).

Exl: Figure 7.6 shows a correspondence $F : X \rightrightarrows \mathbb{R}$ which is unbounded at x_0 . Therefore the graph $\Gamma(F(x_0))$ is not closed in \mathbb{R} at x_0 and the correspondence is not uhc. It is not uhc at x_2

either because, again, $\Gamma(F(x_2))$ fails to be closed. $F(\cdot)$ is not lhc at x_0 because the limit of the sequence $\{y_m\} \rightarrow \infty$ is not in \mathbb{R} . It is not lhc at x_3 for the same reason as above—no interior point in $F(x_3)$ can be made the limit of a sequence $\{y_m\}$. $F(\cdot)$ is not lhc at x_2 because the lower branch's limit does not exist for any sequence $\{y_m\}$. It is lhc at x_1 , though, because both limits exist in the graph $\Gamma(F(x_1))$. \triangleleft




Figure 7.6: A correspondence that is neither uhc nor lhc at x_0 (where $F(x_0)$ goes off to ∞), lhc but not uhc at x_1 , neither uhc nor lhc at x_2 , and uhc but not lhc at x_3 .

Def: An n -dimensional **simplex** $\langle v_0, v_1, \dots, v_n \rangle$ is defined to be the set

$$\Delta_n = \left\{ \mathbf{x} \in \mathbb{R}^{n+1} \mid x = \sum_{j=0}^n \theta_j v_j, \sum_{j=0}^n \theta_j = 1, \theta_j \geq 0 \right\}.$$

7.2 Existence Theorems

The theorems we will discuss in this section are the ones you will encounter in the classical proofs of both the existence of Arrow-Debreu and Nash equilibrium. This does not mean that there are no others. There are, but that is an entirely different story ...

7.2.1 Brouwer's Theorem

Brouwer's Theorem says that if a ball (or its topological equivalent) is mapped continuously onto itself, then at least one point must be mapped onto itself. More formally

Theorem 40. (Brouwer 1910) *Let K be a non-empty, compact and convex set in \mathbb{R}^n and $f : K \mapsto K$ a continuous function. Then f has a fixed point $\mathbf{x}^* \in K$, i.e. a point \mathbf{x}^* such that $f(\mathbf{x}^*) = \mathbf{x}^*$.*

Proof. The involved apparatus to show the general case is too massive to be introduced here. If you are interested, please consult a good book on topology or fixed-point theory as e.g. (Hurewicz and Wallman 1948). Alternative and more approachable proofs of Brouwer's Theorem can be found in (Binmore 1992, 223) and (Munkres 2000, 351). \square

Discussion For the case of $n=1$, it is easy to see why the theorem holds. We look at the unit interval $x \in [0, 1] \equiv K$ mapped into itself $f(x) \in [0, 1] = K$. Brouwer's Theorem says that we

will get at least one fixed point with $f(x) - x = 0$. We know that $f(x) - x$ is continuous because $f(x)$ is continuous. By definition of K we have both

$$0 \leq f(x) \leq 1 \text{ and } 0 \leq x \leq 1.$$

Therefore, we must have $f(x) - x \geq 0$ for $x = 0$ and $f(x) - x \leq 0$ for $x = 1$. Since the fixed point function's graph starts out above the abscissa and ends up below the abscissa and is continuous, we know that there exists a point where it intersects the abscissa. At such a point we have a fixed point

$$f(x) - x = 0.$$

If we only graph $f(x)$, we obtain figure 7.7.




Figure 7.7: Indispensability of the continuity assumption in Brouwer's fixed point theorem.

We can also capture the intuition for the case of $n=2$ graphically. Let's define K to be the unit disk $K \equiv |\mathbf{x}| \leq 1$. Then we get a mapping like in figure 7.2. Imagine that the function $f(x)$ continuously maps x into $f(x)$ by a rotation of, say, 45° . Since the domain and the range must be the same unit disk, at least one point must remain fixed—in the case of our example it is the centre.

Another intuitive argument given by (Luce and Raiffa 1957, 392) goes as follows: Again let K be the unit ball (i.e. the unit disk in the plane) and let S be a sphere (i.e. a circle in the plane); further let $f : K \mapsto K$ be a mapping that takes $x \in K$ into $f(x) \in K$. If f had no fixed point, i.e., if the image of every point were to be distinct from the point itself, then we could perform the following trick. For each x , let $g(x)$ denote the point where the ray beginning at $f(x)$ and passing through x intersects the boundary of the disk, i.e. the circle S ; hence $g : K \mapsto S$. Let's leave point of K already on the boundary fixed, i.e. for $x \in S$ let $g(x) = x$. Since $f(x) \neq x$ and since f is continuous, it follows that g is also continuous. But a function which maps the whole disk onto its boundary (keeping boundary points fixed) necessitates 'ripping' the interior of the disk, i.e. there must be points 'close' together in the interior of S which under the mapping are shoved 'far apart.' Consider, for example, the two points x, y close together at the centre of the disk shown in figure 7.8. Thus g is *not* continuous, contrary to what we have shown above. The

assumption that got us into this contradiction was that f had no fixed point, so we must conclude that $f(x) = x$ for some x .




Figure 7.8: A trick that does not work: $g : K \mapsto S$ cannot be continuous if it does not have a fixed point.

Rem: The case of $n=1$ is really Bolzano's **Intermediate Value Theorem** that you probably know from school. It says that a continuous function $f : [a, b] \mapsto \mathbb{R}$ takes all intermediate values between the extreme ones min, max, i.e. for each $y \in [\min, \max]$ there is a $x \in [a, b]$ such that $y = f(x)$.

7.2.2 Kakutani's Theorem

Theorem 41. (Kakutani) A correspondence $F : X \rightrightarrows X$ satisfying the following assumptions:

1. X is a convex, closed, and bounded subset of \mathbb{R}^n ,
2. $F(\cdot)$ is non-empty and convex-valued, and
3. $F(\cdot)$ has a closed graph

has a fixed point $\mathbf{x}^* \in F(\mathbf{x}^*)$.

Proof. See (Hurewicz and Wallman 1948). □

Discussion We only represent the intuition for the $n=1$ -dimensional case. All elements that we identified as crucial in the case of Brouwer's Theorem are still in place. In addition, we require a convex-valued graph. This is illustrated in figure 7.9.

The (somewhat minimal) property of the correspondence that does not allow it to vanish as it approaches the diagonal is the uhc property which plays the same role as the full continuity assumption of the function in Brouwer's Theorem. Notice that uhc alone is insufficient because the correspondence could 'jump' over the diagonal (both up and down) while still possessing a closed graph. The requirement for the graph $\Gamma(F(\cdot))$ to be convex-valued excludes such jumps.

7.2.3 Application: Existence of Nash equilibria

Reading: (Nash 1950), (Fudenberg and Tirole 1991, 29f), (Sundaram 1996, 246f).

Now we are ready to do some 'real' economics. The basic question we will answer in this section is "Does every game have a Nash equilibrium?" The answer is affirmative, but our procedure




Figure 7.9: The assumption of the convex-valuedness of the graph is indispensable.

to establish this will be a bit involved. First, again, we'll need some additional economics. (We stick to the same notation as in the dominance proof we did previously.)

Def: We call the set of player i 's best responses to the opponents' strategy profile σ_{-i} the **best reply** correspondence $B_i(\cdot)$. This correspondence gives player i 's set of best responses

$$\sigma_i^* = B_i(\sigma_{-i}) = \{\sigma_i^* \in \Delta(S_i) \mid \forall \sigma_i \in \Delta(S_i), \Pi_i(\sigma_i^*, \sigma_{-i}) \geq \Pi_i(\sigma_i, \sigma_{-i})\}.$$

Def: We define the set of all best response correspondences as $B(\sigma) \equiv B_1(\sigma_{-1}) \times \dots \times B_n(\sigma_{-n})$.

Def: We call a strategy profile $\sigma^* = \{\sigma_1^*, \dots, \sigma_n^*\}$ **Nash equilibrium** if each player i finds it optimal to play σ_i^* given his opponents play σ_{-i}^* . Formally, if for all players $i = 1, \dots, n$, and all out of equilibrium strategies $\sigma_i \in \Delta(S_i)$ we have

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

Rem: Hence the expected equilibrium payoff for i is $u_i(B_i(\sigma_{-i}^*), \sigma_{-i}^*)$.

Theorem 42. (Nash) Any game with a finite number of pure strategies and a finite number of players n has at least one Nash equilibrium in mixed strategies.

Proof. The proof consists of showing that our setting satisfies all of Kakutani's assumptions and then concluding that there exists a σ^* such that lies in the set of best responses $B(\sigma^*)$. We will do this in turn

1. $\Delta(S_i)$ is a convex, closed, and bounded subset of \mathbb{R}^n :
 - (a) $\Delta(S_i)$ is convex because if $\sigma_i, \sigma'_i \in \Delta(S_i)$ then any mixture $\lambda\sigma_i + (1-\lambda)\sigma'_i \in \Delta(S_i)$ for all $\lambda \in [0, 1]$.
 - (b) $\Delta(S_i)$ is closed because its boundary points (assigning probability 0 and 1 to some pure strategy) are included in the set.
 - (c) $\Delta(S_i)$ is bounded because all probabilities are bounded between $[0, 1]$.
 - (d) $\Delta(S_i)$ is finite since every finite dimensional probability simplex is finite.
2. $B_i(\Delta(S_i))$ is non-empty and convex-valued:

- (a) There is a finite number of pure strategies (by assumption), therefore there exists at least one best reply to a given profile. Thus the best response correspondence cannot be empty.
- (b) Let $\sigma_i, \sigma'_i \in \Delta(S_i)$ be best replies (e.g. each yielding 10), then the mixture is a best reply as well because it will yield the same amount. Hence, $B_i(\Delta(S_i))$ is convex valued.
3. $B_i(\Delta(S_i))$ has a closed graph; thus any boundary point of the graph is included.
- (a) Let α^n, β^n with $\beta^n \in B_i(\alpha^n)$ be interior points of the graph. Then, if a sequence in the graph $\{(\alpha^n, \beta^n)_{n=1}^\infty\}$ converges to a point (α, β) , we want to show that this point also belongs to the graph. $\beta^n \in B_i(\alpha^n)$ means that, for all i , β'_i

$$u_i(\beta_i^n, \alpha_{-i}^n) \geq u_i(\beta'_i, \alpha_{-i}^n).$$

But since $u_i(\cdot)$ is continuous (it is a polynomial), taking limits of the above preserves the weak inequality. Therefore, for all i , β'_i

$$u_i(\beta_i, \alpha_{-i}) \geq u_i(\beta'_i, \alpha_{-i})$$

and therefore $\beta_i \in B(\alpha_{-i})$. Which is what we wanted to show. This concludes the proof. □

7.2.4 Tarski's Theorem

This theorem and the setting in which it is used—*lattice theory*—is important for the theory of *supermodular* games. For an advanced discussion of lattice theory consult (Birkhoff 1979). Intuitively, supermodular games can be viewed as strategic interactions where the players' strategies are a generalisation of the 'complements' case in the Slutsky equation (3.4.4). The development of supermodular games in an economic context is mainly due to the work of (Topkis 1976), (Milgrom and Roberts 1990), and (Vives 1990). An introduction to the subject from the point of view of game theory can be found in (Fudenberg and Tirole 1991, 12.3). From this angle, the main virtue of lattice theory is that in games of supermodular structure, equilibrium existence can be assured by the mere assumption of nondecreasing best responses and completeness of the underlying lattice.

We start by recalling some ideas from set theory.

Rem: Consider the relation ' \leq ', to be interpreted as 'is contained in,' 'is part of,' or 'is less or equal to.' If this relation partially orders the set M , M is called a *poset*. By greatest lower bound (inf or 'meet' or \wedge) of the poset M , we mean an element z , not necessarily included in M such that, for any $x, y \in M$, we have both $z \leq x$ and $z \leq y$ and there is no $l \in M$ such that $z < l \leq x$ while still $l \leq y$. The least upper bound (sup or 'join' or \vee) is interpreted accordingly.⁶ The relation $x \leq y$ is identical to $y \geq x$. If $x \leq y$ and $x \neq y$ we write $x < y$.

⁶ You may want to consult figure 2.4 on page 28 and its discussion to refresh your memory with respect to bounds.

Def: Let M be a nonempty partially ordered set ('poset') wrt to the relation ' \leq '. M is called a **lattice** if both $\inf M$ and $\sup M$ are contained in M . Sometimes a lattice is denoted $\langle M, \leq \rangle$.

Rem: Notice that $x \leq y$ is ambiguous if $n > 1$ in $x, y \in \mathbb{N}^n$. We will use the interpretation that $x \leq y$, if $x_i \leq y_i$ for all $i = 1, \dots, n$, but $x_j < y_j$ for some j .

Exl: Let $M \subset \mathbb{N}^2 = \{(1, 0), (0, 1)\}$. We can visualise such a set as in figure 7.10. (Full dots are supposed to be elements of the lattice, empty circles are not.) Notice that M is partially ordered wrt the relation ' \leq ' but it is *not* a chain. Observe that $\inf_x M = 0$, $\sup_x M = 1$ and $\inf_y M = 0$, $\sup_y M = 1$, therefore $\inf M = (0, 0)$, $\sup M = (1, 1)$ which are not contained in M . Hence M is not a lattice. \triangleleft



Figure 7.10: $M \subset \mathbb{N}^2 = \{(1, 0), (0, 1)\}$ is not a lattice but $M = \{(0, 0), (1, 1)\}$ is one.

Def: A lattice M is called a **complete** lattice if for every $L \subset M$, both $\inf L$ and $\sup L$ are contained in M .

Def: A subset L of the lattice M is called a **sublattice** if both $\inf L \in L$ and $\sup L \in L$. (A sublattice is therefore a lattice in its own right.)

Def: A sublattice L has a **greatest** element $\bar{x} \in L$ if $x \leq \bar{x}$ for all $x \in L$. It has a **least** element $\underline{x} \in L$ if $x \leq \underline{x}$ for all $x \in L$.

Lemma 43. (Zorn) *If M is a poset and every chain (non-decreasing sequence) has an upper bound in M , then M has a maximal element.*

Theorem 44. *If L is a nonempty, compact sublattice of \mathbb{R}^n then it has a greatest and smallest element.*

Proof. Since L is compact, every chain (non-decreasing sequence) has an upper bound in L , namely the limit of the sequence. By Zorn's lemma there exists a maximal element. If this is not the greatest element, then there exists a maximal element $x' \in L$ and some other element $x \in L$ such that $x' \not\geq x$. Since L is a sublattice, $\sup L \in L$ and by definition $\sup L \geq x$ which contradicts the assumption that x' is maximal. Hence x' is a greatest element. \square

Def: A function $f(x_1, \dots, x_n)$ is **nondecreasing** if $x'_i \geq x_i$ for each i implies that $f(x'_1, \dots, x'_n) \geq f(x_1, \dots, x_n)$. It is called **increasing** if $x'_i > x_i$ for each i implies $f(x'_1, \dots, x'_n) > f(x_1, \dots, x_n)$.

Theorem 45. (Tarski 1955) *Every nondecreasing function f mapping a lattice $\langle S, \leq \rangle$ into itself has a nonempty set of fixed points P forming itself a complete lattice.*

Proof. Consider the 'primary' lattice of fixed points $\langle P, \leq \rangle$ and let

$$u = \sup\{f(x) \geq x\} \tag{7.2.1}$$

denote the least upper bound of the set of values of f above the diagonal. We clearly have $x \leq u$ for every x with $f(x) \geq x$. Since the function is nondecreasing,

$$f(x) \leq f(u) \text{ and } x \leq f(u).$$

By (7.2.1), we conclude that

$$u \leq f(u). \tag{7.2.2}$$

Therefore

$$f(u) \leq f(f(u)),$$

so that $f(u)$ belongs to the set above the diagonal $\{f(x) \geq x\}$; consequently, by (7.2.1),

$$f(u) \leq u. \tag{7.2.3}$$

Together, (7.2.2) and (7.2.3) imply that u is a fixed point of f . Hence we conclude by (7.2.1) that u is the least upper bound of P and

$$\sup P = \sup\{f(x) \geq x\} \in P. \tag{7.2.4}$$

Now consider the ‘dual’ lattice $S' = \langle S, \geq \rangle$. As the primary lattice, it is complete. Obviously the supremum of the dual lattice coincides with the infimum of the primary lattice. Hence, by applying the result established in (7.2.4) to the dual lattice, we obtain

$$\inf P = \inf\{f(x) \leq x\} \in P. \tag{7.2.5}$$

Hence P is a lattice.

To check completeness, consider a subset Y of P , and define

$$B = \langle [\sup Y, \sup S], \leq \rangle$$

which is a complete lattice. For any $x \in Y$, we have $x \leq \sup Y$ and hence

$$x = f(x) \leq f(\sup Y);$$

therefore $\sup Y \leq f(\sup Y)$. Consequently, $\sup Y \leq z$ implies

$$\sup Y \leq f(\sup Y) \leq f(z).$$

Thus, by restricting the domain of f to $[\sup Y, \sup S]$, we obtain an increasing function $f' : [\sup Y, \sup S] \mapsto [\sup Y, \sup S]$. By applying (7.2.5) to the lattice B and the function f' , we conclude that the greatest lower bound v of all fixed points of f' is itself a fixed point of f' . Obviously, v is also a fixed point of f , and in fact the least fixed point of f which is an upper bound of all elements of Y . In other words, v is the least upper bound of Y in the system $\langle P, \leq \rangle$.

Hence, by passing to the dual lattices S' and B' , we see that there exists also a greatest lower bound of Y in $\langle P, \leq \rangle$. Since Y is an arbitrary subset of P , we finally conclude that $\langle P, \leq \rangle$ is complete. \square

Discussion Figure 7.11 shows the logic of the argument for a single dimensional lattice $[0, 1] \subset \mathbb{R}$.




Figure 7.11: The indispensability of $f : S \mapsto S$ to be nondecreasing in Tarski's fixed point theorem.

Exl: Consider again the relation ' \leq ' and let the poset $S \subset \mathbb{N}^2 = \{(0, 0), \dots, (3, 3)\}$. Let $L \subset S$, $L = \{x = (1, 1), y = (1, 2), z = (2, 1)\}$. Furthermore let the function $f : S \mapsto S$ be defined as $f(x) = f(y) = f(z) = (2, 2)$; all other points in S are left unchanged. This situation is illustrated in the left panel of figure 7.12. As is easily checked, the original poset is a complete lattice and f is nondecreasing. Hence Tarski's Theorem predicts that the set of fixed points is a complete lattice. This is indeed the case.

The right panel shows the same initial lattice. This time let f be defined only as $f(x) = (2, 2)$, all other points are left unchanged. The requirement of f to be nondecreasing is violated because although $x \leq y$, $f(y) \leq f(x)$ and $x \leq z$, $f(z) \leq f(x)$. Let us check whether the lattice of fixed points is complete. Take the set $S' = \{f(y) = (1, 2), f(z) = (2, 1)\}$. We have no trouble establishing that $\sup S' = (2, 2)$ but the infimum is problematic: *Both* $(0, 1)$ and $(1, 0)$ are candidates. Recalling that for any set only one element can have the infimum property, we know that the two points are *not* infima. But then there is no infimum at all in L because $(0, 0)$ is smaller than both $(0, 1)$ and $(1, 0)$. Hence the lattice of fixed points is not complete and we have verified Tarski's theorem for this case. \triangleleft

7.2.5 Supermodularity

We now proceed to the discussion of supermodular games (which have nothing to do with fixed points—we just discuss them here because we want to reap the benefits of introducing lattices). The theory of supermodular games provides a framework for the analysis of strategic complementarities. This class of game includes models of oligopolistic competition, macroeconomic coordination failure, bank runs and R&D competition. Supermodular games are characterised




Figure 7.12: Two mappings of a lattice into itself.

by the fact that each agent's strategy set is partially ordered, the marginal returns to increasing one's strategy rise with increases in the opponents' strategies and, in the case of multidimensional actions, the marginal returns to any one component of the agent's strategy rise with increases in the other components. As a result, these games exhibit strategic complementarities that yield monotone increasing individual best responses as strategies are completely ordered.

Let us begin by recalling from section 2 that a function $f : \mathbb{R} \mapsto \mathbb{R}$ is said to be *subadditive* if it satisfies the inequality $f(x + y) \leq f(x) + f(y)$ whenever $x, y \in \mathbb{R}$. It is called *superadditive* if $f(x + y) \geq f(x) + f(y)$.

Def: A function f on the lattice M is said to exhibit **increasing first differences**⁷ if, for $x, y \in M$

$$f(x', y') - f(x, y') \geq f(x', y) - f(x, y)$$

for $x' \geq x, y' \geq y, x, x', y, y' \in M$. The relationship is called strict if all involved inequalities are strict.

Def: A function f on the lattice M is called **supermodular** if for $x, y \in M$

$$f(\inf(x, y)) + f(\sup(x, y)) \geq f(x) + f(y).$$

Again there is also a strict version of the definition.

Unfortunately, increasing first differences does not imply strategic complementarity. It is precisely the aim of introducing supermodularity to show that this generalised version of increasing first differences implies strategic complementarity. (The below example is supposed to illustrate this.)

It is worth mentioning that supermodularity is automatically satisfied in the single-dimensional (scalar) case. You will need to go to higher dimensions to find counterexamples. Similarly, if the lattice the concept is defined on can be ordered by ' \leq ', supermodularity is always satisfied.

Since we want to apply our definition of supermodularity to games, we look at payoff functions $f = u_i$ and strategy lattices $M = X \equiv X_1 \times \dots \times X_n$ in the remainder of this section. We repeat the definition for supermodularity applied to this setting. The payoff function u_i is *supermodular*

⁷ If you happen to know the Spence-Mirrlees 'single-crossing' condition, you will find striking similarities to the concept of increasing first differences.

if, for all $x, x' \in X$,

$$u_i(\sup(x, x')) + u_i(\inf(x, x')) \geq u_i(x) + u_i(x').$$

Exl: (**D. Gale**) There are 2 players $i = 1, 2$, with strategy spaces $X_1 = \{0, 1\} \times \{0, 1\}$ and $X_2 = \{0, 1\}$. The payoffs to player 1 are given by the following matrices:

$(c = 0)$	$b = 0$	$b = 1$		$(c = 1)$	$b = 0$	$b = 1$
$a = 0$	0	4		$a = 0$	0	10
$a = 1$	5	3		$a = 1$	8	9

Figure 7.13: The payoff matrices; 1 chooses (a, b) and 2 chooses c .

By direct calculation, for any $(a', b') \geq (a, b)$, the first differences are:

$(a, b) \rightarrow (a', b')$	$c = 0$	$c = 1$
$(0, 0) \rightarrow (0, 1)$	4	10
$(0, 0) \rightarrow (1, 0)$	5	8
$(0, 0) \rightarrow (1, 1)$	3	9
$(0, 1) \rightarrow (1, 1)$	-1	-1
$(1, 0) \rightarrow (1, 1)$	-2	-1

It is clear that the payoff function u_1 satisfies increasing first differences. However, the best response when $c = 0$ is $(a, b) = (0, 1)$ and the best response when $c = 1$ is $(a, b) = (1, 0)$. It is easy to check that u_1 is not supermodular in x_1 . From the definition, we should have

$$u_1(\inf(x_1, x'_1), c) + u_1(\sup(x_1, x'_1), c) \geq u_1(x_1, c) + u_1(x'_1, c)$$

for every c , but by putting $x_1 = (0, 1)$ and $x'_1 = (1, 0)$ it is easy to see that

$$\begin{aligned} u_1(\inf(x_1, x'_1), c) + u_1(\sup(x_1, x'_1), c) &= u_1(0, 0, c) + u_1(1, 1, c) \\ &< u_1(0, 1, c) + u_1(1, 0, c) \\ &= u_1(x_1, c) + u_1(x'_1, c), \end{aligned}$$

for every c . Hence the best response is not monotonically non-decreasing in the opponent's strategy. \triangleleft

Theorem 46. *If u_i is supermodular, then u_i is supermodular in $x_i \in X_i$ and exhibits increasing first differences.*

Proof. Supermodularity in x_i follows immediately from the definition of supermodularity, if we put $x = (x_i, x_{-i})$, $x' = (x'_i, x'_{-i})$, $x_{-i} = x'_{-i}$.

To prove that u_i exhibits increasing first differences, let $x'_i > x_i$ and $x'_{-i} > x_{-i}$, where $x_i, x'_i \in X_i$ and $x_{-i}, x'_{-i} \in X_{-i}$. Then the definition of supermodularity can be rewritten as

$$u_i(x'_i, x'_{-i}) + u_i(x_i, x_{-i}) \geq u_i(x_i, x'_{-i}) + u_i(x'_i, x_{-i});$$

rearranging terms gives

$$u_i(x'_i, x'_{-i}) - u_i(x_i, x'_{-i}) \geq +u_i(x'_i, x_{-i}) - u_i(x_i, x_{-i})$$

as required. \square

Theorem 47. (Topkis) *If $X_i = \mathbb{R}^\ell$ and u_i is C^2 in $x_i \in X_i$, u_i is supermodular iff*

$$\frac{\partial^2 u_i}{\partial x_{ik} \partial x_{ij}}(x_i, x_{-i}) \geq 0, \quad \forall j, k = 1, \dots, \ell.$$

Theorem 48. *If $X = \mathbb{R}^{\ell n}$ and u_i is C^2 , u_i is supermodular iff, for any two components x_j and x_k of $x \in X$,*

$$\frac{\partial^2 u_i}{\partial x_j \partial x_k}(x) \geq 0, \quad \forall j, k = 1, \dots, \ell n.$$

Proof. Let $e_k = (0, \dots, 0, 1, 0, \dots, 0)$ be an ℓn -vector with the unit in the k^{th} place. Then, for any small numbers $\varepsilon, \eta > 0$ and indexes $k \neq j$, put $u = (x + \varepsilon e_k)$ and $v = (x + \eta e_j)$. Supermodularity of u_i implies that

$$u_i(\sup(u, v)) + u_i(\inf(u, v)) \geq u_i(u) + u_i(v).$$

Substituting the expressions for u and v gives

$$u_i(\sup(x + \varepsilon e_k, x + \eta e_j)) + u_i(\inf(x + \varepsilon e_k, x + \eta e_j)) \geq u_i(x + \varepsilon e_k) + u_i(x + \eta e_j)$$

which implies that

$$\eta \varepsilon \frac{\partial^2 u_i}{\partial x_j \partial x_k}(x) \geq 0, \quad \forall j, k = 1, \dots, \ell n.$$

Conversely, if $X = \mathbb{R}^2$ and

$$\frac{\partial^2 u_i}{\partial x_j \partial x_k}(x) \geq 0$$

everywhere, then this implies that

$$u_i(\sup(x, y)) - u_i(y) \geq u_i(x) - u_i(\inf(x, y))$$

as required. \square

Rem: A normal form game $G = \{N, M, u\}$ ⁸ and u are payoffs), is supermodular if (i) M is a sublattice of \mathbb{R}^n , (ii) u_i exhibits increasing first differences, and (iii) u_i is supermodular in x_i .

Exl: (D. Gale) There is a neat trick that allows us to treat a 2-player **Cournot duopoly** with *strategic substitutes* as a supermodular game. We simply define the strategy of one producer to be the *negative* of his output.⁹ Then an increase in his strategy reduces his output and under the usual conditions this encourages his opponent to increase his output and his strategy.

⁸ In a normal form game $G = \{N, M, u\}$, N is the set of players, M is the set of (product) strategies for all players, and u are payoffs.

⁹ This trick only works in the case of $n = 2$ players. For cases $n > 2$, we do not have a theory for strategic substitutes.

We have two producers $i = 1, 2$; producer 1's strategy set is $X_1 = [0, \text{infy})$ and player 2 chooses from $X_2 = (-\infty, 0]$. The total output is $Q = q_1 + q_2 = x_1 - x_2$ and the payoff functions are defined for $x_i \in X_i$ as

$$u_1(x) = P(x_1 - x_2)x_1$$

and

$$u_2(x) = -P(x_1 - x_2)x_2.$$

Then

$$\frac{\partial^2 u_1(x)}{\partial x_1 \partial x_2} = -P''(x_1 - x_2)x_1 - P'(x_1 - x_2)$$

and

$$\frac{\partial^2 u_2(x)}{\partial x_1 \partial x_2} = P''(x_1 - x_2)x_2 - P'(x_1 - x_2).$$

If $P''(x_1 - x_2) \leq 0$, it is easy to check that both cross partials are positive. This would be, for example, the case if the inverse demand function were affine. \triangleleft

Exercises

Exc 7.1: Imagine you stand in front of a $(\frac{3}{4})$ -full barrel of water and you stir in it with a stick. What would a fixed point be? Is there necessarily a fixed point?

Exc 7.2: Try to solve the fixed point equation $x = 1 - x^5$ on the set $M = [0, 1]$ by what is called successive approximation. This means you compute a sequence for $n = 0, 1, 2, \dots$

$$x_{n+1} = 1 - x_n^5.$$

Start with $x_0 = 0.5$ and write down a sequence $\{x_n\}$ for a couple of n . What goes wrong?

Exc 7.3: Let the letter E be mapped continuously into itself. Must there be a fixed point? What about O or I ?

Exc 7.4: Let X be the interval $-1 \leq x \leq 1$. Let $F(x)$ be the set of numbers y satisfying $x^2 + y^2 \geq \frac{1}{4}$. Draw the graph $G = \{(x, y) : x \in X, y \in F(x)\}$. Is G closed? Characterise G in terms of correspondences. Draw the sets $F(-1), F(0), F(1/3), F(1/2), F(3/4)$. Which of them are convex?

Exc 7.5: For $0 \leq x \leq 1$ define

$$f(x) = \begin{cases} 17 & \text{if } x = 0 \\ \frac{1}{x} & \text{otherwise} \end{cases}$$

Draw the graph $G = \{(x, y) : 0 \leq x \leq 1, y = f(x)\}$. Characterise f in terms of closedness and G in terms of convexity. Define

$$F(x) = \{y : 0 \leq y \leq f(x)\}.$$

Draw the graph $G' = \{(x, y) : 0 \leq x \leq 1, y \in F(x)\}$. Characterise F in terms of closedness and G' in terms of closedness and convexity.

Exc 7.6: Here is a quick proof of Kakutani's Theorem: The assumptions imply that there exists a continuous point-valued function $f(x) \in F(x)$. Now Brouwer's Theorem guarantees a fixed point $x^* = f(x^*)$. The problem is the proof is false. Why? Construct a counterexample.

Exc 7.7: Which of the following are sub-lattices of \mathbb{R}^2 ?

1. $X = \{(0, 0), (\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, 1), (1, \frac{1}{2}), (1, 1)\}$;
2. $X = \{(0, 0), (0, \frac{1}{2}), (1, 0), (0, 1), (1, 1)\}$;
3. $X = X_1 \cup X_2 \cup X_3$ where
 - (a) $X_1 = \{(1, 1 - \frac{1}{n}) | n = 1, 2, \dots\}$,
 - (b) $X_2 = \{(1 - \frac{1}{n}, 1) | n = 1, 2, \dots\}$,
 - (c) $X_3 = \{(1 - \frac{1}{n}, 1 - \frac{1}{n}) | n = 1, 2, \dots\}$;
4. $X = \{(1, 0), (1, 1), (2, 1), (2, 2), \dots\}$;

Exc 7.8: (Optional) What happens to our existence proof if we do not allow for mixed strategies, i.e. if we force the players to use pure strategies? If you can work out the answer if we were to use

Tarski's theorem (i.e. if we restrict ourselves to the study of equilibria in supermodular games) you will discover a surprising property of equilibria in supermodular games.

Exc 7.9: Which of the following utility functions are supermodular?

1. The Cobb-Douglas utility function $u(\mathbf{x}) = x_1^\alpha x_2^{1-\alpha}$, where $\alpha \in (0, 1)$.
2. The CES utility function $u(\mathbf{x}) = (x_1^\rho + x_2^\rho)^{\frac{1}{\rho}}$, where $\rho > 0$.
3. The Leontief utility function $u(\mathbf{x}) = \min \{\alpha_1 x_1, \alpha_2 x_2\}$, where $\alpha_1, \alpha_2 > 0$.

Part II

Dynamic analysis

Chapter 8

Introduction to dynamic systems

Reading: The general reference for dynamic systems is (Kamien and Schwartz 1991) and a comprehensive textbook on ordinary differential equations is (Hirsch and Smale 1974). We will give more specific references on our way.

The second part of the notes is meant to prepare you for macroeconomics. Since most modern macro-models deal with dynamical systems and the classical way to represent (continuous time) dynamical systems is by means of a system of differential equations we start with these. For a variety of reasons we will focus on autonomous first-order linear ordinary differential equations (ODEs)—after all these are the only ones we can be certain to find a closed solution for.¹ The principal reasons are difficulty and time. The general treatment of partial differential equations (PDEs) is, however interesting, well beyond our grasp and we certainly do not have the time for a serious treatment of both difference and differential equations. But beyond these pragmatic reasons there is also the element that most of the higher order systems you may encounter in macro can be reduced to the simpler first-order case.² And the mathematical intuition you gain on the study of differential equations brings you a long way to understanding difference equations. So it is a reasonable strategy to cover only this special case—and we will do it with care. The best book I know on the subject is (Hirsch and Smale 1974)—it is not exactly an easy read and if you do not like it there are a host of further references later on. The remainder of this introduction contains some motivational examples of dynamical systems that hopefully whet your appetite for the more serious stuff in the following sections.

In most real economic problems we have to study the development of variables over time. The intertemporal development of these variables is usually expressed in the form of equations. If time is taken to be a discrete variable (i.e. we look at non-zero length periods) these equations are called *ordinary difference equations* (ODEs) and if time is taken to be continuous, we speak of *ordinary differential equations* (ODEs).³ Both techniques solve essentially the same problem

¹ I leave a previous version of my notes on difference equations as a separate section in the notes although I'm doubtful about their usefulness. It gives, however, a set of references to more useful treatments.

² Every n^{th} -order general ODE can be written as a system of n first-order ODEs. This is discussed in section 8.1.5.

³ *Partial differential equations* (PDEs) or *partial difference equations* (PDEs) are equations that depend on the partial derivatives of a number of variables. Their solution is an entirely different business and we will try to stay away from them as far as possible. Probably the simplest partial differential equation involving a function $u(\zeta, \eta)$ of two independent variables

but depending on how the model is developed, one technique may lend itself more naturally to the solution than the other.

Exl: Radioactive decay. The law describing radioactive disintegration states that the rate of decay with time—that is, the amount of a substance changing per time interval—is proportional to the amount of substance present. Initially, when the material is a pure radioactive element, this rate of decay is high. As the radioactive element changes into a stable element, however, the rate of change falls because there is less radioactive material. Therefore, the rate of decay decreases continuously with time.

In this process the quantity of interest is the amount of substance n remaining after a given time—i.e., the number of atoms or the mass. Clearly this quantity will depend in some way on the time that has elapsed; in other words, it is a mathematical function of time and can be denoted by $n(t)$. A possible formulation of this law is

$$\dot{n} \equiv \frac{dn}{dt} = -\lambda n. \quad (8.0.1)$$

In this case n represents the amount of radioactive material, λ is a constant, and $\dot{n} \equiv \frac{dn}{dt}$ is the notation for the rate at which n increases with time (it decreases here with time because of the minus sign). ◁

Exl: When the independent variable of a function is time and the dependent variable distance, the first derivative is velocity, the rate at which distance changes with time. The second derivative is acceleration, the rate at which velocity changes. ◁

Notation: As a notational convention, we denote the time dependent discrete variable x as x_t and the time dependent continuous variable x as $x(t)$. We will exclusively look at systems where the independent variable is t .

8.1 Elements of the theory of ordinary differential equations

Reading: Both (Chiang 1984, chapters 14, 15) and (Sydsæter and Hammond 1995, chapter 21) give bread and butter, undergraduate-level introductions. A more advanced introduction is presented by (Cronin 1980) and a concise but yet directly useful version is in the first two chapters of (Bender and Orszag 1978).⁴ (Kamien and Schwartz 1991, Appendix B) is a self-contained introduction to differential equations and finally (Hirsch and Smale 1974) is an excellent full treatment that goes well beyond our immediate needs.

can be formalised

$$\frac{du}{d\zeta} = 0,$$

which may readily be seen to have the general solution of an arbitrary function $f(y)$ of the other variable only:

$$u = f(\eta).$$

A comparison with the case of ODEs in which the equation $u'(x) = 0$ has solution $u(x) = c$ illustrates the point that the role played by arbitrary constants in the theory of ODEs is taken over by arbitrary functions in the theory of PDEs. This makes their solution much more difficult.

⁴ If you prefer an introduction to ODEs that focusses on micro and evolutionary systems, a look at (Weibull 1995, chapter 6) may be an alternative.

Def: Ordinary differential equations (ODEs) are equations for which the unknown x is a function of only one variable t .

ODEs differ from ordinary equations in two aspects:

1. the unknown is a function $F(\cdot)$ and not a number, and
2. the equation includes at least one of the derivatives of the unknown function.

Solving an ODE means to find a function and its derivative that satisfy the given differential equation. A system of ODEs represents a deterministic dynamic system in continuous time. The simplest form of an ODE is

$$\dot{x} = F(t).$$

It can be solved by ordinary integration because

$$\dot{x} = F(t) \Leftrightarrow x = \int F(t) dt + C$$

where C is an arbitrary constant.

Exl: Let's look at $F(t) = t^2 - 1$. By integration we get

$$\dot{x} = t^2 - 1 \Leftrightarrow x = \int (t^2 - 1) dt + C = \frac{1}{3}t^3 - t + C.$$

C represents, of course, the constant of integration. Since we cannot pin it down, we really obtain a whole set of solutions (called a 'general' solution) only differing by C . \triangleleft

As usual, we start with a couple of definitions.

Def: The **order** of an ODE is the highest order of derivatives appearing.

Def: The **degree** of an ODE is the highest power of the derivative in the equation.

In general, the ODEs we look at may change in time as the system evolves such as $\dot{x} = F(x, t)$. We will, however, not look at such dynamics but only at ones that do not depend on time such as $\dot{x} = F(x)$.

Def: A differential equation that does not depend on time is called **autonomous** or **time homogeneous**.

Def: A differential equation is called **linear** if the unknown function \dot{x} appears only in its first degree and it can be written in the form

$$\dot{x} + u(t)x = w(t)$$

where $u(t)$ and $w(t)$ are arbitrary functions. We will refer to $u(t)$ as the **coefficient** of the dependent variable x and $w(t)$ as the **term** of the ODE.

Def: An equation is called **homogeneous** if $w(t) \equiv 0$.

Def: A **general** solution to an ODE is the set of all solutions to the ODE.

Def: A **particular** solution to an ODE is the solution to the ODE that is obtained by specifying particular values for the constants of integration.

8.1.1 Existence of the solution

The unknown function $F(x, t)$ is a mapping from an open set $X \subset \mathbb{R}^n$ to \mathbb{R}^n . The mapping $F(x, t)$ itself is called a **vector field**, defining at each point x the direction and velocity of the flow at x . The $x = (x_1, \dots, x_n) \in X$ are called the **state vector**, X is called the **state space**, and the *rhs* of $\dot{x} = F(x, t)$ specifies the direction and velocity of change of the state vector in the state space X .

For each component x_i of the state x , the autonomous function $F_i(x) \in \mathbb{R}$ is its time derivative. See figure 8.1 for the illustration of a vector field in \mathbb{R}^2 . The dashed curve in the same figure represents a solution $\xi(t, x_0)$ through x_0 to our ODE.




Figure 8.1: A vector field $F(\cdot)$ in \mathbb{R}^2 .

Def: A (local) solution through a point $x_0 \in X$ is a function $\xi(\cdot, x_0) : T \mapsto X$, where T is an open interval containing $t = 0$, such that $\xi(0, x_0) = x_0$, and such that for all $t \in T$

$$\frac{d}{dt}\xi(t, x_0) = F[\xi(t, x_0)].$$

As it turns out, the condition that guarantees the existence of a solution $\xi(t, x_0)$ for all vector fields $F(x)$ is that $F(x)$ is (locally) Lipschitz continuous as defined in chapter 2. Let us state this classical result formally.

Theorem 49. (Picard-Lindelöf) *If $X \subset \mathbb{R}^n$ is open and the vector field $F : X \mapsto \mathbb{R}^n$ is Lipschitz continuous, then the system $\dot{x} = F(x)$ has a unique solution $\xi(\cdot, x_0) : T \mapsto X$ through every state $x_0 \in X$. Moreover the solution $\xi(t, x_0)$ is continuous in t and $x_0 \in X$.*

Proof. See (Hirsch and Smale 1974, 164ff). \square

After this discussion of the basic ideas, let us turn to our cookbook and state a set of rules that should enable you to solve anything that comes your way.

8.1.2 First-order linear ODEs

1) **Constant coefficient** $a = u(t)$ **and constant term** $w(t) = b$.

This means we are looking at ODEs of the form⁵

$$\dot{x} + ax = b. \quad (8.1.1)$$

1.a) In the homogeneous case ($b \equiv 0$), direct integration gives the **general** solution of (8.1.1) as

$$x = Ae^{-\int a dt} = Ae^{-at} \quad (8.1.2)$$

for arbitrary A .

Exl: Let's look at how we do that in detail: The recipe is to integrate by applying the substitution rule

$$\int f(x) \frac{dx}{dt} dt = \int f(x) dx = F(x) + C, \quad (8.1.3)$$

and then the log rule

$$\int \frac{1}{x} dx = \ln |x| + C,$$

to obtain $F(x) = \ln |x| + C$. In our case this boils down to (i) transforming

$$\frac{\dot{x}}{x} = -a,$$

(ii) integrating both sides

$$\int \frac{\dot{x}}{x} dt = - \int a dt, \quad (8.1.4)$$

(iii) applying the substitution rule to the *lhs* of (8.1.4) (with $f(x) = \frac{1}{x}$)

$$\int \frac{\dot{x}}{x} dt = \int \frac{dx}{x} = \ln |x| + C_1, \quad (8.1.5)$$

(iv) integrating the constant a on the *rhs* of (8.1.4)

$$- \int a dt = -at + C_2, \quad (8.1.6)$$

(v) setting (8.1.5)=(8.1.6) to obtain ($C \equiv C_2 - C_1$)

$$\ln |x| = -at + C, \quad (8.1.7)$$

(vi) by realising that $e^{\ln x} = x$, we raise each side of (8.1.7) to become the exponent of e (this

⁵ Notice that this is a version of our radioactive decay analysis in (8.0.1).

procedure is called ‘taking antilogs’) and obtain

$$x = e^{-at}e^C = Ae^{-at} \quad (8.1.8)$$

by defining $A \equiv e^C$. ◁

The **definite** solution to (8.1.1) is

$$x = x(0)e^{-at}$$

that pins down a specific initial condition $x(0) = A$. Any particular value of A that is supplied in the general solution leads to a **particular** solution. Notice that all these solution are functions that do *not* contain time derivatives any more.

1.b) In the non-homogeneous case, we have to examine two parts of the **general** solution x_c and x_p ; the first one is the solution to the homogeneous equation above

$$x_c = Ae^{-at} + C$$

but this time every particular solution will also depend on b . In the simplest case that our dependent variable is just a constant $x = k$ we know that the time derivative will be zero. Therefore (8.1.1) reduces to $ax = b$ implying the **particular integral** to be

$$x_p = \frac{b}{a}.$$

Therefore we get the **complete general** solution

$$x(t) = x_c + x_p = Ae^{-at} + \frac{b}{a}.$$

Setting $t = 0$ above we get the **initial condition**

$$\begin{aligned} x(0) &= A + \frac{b}{a} \Leftrightarrow \\ A &= x(0) - \frac{b}{a} \end{aligned}$$

and therefore the **definite** solution is

$$x(t) = \left[x(0) - \frac{b}{a} \right] e^{-at} + \frac{b}{a}.$$

2) Variable coefficient $u(t)$ and term $w(t)$

We are back in the more general case of

$$\dot{x} + u(t)x = w(t). \quad (8.1.9)$$

2.a) In the homogeneous case ($w(t) \equiv 0$), we follow the same procedure as in (8.1.2): We rearrange to

$$\frac{\dot{x}}{x} = -u(t)$$

and get the **general** solution of this by integrating first the *lhs* wrt t

$$\int \frac{\dot{x}}{x} dt = \int \frac{dx}{x} = \ln x + c,$$

then the *rhs*

$$\int -u(t) dt = - \int u(t) dt$$

(we cannot really do more without specifying the functional form of $u(\cdot)$) and equating the two

$$\ln x = -C - \int u(t) dt.$$

Taking antilogs (cf. 8.1.8) and defining $A \equiv e^{-C}$ finally gives the **general** solution of (8.1.9) as

$$x(t) = e^{\ln x} = e^{-C} e^{-\int u(t) dt} = A e^{-\int u(t) dt}. \quad (8.1.10)$$

The **definite** solution is obtained exactly as above using the initial condition at $t = 0$.

2.b) In the non-homogeneous case, we end up with a **general** solution of the form

$$x(t) = e^{-\int u(t) dt} \left(A + \int w(t) e^{-\int u(t) dt} dt \right)$$

for arbitrary A . Looking at the term in brackets, we see that the first term, $A e^{-\int u(t) dt}$, is the same as in (8.1.10) and carries the same intuition as above. The second part stems from the variable term $w(t)$ —its derivation is straightforward but time constraints do not allow us to do this here (cf. (Chiang 1984, 482ff)).

8.1.3 First-order non-linear ODEs

Whenever x appears in a power higher than one, our ODE becomes non-linear—even if \dot{x} is contained only in the first degree. In general, we look at equations of the form

$$\dot{x} = F(x, t)$$

that are unrestricted in terms of the powers of x and t and the function $F(\cdot)$. Non-linear ODEs are not generally easily solved because they may give rise to additional solutions that cannot be obtained from the general solution. Four economically important cases, however, are soluble and should be mentioned here:

1. So called *exact* ODEs. An exact ODE is a derivative of an ODE of lower order

$$Lx = \frac{d}{dt}(Mx) = 0 \quad (8.1.11)$$

that can be simplified by integrating wrt t to obtain the (nonhomogeneous) ODE

$$Mx = C.$$

An *integrating factor* is a function of t and x which, when multiplying an ODE, makes it exact.

2. *Separable* variables of the form that $F(x, t)$ is separable into $G(x) + H(t)$. Then we can solve the non-linear ODE by simple integration.
3. So called *Bernoulli* equations. These are forms such that $\dot{x} = F(x, t) = T(t)x^m - R(t)x$, $m > 1$. Bernoulli equations can always be reduced to linear ODEs.
4. *Reduction of order* is a technique for solving linear ODEs to ones of lower order by factoring off a solution that one has been lucky enough to find.

We will look at none of these *quantitative* techniques in any detail. Please confer to a textbook on dynamics on how to develop these methods (e.g. (Kamien and Schwartz 1991), (Bender and Orszag 1978, 20) or (Chiang 1984, 489)). Instead, we will more carefully study a generally applicable *qualitative* solution method called Phase diagrams.

8.1.4 Stability and Phase diagrams

Reading: See (Kamien and Schwartz 1991, 95) for a more thorough introduction.

One advantage of the qualitative approach we will discuss in this section is that the behaviour of both linear and non-linear ODEs can be analysed. We loose, however, the precise quantitative specification of the time path of the ODE. Let us look at the qualitative time behaviour of

$$\dot{x} = F(x) \tag{8.1.12}$$

where $F(x)$ may be linear or non-linear. Since $F(x)$ does not explicitly depend on t , it is autonomous— $F(x)$, of course, still depends on t through $x(t)$.

A solution to an ODE that does not change over time is called an equilibrium state. More formally

Def: A point a is called **equilibrium state** (or **steady state** or **stationary state**) of the ODE $\dot{x} = F(x)$ if $F(a) = 0$. In this case, $x(t) = a$ for all t is a solution of the ODE $\dot{x} = F(x)$.

Def: We obtain a **phase diagram** showing x on the abscissa by putting $y = \dot{x}$ on the ordinate and graphing $y = F(x)$ in the resulting $x \times \dot{x}$ -plane.

The question we want to address in our phase diagrams is whether $F(x)$ has any stable equilibrium states. We will do this in the following example but first we have to discuss what precisely we mean by ‘stable’.⁶ We will introduce two distinct classical notions of stability: Lyapunov and asymptotic stability. Let C be a compact subset of the state space X . Intuitively, a state $x \in C$ is *Lyapunov stable* if no small perturbation of the state induces a movement away from x . While Lyapunov stability thus requires there to be no pushes away from the state, asymptotic stability requires there to be a pull towards the state. A state $x \in C$ is *asymptotically stable* if it is Lyapunov stable and all sufficiently small perturbations of the state induce a movement back toward x . Formally

⁶ Figure 8.6 shows some intuitive pictures for the discrete case.

Def: A state $x \in C$ is **Lyapunov stable** if every neighbourhood B of x contains a neighbourhood B_0 of x such that the solution to the ODE lies in B for all $x_0 \in B_0 \cap C$.

Def: A state $x \in C$ is **asymptotically stable** if it is Lyapunov stable and there exists a neighbourhood B^* such that for all $x_0 \in B^* \cap C$

$$\lim_{t \rightarrow \infty} \xi(t, x_0) = x.$$

Clearly, a state x has to be stationary in order to be stable. Otherwise the solution ξ would lead away from x even in the absence of a perturbation.

Let us now turn to the qualitative, graphical analysis.

Exl: Any solution $x = x(t)$ of (8.1.12) has an associated $\dot{x} = \dot{x}(t)$ and for every t , we can identify a pair $(x(t), \dot{x}(t))$ on the curve in the diagram in figure 8.2:




Figure 8.2: The steady state a_1 is stable, a_2 is not.

The key is the observation that $F'(x)$, the slope the ODE, is negative to the left of \hat{x} and positive to the right of \hat{x} . Thus it is zero at \hat{x} . Let us look at the behaviour of any point $(x(t), \dot{x}(t))$ on the curve above the abscissa (that is, $F(\cdot) > 0$) and to the left of \hat{x} . For all points above the abscissa, we have $F(x(t)) > 0$ and therefore $\dot{x}(t) > 0$, so that $x(t)$ increases with t . It follows that the particular point $(x(t), \dot{x}(t))$ we chose, moves from left to right in the diagram. For all points to the left of \hat{x} , the slope of $F(x)$ is negative and therefore $(x(t), \dot{x}(t))$ moves downwards. If we look at any point below the abscissa and to the right of \hat{x} , we will get $\dot{x}(t) = F(x(t)) < 0$ and the point moves from right to left and upwards.⁷ In figure 8.3, these movements are indicated through arrows on the graph of $F(x)$.

We can identify two steady states in the above example at the locus $\dot{x}(t) = 0$ —the abscissa. But while a_1 is stable in the sense that points close to the steady state tend towards the steady state, a_2 is unstable. The reason for the stability behaviour of $F(\cdot)$ is easy to see: $F'(a_1) > 0$

⁷You will have noticed that the movements we are talking about are determined by the vector fields we introduced above.

to the left of a_1 , and $F'(a_1) < 0$ to the right. Therefore, both increasing x at the left of a_1 and decreasing x at the right of a_1 will move a point closer to a_1 . The opposite is true for a_2 . \triangleleft

Our observation in the above example generalises as follows.

Proposition 50. $F(a) = 0$ and $F'(a) < 0 \implies a$ is a stable steady state for $\dot{x} = F(x)$. $F(a) = 0$ and $F'(a) > 0 \implies a$ is an unstable steady state for $\dot{x} = F(x)$.

Exl: The Solow growth model. Consider the following workhorse of neoclassical growth theory

$$\dot{k} = sf(k) - \lambda k.$$

$\dot{k} = k(t)$ is capital per worker, s the constant savings rate, $f(\cdot)$ a production function and λ the constant growth rate per worker. Given that $f(\cdot)$ is zero at $k = 0$, increasing and concave and fulfills the so-called **Inada** conditions

$$\lim_{k \rightarrow \infty} f'(k) = 0, \lim_{k \rightarrow 0} f'(k) = \infty, f(0) = 0, \quad (8.1.13)$$

we can construct the phase diagram in Figure 8.3.




Figure 8.3: Phase diagram for the Solow model $\dot{k} = sf(k) - \lambda k$.

Following exactly the same analysis as above, we obtain a unique and stable steady state. We leave this story, however, to your first lecture in growth theory. Alternatively refer to (Barro and Sala-i Martin 1995, chapter 1). \triangleleft

Exl: (Kamien and Schwartz 1991, 95) Let's raise the stakes a bit and consider the following (quasi-autonomous) problem solutions to which give rise to a second order differential (Euler) equation

$$\begin{aligned} \min_{x(t)} \int_0^T e^{-rt} [f(\dot{x}) + g(x)] dt \\ \text{s.t.: } x(t) \geq x_0; x(T) = x_T, \end{aligned} \quad (8.1.14)$$

where $f''(x) > 0$ and $g''(x) > 0$ and f, g is C^2 and convex. As always, we assume the soc to be

satisfied. The Euler equation (see below) implies

$$\ddot{x} = \frac{r f'(\dot{x}) + g'(x)}{f''(\dot{x})}. \quad (8.1.15)$$

Since the integrand of (8.1.14) is convex in x and \dot{x} , we can be sure that a solution to the above will be indeed minimising. Although we do not further specify $f, g, f', g',$ and f'' , we still can apply our diagrammatic analysis in the (x, \dot{x}) plane. (i) First notice that for any solution consistent with our law of motion (8.1.15), whenever $\dot{x} > 0$, x is increasing. Therefore, from any point (x, \dot{x}) above the $\dot{x} = 0$ axis, the x coordinate of a solution to (8.1.15) must be increasing. Thus the $\dot{x} = 0$ locus divides the plane as shown in figure 8.4.




Figure 8.4: A phase diagram for a 2^{nd} order differential equation.

(ii) Next we consider the locus of points (x, \dot{x}) such that $\ddot{x} = 0$. From (8.1.15), we know that such points satisfy

$$0 = r f'(\dot{x}) + g'(x) \quad (8.1.16)$$

which implicitly describes a curve with slope

$$\frac{d\dot{x}}{dx} = -\frac{g''(x)}{r f''(\dot{x})} < 0. \quad (8.1.17)$$

Therefore, along the $\ddot{x} = 0$ locus, a small increase in the x coordinate must be accompanied by a small decrease in the \dot{x} coordinate. Therefore, the $\ddot{x} = 0$ locus divides the (x, \dot{x}) plane in two regions as shown in figure 8.4 above—it represents the second axis in our phase diagram.

The regions the $\ddot{x} = 0$ locus creates are determined by values for (x, \dot{x}) that cause the *rhs* of (8.1.15) to be positive (and hence $\ddot{x} > 0$, therefore \dot{x} increasing) and those values of (x, \dot{x}) for which the *rhs* of (8.1.15) is negative (and thus $\ddot{x} < 0$, and \dot{x} decreasing).

The direction of movement to the right of the $\ddot{x} = 0$ locus is determined by a pair (x_a, \dot{x}_a) satisfying (8.1.16), so $(x_a + k, \dot{x}_a)$ is on the right of it for $k > 0$. Since $g'' > 0$, g' is an increasing

function and therefore

$$rf'(\dot{x}_a) + g'(x_a + k) > rf'(\dot{x}_a) + g'(x_a) = 0.$$

From (8.1.15), this means that $\ddot{x} > 0$ and thus \dot{x} is increasing at $(x_a + k, \dot{x}_a)$. Similarly, for any point to the left of the $\ddot{x} = 0$ locus, we have $\ddot{x} < 0$ and the \dot{x} coordinate is decreasing. This gives the time-paths of points in the four regions created by the $\dot{x} = 0$ and $\ddot{x} = 0$ loci as indicated by the arrows in the above figure 8.4.⁸ These paths are vertical (i.e. $x = 0$) as they cross the $\dot{x} = 0$ locus and horizontal (i.e. $\dot{x} = 0$) as they cross the $\ddot{x} = 0$ locus.

The below figure 8.5 shows how to select among possible sets of boundary conditions for different T . ◁




Figure 8.5: Different possible time paths for different boundary conditions.

8.1.5 Higher-order ODEs

Higher order ODEs frequently appear in economic models. An general example of a second-order ODE is

$$\ddot{x} = f(t, x, \dot{x})$$

where $f(\cdot)$ is a given function and $x = x(t)$ the unknown function. We will be unable to find explicit solutions for more than a small number of rather peculiar functions but an example will make clear how to approach these cases.

Exl: Consider

$$\ddot{x} = \dot{x} + t.$$

⁸ Since the solution to the above differential equation with given suitable boundary conditions is unique, each point in the plane lies on exactly one path.

Note that we do not have an explicit dependency on x in the above ODE. We introduce a new variable $u = \dot{x}$ and rewrite

$$\dot{u} = u + t$$

which is a first-order ODE which has the general solution

$$u = Ae^t - t - 1 = \dot{x}$$

for an arbitrary constant A . Integrating this equation yields

$$x = \int (Ae^t - t - 1) dt = Ae^t - \frac{1}{2}t^2 - t + B$$

for another arbitrary constant B . ◁

This behaviour generalises. In principle, we should always be able to eliminate higher-order dependencies by integrating. Therefore, existence is guaranteed, but most of the time we will not be able to obtain closed-form solutions.

8.2 Elements of the theory of ordinary difference equations

Reading: (Chiang 1984, chapters 16, 17), (Sydsæter and Hammond 1995, chapter 20), (Hirsch and Smale 1974).

We develop the same taxonomy as in the previous section on ODEs. Since the ideas are really the same, we will proceed a little faster.

8.2.1 First-order (linear)

Def: Let $f(t, x)$ be a function defined for all positive integers $t \in \mathbb{N}_+$ and all real numbers $x \in \mathbb{R}$. A general **first order** OΔE relates the value of the function $f(\cdot)$ in each period to the value of the same function in the previous period. Symbolically

$$x_t = f(t, x_{t-1}). \tag{8.2.1}$$

Theorem 51. (*Existence and uniqueness*) Consider the OΔE (8.2.1) where $f(\cdot)$ is defined for all values of the variables. If x_0 is an arbitrary fixed number (i.e. the initial value of x is defined), then there exists a uniquely determined function x_t that is a solution of the equation.

Rem: Apparently, there is a different solution to the above problem for each different x_0 . Since the above x_0 is totally arbitrary, the problem (8.2.1) has infinitely many solutions.

Exl: Let us look at the special ('constant coefficient' or stationary) case $a \equiv a_t$ for all t and let $f(\cdot)$ be defined as a first-order linear OΔE for $t = 1, 2, \dots$

$$x_t = ax_{t-1} + b_t. \tag{8.2.2}$$

Then, for a given x_0 , we can compute the values of

$$\begin{aligned}x_1 &= ax_0 + b_1 \\x_2 &= ax_1 + b_2 = a(ax_0 + b_1) + b_2 = a^2x_0 + ab_1 + b_2 \\x_3 &= ax_2 + b_3 = a(a^2x_0 + ab_1 + b_2) + b_3 = a^3x_0 + a^2b_1 + ab_2 + b_3\end{aligned}$$

The pattern is clear. We can immediately conjecture that the solution might look like

$$x_t = a^t x_0 + \sum_{k=1}^t a^{t-k} b_k.$$

And indeed this hypothesis is correct. If we consider the special ('constant term') case $b \equiv b_t$ for all t , we know from the summation formula for finite geometric series (B.0.2) that we can rewrite this to

$$x_t = a^t \left(x_0 - \frac{b}{1-a} \right) + \frac{b}{1-a}. \quad (8.2.3)$$

provided, of course, that $a \neq 1$. \triangleleft

Rem: In general, of course, there is no reason for the constant a in (8.2.2) to be stationary. Therefore the more general form is

$$x_t = a_t x_{t-1} + b_t$$

which is the basis of, for instance, all compound interest and present discounted value calculations. We will, however, not go into the more general analysis here (cf. (Chiang 1984, chapter 17)).

Def: A solution x^* to (8.2.2) with the property that $x^* = x_t$ for all t is called **steady state** (or **stationary state** or **equilibrium state**).

Exl: If we plug the definition of a steady state into the example in (8.2.3), we obtain

$$x_t - x^* = a(x_{t-1} - x^*) \Leftrightarrow x_t - x^* = a^t(x_0 - x^*).$$

If $|a| < 1$, this means that a^t goes to zero as t becomes large, so we deduce that for $t \mapsto \infty$

$$x_t \mapsto x^* = \frac{b}{1-a}$$

which is a stable steady state. \triangleleft

Def: If the solution to an OΔE converges to a steady state as $t \mapsto \infty$, the equation is called **stable**. This is the case if $|a| < 1$. If, on the other hand, the solution to the equation tends away from the steady state as t increases, (the case for $|a| > 1$), we call the equation **explosive**. If x_t exhibits decreasing fluctuations, we say that the equation has **damped oscillations**, and if x_t exhibits increasing fluctuations, we say that the equation has **explosive oscillations**. This is illustrated in figure 8.6.

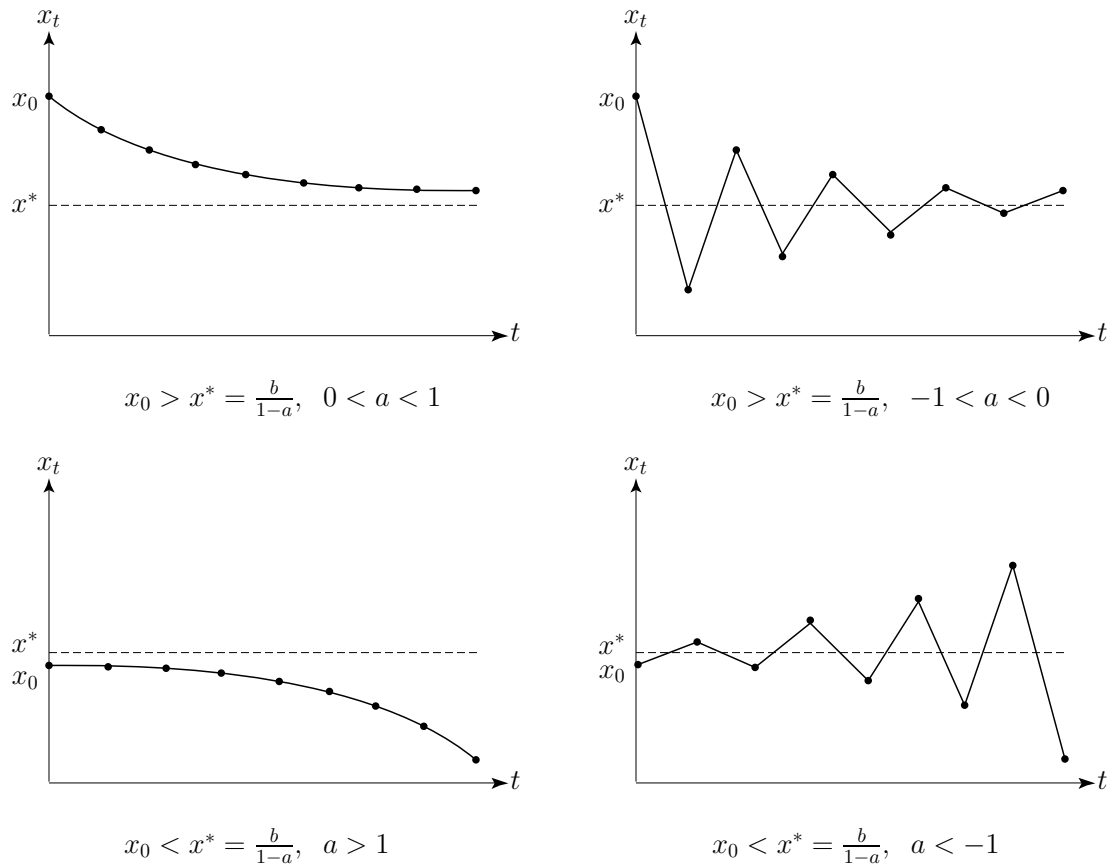


Figure 8.6: Stability in discrete time.

8.2.2 Second-order (linear) OΔEs

Def: Let $f(t, x)$ be a function defined for all positive integers $t \in \mathbb{N}_+$ and all real numbers $x \in \mathbb{R}$. A general **second order** OΔE relates the value of the function $f(\cdot)$ in each period to the value of the same function in the two previous periods. Symbolically

$$x_{t+2} = f(t, x_t, x_{t+1}). \quad (8.2.4)$$

Theorem 52. (*Existence and uniqueness*) Consider the OΔE (8.2.4) where $f(\cdot)$ is defined for all values of the variables. If x_0 and x_1 are arbitrary fixed numbers (i.e. the initial values of x are defined at $t = 0$ and $t = 1$), then there exists a uniquely determined function x_t that is a solution of the equation.

Again we look at a general linear form of a second-order OΔE

$$x_{t+2} + a_t x_{t+1} + b_t x_t = c_t. \quad (8.2.5)$$

Def: The above equation (8.2.5) for the case of $c_t = 0$ is called **homogeneous** OΔE.

Theorem 53. *The homogeneous OΔE*

$$x_{t+2} + a_t x_{t+1} + b_t x_t = 0$$

has the **general solution**

$$x_t = Au_t^{(1)} + Bu_t^{(2)}$$

where $u_t^{(1)}$ and $u_t^{(2)}$ are any two linearly independent solutions and A, B are arbitrary constant.

Def: Suppose we are able to find a **particular solution** u_t^* of (8.2.5).

Theorem 54. *Then the non-homogeneous OΔE (8.2.5) has the general solution*

$$x_t = Au_t^{(1)} + Bu_t^{(2)} + u_t^*$$

where $Au_t^{(1)} + Bu_t^{(2)}$ is the general solution of the associated homogeneous equation and u_t^* is a particular solution of the non-homogeneous equation.

The homogeneous case

Again we only look at the simpler case of the 2^{nd} -order OΔE with constant coefficients and develop a solution for those first. Consider the homogeneous equation

$$x_{t+2} + ax_{t+1} + bx_t = 0 \tag{8.2.6}$$

where the constant a is arbitrary and $b \neq 0$.

According to Theorem (54), we need to find the linearly independent solutions $u_t^{(1)}$ and $u_t^{(2)}$. To find these, we have to look at the solutions of the so called **characteristic equation**⁹

$$m^2 + am + b = 0 \tag{8.2.7}$$

by adopting the following strategy: The general solution of (8.2.6) is

1. For $\frac{1}{4}a^2 - b > 0$, the characteristic equation (8.2.7) has two different real roots, namely

$$x_t = Am_1^t + Bm_2^t, \text{ with } m_{1,2} = -\frac{1}{2}a \pm \sqrt{\frac{1}{4}a^2 - b}.$$

2. For $\frac{1}{4}a^2 - b = 0$, the characteristic equation (8.2.7) has two identical real roots, namely

$$x_t = (A + Bt)m^t, \text{ with } m = -\frac{1}{2}a.$$

3. For $\frac{1}{4}a^2 - b < 0$, the characteristic equation (8.2.7) has two complex roots for $\theta \in [0, \pi]$,

⁹ The characteristic equation does not drop from the heavens. Actually, if you consider the form $x_t = m^t$, then x_{t+1} turns out to be $m \cdot m^t = m^{t+1}$ and $x_{t+2} = m^{t+2}$. For $m \neq 0$ this gives the above characteristic equation (8.2.7).

namely¹⁰

$$x_t = Ar^t \cos(\theta t + \omega), \text{ with } r = \sqrt{b} \text{ and } \cos(\theta) = -\frac{a}{2\sqrt{b}}.$$

Exl:(Sydsæter and Hammond 1995, 749) Find the general solution of

$$x_{t+2} - 5x_{t+1} + 6x_t = 2t - 3. \quad (8.2.8)$$

The associated homogeneous equation is

$$x_{t+2} - 5x_{t+1} + 6x_t = 0.$$

The above procedure implies that there are particular solutions

$$u_t^{(1)} = 2^t, \text{ and } u_t^{(2)} = 3^t$$

which are linearly independent because

$$\begin{vmatrix} u_0^{(1)} & u_0^{(2)} \\ u_1^{(1)} & u_1^{(2)} \end{vmatrix} = \begin{vmatrix} 1 & 1 \\ 2 & 3 \end{vmatrix} = 3 - 2 = 1 \neq 0.$$

Therefore theorem (54) implies that

$$x_t = A2^t + B3^t$$

is a general solution. This suggests that

$$u_t^* = at + b$$

might be a solution. We use the “method of undetermined coefficients” in order to find a and b :

We rewrite

$$\begin{aligned} u_{t+1}^* &= a(t+1) + b \\ u_{t+2}^* &= a(t+2) + b \end{aligned}$$

which, inserted into (8.2.8) yields for all t

$$2a_t - 3a + 2b = 2t - 3.$$

this gives $2a = 2$, and $-3a + 2b = -3$. We conclude that $a = 1$, and $b = 0$. So $u_t^* = t$ is a particular solution of (8.2.8) and we find

$$x_t = A2^t + B3^t + t$$

to be the general solution. \triangleleft

¹⁰ If the characteristic equation contains complex roots, then the solution we are looking at involves oscillations. The number Ar^t is the amplitude at time t , $\frac{\theta}{2\pi}$ is the frequency, and ω is the phase of these oscillations. If $|r| > 0$ they are explosive, if $|r| < 0$ they are damped, and for $r = 1$, we have a stable orbit.

The non-homogeneous case

We go back to the more general non-homogeneous case of

$$x_{t+2} + a_t x_{t+1} + b x_t = c_t \quad (8.2.9)$$

From the discussion above, we know how to find $Au_t^{(1)}$ and $Bu_t^{(2)}$. The problem is how to find u_t^* . In some cases this is easy, e.g. for a constant $c_t = c$ or a c_t that is a linear combination of known terms. In general, however, we cannot solve this problem.

8.2.3 Stability

The intuition we develop here is closely related to definition (8.2.1). We know that if we can pin down suitable initial conditions for our (system of) OΔE(s) we will get a unique solution (of the system). Moreover, if we slightly change these initial conditions, we can look at the behaviour of the system and define that (i) it is **stable** if small changes in the initial conditions die out over time and (ii) **unstable** if our slight changes multiply over time and significantly change the solution.

A general solution of the second-order non-homogeneous OΔE we considered in (8.2.9) is

$$x_t = Au_t^{(1)} + Bu_t^{(2)} + u_t^*. \quad (8.2.10)$$

Def: The solution (8.2.10) is called **globally asymptotically stable** if the solution $Au_t^{(1)} + Bu_t^{(2)}$ of the associated homogeneous equation approaches zero as $t \mapsto \infty$, for all values of A and B . Then any solution of the equation approaches the particular solution u_t^* which is independent of the initial conditions. Therefore, the effect of perturbations of the initial conditions dies out over time.

The following proposition gives a reliable criterion to find out if an equation is stable.

Proposition 55. *Equation (8.2.9) is stable iff both roots of the characteristic equation $m^2 + am + b = 0$ have moduli¹¹ strictly less than 1. This is equivalent to saying that (8.2.9) is stable iff*

$$|a| < 1 + b, \quad b < 1.$$

As in the case of ODEs, we have to take to qualitative analysis in order to understand the dynamics of more complex systems. The technique of the phase diagram that was introduced in the last section is directly applicable.

¹¹ Moduli of real numbers $m \in \mathbb{R}$ are just their absolutes $|m|$. The modulus of a complex number $z \in \mathbb{C}$

$$z = a + bi = r(\cos \theta + i \sin \theta)$$

is given by the length of the vector z from the origin (its Euclidean norm)

$$\|z\| = \sqrt{a^2 + b^2}.$$

Exercises

Exc 8.1: Find the general solutions of

1. $\dot{x} = x + t$,
2. $\dot{x} - 3x = 5$,
3. $3\dot{x} + 2x + 16 = 0$,
4. $\dot{x} + 2x = t^2$.

Exc 8.2: Solve $\ddot{x} = \dot{x} + t$ for $x(0) = 1$, $\dot{x}(0) = 2$.

Exc 8.3: Draw phase diagrams for

1. $\dot{x} = x - 1$,
2. $\dot{x} = 24 - 2x$,
3. $\dot{x} = x^2 - 9$.

Exc 8.4: For $t = 0, 1, \dots$, find the general solution of $x_{t+2} - 5x_{t+1} + 6x_t = 2t - 3$.

Exc 8.5: The Solow model.¹² In a constant population of L identical workers, the per capita consumption is $c(t)$. Each worker derives a strictly concave utility $U(c)$ from consumption. We assume the marginal utility to be increasing without bound as consumption shrinks to zero: $\lim_{c \rightarrow 0} U'(c) = \infty$. A central planner is to select a consumption plan to maximise aggregate discounted utility

$$\max_{c(t)} \int_0^{\infty} e^{-\rho t} L U(c(t)) dt \quad (8.2.11)$$

taking into account the production possibilities of the economy. A single output is produced using capital K and labour L in the production function $F(K, L)$ which is homogeneous of degree 1. The output can be consumed or used to increase the capital stock. Thus

$$F(K, L) = Lc + K', \quad (8.2.12)$$

and $K(0) = K_0$.

1. We assume $F(0, L) = F(K, 0) = 0$. Why?
2. We assume F_K, F_L to be positive and decreasing. Why?
3. Now choose (c, K) to maximise (8.2.11) s.t.: (8.2.12).

¹² Wait with answering this question until the Solow model was covered in macro.

Chapter 9

Introduction to the calculus of variation

Reading: (Kamien and Schwartz 1991, chapters).

From an optimisation theory point of view, the *calculus of variation* (which is the subject of this section) is analogous to classical calculus. That is, it identifies solutions given that all functions describing the problem are well-behaved and all solution candidates are strictly inside the feasible region. The task of the calculus of variation is to select the optimal path for the considered problem among the set of admissible paths. *Optimal control theory* (which is the subject of the next two sections) and its central *Maximum Principle* were developed as a generalisations of the calculus of variations by L.S. Pontryagin in the late 50's. It is applicable in the above circumstances for continuous time problems. It can also accommodate complications such as boundary solutions (in that respect it may be viewed as analogous to Kuhn-Tucker theory).

Dynamic programming (to which we will occasionally refer) was developed by R. Bellman about the same time for discrete time problems.¹ It provides an intuitive way of setting up the optimal control problem using a *value function*.

As with all optimisation problems, we have to account for three questions:

1. existence,
2. necessary conditions, and
3. sufficient conditions.

In this and the following two sections we (or really Maurice Obstfeld) will exclusively address necessity. In this introduction we will look at the motivation for the calculus of variation and discuss a number of key questions that should help you survive macro.² That means we will not at all address (1) and (3) in the above enumeration—you will have to make do with what you know from the last term.

¹(Sundaram 1996, chapters 11, 12) develops dynamic programming out of the micro we did in the first term. A great reference for Macro-applications is (Ljungqvist and Sargent 2000).

² Most of it is taken from (Kamien and Schwartz 1991).

9.1 Discounting

If A pounds are invested at r percent per year, the amount would grow to $A + rA = (1 + r)A$ after one year and after t years, we'd have

$$X \equiv A(1 + r)^t.$$

If the interest is compounded m times per year, then the rate per period is $\frac{r}{m}$. The amount A would grow to $A(1 + \frac{r}{m})^m$ after one year and after t years, we get

$$X = A(1 + \frac{r}{m})^{mt}.$$

Letting m go to infinity, we obtain for continuous compounding in the limit³

$$\lim_{m \rightarrow \infty} (1 + \frac{r}{m})^{mt} = e^{rt}.$$

Therefore, if, on the one hand, we invest A pounds at an annual rate of r under continuous compounding, we'd get after t years the amount

$$X = Ae^{rt}.$$

If, on the other hand, we want to know to what amount X an investment of B grows to after t years under continuous compounding with interest rate r , we know from the above that $B = Xe^{-rt}$ and therefore obtain the *present value* (PV) of $B\pounds$ as $X = Be^{-rt}$.

Exl: The compounded value of $\pounds 10$ invested at $r = 10\%$ over $t = 10$ years under continuous compounding is $X = 10e^{1(10)} \approx \pounds 27.18$. \triangleleft

Exl: The PV of a terminal amount of $\pounds 10$ invested at $r = 10\%$ over $t = 10$ years under continuous compounding is $X = 10e^{-1(10)} \approx \pounds 3.68$. That is, an amount of $\pounds 3.68$ will grow to $\pounds 10$ after 10 years. \triangleleft

9.2 Depreciation

If a stock K depreciates at a constant rate $b > 0$ and is not restocked, then per period the stock reduces by $\dot{K} = -Kb$. Formulated as a *growth-rate* this is

$$\frac{\dot{K}}{K} = -b.$$

We know from the previous section that the solution to this differential equation is

$$K(t) = K(0)e^{-bt}.$$

³ To understand this result it helps to remember that e is actually defined as $\lim_{m \rightarrow \infty} (1 + \frac{1}{m})^m$.

9.3 Calculus of variation: Derivation of the Euler equation

In principle, we will look at dynamic problems like

$$\begin{aligned} \max_{x_t} \quad & \sum_{t=1}^{\infty} F(t, x_t, x_{t-1}) \\ \text{s.t.:} \quad & x(t) \geq 0, t = 1, \dots \end{aligned}$$

which we will interpret as a succession of a firm's optimal production choices that depends on previous levels of output; therefore x_t is the output level and $F(t, x_t, x_{t-1})$ is the profit associated with that particular output at time t ; it also depends on last period's output. We know from last term that the solution to our static micro problem (if it exists) is a particular *number* x^* . Now the solution to our dynamic problem is a *function* $x(t)$ that depends on at least one time derivative: a differential equation. Looking at the same problem in its continuous form,⁴ we can rewrite the above problem as

$$\begin{aligned} \max_{x(t)} \quad & \int_0^{\infty} F(t, x(t), \dot{x}(t)) dt \\ \text{s.t.:} \quad & x(t) \geq 0; x(0) \text{ given.} \end{aligned}$$

We assume $F(\cdot)$ to be continuous in all of its arguments and to be C^1 w.r.t x and \dot{x} .

Although this is the correct statement of what is to come, we start by looking at an easier problem first.^{5,6}

$$\begin{aligned} \max_{x(t)} \quad & \int_{t_0}^{t_1} F(t, x(t), \dot{x}(t)) dt \\ \text{s.t.:} \quad & x(t_0) = x_0, x(t_1) = x_1. \end{aligned} \tag{9.3.2}$$

We know that the solution to this problem is a differential equation $x(t)$ and we know the initial and terminal values of this function. But we know nothing else.

Def: The class of *admissible functions* among which we are looking for a maximum is defined as the set of *all* continuously differentiable functions on $[t_0, t_1]$.

We are simply looking for any function that is C^1 and goes through both boundary points. Now suppose that $x^*(t)$, $t \in [t_0, t_1]$ provides a maximal solution to (9.3.2) and satisfies our boundary conditions. Let $x(t)$ be some arbitrary other admissible function and let

$$h(t) \equiv x^*(t) - x(t)$$

be the gap between the two for each t . Since both functions are admissible, $h(t_0) = h(t_1) = 0$.

⁴ Again, notice our convention to switch to the functional notation $x(t)$ when looking at continuous variables.

⁵ We will solve the infinite horizon problem below using transversality conditions but we will not look at inequality constraints at all. If you are interested in how to handle these, consult (Kamien and Schwartz 1991, chapters 12, 14), or (Chiang 1989, chapter 6).

⁶ For future reference, here is the equivalent statement as an optimal control problem. Let $u(t) = \dot{x}(t)$, then

$$\begin{aligned} \max_{x(t)} \quad & \int_{t_0}^{t_1} F(t, x(t), u(t)) dt \\ \text{s.t.:} \quad & \dot{x}(t) = u(t), x(t_0) = x_0. \end{aligned} \tag{9.3.1}$$

We call a *deviation admissible* if $x^* + h$ is admissible. This implies that for any constant a , the function $y(t) = x^*(t) + ah(t)$ will also be admissible. Now holding x^* and h both fixed, we compute the integral in (9.3.2) for $y(t, a)$. This is illustrated below.




Figure 9.1: An optimal time path x^* and a deviation from it $x^* + h$.

$$\begin{aligned} g(a) &\equiv \int_{t_0}^{t_1} F(t, y(t, a), \dot{y}(t, a)) dt \\ &= \int_{t_0}^{t_1} F(t, x^*(t) + ah(t), \dot{x}^*(t) + ah'(t)) dt \end{aligned}$$

Since x^* maximises (9.3.2), $g(a)$ must assume its maximum at $a = 0$. But this implies that $g'(a = 0) = 0$ which we exploit (by applying the chain and Leibnitz' (B.0.4) rules) to establish

$$\begin{aligned} g'(0) &= \int_{t_0}^{t_1} \left[\frac{dF(t, x^*(t), \dot{x}^*(t))}{da} \right] dt \\ &= \int_{t_0}^{t_1} [F_x(t, x^*(t), \dot{x}^*(t)) h(t) + F_{\dot{x}}(t, x^*(t), \dot{x}^*(t)) h'(t)] dt \\ &= 0. \end{aligned}$$

Integrating the second term by parts (setting $f = F_{\dot{x}}$ and $g' = h'(t)dt$ in (B.0.3)) yields an easier form (assuming that the *total* derivative $\frac{dF_{\dot{x}}(\cdot)}{dt}$ exists)

$$\int_{t_0}^{t_1} F_{\dot{x}}(\cdot) h'(\cdot) dt = F_{\dot{x}}(\cdot) h \Big|_{t_0}^{t_1} - \int_{t_0}^{t_1} h(t) \left(\frac{dF_{\dot{x}}(\cdot)}{dt} \right) dt.$$

Substituting this and our initial and terminal conditions from (9.3.2), into the above condition $g'(0) = 0$ we obtain

$$g'(0) = \int_{t_0}^{t_1} \left(F_x(t, x^*(t), \dot{x}^*(t)) - \frac{dF_{\dot{x}}(t, x^*(t), \dot{x}^*(t))}{dt} \right) h(t) dt = 0. \quad (9.3.3)$$

Since this condition (9.3.3) holds if x^* maximises (9.3.2), it must hold for *every* continuously differentiable function h that is zero at the beginning and the end. It certainly also holds if the

coefficient of h is zero for every t . This insight gives us

$$F_x(t, x^*(t), \dot{x}^*(t)) = \frac{dF_{\dot{x}}(t, x^*(t), \dot{x}^*(t))}{dt}, \text{ for } t_0 \leq t \leq t_1, \quad (9.3.4)$$

or, suppressing the variables $(t, x^*(t), \dot{x}^*(t))$, we get the simpler notation

$$F_x = \frac{dF_{\dot{x}}}{dt}, \text{ for } t_0 \leq t \leq t_1. \quad (9.3.5)$$

This equation is called the *Euler equation* corresponding to problem (9.3.2). It may be viewed as a generalisation of our standard necessary condition $f'(x^*) = 0$ to maximise the function $f(x)$.⁷ A proof that (9.3.4) is a necessary condition for all $t \in [t_0, t_1]$ can be found in (Kamien and Schwartz 1991, 16f). Computing the *total* derivative of the time partial $\frac{dF_{\dot{x}}}{dt} = F_{\dot{x}t} + F_{\dot{x}x}\dot{x} + F_{\dot{x}\dot{x}}\ddot{x}$, we can rewrite the above into an alternative version of the Euler equation that is often more convenient than (9.3.4)

$$F_x = F_{\dot{x}t} + F_{\dot{x}x}\dot{x} + F_{\dot{x}\dot{x}}\ddot{x}, \text{ for } t_0 \leq t \leq t_1. \quad (9.3.6)$$

We conclude that the Euler equation is a second order differential equation for $x(t)$ that has to be solved for the initial and terminal conditions in (9.3.2). As we know from the previous section there may well exist no closed form solution to such equations.

More realistic problem settings than the one above will give rise to a whole system of Euler equations. Solution candidates to the (system of) Euler equation(s) are called *extremals*—these are analogous to the solutions to $f'(x) = 0$ in classical calculus. As you may suspect, there are also corresponding sufficient conditions: If $F(t, x, \dot{x})$ is concave in (x, \dot{x}) then an admissible function that solves the Euler equation(s) solves our problem. But we will not investigate these issues here any further. If you're interested, please take a look at (Kamien and Schwartz 1991, 40ff) or (Chiang 1989, 79ff).

9.4 Solving the Euler equation

As we have stated rather casually above, we do not know whether there are closed form solutions to the general form of the Euler equation(s) or not. If we cannot find a useful solution by analytical means, we may want to employ the qualitative analysis technique of phase diagrams. There is a whole collection of solutions, however, for certain standard forms of Euler equations. We will investigate the most common of these now. Again, most of this is taken from (Kamien and Schwartz 1991, 28ff). Consult the examples and exercises there if you need some practice. If you know how to solve the Euler equation in the cases below, you are safe; do not expect any more horrible things to show up in you macro problem sets.

1. $F = F(t, \dot{x})$.

⁷ You may want to convince yourself that if $dF_{\dot{x}}/dt = 0$, then (9.3.4) reduces to the standard calculus condition.

Since F does not depend on x , the Euler equation (9.3.5) reduces to

$$\frac{dF_{\dot{x}}}{dt} = 0,$$

from which follows that $F_{\dot{x}} = \text{const}$. This is a first order differential equation in (t, \dot{x}) only and we can solve it employing the techniques we studied in the previous section.

2. $F = F(x, \dot{x})$.

Since F does not depend on t , the Euler equation (9.3.6) reduces to

$$F_x - F_{\dot{x}x}\dot{x} - F_{\dot{x}\dot{x}}\ddot{x} = 0,$$

for which \dot{x} is an *integrating factor*. Therefore, multiplying this Euler equation through by \dot{x} yields an exact differential, namely,

$$\dot{x} [F_x - F_{\dot{x}x}\dot{x} - F_{\dot{x}\dot{x}}\ddot{x}] = \frac{d(F - \dot{x}F_{\dot{x}})}{dt} = 0.$$

The middle term being equal to zero implies that

$$F - \dot{x}F_{\dot{x}} = \text{const}, \text{ for } t_0 \leq t \leq t_1$$

which, being a first order differential equation, we can happily solve.

3. $F = F(\dot{x})$.

Since F does only depend on \dot{x} , the Euler equation (9.3.6) reduces to

$$F_{\dot{x}\dot{x}}\ddot{x} = 0.$$

Thus, along the extremal at each t , either $F_{\dot{x}\dot{x}}(\dot{x}) = 0$ or $\ddot{x}(t) = 0$. In the latter case, this indicates that the extremal is linear of the form $x(t) = c_1t + c_2$. In the former case, either $F_{\dot{x}\dot{x}} \equiv 0$ or else \dot{x} is constant (which is the case we just considered). If F is linear in \dot{x} , the Euler equation is an identity and any x satisfies it trivially.

We conclude that if the integrand F depends solely on \dot{x} but is not linear, then the graphs of extremals are straight lines. The boundary conditions determine the constants. (This is a heavy hint towards the answer of exercise 9.6.)

4. $F = F(t, x)$.

This Euler equation is not a differential equation. Therefore, we do not (really) look at a dynamic problem.

5. $F = A(t, x) + B(t, x)\dot{x}$.

In this *separable* case, the Euler equation (9.3.6) is either

$$A_x + B_x\dot{x} = B_t + B_x\dot{x},$$

that is $A_x(t, x) = B_t(t, x)$ which is not a differential equation, or

$$A_x(t, x) \equiv B_t(t, x)$$

which is satisfied by any function x . We do not go into this any further here, but be assured that something called the *Integrability Theorem* of differential equations says that in such cases the value of the integral only depends on its endpoints—the path is irrelevant. Consequently in this case, any feasible path is optimal.

9.5 Transversality condition

When we look at problems with no fixed end date (or no end date at all) we cannot apply the above solution method. Regardless, we can set the problem up as

$$\begin{aligned} \max_{x(t)} \int_{t_0}^{t_1} F(t, x(t), \dot{x}(t)) dt \\ \text{s.t.: } x(t_0) = x_0 \text{ given} \end{aligned} \tag{9.5.1}$$

while leaving t_1 unspecified. The intuition of what we are doing now is captured in the following picture: Rather than pasting a curve in between two points we now have a single initial point and a vertical line as the terminal condition where we want to paste a smooth curve in between.




Figure 9.2: The free horizon problem.

We proceed similarly to before. Assume we know that $x^*(t)$ is an optimal path on $t_0 \leq t \leq t_1$ and there exists a comparison function $x(t)$ on $t_0 \leq t \leq t_1 + \delta t_1$ with a (slightly) different domain.⁸ As always, x^* and x are taken to be C^2 and to satisfy the initial condition. We define a function $h(t)$ as the difference between the two functions

$$h(t) \equiv x^*(t) - x(t), \text{ for all } t_0 \leq t \leq \max(t_1, t_1 + \delta t_1)$$

and, as before, we know $h(0) = 0$.

⁸ Yet we consider the functions to have the same domains. Therefore, we have to extend whichever domain is shorter to the longer one and define the value of the ‘shorter’ function to be of the value indicated by the slope of the last defined point.

Def: We define the *distance* between the optimum and the comparison function to be

$$\|x^*(t) - x(t)\| \equiv \max_t |h(t)| + \max_t |h'(t)| + |\delta t_1| + |x(t_1 + \delta t_1) - x^*(t_1)|.$$

This is different from the standard Euclidean metric that we studied last term. We indicate in figure 9.3 below what idea this is supposed to capture.

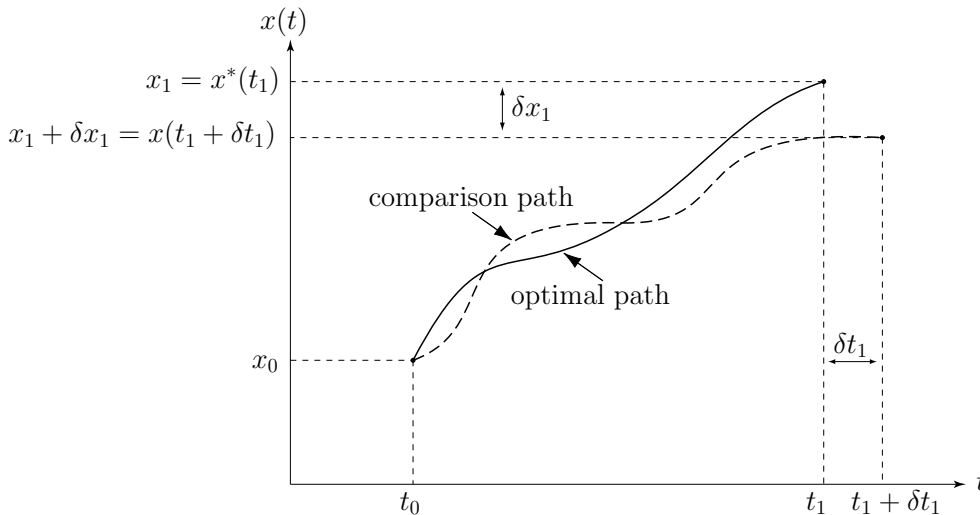


Figure 9.3: Two distance paths.

According to the above definition, two C^2 functions are close if, for each point of the (extended) domain, both slopes are similar, their values are close and their terminal values are similar. As before, we define a function $g(a)$

$$g(a) \equiv \int_{t_0}^{t_1 + \delta t_1} F(t, x^*(t_1) + ah(t), \dot{x}^*(t) + ah'(t)) dt$$

that assumes a maximum at $a = 0$ so that $g'(0) = 0$. Leibnitz' rule and integration by parts (B.0.4, B.0.3) gives

$$g'(0) = F(t_1, x^*(t_1), \dot{x}^*(t_1)) \delta t_1 + F_{\dot{x}}(t_1, x^*(t_1), \dot{x}^*(t_1)) h(t_1) + \int_{t_0}^{t_1} h \left(F_x - \frac{dF_{\dot{x}}}{dt} \right) dt \quad (9.5.2)$$

where the integrand is evaluated at the optimal $(t, x^*(t), \dot{x}^*(t))$. As depicted above, we call the difference of the values of the two functions at their terminal points

$$\delta x_1 \equiv x(t_1 + \delta t_1) - x^*(t_1).$$

Now if the optimal function x^* and the comparison function $x = x^* + h$ are *close* in the above sense, we have approximately

$$\delta x_1 \approx x(t_1) - x^*(t_1) + \dot{x}^*(t_1) \delta t_1 = h(t_1) + \dot{x}^*(t_1) \delta t_1.$$

We rearrange to

$$h(t_1) \approx \delta x_1 - \dot{x}^*(t_1)\delta t_1 \quad (9.5.3)$$

and by that essentially construct a parallelogram as shown below.

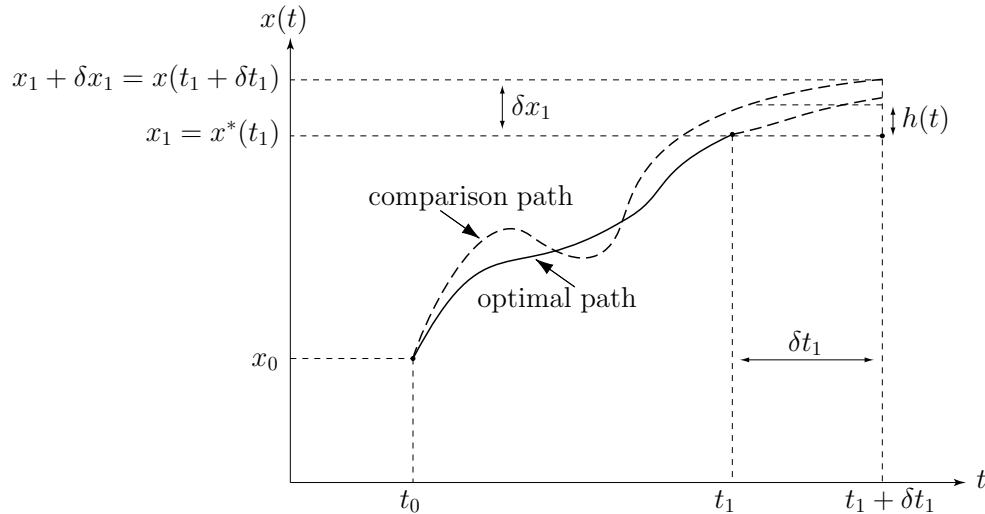


Figure 9.4: The distance parallelogram.

Collecting terms and substituting (9.5.3) into (9.5.2) we get the expression we are after

$$\int_{t_0}^{t_1} h \left(F_x - \frac{dF_{\dot{x}}}{dt} \right) dt + F_{\dot{x}} \Big|_{t_1} \delta x_1 + (F - \dot{x}F_{\dot{x}}) \Big|_{t_1} \delta t_1 = 0. \quad (9.5.4)$$

Since the comparison curve *could* terminate at exactly the same point that x^* does, with $\delta t_1 = 0$ and $\delta x_1 = 0$, it follows that

$$\int_{t_0}^{t_1} h \left(F_x - \frac{dF_{\dot{x}}}{dt} \right) dt = 0$$

must hold for all admissible functions satisfying $h(t_0) = h(t_1) = 0$. Therefore we again know that the Euler equation

$$F_x = \frac{dF_{\dot{x}}}{dt}$$

is satisfied and (9.5.4) reduces to

$$F_{\dot{x}} \Big|_{t_1} \delta x_1 + (F - \dot{x}F_{\dot{x}}) \Big|_{t_1} \delta t_1 = 0.$$

Supposing t_1 to be free but $x(t_1) = x_1$ fixed, we have $\delta x_1 = 0$ and we can rewrite the above into the more usual form of

$$(F - \dot{x}F_{\dot{x}}) \Big|_{t_1} \delta t_1 = 0 \text{ for all } \delta t_1.$$

Therefore we require

$$F - \dot{x}F_{\dot{x}} = 0 \text{ at } t_1 \text{ if } t_1 \text{ is free.}$$

This expression is called the *transversality condition* for (9.5.1). Together with the initial value

it determines the two constants of integration in the Euler equation and the terminal time t_1 .

9.6 Infinite horizon problems

Consider a problem of the following form

$$\begin{aligned} \max_{x(t)} \int_{t_0}^{\infty} e^{-rt} F(x(t), \dot{x}(t)) dt \\ \text{s.t.: } x(0) = x_0. \end{aligned} \quad (9.6.1)$$

Notice that since we do not have an explicit dependency on time we are looking at an *autonomous* problem⁹—we will not look at anything more complicated. The problem with infinite horizons is that there may be no necessary transversality condition. Even if the transversality condition holds, it may be of little help in determining the constants of integration. Be that how it may, without much ado we argue that in problems such as the one above that depend only on time through the discount rate, it *may* be reasonable to expect that the solution $x(t)$ tends towards a *steady state* x_s for which, by definition, $\dot{x} = \ddot{x} = 0$.¹⁰ We solve for (the number) x_s by setting $\dot{x} = \ddot{x} = 0$ in the Euler equation (9.3.6) and obtain the implicit condition

$$F_x(x_s, 0) + rF_{\dot{x}}(x_s, 0) = 0.$$

Given our existence assumption on the steady state, we can define an additional boundary condition on problem (9.6.1) by writing

$$\lim_{t \rightarrow \infty} x(t) = x_s.$$

For example (taken from (Kamien and Schwartz 1991, 89))

$$\begin{aligned} \min_{x(t)} \int_{t_0}^{\infty} e^{-rt} (x^2 + ax + b\dot{x} + c\dot{x}^2) dt \\ \text{s.t.: } x(0) = x_0, \end{aligned}$$

where $c > 0$ and $r > 0$ give rise to the Euler equation

$$\ddot{x} - r\dot{x} - \frac{x}{c} = \frac{a + rb}{2}$$

which has the particular solution

$$x_s = -\frac{a + rb}{2}.$$

The general solution to the Euler equation is therefore

$$x(t) = Ae^{w_1 t} + Be^{w_2 t} + x_s$$

where $w_1, w_2 = \frac{r}{2} \pm \sqrt{(\frac{r}{2})^2 + \frac{1}{c}}$ are the (real) roots of the associated characteristic equation.

⁹ Of course, this is a lie: The term e^{-rt} represents an explicit time dependency. But since it is just through the constant discount term, we can still consider the problem autonomous because the Euler equation will be truly autonomous.

¹⁰ Notice that the steady state need *not* be approached in finite time.

Exercises

All exercises are taken from the first 5 chapters of (Kamien and Schwartz 1991). You should be able to solve them even if you did not yet cover optimisation in macro.

Exc 9.1: A firm has to fulfill an order for B units of its product to be delivered at time T . It seeks a production schedule for filling this order at the specified delivery date at minimum cost. Unit production cost rises linearly with the production rate and the unit cost of holding inventory (per time unit) is constant.

1. Think about the problem and try to determine (the behaviour of) all variables involved.
2. Set up the firm's cost minimisation problem.
3. Find a production plan that is feasible (you do not have to show it is minimal).

Exc 9.2: Find the Euler equation and its solutions for

$$\begin{aligned} \max_{x(t)} \int_1^2 [x(t) + tx'(t) - (\dot{x}(t))^2] dt \\ \text{s.t.: } x(1) = 3, x(2) = 4. \end{aligned}$$

Exc 9.3: Find the extremals for

$$\begin{aligned} \max_{x(t)} \int_0^T \{c_1 [\dot{x}(t)]^2 + c_2 x(t)\} dt \\ \text{s.t.: } x(0) = 0, x(T) = B. \end{aligned}$$

Exc 9.4: Find the shortest distance in the plane between two points (a,A) and (b,B) . To state the problem more formally, recall that the square of the length of the hypotenuse of a right triangle is equal to the sum of squares of lengths of the other two sides. Thus a small distance ds can be expressed as

$$ds = \sqrt{(dt)^2 + (dx)^2} = \sqrt{1 + \dot{x}(t)^2} dt.$$

Thus our problem is to minimise the length of the path or

$$\begin{aligned} \min_{x(t)} \int_a^b \sqrt{1 + \dot{x}(t)^2} dt \\ \text{s.t.: } x(a) = A, x(b) = B. \end{aligned}$$

Chapter 10

Introduction to discrete Dynamic Programming

Reading: (Bertsekas 2000).

This chapter is an (barely started) attempt of introducing dynamic programming in a general decision theoretic context. The idea is based on K. Levine's *Discrete Programming Notes* of which it is in the current state little more than an annotated copy. In the future at least I'd like to integrate a user friendly form of Denardo (1967).

10.1 Assumptions

There is a single player who takes a single action $\alpha \in A$ at every stage of the game. All information is aggregated into the single state variable $y \in Y$. The transition dynamics of y are determined by the transition probability $\pi(y^*|y, \alpha)$. The set of states which can be reached with positive probability from some given state y is defined as

$$S(y) = \{y' | \exists \alpha \text{ with } \pi(y'|y, \alpha) > 0\}.$$

Ass: We assume that

1. A is a compact subset of some finite dimensional space;
2. Y is countable;
3. $\pi(y'|y, \alpha)$ is continuous in α , and
4. $S(y)$ is finite.

There are two main cases of interest; case (a), where Y and A are finite, and case (b), where Y is a tree and only immediate successors have positive probability.

Preferences for this decision problem are given by a period utility $u(\alpha, y)$ which is assumed to be additively separable over time (and states of nature) with discount factor $0 \leq \delta < 1$.

We further assume that

5. u is bounded by \bar{u} and continuous in α .

10.2 Definitions

Def: The finite **history** of the game is denoted $h = (y_1, y_2, \dots, y_t)$. For any given history h , we can recover:

- i. the length of the history at $t(h) = t$,
- ii. the final state $y(h) = y_t$,
- iii. the history through the previous period $h - 1 = (y_1, y_2, \dots, y_{t-1})$, and
- iv. the initial state $y_1(h)$.

Def: Histories are naturally **ordered** according to whether or not a history can logically follow from one another. If a history h' follows from h we denoted this by $h' \geq h$.

Def: We say that a history is **feasible** if $y_\tau \in S(y_{\tau-1})$; the set of histories that is not feasible has probability zero. The set of all feasible histories is denoted by H .

Rem: Since we have assumed $S(y)$ to be finite, it follows that H is countable.

Def: An **action** α is a map from histories to actions $\alpha : H \mapsto A$.

Def: The object of choice is a **strategy** σ which is a map from all histories to actions for each stage $\sigma : H \mapsto \Sigma$. The space of all strategies the player can choose from is $\Sigma \equiv A^\infty$.

Def: A strategy is called **strong Markov** if $\sigma(h) = \sigma(h')$ if $y(h) = y(h')$, that is if actions are determined entirely by the state. Any strong Markov strategy is equivalent to a map $\sigma : Y \mapsto A$. Given a strategy we can define the probabilities of histories by

$$\pi(h|y_1, \sigma) = \begin{cases} \pi(y(h)|y(h-1), \sigma(h-1))\pi(h-1|y_1, \sigma) & \text{if } t(h) > 1 \\ 1 & \text{if } t(h) = 1 \text{ and } y_1(h) = y_1 \\ 1 & \text{if } t(h) = 1 \text{ and } y_1(h) \neq y_1. \end{cases}$$

We may also—for any given initial state and strategy—compute the **expected average present value utility** as

$$V(y_1, \sigma) \equiv (1 - \delta) \sum_{h \in H} \delta^{t(h)-1} u(\sigma(h), y(h)) \pi(h|y_1, \sigma).$$

10.3 The problem statement

The problem is to

$$\begin{aligned} & \max_{\sigma} V(y_1, \sigma) \\ & \text{s.t.: } \sigma \in \Sigma. \end{aligned} \tag{10.3.1}$$

Def: A (not *the*) **value function** is any map $v : Y \mapsto \mathbb{R}$ bounded by \bar{u} .

Essential to the study of dynamic programming are two infinite dimensional objects: strategies and value functions. These naturally lie in two different spaces. Strategies naturally lie \mathbb{R}^∞ , the space of infinite sequences of numbers endowed with the product topology. Value functions naturally lie in ℓ_∞ , the space of bounded functions in the sup norm. From the fact that the space of strategies is compact and utility continuous, it follows that

Lemma 56. *A solution to (10.3.1) exists.*

This enables us to define *the* value function as

$$v(y_1) \equiv \max_{\sigma \in \Sigma} V(y_1, \sigma). \quad (10.3.2)$$

10.4 The Bellman equation

We define a map $T : \ell_\infty \mapsto \ell_\infty$ by $w' = T(w)$ if

$$w'(y_1) = \max_{\alpha \in A} (1 - \delta)u(\alpha, y_1) + \delta \sum_{y'_1 \in S(y_1)} \pi(y'_1 | y_1, \alpha)w(y'_1). \quad (10.4.1)$$

Def: We refer to the operator T as the **Bellman operator**.

Lemma 57. *The value function is a fixed point of the Bellman equation $T(v) = v$.*

Lemma 58. *The Bellman equation is a contraction mapping*

$$\|T(w) - T(w')\| \leq \delta \|w - w'\|.$$

Corollary: The Bellman equation has a unique solution.

Rem: The unique solution to the Bellman equation is the value function.

Lemma 59. *There is a strong Markov optimum that may be found from the Bellman equation.*

Proof. We define the strong Markov plan in the obvious way, and show recursively that it yields a present value equal to the value function

$$\begin{aligned} v(y_1(h)) &= \\ & (1 - \delta) \sum_{t(h) < T} \delta^{t(h)-1} \pi(h | y_1(h), \sigma) u(\sigma(h), y(h)) \\ & + (1 - \delta) \sum_{t(h) = T} \delta^T \pi(h | y_1(h), \sigma) v(h). \end{aligned}$$

□

Chapter 11

Deterministic optimal control in continuous time

Reading: (Obstfeld 1992) and (Kamien and Schwartz 1991). Alternative sources are (Chiang 1989) and (Seierstad and Sydsæter 1987). A derivation of the Maximum Principle is also contained in (Sundaram 1996).

The text that forms chapters (11) and (12) is Obstfeld, M., *Dynamic Optimisation in Continuous-Time Economic Models (A Guide for the Perplexed)*, NBER, 1992. In addition to being concise and very well written, it has the added advantage to be downloadable for free (from <http://elsa.berkeley.edu/users/obstfeld/>) while Kamien-Schwartz represents quite a financial investment. Here I reprint the original Obstfeld paper with some annotations, additional explanations and pictures.¹

11.1 Theory I

The assumption that economic activity takes place continuously is a convenient abstraction in many applications. In others, such as the study of financial-market equilibrium, the assumption of continuous trading corresponds closely to reality. Regardless of motivation, continuous-time modeling allows application of a powerful mathematical tool, the theory of optimal dynamic control.

The basic idea of optimal control theory is easy to grasp—indeed it follows from elementary principles similar to those that underlie standard static optimisation problems. The purpose of these notes is twofold. First, I present intuitive derivations of the first-order necessary conditions that characterise the solutions of basic continuous-time optimisation problems.² Second, I show

¹ I am grateful to Maurice Obstfeld, University of California at Berkeley, for allowing me to use his paper in this way. Since I typed in the paper manually, I guarantee for the addition of typos—for which I take full responsibility. My annotations are usually given in footnotes and marked by a double dagger.[‡]

² When the optimisation is done over a finite time horizon, the usual second-order sufficient conditions generalise immediately. (These second-order conditions will be valid in all problems examined here.) When the horizon is infinite, however, some additional ‘terminal’ conditions are needed to ensure optimality. I make only passing reference to these conditions below, even though I always assume (for simplicity) that horizons are infinite. Detailed treatment of such technical questions can be found in some of the later references.

why very similar conditions apply in deterministic and stochastic environments alike.

A simple unified treatment of continuous-time deterministic and stochastic optimisation requires some restrictions on the form that economic uncertainty takes. The stochastic models I discuss below will assume that uncertainty evolves continuously according to a type of process known as an *Itô* (or *Gaussian diffusion*) process. Once mainly the province of finance theorists, Itô processes have recently been applied to interesting and otherwise intractable problems in other areas of economics, for example, exchange-rate dynamics, the theory of the firm, and endogenous growth theory. Below, I therefore include a brief and heuristic introduction to continuous-time stochastic processes, including the one fundamental tool needed for this type of analysis, Itô's chain rule for stochastic differentials. Do not be intimidated: the intuition behind Itô's Lemma is not hard to grasp, and the mileage one gets out of it thereafter truly is amazing.

The basic problem to be examined takes the form³

$$\begin{aligned} \max_{c(t)} \int_0^{\infty} e^{-\delta t} U[c(t), k(t)] dt \\ \text{s.t.: } \dot{k}_t = G[c(t), k(t)] \end{aligned} \tag{11.1.1}$$

where $k(0)$ is given and $U(c, k)$ is a strictly concave function and $G(c, k)$ is concave. In practice there may be some additional inequality constraints on c and/or k ; for example, if c stands for consumption, c must be non-negative. While I will not deal in any detail with such constraints, they are straightforward to incorporate.⁴ In general, c and k can be vectors, but I will concentrate on the notationally simpler scalar case. I call c the *control variable* for the optimisation problem and k the *state variable*. You should think of the control variable as a flow—for example, consumption per unit time—and the state variable as a stock—for example, the stock of capital, measured in units of consumption.

³ $\delta \geq 0$ is the rate of time preference (or the subjective discount rate).[‡]

⁴ The best reference work on economic applications of optimal control is still (Arrow and Kurz 1970).

Here is a *rough* guide of how to choose the control and state variables:[‡]

The *state* variables

1. If there is only one asset or consumption good, then the number of state variables equals the number of Euler equations.
2. The state variable enters the objective.
3. The change of each state variable is determined by a first order differential equation, the so called *state* (or *transition*) equation.
4. State variables must be predetermined at any given point in time, that is, they are *not* an object of choice. Therefore, only the *rhs* of the rules of consumption, saving, or accumulation are candidates for state variables.
5. State variables can only use sufficient statistics.
6. Avoid using more than one state variable like hell.

The *control* variables

1. Control variables are always exogenous, i.e. in contrast to the state variables they are subject to discretionary choice.
2. They influence the state variable(s) (a) directly through their own value and (b) indirectly through their impact on the evolution of the state variables that enters the objective.
3. Look first at the typical rules of consumption (c_t) saving (b_t) and accumulation (k_t, k_{t+1}).

Notice that the number of state variables may be greater or smaller than the number of control variables.

The problem set out above has a special structure that we can exploit in describing a solution. In the above problem, planning starts at time $t = 0$. Since no exogenous variables enter (11.1.1), the maximised value of (11.1.1) depends only on $k(0)$, the predetermined initial value of the state variable. In other words, the problem is stationary, i.e., it does not change in form with the passage of time.⁵ Let's denote this maximised value by $J[k(0)]$, and call $J(k)$ the *value function* for the problem. If $\{c^*(t)\}_{t=0}^{\infty}$ stands for the associated optimal path of the control and $\{k^*(t)\}_{t=0}^{\infty}$ for that of the state,⁶ then by definition,

$$J[k(0)] \equiv \int_0^{\infty} e^{-\delta t} U [c^*(t), k^*(t)] dt.$$

The nice structure of this problem relates to the following property. Suppose that the optimal plan has been followed until a time $T > 0$, so that $k(T)$ is equal to the value $k^*(T)$ given in the last footnote. Imagine a new decision maker who maximises the discounted flow of utility from

⁵ Non-stationary problems often can be handled by methods analogous to those discussed below, but they require additional notation to keep track of the exogenous factors that are changing.

⁶ These are related by

$$k^*(t) = \int_0^t G [c^*(s), k^*(s)] ds + k(0)$$

time $t = T$ onward,

$$\int_T^\infty e^{-\delta(t-T)} U[c(t), k(t)] dt \quad (11.1.2)$$

s.t.: $\dot{k}_t = G[c(t), k(t)]$

but with the initial value of k given by $k(T) = k^*(T)$. Then the optimal program determined by this new decision maker will coincide with the continuation, from time T onward, of the optimal program determined at time 0, given $k(0)$. You should construct a proof of this fundamental result, which is intimately related to the notion of *dynamic consistency*. This basic idea of dynamic programming—to reduce the general problem to a two period problem—is illustrated in figure ??.




Figure 11.1: Reducing the infinite horizon problem to a two-period problem.[‡]

You should also convince yourself of a key implication of this result, namely that

$$J[k(0)] = \int_0^T e^{-\delta t} U[c^*(t), k^*(t)] dt + e^{-\delta T} J[k^*(T)], \quad (11.1.3)$$

where $J[k^*(T)]$ denotes the maximised value of (11.1.2) given that $k(T) = k^*(T)$ and the original constraint $\dot{k}_t = G[c(t), k(t)]$ is respected. Equation (11.1.3) implies that we can think of our original, $t = 0$, problem as the finite-horizon problem of maximising

$$\int_0^T e^{-\delta t} U[c(t), k(t)] dt + e^{-\delta T} J[k(T)] \quad (11.1.4)$$

s.t.: $\dot{k}_t = G[c(t), k(t)]$ for $0 < t < T$.

Of course, in practice it may not be so easy to determine the correct functional form for $J(k)$, as we shall see below! Nonetheless, this way of formulating our problem—which is known as *Bellman's principle of dynamic programming*—leads directly to a characterisation of the optimum. Because this characterisation is derived most conveniently by starting in discrete time, I first set up a discrete-time analogue of our basic maximisation problem and then proceed to the limit of continuous time.

Let's imagine that time is carved up into discrete intervals of length h . A decision on the

control variable c , which is a flow, sets c at some fixed level per unit time over an entire period of duration h . Furthermore, we assume that changes in k , rather than accruing continuously with time, are ‘credited’ only at the very end of a period, like monthly interest on a bank account. We thus consider the problem:

$$\begin{aligned} & \max_{c(t)} \sum_{t=0}^{\infty} e^{-\delta t h} U[c(t), k(t)] h \\ \text{s.t.: } & k(t+h) - k(t) = hG[c(t), k(t)], k(0) \text{ given.} \end{aligned} \quad (11.1.5)$$

Above, $c(t)$ is the fixed rate of consumption over period t while $k(t)$ is the given level of k that prevails from the very end of period $t-1$ until the very end of t . In (11.1.5) I have multiplied $U(c, k)$ [resp. $G(c, k)$] by h because the cumulative flow of utility [resp. change in k] over a period is the product of a fixed instantaneous rate of flow [resp. rate of change] and the period’s length. Bellman’s principle provides a simple approach to the preceding problem. It states that the problem’s value function is given by

$$\begin{aligned} J[k(t)] &= \max_{c(t)} \{U[c(t), k(t)]h + e^{-\delta h} J[k(t+h)]\} \\ \text{s.t.: } & k(t+h) - k(t) = hG[c(t), k(t)], k(t) \text{ given.} \end{aligned} \quad (11.1.6)$$

It implies, in particular, that the optimal $c^*(t)$ must be chosen to maximise the term in braces. By taking the functional relationship (11.1.6) to the limit as $h \mapsto 0$, we will find a way to characterise the continuous-time optimum.⁷

We will make four changes in (11.1.6) to get it into useful form.

1. First, subtract $J[k(t)]$ from both sides.
2. Second, impose the constraint from (11.1.4) by substituting for $k(t+h)$, $k(t) + hG[c(t), k(t)]$.
3. Third, replace $e^{-\delta h}$ by its power-series representation,

$$e^{-\delta h} = \left(\sum_{n=1}^{\infty} \frac{1}{n!} \right)^{-\delta h} = \sum_{n=0}^{\infty} \frac{(-\delta h)^n}{n!} = 1 - \frac{\delta h}{1} + \frac{\delta^2 h^2}{2} - \frac{\delta^3 h^3}{6} + \dots$$

4. Finally, divide the whole thing by h .

The result is

$$0 = \max_c \left\{ U(c, k) - \left[\delta - \frac{(\delta^2 h)}{2} + \frac{\delta^3 h^2}{6} - \dots \right] J[k + hG(c, k)] + \frac{J[k + hG(c, k)] - J[k]}{h} \right\} \quad (11.1.7)$$

where implicitly all variables are dated t . Notice that

$$\frac{J[k + hG(c, k)] - J[k]}{h} = \frac{\{J[k + hG(c, k)] - J(k)\} G(c, k)}{G(c, k)h}.$$

⁷ All of this presupposes that a well-defined value function exists—something which in general requires justification. (See the extended example in this section for a concrete case.) I have also not proven that the value function, when it does exist, is differentiable. We know that it will be for the type of problem under study here, so I’ll feel free to use the value function’s first derivative whenever I need it below. With somewhat less justification, I’ll also use its second and third derivatives.

It follows that as $h \mapsto 0$, the left-hand side above approaches $J'(k)G(c, k)$. Accordingly, we have proved the following

Proposition 60. *At each moment, the control c^* optimal for maximising (11.1.1) satisfies the Bellman equation*

$$0 = U(c^*, k) + J'(k)G(c^*, k) - \delta J(k) = \max_c \{U(c, k) + J'(k)G(c, k) - \delta J(k)\}. \quad (11.1.8)$$

This is a very simple and elegant formula. What is its interpretation? As an intermediate step in interpreting (11.1.8), define the *Hamiltonian* for this maximisation problem as

$$\mathcal{H}(c, k) \equiv U(c, k) + J'(k)G(c, k). \quad (11.1.9)$$

In this model, the intertemporal tradeoff involves a choice between higher current c and higher future k . If c is consumption and k wealth, for example, the model is one in which the utility from consuming now must continuously be traded off against the utility value of savings. The Hamiltonian $\mathcal{H}(c, k)$ can be thought of as a measure of the flow value, in current utility terms, of the consumption-savings combination implied by the consumption choice c , given the predetermined value of k . The Hamiltonian solves the problem of ‘pricing’ saving in terms of current utility by multiplying the flow of saving, $G(c, k) = \dot{k}$, by $J'(k)$, the effect of an increment to wealth on total lifetime utility. A corollary of this observation is that $J'(k)$ has a natural interpretation as the shadow price (or marginal current utility) of wealth.⁸ More generally, leaving our particular example aside, $J'(k)$ is the shadow price one should associate with the state variable k .

This brings us back to the Bellman equation, equation (11.1.8). Let c^* be the value of c that maximises $\mathcal{H}(c, k)$, given k , which is arbitrarily predetermined and therefore might not have been chosen optimally.⁹ Then (11.1.8) states that

$$\mathcal{H}(c^*, k) = \max_c \{\mathcal{H}(c, k)\} = \delta J(k). \quad (11.1.10)$$

In words, the maximised Hamiltonian is a fraction δ of an optimal plan’s total lifetime value. Equivalently, the instantaneous value flow from following an optimal plan divided by the plan’s total value—i.e., the plan’s rate of return—must equal the rate of time preference, δ . Notice that if we were to measure the current payout of the plan by $U(c^*, k)$ alone, we would err by not taking proper account of the value of the current increase in k . This would be like leaving investment out of our measure of GNP! The Hamiltonian solves this accounting problem by valuing the increment to k using the shadow price $J'(k)$.

To understand the implications of the Bellman equation (11.1.8) for optimal consumption we must go ahead and perform the maximisation that it specifies (which amounts to maximising the Hamiltonian). As a by-product, we will obtain the Pontryagin necessary conditions for optimal

⁸ For reasons that will become clearer later on, $J'(\cdot)$ is also called the *costate variable*.[‡]

⁹ It is important to understand clearly that at a given point in time t , $k(t)$ is not an object of choice (which is why we call it a state variable). Variable $c(t)$ can be chosen freely at time t (which is why it is called a control variable), but its level influences the change in $k(t)$ over the next infinitesimal time interval, $k(t + dt) - k(t)$, not the current value $k(t)$.

control.

Maximising the term in braces in (11.1.8) over c , we get¹⁰

$$U_c(c^*, k) = -G_c(c^*, k)J'(k). \quad (11.1.11)$$

The reason this condition is necessary is easy to grasp. At each moment the decision maker can decide to ‘consume’ a bit more, but at some cost in terms of the value of current ‘savings.’ A unit of additional consumption yields a marginal payoff of $U_c(c^*, k)$, but at the same time, savings change by $G_c(c^*, k)$. The utility c value of a marginal fall in savings thus is $-G_c(c^*, k)J'(k)$; and if the planner is indeed at an optimum, it must be that this marginal cost just equals the marginal current utility benefit from higher consumption. In other words, unless (11.1.11) holds, there will be an incentive to change c from c^* , meaning that c^* cannot be optimal.

Let’s get some further insight by exploiting again the recursive structure of the problem. It is easy to see from (11.1.11) that for any date t the optimal level of the control, $c^*(t)$, depends only on the inherited state $k(t)$ (regardless of whether $k(t)$ was chosen optimally in the past). Let’s write this functional relationship between optimal c and k as $c^* = c(k)$, and assume that $c(k)$ is differentiable. (For example, if c is consumption and k total wealth, $c(k)$ will be the household’s consumption function.) Functions like $c(k)$ will be called *optimal policy functions*, or more simply, *policy functions*. Because $c(k)$ is defined as the solution to (11.1.8), it automatically satisfies

$$0 = U[c(k), k] + J'(k)G[c(k), k] - \delta J(k).$$

Equation (11.1.11) informs us as to the optimal relation between c and k at a point in time. To learn about the implied optimal behaviour of consumption over time, let’s differentiate the preceding equation wrt k :

$$0 = [U_c(c^*, k) + J'(k)G_c(c^*, k)] c'(k) + U_k(c^*, k) + [G_k[c^*, k] - \delta] J'(k) + J''(k)G(c^*, k).$$

The expression above, far from being a hopeless quagmire, is actually just what we’re looking for. Notice first that the left-hand term multiplying $c'(k)$ drops out entirely thanks to (11.1.11): another example of the Envelope Theorem. This leaves us with the rest,

$$U_k(c^*, k) + J'(k) [G_k[c^*, k] - \delta] + J''(k)G(c^*, k) = 0. \quad (11.1.12)$$

Even the preceding simplified expression probably is not totally reassuring.¹¹ Do not despair, however. A familiar economic interpretation is again fortunately available. We saw earlier that $J'(k)$ could be usefully thought of as the shadow price of the state variable k . If we think of k as an asset stock (capital, foreign bonds, whatever), this shadow price corresponds to an asset price. Furthermore, we know that under perfect foresight, asset prices adjust so as to equate the asset’s total instantaneous rate of return to some required or benchmark rate of return, which

¹⁰ I assume interior solutions throughout.

¹¹ To us (11.1.12) should be familiar: It is the *Euler* equation we derived as necessary condition in the last chapter.[‡]

in the present context can only be the time-preference rate, δ . As an aid to clear thinking, let's introduce a new variable, λ , to represent the shadow price $J'(k)$ of the asset k :

$$\lambda \equiv J'(k).$$

Our next step will be to rewrite (11.1.12) in terms of λ . The key observation allowing us to do this concerns the last term on the right-hand side of (11.1.12), $J''(k)G(c^*, k)$. The chain rule of calculus implies that

$$J''(k)G(c^*, k) = \frac{dJ'(k)}{dk} \frac{dk}{dt} = \frac{d\lambda}{dk} \frac{dk}{dt} = \frac{d\lambda}{dt} = \dot{\lambda};$$

and with this fact in hand, it is only a matter of substitution to express (11.1.12) in the form

$$\delta = \frac{U_k + \lambda G_k + \dot{\lambda}}{\lambda}. \quad (11.1.13)$$

This is just the asset-pricing equation promised in the last paragraph.

Can you see why this last assertion is true? To understand it, let's decompose the total return to holding a unit of stock k into 'dividends' and 'capital gains.' The 'dividend' is the sum of two parts, the direct effect of an extra unit of k on utility, U_k , and its effect on the rate of increase of k , λG_k . (We must multiply G_k by the shadow price λ in order to express the *physical* effect of k on \dot{k} in the same terms as U_k , that is, in terms of utility.) The 'capital gain' is just the increase in the price of k , $\dot{\lambda}$. The sum of dividend and capital gain, divided by the asset price λ , is just the rate of return on k , which, by (11.1.13) must equal δ along an optimal path.

11.2 Theory II

We will see in the following three sections how to solve continuous-time deterministic problems using Bellman's method of dynamic programming, which is based on the value function $J(k)$. We have seen how to interpret the derivative of the value function, $J'(k)$, as a sort of shadow asset price, denoted by λ . The last order of business is to show that we have actually proved a simple form of Pontryagin's Maximum Principle (Pontryagin, Boltyanskii, Gamkrelidze, and Mishchenko 1962):

Proposition 61. (*Maximum Principle*) *Let $c^*(t)$ solve the problem (11.1.1). Then there exist variables $\lambda(t)$ —called costate variables—such that the Hamiltonian*

$$\mathcal{H}(c, k(t), \lambda(t)) \equiv U[c, k(t)] + \lambda(t)G[c, k(t)]$$

is maximised at $c = c^(t)$ given $\lambda(t)$ and $k(t)$; that is*

$$\frac{\partial \mathcal{H}}{\partial c}(c^*, k, \lambda) \equiv U_c(c^*, k) + \lambda G_c(c^*, k) = 0 \quad (11.2.1)$$

at all times (assuming, as always, an interior solution). Furthermore, the costate variable obeys the differential equation

$$\dot{\lambda} = \lambda\delta - \frac{\partial \mathcal{H}}{\partial k}(c^*, k, \lambda) = \lambda\delta - [U_k(c^*, k) + \lambda G_k(c^*, k)] \quad (11.2.2)$$

for $\dot{k} = G(c^*, k)$ and $k(0)$ given.¹²

You can verify that if we identify the costate variable with the derivative of the value function, $J'(k)$, the Pontryagin necessary conditions are satisfied by our earlier dynamic-programming solution. In particular, (11.2.1) coincides with (11.1.11) and (11.2.2) coincides with (11.1.13). So we have shown, in a simple stationary setting, why the Maximum Principle ‘works.’ The principle is actually more broadly applicable than you might guess from the foregoing discussion—it easily handles non-stationary environments, side constraints, etc. And it has a stochastic analogue, to which we shall turn in the next chapter.¹³

11.3 Example I: Deterministic optimal control in the Ramsey model

Let’s step back for a moment from this abstract setting to consolidate what we’ve learned through an example. Consider the standard Ramsey problem of a consumer social planner who maximises (Ramsey 1926)¹⁴

$$\int_T^\infty e^{-\delta t} U[c(t)] dt \quad (11.3.1)$$

s.t.: $\dot{k} = f(k) - c$

(where c is consumption, k capital, and $f(k)$ the production function). Now $U_k = 0$, $G(c, k) = f(k) - c$, $G_c = -1$, and $G_k = f'(k)$. In this setting, (11.1.13) becomes the statement that the rate of time preference should equal the marginal product of capital plus the rate of accrual of utility capital gains,

$$\delta = f'(k) + \frac{\dot{\lambda}}{\lambda}.$$

Condition (11.1.11) becomes

$$U'(c) = \lambda.$$

¹² You should note that if we integrate differential-equation (11.2.2), we get the general solution

$$\lambda(t) = \int_t^\infty e^{-\delta(s-t)} \frac{\partial \mathcal{H}}{\partial k}[c^*(s), k(s), \lambda(s)] ds + Ae^{\delta t},$$

where A is an arbitrary constant. [To check this claim, just differentiate the foregoing expression with respect to t : If the integral in the expression is $I(t)$, we find that $\dot{\lambda} = \delta I - \frac{\partial \mathcal{H}}{\partial k} + \delta Ae^{\delta t} = \delta \lambda - \frac{\partial \mathcal{H}}{\partial k}$.] I referred in the prior example to an additional terminal condition requiring the present value of the capital stock to converge to zero along an optimal path. Since $\lambda(t)$ is the price of capital at time t , this terminal condition usually requires that $\lim_{t \rightarrow \infty} e^{-\delta t} \lambda(t) = 0$, or that $A = 0$ in the solution above. The particular solution that remains equates the shadow price of a unit of capital to the discounted stream of its shadow ‘marginal products,’ where the latter are measured by partial derivatives of the flow of value, \mathcal{H} , with respect to k .

¹³ For more details and complications on the deterministic Maximum Principle, see (Arrow and Kurz 1970).

¹⁴ More recent versions of the same model come under the names of (Cass 1965) and (Koopmans 1965). The Ramsey-Cass-Koopmans model looks at a micro-founded (i.e. based on maximising behaviour), closed-economy growth model.

Taking the time derivative of this last equality we get

$$U''(c)\dot{c} = \dot{\lambda},$$

and we can express the optimal dynamics of c and k as a non-linear differential equation system:

$$\begin{aligned} \dot{c} &= -\frac{U'(c)}{U''(c)} [f'(k) - \delta] \\ \dot{k} &= f(k) - c. \end{aligned} \tag{11.3.2}$$

Assuming the *Inada* conditions (8.1.13), we obtain the neoclassical production function as graphed below.¹⁵




Figure 11.2: A neoclassical production function.[‡]

You can see the phase diagram for this system in figure 11.3.¹⁶

The diagram shows that, given k , any positive initial c initiates a path along which the two preceding differential equations for c and k are respected. But not all of these paths are optimal, since the differential equations specify conditions that are merely necessary, but not sufficient, for optimality.

Indeed, only one path will be optimal in general: we can write the associated policy function as $c^* = c(k)$. For given k_0 , paths (like D' in figure 11.3) with initial consumption levels exceeding $c_0 = D$ imply that k becomes negative after a finite time interval. Since a negative capital stock is nonsensical, such paths are not even feasible, let alone optimal.

Paths (like D'' in figure 11.3) with initial consumption levels below $c_0 = D$ imply that k gets to be too large, in the sense that the individual could raise lifetime utility by eating some capital and never replacing it. These ‘overaccumulation’ paths violate a sort of terminal condition stating that the present value of the capital stock should converge to zero along an optimal path. I shall

¹⁵ GR denotes the ‘Golden Rule’ and mGR the ‘modified Golden Rule’ capital stocks k . For a discussion of what these are see e.g. (Blanchard and Fischer 1989, chapter 2).[‡]

¹⁶ Be sure you can derive it yourself! For the necessary economics you may want to consult (Barro and Sala-i Martin 1995, chapter 2) or (Blanchard and Fischer 1989, chapter 2).[‡]




Figure 11.3: The phase diagram for the Ramsey model. A population growth term has been added to the transition equation, i.e. $\dot{k} = f(k) - c - nk$, in order to have the $\dot{k} = 0$ locus go back to zero. Otherwise, the $\dot{k} = 0$ locus would look like the dashed curve.

not take the time to discuss such terminal conditions here.

If we take the CES case

$$U(c) = \frac{c^{1-\frac{1}{\varepsilon}} - 1}{1-\frac{1}{\varepsilon}} \quad (11.3.3)$$

$$f(k) = rk$$

where ε and r are positive constants, we can actually find an algebraic formula for the policy function $c(k)$.

Let's conjecture that optimal consumption is proportional to wealth, that is, that $c(k) = \eta k$ for some constant η to be determined. If this conjecture is right, the capital stock k will follow $\dot{k} = (r - \eta)k$, or, equivalently,

$$\frac{\dot{k}}{k} = r - \eta.$$

This expression gives us the key clue for finding η . If $c = \eta k$, as we've guessed, then also

$$\frac{\dot{c}}{c} = r - \eta.$$

But necessary condition (11.3.2) requires that

$$\frac{\dot{c}}{c} = \varepsilon(r - \delta),$$

which contradicts the last equation unless

$$\eta = (1 - \varepsilon)r + \varepsilon\delta. \quad (11.3.4)$$

Thus, $c(k) = [(1 - \varepsilon)r + \varepsilon\delta]k$ is the optimal policy function. In the case of log utility ($\varepsilon = 1$), we

simply have $\eta = \delta$. We get the same simple result if it so happens that r and δ are equal.

Equation (11.3.4) has a nice interpretation. In Milton Friedman's permanent-income model (Friedman 1957) where $\delta = r$, people consume the annuity value of wealth, so that $\eta = \delta = r$. This rule results in a level consumption path. When $\delta \neq r$, however, the optimal consumption path will be tilted, with consumption rising over time if $r > \delta$ and falling over time if $r < \delta$. By writing (11.3.4) as

$$\eta = r - \varepsilon(r + \delta)$$

we can see these two effects at work. Why is the deviation from the Friedman permanent-income path proportional to ε ? Recall that ε , the elasticity of intertemporal substitution, measures an individual's willingness to substitute consumption today for consumption in the future. If ε is high and $r > \delta$, for example, people will be quite willing to forgo present consumption to take advantage of the relatively high rate of return to saving; and the larger is ε , ceteris paribus, the lower will be η . Alert readers will have noticed a major problem with all this. If $r > \delta$ and ε is sufficiently large, η , and hence 'optimal' consumption, will be negative. How can this be? Where has our analysis gone wrong?

The answer is that when $\eta < 0$, no optimum consumption plan exists! After all, nothing we've done demonstrates existence: our arguments merely indicate some properties that an optimum, if one exists, will need to have. No optimal consumption path exists when $\eta < 0$ for the following reason. Because optimal consumption growth necessarily satisfies $\frac{\dot{c}}{c} = \varepsilon(r - \delta)$, and $\varepsilon(r - \delta) > r$ in this case, optimal consumption would have to grow at a rate exceeding the rate of return on capital, r . Since capital growth obeys $\frac{\dot{k}}{k} = r - \frac{c}{k}$, however, and $c \geq 0$, the growth rate of capital, and thus that of output, is at most r . With consumption positive and growing at 3 percent per year, say, but with capital growing at a lower yearly rate, consumption would eventually grow to be greater than total output—an impossibility in a closed economy. So the proposed path for consumption is not feasible. This means that no feasible path—other than the obviously suboptimal path with $c(t) = 0$, for all t —satisfies the necessary conditions for optimality. Therefore no feasible path is optimal: no optimal path exists.

Let's take our analysis a step further to see how the value function $J(k)$ looks. Observe first that at any time t ,

$$k(t) = k(0)e^{(r-\eta)t} = k(0)e^{\varepsilon(r-\delta)t},$$

where $k(0)$ is the starting capital stock and η is given by (11.3.4). Evidently, the value function at $t = 0$ is just

$$\begin{aligned} J[k(0)] &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} \left\{ \int_0^{\infty} e^{-\delta t} [\eta k(t)]^{1-\frac{1}{\varepsilon}} dt - \frac{1}{\delta} \right\} \\ &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} \left\{ \int_0^{\infty} e^{-\delta t} [\eta k(0)e^{\varepsilon(r-\delta)t}]^{1-\frac{1}{\varepsilon}} dt - \frac{1}{\delta} \right\} \\ &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} \left\{ \frac{[\eta k(0)]^{1-\frac{1}{\varepsilon}}}{\delta - (\varepsilon-1)(r-\delta)} - \frac{1}{\delta} \right\}. \end{aligned}$$

So the value function $J(k)$ has the same general form as the utility function, but with k in place of c (compare with (11.3.5) with (11.3.3) if you do not believe me). This is not the last time we'll encounter this property. Alert readers will again notice that to carry out the final step of the last calculation, I had to assume that the integral in braces above is convergent, that is, that $\delta - (\varepsilon - 1)(r - \delta) > 0$. Notice, however, that $\delta - (\varepsilon - 1)(r - \delta) = r - \varepsilon(r - \delta) = \eta$, so the calculation is valid provided an optimal consumption program exists. If one does not, the value function clearly does not exist either: we cannot specify the maximised value of a function that does not attain a maximum.

This counterexample should serve as a warning against blithely assuming that all problems have well-defined solutions and value functions.

11.4 Example II: Extending the first Ramsey example[‡]

Reading: (Blanchard and Fischer 1989, 37-50).

Consider the same Ramsey problem as above

$$\begin{aligned} & \int_T^\infty e^{-\delta t} U[c] dt \\ \text{s.t.: } & \dot{k} = f(k) - c \end{aligned} \tag{11.4.1}$$

where c is consumption, k capital, and $f(k)$ the production function—again all variables are in per capita terms and implicitly understood to be dated at t .

We enjoy the same simplifications as before:

- Utility only depends on consumption—not on wealth; $U_k = 0$. We will again employ the CES-utility function from (11.3.3) that we will later reduce to the logarithmic case ($\varepsilon = 1$).
- The production function is *c.r.s.*; hence we can again write per capita output as $f(k) \equiv F(\frac{K}{N}, 1)$.
- The transition equation leaves all newly produced wealth that is not directly consumed for the future. All output that is not eaten can be converted into (storable) wealth. $G(c, k) = f(k) - c$; $G_c = -1$; $G_k = f'(k)$.

Contrary to what we've done above, we will now generalise our approach and set-up the *Hamiltonian* from the principles we derived in our theoretical treatment of optimal control. That is, we simply apply our insights from (11.1.9) and set-up the Hamiltonian as

$$\mathcal{H}(c, k) = e^{-\delta t} U(c) + \lambda [f(k) - c].$$

We now apply the *Maximum principle* by taking derivatives of this Hamiltonian wrt the control c . (Below we will do the same for the state k to obtain the *Euler* equation. By these means we actually derive the equivalents of the two conditions for optimality (11.1.8 and 11.1.11) that we

more or less blindly applied in the last example.)¹⁷

$$\frac{\partial \mathcal{H}(c, k)}{\partial c} = 0 \Leftrightarrow e^{-\delta t} U_c = \lambda.$$

Since we already know that we will also need the time derivative of the costate, we take *log-derivatives*, that is, we take logs and derivatives in one step. First the *rhs*

$$\ln \frac{\partial(e^{-\delta t} U_c)}{\partial t} = \frac{\partial}{\partial t} (-\delta t + \ln U_c) = -\delta + \frac{U_{cc} \dot{c}}{U_c} = -\delta + \frac{c U_{cc} \dot{c}}{U_c c},$$

and now the *lhs*

$$\ln \frac{\partial(\lambda)}{\partial t} = \frac{\dot{\lambda}}{\lambda}.$$

To make this less painful, let's define $\theta \equiv -\frac{c U_{cc}}{U_c}$, the coefficient of relative risk aversion (CoRRA). Again, this is related to the elasticity of intertemporal substitution as: $\varepsilon = 1 - \frac{1}{\theta}$. Assuming a CES utility function, we thus obtain

$$\frac{\dot{\lambda}}{\lambda} = -\delta - \theta \frac{\dot{c}}{c}.$$

Going one step further and assuming log-utility (which exhibits $\theta = -1$), our expression then simplifies to

$$\frac{\dot{\lambda}}{\lambda} = -\delta + \frac{\dot{c}}{c}.$$

This leaves us with our *Euler* equation

$$\frac{\partial \mathcal{H}(c, k)}{\partial k} = -\dot{\lambda}$$

into which we plug the above result for $\dot{\lambda}$. This immediately yields

$$\delta - f'(k) = \frac{\dot{c}c}{c}.$$

We thus obtain the following dynamics of c and k

$$\begin{aligned} \dot{c} &= -\frac{U'(c)}{U''(c)} [f'(k) - \delta] \\ \dot{k} &= f(k) - c. \end{aligned} \tag{11.4.2}$$

You will agree that this is identical to the dynamics we derived in the previous example. Therefore the same solutions (& phase diagrams) apply as we derived above.

11.5 Example III: Centralised / decentralised equivalence results[‡]

Reading: (Blanchard and Fischer 1989, 39-55) or, alternatively, (Barro and Sala-i Martin 1995, 59-87).

¹⁷ You will notice that in the Obstfeld paper, the below $e^{-\delta t}$ term is missing. This is the case because he defines his costate variable in 'present value' terms as $\tilde{\lambda} = \frac{\lambda}{e^{-\delta t}}$.

In (11.3.1, 11.4.1) we have set up our optimisation problem in terms of a social planner trying to achieve a social optimum (which we did not explicitly mention at that point) by ‘artificially’ equilibrating the marginal product of capital and the interest rate and the marginal product of labour and wages.¹⁸ Although this is a useful exercise because it is often easier to show certain properties of an economy from a planner’s point of view, it does not necessarily imply that we will also be able to find a decentralised Walrasian equilibrium. But this is what we really want to show in order to keep microeconomists happy.

The point of the discussion in section is that the information available to the decision maker and the timing of this information matters a great deal. Although we will be able to show the equivalence of the centralised and decentralised solutions for the two particular version of the Ramsey model we choose, we will see that this is by no means an immediate result for a greater family of models.

To set up the household’s maximisation problem, we will need to look at the role of uncertainty about the precise paths of future wages and interest rates. We will do this in a separate subsection. Since it would be simply unbearable to use exactly the same setting as above one more time we now account for depreciation and population growth. That is, we look at a population of workers or households N (who supply their labour inelastically) that grows at the rate n and capital depreciates at the rate β . In order to stick with Obstfeld’s notation, we keep using δ for time preference (although this is really standard for depreciation) and use β for depreciation. Sorry if this causes confusion when comparing the text with the literature. All the other features of the model remain unchanged; in particular we stick with all the neoclassical standard assumptions. Output is either consumed or invested, i.e.

$$Y_t = F(K_t, N_t) = C_t + \dot{K}_t + \beta K_t,$$

which, given our assumption of c.r.s., equals

$$y_t = f(k_t) = c_t + \dot{k}_t + (n - \beta)k_t.$$

In the previous example, household preferences at time s were represented by

$$U_s \equiv \int_s^\infty e^{-\delta(t-s)} u(c_t) dt$$

and we assumed $u(\cdot)$ to be non-negative, increasing and concave.¹⁹ As before, the $e^{-\delta t}$ -multiplier represents time preference. A positive value of the time preference parameter δ implies that future utils are valued less the later they are received.²⁰ Since we also have to account for family growth, we now also have to multiply $u(c)$ by the family size, $N = e^{nt}$. We have to share the available utils among more and more family members, so the influence of n is positive. To avoid

¹⁸ In actual fact we did not introduce these prices of capital and labour explicitly.

¹⁹ As we did above, when no confusion is to be feared we shall suppress time subscripts in the further analysis.

²⁰ Ramsey set $\delta = 0$ in 1928 because he considered a social planner who would value future generations’ utility less than current generations utility ‘ethically indefensible’.

obvious stability problems, we shall assume that $\delta > n$ and write the combined multiplier as $e^{-\delta t} e^{nt} = e^{-(\delta-n)t}$ and the ensuing utility function as²¹

$$U_s \equiv \int_s^\infty e^{-(\delta-n)(t-s)} u(c_t) dt$$

11.5.1 The command optimum

The most useful interpretation of the basic unit of our economy for the discussion of the centralised outcome is that of a single household. That is, we conceptualise our economy to consist of precisely one family. We further assume the existence of a benevolent social planner who has access to all household information and who maximises social (i.e. family) welfare. His role is to determine the optimal path of the control variable and thereby make the consumption saving choice for the family.

Let's start by setting the social planner's problem at $t = 0$ is then to

$$\begin{aligned} \arg \max_c U_0 &\equiv \int_0^\infty e^{-(\delta-n)t} u(c) dt \\ \text{s.t.: } \dot{k} &= f(k) - c - (\beta + n)k, \quad k(0) \text{ given.} \end{aligned}$$

We set up the (present value) Hamiltonian from the planner's perspective as

$$\mathcal{H}_p(c, k) \equiv u(c)e^{-(\delta-n)t} + \mu [f(k) - (\beta+n)k - c].$$

We replace μ with its present value representation $\lambda \equiv \mu e^{(\delta-n)t}$ and obtain

$$\mathcal{H}_p(c, k) = \{u(c) + \lambda [f(k) - (\beta+n)k - c]\} e^{-(\delta-n)t}.$$

The *foc* are given by

$$\begin{aligned} \frac{\partial \mathcal{H}_p}{\partial c} = 0 &\Leftrightarrow e^{-(\delta-n)t} u' = \lambda, \\ \frac{\partial \mathcal{H}_p}{\partial k} = -\dot{\lambda} &\Leftrightarrow \lambda [f'(k) - (n + \beta)] = -\dot{\lambda}. \end{aligned}$$

Which implies for logarithmic utility $u(c) = \ln(c)$ the system:

$$\begin{aligned} \frac{\dot{c}}{c} &= f'(k) - \beta - \delta, \\ \dot{k} &= f(k) - c - (\beta + n)k. \end{aligned} \tag{11.5.1}$$

11.5.2 The decentralised optimum

We will now replace the somewhat artificial centralised problem with a ('micro-founded') decentralised problem. To do this convincingly, we need three additional markets: one for labour (the rental price of which is the wage w), one for capital services (the rental price of which is the interest rate r), and one for debt in which the household can borrow and lend (as it turns out, the rate of return on debt will equal the interest rate r).

²¹ The awkward business of using the notation $e^{-(\delta-n)t}$ rather than $e^{(n-\delta)t}$ stems from the desire to constantly reassure ourselves that we are dealing with a negative exponent.

Households

The households own all capital and labour and both can be rented out to the firms. We would also like to model inter-household borrowing and lending which would be rather unconvincing if there is only a single household. Therefore, we change our interpretation of the maximising unit to that of a representative household among many identical households. For now we will assume that the household knows the future paths of w and r perfectly well in advance. That is, under ‘perfect foresight’ $\{w_t, r_t\}_{t=0}^{\infty}$ are known and can be used in the maximisation problem. Later we will relax this to the case of uncertainty.

The household’s problem at $t = s$ is to decide how much labour and capital it wishes to rent out in order to obtain wages it can consume. Since both decisions are inelastic, the household’s problem at each point in time is to decide upon how much to consume and how much to save:

$$\arg \max_c U_0 \equiv \int_0^{\infty} e^{-(\delta-n)t} u(c) dt \quad (11.5.2)$$

s.t.:

$$\dot{a} = w + (r - n)a - c, \text{ where } a \equiv k - b_p, \text{ and } k(0) \text{ is given.} \quad (11.5.3)$$

It is useful to distinguish between non-human wealth a , which consists of the household’s capital holdings k minus its debt services b , and human wealth which is the present value of the household’s wages w .

Firms

We assume that there are many small, profit maximising firms that all employ the same c.r.s. technology fully characterised by the production function $F(K, N)$. This amounts to assuming perfect competition. The firm’s problem is to

$$\max_{K,L} F(K, L) - rK - \beta K - wL. \quad (11.5.4)$$

Market clearing

The *foc* for profit maximisation give the market clearing rates for the labour and capital markets as

$$\begin{aligned} r &= f'(k) - \beta, \\ w &= f(k) - kf'(k). \end{aligned}$$

No-Ponzi-game condition

In the absence of restrictions on borrowing the solution to the household’s maximisation problem is trivial: Borrow as much as possible and spend on consumption while deferring repayment forever by borrowing even more for debt service. Such behaviour is prevented by the introduction

of the so called no-Ponzi-game²² condition (which is a particular TVC)

$$\lim_{t \rightarrow \infty} a e^{-\int_0^t (r\nu - n) d\nu} \geq 0. \quad (11.5.5)$$

The interpretation is that family debt may increase (asymptotically) no faster than the interest rate.

Hamiltonian

We maximise (11.5.2) s.t.: (11.5.3) and (11.5.5). We realise that while the control is still c , our new state must be a . This means we set up our (decentralised) Hamiltonian as

$$\mathcal{H}_d(c, k) \equiv u(c)e^{-(\delta-n)t} + \lambda[(r-n)a + w - c], \quad (11.5.6)$$

and obtain:

1. The Euler equation as

$$\frac{\partial \mathcal{H}_d(c, k)}{\partial a} = -\dot{\lambda} \Leftrightarrow \dot{\lambda} = (n-r)\lambda. \quad (11.5.7)$$

2. The maximum principle as

$$\frac{\partial \mathcal{H}_d(c, k)}{\partial c} = 0 \Leftrightarrow u'(c)e^{-(\delta-n)t} = \lambda. \quad (11.5.8)$$

3. While still respecting the no-Ponzi-game condition

$$\lim_{t \rightarrow \infty} a\lambda = 0.$$

Log-differentiating (11.5.8) gives

$$\frac{\dot{\lambda}}{\lambda} = -\delta + n + \frac{u''(c)\dot{c}}{u'(c)}$$

and equalising this to the *rhs* of our Euler equation (11.5.7) yields

$$n - r = -\delta + n + \frac{u''(c)\dot{c}}{u'(c)}.$$

Assuming CES utility with the reciprocal of the elasticity of intertemporal substitution

$\theta \equiv -\frac{u''(c)c}{u'(c)}$, we obtain

$$\frac{\dot{c}}{c} = \frac{1}{\theta} [r - \delta].$$

²² “Charles Ponzi, one of Boston’s sons made a quick fortune in the 1920s using chain letters. He was sent to prison and died poor.” (Blanchard and Fischer 1989, 84). Therefore, if you want to escape a similar fate, better obey condition (11.5.5).

Equilibrium

Assuming logarithmic utility $u(c) = \ln(c)$ with $(\theta = 1)$ and imposing market clearing, this yields

$$\frac{\dot{c}}{c} = f'(k) - \beta - \delta.$$

Since in equilibrium *aggregate* private debt b_p must be zero, there is no equilibrium lending or borrowing. Thus we conclude that $a = k$ and by inserting the market clearing conditions into (11.5.3), we obtain

$$\dot{k} = (r - n)k + w - c.$$

The last two equations characterise precisely the same system that we derived for the command optimum (11.5.1). Therefore our analysis of the dynamics of the command economy carries over to the decentralised setting.

Exercises

Exc 11.1: (Kiyotaki) Consider a household maximising the utility of consumption from present to future under certainty

$$\sum_{t=0}^{\infty} \beta^t \ln c_t$$

s.t.: $b_{t+1} = Rb_t + y_t - c_t$

taking as given the initial asset Rb_0 , the real gross interest rate R and the path of the wage income y_t .

1. Derive the relationship between permanent income and the initial asset and the wage income from present to future.
2. Write down the Bellman equation. Derive the value function.
3. Show that the consumption is proportional to permanent income.
4. Suppose that wages grow at the constant rate

$$y_{t+1} = (1 + \gamma)y_t$$

where $1 + \gamma < R$. Derive the relationship between consumption and the initial asset and the current wage explicitly.

Exc 11.2: (Pissarides) Express the neoclassical production function $Y = F(K, N)$ in per capita terms. Let $k \equiv \frac{K}{N}$ and $y \equiv \frac{Y}{N}$.

1. On a diagram of y against k , show the marginal products of labour and capital.
2. Show that if the factors are paid according to their marginal product, the wage rental ratio ω (the ratio of the wage rate to the capital rental rate) may be written as a function of k . Derive this function, sign its first and second derivatives and graph it on a diagram against k .
3. Calculate the inverse elasticity $\frac{dk}{d\omega} \frac{\omega}{k}$ of the function derived above for the Cobb-Douglas and CES functions.

Exc 11.3: (Pissarides) Consider the effects of a capital depreciation allowance equal to a fraction a of total depreciation on a steady-state capital stock and interest rate in *the* Solow model. The capital depreciation allowance is financed by lump-sum taxes.

Exc 11.4: (Pissarides) Derive the equilibrium of the Solow growth model when instead of consumption equal to a constant fraction of income, it is equal to all wage income.

Exc 11.5: (Pissarides based on (Mankiw, Romer, and Weil 1992) Consider the production function

$$Y = K^\alpha H^\lambda (AL)^{1-\alpha-\lambda},$$

where Y is output, K is physical capital, H is human capital, A is the level of technology and L is labour. The parameters α and λ are positive and $\alpha + \lambda < 1$. L and A grow at constant rates n and x respectively. Output can be used on a one-to-one basis for consumption or investment in either type of capital. Both types of capital depreciate at the rate δ . Assume that gross investment in

physical capital is a fraction s_k of output and that gross investment in human capital is a fraction s_h of output.

1. Obtain the laws of motion for physical and human capital per units of effective labour.
2. What are the steady-state values of physical and human capital and output, all in units of effective labour?
3. This augmented Solow-Swan model can be tested empirically with cross-country data if we assume that all countries are in steady-state. Derive a *log*-linear regression equation for output per worker. What problems would arise in estimating this regression by OLS?
4. Derive an equation for the growth rate of output per unit of effective labour. How does this equation look when expressed as a linear approximation in the neighbourhood of the steady-state? If $\alpha = 0.3$, $\lambda = 0.5$, $\delta = 0.05$, $n = 0.01$, and $x = 0.02$, what is the rate of convergence near the steady-state? Compare the convergence rate in the augmented Solow-Swan model with that in the standard model.
5. Use the result from (d) to derive a regression equation for the average growth rate of output per worker, $\frac{1}{T} \log \left[\frac{y(t+T)}{y(t)} \right]$ where T is the length of the observation interval. What problems arise in the econometric estimation of the rate of convergence, for example, if the levels of technology differ across countries?

Exc 11.6: (Pissarides) The government imposes a proportional tax on output and uses it to finance a lump sum transfer of goods x to each individual. Compute the economy's adjustment path to this tax from the day the tax is announced for the Ramsey model with $u(c) = \log c$. Assume that both the tax and transfer payments are made in goods.

Chapter 12

Stochastic optimal control in continuous time

Reading: (Obstfeld 1992) and (Kamien and Schwartz 1991).

The optimisation principles set forth in the last chapter extend directly to the stochastic case. The main difference is that to do continuous-time analysis, we will have to think about the right way to model and analyse uncertainty that evolves continuously with time. To understand the elements of continuous-time stochastic processes requires a bit of investment, but there is a large payoff in terms of the analytic simplicity that results.

12.1 Theory III

Let's get our bearings by looking first at a discrete-time stochastic model.¹ Imagine now that the decision maker maximises the von Neumann-Morgenstern expected-utility indicator

$$E_0 \left[\sum_{t=0}^{\infty} e^{-\delta t h} U [c(t), k(t)] h \right], \quad (12.1.1)$$

where $E_t[X]$ is the expected value of the random variable X conditional on all information available up to (and including) time t .² Maximisation is to be carried out subject to the constraint that

$$k(t+h) - k(t) = G[c(t), k(t), \theta(t+h), h], \quad k(0) \text{ given}, \quad (12.1.2)$$

where $\{\theta(t)\}_{t=-\infty}^{\infty}$ is a sequence of exogenous random variables with known joint distribution, and such that only realisations up to and including $\theta(t)$ are known at time t . For simplicity I will assume that the θ process is *first-order Markov*, that is, that the joint distribution of $\{\theta(t+h), \theta(t+2h), \dots\}$ conditional on $\{q(t), q(t-h), \dots\}$ depends *only* on $\theta(t)$. For example,

¹ An encyclopedic reference on discrete-time dynamic programming and its applications in economics is (Stokey and Lucas 1989). The volume pays special attention to the foundations of stochastic models.

² Preferences less restrictive than those delimited by the von Neumann-Morgenstern axioms have been proposed, and can be handled by methods analogous to those sketched below.

the $AR(1)$ process

$$q(t) = \rho\theta(t - h) + v(t),$$

where $v(t)$ is distributed independently of past θ 's, has this first-order Markov property.

Constraint (12.1.2) differs from its deterministic version in (11.1.4) in that the time interval h appears as an argument of the transition function, but not necessarily as a multiplicative factor. Thus, (12.1.2) is somewhat more general than the constraint in (11.1.4). The need for this generality arises because $\theta(t + h)$ is meant to be 'proportional' to h in a sense that will become clearer as we proceed.

Criterion (12.1.1) reflects inherent uncertainty in the realisations of $c(t)$ and $k(t)$ for $t > 0$. Unlike in the deterministic case, the object of individual choice is not a single path for the control variable c . Rather, it is a *sequence of contingency plans* for c . Now it becomes really essential to think in terms of a policy function mapping the 'state' of the program to the optimal level of the control variable. The optimal policy function giving $c^*(t)$ will not be a function of the state variable $k(t)$ alone, as it was in the last chapter; rather, it will depend on $k(t)$ and $\theta(t)$, because $\theta(t)$ (thanks to the first-order Markov assumption) is the piece of current information that helps forecast the future realisations $\theta(t + h)$, $\theta(t + 2h)$, etc. Since $k(t)$ and $\theta(t)$ evolve stochastically, writing

$$c^*(t) = c[k(t); \theta(t)]$$

makes it clear that from the perspective of any time before t , $c^*(t)$ will be a random variable, albeit one that depends in a very particular way on the realised values of $k(t)$ and $\theta(t)$.

Bellman's principle continues to apply, however. To implement it, let us write the value function—again defined as the maximised value of (12.1.1)—as $J[k(0); \theta(0)]$. Notice that $\theta(0)$ enters the value function for the same reason that $\theta(t)$ influences $c^*(t)$. If θ is a positive shock to capital productivity (for example), with θ positively serially correlated, then a higher current value of θ leads us to forecast higher θ 's for the future. This higher expected path for θ both raises expected lifetime utility and influences the optimal consumption choice.

In the present setting we write the Bellman equation as

$$J[k(t); \theta(t)] = \max_{c(t)} \left\{ U[c(t), k(t)]h + e^{-\delta h} E_t [J[k(t+h); \theta(t+h)]] \right\} \quad (12.1.3)$$

where the maximisation is done subject to (12.1.2). The rationale for this equation basically is the same as before for (11.1.6). The contingent rules for $\{c(s)\}_{s=t+1}^{\infty}$ that maximise

$$E_t \left[\sum_{s=t}^{\infty} e^{-\delta sh} U[c(s), k(s)] h \right]$$

subject to (12.1.2), given $k(t)$ and the optimal choice $c^*(t)$, will also maximise

$$E_t \left[\sum_{s=t+1}^{\infty} e^{-\delta sh} U[c(s), k(s)] h \right]$$

subject to (12.1.2), given the probability distribution for $k(t+h)$ induced by $c^*(t)$.

Equation (12.1.3) is the stochastic analogue of (11.1.6) for the case of first-order Markovian uncertainty. The equation is immediately useful for discrete-time analysis: just use (12.1.2) to eliminate $k(t+h)$ from (12.1.3) and differentiate away. But our concern here is with continuous-time analysis. We would like to proceed as before, letting the market interval h go to zero in (12.1.3) and, hopefully, deriving some nice expression analogous to (11.1.8). Alas, life is not so easy. If you try to take the route just described, you will end up with an expression that looks like the expected value of

$$\frac{J[k(t+h); \theta(t+h)] - J[k(t); \theta(t)]}{h}.$$

This quotient need not, however, converge (as $h \mapsto 0$) to a well-defined random variable. One way to appreciate the contrast between the present setup and the usual setup of the calculus is as follows. Because $J[k(t); \theta(t)]$ is a *random variable*, a plot of its realisations against time—a sample path—is unlikely to be differentiable. Even after time is carved up into very small intervals, the position of the sample path will change abruptly from period to period as new realisations occur. Thus, expressions like the quotient displayed above may have no well-defined limiting behaviour as $h \mapsto 0$. To proceed further we need a new mathematical theory that allows us to analyse infinitesimal changes in random variables. The stochastic calculus is designed to accomplish precisely this goal.

12.2 Stochastic Calculus

Let $X(t)$ be a random variable whose change between periods $t-1$ and t , $\Delta X(t) = X(t) - X(t-1)$, has mean μ and variance σ^2 . To simplify matters I'll assume that $\Delta X(t)$ is normally distributed, although this is not at all necessary for the argument.³

We are interested in the case where $\Delta X(t)$, the change in random variable X over the period of length 1 between $t-1$ and t , can be viewed as a sum (or integral) of very small (in the limit, infinitesimal) random changes. We would also like each of these changes, no matter how small, to have a normal distribution. Our method, as in the usual calculus, is to divide the time interval $[t-1, t]$ into small segments. But we need to be sure that no matter how finely we do the subdivision, $\Delta X(t)$, the sum of the smaller changes, remains $N(\mu, \sigma^2)$. To begin, carve up the interval $[t-1, t]$ into n disjoint subintervals, each of length $h = 1/n$. For every $i \in \{1, 2, \dots, n\}$, let $\nu(i)$ be a $N(0, 1)$ random variable with $E[\nu(i)\nu(j)] = 0$ for $i \neq j$. Suppose that ν can be written as

$$\Delta X(t) = \sum_{i=1}^n \mu h + \sigma h^{\frac{1}{2}} \nu(i) \tag{12.2.1}$$

Then since $\nu h = 1$, (12.2.1) is consistent with our initial hypothesis that $E[\Delta X(t)] = \mu$ and

³ For a simplified yet rigorous exposition of these matters, see (Merton 1974).

$V[\Delta X(t) = s]$.⁴ For example,

$$V[\Delta X(t)] = \sigma^2 \sum_{i=1}^n \sum_{j=1}^n E \left[\frac{\nu(i)\nu(j)}{n} \right] = \sum_{i=1}^n E \left[\frac{\nu(i)^2}{n} \right] = \sigma^2.$$

Equation (12.2.1) expresses the finite change $\Delta X(t)$ as the sum of tiny independent normal increments of the form $\mu h + \sigma h^{\frac{1}{2}} \nu$. It is customary to denote the limit of such an increment as $h \mapsto 0$ by $\mu dt + \sigma dz$, where for any instant τ , $dz(\tau) = \lim_{h \rightarrow 0} h^{\frac{1}{2}} \nu(\tau)$. When this limit is well-defined, we say that $X(t)$ follows the **Gaussian diffusion process**

$$dX(t) = \mu dt + \sigma dz(t), \quad (12.2.2)$$

which means, in notation that is suggestive but that I will not attempt to define rigorously, that

$$X(t) = X(\tau) + \mu(t - \tau) + \sigma \int_{\tau}^t dz(s) = X(\tau) + \mu(t - \tau) + \sigma[z(t) - z(\tau)]$$

for all $\tau \leq t$.⁵

Think of $X(t)$ as following a Brownian motion, i.e. a continuous-time random walk with a predictable rate of drift μ and an instantaneous rate of variance (variance per unit of time) σ^2 . When $\sigma = 0$, we are back in the deterministic case and are therefore allowed to assert that $X(t)$ has time derivative μ : $\frac{dX(t)}{dt} = \mu$. But when $\sigma > 0$, $X(t)$ has sample paths that are differentiable nowhere. So we use a notation, (12.2.2), that does not require us to ‘divide’ random differences by dt . Because we are looking at arbitrarily small increments over arbitrarily small time intervals, however, the sample paths of $X(t)$ are continuous.

Now that we have a sense of what (12.2.2) means, I point out that this process can be generalised while maintaining a Markovian setup in which today’s X summarises all information useful for forecasting future X ’s. For example, the process

$$dX = \mu(X, t)dt + \sigma(X, t)dz \quad (12.2.3)$$

allows the drift and variability of dX to be functions of the level of $X(t)$ itself, which is known at time t , and of time.

There is a further set of results we’ll need before tackling the one major theorem of stochastic analysis applied below, Itô’s chain rule. We need to know the rules for multiplying stochastic differentials. We’re familiar, from the usual differential calculus, with the idea that quantities

⁴ $V_t[X]$ gives the variance of the random variable X using all information available at time t .

⁵ Again, see (Merton 1974) for a more rigorous treatment. To make all this more plausible, you may want to write (12.2.1) (for our earlier case of $\tau = t - 1$) as

$$\Delta X(t) - \mu = \sum_{i=1}^n \frac{\nu(i)}{\sqrt{n}},$$

where $n = \frac{1}{h}$ is the number of increments in $[t - 1, t]$. We know from the Central Limit Theorem 10 that as $n \mapsto \infty$, the right-hand side above is likely to approach a limiting normal distribution even if the $\nu(i)$ ’s are not normal (so my assumptions above were stronger than necessary). Obviously, also, $X(t) - X(t - h)$ will be normally distributed with variance $h\sigma^2$ no matter how small h is. But $X(t) - X(t - h)$ divided by h therefore explodes as $h \mapsto 0$ (its variance is $\frac{\sigma^2}{h}$). This is why the sample paths of diffusion processes are not differentiable in the usual sense.

of order dt are important, whereas quantities of order dt^m , $m > 1$, are not. For example, in calculating the derivative of the function y^2 , we compute h^{-1} times the limit of $(y+h)^2 - y^2 = 2yh + h^2$ as $h \mapsto 0$.

The derivative is simply $2y$, because the term h^2 goes to zero even after division by h . The same principle will apply in stochastic calculus. Terms of order greater than h are discarded. In particular $dt^2 = \lim_{h \rightarrow \infty} h^2$ will be set to zero, just as always.

What about something like the product $dzdt$? Since this is the limit of $h^{\frac{3}{2}}\nu$ as $h \mapsto \infty$, it shrinks faster than h and accordingly will be reckoned at zero:

$$dzdt = 0 \tag{12.2.4}$$

Finally, consider $dz^2 = \lim_{h \rightarrow \infty} h\nu^2$. This is of order h , and thus does not disappear as h gets very small. But the variance of this term can be shown to be $2h^2$, which is zero asymptotically.⁶

By Chebyshev's inequality (cf. B.0.6), $h\nu^2$ thus converges in probability to its expected value, h , as $h \mapsto 0$, and so we write

$$dz^2 = dt. \tag{12.2.5}$$

Let's turn now to Itô's famous lemma. Suppose that the random variable $X(t)$ follows a diffusion process such as (12.2.3). The basic idea of Itô's Lemma is to help us compute the stochastic differential of the random variable $f[X(t)]$, where $f(\cdot)$ is a differentiable function. If $\sigma(X, t) \equiv 0$, then the chain rule of ordinary calculus gives us the answer: the change in $f(X)$ over an infinitesimal time interval is given by $df(X) = f'(X)dX = f'(X)\mu(X, t)dt$. If $\sigma(X, t) \neq 0$ but $f(\cdot)$ is linear, say $f(X) = aX$ for some constant a , then the answer is also quite obvious: in this special case, $df(X) = f'(X)dX = a\mu(X, t)dt + a\sigma(X, t)dz$.

Even if $f(\cdot)$ is non-linear, however, there is often a simple answer to the question we've posed:

Lemma 62. (*Itô's Lemma*) *Let $X(t)$ follow a diffusion process, and let $f : \mathbb{R} \mapsto \mathbb{R}$ be twice continuously differentiable (C^2). The stochastic differential of $f(X)$ is*

$$df(X) = f'(X)dX + \frac{1}{2}f''(X)dX^2. \tag{12.2.6}$$

Rem: If X follows the diffusion process (12.2.3), then, using rules (12.2.4) and (12.2.5) to compute dX^2 in (12.2.6), we get $\sigma(X, t)$

$$df(X) = \left[\mu(X, t)f'(X) + \frac{\sigma(X, t)^2}{2}f''(X) \right] dt + \sigma(X, t)f'(X)dz. \tag{12.2.7}$$

You'll notice that (12.2.7) differs from the 'naïve' chain rule only in modifying the expected drift in $f(X)$ by a term that depends on the curvature of $f(\cdot)$. If $f''(X) > 0$ so that $f(\cdot)$ is strictly

⁶ To prove this, note that because $n \sim N(0, 1)$

$$V[h\nu^2] = E[(h\nu^2 - h)^2] = E[h^2\nu^4 - 2h^2\nu^2 + h^2] = 3h^2 - 2h^2 + h^2 = 2h^2.$$

convex, for example, (12.2.7) asserts that

$$\begin{aligned} E_t[df(X)] &= E_t f[X(t+dt)] - f[X(t)] \\ &> \\ f'(X)\mu(X,t)dt &= f'(X)E_t[dX] = f[E_t[X(t+dt)]] - f[X(t)]. \end{aligned}$$

But anyone who remembers Jensen's Inequality (5.3.1) knows that

$$E_t[f[X(t+dt)]] > f[E_t[X(t+dt)]]$$

for convex $f(\cdot)$, and that the opposite inequality holds for concave $f(\cdot)$. So Itô's Lemma should not come as a surprise.⁷

Motivation for Itô's Lemma. The proof of Itô's Lemma is quite subtle, so a heuristic motivation of this key result will have to suffice. Once again I'll rely on a limit argument. For an interval length h , Taylor's Theorem⁸ (B.0.5) implies that

$$\begin{aligned} f[X(t+h)] - f[X(t)] &= f'[X(t)][X(t+h) - X(t)] \\ &+ \frac{1}{2}f''\{X(t) + \xi(h)[X(t+h) - X(t)]\}[X(t+h) - X(t)]^2 \end{aligned}$$

where $\xi(h) \in [0, 1]$. It may look 'obvious' to you that this converges to (12.2.6) as $h \mapsto 0$. Beware. It turns out to be quite a chore to ensure that the right-hand side of this expression is well behaved as $h \mapsto 0$, largely because of the complicated dependence of the term $f''\{X(t) + \xi(h)[X(t+h) - X(t)]\}$ on h . Fortunately, as $h \mapsto 0$, the randomness in this term does disappear quickly enough that we can safely equate it to $f''[X(t)]$ in the limit. The result is (12.2.6). It should now be clear how one would motivate a multivariate version of Itô's Lemma using the multivariate Taylor expansion.

12.3 Theory IV

The preceding digression on stochastic calculus has equipped us to answer the question raised at the outset: What is the continuous-time analogue of (12.1.3), the stochastic Bellman equation?

To make matters as simple as possible, in analogy with section (12.1)'s time-stationary setup, I'll assume that $q(t+h) = X(t+h) - X(t)$, where $X(t)$ follows the simple diffusion process (12.2.2), $dX = rdt + \sigma dz$, for constant r and σ . Under this assumption $E_t[\theta(t+h)] = rh$ always, so knowledge of $\theta(t)$ gives us no information about future values of θ . Thus the value function

⁷ In case you do not remember Jensen's Inequality, here is a quick sketch of a proof. Recall that a convex function has the property that

$$\gamma f(X_1) + (1-\gamma)f(X_2) \geq f[\gamma X_1 + (1-\gamma)X_2]$$

$\forall \gamma \in [0, 1]$. It is easy to extend this to the proposition that $\sum \pi_i f(X_i) \geq f(\sum \pi_i X_i)$ for (π_1, \dots, π_n) in the unit simplex. (Try it.) So for finite discrete probability distributions we're done. (Obviously concave functions work the same way, with the inequalities reversed.) Now consider the case in which the random variable X has an arbitrary continuous density function $\pi(X)$. We can approximate $E[f(X)]$ by sums of the form $\sum f(X_i)\pi(X_i)h$, each of which must be at least as great as $f[\sum X_i\pi(X_i)h]$ if we choose the ??

⁸ For Taylor's theorem with remainder, see any good calculus text.

depends on the state variable k alone. Now (12.1.3) becomes

$$J[k(t)] = \max_{c(t)} \left\{ U[c(t), k(t)]h + e^{-\delta h} E_t [J[k(t+h)]] \right\}. \quad (12.3.1)$$

Let's carry on by adapting the last section's strategy of subtracting $J[k(t)]$ from both sides of (12.1.3) and replacing $e^{-\delta h}$ by $1 - \delta h$. (We now know we can safely ignore the terms in h_m for $m \geq 2$.) The result is

$$0 = \max_{c(t)} \{ U[c(t), k(t)]h + E_t [J[k(t+h)]] - J[k(t)] - \delta E_t [J[k(t+h)]h] \}.$$

Now let $h \mapsto 0$. According to (12.1.2), $dk = G(c, k, dX, dt)$, and I assume that this transition equation defines a diffusion process for k . Itô's Lemma then tells us that

$$dJ(k) = J'(k)dk + \frac{1}{2}J''(k)dk^2. \quad (12.3.2)$$

Thus as $h \mapsto 0$,

$$E_t [J[k(t+h)]] - J[k(t)] \mapsto J'[k(t)]E_t [dk(t)] + \frac{1}{2}J''[k(t)]E_t [dk(t)^2].$$

Furthermore, as $h \mapsto 0$, $E_t [J[k(t+h)]] \mapsto J[k(t)]$. So we end up with the following:

Proposition 63. (*Continuous-Time Stochastic Bellman Equation*) Consider the problem of maximising

$$E_0 \left[\int_0^\infty e^{-\delta t} U(c, k) dt \right]$$

subject to a diffusion process for k controlled by c , and given $k(0)$. At each moment, the optimal control c^* satisfies the Bellman equation

$$\begin{aligned} 0 &= U[c^*, k]dt + J'(k)E_t [G(c^*, k, dX, dt)] \\ &+ \frac{1}{2}J''(k)E_t [G(c^*, k, dX, dt)^2] - \delta J(k)dt \\ &= \max_{c(t)} \left\{ U[c, k]dt + J'(k)E_t [dk] + \frac{1}{2}J''(k)E_t [dk^2] - \delta J(k)dt \right\}. \end{aligned} \quad (12.3.3)$$

Equation (12.3.3) is to be compared with equation (11.1.8), given in Proposition (61). Indeed, the interpretation of Proposition (63) is quite similar to that of Proposition (61). Define the stochastic Hamiltonian [in analogy to (11.1.9)] as

$$\mathcal{H}(c, k) \equiv U(c, k) + J'(k) \frac{E_t [dk]}{dt} + \frac{1}{2}J''(k) \frac{E_t [dk^2]}{dt}. \quad (12.3.4)$$

The Hamiltonian has the same interpretation as (11.1.9), but with a stochastic twist. The effect of a given level of 'savings' on next period's 'capital stock' now is uncertain. Thus the Hamiltonian measures the expected flow value, in current utility terms, of the consumption-savings combination implied by the consumption choice c , given the predetermined (and known)

value of k . The analogy will be clearer if you use (12.3.2) to write (12.3.3) as

$$\mathcal{H}(c, k) \equiv U(c, k) + \frac{E_t [dJ(k)]}{dt},$$

and if you use the ordinary chain rule to write the deterministic Hamiltonian (11.1.9) as⁹

$$U(c, k) + J'(k)\dot{k} = U(c, k) + \frac{dJ(k)}{dt}.$$

The stochastic Bellman equation therefore implies the same rule as in the deterministic case, but in an expected-value sense. Once again, optimal consumption c^* satisfies (11.1.10),

$$\mathcal{H}(c^*, k) = \max_c \{\mathcal{H}(c, k)\} = \delta J(k).$$

Rather than proceeding exactly as in our deterministic analysis, I will sacrifice generality for clarity and adopt a specific (but widely used) functional form for the continuous-time version of (12.1.2), $dk = G(c, k, dX, dt)$. I will assume the linear transition equation

$$dk = kdX - cdt = (rk - c)dt + \sigma kdz \quad (12.3.5)$$

(since $dX = rdt + \sigma dz$). What form does (12.3.3) now assume? To see this we have to calculate $E_t [dk]$ and $E_t [dk^2]$. It is clear from (12.3.5) that $E_t [dk] = (rk - c)dt$. Invoking (12.2.4) and (12.2.5), and recalling that $dt^2 = 0$, we see that

$$dk^2 = E_t [dk^2] = k^2 dX^2 - 2ckdXdt + c^2 dt^2 = \sigma^2 k^2 dt.$$

We thus conclude that c^* must solve

$$0 = \max_{c(t)} \left\{ U(c, k) + J'(k)(rk - c) + \frac{1}{2} J''(k) k^2 \sigma^2 - \delta J(k) \right\}. \quad (12.3.6)$$

In principle this equation is no harder to analyse than was (11.1.8): the two are identical [if $G(c, k) = rk - c$] aside from the additional second derivative term in (12.3.6), due to Itô's Lemma. So we proceed as before, starting off by maximising the Hamiltonian.

Since k is predetermined and known at each moment, the necessary condition for c^* to maximise the right hand of (12.3.6) is

$$U_c(c^*, k) = J'(k), \quad (12.3.7)$$

which is the same as (11.1.11) because I've assumed here that $G_c = -1$.

We can also define the optimal policy function $c^* = c(k)$, just as before. By definition $c(k)$

⁹ The notation in (12.3.4) and in the next line below is common. Since $E_t [dk]$, for example, is deterministic, $\frac{E_t [dk]}{dt}$ can be viewed as the expected rate of change in k . Since diffusion processes are not differentiable, $E_t \left[\frac{dk}{dt} \right]$ is in contrast a nonsensical expression.

satisfies the equation

$$0 = U[c(k), k] + J'(k)[rk - c(k)] + \frac{1}{2}J''(k)k^2\sigma^2 - \delta J(k). \quad (12.3.8)$$

One would hope to understand better the implied dynamics of c by differentiating with respect to the state variable. The result is

$$U_k(c^*, k) + J'(k)(r - \delta) + J''(k)k\sigma^2 + J''(k)(rk - c^*) + \frac{1}{2}J'''(k)k^2\sigma^2 = 0, \quad (12.3.9)$$

where I've already applied the envelope condition (12.3.7). It is tempting to give up in the face of all these second and third derivatives; but it is nonetheless possible to interpret (12.3.9) in familiar economic terms. Let's again define the shadow price of k , λ , by

$$\lambda \equiv J'(k).$$

This shadow price is known at time t , but its change over the interval from t to $t + dt$ is stochastic. Equation (12.3.9) differs from (11.1.12) only by taking this randomness into account; and by writing (12.3.9) in terms of λ , we can see precisely how this is done.

To do so we need two observations. First, Itô's Lemma discloses the stochastic differential of λ to be

$$d\lambda = dJ'(k) = J''(k)(kdX - cdt) + \frac{1}{2}J'''(k)k^2\sigma^2 dt \quad (12.3.10)$$

(verify this), so that

$$\frac{E_t[d\lambda]}{dt} = J''(k)(rk - c) + \frac{1}{2}J'''(k)k^2\sigma^2. \quad (12.3.11)$$

Second, the term $J''(k)ks$ in (12.3.9) can be expressed as

$$J''(k)k\sigma^2 = -J'(k)R(k)\sigma^2, \quad (12.3.12)$$

where $R(k) \equiv -\frac{J''(k)k}{J'(k)}$ should be interpreted as a coefficient of relative risk aversion.

Using (12.3.11) and (12.3.12), rewrite (12.3.9) in terms of $\lambda = J'(k)$ as

$$U_k(c^*, k) + \lambda[r - R(k)\sigma^2 - \delta] + \frac{E_t[d\lambda]}{dt},$$

or, in analogy to (11.1.13), as

$$U_k + \lambda \left[r - \frac{R(k)\sigma^2}{2} \right] + \left[\frac{E_t[d\lambda]}{dt} - \frac{\lambda R(k)\sigma^2}{2} \right] = \delta. \quad (12.3.13)$$

To compare (12.3.13) with (11.1.13), notice that under the linear transition equation (12.3.5), r corresponds to the expected value of G_k ; we adjust this expectation downward for risk by subtracting the product of the risk-aversion coefficient and $\frac{\sigma^2}{2}$. An identical risk adjustment is made to the expected 'capital gains' term, $\frac{E_t[d\lambda]}{dt}$. Otherwise, the equation is the same as (11.1.13), and has a corresponding 'efficient asset price' interpretation.

12.4 Example of stochastic optimal control

An individual maximises the expected discounted utility of consumption,

$$E_0 \left[\int_0^\infty e^{-\delta t} U(c) dt \right]$$

subject to a stochastic capital accumulation constraint that looks like (12.3.5):

$$dk = rkdt + \sigma kdz - cdt, \quad k(0) \text{ given.}$$

What is the meaning of this savings constraint? Capital has a mean marginal product of r , but its realised marginal product fluctuates around r according to a white-noise process with instantaneous variance σ^2 . The flow utility function is

$$U(c) = \frac{c^{1-\frac{1}{\varepsilon}} - 1}{1 - \frac{1}{\varepsilon}},$$

as in the second part of the last section's example. To solve the problem I'll make the same guess as before, that the optimal consumption policy function is $c(k) = \eta k$ for an appropriate η . As will be shown below—and as was the case in a deterministic setting—the value function $J(k)$ is a linear function of $k^{1-\frac{1}{\varepsilon}}$, making the risk aversion coefficient $R(k)$ defined after (12.3.12) a constant, $R \equiv \frac{1}{\varepsilon}$. For now I will assume this, leaving the justification until the end.

How can we compute η in the policy function $c(k) = \eta k$? The argument parallels our earlier discussion of the non-stochastic case, which you may wish to review at this point. Start by thinking about the implications of the postulated policy function for the dynamics of capital. If $c(k) = \eta k$, then

$$dk = rkdt + \sigma kdz - c(k)dt = (r - \eta)kdt + \sigma kdz.$$

But as optimal c is proportional to k ,

$$dc = (r - \eta)c dt + \sigma c dz.$$

Above we defined λ as $J'(k)$; but first-order condition (12.3.7) implies that $\lambda = U'(c) = c^{-\frac{1}{\varepsilon}}$. Application of Itô's Lemma to $\lambda = c^{-\frac{1}{\varepsilon}}$ leads to

$$d\lambda = - \left(\frac{1}{\varepsilon} \right) c^{-1-\frac{1}{\varepsilon}} dc + \left(\frac{1}{2} \right) \left(\frac{1}{\varepsilon} \right) \left(1 + \frac{1}{\varepsilon} \right) c^{-2-\frac{1}{\varepsilon}} dc^2.$$

Because we've already established that $E_t[dc] = (r - \eta)c dt$ and that $dc^2 = \sigma^2 c^2 dt$, we infer from the equation above that

$$\frac{E_t[d\lambda]}{dt} = \frac{c^{-1-\frac{1}{\varepsilon}}}{\varepsilon} \left[\eta - r + \left(\frac{1}{2} \right) \left(1 + \frac{1}{\varepsilon} \right) \sigma^2 \right].$$

But there is an alternative way of describing the dynamics of λ : equation (12.3.13) can be written

here as

$$\frac{E_t[d\lambda]}{dt} = \lambda [\delta - (r - R\sigma^2)] = c^{-\frac{1}{\varepsilon}} \left[\delta - \left(r - \frac{\sigma^2}{\varepsilon} \right) \right].$$

So we have derived two potentially different equations for $\frac{E_t[d\lambda]}{dt}$; clearly the two are mutually consistent if and only if

$$\delta - \left(r - \frac{\sigma^2}{\varepsilon} \right) = \frac{[\eta - r + (\frac{1}{2})(1 + \frac{1}{\varepsilon})\sigma^2]}{\varepsilon},$$

or, solving for η , if and only if

$$\eta = r - \varepsilon(r - \delta) + \frac{\varepsilon - 1}{2\varepsilon}\sigma^2.$$

The implied consumption rule is similar to the one that arose in the non-stochastic example analysed earlier, but it corrects for the unpredictable component of the return to capital. (Notice that we again obtain $\eta = \delta$ if $\varepsilon = 1$.) The analogy with (11.3.4) will be clearest if the rule is written as

$$\eta = (1 - \varepsilon)(r - \frac{1}{2}R\sigma^2) + \varepsilon\delta. \quad (12.4.1)$$

In (12.4.1), η appears as the weighted average of the time-preference rate and a risk-adjusted expected return on investment. Problems still arise if $\eta \leq 0$. In these cases an optimum fails to exist, for reasons essentially the same as those discussed in the last section's example.

As a final exercise let's calculate the value function $J(k)$ and confirm the assumption about its form on which I've based my analysis of the optimal consumption policy function. In the process we'll learn some more about the importance of Itô's Lemma. One way to approach this task is to calculate the (random) path for k under an optimal consumption plan, observe that the optimal contingency rule for consumption is $c = \eta k$, and then use this formula to compute the optimal (random) consumption path and lifetime expected utility. Indeed, we took a very similar tack in the deterministic case. So we start by asking what the optimal transition equation for the capital stock, $dk = (r - \eta)kdt + \sigma dz$, implies for the level of k . [Throughout the following discussion, you should understand that η is as specified by (12.4.1).]

Observe first that the optimal capital-stock transition equation can be written as

$$\frac{dk}{k} = (r - \eta)dt + \sigma dz.$$

A crucial warning. You might think that $\frac{dk}{k}$ is the same thing as $d \log(k)$, as in the ordinary calculus. If this were true, we would conclude that the capital stock follows the stochastic process

$$\log[k(t)] = \log[k(0)] + (r - \eta)t + \sigma \int_0^t dz(s),$$

or, equivalently, that

$$k(t) = k(0)e^{(r-\eta)t + \sigma[z(t)-z(0)]}.$$

But this is incorrect. Itô's Lemma tells us that

$$d \log(k) = \frac{dk}{k} - \frac{1}{2} \sigma^2 dt = (r - \eta - \frac{1}{2} \sigma^2) dt + \sigma dz.$$

[The reason for this divergence is Jensen's inequality— $\log(\cdot)$ is a strictly concave function.] It follows that the formula for $k(t)$ below is the right one:

$$k(t) = k(0) e^{(r - \eta - \frac{\sigma^2}{2})t + \sigma[z(t) - z(0)]}. \quad (12.4.2)$$

At an optimum, $k(t)$ will be conditionally log-normally distributed, with an expected growth rate of $r - \eta$ ¹⁰

$$\frac{E_0 [k(t)]}{k(0)} = e^{(r - \eta)t}.$$

As a result of (12.4.2), the value function at $t = 0$ is

$$\begin{aligned} J[k(0)] &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} E_0 \left[\int_0^\infty e^{-\delta t} [\eta k(t)]^{1 - \frac{1}{\varepsilon}} dt - \frac{1}{\delta} \right] \\ &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} \left\{ \int_0^\infty e^{-\delta t} E_0 \left[\eta k(0) e^{(r - \eta - \frac{\sigma^2}{2})t + \sigma[z(t) - z(0)]} \right]^{1 - \frac{1}{\varepsilon}} dt - \frac{1}{\delta} \right\} \\ &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} \left\{ [\eta k(0)]^{1 - \frac{1}{\varepsilon}} \int_0^\infty e^{-\delta t} e^{(1 - \frac{1}{\varepsilon})(r - \eta - \frac{\sigma^2}{2})t} dt - \frac{1}{\delta} \right\} \\ &= \left[1 - \frac{1}{\varepsilon}\right]^{-1} \left\{ \frac{[\eta k(0)]^{1 - \frac{1}{\varepsilon}}}{\delta - (\varepsilon - 1)(r - \frac{R\sigma^2}{2} - \delta)} - \frac{1}{\delta} \right\}. \end{aligned}$$

You'll recognise the final product above as the same formula for $J[k(0)]$ that we encountered in (11.3.3), with the sole amendment that the risk-adjusted expected return $r - \frac{R\sigma^2}{2}$ replaces r everywhere [including in η ; recall (12.4.1)].¹¹ Because $\delta - (\varepsilon - 1)(r - \frac{R\sigma^2}{2} - \delta) = \eta$, $\eta > 0$ ensures convergence of the integral defining $J(k)$. Finally, $J(k)$ is a linear function of $k^{1 - \frac{1}{\varepsilon}}$, as claimed earlier.

There is another, more direct way to find the value function, one that also applies in the deterministic case. [Had we known the value function in advance, we could have used (12.3.7) to compute the consumption function without trial-and-error guesses.] By (12.3.7), the optimal control must satisfy

$$c(k) = J'(k)^{-\varepsilon}.$$

Thus by (12.3.6),

$$0 = \frac{[J'(k)]^{1 - \varepsilon}}{1 - \frac{1}{\varepsilon}} + J'(k) [rk - J'(k)^{-\varepsilon}] + \frac{1}{2} J''(k) k^2 \sigma^2 - \delta J(k).$$

¹⁰ If X is a normal random variable with mean μ and variance σ^2 , e^X is said to be log-normally distributed. The key fact about log-normals that is used repeatedly is that when X is normal,

$$E[e^X] = e^{\frac{\mu + \sigma^2}{2}}.$$

For a proof, see any good statistics text.

¹¹ To move from the second to the third equality above, I used the fact that the normal random variable $[1 - \frac{1}{\varepsilon}]\sigma[z(t) - z(0)]$ has mean zero and variance $[1 - \frac{1}{\varepsilon}]^2 \sigma^2$ conditional on $t = 0$ information.

This is just an ordinary second-order differential equation which in principle can be solved for the variable $J(k)$. You may wish to verify that the value function $J(k)$ we derived above is indeed a solution. To do the non-stochastic case, simply set $\sigma^2 = 0$.

12.5 Theory V

The similarities between this example and its deterministic analogue are striking. They are not always so direct. Nonetheless, it is noteworthy that for the linear state transition equation considered above, there exists a stochastic version of Pontryagin's Maximum Principle. One could attack the problem in full generality,¹² but as my goal here is the more modest one of illustrating the basic idea, I will spare you this.

Proposition 64. (*Stochastic Maximum Principle*) *Let $c^*(t)$ solve the problem of maximising*

$$E_0 \left[\int_0^\infty e^{-\delta(s-t)} U(c(s), k(s)) ds \right]$$

subject to the transition equation

$$dk(t) = rk(t)dt + \sigma k(t)dz(t) - c(t)dt, \quad k(0) \text{ given}$$

where $z(t)$ is a standard Gaussian diffusion. Then there exist costate variables $\lambda(t)$ such that if $\zeta(t)$ is the instantaneous conditional covariance of $\lambda(t)$ and $z(t)$, the risk-adjusted Hamiltonian

$$\mathcal{H}[c, k(t), \lambda(t), \zeta(t)] \equiv U[c, k(t)] + \lambda(t)[rk(t) - c] + \zeta(t)\sigma k(t)$$

is maximised at $c^ = c(t)$ given $\lambda(t), \zeta(t)$, and $k(t)$; that is,*

$$\frac{\partial \mathcal{H}}{\partial c}(c^*, k, \lambda, \zeta) = U_c(c^*, k) - \lambda = 0 \quad (12.5.1)$$

at all times (assuming an interior solution). Furthermore, the costate variable obeys the stochastic differential equation

$$\begin{aligned} d\lambda &= \lambda \delta dt - \frac{\partial \mathcal{H}}{\partial c}(c^*, k, \lambda, \zeta) + \zeta dz \\ &= \lambda \delta dt - [U_k(c^*, k) + \lambda r + \zeta \sigma] dt + \zeta dz \end{aligned} \quad (12.5.2)$$

*for $dk = rkdt - c^*dt + \sigma kdz$ and $k(0)$ given.*

To understand how this proposition follows from our earlier discussion, observe first that because λ will again equal $J'(k)$, the instantaneous conditional covariance of $\lambda(t)$ and $z(t)$ can be seen from (12.2.4), (12.2.5), and (12.3.10) to be

$$\zeta = \frac{E_t [d\lambda dz]}{dt} = J''(k)\sigma k \quad (12.5.3)$$

Thus, with reference to the definition (12.3.4) of the unadjusted stochastic Hamiltonian, given

¹² As does (Bismut 1975).

here by

$$\mathcal{H}(c, k) = U(c, k) + J'(k)(rk - c) + \frac{1}{2}J''(k)\sigma^2k^2,$$

we have

$$\tilde{\mathcal{H}}(c, k, \lambda, \zeta) = \mathcal{H}(c, k) + U(c, k) + \frac{1}{2}J''(k)\sigma^2k^2 = \mathcal{H}(c, k) - \frac{\lambda R(k)\sigma^2}{2},$$

where $R(k)$ is the relative risk-aversion coefficient defined above. Accordingly, we can interpret $\tilde{\mathcal{H}}$ as the expected instantaneous flow of value minus a premium that measures the riskiness of the stock of capital currently held.

With (12.5.3) in hand it is easy to check the prescriptions of the Stochastic Maximum Principle against the results we've already derived through other arguments. Clearly (12.5.1) corresponds directly to (12.3.7). Likewise, if you multiply (12.3.9) by dt and combine the result with (12.3.10), you will retrieve (12.5.2).

12.6 Conclusion

These notes have offered intuitive motivation for the basic optimisation principles economists use to solve deterministic and stochastic continuous-time models. My emphasis throughout has been on the Bellman principle of dynamic programming, which offers a unified approach to all types of problems. The Maximum Principle of optimal control theory follows from Bellman's approach in a straightforward manner. I have only been able to scratch the surface of the topic. Methods like those described above generalise to much more complex environments, and have applications much richer than those I worked through for you. The only way to gain a true understanding of these tools is through 'hands on' learning: you must apply them yourself in a variety of situations. As I noted at the outset, abundant applications exist in many areas of economics. I hope these notes make this fascinating body of research more approachable.

Exercises

Just kidding.

Appendix A

Some more interesting problems

Exc 1.1:[Bicycle tracks] Peter Doyle¹ called attention to the following section from *The Adventure of the Priory School*, by Sir Arthur Conan Doyle: “This track, as you perceive, was made by a rider who was going from the direction of the school.” “Or towards it?” “No, no, my dear Watson. The more deeply sunk impression is, of course, the hind wheel, upon which the weight rests. You perceive several places where it has passed across and obliterated the more shallow mark of the front one. It was undoubtedly away from the school.”

Does Holmes know what he is talking about? The impression Holmes and Watson are looking at is this




Figure A.1: Which way did the bicycle go?

Exc 1.2:[Bicycle pedals] It is fairly commonly known, at least among bicyclists, that there is something funny about the way that the pedals of a bicycle screw into the cranks. One of the pedals has a normal ‘right-hand thread’, so that you screw it in clockwise—the usual way—like a normal screw or light bulb, and you unscrew it counter-clockwise. The other pedal has a ‘left-hand thread’, so that it works exactly backwards: You screw it in counter-clockwise, and you

¹ The first three problems are from Peter Doyle’s notes *Geometry and the Imagination*. They can be downloaded from <http://www.math.dartmouth.edu/~doyle>.

unscrew it clockwise.

This ‘asymmetry’ between the two pedals—actually it is a surfeit of symmetry we have here, rather than a dearth—is not just some whimsical notion on the part of bike manufacturers. If the pedals both had normal threads, one of them would fall out before you got to the end of the block.

If you are assembling a new bike out of the box, the fact that one of the pedals screws in the wrong way may cause some momentary confusion, but you can easily figure out what to do by looking at the threads. The real confusion comes when for one reason or another you need to unscrew one of your pedals, and you can not remember whether this pedal is the normal one or the screwy one, and the pedal is on so tightly that a modest torque in either direction fails to budge it. You get set to give it a major twist, only which way do you turn it? You worry that if you to turn it the wrong way you will get it on so tight that you will never get it off.

If you try to figure out which pedal is the normal one using common sense, the chances are overwhelming that you will figure it out exactly wrong. If you remember this, then you are all set: Just figure it out by common sense, and then go for the opposite answer. Another good strategy is to remember that ‘right is right; left is wrong.’

1. What is the difference between a screw and a bolt?
2. Do all barber poles spiral the same way? What about candy canes? What other things spiral? Do they always spiral the same way?
3. Take two identical bolts or screws or light bulbs or barber poles, and place them tip to tip. Describe how the two spirals meet.
4. Take a bolt or a screw or a light bulb or a barber pole and hold it perpendicular to a mirror so that its tip appears to touch the tip of its mirror image. Describe how the two spirals meet.
5. When you hold something up to an ordinary mirror you can not quite get it to appear to touch its mirror image. Why not? How close can you come? What if you use a different kind of mirror?
6. Why is a right-hand thread called a ‘right-hand thread’? What is the ‘right-hand rule’?
7. Which way do tornados and hurricanes rotate in the northern hemisphere? Why?
8. Which way does water spiral down the drain in the southern hemisphere, and how do you know?
9. Use common sense to figure out which pedal on a bike has the normal, right-hand thread. If you come up with the correct answer that ‘right is right; left is wrong’ then we offer you our humblest apologies.
10. Now see if you can figure out the correct explanation.
11. You can simulate what is going on here by curling your fingers loosely around the eraser end of a nice long pencil (a long thin stick works even better), so that there is a little extra room for the pencil to roll around inside your grip. Get someone else to press down gently

- on the business end of the pencil, to simulate the weight of the rider's foot on the pedal, and see what happens when you rotate your arm like the crank of a bicycle.
12. The best thing is to make a wooden model. Drill a block through a block of wood to represent the hole in the crank that the pedal screws into, and use a dowel just a little smaller in diameter than the hole to represent the pedal.
 13. If your pedal is on really, really tight, you might be tempted to use a 'cheater', which is a pipe that you slip over the end of a wrench to increase the effective length of the handle. If it takes a force of 150 pounds on the end of a 9-inch long adjustable wrench to loosen the pedal, how much force will be required on the end of an 18-inch long cheater?
 14. Wrench manufacturers, pipe manufacturers, bicycle manufacturers, your insurance underwriter, your family doctor, and your geometry teacher all maintain that using a cheater is a bad idea. Do you understand why? Discuss some of the dangers to which you might expose yourself by using a cheater. If despite all this well-meaning advice you should go ahead and use a cheater, and some harm should come to you, who will be to blame?

Exc 1.3:[Pulling back on a pedal] Imagine that I am steadying a bicycle to keep it from falling over, but without preventing it from moving forward or back if it decides that it wants to. The reason it might want to move is that there is a string tied to the right-hand pedal (which is to say, the right-foot pedal), which is at its lowest point, so that the right-hand crank is vertical. You are squatting behind the bike, a couple of feet back, holding the string so that it runs (nearly) horizontally from your hand forward to where it is tied to the pedal.

1. Suppose you now pull gently but firmly back on the string. Does the bicycle go forward, or backward? Remember that I am only steadying it, so that it can move if it has a mind to. No, this is not a trick; the bike really does move one way or the other. Can you reason it out? Can you imagine it clearly enough so that you can *feel* the answer intuitively?
2. Try it and see.

John Conway makes the following outrageous claim. Say that you have a group of six or more people, none of whom have thought about this problem before. You tell them the problem, and get them all to agree to the following proposal. They will each take out a dollar bill, and announce which way they think the bike will go. They will be allowed to change their minds as often as they like. When everyone has stopped waffling, you will take the dollars from those who were wrong, give some of the dollars to those who were right, and pocket the rest of the dollars yourself. You might worry that you stand to lose money if there are more right answers than wrong answers, but Conway claims that in his experience this *never* happens. There are always more wrong answers than right answers, and this despite the fact that you tell them in advance that there are going to be more wrong answers than right answers, and allow them to bear this in mind during the waffling process. (Or is it *because* you tell them that there will be more wrong answers than right answers?)

Exc 1.4:[The two envelopes problem] You are offered a choice between two envelopes. Each envelope contains an unstated amount of cash and you are told that one of the envelopes contains

twice as much money as the other. You pick one of them and look inside and find ten dollars. That means the other envelope has either £5 or £20. Now you are offered the choice of keeping what you have or exchanging it for the other (unopened) envelope.

Should you switch?

On the one hand one may think “Well, the odds that I picked the envelope with the higher amount were 50-50 at the start and nothing has changed. It makes no difference that I looked inside of this one, thus my odds are still 50-50 and It makes no difference whether I keep what I have or switch. On the average I cannot do any better”.

On the other hand, one might consider the expected value of switching: “There is a 50-50 chance that the envelope contains £20 and a 50-50 chance that it contains £5. If we repeated this game many times and I switched when I saw £10 in my envelope then half the time I wind up with £5 and half the time I would wind up with £20. Thus on the average the other envelope will be worth $(.5)5 + (.5)20 = £12.5$. Since it will cost me £10 to make the switch (you must give up the envelope you have), the expected net gain for switching is £2.5. So it makes sense to switch!”

Or does it? Do either of these two arguments make sense? Why?

Now suppose there had been another person in the game. This person gets the envelope you do not choose. You both look inside your envelopes but do not tell the other person how much you have got. Then you both are offered the opportunity to exchange envelopes. And by the above logic you both want to switch! Worse still you both would be willing to pay for the privilege of switching: in your case, seeing £10, you would be willing to pay anything up to £2.5 for the privilege of switching. Is this possible? After all you had complete free will over your choice at the start.

But actually things are a lot worse than that. Suppose you do not look into the envelope you picked. You know there is some unknown amount in the envelope, say X . The other envelope has either $2X$ or $.5X$. The expected value of the other envelope is therefore $(.5)(.5)X + (.5)(2)X = 1.25X$. Thus if you switch you expect to gain $.25X$. So which ever envelope you choose to start with, you immediately want to switch to the other before you even look at it. And of course once you get that envelope you want to switch back. Well ...

Can you devise a Bayesian approach to rescue the decision maker(s)?

Exc 1.5: [Brachistochrone problem from (Sengupta and Fanchon 1997)] The Brachistochrone problem is the study of the ‘fastest’ path a small object can take when moving between two fixed points under the influence of gravity. It was analysed and solved by the Bernoulli brothers in the late 17th century; this was the birth of the calculus of variations. Let f be a continuous function of x , $\dot{x} \equiv \frac{dx}{dt}$ and t and consider the problem of finding the extreme values of the function

$$J = \int_a^b f(x, \dot{x}, t) dt.$$

What would you suggest as sensible conditions on optimality? Can you derive and interpret the

first order conditions? What can go wrong?

The next 5 problems are taken from Tom Ferguson's web-page <http://www.math.ucla.edu/~tom/>.

Exc 1.6: The joint distribution of two Bernoulli random variables, X and Y , is determined by the 2 by 2 matrix of probabilities, $p(i, j) = P(X = i, Y = j)$ for $i = 0, 1$ and $j = 0, 1$. Suppose the (prior) distribution of $(p(0, 0), p(0, 1), p(1, 0), p(1, 1))$ is uniform on the set $p(0, 0) > 0, p(0, 1) > 0, p(1, 0) > 0, p(1, 1) > 0, p(0, 0) + p(0, 1) + p(1, 0) + p(1, 1) = 1$. (This is the Dirichlet distribution with parameters 1,1,1,1.) Find the conditional distribution of the $p(i, j)$ given that X and Y are independent (conditionally, given the $p(i, j)$).

Hint: It is natural that two Bayesians should disagree when they disagree on the prior. But when they agree on the prior, how can they disagree on the posterior?

Let $U = p(0, 0) + p(0, 1) = P(X = 0)$, and $V = p(0, 0) + p(1, 0) = P(Y = 0)$. Both Bayesians took a Dirichlet, $D(1, 1, 1, 1)$, as the prior distribution of the $p(i, j)$, and both computed the posterior distribution of U and V given that X and Y are independent, that is given $p(0, 0) = UV$, or $p(0, 0)p(1, 1) = p(0, 1)p(1, 0)$. The marginal distribution of both U and V is the beta distribution, $Be(2, 2)$ (density $6x(1 - x)$ on $(0, 1)$).

One Bayesian discovered a simple way to find the conditional distribution of U and V given that X and Y are independent. One way of generating the $D(1, 1, 1, 1)$ distribution is to choose U from the beta distribution, $Be(2, 2)$, and then choose P and Q are independent uniforms on $(0, 1)$, independent of U . Then, $p(0, 0) = UP, p(0, 1) = U(1 - P), p(1, 0) = (1 - U)Q$, and $p(1, 1) = (1 - U)(1 - Q)$ have the $D(1, 1, 1, 1)$ distribution, and of course, $V = UP + (1 - U)Q$. Here, P represents the conditional probability that Y is 0 given $X = 0$, and Q represents the conditional probability that Y is 0 given $X = 1$. It is clear that X and Y are independent if and only if $P = Q$.

In this form, it is a very easy matter to compute the conditional distribution of U and V given $P = Q$. First of all it is clear that U is not affected when one conditions on $P = Q$ since U is independent of P and Q , and so U still has the $Be(2, 2)$ distribution. Moreover, when $P = Q$, then V is equal to this common value, and the conditional distribution of the common value of two independent uniforms given they are equal is still uniform. Thus, we have answered problem 1: The conditional distribution of U and V given X and Y are independent is that U and V are independent, U is $Be(2, 2)$ and V is uniform on $(0, 1)$.

Or have we? The answer is not symmetric in U and V ! The other Bayesian modeled the joint distribution of the $p(i, j)$ as above starting with V rather than U . He/she then arrived at the conditional distribution of U and V given X and Y are independent as that of independent U and V with V being $Be(2, 2)$ and U being uniform on $(0, 1)$. Of course it is best not to take sides when two Bayesians disagree. But can you help resolve the problem?

Exc 1.7:[Another two envelopes problem] A fair coin is tossed until it comes up heads and the number, X , of tails before the first heads appears is recorded. Then two slips of paper are taken and the number 3 to the power X is written on one of the slips and 3 to the power $X + 1$ is written on the other. One of these two slips is chosen at random (probability $\frac{1}{2}$ each) and given

to you. It will have one of the numbers 1, 3, 9, 27, 81, &c written on it. After observing this number, call it z , you are given the option of exchanging your slip for the other, unseen slip. Then you are paid in dollars the number on the slip you possess. For which values of z would you make the exchange? For which values of z would you make the exchange if you had to pay a 10% commission to the broker who effects the exchange?

Hint: If $z = 1$, you know the other envelope contains 3 so you should exchange. For an arbitrary $z > 1$, we may compute the posterior probability, call it $p(z)$, that the other envelope contains the number $3z$. This is for all $z > 1$

$$p(z) = \frac{P(X = x \mid \text{observe } z = 3^x) = \frac{P(X=x \text{ and observe } 3^x)}{P(\text{observe } 3^x)} = \frac{.5^x \frac{1}{2}}{.5^x \frac{1}{2} + .5^{x-1} \frac{1}{2}} = \frac{1}{3}.$$

Your expected return if you exchange is

$$\frac{2}{3}z + \frac{1}{3}3z = \frac{11}{9}z.$$

Since this is greater than z , you should exchange. Moreover, you should be willing to give up a certain percentage of your winnings, say 10%, in order to make this exchange, since your expected return will still be $\frac{11}{10}z$ which is greater than z . Thus you should make the exchange no matter what z is observed.

But there is something called the sure-thing principle that says that in this case there is no need to look at z . One can make the exchange anyway. But then you should be willing to exchange twice or three times, paying 10% at each exchange. Now you are really losing money. How do you resolve this problem?

Exc 1.8: Kriegspiel (Ferguson, Shapley, and MacQueen 1996) is the game of chess in which the players are not allowed to see the moves of the opponent. Instead, there is a "referee" who keeps track of the each move on his own private board and tells the players if an attempted move is legal. If an attempted move is legal, the player must play it; otherwise, the player may try again. In addition, the referee announces all captures and whether the captured object is a piece or a pawn. In the RAND rules (see below), the referee also announces checks, and possible captures by a pawn.

The question is: Can a blind person, hearing only the announcements of the referee, deduce a mate for one of the players?

Hint: Here is an attempt to answer this question in the affirmative. However, this requires a clarification of the definition of a "No". In the following problem, we use the most liberal definition of a "No" as an attempt by a player to make a move that appears legal only taking into consideration his pieces on his own board, but is in fact illegal according to the actual position as seen by the referee. Some players require, in addition, an attempted move by a player to be "consistent with the current announcement". In fact most players, and also the RAND rules,

consider this stricter definition of a “No” to be in force if the present announcement is a check of some sort. Under this stricter definition, the number of “No”s White receives in trying to make his 10th move must be decreased to, say 14.

Announcement	Announcement	
1	White	Black
2	White	Black
3	White	Black
4	White	Black
5	White	Try Q5
6	Pawn gone, Q4	Try QB6
7	Piece gone, QB3. Try there	Black
8	Check long diagonal	Black
9	Tries, Q7 and KB7	Try QN4
10	Tries, Q7 and KB7. No. No. . . .No. (16 or 17 Nos).	

At which time, a person sitting in the next room listening to the game, shouts, “White to mate in 3”.

The announcements are made according to the RAND rules which are for this problem:

1. Pawn tries (i.e. possible pawn captures) are announced by specifying the square of the capture. Thus the announcement preceding black’s 5th move tells the players that black has a pawn that can make a capture by moving to black’s Q5.
2. Captures are announced by specifying the square and whether a pawn or a piece was captured. Thus the announcement preceding white’s 6th move tells the players that white has lost a pawn that was on his Q4.
3. Checks and the direction of the check (horizontal, vertical, one of the two diagonals or by a knight) are announced. Thus, the announcement preceding white’s 8th move tells the players that white is in check by a diagonally capturing piece (queen, bishop, or pawn) and that the check is along the longer of the two diagonals that meet at the square the white king occupies.
4. There are two other types of ‘No’s. There is the ‘Repeat No’, announced when a player tries a move that has already received a ‘No’, and there is the ‘Hell No’, announced when a player tries a move that is illegal on his own board.

Exc 1.9: Two positive integers are chosen. The sum is revealed to logician A, and the sum of the squares is revealed to logician B. Both A and B are given this information and the information contained in this sentence. The conversation between A and B goes as follows: B starts. B: “I can’t tell what the two numbers are.” A: “I can’t tell what the two numbers are.” B: “I can’t tell

what the two numbers are." A: "I can't tell what the two numbers are." B: "I can't tell what the two numbers are." A: "I can't tell what the two numbers are." B: "Now I can tell what the two numbers are."

What are the two numbers?

Hint: (Ferguson 1983) After B's first announcement, the sum of squares must be attainable in at least two ways, and the possibilities are reduced to a table that begins as follows.

Sum	Sum of Squares
8	50
9	65
10	50
11	65, 85
13	85, 125, 145
14	130, 170
15	125
16	130, 200
17	145, 185, 205
...	...

The only appearances of 50, 65, 85, 125, and 145 as a sum of squares in this table are already displayed. Keep going and you will find one possible answer. But how do you know it's the only possible one?

B's first announcement gives A a lot of information. But when A makes his first announcement, B already knows that A cannot possibly tell what the two numbers are. What good does it do B to listen to A?

Exc 1.10: The outcome of a certain experiment results in three (possibly dependent) random variables, X, Y, Z, each taking values either zero or one. (So there are 8 possible outcomes to the experiment.) Is it possible to set up such an experiment so that

$$P(X = Y) + P(X = Z) + P(Y = Z) < 1?$$

Hint: It is impossible. Since among X, Y, Z, there is bound to be at least two zeros or two ones, at least one of the events X=Y, X=Z, Y=Z is certain to happen. Then since the probability of the union is less than or equal to the sum of the probabilities,

$$1 \leq P(X = Y) + P(X = Z) + P(Y = Z).$$

Or is it? Bell's analysis of the Einstein-Podolsky-Rosen experiment in quantum mechanics shows that things are not quite so simple. Here is a probabilist's view of Bell's paradox

1. There is an actual physical experiment we may perform that will produce any of six (de-

- pendent) random variables, X, Y, Z, X', Y', Z' , each taking values 0 or 1. This experiment may be replicated independently any number of times.
2. We cannot observe more than two of the variables at the same time. In fact, we can only observe one primed and one unprimed variable together.
 3. When we observe X and X' together, they always have the same value, both zero or both one. Probabilistically, this means $P(X = X') = 1$. Similarly, $P(Y = Y') = 1$ and $P(Z = Z') = 1$. (Using this, we may reduce the problem to X, Y, Z , by identifying X with X', Y with Y' , and Z with Z' .)
 4. When X and Y' (= Y) are observed together, we find that they are the same about $\frac{1}{4}$ of the time, and different about $\frac{3}{4}$ of the time. So we may presume $P(X = Y) = \frac{1}{4}$ approximately. Similarly, $P(X = Z) = \frac{1}{4}$ approximately, and $P(Y = Z) = \frac{1}{4}$ approximately.
 5. But we must have $1 \leq P(X = Y) + P(X = Z) + P(Y = Z)$, because at least one of the three events $X = Y, X = Z, Y = Z$ must occur. (This is a special case of Bell's inequality.) So we have

$$1 \leq P(X = Y) + P(X = Z) + P(Y = Z) = \frac{3}{4} \text{ approximately?????!}$$

Comment: This paradox does not depend on the assumption that the experiments are independent and identically distributed. In fact, we may imagine a person writing his own personal choices of X, Y, Z on one slip of paper and of X', Y', Z' on another and giving them in sealed envelopes to two experimenters. Each experimenter then decides which of the three variables to observe. If the experimenter observing X, Y, Z decides to observe X , then Y and Z vanish and are no longer observable, &c. The two experimenters make the decisions of which variables to observe at random, independently (say probability $\frac{1}{3}$ for each variable). Then they compare the results. It always turns out that if X and X' were observed, then $X = X'$; if X and Y' were observed, then $X = Y'$ about $\frac{1}{4}$ of the time; &c. How can the guy writing down the variables do it? He seems to be able to guess whether or not the experimenters are going to choose to observe the same lettered variables. He must have ESP!

Appendix B

Some useful results

1. **Algebra** is the branch of mathematics in which arithmetic relations are generalised and explored by using letter symbols to represent numbers, variable quantities, or other mathematical entities (as vectors and matrices), the letter symbols being combined, especially in forming equations, in accordance with assigned rules.
2. **Analysis** is the branch of mathematics which embraces the theory of functions, the use of limits, continuity, and the operations of calculus.
3. **Geometry** is the branch of mathematics that deals with the measurement, properties, and relationships of points, lines, angles, surfaces, and solids.
4. Some **numbers**: natural numbers: $\mathbb{Z}_+ = \{1, 2, 3, \dots\}$; negative integers: $\mathbb{Z}_- = \{-1, -2, -3, \dots\}$; integers: $\mathbb{Z}_+ \cup \mathbb{Z}_- \cup \{0\}$; rationals: $\mathbb{Q} = \mathbb{Z} \cup \{\frac{p}{q} \mid p, q \in \mathbb{Z}, q > 0\}$; real numbers: $\mathbb{R} = \mathbb{Q} \cup \{\text{irrational numbers}\}$; complex numbers: $\mathbb{C} = \mathbb{R} \cup \{\text{imaginary numbers}\}$.
5. The **Implicit function theorem** is a useful way of finding out how two variables x_1, x_2 are related although there is no closed form solution of the form $x_2 = h(x_1)$. Given, however, is an implicit function that satisfies $f(x_1, x_2) = 0$ which is C^2 and $f_{x_2} \neq 0$. Then

$$\frac{dx_2}{dx_1} = -\frac{f_{x_1}(x_1, x_2)}{f_{x_2}(x_1, x_2)}.$$

It is easy to see why this is true; totally differentiating $f(x_1, x_2) = 0$ gives $f_{x_1}dx_1 + f_{x_2}dx_2 = 0$ which only needs to be rearranged to imply the theorem.

6. **Vieta's rule**:

$$ax^2 + bx + c = 0; x_{1,2} = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

7. The **Hessian** of a scalar function $f(\mathbf{x})$; $\mathbf{x} \in \mathbb{R}^n$

$$\frac{d^2 f}{d\mathbf{x}^2} = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \vdots & & \vdots \\ \frac{\partial^2 f}{\partial x_1 \partial x_n} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix} = H[f(\mathbf{x})]$$

Remark: As a consequence of **Young's Theorem**, the Hessian $H[f(\mathbf{x})]$ is symmetric.

8. The **Jacobian** of a vector function $\mathbf{f}(\mathbf{x})$; $\mathbf{x} \in \mathbb{R}^n$; $\dim \mathbf{f} = m$

$$\frac{d\mathbf{f}(\mathbf{x})}{d\mathbf{x}} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{bmatrix} = \begin{bmatrix} \left(\frac{df_1}{d\mathbf{x}}\right)^T \\ \vdots \\ \left(\frac{df_m}{d\mathbf{x}}\right)^T \end{bmatrix}$$

9. The **transpose** of a matrix

$$(A \cdot B)^T = B^T \cdot A^T$$

$$(A \cdot B)^{-1} = B^{-1} \cdot A^{-1}$$

$$(A^{-1})^T = (A^T)^{-1}$$

10. Lemma for **matrix inversion**

Let \mathbf{u}, \mathbf{v} be column vectors and $A + \mathbf{u}\mathbf{v}^T$ regular:

$$(A + \mathbf{u}\mathbf{v}^T)^{-1} = A^{-1} - \frac{(A^{-1}\mathbf{u})(\mathbf{v}^T A^{-1})}{1 + \mathbf{v}^T A^{-1}\mathbf{u}}$$

11. Lemma for general $A \in \mathbb{R}^{m \times m}, C \in \mathbb{R}^{n \times n}$:

$$(A + BCD)^{-1} = A^{-1} - A^{-1}B(DA^{-1}B + C^{-1})^{-1}DA^{-1}.$$

12. Result

$$\delta^t = x \iff t \ln(\delta) \iff t = \frac{\ln(x)}{\ln(\delta)}$$

13. The **probability** of $r \leq x$ on the uniform distribution with support $[a, b]$ is

$$\frac{x - a}{b - a}.$$

14. The expectation of the **lowest draw** from the set $[a, b]$ among n independent draws is

$$\frac{b - a}{n + 1}.$$

15. The expectation of the **highest draw** from the set $[a, b]$ among n independent draws is

$$\frac{(b - a)n}{n + 1}.$$

16. The expectation of the **second highest draw** from the set $[a, b]$ among n independent draws is

$$\frac{(b - a)(n - 1)}{n + 1}.$$

17. **Fundamental Theorem of Calculus**, $F'(x) = f(x)$:

$$\int_a^b f(x) dx = F(b) - F(a).$$

18. Result

$$\int_a^b x^a dx = \frac{1}{a+1} x^{a+1} + C; a \neq -1.$$

19. The sum of the **infinite geometric power series** $a + ak + ak^2 + \dots$ ($|k| < 1$) is:

$$\frac{a}{1-k}.$$

20. The sum of the **infinite geometric power series with initial term** $ak^n + ak^{n+1} + \dots$ ($|k| < 1$) is:

$$\frac{k^n}{1-k}. \quad (\text{B.0.1})$$

21. The sum of the **finite geometric power series** $a + ak + ak^2 + \dots + ak^{n-1}$ ($k \neq 1$) is:

$$a \frac{1-k^n}{1-k}. \quad (\text{B.0.2})$$

22. The **normal** distribution $f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{1}{2\sigma^2}(x-m)^2)$ with mean m and variance σ^2 has

$$f(m) = \frac{1}{\sqrt{2\pi\sigma^2}}.$$

23. $A \sim N(0, \sigma^2) \implies (A \pm A) \sim N(0, 2\sigma^2)$.

24. **CARA** utility (e.g.: w ... wage, $\psi(e)$... cost of effort e); $r = \frac{u''}{u'}$:

$$-e^{-r(w-\psi(e))}$$

25. All polynomials are continuous.

26. A **permutation** is an arrangement of n elements in a definite sequence. The number of permutations with repetition is $n!$ and without repetition $\frac{n!}{k!}$.

27. A **combination** selects k elements from within a total set of n elements without looking at the sequence. The number of combinations with repetition is $\binom{n}{k}$ and without repetition it is $\binom{n+k-1}{k}$.

28. A **variation** selects k elements from within a total set of n elements with looking at the sequence. The number of variations with repetition is $k! \binom{n}{k}$ and without repetition it is n^k .

29. **Jensen's inequality** says

$$\int_{-\infty}^{+\infty} u(x) dF(x) \leq u\left(\int_{-\infty}^{+\infty} x dF(x)\right).$$

30. The **unit normalisation** of a vector $\mathbf{p} \in \mathbb{R}^n$ is a rule that sets the sum

$$\sum_{i=1}^n p_i = 1.$$

Therefore, we get the normalised vector as

$$\frac{\sum p_i}{\sum p_i} = 1 = \frac{p_1}{\sum p_i} + \frac{p_2}{\sum p_i} + \dots + \frac{p_n}{\sum p_i}.$$

31. Quotient rule

$$\frac{d}{dx} \frac{f(x)}{g(x)} = \frac{f'(x)g(x) - f(x)g'(x)}{g^2(x)}$$

32. Inverse function rule

$$\begin{aligned} \frac{dx}{dy} &= \frac{1}{dy/dx} \\ f^{-1}(x) &= \frac{1}{f'(y)} \end{aligned}$$

33. Exponential function rule

$$\begin{aligned} \frac{d}{dx} e^t &= e^t \\ \frac{d}{dt} \ln(f(t)) &= \frac{f'(t)}{f(t)} \end{aligned}$$

34. Logarithms with base b

$$\frac{d}{dt} \log_b t = \frac{1}{t \ln(b)}$$

35. **Scalar product** of two vector functions $\mathbf{f}(\mathbf{x}), \mathbf{g}(\mathbf{x})$

$$\frac{d[\mathbf{g}(\mathbf{x})^T \mathbf{f}(\mathbf{x})]}{d\mathbf{x}} = \left[\frac{d\mathbf{g}(\mathbf{x})}{d\mathbf{x}} \right]^T \mathbf{f}(\mathbf{x}) + \left[\frac{d\mathbf{f}(\mathbf{x})}{d\mathbf{x}} \right]^T \mathbf{g}(\mathbf{x})$$

36. Product of a vector function $\mathbf{f}(\mathbf{x})$ and a scalar function $g(\mathbf{x})$

$$\frac{d[g(\mathbf{x})\mathbf{f}(\mathbf{x})]}{d\mathbf{x}} = \left[\frac{dg(\mathbf{x})}{d\mathbf{x}} \right]^T \mathbf{f}(\mathbf{x}) + \frac{d\mathbf{f}(\mathbf{x})}{d\mathbf{x}} g(\mathbf{x})$$

37. Linear terms

$$\begin{aligned} \frac{d(A\mathbf{x})}{d\mathbf{x}} &= A; \quad \frac{d(\mathbf{x}^T A)}{d\mathbf{x}} = A \\ \frac{d\mathbf{a}^T \mathbf{x}}{d\mathbf{x}} &= \mathbf{a}; \quad \frac{d\mathbf{x}^T \mathbf{a}}{d\mathbf{x}} = \mathbf{a} \end{aligned}$$

38. Quadratic terms

$$\frac{d(\mathbf{x}^T A \mathbf{x})}{d\mathbf{x}} = (A + A^T) \mathbf{x}$$

39. **Chain rule** for scalar functions

$$\frac{df[\mathbf{y}(\mathbf{x})]}{d\mathbf{x}} = \frac{d\mathbf{y}^T}{d\mathbf{x}} \cdot \frac{df}{d\mathbf{y}}$$

40. **Chain rule** for a vector function

$$\frac{df[\mathbf{z}(\mathbf{w}(\dots(\mathbf{y}(\mathbf{x}))))]}{d\mathbf{x}} = \left(\frac{d\mathbf{f}}{d\mathbf{z}}\right) \left(\frac{d\mathbf{z}}{d\mathbf{w}}\right) \dots \left(\frac{d\mathbf{y}}{d\mathbf{x}}\right)$$

41. Time dependent A, B and \mathbf{x}, \mathbf{y}

$$\frac{d(\mathbf{x}^T \mathbf{y})}{dt} = \dot{\mathbf{x}}^T \mathbf{y} + \mathbf{x}^T \dot{\mathbf{y}}$$

$$\frac{d(A\mathbf{x})}{dt} = \dot{A}\mathbf{x} + A\dot{\mathbf{x}}$$

$$\frac{d(AB)}{dt} = \dot{A}B + A\dot{B}$$

$$\frac{d(\mathbf{x}^T A\mathbf{x})}{dt} = \dot{\mathbf{x}}^T A\mathbf{x} + \mathbf{x}^T \dot{A}\mathbf{x} + \mathbf{x}^T A\dot{\mathbf{x}}$$

42. The **log rule**

$$\int \frac{1}{x} dx = \ln x + C$$

43. The **substitution rule**

$$\int f(x) \frac{dx}{dt} dt = \int f(x) dx = F(x) + C$$

44. **Integration by parts**

$$\int_a^b f(\theta)g'(\theta) d\theta = f(\theta)g(\theta)\Big|_a^b - \int_a^b f'(\theta)g(\theta) d\theta. \quad (\text{B.0.3})$$

45. **Leibnitz'** rule: Let $F(\theta) = \int_{a(\theta)}^{b(\theta)} f(\theta, x) d\theta$, then the derivative of F is given by¹

$$F'(\theta) = f(\theta, b(\theta))b'(\theta) - f(\theta, a(\theta))a'(\theta) + \int_{a(\theta)}^{b(\theta)} \frac{\partial f(\theta, x)}{\partial \theta} dx. \quad (\text{B.0.4})$$

What happens to the value of the integral if the parameter θ is changed a bit? We have three effects:

- (a) The entire curve shifts up or down depending on the derivative $\frac{\partial f}{\partial \theta}$.
- (b) The upper boundary of the integral changes by $b'(\theta)d\theta$. To get the change of the area under the curve, we multiply this by $f(b(\theta))$.
- (c) Analogous for the lower bound.

46. **Taylor's** Theorem says that

$$f(x) = f(0) + \frac{1!}{1!}f'(0)x + \dots + \frac{1}{n!}f^{(n)}(0)x^n + \frac{1}{(n+1)!}f^{(n+1)}(c)x^{n+1} \quad (\text{B.0.5})$$

for some number $c \in [0, 1]$. (The last term is in the Lagrange form of the remainder.)

¹ Take a look at (Feynman 1986) and you will see that this can be actually quite useful.

For the following let a, b, c be constants and x, y random variables.

1. $\mathbb{E}_t[xy] = \mathbb{E}_t[x] \mathbb{E}_t[y] \text{cov}[x, y]$.
2. $\mathbb{E}[ax + by + c] = a\mathbb{E}[x] + b\mathbb{E}[y] + c$.
3. $\text{var}[ax + by + c] = a^2\text{var}[x] + b^2\text{var}[y] + 2\text{cov}[ax, by] = a^2\text{var}[x] + b^2\text{var}[y] + 2ab\text{cov}[x, y]$.
4. $\text{cov}[ax + by, cx + dy] = ac\text{var}[x] + bd\text{var}[y] + (ad + bc)\text{cov}[x, y]$.
5. $\mathbb{E}[e^{cx}] = e^{c^2 \frac{\sigma^2}{2}}$.
6. $\mathbb{E}_t[x_{t+1}] = x_{t+1}$ if x_{t+1} is planned and therefore known.
7. A **random walk** $x_t = x_{t-1} + \varepsilon_t$ is a non-stationary process. It is a special form of a Markov process; Markov processes have the property that $\text{prob}(\text{var}_t = i \mid Z_{t-1}, Z_{t-2}, \dots, Z_0 = i \mid Z_{t-1})$. That is, the last periods' information is sufficient for the whole history.
8. A **martingale** is a stochastic process which has increments that are unpredictable, i.e., the expected value of the increments are unpredictable. If x follows a martingale, then the expected value of x_{t+1} conditional on *all* available time- t information Ω_t (and *not* just the history of x) is x_t

$$E_t[x_{t+1} | \Omega_t] \equiv E_t[x_{t+1}] = x_t.$$

Hence x_t is the best predictor for all future x .

9. **Chebyshev's inequality** says that for $\mathbb{E}[x] = \mu$ and $V[x] = \sigma^2$,

$$P(|x - \mu| \geq \varepsilon) \leq \frac{\sigma^2}{\varepsilon^2}. \quad (\text{B.0.6})$$

10. The **Central Limit Theorem** states that if x_1, \dots, x_n are a random sample from *any* probability distribution with finite mean μ and finite variance σ^2 , and $\bar{x}_n = \frac{1}{n} \sum_i x_i$, then

$$\sqrt{(n)}(\bar{x}_n - \mu) \mapsto^d N[0, \sigma^2].$$

11. **Bernoulli's theorem** (or **weak law of large numbers**). Let $\{X_1, \dots, X_n\}$ be a sequence of independently and identically distributed (iid) random variables, each with mean μ and standard deviation σ . Define a new variable

$$\bar{x} \equiv \frac{X_1 + \dots + X_n}{n}.$$

Then, as $n \rightarrow \infty$, the sample mean \bar{x} tends to the population mean μ

$$\bar{x} = \frac{\mu n}{n} = \mu. \quad (\text{B.0.7})$$

Moreover, the sample variance $\text{var}(X)$ equals

$$\begin{aligned} \text{var}(X) &= \text{var}\left(\frac{X_1 + \dots + X_n}{n}\right) \\ &= \text{var}\left(\frac{X_1}{n}\right) + \dots + \text{var}\left(\frac{X_n}{n}\right) \\ &= \frac{\sigma^2}{n^2} + \dots + \frac{\sigma^2}{n^2} \\ &= \frac{\sigma^2}{n}. \end{aligned}$$

Therefore, by Chebyshev's inequality (B.0.6), for all $\varepsilon > 0$,

$$P(|X - \mu| \geq \varepsilon) \leq \frac{\text{var}(X)}{\varepsilon^2} = \frac{\sigma^2}{n\varepsilon^2}$$

and as $n \rightarrow \infty$,

$$\lim_{n \rightarrow \infty} P(|X - \mu| \geq \varepsilon) = 0.$$

Appendix C

Notational conventions

$;$	such that
wlg	without loss of generality
wrt	with respect to
ae	almost everywhere
$\{ \}$	braces
$[\]$	brackets
$(\)$	parentheses
\sum	sum
\prod	product
\mathbb{R}	the set of real numbers
\mathbb{Z}	the set of integer numbers
\mathbb{N}	the set of natural numbers
\mathbb{Q}	the set of rational numbers
\mathbb{C}	the set of complex numbers
∞	infinity
\equiv	identity; also used for the initial definition of some notation
$a \implies b$	a implies b
$a \impliedby b$	a is implied by b
$a \iff b$	a implies b and a is implied by b (<i>iff</i>)
<i>iff</i>	if and only if
\subset	set inclusion; $A \subset B$ means A is a subset of B (by convention, $A \subset A$ is true)
\supset	set inclusion; $A \supset B$ means B is a subset of A (by convention, $A \supset A$ is true)
\cap	intersection $A \cap B = \{x : x \in A \text{ and } x \in B\}$; on collections $\bigcap_{\alpha \in A} A_\alpha = \{x : x \in A_\alpha, \forall \alpha \in \Lambda\}$
\cup	union $A \cup B = \{x : x \in A \text{ or } x \in B\}$; on collections $\bigcup_{\alpha \in A} A_\alpha = \{x : x \in A_\alpha, \exists \alpha \in \Lambda\}$
	difference is defined as $B \setminus A = \{x \in B : x \notin A\}$
$\exists x$	there exists (at least) one x
$\forall x$	for all x
\vee	or
\wedge	and
\emptyset	empty set; by convention, \emptyset is an element of every set except for the empty set itself
\times	Cartesian product—result is a set of ordered pairs $A \times B = \{(a, b) : a \in A, b \in B\}$
\cdot	scalar (dot) product of 2 N -dimensional vectors \mathbf{a}, \mathbf{b} —result is a scalar $\mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^N a_i b_i$
\mathbf{a}	bold symbols denote vectors
\mathcal{F}	script symbols denote systems of sets
x_1	subscripts denote different components of a vector
x^i	superscripts denote different vectors
A	capital letters denote matrices (or sets)
$[n \times m]$	Matrices are always given in [rows \times columns]

Index

- L_1 -norm, 32
- L_2 -norm, 33
- ℓ_2 -norm, 32
- σ -algebra, 99
- sup-norm, 33
- ε -neighbourhood, 36
- ε -ball, 127
- Lebesgue-integral, 95

- Abelian group, 24
- affine function, 30
- algebraically closed field, 24
- asymptotically stable, 153
- atom, 95
- autonomous, 147, 174

- Banach space, 35
- basis, 48
- Bayes' Rule, 100
- Bayes' Rule, continuous case, 102
- Bayesian rationality, 57
- belief, 57
- Bellman's Principle, 184
- Bernoulli equation, 152
- best reply, 132
- bijective, 30
- bilinear function, 30
- binary operation, 24
- binding constraint, 17
- Bolzano-Weierstrass Theorem, 37
- bordered Hessian, 75
- Borel sets, 99
- Borel-algebra, 99
- bounded, 28, 36, 37

- Brachistochrone problem, 220
- Brouwer's Theorem, 129

- Calculus of Variation, 165
- Cantor set, 90
- CARA utility, 229
- cardinality, 28
- Cartesian plane, 32
- Cartesian product, 32
- Cauchy sequence, 36, 127
- Cauchy-Schwartz inequality, 34
- Central Limit Theorem, 206
- certainty equivalent, 103
- Chebyshev's inequality, 232
- closed, 127
- closed set, 37
- closure, 37
- coefficient of absolute risk aversion, 104
- combination, 31, 229
- commutative ring, 24
- compact, 37, 126
- complete, 36
- complete general solution (ODE), 150
- complete metric space, 127
- concave function, 44
- conditional probability, 100
- cone, 46
- constraint qualification, 16, 77
- contact point, 37
- continuous function, 41, 122, 126
- continuously differentiable, 43
- contraction mapping, 123
- Contraction Mapping Theorem, 123

- control variable, 182
- convergent, 36
- convex hull, 38
- convex set, 38
- correspondence, 47, 127
- costate variable, 186
- countable, 26
- countably infinite, 26
- Cournot duopoly, 139
- Cramer's rule, 50
- cross product, 32
- cumulative distribution function (cdf), 101

- damped oscillations, 158
- Darboux sums, 43
- definite, 52
- definite solution (ODE), 150
- degree (ODE), 147
- depreciation, 166
- derivative, 42
- determinant, 49
- differentiable function, 42
- dimension, 49
- discounting, 166
- discrete topology, 126
- distance, 33
- domain, 29
- dot product, 33
- dual program, 11, 85
- duality, 85
- Duality Theorem, 86
- dynamic consistency, 184
- Dynamic Programming, 165

- eigenspace, 50
- eigenvalues, 50
- eigenvectors, 50
- Envelope Theorem, 60, 64
- equilibrium state, 152
- equivalence relation, 25

- Euclidean distance, 33
- Euclidean inner product, 33
- Euclidean metric, 33
- Euclidean norm, 32
- Euclidean space, 34
- Euler equation, 169
- Euler Theorem, 46
- Euler's Theorem, 47
- expected utility, 103
- expected value, 102
- expenditure minimisation problem, 68
- explosive, 158
- explosive oscillations, 158

- factorial, 31
- field, 24, 98
- field axioms, 98
- finite, 26
- finite geometric power series, 229
- first order ODE, 157
- first-order stochastic dominance (fosd), 105
- fixed point, 121
- function, 29
- Fundamental Theorem of Calculus, 96, 228

- Gaussian diffusion process, 206
- general solution (ODE), 147, 149
- global asymptotic stability, 162
- gradient, 43
- graph, 128
- group, 24

- Hamiltonian function, 186
- Hausdorff space, 35, 126
- Heine-Borel Theorem, 37, 126
- Hessian, 52
- Hicks-Slutsky problem, 69
- Hicksian (compensated) demand, 68
- Hilbert space, 35
- homeomorphism, 121, 122, 126

- homogeneity, 45
- homogeneous equation, 147
- homothetic function, 47
- Hotelling's Lemma, 70
- hyperplane, 55

- ideal, 24
- Implicit function theorem, 227
- Inada conditions, 154
- increasing differences, 137
- indefinite, 52
- indicator function, 96
- indirect utility function, 69
- infimum, 28
- infinite geometric power series, 229
- infinite set, 89
- initial condition, 150
- injective, 29
- inner product, 33
- inner product space, 35
- integral domain, 24
- integrating factor, 152
- integration by parts, 108
- interior point, 37
- Intermediate Value Theorem, 131
- into (onto), 29
- inverse, 30
- isolated point, 37
- isomorphism, 30
- Itô's Lemma, 208

- Jacobian, 228
- Jensen's inequality, 103, 229

- Kakutani's Theorem, 131
- Kuhn-Tucker, 76

- Lagrange Theorem, 74
- Lagrangian, 12, 15
- lattice, 133, 134
- Lebesgue measure, 89

- Leibnitz' rule, 168
- limit, 36
- limit point, 37
- linear (in)dependence, 48
- linear extension, 26
- linear function, 30
- linear program, 11
- linear space, 32
- linear subspace, 48
- Lipschitz continuous, 41
- log-derivative, 194
- lottery, 102
- lower hemi-continuous (lhc), 128
- Lyapunov stable, 153

- Markov process, 232
- Martingale, 232
- Maximum Principle, 165
- mean-preserving spread, 107
- measurable, 93
- measurable function, 97
- measure theory, 89
- metric, 33
- metric space, 35
- MiniMax Theorem, 86
- Minkowski's Theorem, 57
- mixed strategy, 57
- monoid, 24
- monotonic, 31

- Nash equilibrium, 132
- No-Ponzi-game, 198
- nondecreasing, 31
- nonexpansive mapping, 125
- nonincreasing, 31
- norm, 32
- normal distribution, 229
- normed linear space, 34
- null-set, 89

- open, 37, 127

- open cover, 126
- Optimal Control Theory, 165
- order (ODE), 147
- order preserving mapping, 30
- ordinary difference equation (ODE), 145
- ordinary differential equation (ODE), 145
- orthant, 34
- outer measure, 91

- partial ordering, 26
- partial derivative, 43
- partial differential equation (PDE), 145
- particular solution (ODE), 147, 150
- partition, 28
- payoff, 57
- permutation, 31, 229
- Perron-Frobenius Theorem, 50
- phase diagram, 152
- Picard-Lindelöf Theorem, 148
- Pontryagin Maximum Principle, 165
- poset, 26, 133
- power set, 28
- present value, 166
- primary program, 85
- probability density function (pdf), 101
- probability measure space, 100
- pure strategy, 57

- quadratic form, 52
- quasiconcave, 44

- Ramsey model, 189, 193, 194
- random walk, 232
- range, 29
- rank, 49
- Riemann integral, 44, 95
- Riemann sum, 43
- ring, 24, 99
- Roy's Identity, 69
- scalar product, 33
- second order ODE, 159
- second-order stochastic dominance (sosd), 107
- semi-definite, 52
- semigroup, 24
- separable variables, 152
- separating hyperplane, 56
- sequence, 35
- shadow price, 67
- Shephard's Lemma, 68
- simplex, 129
- singleton, 28
- singular matrix, 49
- slack constraint, 17
- Slutsky equation, 69
- smooth, 43
- Solow model, 154
- span, 48
- stable, 158
- state variable, 182
- stationary state, 152
- steady state, 152
- stochastic Calculus, 205
- strict domination, 57
- strictly concave function, 44
- strictly convex set, 38
- subadditive, 41
- sublattice, 134
- subsequence, 36
- substitution rule, 231
- superadditive, 41
- supermodular, 133, 137
- support, 41
- supporting hyperplane, 56
- supremum, 28
- surjective, 30

- Tarski's Theorem, 134
- Taylor's Theorem, 208, 231
- time homogeneous, 147
- topological space, 35, 126

topology, 35, 126
total derivative, 43
trace, 49
transversality condition, 171, 173
triangle inequality, 33

unit normalisation, 230
upper bound, 28
upper contour set, 44
upper hemi-continuous (uhc), 128
upper level set, 44
upper semi-continuous (usc), 127
utility maximisation problem, 67

value function, 62, 165
variation, 229
variations, 31
vector space, 32
Vieta's rule, 227
von Neumann-Morgenstern expected utility,
 102

Weak law of large numbers, 232
Weierstrass Theorem, 42

zero divisor, 24

Bibliography

- ALIPRANTIS, C. D., AND K. C. BORDER (1999): *Infinite Dimension Analysis: A Hitchhiker's Guide*. Springer, New York, 2 edn.
- ARROW, K., AND M. KURZ (1970): *Public investment, the rate of return, and optimal fiscal policy*. Johns Hopkins UP.
- BARON, D. P., AND R. MYERSON (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–930.
- BARRO, R., AND X. SALA-I MARTIN (1995): *Economic growth*. McGraw-Hill.
- BENDER, C., AND S. ORSZAG (1978): *Advanced mathematical methods for scientists and engineers*. McGraw-Hill.
- BERTSEKAS, D. P. (2000): *Dynamic Programming and Optimal Control*. Athena Scientific, Cambridge, Mass., 2 vols, 2nd edn.
- BINMORE, K. (1992): *Fun and Games*. D. C. Heath.
- BIRKHOFF, G. (1979): *Lattice Theory*, Colloquium publications XXV. American Mathematical Society, Providence, Rhode Island, third edn.
- BISMUT, J.-M. (1975): "Growth and the Optimal Intertemporal Allocation of Risks," *Journal of Economic Theory*, 10, 239–57.
- BLANCHARD, O., AND S. FISCHER (1989): *Lectures in macroeconomics*. MITP.
- BORDER, K. C. (1985): *Fixed point theorems with applications to economics and game theory*. Cambridge University Press, Cambridge, U.K.
- CAPIŃSKI, M., AND E. KOPP (1999): *Measure, Integral and Probability*. Springer, London.
- CASS, D. (1965): "Optimum Growth in an Aggregative Model of Capital Accumulation," *Review of Economic Studies*, 32, 233–40.
- CHIANG, A. (1984): *Fundamental methods of mathematical economics*. McGraw-Hill.
- (1989): *Dynamic optimisation*. McGraw-Hill.

- CORNES, R. (1992): *Duality and modern economics*. Cambridge University Press, Cambridge, Mass.
- CRONIN, J. (1980): *Differential equations: Introduction and qualitative theory*. Dekker.
- DENARDO, E. V. (1967): "Contraction mappings in the theory underlying dynamic programming," *SIAM Review*, 9(2), 165–77.
- FERGUSON, T. S. (1983): "Bayesian density estimation by mixtures of normal distributions," *Annals of Statistics*, 1, 209–30.
- FERGUSON, T. S., L. S. SHAPLEY, AND J. B. MACQUEEN (eds.) (1996): *Probability and Game Theory*, vol. 30 of *Lecture Notes-Monograph series*. IMS.
- FERREIRA, P. J. S. G. (1996): "Fixed point problems—an introduction," *Revista do Detua*, 1(6), 505–13.
- FEYNMAN, R. (1986): *Surely you're joking, Mr. Feynman*. Bantam Books.
- FRANKLIN, J. (1980): *Methods of mathematical economics*. Springer.
- FRIEDMAN, M. (1957): *A theory of the consumption function*. Princeton UP.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press, Cambridge, Mass.
- GIBBONS, R. (1992): *A primer in Game Theory*. Harvester Wheatsheaf, Herfordshire.
- HIRSCH, M. W., AND S. SMALE (1974): *Differential equations, dynamical systems, and linear algebra*. Academic Press, Orlando.
- HUMPHREYS, J. F., AND M. Y. PREST (1989): *Numbers, groups and codes*. Cambridge University Press, Cambridge, Mass.
- HUREWICZ, W., AND H. WALLMAN (1948): *Dimension Theory*. Princeton University Press, Princeton.
- KAMIEN, M., AND N. SCHWARTZ (1991): *Dynamic optimization—the calculus of variations and optimal control in economics and management*. North-Holland, 2nd edn.
- KOLMOGOROV, A., AND S. FOMIN (1970): *Introductory real analysis*. Dover.
- KOOPMANS, T. (1965): *On the Concept of Optimal Growth*. in: *The Econometric Approach to Development Planning*, Amsterdam, North Holland.
- KREPS, D. M. (1998): *A Course in Microeconomic Theory*. Harvester Wheatsheaf, London.
- LJUNGQVIST, L., AND T. J. SARGENT (2000): *Recursive Macroeconomic Theory*. Massachusetts Institute of Technology Press, Cambridge, Mass.
- LUCE, R., AND H. RAIFFA (1957): *Games and Decisions*. John Wiley and Sons, London.

- LUENBERGER, D. G. (1984): *Linear and Nonlinear Programming*. Addison-Wesley, Reading, Mass.
- MANKIW, G., D. ROMER, AND D. WEIL (1992): "A contribution to the empirics of economic growth," *Quarterly Journal of Economics*, 107, 407–34.
- MAS-COLELL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic Theory*. Oxford University Press.
- MERTON, R. (1974): "An asymptotic theory of growth under uncertainty," *Review of Economic Studies*, pp. 375–93.
- MILGROM, P. (1999): "The Envelope Theorems," Stanford WP99-016.
- MILGROM, P., AND J. ROBERTS (1990): "Rationalizability, Learning, and Equilibrium in Games with Strategic Complements," *Econometrica*, 58, 1255–77.
- MUNKRES, J. (2000): *Topology*. Prentice-Hall, 2nd edn.
- NASH, J. F. (1950): "Equilibrium points in n -person games," *Proceedings of the National Academy of Sciences*, 36, 48–9.
- NICHOLSON, W. (1990): *Microeconomic Theory: Basic Principles and Extensions*. Dryden Press, London, 7 edn.
- OBSTFELD, M. (1992): "Dynamic Optimization in Continuous-Time Economic Models (A Guide for the Perplexed)," *NBER*, 47.
- O'HARA, M. (1995): *Market microstructure theory*. Blackwell Publishers, Oxford.
- PONTRYAGIN, L. S., V. G. BOLTYANSKII, R. V. GAMKRELIDZE, AND E. F. MISHCHENKO (1962): *The Mathematical Theory of Optimal Processes*. Interscience Publishers.
- RAMSEY, F. (1926): "A mathematical theory of saving," *Economic Journal*, 38, 543–59.
- ROYDEN, H. L. (1988): *Real analysis*. Macmillan, 3rd edn.
- RUDIN, W. (1975): *Principles of mathematical analysis*. McGraw-Hill, 3rd edn.
- SEIERSTAD, A., AND K. SYDSÆTER (1987): *Optimal control theory with economic applications*. North-Holland.
- SENGUPTA, J., AND P. FANCHON (1997): *Control theory methods in economics*. Kluwer.
- SIMON, C. P., AND L. BLUME (1994): *Mathematics for Economists*. Norton, New York.
- STARR, R. M. (1997): *General Equilibrium Theory*. Cambridge University Press, Cambridge, Mass.

- STOKEY, N., AND E. LUCAS, ROBERT (AND PRESCOTT (1989): *Recursive Methods in Economic Dynamics*. HUP.
- SUNDARAM, R. K. (1996): *A first course in optimisation theory*. Cambridge University Press.
- SYDSÆTER, K., AND P. HAMMOND (1995): *Mathematics for economic analysis*. Prentice Hall.
- TARSKI, A. (1955): "A lattice-theoretical fixedpoint theorem and its applications," *Pacific Journal of Mathematics*, 5, 285–309.
- TOPKIS, D. M. (1976): "The Structure of Sublattices of the Product of n Lattices," *Pacific Journal of Mathematics*, 65, 525–32.
- VARIAN, H. R. (1992a): *Economic and Financial Modeling with Mathematica*. Telos, Springer, New York, 2nd edn.
- (1992b): *Microeconomic Analysis*. W. W. Norton, London, 3 edn.
- VIVES, X. (1990): "Nash Equilibrium with strategic complementaities," *Journal of Mathematical Economics*, 19, 305–21.
- WEIBULL, J. (1995): *Evolutionary Game Theory*. MIT Press, Cambridge, Mass.
- WEIBULL, J. (2001): *Lecture notes in Mathematics*. Stockholm School of Economics.
- WHEEDEN, R., AND A. ZYGMUND (1977): *Measure and Integral: An Introduction to Real Analysis*. Dekker.
- WOLFRAM, S. (1999): *The Mathematica Book*. Cambridge University Press, 4th edn.