The Microeconomics of Risk and Information

Richard Watt

This book is dedicated, with my deepest gratitude, to my wife Marta and our children, Daniel and Olivia, for making life a pleasure, and hard work worthwhile.

List of Figures

Preface

This is a book about microeconomic theory. More specifically, it is a book concerning the way the presence of risk affects optimal decision making. The study of decision making under risk is certainly not new, but due to certain mathematical complexities that it involves, it is often left out of undergraduate microeconomics programs on the grounds of being too difficult for students to manage. However, the increase in complexity that is introduced in models with risk is often as much due to increases in dimensionality as the addition of risk to the modelling environment. The basic modelling that is required in order to solve problems in decision making under risk is no different to what is done under an assumption of certainty, so long as the dimensionality of the problem is not altered. All that needs to be done is to re-interpret the basic elements of the model – the variables, and the graphical curves and lines that are used to analyse the problem. With that in mind, this book offers a short course in choice under risk, packaged in exactly the same environment as a typical undergraduate course in choice under certainty, that is, an environment with two choice variables (at most). The principal idea is to show students how to handle scenarios with risk, and to point out some of the mathematical toolkit that is useful in that environment (indeed, in microeconomic theory generally), without actually leaving the comfort zone of a simple two-dimensional graphical setting.

As with most text books, the material presented here derives from a fairly long history of teaching the subject. Over the past 20 years or so, I have taught this information to students in their final year of undergraduate economics on two continents. The problems that arise are the same everywhere, and those problems generally involve a difficulty in visualising microeconomic problems in mathematical guise. In essence, what is of issue is not a lack of mathematical ability or knowledge, but rather a shortfall in the understanding of microeco-

nomic processes. Once a student can see what a problem involves, and how it should be tackled graphically, then it is a relatively easy step to apply the correct mathematical techniques to it. The underlying theme of the present book is to attempt to achieve this by sticking rigorously with problems in only two dimensions, and showing as much as possible both mathematical and graphical treatments side by side.

I am indebted to a great many individuals both for fostering my own interest in the topic of the microeconomics of decision making under risk, and for turning my rough-and-ready lecture notes into what I hope is now a coherent and sensible treatment of the topic. I was initially lured by problems in choice under risk by the late Prof. Richard Manning in classes that were taught at the University of Canterbury some 25 years ago. Since then, the main impetus to my interest in the topic has come from the many vibrant discussions that are so typical at the annual meetings of the European Group of Risk and Insurance Economists (EGRIE), which I habitually attend. I owe a huge debt of gratitude to Jasper Mackenzie who took on the arduous task of preparing so professionally the graphs that appear in the book. I also thank Nick Sanders who helped me with an earlier set of graphs, which allowed deadlines to be reached. Aleta Bezuidenhout and Jaime Marshall at Palgrave Macmillan have been a pleasure to work with.

Contents

Chapter 1

Introduction

The standard theory of choice that is taught in all introductory and intermediate microeconomic theory courses posits a consumer who would like to make a choice of how much of each of two goods to consume, given the prices of those goods and a level of wealth that can be dedicated to the purchase. This is a typical constrained optimisation problem – the choice variables are the amounts of each good to consume, the objective is to maximise welfare, and the constraints are determined by, on the one hand, the fact that neither good can be consumed in negative amounts and, on the other, the restriction that the cost of the choice (quantities demanded multiplied by prices) cannot exceed the financial resources available.

In spite of its radical simplicity, as a theoretical construction, this standard consumer choice model is able to provide logically persuasive solutions to questions of some importance. For example, the model predicts (always) that welfare is increasing in wealth and decreasing in prices, and that if wealth is increased in compensation for a price increase then the demand for each good is decreasing in its own price and increasing in the price of the other good. The model also predicts (usually, but not always) that the uncompensated demand for each good is increasing in wealth and decreasing in its own price. Any number of other results can also be obtained, related to such things as changes in preferences, introduction of taxes of different types, and even non-linear pricing.

However there are aspects of the model that many students find to be overly simplified. Perhaps the simplification that is most often noted is the fact that the model is usually presented in only two

dimensions. That is, the assumption is that there are only two goods present in the choice problem. This, however, should not be a concern. The restricted number of dimensions is in place only in order that the visual apparatus of a graph can be used. There is not doubt that a graphical exposition of the solution helps enormously to capture most of the essential elements of the solution, and for that reason two-dimensional analysis is often used. But the model itself is robust to an extension to any number of goods, and indeed it is often solved in its multi-dimensional version in more advanced courses.

The other most often cited simplification that is important for the model to be a faithful representation of real-world decision making is the fact that everything that the decision maker needs to know, he does know. In particular, he is fully informed of the availability of all goods, of the prices of all goods, and, of course, of his own income and preferences. It is likely that none of these things are really quite so certain. Prices and availability of goods differ over sellers, and it is often very difficult (or at least, very costly) to know exactly where to go to get any particular item at any particular price. Even personal attributes such as the disposable income and preferences of the decision maker are known only approximately. One way to deal with income uncertainty might be to set a budget for purchases that is small enough to be guaranteed to be available, and then any surplus income that results is simply retained as a random element of savings. But then, we should ask what would be the optimal size of the consumption budget that should be established? More generally, we would be better to enquire about how the risks and uncertainties that undoubtedly surround a decision-making environment can be best catered for. This is the underlying theme of this book.

Risk is an ever-present element in decision making. It is often related to time, because the final consequences of the decisions that we make often do not occur simultaneously with the decision. Between the moment of the decision, and the moment of the consequence, other random elements in the problem environment might be playing out, affecting the consequences of our decisions. That is, a given decision can, feasibly, lead to more than one outcome or consequence, depending on the outcome of other relevant stochastic elements. What we need, therefore, is a convincing theory of how to best take such stochastic elements into account when the decisions are made.

One obvious way in which the existence of risk affects economic transactions is the existence of markets and institutions in which

risk itself can be traded. Take, for example, the insurance industry, which clearly offers a service designed to shift risk from insurance consumers to insurance companies, in exchange for a premium payment. However, many other examples exist, including (but certainly not restricted to) markets for financial products like shares in businesses, and, of course, options and futures on those shares, contracts between employers and employees that shift risks from the former to the latter, fixed rather than variable interest rate contracts that shift risk from borrowers to lenders, and so on. Achieving an understanding of how such markets and institutions work to the mutual benefit of all concerned, and how they affect decision making, is a fundamental purpose of this book.

1.1 Focus of the book

This book contains a short course, designed to be completed in a single semester of study, in the economic theory of risk and information. These two intimately related topics are now standard inclusions in the economic theory curriculum at universities all over the world. This, of course, reflects the now generally recognised importance of risk and information as integral aspects of almost any economic analysis. Knowing how to handle risky, or stochastic, environments, and above all, how to deal with scenarios in which the parties to a transaction have different information sets, is of primordial importance in the education of economists.

That said, the supply of specialist text books designed to cater to the need to learn about risk and information has typically been restricted to texts at the post-graduate level. This is a natural course of events since the norm has been to teach economics undergraduates the standard theory of consumer and producer choice, equilibrium and markets, all under certainty, and then to move onto the extension to stochastic environments only in post-graduate courses. At most, undergraduates will have seen a single chapter in their general microeconomics text on choices under risk, and another for the economics of asymmetric information.

While it is true that risk should be studied only after successfully following a course in choice under certainty, it has increasingly been the case that final year undergraduates are offered a one-semester elective in the economics of risk and information, but as yet there has been no specific text book that caters to such a course. The present text is an attempt to fill that gap.

1.2   Basic objectives

The book has several objectives. First and foremost, it offers students
a minimal content of topics in the economics of risk and information.
However, the book has also been designed to be able to be studied in a
single semester course, with perhaps between 24 and 36 lecture hours
only. Thus some selection of possible topics has taken place, and I
hope that the final choice of included topics is a fair reflection of what
the profession has deemed, by revealed preference, to be important.

The primary objective of the book is to provide understanding
rather than to simply inform. This is a very difficult task, as anyone
who has ever attempted to lecture a theory topic will attest. How-
ever, in order to fulfil the objective of comprehension, with only one
or two exceptions, the book sticks entirely with a two-dimensional
setting, one that should be intimately familiar to any student of
microeconomics who has completed at least an introductory 101-
type course. In particular, I have purposefully avoided the technique
of providing mathematical analysis at n-dimensional level and only
illustrative examples in two dimensions. By having full correspondence
between the two-dimensional analysis and the two dimensional graphs
throughout, a student gets two looks at each and every critical point
that is brought up, and it is hoped that this eventually leads to greater
understanding rather than just learning.

The second fundamental objective of the book is comprehension
of the use of constrained optimisation techniques in microeconomic
theory generally. Again, by retaining a strict two dimensional analysis
throughout, it is hoped that students will ultimately see that what is
being taught in this book is really no different to what was taught in,
say, consumer theory under certainty. All that has happened is that
the axes of the graphical environment have been re-labeled to mea-
sure different (all-be-it very similar) variables, the budget constraint
has been re-interpreted, and the indifference curves correspond to a
particular case of utility. Noting the huge similarity between how risk
is handled and how certainty is analysed is an important step in fully
understanding the robustness of the standard consumer theory setting
to the analysis of different problems. What is more, it is hoped that
from the present text, if it had not already been noted, the student
can clearly relate a graph of indifference curves and a choice set to
a problem of constrained optimisation, and the characteristics of the
solution point (a tangency on the frontier of the choice set) to the first

and second order conditions of that optimisation problem.

Third, the book works simultaneously with mathematical treatments of problems and their graphical representations. Once again, this is achievable only by retaining a two dimensional setting throughout. All too often it is the case that economics students are asked to provide mathematical analyses of problems without really fully understanding what that mathematics is doing. Students are taught that utility functions should be concave, however the relationship between concavity of utility and the shape of the indifference curves is often not understood. Indeed, second-order conditions of maximisation problems generally are poorly covered and often not always grasped by students. It is hoped that the technique used in this book will help to ease the burden of moving into a fully mathematical analysis of economic problems, by providing such an analysis *side-by-side* with a full graphical intuition for a series of models.

Given the objective of providing a gentle mathematical treatment of microeconomic problems, a certain degree of mathematical sophistication is assumed throughout. However, at no point will any mathematical technique be used that is not a standard inclusion in high school mathematics the world over. Indeed, in terms of mathematics, all that is needed is the ability to derive (derivatives of the first and second degree), and algebraic manipulation of equations. In any case, in Appendix A, a set of important mathematical tools is carefully explained. It is recommended that readers who are not confident at applying mathematics to problem solving should begin by carefully reading the mathematical appendix (Appendix A), and that this appendix be consulted whenever an unfamiliar mathematical technique appears. Above all, this is a thinking course, and as such you can get a very long way if you search for economic intuition and logic in the results that are derived mathematically.

1.3 Content and structure

The book is divided into two main parts and several chapters. Part I deals with individual decision making (i.e., scenarios in which there is a single decision maker active in the model) and Part II deals with scenarios with two decision makers. The chapters themselves are ordered such that a coherent story is able to be told. The story in question is about a decision maker who is exposed to scenarios of ever-increasing complexity in which choices must be made. To start

with, the decision maker lives in a world in which the only person
of relevance is himself, just like Robinson Crusoe living alone on his
tropical island. His choices and decisions are made in an environment
in which other important things may change – the weather, the tides,
the appearance of ships on the horizon – but those other changes are
not controlled directly by any other decision maker. They are, as it
were, acts of nature. The objective of these decisions is to provide the
decision maker with the greatest possible welfare (or utility), given
the fact that some other important values have yet to be fixed. Part
I of this book deals with this type of single-person decision problem.

Then, another decision maker turns up. Just like Friday, whom
Robinson Crusoe meets on the island. Now, with two decision makers
on the island, a small economy emerges in which meaningful transac-
tions can take place between the two. In Part II of this book we look
at how these small economies may work in as much as risk sharing
goes. In particular, we are interested in how the two individuals can
join together in an effort to affront the risks that they face, the risks
posed by the whims of nature. The main thing at this intermediate
stage of existence is that both Robinson and Friday are both fully
informed about the exact nature of the risks that they face. They
both know what outcomes would result under each and every feasible
state of nature, and (importantly) the likelihood of each and every
state of nature.

The climax of the story is when, perhaps after some time on
the island, Friday begins to understand that there is a fundamental
difference between himself and Robinson. Their "information sets" are
different, and this will have a profound effect upon the way they work
together. Perhaps we can think that Robinson (as the master) is busy
writing the story of his adventures, and so he sends Friday (as the
servant) off to labour each day in the jungles and oceans to get food
for the two of them. Assume, for example, that Robinson really likes
to eat fish, and that Friday is happy to eat only fruit. Fruit is easily
available and in plentiful supply all over the island, and so there is no
problem about gathering all the fruit that the two may ever require.
But fishing is different. It is inherently risky, and the outcome of how
many fish are caught depends upon many random factors. Maybe it
turns out that the best place to fish is a cove that is very far away,
and Friday would rather not walk so far, and instead he prefers to fish
at a closer location in spite of it not being such a plentiful supplier
of fish. Robinson, who does not want to have to accompany Friday

fishing each day to see where he goes, must think up an agreement with Friday that convinces Friday that indeed he should go to the far away fishing spot, in spite of the personal costs to him of getting there. After all, when Friday comes back in the evening with only one fish rather than many after having fished at the closer spot, he could just tell Robinson that he was at the far away cove but that it was just a poor day for fishing. Working out exactly how Robinson should go about convincing Friday to fish where Robinson would like him to rather than where Friday would prefer is the final part of our book. It is the point to which all of the earlier work leads, as it will happen that Robinson cannot solve the informational problem with Friday without appealing to his previous experiences, first alone and then with Friday but under common information sets.

With that in mind, in the next chapter, a detailed investigation into the very concept of risk is provided, together with what we know about how the existence of risk should be incorporated into preferences, and ultimately into decision making. Chapter 3 sets out a more detailed analysis of the concept of "risk aversion", which shapes all of the rest of the topics that are analysed in the book. Once risk aversion is clearly analysed, Chapter 4 moves on to look at a series of applications of decision making under risk that has been the subject of economists' attention. Specifically, in Chapter 4 the reader will find the applications of portfolio demand, insurance demand, precautionary savings, and producer theory.

These first four chapters of the book, grouped together in Part I, deal with situations in which there is only one active party in the model, making optimal decisions in a risky environment.[1] The last three chapters (grouped as Part II) bring a second player into the model. It is at this point that the assumption of "risk" rather than "uncertainty" becomes more important. The risk assumption implies that both of the parties to the transactions that we analyse agree upon the probabilities of the different outcomes. This assumption is made only because the models then become more user-friendly, in the sense that otherwise they would become excessively cluttered. If the two parties had different probability beliefs in a two-dimensional

---

[1] The possible exception is the case of insurance demand, where one could argue that not only the insurance consumer, but also the insurer is present. However, when we analyse the insurance demand model, our primary attention is placed upon the decision of the consumer, and the insurer is really present only as a parameter set in the demander's decision problem.