The Analytics of Uncertainty and Information
Second Edition

SUSHIL BIKHCHANDANI
University of California, Los Angeles

JACK HIRSHLEIFER

JOHN G. RILEY
University of California, Los Angeles

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

32 Avenue of the Americas, New York, NY 10013-2473, USA

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9780521541961

First published 1992
Second edition 2013

Printed in the United States of America

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging in Publication Data
Bikhchandani, Sushil.
The analytics of uncertainty and information / Sushil Bikhchandani, University of California – Los Angeles, Jack Hirshleifer, John G. Riley, University of California, Los Angeles. – Second Edition.
pages cm
Includes bibliographical references and index.
ISBN 978-0-521-83408-7 (hard covers) – ISBN 978-0-521-54196-1 (pbk.)
1. Uncertainty. 2. Equilibrium (Economics) 3. Decision making. 4. Information theory in economics. I. Hirshleifer, Jack. II. Riley, John G. III. Title.
HB615.H568 2013
339.5–dc23 2013005525

ISBN 978-0-521-83408-7 Hardback
ISBN 978-0-521-54196-1 Paperback

Contents

* Starred sections represent more difficult or specialized materials that can be omitted without significant loss of continuity.

* Starred sections represent more difficult or specialized materials that can be omitted without significant loss of continuity.

* Starred sections represent more difficult or specialized materials that can be omitted without significant loss of continuity.

Acknowledgments

The number of people who have made helpful contributions is very great indeed. In particular, we are indebted to the following current or former UCLA graduate students for their assistance: Svend Albaek, Nigel Chalk, Jen-Wen Chang, Alex David, Maxim Engers, Deborah Frohman, Gerry Garvey, Geoffrey Hosta, Kwanho Kim, Son Ku Kim, Fernando Losada, Wai-Nang Po, and Eduardo Siandra. Jen-Wen Chang wrote solutions to the Exercises and Excursions at the end of each chapter. Special thanks to Daniel A. Graham for informing us of errors he discovered in using the book as a course text.

In Memoriam: Jack Hirshleifer

The origins of this book go back to the mid-1970s at UCLA, when Jack and I decided to team teach a course on uncertainty and information. The first year we sat in on each other's lectures, and I certainly learned a lot, not just about how to extract insights from formal economics models, but also how to teach the material. After some years we were invited to write on these topics for the *Journal of Economic Literature*, and the first edition grew out of this effort. We both worked hard on the project, but even when I was satisfied, Jack demanded more. If an argument was not immediately transparent, it was sent back to the drawing board.

As a result of our joint project we became friends, and it was always a pleasure to join Jack for lunch at UCLA or in one of our homes. Even at the end of his life, when his body was ravaged by sickness, Jack's mind was as curious as ever. His range of interests stretched to the boundaries of economics and beyond. I never fully lost my sense of awe.

Originally, we were going to work on the second edition together with our colleague and friend "Bikh." Our goal was to add new applications and insights. Alas, Jack left us too soon. Given the clarity of much of the first edition, many of the original chapters have been only lightly edited. Therefore Jack's words and wisdom continue to resonate. It will be up to the reader to decide if we have retained the level of clarity demanded by Jack. Certainly, this was our aspiration.

John G. Riley

Introduction

The Economics of Uncertainty and Information

All human endeavors are constrained by our limited and uncertain knowledge – about external events past, present, and future; about the laws of nature, God, and man; about our own productive and exchange opportunities; about how other people and even we ourselves are likely to behave. Economists have, of course, always recognized the all-pervasive influence of inadequate information, and its correlate of risk, on human affairs. But only in the period after the Second World War did an accepted *theory of uncertainty and information* begin to evolve. This theory provides a rigorous foundation for the analysis of individual decision making and of market equilibrium, under conditions where economic agents are unsure about their own situations and/or about the opportunities offered them by market dealings.

With recent explosive progress in the analysis of uncertainty, the topic can no longer be described as neglected. Nor have the advances been "merely academic." The economic theory of uncertainty and information now flourishes not only in departments of economics but also in professional schools and programs oriented toward business, government and administration, and public policy. In the world of commerce, stock market analysts now regularly report measures of share-price uncertainty devised by economic theorists. Even in government and the law, formal analysis of uncertainty plays a role in dealing with issues like safety and health, allowable return on investment, and income distribution.

Unfortunately, these new advances have not always taken a form comprehensible to the general economic reader. Brilliant intellectual progress often appears in erratic and idiosyncratic guise; novel terminologies, approaches, and modes of thought can easily hamper understanding. That has certainly been the case here. Even specialists in some areas of the economics of uncertainty and information often find it hard to grasp the import of closely related research originating from a slightly different angle. As a related point, early explorers may have mistaken the part for the whole – a foothill for the mountain, an outlying peninsula for the mainland. Specifically, some scientific contributions that have appeared under ambitious titles like "the economics of information" or the "economics of uncertainty" actually deal only with tiny portions of those large subjects.

We view our task mainly as one of integration: unifying these important though partial new results and concepts into a satisfying single picture. We would not want to claim that our own view of the whole is the only one logically possible or useful. But we believe that it is an outlook with many appealing and satisfying features: (1) it goes far in de-mystifying the topic; (2) with certain significant exceptions, it provides a natural taxonomy for most of the major problems that have been studied; and (3) most important of all, our approach makes it clear that the economics of uncertainty and information is not a totally new field utterly disconnected from previous economic reasoning, but is rather a natural generalization and extension of standard economic analysis.

A fundamental distinction is between the *economics of uncertainty* and the *economics of information*. In the economics of uncertainty, each person adapts to his or her given state of limited information by choosing the best "terminal" action available. In the economics of information, in contrast, individuals can attempt to *overcome* their ignorance by "informational" actions designed to generate or otherwise acquire new knowledge before a final decision is made. Put another way, in the economics of uncertainty the individual is presumed to act on the basis of *current fixed beliefs* (e.g., deciding whether or not to carry an umbrella in accordance with one's present estimate of the chance of rain). In the economics of information, a person typically is trying to arrive at improved beliefs – for example, by studying a weather report or by looking at a barometer before deciding to take the umbrella.

Another crucial element is *strategic uncertainty*. If there are a large number of individuals, then each acts as price-taker. In contrast, in economic interactions between only a few individuals, each individual may have an appreciable impact on the terms of trade through his or her actions. There are gains from behaving strategically. Consequently, in addition to possibly limited knowledge about preferences and endowments of others, each individual cares about, and is uncertain about, actions other individuals may take. There is strategic uncertainty. The best course of action available to individual *A* depends on what individual *B* might do, and vice versa. Game-theoretic reasoning cuts through this morass.

The sequence of topics in this book is guided by the pedagogical principle of advancing from the easy to the difficult, from the familiar to the more strange and exotic. Part I deals with terminal actions only – the economics of uncertainty. The first three chapters analyze the optimal risk-involved decisions of the individual. Chapter 4 moves on to the market as a whole, showing how the overall equilibrium that determines the prices of risky assets also distributes social risks among all individuals in the economy.

Part II turns to the economics of information and to strategic uncertainty. Starting with a discussion of the value of better information in Chapter 5, we then explore the effect of autonomously *emergent* information upon the market equilibrium solution (Chapter 6). The issue of information leakage via changes in asset prices is also considered.

In preparation for analyzing strategic uncertainty, Chapter 7 provides an introduction to game theory. The standard Nash equilibrium concept often produces multiple equilibria, some of which seem intuitively implausible. Chapter 7 reviews various efforts to refine the notion of equilibrium. Chapter 8 then analyzes contracting between two agents, one of whom has only imperfect information about the other's preferences (hidden knowledge) or is unable to observe the other's behavior (hidden actions). The former condition leads to *adverse selection* in markets while the latter results in *moral hazard*. Chapter 9 examines market equilibrium under adverse selection.

In Chapter 10 we analyze auctions and other market mechanisms. Issues that arise when interactions among agents are repeated over long or indefinite time periods are considered in Chapter 11. We end with an analysis of information transmission, acquisition, and aggregation in Chapter 12.

Our mode of exposition is highly eclectic. "Literary" reasoning, geometrical demonstration, and analytical proofs are all employed from time to time – as called for by the nature of the topic, by the psychological need for variety, and by our desire to illustrate all the major forms of economic argument arising in these contexts. In addition, certain more advanced topics are separated from the main text in specially marked starred sections that can be omitted with minimal loss of continuity. Finally, mixed with the more purely formal portions of our analysis will be applications to important real-world phenomena such as insurance, securities markets, corporate financial structures, the use of experts and agents, group decisions where returns and risks are shared, and the value of education.

Over the last 20 years, game-theoretic reasoning has become widespread in economics. Therefore, in this second edition, we have placed greater emphasis on game theory. Consequently, most of the changes are in part II of the book (although every chapter has at least some modifications to improve the logical flow of material). The chapter on game theory (Chapter 7) has been rewritten and appears earlier. New topics in Part II include posted-price markets, mechanism design, common-value auctions, and the one-shot deviation principle for repeated games. Chapter 12 is entirely new; the results on information aggregation and acquisition that are described here were published after the first edition.

PART I

PART I

1

Elements of Decision under Uncertainty

We introduce a model for decision making under uncertainty that will be our workhorse throughout the book. Uncertainty is modeled with a set of states of nature, one of which will occur. The decision maker or individual has a probability distribution over the states of nature that represents his (or her) subjective beliefs about the likelihood of different states of nature. This individual chooses actions and actions have consequences. The consequence for the individual depends on the state of nature and his choice of action. The states of nature are represented in a way that the probabilities of states are unaffected by the individual's actions. The individual's preferences over consequences are captured by a utility function. The probability distribution over states of nature and the utility function over consequences, both of which are subjective,[1] are combined by the expected-utility rule to induce an expected utility over actions.

An individual must choose among *acts* – or synonymously, he or she must make *decisions*, or select among *actions*, *options*, or *moves*. And, where there is uncertainty, nature may be said to "choose" the *state of the world* (or *state*, for short). You decide whether or not to carry an umbrella; nature "decides" on rain or sunshine. Table 1.1 pictures an especially simple 2×2 situation. Your alternative acts $x = (1, 2)$ are shown along the left margin, and nature's alternative states $s = (1, 2)$ across the top. The body of the table shows the *consequences* c_{xs} resulting from your choice of act x and nature's choice of state s.

[1] Subjective in the sense that another individual, faced with the same decision problem, may have a different probability distribution and a different utility function: beliefs and tastes may differ from person to person.

Table 1.1. *Consequences of alternative acts and states*

		States	
		$s = 1$	$s = 2$
Acts	$x = 1$	c_{11}	c_{12}
	$x = 2$	c_{21}	c_{22}

More generally, the individual under uncertainty will, according to this analysis, specify the following elements of his decision problem:

(1) a *set of acts* $(1, \ldots, x, \ldots, X)$ available to him;
(2) a *set of states* $(1, \ldots, s, \ldots, S)$ available to nature;
(3) a *consequence function* c_{xs} showing outcomes under all combinations of acts and states.

And, in addition:

(4) a *probability function* $\pi(s)$ expressing his beliefs (as to the likelihood of nature choosing each and every state);
(5) an *elementary-utility function* $v(c)$ measuring the desirability of the different possible consequences to him.

We will explain below how the "expected-utility rule" integrates all these elements so as to enable the individual to decide upon the most advantageous action. Put another way, we will show how the economic agent can derive a personal preference ordering of his possible *acts* from his given preference scaling over *consequences.*

COMMENT: The approach here does not allow for the psychological sensations of vagueness or confusion that people often suffer in facing situations with uncertain (risky) outcomes. In our model, the individual is neither vague nor confused. While recognizing that his knowledge is imperfect, so that he cannot be sure which state of the world will occur, he nevertheless can assign exact numerical probabilities representing his degree of belief as to the likelihood of each possible state. Our excuse for not picturing vagueness or confusion is that we are trying to model economics, not psychology. Even the very simplest models in economic textbooks, for example, indifference-curve diagrams, implicitly postulate a degree of precise self-knowledge that is descriptively unrealistic. The ultimate justification, for indifference-curve diagrams or for theories of decision under uncertainty, is the ability of such models to help us understand and predict behavior.

1.1 The Menu of Acts

There are two main classes of individual actions: *terminal* moves versus *informational* moves. Here, in Part I of the book, we consider a simplified world where only terminal acts are available, so that the individual is limited to making the best of his or her existing combination of knowledge and ignorance. An example of terminal action under uncertainty is the statistical problem of coming to a decision on the basis of sample evidence now in hand: for instance, when a regulatory agency has to decide whether or not to approve a new drug on the basis of experimental test results. We will be considering terminal actions of this type, and especially the risk-involved decisions of *individuals in markets*: whether or not to purchase insurance, to buy or sell stocks and bonds, to participate in a partnership, etc. Anticipating a bit, a key theme of our analysis will be that markets allow decision makers to share risks and returns in ways that accord with the particular preferences and opportunities of the different transactors.

Part II of the book will be covering *informational* actions – decisions concerning whether and how to improve upon one's state of knowledge before making a terminal move. In the class of informational actions would fall statistical choices such as how much additional evidence to collect before coming to a terminal decision, what sampling technique to employ, etc. Once again, our emphasis will be on ways of acquiring new information *through markets*. Knowledge can be acquired by direct market purchase – by buying newspapers for weather and stock market reports, by undergoing a course of training to gain "know how" in a trade, or by employing an expert for private advice. Rather less obviously, markets open up an indirect means of acquiring information: for example, a person can observe the market choices of better-informed traders, or might draw inferences from people's reputations acquired in the course of their previous market dealings. Or, a producing firm might imitate other commercially successful firms. But these interesting phenomena involving information-involved actions will have to be set aside until Part II.

1.2 The Probability Distribution

We assume that each person is able to represent his beliefs as to the likelihood of the different states of the world (e.g., as to whether nature will choose rain or shine) by a "subjective" probability distribution (Savage, 1954). Assuming discrete states of the world, the individual is supposed to be able to assign to each state s a degree of belief, in the form of numerical weights

π_s lying between zero and one inclusive, and summing to unity: $\Sigma_s \pi_s = 1$. In the extreme case, if the person were certain that some particular state s would be occurring, the full probabilistic weight of unity would be assigned to that state. Then $\pi_s = 1$, so that zero probability is attached to every other state in the set $1, \ldots, s, \ldots, S$. More generally, a high degree of subjective assurance will be reflected by a relatively "tight" probability distribution over the range of possible states; a high degree of doubt would be reflected by a wide dispersion.

At times, we shall find it will be more convenient to assume that the variable or variables defining the state of the world vary continuously (rather than discretely) so that the number of distinct states is uncountably infinite. Here the probability of any exact single state coming about is regarded as zero ("infinitesimal"), although the event is not *impossible.* Making use of a continuous state-defining variable s, where s can be any real number between 0 and S, the individual's subjective probability beliefs would be represented by a probability density function $\pi(s)$ such that $\int_0^S \pi(s)\ ds = 1$.

1.2.1 Risk versus Uncertainty

A number of economists have attempted to distinguish between risk and uncertainty, as originally proposed by Frank H. Knight (1921, pp. 20, 226). (1) "Risk," Knight said, refers to situations where an individual is able to calculate probabilities on the basis of an *objective* classification of instances. For example, in tossing a fair die the chance of any single one of the six faces showing is exactly one-sixth. (2) "Uncertainty," he contended, refers to situations where no objective classification is possible, for example, in estimating whether or not a cure for cancer will be discovered in the next decade.

In this book, we disregard Knight's distinction. For our purposes, risk and uncertainty mean the same thing. It does not matter, we contend, whether an "objective" classification is or is not possible. For we will be dealing throughout with a "subjective" probability concept (as developed especially by Savage, 1954): probability is simply *degree of belief.* In fact, even in cases like the toss of a die where assigning "objective" probabilities appears possible, such an appearance is really illusory. That the chance of any single face turning up is one-sixth is a valid inference *only if the die is a fair one* – a condition about which no one could ever be "objectively" certain. Decision makers are therefore never in Knight's world of risk but instead always in his world of uncertainty. That this approach, assigning probabilities on the basis of subjective degree of belief, is a workable and fruitful procedure will be shown constructively throughout the book.[2]

1.2.2 "Hard" versus "Soft" Probabilities

While we do not distinguish between what Knight called risk and uncertainty, he was getting at – though imperfectly expressing – an important and valid point. In his discussion, Knight suggested that a person's actions may well depend upon his "estimate of the chance that his estimates are correct," or, we shall say, upon his *confidence in his beliefs*. This brings us to a distinction between "hard" versus "soft" probability estimates.

Suppose that for purposes of an immediate bet you had to estimate the probability of heads coming up on the next toss of coin A – the coin having been previously tested many times by you and found to have historically come up heads and tails with just about equal frequency. If you are a reasonable person, you would assign a degree of belief (subjective probability) of about 0.5 to heads, and you would be rather confident about that number. In contrast, imagine instead that you are dealing with coin B, about which you know absolutely nothing. You have not even been able to inspect it to verify whether it is possibly two tailed or two headed. Nevertheless, if you *had* to pick some single number you would be compelled again to assign 0.5 probability to heads coming up on the next toss, since as a reasonable person you lack any basis for a greater or lesser degree of belief in heads than tails. But, your *confidence* in the 0.5 figure for coin B would surely be much less.

It is not the psychological sensation of confidence or doubt that interests us, but the possible implications for decisions. If the same probability assignment of 0.5 will be made either way, as has just been argued, is there any action-relevant difference between the two cases? For our purpose, the answer is NO, if you are committed to *terminal* action.[3] If you must bet now on the basis of your current information, 0.5 is the relevant probability for guiding your choice of heads or tails. In either situation, you have no grounds for thinking heads more likely or tails more likely. But the answer is YES, there is indeed a difference between the two situations if you have the

[2] See Schmeidler (1989) for the foundations of an alternative approach that explicitly models Knightian uncertainty and individuals' attitudes to it.

[3] Later in this chapter, we describe the Ellsberg paradox, which is an experiment indicating that individuals may react differently to hard and soft probabilities in a setting with terminal actions.

option of *informational* action. When this option is available, you should be more willing to invest money or effort to obtain additional information about coin B than about coin A. In short, greater prior doubt (lesser degree of confidence) makes it more important to acquire additional evidence before making a terminal move. So, we see that a person's *informational* actions, though not his *terminal* actions, do depend upon his confidence in his beliefs – in Knight's language, upon his "estimate of the chance that his estimates are correct." Confidence will be an important topic in Part II of the book, where we cover the economics of information, but will not be involved in our more elementary treatment of the economics of uncertainty in Part I.

Exercises and Excursions 1.2

1 Consistency of Probability Beliefs

An individual believes that credible information will soon arrive in the form of news about the probability of rain. He believes there is a 50% chance that the news will be "rain certain," a 30% chance that the news will be "no rain," and a 20% chance that the news will be "rain with probability 0.5." Is this consistent with his currently believing that the odds in favor of rain are 2:1?

2 Information and Confidence

In terms of the chances of a coin coming up heads, suppose there are three states of the world regarded as possible:

State 1: chance of heads is 100% [coin is two headed]
State 2: chance of heads is 50% [coin is fair]
State 3: chance of heads is 0% [coin is two tailed].

An individual initially assigns equal probabilities $(\pi_1, \pi_2, \pi_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ to all three states.

(A) For an immediate bet (terminal action), what is his best estimate for the probability p of heads on the next toss?
(B) Suppose new information were now to change his probability vector to $(\pi_1, \pi_2, \pi_3) = (0, 1, 0)$. What can you now say about his best estimate for p? What has happened to his *confidence* in that estimate?
(C) Same question if, instead, the new information changed his probability vector to $(\frac{1}{2}, 0, \frac{1}{2})$.

1.3 The Utility Function

As shown in Table 1.1, each *consequence* is the outcome of an economic agent's choice of action combined with nature's "choice" of the state of the world. In principle, the consequence is a full description of all aspects of the individual's environment resulting from such an interaction. For example, if someone decides not to carry an umbrella and nature chooses rain, the consequences might include getting wet, being late for work, and a variety of other discomforts. But we shall mainly be concerned with consequences describable in terms of alternative *baskets of consumption goods* that enter into individuals' utility functions. Very frequently we shall deal with an even simpler picture in which consequences take the form of entitlements to a single summary variable like monetary *income.*

Consequences might be quantities that are certain, or might themselves be probabilistic – depending upon how states of the world are described. If the states are defined deterministically, as in "Coin shows heads," and supposing the action chosen was "Bet $1 at even money on heads," then the consequence would be "Win one dollar." But states of the world can sometimes be defined as probabilistic processes. The relevant states might be "Coin has 50% chance of coming up heads" versus "Coin is biased to have 75% chance of coming up heads." Here the act "Bet on heads" will be reflected, in either state of the world, by an uncertain consequence taking the form of a specified chance of winning the dollar.

We shall use the notation $v(c)$ to represent a person's *utility function* (or *elementary-utility function*) over the consequences c.

1.4 The Expected-Utility Rule

Utility attaches directly to consequences, and only derivatively to actions. To emphasize this distinction, we shall use the notation $U(x)$ for a person's derived preference ordering over actions x. The expected-utility rule is used to derive $U(x)$ from the utility function $v(c)$, as explained below.

A CRUCIAL DISTINCTION
$v(c)$ is a utility function defined over consequences
$U(x)$ is the expected-utility function defined over actions

The analytical problem is to explain and justify this derivation, that is, to show how, given his direct preferences over *consequences*, the individual can order the desirability of the *actions* available to him.

To choose an act is to choose one of the rows of a consequence matrix like Table 1.1. As the individual is also supposed to have attached a probability (degree of belief) to the occurrence of every state, each such row can be regarded as a probability distribution. We may therefore think of a person as choosing among probability distributions or "prospects." A convenient notation for the "prospect" associated with an act x, whose uncertain consequences $c_x = (c_{x1}, c_{x2}, \ldots, c_{xS})$ are to be received with respective state probabilities $\pi = (\pi_1, \pi_2, \ldots, \pi_S)$ – the probabilities summing, of course, to unity – is:

$$X \equiv (c_{x1}, c_{x2}, \ldots, c_{xS}; \pi_1, \pi_2, \ldots, \pi_S)$$

The crucial step is to connect the $v(c)$ function for consequences with the utility ordering $U(x)$ of acts. We can take this step using the famous "expected-utility rule" of John von Neumann and Oskar Morgenstern (1944, pp. 15–31):

EXPECTED-UTILITY RULE

$$\begin{aligned} U(x) &\equiv \pi_1 v(c_{x1}) + \pi_2 v(c_{x2}) + \cdots + \pi_S v(c_{xS}) \\ &\equiv \sum_{s=1}^{S} \pi_s v(c_{xs}) \end{aligned} \tag{1.4.1}$$

This says that the expected utility $U(x)$ of act x is calculable in an especially simple way: to wit, as the mathematical expectation (the probability-weighted average) of the elementary utilities $v(c_{xs})$ of the associated consequences. Note that Equation (1.4.1) is simply additive over states of the world, which means that the consequence c_{xs} realized in any state s in no way affects the preference scaling $v(c_{xs^0})$ of consequences in any other state s^0. Equation (1.4.1) is also linear in the probabilities, another very specific and special functional form. As the von Neumann–Morgenstern expected-utility rule is absolutely crucial for our theory of decision under uncertainty, we shall be devoting considerable space to it.

It turns out that the expected-utility rule is applicable *if and only if the $v(c)$ function has been determined in a particular way that has been termed the assignment of "cardinal" utilities to consequences.* More specifically, the proposition that we will attempt to explain and justify (though not rigorously prove) can be stated as follows:

Given certain "postulates of rational choice," there is a way of assigning a cardinal utility function $v(c)$ *over consequences* such that the expected-utility rule determines the individual's preference ranking $U(x)$ *over actions.*

A "cardinal" variable is one that can be measured quantitatively, like altitude, time, or temperature. While different measuring scales might be employed, such scales can diverge only in zero-point and unit-interval. Temperature, for example, can be measured according to the Celsius or the Fahrenheit scales; 32° Fahrenheit is 0° Celsius, and each degree up or down of Celsius is 1.8° up or down of Fahrenheit. Similarly, altitude could be measured from sea level or from the center of the Earth (shift of zero-point) and in feet or meters (shift of unit-interval). Cardinal variables have the following property: regardless of shift of zero-point and unit-interval, the relative magnitudes of *differences* remains unchanged. The altitude difference between the base and crest of Mount Everest exceeds the difference between the foundation and roof of even the tallest man-made building – whether we measure in feet above sea level or in meters from the center of the Earth.

In dealing with certainty choices, standard economic theory treats utility (intensity of preference) as an *ordinal* rather than a cardinal variable. The individual, it is postulated, can say "I prefer basket A to basket B." He is not required to quantify *how much* he prefers A to B. Put another way, if any given utility function in the form of an assignment of cardinal numbers to consequences (consumption baskets) correctly describes choices under certainty, so will any ordinal (positive monotonic) transformation of that function. Suppose that, for choices not involving risks, some scale u of cardinal numbers was attached as preference labels to consequences – where, of course, higher u indicates greater level of satisfaction. Then any positive monotonic transformation of those numbers would lead to the same decisions. For example, suppose an individual always prefers more consumption income c to less. Then we might say, "He is trying to maximize the function $u = c$." But the income level that maximizes u also maximizes $\log u$ or e^u, both of which are positive monotonic transformations of u. So $u = e^c$ or $u = \log c$ could equally well have served to indicate the preference scaling. More formally, if u is a satisfactory function for choices under certainty, then so is $\hat{u} \equiv F(u)$, provided only that the first derivative is positive: $F'(u) > 0$.

In contrast, when it comes to choices under *uncertainty*, the expected-utility rule is applicable only if the utility function $v(c)$ has been constructed in a particular way that provides fewer degrees of freedom. In fact, as will shortly be seen, given any initially satisfactory $v(c)$ function, only the *cardinal* (positive linear, rather than positive monotonic) transformations of $v(c)$ will leave preference rankings unchanged. Formally, if $v(c)$ satisfactorily describes the individual's choices under uncertainty, then so does $\hat{v} = \alpha + \beta v$, where α is any constant and β is any positive constant.

Why are all the positive monotonic transformations of the utility function permissible in the riskless case, while only the positive *linear* transformations are allowed when it comes to risky choices? In the absence of uncertainty, deciding upon an action is immediately equivalent to selecting a single definite consequence. It follows that if someone can rank *consequences* in terms of preferences he has already determined the preference ordering of his *actions* – which is all that is needed for purposes of decision. But in dealing with risky choices it is not immediately evident how a ranking of consequences leads to an ordering of actions, since each action will in general imply a probabilistic mix of possible consequences. The great contribution of von Neumann and Morgenstern was to show that, given plausible assumptions about individual preferences, it is possible to construct a $v(c)$ function – "cardinal" in that only positive linear transformations thereof are permissible – whose *joint* use with the expected-utility rule (1.4.1) will lead to the correct ordering of actions.

1.4.1 An Informal Presentation

To formally justify the joint use of a cardinal utility function and the expected-utility rule, for dealing with choices among risky prospects, involves a somewhat higher order of technical difficulty. What follows here is an informal presentation (based mainly upon Schlaifer, 1959) illustrating, by direct construction, how the required type of utility function can be developed.

For the purpose of this discussion, assume that the consequences c are simply amounts of income a person might receive. Let m represent the worst possible consequence (the smallest amount of income) that can occur with positive probability, and M the best possible consequence (the largest amount of income). More income is preferred to less – so the individual already has, to begin with, an *ordinal* utility scale. The problem is to "cardinalize" this scale, that is, to show that there is a way of assigning numerical values (arbitrary only with respect to zero-point and unit-interval) to the degrees of preference associated with all levels of income. These values must be rising with income, else they would not be consistent with the given ordinal preference ("more is preferred to less"). But the chosen scale must also lead to correct answers when used in conjunction with the expected-utility rule. The method we shall employ to establish such a cardinal scale is called "the reference-lottery technique."

Consider any level of income c^* between m and M. Imagine that the individual is faced with the choice between c^* and some "reference lottery" having the form $(M, m; \pi, 1 - \pi)$ in prospect notation. That is, he has a choice between c^* for certain versus a gamble yielding the best possible outcome M with probability π and the worst possible outcome m with probability $1 - \pi$. We shall suppose that the individual can say to himself: "When π becomes very close to unity, I surely will prefer the gamble; for lotteries with π very close to zero, I surely prefer the certainty of c^*. Consequently, in between there must be some intermediate probability π^* of success in the reference lottery, such that I am exactly indifferent between the certain income c^* and the prospect $(M, m; \pi^*, 1 - \pi^*)$." After due introspection, we assume, the individual can in fact specify this π^*. We may set $v(m) = 0$ and $v(M) = 1$. Then the π^* so derived is a cardinal measure of the utility of income level c^* for him. That is: $v(c^*) = \pi^*$.[4] Or, more elaborately:

$$v(c^*) \equiv U(M, m; \pi^*, 1 - \pi^*) \equiv \pi^* \tag{1.4.2}$$

An individual proceeding to assign cardinal preference values to income in this way will generate a $v(c)$ function over the range $m \leq c \leq M$, which can be employed with the expected-utility rule (1.4.1) to order his choices among actions.

Figure 1.1 illustrates a hypothetical individual situation. Let $m = 0$ and $M = 1{,}000$ (in dollars, we can suppose) be the extremes of income that need be considered. For the specific income $c^* = 250$, the person's success-in-equivalent-reference-lottery probability is assumed to be $\pi^* = 0.5$ – meaning that he finds himself indifferent between a sure income of \$250 and a 50% chance of winning in an income lottery whose alternative outcomes are \$1,000 or nothing. Then the utility assigned to the sure consequence \$250 is just $\frac{1}{2}$ – that is, $v(250) = 0.5$, determining the location of point Q on the $v(c)$ curve. Repeating this process, the reference-lottery technique generates the entire $v(c)$ curve between $m = 0$ and $M = 1{,}000$.

A full justification, showing why this particular procedure works to derive a suitable cardinal scale, requires a more formal analysis (to be touched on in Section 1.4.2). But we can give a geometric intuition here. The essential point is that the $v(c)$ measure obtained via the reference-lottery technique is in the form of a *probability*, so that the expected-utility rule (1.4.1) becomes equivalent to the standard formula for compounding probabilities.

[4] Because shifts of zero-point and unit-interval are permissible for cardinal scaling, more generally we can write $v(c^*) = \alpha + \beta\pi^*$, for arbitrary α and $\beta > 0$. This is equivalent to assuming $v(m) = \alpha$ and $v(M) = \alpha + \beta$. We will henceforth ignore this uninteresting generalization.

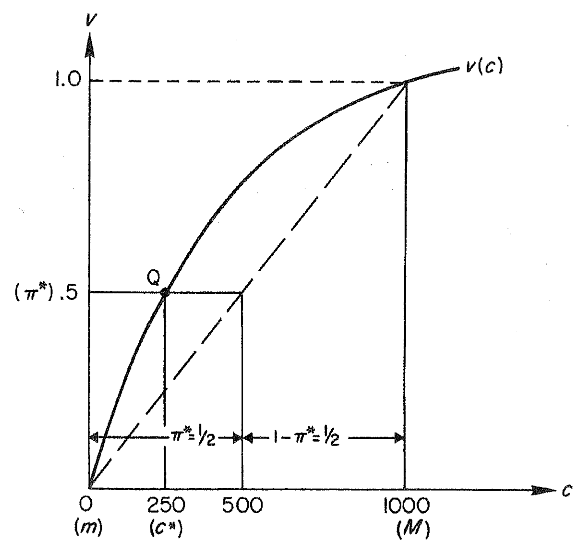

Figure 1.1. Utility function.

A Geometric Interpretation

The expected-utility rule is equivalent to the assumption that indifference curves over lotteries are parallel straight lines. To see this, consider lotteries over three possible consequences m, c^*, and M. Any lottery

$$x = (m, c^*, M; \hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3) = (m, c^*, M; \hat{\pi}_1, 1 - \hat{\pi}_1 - \hat{\pi}_3, \hat{\pi}_3)$$

may be represented as a point in (π_1, π_3) space; see Figure 1.2. The triangle ABC is the set of all possible lotteries with outcomes m, c^*, and M. Point A corresponds to getting c^* for sure, point B is M for sure, and point C is m for sure. In the lottery x, the probabilities $\hat{\pi}_1$ and $\hat{\pi}_3$ (of outcomes m and M, respectively) are the coordinates of the point x. The probability of outcome c* in this lottery, $\hat{\pi}_2$, is the horizontal (or equivalently vertical) distance from point x in Figure 1.2 to the hypotenuse BC of the triangle. The expected utility of x is:

$$\begin{aligned} U(x) &= \hat{\pi}_1 v(m) + (1 - \hat{\pi}_1 - \hat{\pi}_3) v(c^*) + \hat{\pi}_3 v(M) \\ &= v(c^*) - \hat{\pi}_1 v(c^*) + \hat{\pi}_3 (1 - v(c^*)) \end{aligned}$$

where we substitute $v(m) = 0$ and $v(M) = 1$.

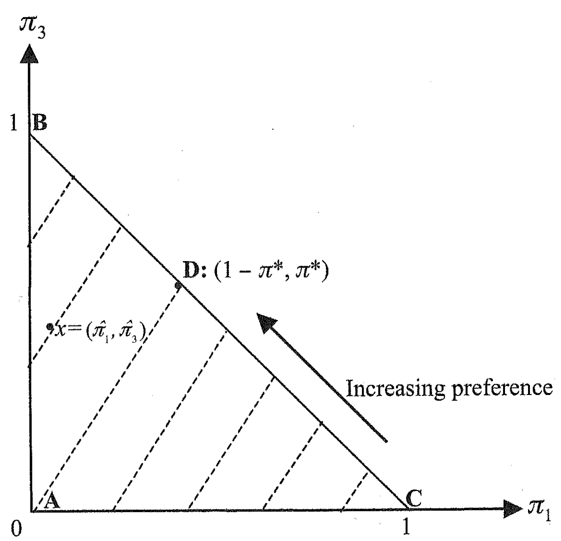

Figure 1.2. Indifference curves under expected-utility rule.

Let $y = (m, c^*, M; \pi_1, 1 - \pi_1 - \pi_3, \pi_3)$ be any other lottery. A similar calculation shows that $U(y) = v(c^*) - \pi_1 v(c^*) - \pi_3(1 - v(c^*))$. If lottery y yields the same expected utility as lottery x, then:

$$\begin{aligned} U(y) &= v(c^*) - \pi_1 v(c^*) + \pi_3(1 - v(c^*)) \\ &= v(c^*) - \hat{\pi}_1 v(c^*) + \hat{\pi}_3(1 - v(c^*)) = U(x) \end{aligned}$$

Re-arranging this we have:

$$\pi_3 = \frac{U(x) - v(c^*)}{1 - v(c^*)} + \frac{v(c^*)}{1 - v(c^*)}\pi_1$$

This is the equation of the (straight line) indifference curve through x. Observe that the slope of the indifference curve, $v(c^*)/(1 - v(c^*))$, does not depend on the lottery x. Thus, all indifference curves are parallel straight lines (shown as broken lines in the Figure 1.2). The indifference lines have positive slope because (i) any rightward movement from x leads to a *less* desirable lottery as it corresponds to increasing π_1 at the expense π_2 and (ii) any upward movement from x leads to a *more* desirable lottery as it corresponds to increasing π_3 at the expense π_2. The direction of increasing preference is northwest, as indicated by the arrow in Figure 1.2.

Point D on the line segment BC in Figure 1.2 corresponds to the lottery $(m, c^*, M; 1-\pi^*, 0, \pi^*)$. This lottery has expected utility π^*. As $v(c^*) = \pi^*$, point D is on the same indifference line as point A.

Example 1.1: Imagine that an individual finds that his reference-lottery utilities over the range $0 \leq c \leq 1{,}000$ satisfy the specific utility function $v(c) = (c/1{,}000)^{1/2}$. (This formula is consistent with the previously obtained point $v(250) = 0.5$ in Figure 1.1.) Suppose he is now offered a choice between option A, representing \$250 for certain once again, and option E taking the form of a three-way prospect: E = (810, 360, 160; 0.1, 0.5, 0.4). Which should he choose?[5]

We already know that $v(250) = 0.5$: option A is equivalent to a reference lottery with 50% chance of success. For the elements of option E, we can readily compute: $v(810) = 0.9$, $v(360) = 0.6$, and $v(160) = 0.4$. That is, in option E the high possible payoff of \$810 is equivalent in preference to a reference lottery with 90% chance of success, the middling payoff \$360 is equivalent to a 60% chance of success, and the poor payoff \$160 to a 40% chance of success. Now we ask ourselves: What is the *overall* equivalent probability of success associated with option E? We can simply compute it by using the rule for compounding probabilities:

$$0.1(0.9) + 0.5(0.6) + 0.4(0.4) = 0.55$$

So prospect E offers, overall, the equivalent of a 0.55 chance of success in the reference lottery whereas option A was equivalent only to a 0.5 chance of success. Evidently, option E is better. The key point is that the equation leading to the 0.55 number, which we presented as the familiar formula for compounding probabilities, is also an instance of applying the expected-utility rule (1.4.1). □

In short, the prescribed way of determining a cardinal $v(c)$ function for use with the expected-utility rule *makes it possible to interpret each $v(c)$ value as a probability* – to wit, the equivalent chance of success in a standardized reference lottery – and therefore to use the laws of compounding probabilities for determining the desirability of more complicated prospects.

A few additional comments:

1. We have been assuming here that consequences take the form of simple quantities of income. More generally, each consequence c might be a

[5] The prospect E cannot be represented in Figure 1.2. Only prospects that yield \$0, \$250, or \$1000 are depicted in Figure 1.2.

basket (vector) of consumption goods. The same technique can be employed so long as the individual has an *ordinal* preference scaling of baskets (an indifference map) to begin with.

2. We have also assumed that the same $v(c)$ scale is applicable in each and every state of the world. But, if the states are defined as "rain versus shine," or "healthy versus sick," it might appear that attitudes toward income and income risks, as reflected in the $v(c)$ function, could differ from state to state. We shall see in Chapter 2, under the heading of "state-dependent utilities," how this difficulty can be handled.
3. Some people find it disturbing that the additive form of the expected-utility rule (1.4.1) excludes any "complementarities," positive or negative, between consequences in different states. For example, if consequences are simple incomes, a higher or lower income in any state s^o is supposed in no way to affect the $v(c)$ number assigned to income received in any other state s^*. The reason is simple: incomes in the distinct states s^o and s^* can never be received *in combination* but only as *mutually exclusive alternatives.* There can be no complementarity where no possibility of jointness exists.
4. There can be confusion over whether or not the von Neumann–Morgenstern analysis proves that utility is "really" cardinal rather than ordinal. Some of the difficulty stems from a mix-up between the $v(c)$ and the $U(x)$ functions. The cardinality restriction applies to the $v(c)$ function – the preference scaling over *consequences.* But we are ultimately interested in the utility rankings of alternative *actions,* and when it comes to actions any ordinal transformation of an acceptable utility measure will always serve equally well. Suppose, for example, that use of the reference-lottery technique provides the needed utility function $v(c)$ such that an individual's *actions* (prospects) are correctly ordered by the expected-utility formula $U(x) = \Sigma_S \pi_S v(c_S)$. Then any positive monotonic transformation of $U(x)$, such as $\hat{U}(x) = e^{U(x)}$, would provide an equally correct ordering of the *actions.* Observe that if $U(x) > U(y)$ then $\hat{U}(x) = e^{U(x)} > e^{U(y)} = \hat{U}(y)$.
5. We have emphasized that the von Neumann–Morgenstern analysis justifies this particular method of constructing a cardinal $v(c)$ scale only when jointly used with the expected-utility rule. Correspondingly, the expected-utility rule has not been "proved" to be true. All that has been shown is that there exists a way of constructing a $v(c)$ function that *makes* the expected-utility rule valid as a way of deriving preferences as to actions from given preferences as to consequences.

1.4.2 The Independence Axiom

We are not providing here a formal proof of the expected-utility rule. Instead, our objective is to clarify the crucial element in the proof, the principle of *non-complementarity of incomes in different states* (see Comment 3 above). The formal postulate expressing this principle, the Independence Axiom, is also known as the substitution axiom or the axiom of complex gambles.

Independence Axiom: Suppose an individual is indifferent between two actions or prospects x and y. Then, for any other prospect z and any fixed probability p, he will be indifferent between a first complex lottery in which he receives x with probability p and z otherwise, versus a second complex lottery yielding y with probability p and z otherwise. Moreover, if he strictly prefers x over y, he will strictly prefer the first complex lottery. Thus, using the symbol $\sim$ to indicate indifference and the symbol $\succ$ for strict preference:

If $x \sim y$, then $(x, z; p, 1 - p) \sim (y, z; p, 1 - p)$
If $x \succ y$, then $(x, z; p, 1 - p) \succ (y, z; p, 1 - p)$

This axiom would be violated if, in a complex prospect, the presence of z differentially affected the attractiveness of x relative to y – i.e., if there were any complementarity effect. It might seem this could happen if, say, x and y were amounts of ordinary commodities like bread and margarine and z were a commodity like butter (since butter is a consumption complement for bread but a substitute for margarine). However, in the complex prospects or lotteries dealt with here, positive or negative complementarity has no role. The rationale behind this is that the occurrence of x in the one case or of y in the other rules out z. An individual can never simultaneously enjoy both *x and z* together, or both *y and z.*

An immediate implication of this axiom is that, for two lotteries x and y such that $x \sim y$, we can *substitute* one for the other in any prospect in which either appears, without changing the relative preference ordering of prospects.

In the reference-lottery process, the $v(c)$ associated with any income level c was determined by finding the probability of success in the reference lottery equally preferred to that income, i.e.:

If $c \sim (M, m; \pi, 1 - \pi)$, then, because $v(m) = 0$ and $v(M) = 1$, we have $v(c) = \pi$

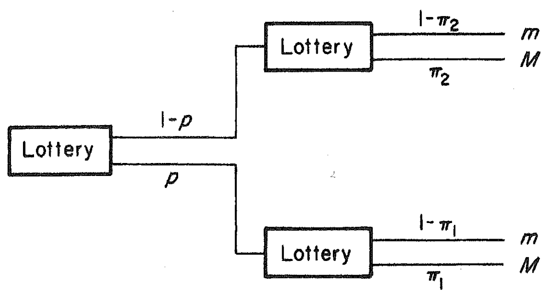

Figure 1.3. Tree diagram of compound lottery.

In what follows, it will be helpful to introduce the notation $l^*(\pi)$ to represent a reference lottery in which M is the outcome with probability π and m is the outcome with probability $1-\pi$:

$$l^*(\pi) = (M, m; \pi, 1-\pi)$$

Thus, if $c \sim l^*(\pi)$ then $v(c) = \pi$.

Consider now two levels of income c_1 and c_2 and their equivalent reference lotteries $l^*(\pi_1)$ and $l^*(\pi_2)$. Then $v(c_1) = \pi_1$ and $v(c_2) = \pi_2$. Suppose we wanted to find the preference equivalent of a lottery $(c_1, c_2; p, 1-p)$ involving consequences c_1 and c_2 with respective probabilities p and $1-p$. Using the ability to *substitute* preference-equivalent prospects:

$$c_1 \sim l^*(\pi_1) \Rightarrow (c_1, c_2; p, 1-p) \sim (l^*(\pi_1), c_2; p, 1-p)$$

Moreover:

$$c_2 \sim l^*(\pi_2) \Rightarrow (l^*(\pi_1), c_2; p, 1-p) \sim (l^*(\pi_1), l^*(\pi_2); p, 1-p)$$

Combining these implications:

$$(c_1, c_2; p, 1-p) \sim (l^*(\pi_1), l^*(\pi_2); p, 1-p) \tag{1.4.3}$$

The lottery on the right-hand side of (1.4.3) is depicted as a "tree diagram" in Figure 1.3. Each box or "node" represents a point at which nature makes a move. Outcomes are indicated at the end of each branch of the tree.

At the initial node, nature "chooses" probabilistically between the two reference lotteries. Then, depending on this choice, one of the reference lotteries is played. Note that there are only two outcomes of this compound lottery, M and m. Adding probabilities, outcome M is reached with probability $p\pi_1 + (1-p)\pi_2$. Then the compound lottery is itself equivalent to a reference lottery:

$$(l^*(\pi_1), l^*(\pi_2); p, 1-p) = l^*(p\pi_1 + (1-p)\pi_2) \tag{1.4.4}$$

Combining (1.4.3) and (1.4.4) it follows that the individual is indifferent between $(c_1, c_2; p, 1-p)$ and a reference lottery in which the probability of success is $p\pi_1 + (1-p)\pi_2$. As $\pi_1 \equiv v(c_1)$ and $\pi_2 \equiv v(c_2)$, it follows that:

$$\begin{aligned} U(c_1, c_2; p, 1-p) &= p\pi_1 + (1-p)\pi_2 \\ &= pv(c_1) + (1-p)v(c_2) \end{aligned}$$

Thus, the independence axiom, which formalizes the principle of non-complementarity of income over states of the world, leads directly to the von Neumann–Morgenstern expected-utility rule.

Exercises and Excursions 1.4

1 *Transformation of Preferences*

An individual claims to be maximizing:

$$U = (1 + c_1)^{\pi_1}(1 + c_2)^{\pi_2}$$

where $(c_1, c_2; \pi_1, \pi_2)$ is a two-state prospect (which means that $\pi_1 + \pi_2 = 1$). Is he a von Neumann–Morgenstern expected-utility (EU) maximizer? Would all his decisions be consistent with those of an EU maximizer?

2 *Indifference Curves in Consequence Space*

(A) If the utility function is $v(c) = c^{\frac{1}{2}}$, where c is income, suppose a person's preference ordering over actions or prospects in a two-state world is given by:

$$U(c_1, c_2; \pi_1, \pi_2) = \pi_1 (c_1)^{\frac{1}{2}} + \pi_2 (c_2)^{\frac{1}{2}}$$

Depict the indifference curves in a diagram with c_1 on the horizontal axis and c_2 on the vertical axis (probabilities held constant). Show that each indifference curve touches the axes and is everywhere bowed toward the origin.

(B) If $U = \sum_1^2 \pi_s v(c_s)$ and $v(.)$ is a strictly concave function, show that if the individual is indifferent between (c_1, c_2) and (c_1', c_2') he will strictly prefer the convex combination $(\lambda c_1 + (1-\lambda)c_1',\ \lambda c_2 + (1-\lambda)c_2')$. Hence draw a conclusion about the shape of the indifference curves in the (c_1, c_2) plane.

3 *The Expected-Utility Rule*

Let $v(c)$ be the utility functions for certain outcomes. Then, for lotteries of the form $(c_1, c_2; \pi_1, \pi_2)$, we have seen that:

$$U(c_1, c_2; \pi_1, \pi_2) = \sum_{s=1}^{2} \pi_s v(c_s)$$

In this exercise, you are asked to generalize this result to lotteries with three outcomes. An inductive argument can then be used to show that for any lottery $(c_1, c_2, \ldots, c_s; \pi_1, \pi_2, \ldots, \pi_s)$:

$$U(c_1, \ldots, c_s; \pi_1, \ldots, \pi_s) = \sum_{s=1}^{S} \pi_s v(c_s)$$

(A) Consider the lottery:

$$\hat{l} \equiv \left(c_1, c_2; \frac{\pi_1}{\pi_1 + \pi_2}, \frac{\pi_2}{\pi_1 + \pi_2} \right)$$

Explain why $\hat{l} \sim l^*(\bar{v})$ where:

$$\bar{v} \equiv \frac{\pi_1}{\pi_1 + \pi_2} v(c_1) + \frac{\pi_2}{\pi_1 + \pi_2} v(c_2)$$

(B) Appeal to the independence axiom to establish that:

$$(\hat{l}, c_3; 1 - \pi_3, \pi_3) \sim (l^*(\bar{v}), c_3; 1 - \pi_3, \pi_3)$$

and

$$(l^*(\bar{v}), c_3; 1 - \pi_3, \pi_3) \sim (l^*(\bar{v}), l^*(c_3); 1 - \pi_3, \pi_3)$$

(C) Depict the two lotteries $(\hat{l}, c_3; 1 - \pi_3, \pi_3)$ and $(l^*(\bar{v}), \ l^*(v(c_3)); 1 - \pi_3, \pi_3)$ in tree diagrams.

(D) Confirm that the first is equivalent to the lottery $(c_1, c_2, c_3; \pi_1, \pi_2, \pi_3)$. Confirm that the second is equivalent to the reference lottery with success probability $\sum_{s=1}^{3} \pi_s v(c_s)$.

(E) Suppose the expected-utility rule is true for prospects with S outcomes. (We have seen that it is true for $S = 2$ and 3.) Show that the above argument can, with only slight modifications, be used to establish that the expected-utility rule must be true for prospects with $S + 1$ outcomes.

1.5 Risk Aversion

In Figure 1.1 the individual pictured was indifferent between a certainty income of \$250 and a prospect yielding equal chances of \$1,000 or nothing. Such a person is termed *risk averse.* More generally:

DEFINITION: A person is *risk averse* (displays *risk aversion*) if he strictly prefers a certainty consequence to any risky prospect whose mathematical expectation of consequences equals that certainty. If his preferences go the other way he is a *risk preferrer or loving* (displays *risk preference*); if he is indifferent between the certainty consequence and such a risky prospect he is *risk neutral* (displays *risk neutrality*).

The risky prospect described above, equal chances of \$1,000 or nothing, has a mathematical expectation of \$500 of income. Since our individual was indifferent between the prospect and a mere \$250 certain, for him \$500 certain is surely preferable to the risky prospect, which verifies that he is indeed risk averse.

The term "fair gamble" is used to describe an uncertain prospect whose mathematical expectation is zero. (A gamble with negative expectation is called "unfavorable"; one with positive expectation is called "favorable.") For example, odds of 5:1 on a roll of a fair die represent a fair gamble: since you lose (say) a dollar if the face you name does not come up, and win five dollars if it does come up, the expectation of gain is $(-1)\frac{5}{6} + 5\left(\frac{1}{6}\right) = 0$. Then a risk-averse person would refuse a fair gamble; a risk preferrer would accept a fair gamble; and a risk-neutral person would be indifferent.[6]

Figure 1.4 displays three possible utility functions: $v_1(c)$ would apply to a risk-averse individual, $v_2(c)$ to someone who is risk neutral, and $v_3(c)$ to a risk preferrer. Consider the fair prospect or gamble $G = (750, 250; \frac{1}{2}, \frac{1}{2})$ whose mathematical expectation is \$500. For the first or risk-averse individual the utility of \$500 certain, $v_1(500)$, is indicated by the height of point T along the $v_1(c)$ curve. The utility he attaches to the risky prospect, choosing the gamble G, is indicated by point L – whose height is the probability-weighted average of the heights of points J and K. This is, of course, the geometrical equivalent of the expected-utility rule, which tells us that $U_1(G) = \frac{1}{2}\, v_1(750) + \frac{1}{2}\, v_1(250)$. Evidently, whenever the utility function has the "concave" shape of $v_1(c)$, points associated with a certainty income (like T in the diagram) will be higher than points (like L) representing a fair gamble with the same expectation of income. By an analogous argument,

[6] However, as we shall see below, a risk-averse individual would accept a fair gamble if it offset *other* risks to which he was exposed. To purchase insurance, for example, is to accept an offsetting (risk-reducing) gamble.

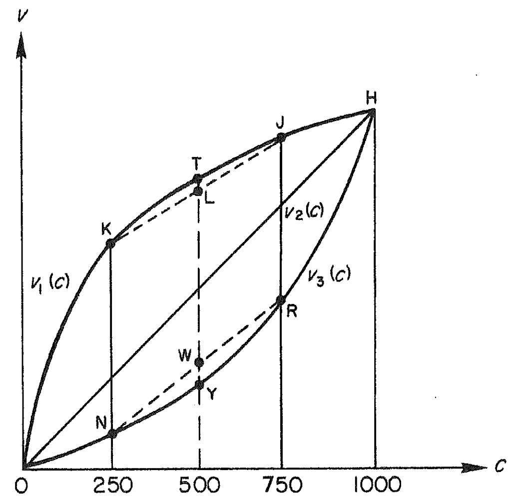

Figure 1.4. Attitudes toward risk.

for the risk-preferring individual, $v_3(500)$ at point Y will be less than at point W; such a person would choose the gamble *G* rather than receive its mathematical expectation of income, \$500, as a certainty. Finally, the $v_2(c)$ curve indicates that the risk-neutral person would be indifferent between the gamble *G* and the certainty of \$500.

We will often have occasion to make use of *Jensen's inequality*: If $\tilde{c}$ is a random variable (taking on at least two values with non-zero probability) and $v(c)$ is a twice-differentiable function:

If $v''(c) < 0$, then $E[v(c)] < v[E(c)]$
If $v''(c) = 0$, then $E[v(c)] = v[E(c)]$
If $v''(c) > 0$, then $E[v(c)] > v[E(c)]$

Evidently, these conditions correspond immediately to the risk-averse, risk-neutral, and risk-preferring cases of Figure 1.4.

It is useful to consider how attitude toward risk is reflected in the triangle diagram introduced in Section 1.4.1. Figure 1.5 below shows the set of lotteries over three income levels, \$0, \$500, \$1,000. Thus, π_1 is the probability of getting \$0 and π_3 is the probability of \$1,000. The origin corresponds to getting \$500 for sure.

Point L in Figure 1.5 is the lottery that gives either \$1,000 or nothing with equal probability. A risk-neutral individual is indifferent between this prospect and \$500 for sure. Hence, the solid line joining L to the origin is the indifference line through L for a risk-neutral person. As indifference

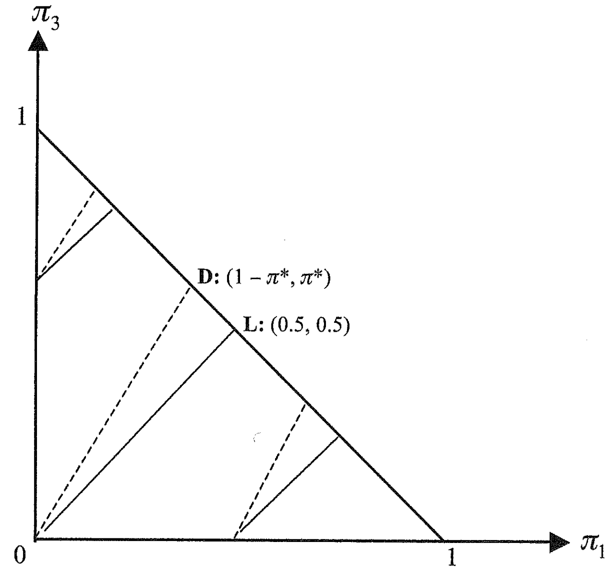

Figure 1.5. Risk-averse and risk-neutral indifference lines.

lines under the expected-utility rule are parallel straight lines, the solid lines inside the triangle are indifference lines under risk neutrality. That is, all lotteries on a solid line have the same expected value *and* the same expected utility for a risk-neutral person.

Any risk-averse individual strictly prefers \$500 for sure to the prospect L. Thus, since the direction of increasing preference is to the northwest, any risk-averse indifference line through the origin must intersect the hypotenuse of the triangle at a point D to the northwest of L ($\pi^* > 0.5$). Hence, indifference lines for a risk-averse person (the broken lines in the triangle) are steeper than the indifference lines for a risk-neutral person (the solid lines in the triangle). Similarly, the indifference lines for a risk-neutral person are steeper than indifference lines for a risk-preferring person.

We now consider what observation of the world tells us about the actual $v(c)$ curves entering into people's decisions. First of all, we have already postulated that *more income is preferred to less,* justified by the observation that only rarely do people throw away income. This implies a rising $v(c)$ function, with positive first derivative $v'(c)$, that is, positive marginal utility of income. The question of risk aversion versus risk preference concerns the second derivative $v''(c)$ – whether marginal utility of income falls or rises with income.

Risk aversion – "concave" curves like $v_1(c)$ displaying diminishing marginal utility – is considered to be the normal case, based upon the observation that individuals typically hold *diversified portfolios*. Suppose someone were merely risk neutral, so that for him $v''(c) = 0$. Then he would ignore the riskiness or variance of different investment options or assets (gambles), and take account only of the mathematical expectation of income associated with each. Such a person would plunge all his wealth into that single asset that, regardless of its riskiness, offered the highest mathematical expectation of income. But we scarcely ever see this behavior pattern, and more commonly observe individuals holding a variety of assets. Since the risks associated with different assets are generally partially offsetting, diversification reduces the chance of ending up with an extremely low level of income. This safety feature is achieved, however, only by accepting a lower overall mathematical expectation of income; some expected income has been sacrificed in order to reduce risk.[7]

What of the seemingly contrary evidence that "unfavorable" (negative mathematical expectation) gambles are cheerfully accepted by bettors at Las Vegas and elsewhere? Even more puzzling, why is it that the same person might behave quite conservatively (insure his house, diversify his asset holdings) in some circumstances, and in other circumstances accept fair or even unfavorable gambles? There have been attempts to construct utility functions $v(c)$ that would be consistent with avoiding gambles (insuring) over certain ranges of income *and* with seeking gambles over other ranges (Friedman and Savage, 1948; Markowitz, 1952). We will briefly discuss the Friedman-Savage version.

Consider the doubly inflected utility function in Figure 1.6. The $v(c)$ curve is concave, reflecting normal risk aversion, in the region OK and once again in the region LN. But it is convex, reflecting risk preference, in the middle region KL. With this sort of $v(c)$ function, risk-taking behavior will vary with wealth. For those whose endowments fall in the first concave segment, the tendency is to insure against relatively small risks but to accept fair (or even mildly adverse) long-shot big-payoff gambles, offering a chance of landing somewhere toward the upper end of the curve. It can be verified that this pattern will particularly apply for those with incomes toward the upper edge of the bottom segment – the less indigent poor, and perhaps the lower-middle class. The *very* poor, in contrast, would be much less inclined to gamble. Looking now toward the top of the scale, those with incomes near the lower edge of the upper concave segment – the rich but not super-rich,

[7] An individual characterized by *risk-preference* might also plunge all of his wealth into a single asset, but this need not be the asset with the highest mathematical expectation of income. He might choose an asset with greater riskiness over the asset with highest income yield (that is, he would sacrifice some expected income in order to *enlarge* his risk).

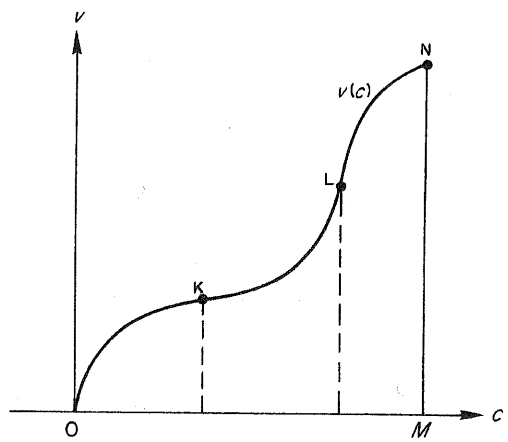

Figure 1.6. Gambling and insuring – doubly inflected utility function.

and perhaps the upper middle class – would seem to have a taste for risks likely to have a favorable payoff but offering a long-shot chance of a really large loss. (But the super-rich, like the super-poor, are very disinclined to gamble at all.) The central group, finally, would be happy to accept almost any fair or not-too-unfavorable gamble.

The doubly inflected utility function of Figure 1.6 does then explain why a person might gamble in some circumstances and insure in others, or accept some fair gambles while rejecting other ones. But it also implies other behavior that is quite inconsistent with common observation. It is hard to believe that people of middling incomes are always great gamblers. If the picture in Figure 1.6 were correct, the middle group in the convex KL segment would be so anxious to gamble as to seek out enormous riches-or-ruin bets. These middle ranges of income would then rapidly be depopulated, which is surely not what is observed. And that the really solid risk avoiders in our society are only the very poor and the super-rich is equally difficult to credit.

A more acceptable explanation, of why people simultaneously gamble and insure, is that most of us engage in gambling as a recreational rather than an income-determining activity. Put another way, gambling is normally more like a consumption good than an investment good. As it happens, it is quite possible operationally to distinguish recreational or pleasure-oriented from serious wealth-oriented gambling. The latter, if efficiently conducted, would take the form of once-and-for-all wagers at enormous stakes. Pleasure-oriented gambling, in contrast, being designed to yield enjoyment over some period of time, will be characterized by repetitive minuscule bets practically guaranteed *not* to change one's wealth status in any drastic way. What is observed at Las Vegas is very much more the repetitive small-stake than the riches-or-ruin huge-stake betting pattern.

Nevertheless, in exceptional situations, risk-preferring behavior does indeed surely occur. Consider the following. As bank cashier you have dipped into the till to the extent of $30,000. The bank examiners are arriving tomorrow, so you have time to replace the missing funds, but you have only $10,000 left on hand. Suppose you value the consumption benefit of spending the remaining $10,000 today far less than you value avoiding the shame and pain of exposure as an embezzler. Then you surely would be willing to risk the $10,000 on a fair gamble today – say, with a $\frac{1}{3}$ chance of winning $20,000. You would probably even take quite an adverse bet if necessary, so long as the possible payoff sufficed to cover the $20,000 of additional funds you need.

What is involved here is a "threshold" phenomenon, a critical level of income where a little bit more can make a big difference. Put another way, there is a range of *increasing* marginal utility – in the extreme, a single discrete step to a higher utility level. Threshold phenomena are quite common in nature. In many species, animals must take risks in accumulating resources or engaging in combat in order to achieve nutritional viability or win the privilege of mating. These phenomena have evident analogs for humans living in primitive societies. To what extent they may explain risk-taking behavior under modern conditions may be left an open question.[8]

This discussion may possibly suggest, contrary to a point made earlier, that it is after all true that utility must "really" be cardinal. A viability threshold, for example, might seem to be a cardinal feature of preference that would apply to riskless as well as to risky decision making. Nevertheless, our original point remains valid. For certainty choices, only ordinal comparisons of consequences are needed. For decisions under uncertainty we can derive, by the reference-lottery technique, a $v(c)$ function that may have convex or concave or mixed curvature, as the case may be. But the shape of this function for any individual is an inseparable blend of two elements: (i) the individual's valuations of the consequences, and (ii) his attitudes toward

[8] See Rubin and Paul (1979). These authors suggest that the propensity of young males to engage in highly risky activities – as evidenced, for example, by their high automobile accident rates – may be the result of natural selection for risk-taking. The evolutionary history of the human species may have instilled risk-preferring attitudes among individuals in age and sex groups liable to encounter viability or mating thresholds. (Note that the threshold argument is also consistent with the observation that risk-taking behavior will be observed predominantly *among the poor*.)

risk. We may therefore interpret a concave $v(c)$ function as reflecting *either* risk aversion (attitude toward risk) or diminishing marginal utility (attitude toward income); similarly, a convex $v(c)$ function can be said to reflect *either* risk preference or increasing marginal utility. Both terminologies are somewhat misleading, since what the curvature of $v(c)$ really represents is the *interaction* of the two factors working together.

Finally, another category of seeming risk-taking behavior may be explainable in terms of *state-dependent utility functions.* An example: Suppose it is very important to me, as the psychological equivalent of having a large sum of money, that the home team wins the big game. Then I might plausibly bet *against* the home team, at fair or even adverse odds! (How this works out in detail will be left for the chapter following.)

Exercises and Excursions 1.5

1 Risk Aversion, Risk Preference, Risk Neutrality

(A) Identify each of the following "cardinal" utility functions with risk-averse, risk-preferring, or risk-neutral behavior:

(i) $v = \ln c$ (ii) $v = ac - bc^2$ (a, b positive constants)
(iii) $v = c^2$ (iv) $v = c^{\frac{1}{2}}$
(v) $v = 100 + 6c$ (vi) $v = 1 - e^{-c}$

(B) The quadratic form (ii) above has an unsatisfactory feature for $c > a/2b$. Explain.

2 Diversification

Three individuals have respective utility functions $v_1 = c$ (risk neutral), $v_2 = c^{0.5}$ (risk averse), and $v_3 = c^2$ (risk preferrer). They each have the option of investing in *any one* of the three following prospects or gambles, with mathematical expectations of income as shown:

G1 = (480, 480; 0.5, 0.5) E[G1] = 480
G2 = (850, 200; 0.5, 0.5) E[G2] = 525
G3 = (1,000, 0; 0.5, 0.5) E[G3] = 500

Notice that, comparing the first two gambles, higher risk is associated with greater mathematical expectation of income. The third gamble has highest risk of all, but intermediate mathematical expectation.

(A) Show that risk-neutral individual 1 will prefer gamble G2 with the highest expectation, while risk-averse individual 2 will prefer gamble

G1 with the lowest risk. Show that the risk-preferring individual 3 is willing to sacrifice some expectation to *increase* his risk, by choosing G3.

(B) If the individuals could "diversify" by choosing any desired mixture of these gambles, which of them would diversify? (Assume that the payoffs of gambles G2 and G3 are perfectly correlated.)

3 *Doubly Inflected Utility Function*

In the doubly inflected $v(c)$ curve shown in Figure 1.6, suppose that the borders of the segments (inflection points) occur at $c = 250$ and at $c = 750$.

(A) Illustrate geometrically that an individual with initial income of \$240 would be likely to accept a (fair) gamble offering a one-sixth chance of a \$600 gain and a five-sixth chance of a \$120 loss. Show that someone with initial income of \$120 would be much less likely to accept the same gamble.

(B) Show that someone with initial endowed income of \$760 would be likely to accept a fair gamble which is the reverse of the above: a five-sixth chance of a \$120 gain and a one-sixth chance of a \$600 loss. What about a person with initial wealth of \$880?

(C) Show that someone with endowed wealth of exactly \$500 would surely accept *any* fair gamble with 50: 50 odds – at least up to a scale of \$250 gain and \$250 loss. He might even accept much larger fair gambles of this type; indicate geometrically the limits of what he would accept.

4 *Linear Risk Tolerance*

Risk aversion is characterized by the condition $v''(c) < 0$. For some purposes, as we shall see below, the ratio $-v''/v'$ is a useful measure of risk aversion. The reciprocal of this ratio, $-v'/v''$, is known as the *risk tolerance.* An interesting class of $v(c)$ functions is defined by the condition of *linear* risk tolerance: $-v'/v'' = \alpha + \beta c$.

(A) Show that, for arbitrary constants M, N with $N > 0$:
  (i) $\beta = 0$ implies $v = M - Ne^{-c/\alpha}$
  (ii) $\alpha = 0, \beta \neq 1$ implies $v = M + Nc^{1-\gamma/(1-\gamma))}$ where $\gamma = 1/\beta$
  (iii) $\alpha = 0, \beta = 1$ implies $v = M + N \ln c$
  (iv) $\alpha > 0, \beta = -1$ implies $v = M - N(\alpha - c)^2$

(B) Some of the above functions are valid only in restricted ranges of c. Indicate the restrictions, if any, that apply in each case. Also explain why N must be positive if v is to be a well-behaved utility function.

5 The Bank Examiner Is Coming

You have stolen \$30,000 from the bank but have the opportunity to replace it by winning a fair gamble. You have at your disposal just \$10,000. Your utility function is such that $v(c) = -B$, where B is a very big number, when $c < 0$ (i.e., should you not replace *all* the missing funds), and otherwise $v(c) = c^{\frac{1}{2}}$. Assuming fair gambles are available at any terms you desire, solve *geometrically* for your optimal fair gamble. Will you surely stake all your \$10,000? Will you look only for a \$20,000 payoff, or would you prefer a bet with a smaller chance of a bigger payoff?

6 Utility Functions with Multiple Goods

The argument in the text above, developing a cardinal utility function $v(c)$ for use with the expected-utility rule, ran in terms of a single desired good or commodity c. Extend the argument to cardinal utility functions of two goods, in the form $v(a, b)$. Show that, starting with an *ordinal* preference function defined over combinations of a and b (that is, starting with an ordinary indifference map on a, b axes), the reference-lottery technique can be used to generate a cardinal scaling that amounts to giving a numerical utility value to each indifference curve.

7 Risk Aversion with Multiple Goods

An individual has a utility function $v(a, b) = a^{\frac{1}{2}} b^{\frac{1}{4}}$. He has income I available for spending on a and b, and faces fixed prices $P_a = P_b = 1$.

(A) Show that he would strictly prefer the certain income of 50 to an equal chance of his income rising or falling by 49 before he makes his consumption choices.

(B) Obtain an expression for the individual's "indirect" utility function. (That is, the maximized level of v given income I and prices P_a and P_b.) Hence show that this individual exhibits aversion to income risks.

(C) Suppose $I = 50$ and $P_b = 16$. Would the individual prefer to face a certain $P_a = 64$ or a stochastically varying P_a that might equal 1 or 81 with equal chances? Does your answer cast doubt upon whether the individual is really risk averse? Explain.

8 Jensen's Inequality (I)

(A) If the utility function $v(c)$ is twice continuously differentiable with $v''(c) \leq 0$, show that for any random variable $\tilde{c}$:

$$\mathrm{E}[v(\tilde{c})] \leq v(\mathrm{E}[\tilde{c}])$$

(B) If $v''(c) < 0$ and $\Pr[\tilde{c} \neq \mathrm{E}[\tilde{c}]] > 0$ show that:

$$\mathrm{E}[v(\tilde{c})] < v(\mathrm{E}[\tilde{c}])$$

9 *Jensen's Inequality (II)*
Suppose $v(c)$ is a concave function (not necessarily differentiable), that is, for any c_1, c_2:

$$v((1-\lambda)c_1 + \lambda c_2) \geq (1-\lambda)v(c_1) + \lambda v(c_2),\ 0 \leq \lambda \leq 1$$

(A) Prove by induction that, for any $c_1, \ldots, c_n$:

$$v\left(\sum_{i=1}^{n} \mu_i c_i\right) \geq \sum_{i=1}^{n} \mu_i v(c_i), \quad \text{for } \mu_i \geq 0,\ \sum_{i=1}^{n} \mu_i = 1$$

(B) Hence derive Jensen's inequality once again.

1.6 Utility Paradoxes and Rationality

A very considerable literature arguing against expected utility as a good descriptive theory has appeared in the last 30 years. Its main thrust has been that actual decision makers do not behave rationally in the face of uncertainty, or at any rate do not consistently follow the expected-utility rule.[9] To some extent, these complaints have been supported by experimental evidence.[10] We provide four illustrations of which the first two are the following.

1.6.1 Probability Matching

You are paid $1 each time you guess correctly whether a red or a white light will flash. The lights flash randomly, but the red is set to turn on twice as often as the white.

It has been found that subjects tend to guess red about two-thirds of the time and white one-third. Yet, obviously, it would be more profitable always to guess red.

[9] Machina (1987) provides a very helpful and clear survey to this literature. A more recent survey is Starmer (2000).

[10] See, for example, Slovic and Lichtenstein (1983), Tversky and Kahneman (1981), and Schoemaker (1982).

1.6.2 Framing the Question

Imagine that you have been given \$200 and are asked to choose between (i) \$50 additional or (ii) a 25% chance of winning \$200 additional (or else, gaining nothing). Alternatively, imagine that you have been given \$400, but you must now choose between (i) giving up \$150 or (ii) a 75% chance of losing \$200 (or else, losing nothing).

Most experimental subjects choose option (i) in the first version of the question, but option (ii) in the second. Yet, obviously, option (i) generates the same income prospect whichever way the question is framed, and similarly for option (ii).

This literature on non-expected utility theory claims that the discrepancies revealed by these results refute the economist's standard assumption of rationality, or at least the expected-utility rule as a specific implication of that assumption. We do not accept this interpretation. A much more parsimonious explanation, in our opinion, is that this evidence merely illustrates certain limitations of the human mind as a computer. It is possible to fool the brain by the way a question is posed, just as optical illusions may be arranged to fool the eye. Discovering and classifying such mental illusions are fruitful activities for psychologists, but these paradoxes have less significance for economics.

We would not go so far as to insist that rationality failures have *no* economic implications. If these shortcomings do indeed represent ways in which people could systematically be fooled, economists would predict that tricksters, confidence men, and assorted rogues would enter the "industry" offering such gambles to naive subjects. For example:

PROBABILITY MATCHING: The trickster could challenge the subject along the following line: "I have a secret method of guessing which light will flash. (His secret method, of course, is always to bet on red.) I will write my guess down on paper each time, and you will write yours down. At the end we will total up our successes. For each time I am right and you are wrong, you will pay me \$1; in the reverse case, I will pay you \$1.50."[11] If the subject really believes that his is the right method, he should surely accept so generous an offer.

And similarly, clever tricksters could win sure-thing income from the inconsistent answers offered by naive individuals in our other illustration. The confidence-man profession does obviously exist, and is unlikely (given the

[11] The maximum or breakeven payment that the trickster could offer is \$2 exactly. Clearly, there will be no payment either way in the two-thirds of the cases where the naive subject bets on red. And when he bets on white, he will be wrong twice as often as he is right.

limitations of the human mind) ever to disappear.[12] But the more important the decision, the more it is worth people's while to learn how not to be fooled.

It will be of interest to analyze some of the parallels and differences among these various rationality failures, and in particular to attempt to identify more precisely the source of the slippage in each case.

If the subjects in PROBABILITY MATCHING did mentally compare the matching rule with "Always bet on red" and chose the former, they committed a straightforward logical error.[13] We know that people do often commit such errors, even in contexts where no uncertainty is involved. Consider the following example from a psychological experiment (adapted from Cosmides, 1989):

> You are faced with a card-sorting task, in which each card has a number on one side and a letter on the other. There is only one rule: "Every card marked with an 'X' on one side should have a '1' on the other." Indicate whether you need to inspect the reverse side of the following cards to detect violation of the rule: (a) a card showing an 'X'; (b) a card showing a 'Y'; (c) a card showing a '1'; (d) a card showing a '2'.

In a large preponderance of cases, while the subjects correctly realized the need to inspect the reverse of card (a), they failed to notice that they should do the same for card (d).

What is instructive for our purposes, however, is that the experimenter went on to investigate a logically identical choice, presented to the subjects more or less as follows:

> You are the bouncer in a Boston bar, assigned to enforce the following rule: "Anyone who consumes alcohol on the premises must be at least twenty years old." To detect violation of the rule, indicate whether you need more information about any of the following individuals: (a) someone drinking whisky; (b) someone drinking soda; (c) an individual aged twenty-five; (d) an individual aged sixteen.

Here almost everyone perceived the need for more information about individual (d) as well as individual (a). Evidently, humans have trouble with purely abstract problems, but do a lot better when the logically equivalent choices are offered in a realistic context – particularly where possible cheating or violations of social norms may be involved. Returning to

[12] An analogous example is the racetrack tout who offers to predict the winning horse for $20, telling you that he will refund your money unless his prediction is correct. His intention, of course, is to tout customers onto all the horses in the race.

[13] Another possibility is that the correct rule never came to mind at all – in effect, the subjects did not think very hard about what was going on. This would not be too surprising if the stakes were trivial in magnitude.

PROBABILITY MATCHING, in our opinion few individuals would be more than momentarily fooled by the trickster described above if some serious issue or some substantial amount of money were at stake.

The second example, FRAMING THE QUESTION, is rather like an optical illusion involving perspective, a nearby small object being made to seem larger than a far-off large object. In the first choice offered the subjects, the risk – the chance of losing $50 – is placed in the foreground, so to speak. From this viewpoint, the 25% chance of gaining an extra $200 does not seem enough recompense. In the second version what is placed in the foreground is the unpleasant option of a $150 loss. Here the risk of losing an additional $50 fades into comparative insignificance, as compared with the 25% hope of recouping the $150 and suffering no loss at all. Notice that these experimental subjects proved to be highly risk averse; they were fooled by a shift in the setting, the same risk being highlighted in the one choice and left in the shadows in the other case.

1.6.3 Allais Paradox

We will provide a more extended discussion of a third example, the ALLAIS PARADOX, which illustrates the powerful effect of just how the choices are framed (Allais (1953)):

You are offered the choice between prospects A and B:

A: with certainty, receive $1,000,000
B: with probability 0.10, receive $5,000,000
with probability 0.89, receive $1,000,000
with probability 0.01, receive zero.

Alternatively, you are offered the choice between C and D:

C: with probability 0.11, receive $1,000,000
with probability 0.89, receive zero
D: with probability 0.10, receive $5,000,000
with probability 0.90, receive zero.

It has been found that most people prefer A to B, but D to C. But it is easy to show that choosing A over B but D over C is inconsistent with the expected-utility rule. According to that theorem:

$$\text{If } A \succ B, \text{ then } v(\$1{,}000{,}000) > 0.10\ v(\$5{,}000{,}000) + 0.89\ v(\$1{,}000{,}000) + 0.01\ v(\$0)$$

Then, by elementary algebra:

$$0.11\,v(\$1{,}000{,}000) + 0.89\,v(\$0) > 0.10\,v(\$5{,}000{,}000) + 0.90\,v(\$0)$$

But the latter inequality is equivalent, according to the expected-utility rule, to $C \succ D$.

The explanation, in perceptual terms, appears to be that the A versus B framing makes the 0.01 chance of receiving zero stand out as a very adverse feature in making option B undesired – but exactly the same chance fades into comparative insignificance, psychologically speaking, as an adverse feature of D in comparison with C.

The question is, does the observed failure of subjects to follow the dictates of the expected-utility rule represent only a logical lapse, akin to an optical illusion? Or is it perhaps that the rule is an incorrect, or at least an excessively narrow, specification of rational behavior? The latter was the position taken by Allais.

Individuals who choose in accordance with the Allais Paradox example are violating the independence axiom: that any complex lottery can be reduced to its elements. The following thought experiment reveals how the choices described by Allais violate this axiom. Consider the prospect X and sure thing Y:

X: with probability 10/11, receive $5,000,000
 with probability 1/11, receive zero.
Y: receive $1,000,000 with probability 1.

Then, the prospects A, B, C, and D may be written as the following complex gambles:

Thus, A is decomposed into a gamble which yields Y with probability 0.11 and $1,000,000 with probability 0.89, B is the complex gamble which leads to X with probability 0.11 and $1,000,000 with probability 0.89, and so on.

According to the independence axiom, the choice between A and B, and between C and D should be determined by choice between Y and X. Thus:

If $Y \succ X$ then $A \succ B$ and $C \succ D$
If $Y \prec X$ then $A \prec B$ and $C \prec D$.

However, as noted above, for most subjects $A \succ B$ and $C \prec D$. Thus, the attractiveness of Y relative to X depends on the lower branches of the complex gambles A, B, C, and D in Figure 1.7, i.e., on what might happen in the

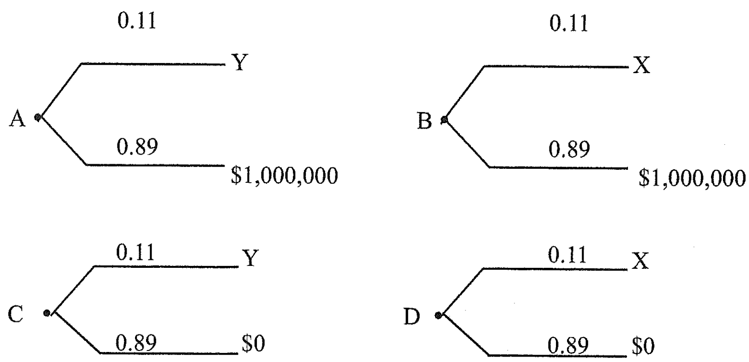

Figure 1.7. Allais Paradox lotteries.

event one does not have to make a choice between Y and X. This complementarity between the lower and upper branches of the complex gambles depicted in Figure 1.7 represents a violation of the independence axiom. As Machina (1987) points out, in comparing A and B, subjects judge the relative attractiveness of Y and X in the event of an opportunity loss of $1,000,000 (i.e., in the event that a 0.89 chance of winning $1,000,000 does not occur); this opportunity loss makes them very risk averse and they prefer $Y \succ X$. In comparing C and D, on the other hand, there is no opportunity loss when faced with a choice between Y and X; subjects are less risk averse and prefer $X \succ Y$. This in turn implies that the indifference curves for Allais Paradox preferences, when depicted as solid straight lines in Figure 1.8, cannot be parallel straight lines. The indifference curves, instead of being parallel, fan out as shown in Figure 1.8.

The triangle diagram represents prospects with outcomes $0, $1,000,000, and $5,000,000. The origin represents getting $1,000,000 with certainty. The right-hand corner ($\pi_1 = 1, \pi_3 = 0$) represents $0 for sure and the third corner of the triangle, ($\pi_1 = 0, \pi_3 = 1$), is $5,000,000 for sure. Thus, the gamble X is on the hypotenuse of the triangle, as shown, and the sure thing Y is at the origin.

Complex gambles A and C are obtained by combining (in proportion 0.11 to 0.89) Y with the origin and with the right-hand corner, respectively. Similarly, B and D are obtained by combining X with the origin and with the right-hand corner, respectively. The solid lines in the triangle are indifference lines consistent with Allais Paradox preferences. The prospect C lies below the indifference line through D and A lies above the indifference line through B. By simple geometry, the line segments AB and CD are parallel. Thus, indifference lines must fan out as shown. At the lower right-hand side

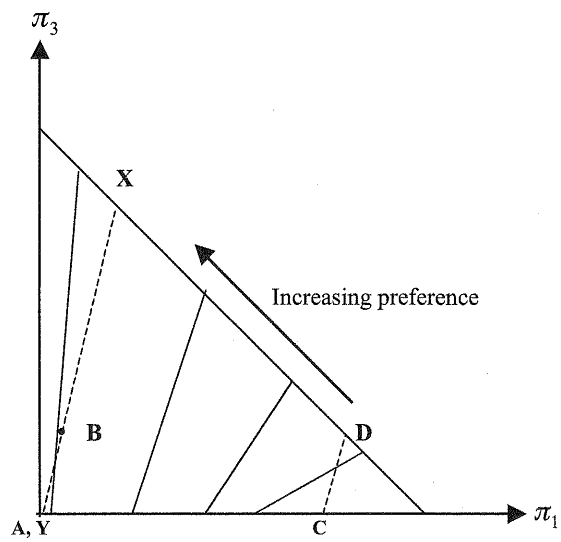

Figure 1.8. Allais Paradox indifference lines.

of the triangle, indifference lines are less steep (less risk averse), and to the upper left-hand side they are steeper (more risk averse).

In order to explain the Allais Paradox and other experimental evidence, various generalizations of the expected-utility rule have been proposed. One generalization, due to Machina (1982), is to allow the utility function for consequences to depend on the prospect being evaluated. Thus, the "expected utility" of x is

$$U(x) = \mathrm{E}[v(c_{xs}, x)] \tag{1.6.1}$$

The preceding discussion implies that $v(c_{xs}, x)$ is less risk averse for x in the lower right-hand side of the triangle and more risk averse for x in the upper left-hand side of the triangle.

However, even this generalization of the expected-utility rule cannot explain the next example.

1.6.4 Ellsberg Paradox

Urn I has 50 red balls and 50 black balls. Urn II also has 100 red and black balls, but in unknown proportions. You are invited to bet on the color of a ball that will be drawn randomly from one of the two urns. You will win $100 in the event of a correct choice.

(A) Of the two red bets R_I or R_{II} (a bet on red if the drawing is made from the first, or alternatively from the second urn), which do you prefer?
(B) Same question, for the two black bets B_I and B_{II}.

It has been found that most subjects prefer R_I over R_{II}, and also prefer B_I over B_{II}. To say that you prefer R_I over R_{II} is to say that you believe that the probability of a red ball from urn II is less than 0.5. But this implies that the probability of a black ball from urn II is more than 0.5 and hence subjects should prefer B_{II} over B_I. Thus, preferences reported by most subjects are inconsistent with the idea that subjects' beliefs about uncertainty can be expressed as probabilities. In particular, the expected-utility rule or even a generalization such as the formula (1.6.1) cannot be used to express these preferences (see Ellsberg 1961).

The ELLSBERG PARADOX plays on the subjects' aversion to ambiguity and vagueness. Recalling the discussion of "hard" versus "soft" probability estimates earlier in the chapter, the subjects have a preference for acting on the basis of a hard probability (the urn known to have 50 black and 50 red balls) than acting on the basis of a soft probability (the urn with an unknown mixture). But if only an immediate *terminal* action is called for, as postulated here, it makes no difference whether the probability is hard or soft. In the absence of any basis for one color being more likely than the other, the subjective probability of success has to be the same for the second as for the first urn – whether betting on black or on red. The subjects seem to associate *higher confidence* (which indeed holds with regard to the probability of success using the first urn) with *lesser risk*. A generalization of the expected-utility rule that is consistent with the Ellsberg Paradox has been proposed by Schmeidler (1989).[14] Soft and hard probabilities are processed differently in this generalization so as to allow for aversion to ambiguity.

We do not want to be dismissive of what is, on a number of grounds, an intellectually significant literature. But we do note that most of the evidence of violations to the expected-utility rule has been experimental evidence gathered in an economics laboratory. Subjects are asked to make certain

[14] There is also an alternative explanation, entirely consistent with expected-utility behavior. In an actual experiment the first urn would presumably be transparent, to allow everyone to see that half the balls are red and half black. But, of course, the second urn could not be transparent, which makes trickery more possible. A subject attaching even a small likelihood to being cheated (by the experimenter shifting the proportions in the second urn after the bet is down) would definitely and quite rationally prefer drawing from the first urn.

choices. In a real setting where individuals make choices repeatedly, there is a greater opportunity and incentive to learn and reconsider the kinds of choices made by subjects in the four examples presented here.

As an empirical matter, such important phenomena as advertising and political persuasion depend very importantly upon clever use of fallacious analogy, irrelevant associations, and other confidence-man tricks. But the analysis of error is only a footnote to the analysis of valid inference. It is only because people have a well-justified confidence in reason that deception, whether artful or unintended, can sometimes occur. Especially when it comes to subtle matters and small differences, it is easy for people to fool themselves, or to be fooled. But less so when the issues are really important, for the economically sound reason that correct analysis is more profitable than error.

Exercises and Excursions 1.6

1 *Framing the Question*

Could a confidence-man or trickster exploit individuals whose choices are as described in the *framing the question* example above?

2 *A Second Ellsberg Paradox*

An urn contains 30 red balls and 60 other balls, some yellow and some black. One ball is to be drawn at random from the urn.

(A) You are offered the opportunity to choose either red or black. If you pick the color of the ball drawn, you win $100. Which color do you choose?

(B) Alternatively, suppose you are offered once again the opportunity to choose either red or black. However, now you win $100 as long as the ball drawn is *not* the color picked. Which color do you choose?

(C) Show that only two of the four possible combinations of choices (for questions A and B, respectively) – red-red, red-black, black-red, and black-black – are consistent with the independence axiom.

(D) If your choices were inconsistent with the axiom, do you wish to change either of them?

3 *The Allais Paradox*

(A) Does Allais Paradox violate the independence axiom? If so, how?

(B) As a confidence-man, how would you exploit an individual whose choices were consistent with Allais Paradox?

4 Risk Aversion – Price or Quantity?
This exercise illustrates a different kind of utility "paradox." Suppose an individual with given wealth W can purchase commodities x and y. Let his utility function be:

$$v(x, y) = x + \alpha \ln y$$

Note that, in terms of our definitions above, for variations in x alone the individual is risk neutral ($\partial^2 v/\partial x^2 = 0$), while for variations in y alone he is risk averse ($\partial^2 v/\partial y^2 < 0$).

(A) Let the price of x be fixed at unity, and let p be the price of y. Show that his "indirect" utility, that is, elementary utility as a function of p, is given by:

$$\hat{v}(p) = \underset{x,y}{\text{Max}} \left\{ v(x, y) | x + py = W \right\} = \underset{y}{\text{Max}}(W + \alpha \ln y - py)$$

(B) Letting y^* denote his optimal consumption of good y, show that:

$$y^*(p) = \alpha/p$$
$$\hat{v}(p) = W - \alpha + \alpha \ln \alpha - \alpha \ln p$$

(C) Show that $\hat{v}(p)$ is a *convex* function of p, that is, $d^2\hat{v}/dp^2 > 0$.

(D) Explain the paradox that, while the $v(x, y)$ function displays risk aversion with respect to quantities of y, the $\hat{v}(p)$ function seems to display risk preference with respect to the price of y.

SUGGESTIONS FOR FURTHER READING: The expected-utility rule for objective probabilities was derived from a set of axioms on individual behavior by von Neumann and Morgenstern (1944). For a proof of the expected-utility rule, see Kreps (1988). Savage (1954) increased the domain of applicability of the expected-utility rule by allowing probabilities to be subjective. In Savage's world, both the utility over consequences *and* probabilities over uncertain events are personalized and may differ among individuals; in this setting Savage derives the expected-utility rule from a set of axioms. Kreps (1988) provides a relatively reader-friendly development of Savage's theory of expected utility with subjective probability. The literature on non-expected utility theory has burgeoned over the last twenty years. In addition to surveys by Machina (1987) and Starmer (2000) mentioned in the chapter, see the books by Schmidt (1998) and Gilboa (2009) for more on this subject.

References

Allais, Maurice, "Le Comportement de l'homme rationnel devant le risque," *Econometrica*, 21 (1953), 505–546.

Cosmides, Leda, "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task," *Cognition*, 31 (1989), 187–276.

Ellsberg, Daniel, "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75 (1961), 643–669.

Friedman, Milton and Savage, Leonard J., "The Utility Analysis of Choices Involving Risks," *Journal of Political Economy*, 56 (1948), 279–304.

Gilboa, Itzhak, *Theory of Decision under Uncertainty*, Econometric Society Monographs, Cambridge: Cambridge University Press, 2009.

Knight, Frank H., *Risk, Uncertainty and Profit*, New York: Houghton Mifflin, 1921.

Kreps, David, *Notes on the Theory of Choice*, Boulder: Westview Press, 1988.

Machina, Mark J., "'Expected Utility' Analysis without the Independence Axiom," *Econometrica*, 50 (1982), 277–323.

———, "Choice under Uncertainty: Problems Solved and Unsolved," *Journal of Economic Perspectives*, 1 (1987), 121–154.

Markowitz, Harry, "The Utility of Wealth," *Journal of Political Economy*, 60 (1952), 151–158.

Neumann, John von and Morgenstern, Oskar, *Theory of Games and Economic Behavior*, Princeton: Princeton University Press, 1944.

Rubin, Paul H., and Paul, Chris W., "An Evolutionary Model of the Taste for Risk," *Economic Inquiry*, 17 (1979), 585–596.

Savage, Leonard J., *The Foundations of Statistics*, New York: Wiley, 1954.

Schlaifer, Robert, *Probability and Statistics for Business Decisions*, New York: McGraw-Hill, 1959.

Schmeidler, David, "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57 (1989), 571–587.

Schmidt, Ulrich, *Axiomatic Utility Theory under Risk: Non-Archimedean Representations and Application to Insurance Economics*, Berlin: Springer-Verlag, 1998.

Schoemaker, Paul J. H., "The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations," *Journal of Economic Literature*, 20 (1982), 529–563.

Slovic, Paul and Lichtenstein, Sarah, "Preference Reversals: A Broader Perspective," *American Economic Review*, 83 (1983), 596–605.

Starmer, Chris, "Developments in Non-expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature*, 38 (2000), 332–382.

Tversky, Amos and Kahneman, Daniel, "The Framing of Decisions and the Psychology of Choice," *Science*, 211 (1981), 453–458.

2

Risk Bearing

The Optimum of the Individual

In this chapter, we address a basic problem in decision making under uncertainty: how should an individual select consumption across different states of nature so as to maximize his expected utility? The trick is to view consumption in one state of nature as a separate good from consumption in another state of nature. We then apply an indispensable technique in microeconomics – indifference curve analysis – to the problem to obtain the Fundamental Theorem of Risk Bearing. This theorem is directly applicable only when all state claims are available, i.e., for each state there is a good that pays if and only if that state obtains. However, trading in state claims is usually not a feasible option. Therefore, we generalize to a model with assets, where each asset is viewed as a vector of payoffs, one for each state of nature. This leads to the Risk-Bearing Theorem for Assets Markets. Next, we investigate risky choices made by an individual who cares only about the mean and standard deviation of his consumption. We end the chapter with a model of state-dependent utility.

The individual's best action under uncertainty – the "risk-bearing optimum" – involves choosing among prospects $x \equiv (c; \pi) \equiv (c_1, \ldots, c_S; \pi_1, \ldots, \pi_S)$ where the c_s are the state-distributed consequences and π_s are the state probabilities. In the realm of the economics of *uncertainty* proper, before turning to the economics of *information*, the individual's probability beliefs π remain constant and so $c_1, \ldots, c_S$ are the only decision variables. In general, each c_S represents the multi-good basket that the individual is entitled to consume if state *s* occurs. For simplicity, however, we will often think in terms of a single generalized consumption good ("corn"). Then c_s would simply be the individual's state-*s* entitlement to corn if state *s* occurs, and the risk-bearing problem is how to choose among alternative vectors $(c_1, \ldots, c_s)$ of "corn incomes" distributed over states of the world. Unless otherwise indicated, when the symbol c_s is described as representing

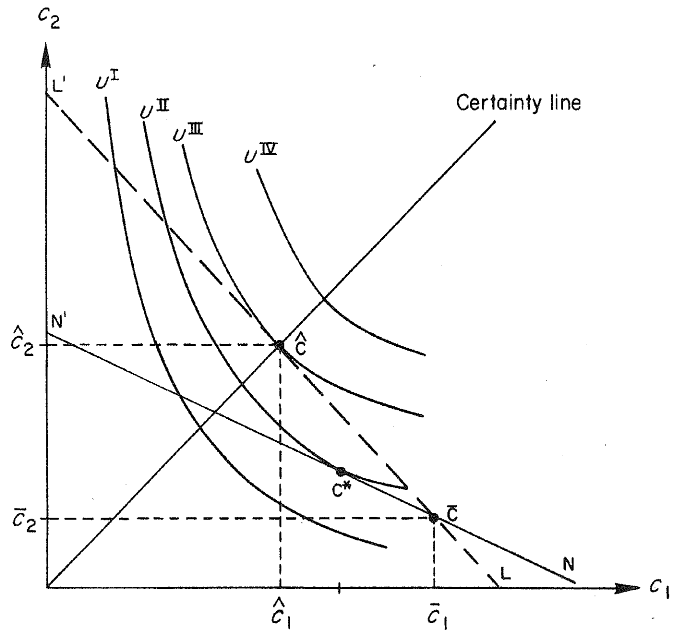

Figure 2.1. Individual optimum.

"income" the implication is that we are using the simplified model of a single consumption good.[1]

2.1 The Risk-Bearing Optimum: Basic Analysis

Suppose there are only two states of the world $s = 1, 2$. The two states might represent war versus peace, or prosperity versus depression. In the state-claim space of Figure 2.1 the axes indicate amounts of the contingent income claims c_1 and c_2.

To represent preferences in this space, we can start with Equation (1.4.1), the expected-utility rule. In a simplified two-state world, this reduces to:

$$U \equiv \pi_1 v(c_1) + \pi_2 v(c_2), \text{ where } \pi_1 + \pi_2 = 1 \qquad (2.1.1)$$

For a given level of U, Equation (2.1.1) describes an entire set of c_1, c_2 combinations that are equally preferred, so this is the equation of an indifference

[1] With multiple consumption goods, only if the price ratios among them were *independent of state* could there be an unambiguous interpretation of "income." Consider an individual whose multi-commodity physical endowment is distributed over two states s^o and s^*. When price ratios vary over states it might be that, valued in terms of good g as numeraire, his endowed "income" is higher in state s^o – while in terms of good h as numeraire instead, the value of endowed "income" in state s^* is higher.

curve. As U varies, the whole family of indifference curves implied by the individual's utility function $v(c)$ and probability beliefs π_1, π_2 is traced out – as indicated by the various curves $U^{\text{I}}, U^{\text{II}}, \ldots$, shown in the diagram.

It is elementary to verify that the absolute indifference-curve slope $M(c_1, c_2)$ in Figure 2.1, the Marginal Rate of Substitution in Consumption, is related to the marginal utilities $v'(c_1)$ and $v'(c_2)$ via:[2]

$$M(c_1, c_2) \equiv -\left.\frac{dc_2}{dc_1}\right|_{U=\text{constant}} \equiv \frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} \tag{2.1.2}$$

The 45° "certainty line" in the diagram connects all the points such that $c_1 = c_2$. Note that any indifference curve, as it crosses the certainty line, has absolute slope equal simply to π_1/π_2 – the ratio of the state probabilities.

Intuitively, risk aversion in state-claim space corresponds to convex ("bowed toward the origin") indifference curves as shown in Figure 2.1. Risk aversion, we know, leads to diversification. Non-convex indifference curves, when juxtaposed against the individual's opportunity set, would lead to a corner optimum – to choice of an all c_1 or an all c_2 state-claim holding, a non-diversified portfolio. More specifically, a risk-averse utility function $v(c)$, one with positive first derivative $v'(c)$ and negative second derivative $v''(c)$, does indeed imply that indifference curves in state-claim space will be bowed toward the origin. That is, $v'(c) > 0$, $v''(c) < 0$ imply that along any indifference curve, the absolute indifference curve slope diminishes moving to the right: $dM(c_1, c_2)/dc_1$ will be *negative*.[3]

It will also be of interest to translate into state-claim space the proposition that a risk-averse individual would never accept a fair gamble in "corn" income (would always prefer a sure consequence to any probabilistic mixture of consequences having the same mathematical expectation). In Figure 2.1, the dashed line LL′ through the point $\bar{C} \equiv (\bar{c}_1, \bar{c}_2)$ shows all the c_1, c_2

[2] Along an iso-utility curve, $0 = dU \equiv \pi_1 v'(c_1)\, dc_1 + \pi_2 v'(c_2)\, dc_2$. Then, $-dc_2/dc_1 = [\pi_1 v'(c_1)]/[\pi_2 v'(c_2)]$.

[3] The sign of $dM(c_1, c_2)/dc_1$ will be the same as that of $d \ln M(c_1, c_2)/dc_1$ where:

$$\begin{aligned}\frac{d}{dc_1} \ln M(c_1, c_2) &= \frac{d}{dc_1}[\ln \pi_1 + \ln v'(c_1) - \ln \pi_2 - \ln v'(c_2)] \\ &= \frac{v''(c_1)}{v'(c_1)} - \frac{v''(c_2)}{v'(c_2)} \frac{dc_2}{dc_1} \\ &= \frac{v''(c_1)}{v'(c_1)} + \frac{v''(c_2)}{v'(c_2)} \frac{\pi_1}{\pi_2} \frac{v'(c_1)}{v'(c_2)}\end{aligned}$$

Since the first derivatives are both positive, $v''(c_1)$ and $v''(c_2)$ both negative imply a diminishing Marginal Rate of Substitution.

combinations having the same mathematical expectation of income $E[c] = \hat{c}$ as the combination $(\bar{c}_1, \bar{c}_2)$. The equation for LL′ is:

$$\pi_1 c_1 + \pi_2 c_2 = \pi_1 \bar{c}_1 + \pi_2 \bar{c}_2 = \hat{c} \tag{2.1.3}$$

Along LL′ the most preferred point must be where the line is just tangent to an indifference curve of expected utility. The slope of LL′ is $dc_2/dc_1 = -\pi_1/\pi_2$, which (we know from 2.1.2) is the same as the slope along any indifference curve where it crosses the 45° line. Hence the tangency must be on the 45° line, to wit, at point $\hat{C}$ where $c_1 = c_2 = \hat{c}$. Thus, the certainty of having income $\hat{c}$ is preferred to any other c_1, c_2 combination whose mathematical expectation is $\hat{c}$.

2.1.1 Contingent-Claims Markets

As discussed in Chapter 1, Section 1.1, we are particularly interested in the risk-involved actions that economic agents can take through *market* dealings. Suppose the individual is a price taker in a market where contingent income claims c_1 and c_2 – each of which offers a unit of "corn income" if and only if the corresponding state obtains – can be exchanged in accordance with the price ratio P_1/P_2. This is indicated in Figure 2.1 by the budget line NN′ through the point $\bar{C} \equiv (\bar{c}_1, \bar{c}_2)$, now interpreted as the individual's endowment position. (The overbar will be used henceforth to represent endowed quantities.) The equation for the budget line NN′ is:

$$P_1 c_1 + P_2 c_2 = P_1 \bar{c}_1 + P_2 \bar{c}_2 \tag{2.1.4}$$

Maximizing expected utility from (2.1.1), subject to the budget constraint (2.1.4), leads (assuming an interior solution) to the indifference-curve tangency[4] condition:

$$\frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} = \frac{P_1}{P_2} \tag{2.1.5}$$

4 The necessary conditions for maximizing expected utility are obtained from the usual Lagrangian expression:

$$L(c_1, c_2, \lambda) = U(c_1, c_2) - \lambda(P_1 c_1 + P_2 c_2 - P_1 \bar{c}_1 - P_2 \bar{c}_2)$$

Using the expected-utility formula (2.1.1), setting the partial derivatives equal to zero implies:

$$\pi_1 v'(c_1) = \lambda P_1 \text{ and } \pi_2 v'(c_2) = \lambda P_2$$

Dividing the first equality by the second, we obtain (2.1.5). Conceivably, however, the tangency conditions cannot be met in the interior (i.e., for non-negative c_1, c_2) in which

Thus, at the individual's risk-bearing optimum, shown as point C* in Figure 2.1 along indifference curve U^{II}, the quantities of state claims held are such that the ratio of the probability-weighted marginal utilities equals the ratio of the state-claim prices.

Making the obvious generalization to S states, we arrive at an equation that will be used repeatedly throughout the book:

FUNDAMENTAL THEOREM OF RISK BEARING

$$\frac{\pi_1 v'(c_1)}{P_1} = \frac{\pi_2 v'(c_2)}{P_2} = \cdots = \frac{\pi_s v'(c_s)}{P_s} \qquad (2.1.6)$$

In words: Assuming an interior solution, at the individual's risk-bearing optimum the expected (probability-weighted) marginal utility per dollar of income will be equal in each and every state. (The interior-solution condition will henceforth be implicitly assumed, except where the contrary is indicated.)

In terms of the simplified two-state optimum condition (2.1.5), we can reconsider once again the acceptance or rejection of fair gambles. If a gamble is fair, then $\pi_1 \Delta c_1 + \pi_2 \Delta c_2 = 0$ – the mathematical expectation of the contingent net gains must be zero. But in market transactions $P_1 \Delta c_1 + P_2 \Delta c_2 = 0$ – the exchange value of what you give up equals the value of what you receive. So if the price ratio P_1/P_2 equals the probability ratio π_1/π_2, the market is offering an opportunity to transact fair gambles. Geometrically, the line NN′ would coincide with LL′ in Figure 2.1. It follows immediately from Equation (2.1.5) that the tangency optimum would be the certainty combination where $c_1 = c_2 = \hat{c}$.

Thus, confirming our earlier result, *starting from a certainty position* a risk-averse individual would never accept any gamble at fair odds. But, if his initial endowment were not a certainty position (if $\bar{c}_1 \neq \bar{c}_2$), when offered the opportunity to transact at a price ratio corresponding to fair odds he would want to "insure" by moving to a certainty position – as indicated by the solution $\hat{C}$ along the fair market line LL′. Thus, an individual with an uncertain endowment might accept a "gamble" in the form of a risky contract offering contingent income in one state in exchange for income in another. But he would accept only very particular risky contracts, those that offset the riskiness of his endowed gamble. (Notice that mere acceptance

case the optimum would be at an intersection of the budget line with one of the axes. If such holds at the c_1-axis, (2.1.5) would be translated to an inequality:

$$\frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} > \frac{P_1}{P_2}$$

where $c_2 = 0$, $c_1 = (P_1\bar{c}_1 + P_2\bar{c}_2)/P_1$.

of a risky contract therefore does not tell us whether the individual is augmenting or offsetting his endowed risk.) Finally, if the market price ratio did not represent fair odds, as in the case of market line NN′ in Figure 2.1, whether or not he starts from a certainty endowment the individual *would* accept some risk; his tangency optimum would lie off the 45° line at a point like C* in the direction of the favorable odds.

2.1.2 Regimes of Asset Markets – Complete and Incomplete

In their risk-bearing decisions, individuals do not typically deal directly with elementary state claims – entitlements to consumption income under different states of the world like war versus peace, prosperity versus depression, etc. Rather, a person is generally endowed with, and may be in a position to trade, *assets* like stocks, bonds, and real estate. An asset is a more or less complicated bundle of underlying pure state claims. A share of stock in some corporation is desired by an individual because it promises to yield him a particular amount of income if state 1 occurs, perhaps a different amount under state 2, and so on through the entire list of states of the world that he perceives. There must then be a relationship between the price of any such marketable asset and the underlying values that individuals place upon the contingent-claim elements of the bundle. This is the relationship we now proceed to analyze.

The income yielded by asset a in state s will be denoted z_{as}. Suppose there are only two states of the world $s = 1, 2$ and just two assets $a = 1, 2$ with prices P_1^A and P_2^A.[5] Then the budget constraint can be written:

$$P_1^A q_1 + P_2^A q_2 = P_1^A \bar{q}_1 + P_2^A \bar{q}_2 \equiv \bar{W} \tag{2.1.7}$$

Here, q_1 and q_2 represent the numbers of units held of each asset, and as usual the overbar indicates endowed quantities. The individual's endowed wealth, $\bar{W}$, is defined as the market value of his asset endowment.

Someone might possibly hold an asset as a "non-diversified" (single-asset) portfolio, in which case $q_a = \bar{W}/P_a^A$ for the single asset held (while $q_{a^o} = 0$ for any other asset $a^o \neq a$). The vector of state-contingent incomes generated by such a single-asset portfolio would be:

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} z_{a1} \\ z_{a2} \end{pmatrix} \bar{W}/P_a^A \tag{2.1.8}$$

More generally, a person in a two-asset world will hold some fraction K_1 of his wealth in asset 1 and $K_2 \equiv 1 - K_1$ in asset 2, so that

[5] *Asset* prices will be written P_1^A, P_2^A, etc. to distinguish them from *state-claim* prices which have numerical subscripts only (P_1, P_2, etc.) Throughout this discussion we continue to assume that individuals are price-takers in all markets.

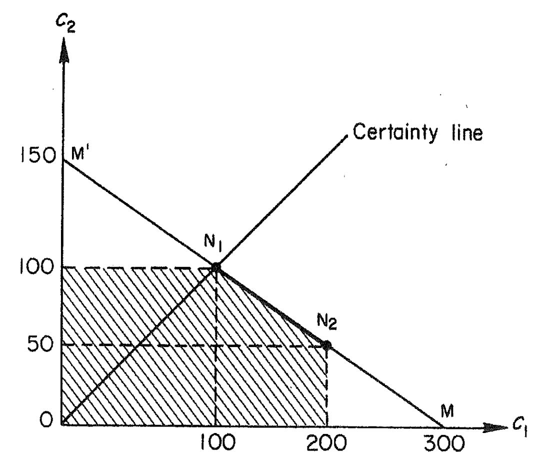

Figure 2.2. Trading in asset markets.

$q_1 = K_1\,\bar{W}/P_1^A$ and $q_2 = K_2\,\bar{W}/P_2^A$. Then the contingent incomes from the portfolio will be:

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \equiv q_1 \begin{pmatrix} z_{11} \\ z_{12} \end{pmatrix} + q_2 \begin{pmatrix} z_{21} \\ z_{22} \end{pmatrix} = K_1\,(\bar{W}/P_1^A) \begin{pmatrix} z_{11} \\ z_{12} \end{pmatrix} + K_2\,(\bar{W}/P_2^A) \begin{pmatrix} z_{21} \\ z_{22} \end{pmatrix} \tag{2.1.9}$$

Equation (2.1.9) expresses the vector of state-contingent portfolio incomes as the share-weighted average of the incomes generated by the two single-asset portfolios.

For concreteness, define a unit of a *certainty asset* (asset 1) so that the contingent returns are unity for each state: $(z_{11}, z_{12}) = (1, 1)$. Let its price be $P_1^A = 1$. Suppose there is also an asset 2 that pays off relatively more heavily in state 1 than in state 2 – specifically, $(z_{21}, z_{22}) = (4, 1)$ – and that its price is $P_2^A = 2$. But imagine that the individual is initially endowed with nothing but 100 units of asset $1(\bar{q}_1 = 100,\ \bar{q}_2 = 0)$. Then point N_1 in Figure 2.2 pictures the implied endowed contingent incomes $(\bar{c}_1, \bar{c}_2) = (100, 100)$. Since the value of the individual's endowment is $\bar{W} = 100$, he could trade away his entire endowment of asset 1 for 50 units of asset 2 and attain the final consumptions $(c_1, c_2) = (200, 50)$ – point N_2 in the diagram. More generally, from Equation (2.1.9) we see that the final consumption vector (c_1, c_2) will lie along the line joining N_1 and N_2, at distances toward N_1 and N_2 in proportion to the relative wealth shares K_1 and K_2.

If the individual were constrained to hold non-negative amounts of assets, the opportunity boundary would be only the line segment between N_1 and N_2 in the diagram. However, it is entirely permissible to let either q_1 or q_2 go negative. A negative q_1 (which implies, of course, a negative $K_1 \equiv P_1^A q_1/\bar{W}$) means that the individual holds *liabilities* rather than assets of type 1; he is, in effect, committed to *deliver* the amount $|q_1\ z_{11}|$ if state 1 occurs and the amount $|q_1\ z_{12}|$ if state 2 occurs. (This is sometimes described as being in a "short" position with regard to asset 1.) We will, however, be imposing a *non-negativity constraint upon the ultimate* c_1, c_2 combinations *arrived* at; the individual cannot end up consuming negative income in either state of the world. He cannot therefore go short on any asset to the extent of violating any of his delivery commitments – in effect, he is not permitted to "go bankrupt" in any state of the world. (And, *a fortiori*, he cannot go short on *all* assets simultaneously!) This means that, while the trading possible along the line MM′ in Figure 2.2 may extend beyond points N_1 and N_2, the attainable combinations remain bounded by the vertical and horizontal axes.

Having described the feasible alternatives, we now consider the individual's actual portfolio-choice decision. Since the consumption vector (c_1, c_2) is generated by his asset holdings as shown in (2.1.9), the individual can be regarded as choosing his portfolio asset shares $(K_1, K_2) \equiv (q_1 P_1^A/\bar{W},\ q_2 P_2^A/\bar{W})$ so as to maximize expected utility subject to his asset-holding constraint, that is:

$$\underset{(K_1,K_2)}{\text{Max}}\ U = \pi_1 v(c_1) + \pi_2 v(c_2) \text{ subject to } K_1 + K_2 = 1$$

From (2.1.9) we know that:

$$\frac{\partial c_s}{\partial K_1} = \frac{\bar{W}}{P_1^A} z_{1s} \quad \text{and} \quad \frac{\partial c_s}{\partial K_2} = \frac{\bar{W}}{P_2^A} z_{2s}$$

Then the endowed wealth cancels out of the first-order condition for an interior optimum, which can be written:[6]

$$\frac{\sum_{s=1}^{2} \pi_s v'(c_s) z_{1s}}{P_1^A} = \frac{\sum_{s=1}^{2} \pi_s v'(c_s) z_{2s}}{P_2^A}$$

6 The Lagrangian expression is:

$$L(K_1, K_2, \lambda) = \pi_1 v(c_1) + \pi_2 v(c_2) - \lambda(K_1 + K_2 - 1)$$

Setting the partial derivatives equal to zero leads to:

$$\pi_1 v'(c_1)\,(\bar{W}/P_1^A)\, z_{11} + \pi_2 v'(c_2)\,(\bar{W}/P_1^A)\, z_{12} = \lambda$$

$$\pi_1 v'(c_1)\,(\bar{W}/P_2^A)\, z_{21} + \pi_2 v'(c_2)\,(\bar{W}/P_2^A)\, z_{22} = \lambda$$

(Henceforth, the maximization calculus will not be spelled out in detail except where points of special interest or difficulty arise.)

This says that, at his risk-bearing optimum, the individual will have adjusted his holdings of the two assets until their given prices become proportional to the expected marginal utilities he derives from the contingent consumptions they generate. Or, we can say: at the optimum, he will derive the same *expected marginal utility per dollar* held in each asset.

An obvious generalization to any number A of assets and S of states leads to an adaptation of (2.1.6), The Fundamental Theorem of Risk Bearing, for a regime of asset markets:

RISK-BEARING THEOREM FOR ASSET MARKETS

$$\frac{\sum_s \pi_s v'(c_s) z_{1s}}{P_1^A} = \frac{\sum_s \pi_s v'(c_s) z_{2s}}{P_2^A} = \cdots = \frac{\sum_s \pi_s v'(c_s) z_{As}}{P_A^A} \qquad (2.1.10)$$

We have now described the individual's optimal risk-bearing decision (i) in a market of elementary state claims and (ii) in a market of more generally defined assets. It is natural to ask if trading in asset markets can replicate the results of a regime in which all state claims are explicitly traded. The answer turns out to depend upon whether the set of tradable assets constitutes a regime of *complete* or *incomplete* markets. Intuitively, markets are complete if a rich enough class of assets is traded so that the equilibrium consumption of the state-claim regime is attained. That is, the consumption vector $c_1, c_2, \ldots, c_S$ that satisfies equilibrium condition (2.1.10) also satisfies the equilibrium condition (2.1.6).

Complete Markets

Returning to the numerical example depicted in Figure 2.2, where A (the number of distinct assets) and S (the number of states) both equal 2, if the individual has endowment N_1 and can trade elementary state claims at prices P_1 and P_2 his budget constraint (line MM′ in the diagram) would be:

$$P_1 c_1 + P_2 c_2 = P_1(100) + P_2(100) = \bar{W}$$

A market regime allowing trading in all the elementary state claims is obviously complete. We will call it a regime of Complete Contingent Markets (CCM). The CCM regime provides a benchmark for measuring the completeness of alternative asset-market regimes.

In any asset-trading regime, the prices of assets can be directly computed if the state-claim prices are known. Specifically in our example, since any asset a has state-contingent yields z_{a1}, z_{a2}, the market values of assets 1 and 2 are:

$$P_1^A = z_{11}P_1 + z_{12}P_2 = P_1 + P_2$$

$$P_2^A = z_{21}P_1 + z_{22}P_2 = 4P_1 + P_2$$

In order to establish whether an asset-market regime is or is not complete, we must invert this analysis. That is, for given asset prices P_a^A $(a = 1, \ldots, A)$, the question is whether or not it is possible to extract the implicit state-claim prices. In our example, knowing that $(z_{11}, z_{12}) = (1, 1)$ and $(z_{21}, z_{22}) = (4, 1)$, we can rewrite the above equations in matrix form:

$$\begin{bmatrix} P_1^A \\ P_2^A \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 4 & 1 \end{bmatrix} \begin{bmatrix} P_1 \\ P_2 \end{bmatrix}$$

As the two rows are not proportional, we can invert the matrix and obtain:

$$\begin{bmatrix} P_1 \\ P_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 4 & 1 \end{bmatrix}^{-1} \begin{bmatrix} P_1^A \\ P_2^A \end{bmatrix} = \begin{bmatrix} -\frac{1}{3} & \frac{1}{3} \\ \frac{4}{3} & -\frac{1}{3} \end{bmatrix} \begin{bmatrix} P_1^A \\ P_2^A \end{bmatrix}$$

So in this case it *is* possible to compute "implicit" state-claim prices from given asset prices. In our example, if (say) $P_1^A = 1$ and $P_2^A = 2$, then $P_1 = 1/3$ and $P_2 = 2/3$. So the implicit budget constraint in state-claim units would be:

$$\frac{1}{3}c_1 + \frac{2}{3}c_2 = \frac{1}{3}(100) + \frac{2}{3}(100) = 100$$

Thus, on the assumption (as already discussed) that traders are allowed to "go short" on either asset so long as they can guarantee delivery, the asset-market equilibrium is the same as would be attained under CCM. We will call a set of asset markets meeting this condition a regime of Complete Asset Markets (CAM).

Generalizing this example, suppose there are S states and A assets, and exactly S of the assets have *linearly independent* yield vectors. That is, suppose it is impossible to express any one of these S yield vectors as a linear sum of the other $S - 1$ asset yields. In economic terms this means that it is not possible to duplicate any of these S assets by buying a combination of the other $S - 1$ assets – i.e., all of the S assets are economically distinct.[7] (Of

[7] In our simple example with $A = S = 2$, the two yield vectors (z_{11}, z_{12}) and (z_{21}, z_{22}) were linearly independent since otherwise one vector would have been a scalar multiple of the other. That is, the two rows of the z-matrix were not proportional, which is what permitted inverting the matrix.

course, this case can only come about if $A \geq S$, although A greater than or equal to S does not *guarantee* the existence of S linearly independent assets.) So, a CAM regime exists, in a world of S states, if among the A assets there are S with linearly independent yield vectors.

Summarizing in compact and general form, given a state-claim price vector $(P_1, \ldots, P_s)$ the market value of asset a is:

$$P_a^A = \sum_s z_{as} P_s \tag{2.1.11}$$

Or, in matrix notation for the entire set of assets:

$$P^A = P[z_{as}] \equiv PZ \tag{2.1.12}$$

This permits us always to generate asset prices from a known state-claim price vector. But the reverse can be done only under linear independence (so that the matrix $Z \equiv [z_{as}]$ can be inverted). If so, the asset-market regime is complete (CAM holds): for any prices $P_1^A, \ldots, P_A^A$ of the A assets there will be a unique implicit state-claim price vector $(P_1, \ldots, P_s)$. It follows that under CAM, the Fundamental Theorem of Risk Bearing (2.1.6) also holds, in addition to the weaker Risk-Bearing Theorem for Asset Markets (2.1.10).

Incomplete Markets

Consider now a three-state world with asset trading. For this trading regime to be complete, there would have to be three assets with linearly independent return vectors. In Figure 2.3 the points N_1, N_2, N_3 represent an individual's three attainable single-asset portfolios, for assets 1, 2, 3, respectively, while E indicates his endowed mixed portfolio of these three assets. The interior of the shaded triangle N_1 N_2 N_3 represents the state-contingent consumption combinations attainable by holding non-negative amounts of all three assets. As in the previous two-asset case, however, there is no reason to exclude "going short" on any asset – so long as the individual ends up with non-negative *consumption* entitlements in all three states, i.e., in the positive orthant. If going short to this extent is allowed, any point in the larger triangle MM′M″ is an income combination attainable by holding some diversified portfolio. The equation of the "budget plane" through MM′M″ is:

$$P_1 c_1 + P_2 c_2 + P_3 c_3 = \bar{W} = P_1^A q_1 + P_2^A q_2 + P_3^A q_3$$

It will be evident that, if the asset-market budget constraint corresponds to the full triangle MM′M″, we have a CAM regime: the choice of an optimal asset portfolio at given asset prices (P_1^A, P_2^A, P_3^A) is equivalent to choosing

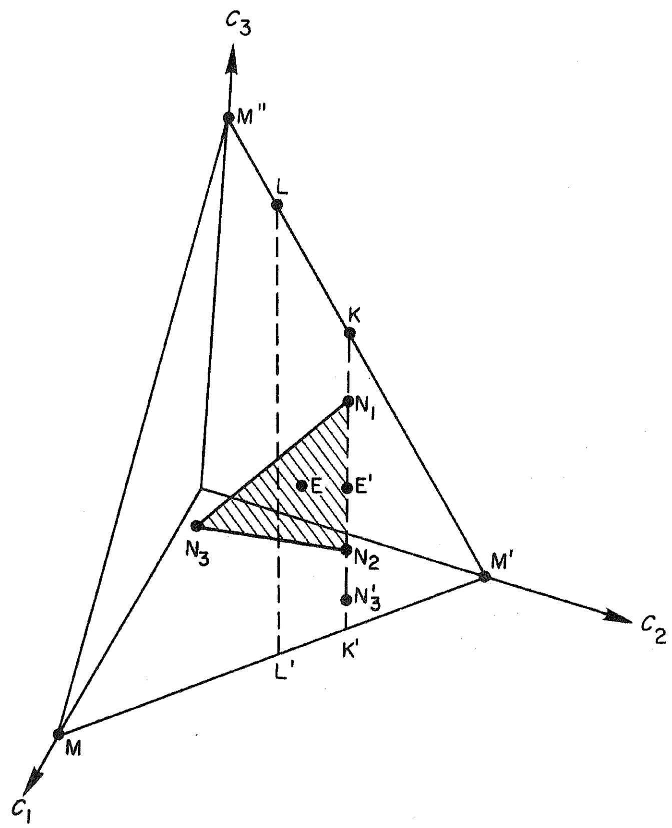

Figure 2.3. Alternative patterns of incomplete markets.

an optimal state-claim consumption vector given some endowment point and a full set of implicit state-claim prices (P_1, P_2, P_3).

Markets in such a three-state world can be *incomplete* in several distinct ways. First, there might simply be fewer than three assets available. In Figure 2.3, if only assets 1 and 2 exist the individual's market opportunities consist only of the state-claim combinations shown by a "degenerate" budget constraint – the market line KK′ through points N_1 and N_2. In this case, the endowment E′ is a mixed portfolio of the two assets. Second, it might be that there is a third asset, in addition to asset 1 and 2, but this third asset is linearly dependent[8] upon 1 and 2 – indicated geometrically by the collinearity of points N_1, N_2, and N_3'. The line KK′ through these three points remains degenerate; once again, not all the c_1, c_2, c_3 state-claim

[8] That is, it is possible to find an α and β such that $(z_{31}, z_{32}, z_{33}) = \alpha(z_{11}, z_{12}, z_{13}) + \beta\,(z_{21}, z_{22}, z_{23})$.

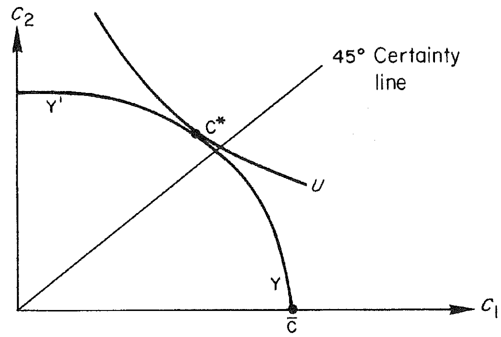

Figure 2.4. Risk-bearing with production.

combinations are attainable by market trading. Third, it might be that the third asset, while present in an individual's endowment and linearly independent of the other two, is *non-tradable.* (An example of such a non-tradable asset might be one's "human capital.") Suppose now that point E in Figure 2.3 represents an endowment containing positive amounts of a non-tradable asset 3 as well as of marketable assets 1 and 2. Here the dotted line LL′ is the "degenerate" budget constraint for the individual. Note that LL′ is parallel to the KK′ line that applied when there was no third asset at all.

In each of these cases there is no longer equivalence between trading in asset markets and trading in CCM. It follows, and this is the crucial point, that, while the Risk-Bearing Theorem for Asset Markets (2.1.10) will always hold, the Fundamental Theorem of Risk Bearing (2.1.6) does not.

When and why it is that *incomplete* trading regimes exist, despite the disadvantages just described, is a question we must leave to Chapter 4.

2.1.3 Productive Opportunities

So far in this section we have considered only the risk-bearing decisions of individuals in markets. But it is also possible to respond to risk by *productive* adjustments.

A Robinson Crusoe isolated from trading can adapt to risk solely by productive transformation. Before he takes productive action, suppose Robinson's corn crop is sure to be good if the weather is moist (state 1) but will fail entirely if the weather is dry (state 2). Thus, Robinson's endowment position $\bar{C}$ is along the c_1-axis of Figure 2.4. However, by installing irrigation systems of greater or lesser extent, Robinson can improve his state-2 crop y_2. On the other hand, the effort required to do so will divert him from ordinary cultivation, and hence reduce his state-1 crop y_1. Then Robinson's feasible

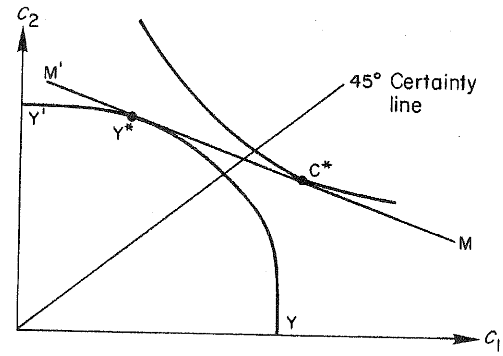

Figure 2.5. Productive and consumptive risk bearing.

state-contingent outputs (y_1, y_2) might lie in the convex set bounded by the axes and the curve YY′. This curve is "bowed away from the origin" reflecting the operation of diminishing returns. Robinson's productive-consumptive risk-bearing optimum will evidently be at point C* (*not* in general on the 45° certainty line) where the production frontier YY′ is tangent to his highest attainable indifference curve.

Writing Robinson's productive opportunity constraint as $F(y_1, y_2) = 0$, his optimum (tangency) condition can be expressed as:

$$\frac{\partial F/\partial y_1}{\partial F/\partial y_2} \equiv -\frac{dy_2}{dy_1}\bigg|_F = -\frac{dc_2}{dc_1}\bigg|_U \equiv \frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} \qquad (2.1.13)$$

By an obvious modification of our earlier argument, Robinson's optimum will lie on the 45° line only in the exceptional case where his Marginal Rate of Technical Substitution (the absolute slope – dy_2/dy_1 along YY′) happens to be exactly equal to the probability ratio π_1/π_2 at the point where YY′ cuts the certainty line. So, in his productive decisions, Robinson will not in general want to avoid all risk, even though it may be possible for him to do so. Some risks are profitable enough to be worth taking, i.e., they may represent sufficiently favorable productive gambles.

Now consider individuals who can combine both market opportunities and physical productive opportunities. In general, any such individual will have a productive optimum (indicated by Y* in Figure 2.5) distinct from his consumptive optimum (indicated by C*). The availability of markets for trading contingent claims makes it possible to separate *productive* risk bearing from *consumptive* risk bearing. An example in everyday terms: a corporation may engage in risky productive activities, yet the shareholders may be able to largely eliminate personal risks by diversifying their individual portfolios.

Without going through the straightforward derivation, we will state the conditions for the individual's productive and consumptive optimum positions – assuming a regime of CCM with explicit state-claim prices P_1 and P_2 (or an equivalent CAM regime for which the corresponding prices are implicitly calculable):

$$-\left.\frac{dy_2}{dy_1}\right|_F \overset{\substack{\text{PRODUCTIVE}\\\text{OPTIMUM}\\\text{CONDITION}}}{=} \frac{P_1}{P_2} \overset{\substack{\text{CONSUMPTIVE}\\\text{OPTIMUM}\\\text{CONDITION}}}{=} -\left.\frac{dc_2}{dc_1}\right|_U \tag{2.1.14}$$

The price ratio here may be said to *mediate* between the individual's productive optimum and consumptive optimum. Whereas the Crusoe condition (2.1.13) required a direct equality of the YY′ slope with an indifference curve at a single common productive-consumptive optimum, the availability of markets makes it possible for a person to separate his Y* and C* positions and thereby attain improved combinations of contingent consumptions.

As no essentially new ideas depend thereon, the generalized productive solutions for any number S of states of the world and the complexities introduced by regimes of incomplete markets will not be detailed here.

Exercises and Excursions 2.1

1 *Linear Independence*

(A) With state yields expressed in the form (z_{a1}, z_{a2}, z_{a3}), the rows below indicate four different three-asset combinations, labeled (i) through (iv). Verify that only asset combinations (i) and (ii) are linearly independent.

	$a = 1$	$a = 2$	$a = 3$
(i)	(1,0,0)	(0,1,0)	(0,0,1)
(ii)	(1,1,1)	(1,4,0)	(0,7,1)
(iii)	(0,2,3)	(1,0,1)	(0,4,6)
(iv)	(1,3,2)	(4,0,5)	(2,2,3)

(B) For each of the combinations above, *if* it is possible to have $P_1^A = P_2^A = P_3^A = 1$ what can you say about the implied state-claim prices P_1, P_2, P_3? For given asset endowment holdings $\bar{q}_1 = \bar{q}_2 = \bar{q}_3 = 1$ solve for and picture the market plane MM′M″ in state-claim space, wherever it is possible to do so. (Where it is not possible to do so, picture the relevant trading opportunity constraint.)

(C) For cases (i) and (ii) only of (A) above, assume instead that the endowment is given in state-claim units as $(\bar{c}_1, \bar{c}_2, \bar{c}_3) = (1, 1, 1)$, and that $P_2^A = P_3^A = 1$ while no trading is possible in asset 1. Picture the trading opportunity constraint in state-claim space.

2 *Non-negativity*

(A) For each of the combinations in 1(A) above, would the asset-holding portfolio $q_1 = -1$, $q_2 = q_3 = 1$ violate the non-negativity constraint on state incomes?

(B) Suppose case (i) above were modified by replacing $a = 1$ with a new $a = 1'$ whose returns are $(-1, 2, 3)$. Would the combination $q_1 = q_2 = q_3 = 1$ be feasible? What if this new asset were instead to replace the first asset in case (ii) above?

3 *Risk-Bearing Optimum*

(A) In cases (i) and (ii) under 1(A), if explicit trading in state claims is ruled out, find the individual's risk-bearing optimum expressed as (q_1^*, q_2^*, q_3^*) in asset units and as (c_1^*, c_2^*, c_3^*) in state-claim units – if: $\pi_1 = \pi_2 = \pi_3 = 1/3$, $P_1^A = P_2^A = P_3^A = 1$, $\bar{q}_1 = \bar{q}_2 = \bar{q}_3 = 1$, and $v(c) = \ln c$.

(B) What can you say about cases (iii) and (iv)?

4 *Consumer Choice*

An individual with utility function $v(c) = \ln c$ must choose a state-contingent consumption bundle $(c_1, \ldots, c_S)$. The price of a state-s claim is P_s and the consumer's initial endowment has a value of $\bar{W}$.

(A) Solve for the individual's optimum in each state.

(B) Hence show that for any pair of states s and s':

$$\frac{c_s}{c_{s'}} = \frac{\pi_s}{\pi_{s'}} \frac{P_{s'}}{P_s}$$

(C) What condition defines the state in which consumption is greatest? Least?

(D) Is the rule derived in (C) true for any concave utility function $v(c)$?

5 Portfolio Choice

Asset 1 and asset 2 both cost \$150. Yields on asset 1 in states 1 and 2 are $(z_{11}, z_{12}) = (100, 200)$ and on asset 2 are $(z_{21}, z_{22}) = (200, 100)$. An individual with an initial wealth of \$150 has a utility function:

$$v(c) = -e^{-c}$$

(A) Show that the state-contingent budget constraint can be expressed as:

$$c_1 + c_2 = 300$$

(B) If the individual believes that state 1 will occur with probability π, show that his optimal consumption in state 1 is:

$$c_1^* = 150 + \frac{1}{2}\ln(\pi/(1-\pi))$$

(C) If q_1 is the number of units of asset 1 purchased show that:

$$c_1^* = 200 - 100q_1^*$$

and hence obtain an expression for q_1^* in terms of π, the probability of state 1.

(D) What values do c_1^* and q_1^* approach as the probability of state 1 becomes very small?

2.2 Choosing Combinations of Mean and Standard Deviation of Income

2.2.1 μ, σ, Preferences

We have described decision making under uncertainty as choice among actions or prospects $x = (c_1, \ldots, c_S; \pi_1, \ldots, \pi_S)$ – probability distributions that associate an amount of contingent consumption in each state of the world with the degree of belief attaching to that state. There is another approach to the risk-bearing decision that has proved to be very useful in modern finance theory and its applications. This alternative approach postulates that, for any individual, the probability distribution associated with any prospect is effectively represented by just two summary statistical measures: the *mean* and the *standard deviation* of income. Specifically, the

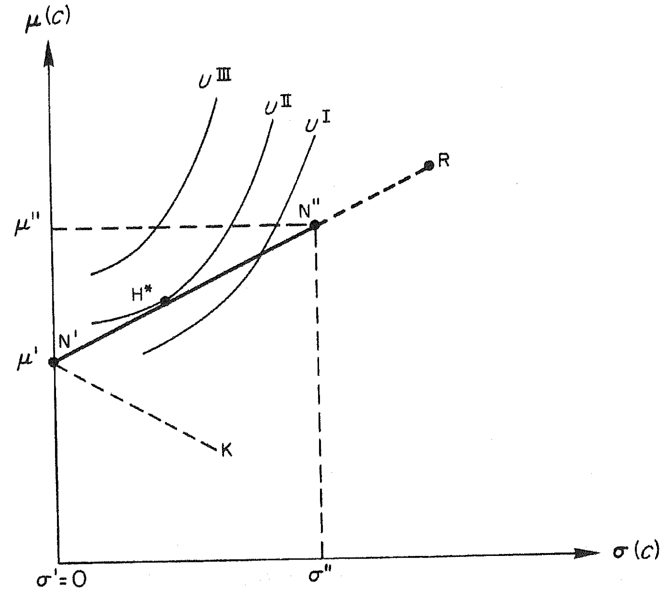

Figure 2.6. Portfolio choice with one riskless and one risky asset.

individual is supposed always to prefer higher average income (measured by the expectation or mean $\mu(c)$ of the probability distribution achieved by holding any particular portfolio of assets) and lower variability of income (measured by the standard deviation $\sigma(c)$).[9] His preferences can therefore be represented by the indifference curves pictured on $\mu(c)$ and $\sigma(c)$ axes as in Figure 2.6.

The approach in terms of preference for high $\mu(c)$ and low $\sigma(c)$ is broadly consistent with the previous analysis. In maximizing expected utility $E[v(c)]$ under uncertainty, other things equal, a higher average level of income is surely to be preferred. And, given risk aversion, the theorem that fair gambles would not be accepted implies that distributions with low $\sigma(c)$ tend to be more desirable. Nevertheless, in moving from a probability distribution that was fully defined in terms of consequences in each and every state of the world to a mere statistical summary of that distribution – one that makes use only of the two parameters $\mu(c)$ and $\sigma(c)$ – some information has been lost. The question addressed here is: when, if ever, is such a reduction valid,

[9] Some analysts prefer to think in terms of the variance of income $\sigma^2(c)$. But for purposes of economic interpretation the standard deviation is more convenient, since $\sigma(c)$ has the same dimensionality as $\mu(c)$ and c itself.

exactly or at least approximately? That is, when can we justifiably convert $U = \mathrm{E}[v(c)]$ into a function only of $\mu(c)$ and $\sigma(c)$?

To indicate the nature of the approximation involved, $v(c)$ can be expanded in a Taylor's series about its expected value $\mathrm{E}[\tilde{c}] = \mu$:[10]

$$v(\tilde{c}) = v(\mu) + \frac{v'(\mu)}{1!}(\tilde{c}-\mu) + \frac{v''(\mu)}{2!}(\tilde{c}-\mu)^2 + \frac{v'''(\mu)}{3!}(\tilde{c}-\mu)^3 + \cdots$$

Taking expectations, remembering that $U = \mathrm{E}[v(\tilde{c})]$ while noticing that the expectation of $(\tilde{c} - \mu)$ is zero and that the expectation of $(\tilde{c} - \mu)^2$ is the variance $\sigma^2(\tilde{c})$, we have:

$$U = v(\mu) + \frac{v''(\mu)}{2!}\sigma^2 + \frac{v''(\mu)}{3!}\mathrm{E}[(c-\mu)^3] + \cdots \qquad (2.2.1)$$

The omitted terms suggested by the dots are functions of the fourth or higher powers of $(\tilde{c} - \mu)$ – higher moments about the mean, in statistical terminology.

Possible justifications for treating U as a function only of the mean and standard deviation of income may be found (i) in the properties we are willing to assume for the utility function $v(c)$ or (ii) in the properties of the probability distribution of $\tilde{c}$.

(i) First of all, suppose that the $v(c)$ function is quadratic, so that it can be written (with K_0, K_1, and K_2 as constants):

$$v(c) = K_0 + K_1 c - \frac{1}{2}K_2 c^2 \qquad (2.2.2)$$

where $K_1, K_2 > 0$. Then the third derivative $v'''(c)$ is always zero, as are all higher derivatives. So (2.2.1) can be expressed more specifically as:

$$U = K_0 + K_1\mu - \frac{1}{2}K_2(\mu^2 + \sigma^2)$$

With U as parameter, this equation represents a family of indifference curves on μ, σ, axes in Figure 2.6. By completing the square it may be verified that the curves constitute a set of concentric circles, the center being $\mu = K_1/K_2$, $\sigma = 0$.

However, the utility function $v(c)$ given by (2.2.2) has an economically unacceptable implication – that the marginal utility of income, $v'(c) = K_1 - K_2\, c$, eventually becomes negative. A quadratic $v(c)$ function thus leads to a highly special indifference-curve map, with acceptable properties only over a limited range.

[10] A tilde overlying any symbol indicates a random variable. We will use this notation only when it is desired to emphasize that feature.

(ii) Turning now to possible justifications that run in terms of probability distributions for $\tilde{c}$, the Central Limit Theorem of probability theory offers a lead. While a full discussion would be out of place here, the Central Limit Theorem essentially says that the distribution of the sum of any large number N of random variables approaches the *normal distribution* as N increases – provided only that the variables are not too correlated. The point is that the overall or portfolio income $\tilde{c}$ yielded by an individual's holdings of assets can be regarded as the sum of underlying random variables, each summand representing the income generated by one of the assets entering into his portfolio. The normal distribution is fully specified by just two parameters, its mean and standard deviation. Then, in Equation (2.2.1), while the terms involving higher moments do not all disappear,[11] the higher moments remaining are functions of the mean and standard deviation.[12] It therefore follows that indifference curves for alternative normal distributions of consumption income $\tilde{c}$ can be drawn on $\mu(c)$, $\sigma(c)$ axes.[13]

The tendency toward normality under the Central Limit Theorem is the stronger, roughly speaking, the closer to normal are the underlying random variables, the more equal are their weights in the summation, and the less correlated they are with one another. Looking at portfolio income as the summation variable, income yields of the assets that comprise portfolios will rarely if ever have normal distributions themselves. In particular, the normal distribution extends out to negative infinity, whereas any "limited liability" asset cannot generate unlimited negative income. (And even without limited liability, personal bankruptcy establishes a lower limit on how large a negative yield the individual need consider.) Furthermore, asset weights in portfolios tend to be highly unequal: a person will likely have more than half his income associated with his wage earnings – the income generated by his single "human capital" asset. And, finally, there typically is

[11] Since the normal distribution is symmetrical about its mean, all the higher *odd* moments are zero, but the even moments do not disappear.

[12] There are other families of statistical distributions, besides the normal, that are fully specified by the mean and standard deviation. However, as we have seen, the Central Limit Theorem leads specifically to the normal as the approximating distribution of portfolio income.

[13] We have not, however, justified the standard shape of the preference map pictured in Figure 2.9. A proof that normally distributed returns imply positive indifference-curve slope and curvature, as shown in the diagram, is provided in Copeland, Weston, and Shastri (2004).

considerable correlation among returns on the different assets making up any portfolio. Portfolios for which the Central Limit Theorem justifies use of the normal distribution as approximation are called "well-diversified"; unfortunately, we have no handy rule for deciding when a portfolio may be considered well-diversified. For all the reasons given above, use of the normal distribution as approximation remains subject to considerable questions.

It is of interest to consider the effect upon utility of the third moment $E[(\tilde{c} - \mu)^3]$– entering into the leading term dropped from (2.2.1) if the normal approximation is adopted. The third moment is a measure of *skewness*: skewness is zero if the two tails of a distribution are symmetrical, positive if the probability mass humps toward the left (so that the right tail is long and thin), and negative in the opposite case. To see the effect of skewness, consider an investor choosing between two gambles with the same means and standard deviations. Specifically, suppose gamble J offers 0.999 probability of losing $1 and 0.001 probability of gaining $999, while gamble K offers 0.999 probability of gaining $1 and 0.001 probability of losing $999. J and K have the same mean (zero) and the same standard deviation, but J is positively skewed while K is negatively skewed. Almost all commercial lotteries and games of chance are of form J, thus suggesting that individuals tend to prefer positive skewness. While the primary purpose of diversification is to reduce the standard deviation of income, diversification also tends to eliminate skewness – since the normal distribution that is approached has zero skewness. We would expect to see, therefore, lesser desire to diversify where skewness of the portfolio held is positive, greater desire to diversify where skewness is negative. But the main point is that the survival of preference for positive skewness suggests that individual real-world portfolios are typically not so well-diversified.

The upshot, then, is that the attempt to reduce preferences for income prospects to preferences in terms of $\mu(c)$ and $\sigma(c)$ falls short of being fully satisfying. But the approach remains an eminently manageable approximation, expressed as it is in terms of potentially measurable characteristics of individual portfolios and (as we shall see shortly) of the assets that comprise portfolios. The ultimate test of any such approximation is, of course, its value as a guide to understanding and prediction.

2.2.2 Opportunity Set and Risk-Bearing Optimum

In examining the individual's opportunities for achieving combinations of mean and standard deviation of portfolio income – $\mu(c)$ and $\sigma(c)$ – in his risk-bearing decisions, we need to show how these statistical properties of consumption income emerge from the yields generated by the separate assets held.

For any asset a, let μ_a represent the mean of the income yield $\tilde{z}_a$ per unit of a held. Let σ_a represent the standard deviation of $\tilde{z}_a$, and σ_{ab} the covariance of $\tilde{z}_a$ and $\tilde{z}_b$. Then, following the usual statistical definitions:

$$\mu_a \equiv \mathrm{E}(\tilde{z}_a)$$
$$\sigma_a \equiv [\mathrm{E}(\tilde{z}_a - \mu_a)^2]^{\frac{1}{2}}$$
$$\sigma_{ab} \equiv \mathrm{E}[(\tilde{z}_a - \mu_a)(\tilde{z}_b - \mu_b)]$$

And, of course, $\sigma_a \equiv (\sigma_{aa})^{\frac{1}{2}}$.

If the individual holds a portfolio consisting of q_a units each of assets $a = 1, \ldots, A$ his portfolio income statistics are related to the asset return parameters above via:

$$\mu(c) \equiv \mu \equiv \sum_a q_a \mu_a$$

$$\sigma(c) \equiv \sigma \equiv \left(\sum_a \sum_b q_a \sigma_{ab} q_b \right)^{\frac{1}{2}} \equiv \left(\sum_a \sum_b \sigma_a q_a \rho_{ab} q_b \sigma_b \right)^{\frac{1}{2}}$$

Here ρ_{ab} is the simple correlation coefficient between the distributions of $\tilde{z}_a$ and $\tilde{z}_b$, using the identity $\sigma_{ab} \equiv \sigma_a \rho_{ab} \sigma_b$ that relates covariance and the correlation coefficient.

The individual's budget constraint can be written in terms of his asset holdings (compare Equation (2.1.7)) as:

$$\sum_a P_a^A q_a = \sum_a P_a^A \bar{q}_a \equiv \bar{W}$$

Drawn on $\mu(c), \sigma(c)$ axes, this budget constraint bounds an opportunity set of feasible combinations of mean and standard deviation of portfolio income. We now need to determine the characteristic shape of this opportunity set. (But notice that, since the individual desires high μ and low σ, he will be interested only in the northwest boundary.)

To start with the simplest case, suppose there are just two assets, and that asset 1 is riskless ($\sigma_1 = 0$) while asset 2 is risky ($\sigma_2 > 0$). For this to be at all an interesting situation, it must, of course, also be true that $\mu_2/P_2^A > \mu_1/P_1^A$–i.e., the risky asset has a higher mean yield per dollar. The portfolio income yield is a random variable given by:

$$\tilde{c} \equiv q_1 z_1 + q_2 \tilde{z}_2 \tag{2.2.3}$$

In this simplest case, with q_2 units of asset 2 purchased at a cost of $P_2^A q_2$, the individual has $\bar{W} - P_2^A q_2$ dollars to invest in the riskless asset 1. For any security a, we can define its rate of return R_a in:

$$\frac{z_a}{P_a^A} \equiv 1 + R_a$$

Then $R_1 \equiv z_1/P_1^A - 1$ is the rate of return on the riskless asset.[14] Expression (2.2.3) can then be rewritten as:

$$\tilde{c} = \left(\bar{W} - P_2^A q_2\right)(1 + R_1) + q_2 \tilde{z}_2$$

And the parameters of the income distribution become:

$$\mu(c) = \bar{W}(1 + R_1) + \left[\mu_2 - (1 + R_1)P_2^A\right]q_2$$
$$\sigma(c) = \sigma_2 q_2$$

It then follows that the budget constraint on μ, σ, axes is a straight line, shown as N′N″ in Figure 2.6. (The opportunity set consists of the line and the area lying below it in the diagram.) Point N′ is the μ, σ, combination μ', σ' attainable by holding a single-asset portfolio consisting of the riskless asset (asset 1) exclusively. For this portfolio, $\mu'(\tilde{c}) = (\bar{W}/P_1^A)z_1 = \bar{W}(1 + R_1)$ while $\sigma'(\tilde{c}) = 0$. Point N″ is the μ, σ, combination μ'', σ'', generated by the single-asset portfolio consisting of the risky asset (asset 2). Here $\mu''(\tilde{c}) = (\bar{W}/P_2^A)\mu_2$ while $\sigma''(\tilde{c}) = (\bar{W}/P_2^A)\sigma_2$.

What about portfolios containing mixtures of the two assets? If the fractional shares of wealth devoted to the riskless and the risky assets are $\alpha = q_1 P_1^A/\bar{W}$ and $\kappa = q_2 P_2^A/\bar{W}$, respectively, where $\alpha + \kappa = 1$, the portfolio statistics μ and σ can each be written in two useful ways:

$$\mu = \alpha\mu' + \kappa\mu'' = \bar{W}(1 + R_1) + \left[\mu_2 - (1 + R_1)P_2^A\right]q_2$$
$$\sigma = \kappa\sigma'' = \sigma_2 q_2$$

[14] This terminology would be appropriate if, as is usually assumed in the finance literature, the assets are purchased (the prices P_a^A are paid out) one time period earlier than the date of the income yields $\tilde{z}_a$. Strictly speaking, such a convention implies an *intertemporal* choice situation, where earlier consumption should be balanced against later consumption, over and above the atemporal risk-bearing choices we have dealt with so far. However, we will follow the finance tradition in using "rate of interest" terminology without necessarily addressing the problem of intertemporal choice.

It follows that the budget constraint can also be written in two ways:

$$\mu = \frac{\mu'' - \mu'}{\sigma''}\sigma + \mu'$$
$$\mu = \bar{W}(1 + R_1) + \left[\mu_2 - (1 + R_1)P_2^A\right]\sigma/\sigma_2$$

It is evident that all the μ, σ combinations satisfying these conditions (meeting the budget constraint) fall along the straight line N′N″ in the diagram. The constant slope $d\mu/d\sigma$ of this budget constraint line is known as the *price of risk reduction*, also known as the *Sharpe ratio*, which we will symbolize as Θ:

THE PRICE OF RISK REDUCTION OR THE SHARPE RATIO

$$\Theta \equiv \frac{d\mu}{d\sigma} = \frac{\mu'' - \mu'}{\sigma''} = \frac{\mu_2 - (1 + R_1)P_2^A}{\sigma_2} \qquad (2.2.4)$$

Note that the steepness of the opportunity line reflects only the market or "objective" data.[15]

To find the individual optimum, the budget constraint must be juxtaposed against the "subjective" data of the individual's preference function. This solution is, of course, a familiar type of tangency: in Figure 2.6, the point H* is the individual's risk-bearing optimum on μ, σ axes.

Could a portfolio be represented by a μ, σ combination along the (dashed) extension of N′N″ lying to the northeast of point N″, which would correspond to holding *negative amounts* of the riskless security? This is sometimes referred to as "selling short" the riskless security, which means incurring a liability requiring *delivery* of the promised amount of income certain. Incurring such a debt can also be thought of as "issuing" units of the riskless security. As explained earlier, doing so would be perfectly reasonable provided that the issuer can really satisfy such an obligation with certainty. Clearly, the opportunity line N′N″ cannot be extended to the northeast without limit: someone with risky assets and riskless liabilities faces some likelihood of bankruptcy, owing to states of the world in which

[15] The parameters involved in the expression for $d\mu/d\sigma$ are "objective" in that they reflect the individual's market opportunities independent of what his personal *preferences* might be. However, a "subjective" element may still enter if *beliefs* about market parameters vary from person to person. While asset prices P_a^A can usually be taken as interpersonally agreed data, there might well be disagreement about some or all of the security yield parameters μ_a, σ_a. If there were such disagreement, the implicit "price of risk reduction" would vary from person to person.

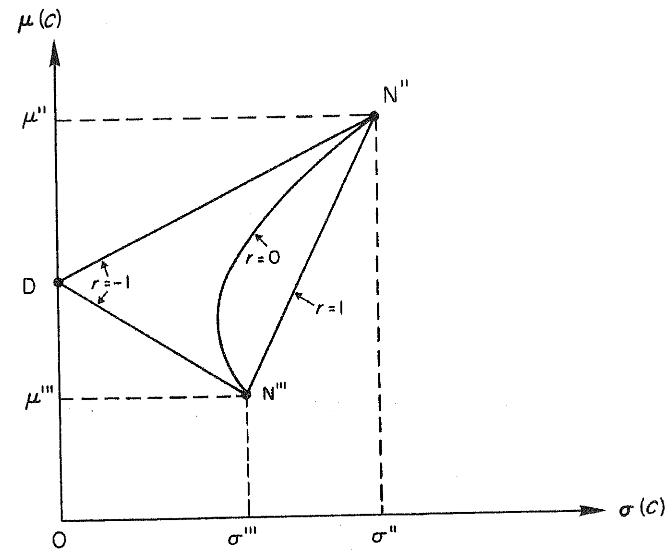

Figure 2.7. Portfolio choice with two risky assets.

the risky earnings fail to cover the fixed obligation of repaying the debt.[16] In the literature of finance theory it is usually assumed that tangencies in the "short sales" region are possible, at least within the range of practical interest. In our diagram, the termination of the dashed extension of N′N″ at point R indicates the limit to which riskless debt can be issued by the investor in order to achieve a larger holding of risky assets.

What about "selling short" the *risky* security instead – i.e., incurring an obligation to meet varying payments over states of the world, while enlarging one's holding of riskless claims? By an analogous argument, this may also be feasible up to a point. It is not difficult to show that such portfolios lead to μ, σ combinations along a line like N′K in Figure 2.6, which clearly cannot be a portion of the efficient opportunity boundary.

Figure 2.7 illustrates the efficient boundary for proper portfolio combinations of two *risky* securities. (We will not consider "short selling" in this discussion.) If asset 2, let us say, has a higher mean return per dollar than asset 3 ($\mu_2/P_2^A > \mu_3/P_3^A$), then for the situation to be interesting once again it must also be that asset 2 involves a greater risk per dollar as well ($\sigma_2/P_2^A > \sigma_3/P_3^A$) In the diagram, points N″ and N‴ represent the

[16] If risky portfolio income were normally distributed, there would always be some non-zero probability of negative returns exceeding *any* preassigned limit. No holder of such a risky distribution could ever issue even the tiniest riskless security obligation, since he could not guarantee to repay a debt with certainty.

single-asset portfolios for securities 2 and 3, respectively, where $\mu'' > \mu'''$ and $\sigma'' > \sigma'''$.

The diagram illustrates how *diversification*, holding a mixture of assets, tends to reduce portfolio standard deviation σ. The power of diversification is a function of the size and sign of the correlation coefficient $\rho_{23} \equiv \sigma_{23}/\sigma_2\sigma_3$ (henceforth, ρ for short) between the asset return distributions $\tilde{z}_2$ and $\tilde{z}_3$.

Consider first the limiting case of *perfect positive correlation* ($\rho = 1$). Here the μ, σ combinations associated with mixtures of assets 2 and 3 would lie along the straight line N′N″ in Figure 2.7, at distances proportionate to the relative budget shares.[17] If instead the yields on the two assets were *uncorrelated* ($\rho = 0$), the attainable μ, σ combinations would fall on a boundary represented by the middle curve connecting N″ and N‴ in the diagram. It is important to notice that in the region of point N‴ the slope $d\mu/d\sigma$ of this curve becomes actually negative, the implication being that the efficient opportunity boundary no longer includes point N‴ itself.[18] Thus, for any portfolio of two uncorrelated risky assets, the single-asset portfolio consisting of the lower-μ, lower-σ (per dollar of cost) security is driven out of the efficient set. More generally, the slope $d\mu/d\sigma$ will be negative if $\sigma''' > \rho\sigma''$.[19] So, for any number of risky assets, if all yields are uncorrelated then only one single-asset portfolio would be located on the efficient boundary, to wit, the portfolio consisting of that asset a^* that offers highest yield per dollar ($\mu_{a^*}/P^A_{a^*} > \mu_a/P^A_a$ for any $a \neq a^*$). But, if asset yields are correlated, any asset with sufficiently high positive correlation with a^* might also be an efficient single-asset portfolio.

Finally, in the limiting case of *perfect negative correlation* ($\rho = -1$), the σ-reducing effect of diversification is so great that the curve N″N‴ breaks

17 If the budget shares are $\kappa = q_2 P^A_2/\bar{W}$ and $1-\kappa = q_3 P^A_3/\bar{W}$ then $\mu = \kappa\mu'' + (1-\kappa)\mu''$, where $\mu'' = \bar{W}\mu_2/P^A_2$ and $\mu''' = \bar{W}\mu_3/P^A_3$ are the mean yields on the respective single-asset portfolios. And σ is:

$$\sigma = [(\kappa\sigma'')^2 + 2\rho\kappa(1-\kappa)\sigma''\sigma''' + ((1-\kappa)\sigma''')^2]^{\frac{1}{2}}$$

If $\rho = 1$, then $\sigma = \kappa\sigma'' + (1-\kappa)\sigma''$. So μ and σ both increase linearly with K, proving the assertion in the text.

18 As κ increases, the slope along any of the curves connecting N″ and N‴ can be written:

$$\frac{d\mu}{d\sigma} = \frac{d\mu/d\kappa}{d\sigma/d\kappa} = \frac{\mu'' - \mu'''}{[\kappa(\sigma'')^2 + (1-2\kappa)\sigma''\rho\sigma''' - (1-\kappa)(\sigma''')^2]/\sigma}$$

If $p \leq 0$, at point N‴ where $\kappa = 0$ the denominator will be negative, hence $d\mu/d\sigma$ will be negative. So there will exist a portfolio with $\kappa > 0$ having lower σ and higher μ than the asset-3 single-asset portfolio.

19 From the preceding footnote, the sign of $d\mu/d\sigma$ when $\kappa = 0$ will be the same as the sign of $\rho\sigma''\sigma''' - (\sigma''')^2$ or $\rho\sigma'' - \sigma'''$.

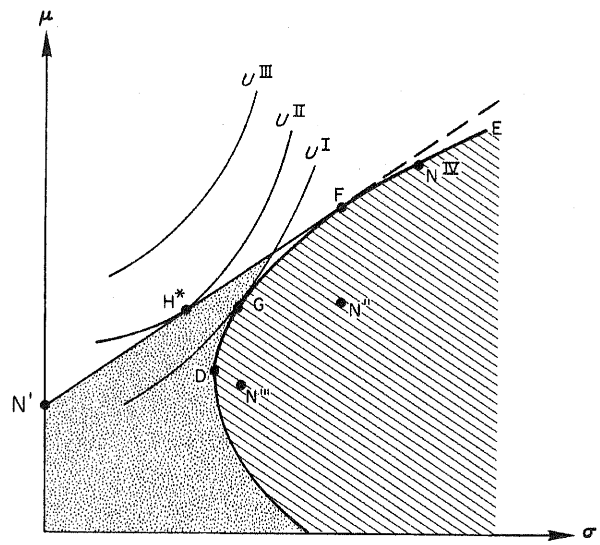

Figure 2.8. Mutual-Fund Theorem.

into two lines meeting at the vertical axis (at point D). Thus, with $\rho = -1$, it is possible to achieve a riskless combination of two risky assets.[20] Of course, here again point N‴ can no longer be in the efficient set.

Generalizing the previous diagram to any number of assets, Figure 2.8 is intended to suggest the nature of the opportunity set and efficient (northwest) boundary. If there are only risky securities available, the opportunity set will be the shaded area whose northwest (efficient) boundary is the curve DE. Here points N″ and N‴ represent single-asset portfolios that are not efficient, whereas point N^{IV} represents an efficient single-asset portfolio. Owing to the power of diversification, almost all of the boundary DE would likely represent multiasset portfolio mixtures.

The introduction of a *riskless* asset, whose single-asset portfolio is represented by point N′, enlarges the opportunity set by the dotted area in the diagram. The efficient boundary now becomes the line from N′ drawn tangent to the DE curve (at point F). In general, as just argued, F would represent a particular mixed portfolio of assets.

In the absence of the riskless asset, the individual's risk-bearing optimum (indifference-curve tangency) would be at point G along the curve DE. But for the opportunity set enlarged by the presence of the riskless asset, the

[20] If $\rho = -1$, then $\sigma = \kappa\sigma'' - (1-\kappa)\sigma'''$. Setting $\kappa = \sigma''/(\sigma'' + \sigma''')$, we have $\sigma = 0$.

optimum is at point H* along the line N′F. As discussed earlier, we also admit the possibility that H* might fall in the dashed extension to the northeast of F along this line, representing an individual who issues riskless obligations in order to hold more than 100% of his endowed wealth in the form of the risky combination F.

The Mutual-Fund Theorem

An important result follows from our previous discussion:

If individuals' preferences are summarized by desire for large μ and small σ, and if there exists a single riskless asset and a number of risky assets, in equilibrium the asset prices will be such that everyone will wish to purchase the *risky* assets in the same proportions.

Thus, if one individual holds risky assets 2 and 3 in the quantities $q_2 = 10$ and $q_3 = 9$, someone who is richer (or less risk averse) might hold larger amounts of each risky asset – but still in the ratio 10:9! This remarkable "Mutual-Fund Theorem" underlies all the main properties of the Capital Asset Pricing Model (CAPM), which constitutes the centerpiece of modern finance theory. (CAPM will be discussed further in Chapter 4.)

Justification of the theorem requires no complicated analytical apparatus. All we need do is to re-interpret the opportunity set of Figure 2.8 in per-dollar-of-wealth terms. Thus, for any individual the vertical axis would now be scaled in units of $\mu/\bar{W}$ and the horizontal axis in units of $\sigma/\bar{W}$ Since both axes are being divided by the same constant, the opportunity set would change only by a scale factor. For example, the single-asset portfolio corresponding to any asset a that formerly was represented by the vector (μ^a, σ^a), where $\mu^a = (\bar{W}/P_a^A)\mu_a$ and $\sigma^a = (\bar{W}/P_a^A)\sigma_a$, would now have the coordinates $\mu^a/\bar{W} = \mu_a/P_a^A$ and $\sigma^a/\bar{W} = \sigma_a/P_a^A$. And, in particular, points N′ and F would similarly maintain their positions, so that the efficient boundary N′F would have the same slope as before.

The significance of this conversion to per-dollar dimensions is that, in these per-dollar units, *every individual in the economy, regardless of wealth, faces exactly the same opportunities*! If asset a offers a mean yield per dollar μ_a/P_a^A to one individual, it offers the same per-dollar mean yield to everyone. And similarly for the standard deviation per dollar σ_a/P_a^A of asset a, and for all combinations of assets as well. Thus, every individual will hold, in whatever fraction of his wealth is devoted to risky securities, the same proportionate mixture of assets represented by point F in the diagram. So, we can say, a "mutual fund" of risky securities set up to meet the needs of any single investor will meet the needs of all.[21] What still does vary among individuals is the *fraction of wealth held in riskless versus risky form*, a decision that will depend upon individuals' varying personal preferences as to risk bearing (the shapes of their indifference curves in the original diagram of Figure 2.8).

In the economy as a whole, there is exactly one unit of this mutual fund F, corresponding to the economy-wide amounts $q_2^F, \ldots, q_A^F$ of the risky assets $a = 2, \ldots, A$. (So the typical individual will be holding a fractional unit of F.) Then the price P_F^A of a unit of the fund is:

$$P_F^A = \sum_{a=2}^{A} q_a^F P_a^A$$

Writing the mean return and standard deviation of return for a unit of portfolio F as μ_F and σ_F, we can obtain an expression for $\Theta \equiv d\mu/d\sigma$ – the price of risk reduction – in terms of the slope of the line N′F in Figure 2.8. For any individual, point N′ (the riskless single-asset portfolio) has μ-coordinate $\mu' = (\bar{W}/P_1^A)\mu_1 = \bar{W}(1+R_1)$ and μ-coordinate $\sigma' = 0$. Point F (the portfolio held entirely in the mutual fund) has μ-coordinate $\mu^F = (\bar{W}/P_F^A)\mu_F$ and σ-coordinate $\sigma^F = (\bar{W}/P_F^A)\sigma_F$. So the steepness of the line is:

$$\frac{d\mu}{d\sigma} = \frac{\mu^F - \mu'}{\sigma^F - \sigma'} = \frac{(\bar{W}/P_F^A)\,\mu_F - \bar{W}(1+R_1)}{(\bar{W}/P_F^A)\,\sigma_F} = \frac{\mu_F/P_F^A - (1+R_1)}{\sigma_F/P_F^A}$$

Note that, consistent with our previous discussion, the individual wealth parameter $\bar{W}$ has cancelled out. Thus, a corollary of the Mutual-Fund Theorem is that *the price of risk reduction is the same for every individual.*

Exercises and Excursions 2.2

1 Ranking of Alternative Wealth Prospects

An individual with utility function $v(c) = c^{0.5}$ has an initial wealth of zero. He must choose one of two jobs. In the first there is an equal probability of earning 1 or 3. In the second there is a probability of 1/9 that he will earn zero, a probability of 7/9 that he will earn 2, and a probability of 1/9 that he will earn 4.

[21] If, however, individuals differed in their personal estimates of the asset characteristics μ_a and σ_a, their perceived opportunity sets would not be identical in per-dollar units and the Mutual-Fund Theorem would not be valid.

(A) Show that both jobs have the same mean income μ, that the standard deviation of income σ is lower in the second job, and that despite this the individual will choose the first job.

(B) Can you explain this result in terms of preferences for skewness? If not, what is the explanation?

2 *Constant Absolute Risk Aversion and Normally Distributed Asset Returns*

(A) Show that:

$$Ac + \frac{1}{2}\left(\frac{c-\mu}{\sigma}\right)^2 = \frac{1}{2}\left(\frac{c-(\mu - A\sigma^2)}{\sigma}\right)^2 + A\left(\mu - \frac{1}{2}A\sigma^2\right)$$

where $A \equiv -v''(c)/v'(c)$ is known as the measure of "absolute risk aversion." Hence, or otherwise, show that if c is distributed normally with mean μ and variance σ^2 and if $v(c) = -\,e^{-Ac}$ then:

$$\mathrm{E}\,[v\,(c)] = -\int_{-\infty}^{\infty} e^{-Ac}\frac{1}{(2\pi)^{\frac{1}{2}}}\exp\left\{-\frac{1}{2}\left(\frac{c-\mu}{\sigma}\right)^2\right\} dc = -e^{-A(\mu-\frac{1}{2}A\sigma^2)}$$

(B) Under the above assumption, it follows that preference can be represented by the indirect utility function:

$$U(\mu, \sigma) = \mu - \frac{1}{2}A\sigma^2$$

Suppose an individual with such preferences must choose between a riskless asset and a normally distributed risky asset. Show that the amount of the risky asset purchased is independent of initial wealth and decreasing in the degree of absolute risk aversion A. What happens if the indicated expenditure on the risky asset exceeds the individual's endowed wealth $\bar{W}$?

3 *The μ, σ Opportunity Locus*

An individual spends a fraction κ of his wealth $\bar{W}$ on asset a and the remainder on asset b. Each asset has a price of unity and the yields $(\tilde{z}_a, \tilde{z}_b)$ have means μ_a and μ_b and covariance matrix $[\sigma_{ab}]$.

(A) Obtain expressions for the mean μ and standard deviation σ of the portfolio as functions of κ.

(B) Hence show that the standard deviation can be expressed as:

$$\sigma = \{[(\mu - \bar{W}\mu_a)^2\sigma_{aa} - 2(\mu - \bar{W}\mu_a)(\mu - \bar{W}\mu_b)\sigma_{ab} + (\mu - \bar{W}\mu_b)^2\sigma_{bb}]/(\mu_a - \mu_b)^2\}^{\frac{1}{2}}$$

(C) Suppose $\mu_a > \mu_b$ and $\sigma_{aa} > \sigma_{bb} > 0 \geq \sigma_{ab}$. Obtain an expression for the rate of change of σ with respect to μ. Hence establish that if the individual begins with all his wealth in asset b, he can increase the mean yield and simultaneously reduce the standard deviation by trading some of asset b for asset a. Illustrate the locus of feasible μ, σ combinations in a diagram.

(D) Assuming the individual's utility is a function only of the mean and standard deviation, can you draw any conclusions as to the composition of his optimal portfolio?

4 Investor's Portfolio Optimum in a μ, σ Model

In a competitive economy there are I investors, all having the same utility function $U = \mu^{10}e^{-\sigma}$. Each individual is endowed with exactly one unit each of assets 1, 2, and 3 with payoff statistics as shown in the table below, all the payoff distributions being uncorrelated ($\sigma_{ab} = 0$, for all $a \neq b$). Given asset prices are also shown:

	μ_a	σ_a	P_a^A
Asset 1	1	0	1.0
Asset 2	1	3	0.46
Asset 3	1	4	0.04

(A) Sketch the indifference curves on $\mu(c)$, $\sigma(c)$ axes. Locate, for any single individual, the three single-asset portfolios he might hold.

(B) Under the assumptions here, each individual's optimum portfolio H* must evidently be the same as his endowed portfolio. Locate this portfolio, and also the mutual fund portfolio *F*. What fraction of his wealth does the individual hold in the mutual fund?

(C) Verify that the price of risk reduction is $\Theta = \frac{3}{10}$. What is the equation of the individual's budget line? What is his Marginal Rate of Substitution (the slope of the indifference curve) at H*?

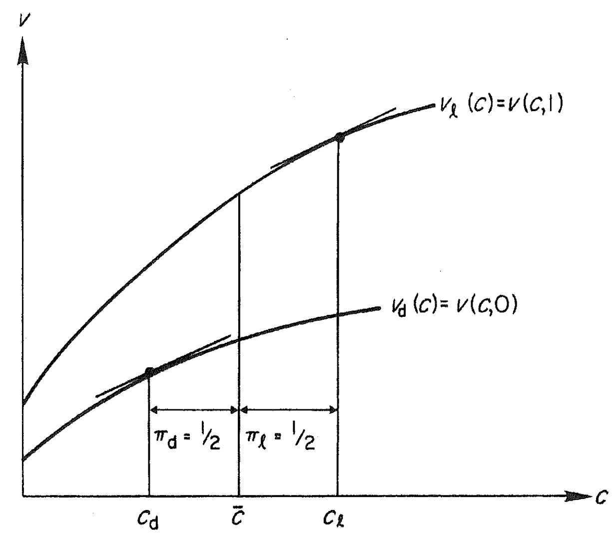

Figure 2.9. State-dependent utility – complementary preferences.

2.3 State-Dependent Utility

In risk-bearing decisions it sometimes appears that the individual's utility function $v(c)$ might itself depend upon the state of the world, contrary to our assumption in Section 1.4 in Chapter 1. To take an extreme case, if the states under consideration were "being alive" versus "being dead", a typical individual would likely value rights to income in the former state more heavily! (Yet he might attach *some* "bequest utility" to income contingent upon his own death.) Similar considerations could apply if states of the world were defined in terms of one's sickness versus health, or life versus death of one's child, or success versus failure at love, or retention versus loss of a unique heirloom.[22]

For concreteness, think in terms of two states $s = \ell$, d corresponding to life versus death of one's child. Then we can imagine a *pair* of utility functions for consumption income, $v_\ell(c)$and $v_d(c)$ as in Figure 2.9. For any given amount of income c, the former curve would definitely be the higher (an *ordinal* comparison). But the question is whether our analysis in Section 1.4 can be extended so that, despite what appear to be two distinct $v(c)$ functions, a single cardinal scaling can be arrived at permitting use of the expected-utility rule.

[22] Our discussion here follows Cook and Graham (1977) and Marshall (1984).

This seemingly difficult problem resolves itself immediately, once it is realized that we are still really dealing with a *single* underlying utility function v. The only change is that v is now to be regarded as a function of two "goods": $v \equiv v(c, h)$. The first argument c still represents amounts of the consumption commodity, just as before. The second argument h represents the amount of the state-determining or "heirloom" good; in our example, if the child lives (state ℓ) then $h = 1$, if she dies (state d) then $h = 0$. The curves $v_l(c)$ and $v_{\mathrm{d}}(c)$ can therefore be more explicitly labeled $v(c, 1)$ and $v(c, 0)$; these two curves are not two separate utility functions for the parent, but two sections of his single overall $v(c, h)$ function. We already know, of course, that there is no difficulty deriving a cardinal utility function, where v is a function of two or more goods (see Exercises and Excursions 1.5.6).

We now turn to the risk-bearing decision under state-dependent utility. Suppose a risk-averse person is endowed with a given quantity $\bar{c}$ of income certain, but faces a gamble involving the heirloom commodity – his child might live or might die. Is it rational to insure one's child's life, at actuarial ("fair") odds? (Doing so means that the parent will end up with higher income if the child dies.) Or, should the parent do the opposite and buy an annuity upon his child's life, a contractual arrangement that provides more income so long as the child lives, but less if the child dies? Here is a less agitating example: if our college team is playing its traditional rivals in a crucial match, is it rational for us as loyal yet risk-averse fans to bet at fair odds *against* our team (equivalent to insuring our child's life), or to bet the other way (buy the annuity instead)?

Since, as argued above, there is no difficulty in developing a cardinal $v(c, h)$ scale for use with the expected-utility rule, then (given Complete Contingent Markets [CCM]) the Fundamental Theorem of Risk Bearing continues to apply. In terms of the states $s = \ell$ and $s = \mathrm{d}$, we can rewrite (2.1.6) in the form:

$$\frac{\pi_\ell v'_\ell(c_\ell)}{P_\ell} = \frac{\pi_{\mathrm{d}} v'_{\mathrm{d}}(c_{\mathrm{d}})}{P_{\mathrm{d}}} \tag{2.3.1}$$

Here $v'_\ell(c_\ell)$, the state-ℓ marginal utility – which could also be written $\partial v(c, 1)/\partial c$ – corresponds to the slope along the upper $v(c)$ curve in Figure 2.9 and is a *partial* derivative of the underlying $v(c, h)$ function. Similarly, $v'_{\mathrm{d}}(c_{\mathrm{d}}) \equiv \partial v(c, 0)/\partial c$ corresponds to the slope along the lower $v(c)$ curve in the diagram.

Suppose that the parent is offered a contract of insurance or annuity on his child's life, at fair odds in either case. Imagine that his optimizing choice is to accept neither, but remain at his endowed certainty-income position.

Then, in Equation (2.3.1) it must be that $v'_\ell(c_\ell) = v'_\mathrm{d}(c_\mathrm{d})$ when $\bar{c}_\ell = \bar{c}_\mathrm{d} = \bar{c}$. Geometrically, for this to occur the two curves in Figure 2.9 would have to be vertically parallel at the endowment level of income. In economic terms we would say that the two goods are "independent in preference": varying the amount of h by having the child live or die does not affect the *marginal* utility of the c commodity (although it certainly affects the parent's *total* utility).

Suppose instead that the two goods are "complements in preference": an increase in h raises $v'(c)$. Then at the endowment level of income the upper curve has steeper slope: $v'_\ell(\bar{c}) > v'_\mathrm{d}(\bar{c})$; this is the situation pictured in Figure 2.9. Because the slope along either curve diminishes as income increases (diminishing marginal utility, reflecting the risk-aversion property), it follows that the optimality Equation (2.3.1) can be satisfied only by having $c_\ell > c_\mathrm{d}$. (And specifically, if the probabilities are $\pi_\ell = \pi_\mathrm{d} = \frac{1}{2}$, as in the diagram, then $c_\ell - \bar{c} = \bar{c} - c_\mathrm{d}$.) To achieve this position, the parent would purchase an annuity on his child's life. The economic interpretation is this: if having your child alive raises the marginal utility of income to you (perhaps because you mainly desire income only in order to meet her needs), then at fair odds you would not insure your child's life but would buy the annuity to generate more income while she lives.

Finally, if c and h are "substitutes in preference," i.e., if $v'_\ell(c) < v'_\mathrm{d}(c)$, you would insure your child's life. This might correspond to a situation where your child, if she lives, would support you in your old age; not having that source of future support raises your marginal utility of income, since should the child die you will have to provide for your declining years yourself.

What about a parent insuring his *own* life on behalf of a child? Here the child's survival is not in question, so now we must let the states $s = \ell$ versus $s = \mathrm{d}$ refer to the *parent's* life as the "heirloom" commodity. In such a situation the upper curve $v_\ell(c) \equiv v(c, 1)$ in Figure 2.9 pictures the parent's "living utility" of consumption income while the lower curve $v_\mathrm{d}(c) \equiv v(c, 0)$ shows the "bequest utility" he attaches to the child's income after his own death. There appears to be a puzzle here. It is reasonable to assume that h and c are complements in preference – in state ℓ there are two persons, parent and child, who need income for purpose of consumption while in state d only the latter requires income. But whereas in our previous discussion complementary preferences led to purchase of an annuity rather than insuring, we know that parents do typically insure their lives on behalf of their children.

The puzzle is resolved when we realize that death of a parent will typically mean not only loss of the heirloom commodity (life), but *also* loss of

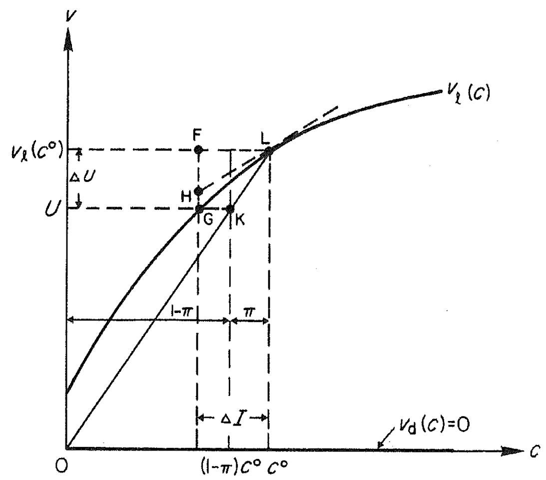

Figure 2.10. Value of a life.

income earnings that the parent would have generated. Consequently, the typical endowment position is not a life risk plus a quantity of income certain ($\bar{c}_\ell = \bar{c}_d$), but rather involves a life risk plus a correlated income risk ($\bar{c}_\ell > \bar{c}_d$). In buying life insurance a risk-averse parent is purchasing an *offsetting gamble*, tending to reduce the overall riskiness of a situation in which both *c* and *h* take on low values should the state d occur.

2.3.1 An Application: The "Value of Life"

Figure 2.10 illustrates the situation of an individual choosing between higher income and higher survival probability, for example, working at a risky high-income profession versus a lower-paying but safer occupation.

As an analytical simplification, suppose that the individual has no dependents, so that his "bequest utility" function $v_d(c)$ can be assumed to be everywhere zero. Thus, the lower curve of our previous diagram now runs along the horizontal axis.[23] Let us suppose that he finds himself initially in the risky situation, with income c^o if he survives. (His non-survival income, if any, is irrelevant since $v_d(c) = 0$ everywhere.) Denote the death probability

[23] The $v_\ell(c)$ curve is shown as intersecting the vertical axis, which represents an arguable assumption that life (even at zero income) is preferable to death. However, only the local shape of $v_\ell(c)$ in the neighborhood of the endowment income plays any role in the analysis, so there is no need to insist upon this assumption.

as π, so the probability of survival is $1 - \pi$. Then the expected utility U of this risky prospect is:

$$U = (1 - \pi)v_\ell(c^\circ)$$

The expected utility of the endowed gamble is indicated in the diagram by the vertical height of point K, which is the probability-weighted average of point L where $(c, h) = (c^\circ, 1)$ and the origin 0 where $(c, h) = (0, 0)$. (Here as before, life is the "heirloom" commodity h.)

The *utility loss* due to the existence of the hazard (that is, in comparison with a situation with the same c° but where the death probability is $\pi = 0$) can be expressed as:

$$\Delta U = v_\ell(c^\circ) - U = \pi v_\ell(c^\circ)$$

This corresponds to the vertical distance between $v_\ell(c^\circ)$ and U as marked off on the vertical axis of the diagram – the vertical distance between points F and G. Along the v_ℓ function, point G has the same utility as the endowed gamble (point K). So ΔI, the income equivalent of the utility loss, is the horizontal difference between points G and L. ΔI may be termed the income-compensation differential: the remuneration reduction (viewed from point L) that this individual would be just willing to accept to be free of his endowed death risk.

We now seek a more general analytical expression for ΔI. The line tangent to the $v_\ell(c)$ curve at L has slope $v'_\ell(c^\circ)$. Then the vertical distance between points F and H equals $\Delta I v'_\ell(c^\circ)$. As long as the death probability is fairly small, this distance is approximately equal to the vertical distance between the points F and G, which we have already seen is the utility loss $\pi v_\ell(c^\circ)$. So, to a first approximation:

$$\Delta I v'_\ell(c^\circ) = \pi v_\ell(c^\circ)$$

or:

$$\Delta I = \frac{\pi v_\ell(c^\circ)}{v'_\ell(c^\circ)} = \frac{\pi c^\circ}{e}$$

where:

$$e = \frac{dU/dI}{U/I} = \frac{v'_\ell(c^\circ)}{v_\ell(c^\circ)/c^\circ}$$

Here e signifies the income elasticity of utility, evaluated at the income level c°. Since π is the probability of loss, the other factor $v_{\ell}(c^{\circ})/v'_{\ell}(c^{\circ}) = c^{\circ}/e$ then represents the *value of life*[24] implied by this analysis.

We must be careful not to misinterpret this value, however. It does not represent the amount that an individual would pay to "buy his life," for example, if he were being held for ransom. It represents the exchange rate at which he would be willing to give up a *small* amount of income for a *small* reduction in the probability of death π (when π is close to zero). This can be shown more explicitly as follows. Since $U = (1 - \pi)v_l(c^{\circ})$, the Marginal Rate of Substitution $M(c, \pi)$ between income c and death probability π is:

$$M(c, \pi) \equiv \left.\frac{dc}{d\pi}\right|_U \equiv \frac{-\partial U/\partial \pi}{\partial U/\partial c} = \frac{v_{\ell}(c^{\circ})}{(1-\pi)\, v'_{\ell}(c^{\circ})}$$

For π close to zero the denominator is approximately $v'_{\ell}(c^{\circ})$. Thus $v_{\ell}(c^{\circ})/v'_{\ell}(c^{\circ}) = c^{\circ}/e$ does not represent the purchase price of a whole life, but the *Marginal Rate of Substitution* between small increments of income and survival probability.

Nevertheless, the interpretation in terms of "value of life" is not wholly unwarranted when we think in terms of society as a whole. Suppose that each member of a large population voluntarily accepts a 0.001 decrease in survival probability in order to earn $500 more income, implying a figure of $500,000 for the "value of life." Again, this does not mean that any single individual would trade his whole life for $500,000; indeed, there might be no one willing to make such a trade, for any amount of income whatsoever. But if everyone in a population of 1,000,000 accepts such a small per-capita hazard, there will be about $500,000,000 more of income and about 1,000 additional deaths. So, in a sense, $500,000 is indeed the "value of a life"!

Exercises and Excursions 2.3

1 Betting for or against the Home Team?

Your endowed income is $\bar{c} = 100$. There is a 50:50 chance that the home team will win the big game. You can bet at fair odds, picking either side to win and for any amount of money. Each of the utility functions (i) through

[24] Of course, since death is ultimately certain, any increased chance of life can only be temporary. If we are dealing with *annual* death probability, we should really speak of the value of an *incremental year of life expectation*. But the dramatic, if misleading, term "value of life" is too firmly established to be displaced.

(iv) below consists of a *pair* of utility functions, which differ depending upon whether the home team wins (W) or loses (L):

(i) $v_W(c) = 2c^{0.5}$ and $v_L(c) = c^{0.5}$
(ii) $v_W(c) = 2 - \exp(-c)$ and $v_L(c) = 1 - \exp(-c)$
(iii) $v_W(c) = 1 - \exp(-2c)$ and $v_L(c) = 1 - \exp(-c)$
(iv) $v_W(c) = \ln(50 + c)$ and $v_L(c) = \ln(c)$

For each utility function:

(A) Verify that, at any level of income *c*, you prefer the home team to win.
(B) Find the optimal *b*, the amount of money bet on the home team (so that *b* is negative if you bet *against* the home team).
(C) Having made the optimal bet, do you still want the home team to win? Explain the differences among the four cases.

2 *Risk Preference under State-Dependent Utility?*
An individual can choose between two suburbs in which to live. The homes in the first suburb are small, while in the second they are large. Utility in the first suburb is:

$$v(c, h_1) = 8c^{\frac{1}{2}}$$

where *c* is spending on goods other than housing (i.e., on "corn," whose price is unity). Utility in the second suburb is:

$$v(c, h_2) = 5c^{\frac{2}{3}}$$

Housing in the first suburb costs \$20 per year and in the second costs \$56.

(A) Sketch the two utility functions. Verify that the utility functions cross at $\bar{c} = 120$, and explain what this signifies.
(B) Suppose that before having invested in housing, the individual's endowed income was \$120. Consider the gamble corresponding to the prospect (181, 56; 0.5, 0.5). Note that this is an adverse gamble; in comparison with the endowment income, the payoff is \$61 and the possible loss is \$64. Compute the individual's utility for each outcome and hence confirm that taking the gamble does raise expected utility.
(C) Indicate, geometrically, the optimal gamble for this individual. Explain why the individual wants to undertake such a gamble.

[HINT: Are *c* and *h* complements here?]

(D) Can this kind of argument explain why some people gamble regularly?

3 "Superstars" and the Value of Life

An individual with endowed income $\bar{c}$ has a concave utility function $v(c)$. He has contracted a disease which, if not treated, will be fatal with probability $1 - p_o$ and will spontaneously cure itself with probability p_o. His "bequest utility" in the event of death is zero everywhere.

(A) Suppose that, when treated by a physician who charges z, his probability of survival rises to p. If z is his maximum willingness to pay for that treatment, show that:

$$pv(\bar{c} - z) = p_o v(\bar{c})$$

(B) Hence show that:

$$\frac{dp}{dz} = \frac{p_o v(\bar{c}) v'(\bar{c} - z)}{v(\bar{c} - z)^2}$$

Depict the relationship between p and z in a figure. Interpret its shape.

(C) Suppose a "superstar" physician can increase the individual's probability of survival by $1 - p_o$, so that he is sure to live, while another physician can increase this probability by only $(1 - p_o)/2$. Indicate in the figure the maximum amounts X and Y that the two physicians could charge.

(D) It has been asserted that "superstars" tend to receive disproportionally high incomes. In this context, this means that the ratio of physicians' fees would exceed the ratio of the survival rate increments that they provide. Assuming that both physicians can charge the maximum, is the assertion $X/Y > 2$ valid here?

(E) The maximum that the individual would be willing to pay a physician, who (in effect) provides him with "a fraction $p - p_o$ of his life," is $z(p)$ where this function is defined implicitly in part (A). Then it can be argued that his implied valuation of his own life is $z(p)/(p - p_o)$. For example, if $p_o = p - p_o = 0.5$ and $X = \$100{,}000$, the value of his life would be \$200,000. Carry this argument to the limit in which p is small and show that, at this limit, the value he places upon his life is $v(\bar{c})/p^{o} v'(\bar{c})$. Compare this conclusion with that reached at the end of the discussion in the text.

SUGGESTIONS FOR FURTHER READING: The state-contingent claims model was introduced by Arrow (1953, reprinted 1964). See also Debreu (1959). A classical treatment of asset-pricing and the CAPM model is provided in Ingersoll (1987). Some recent work in finance extends beyond the traditional μ, σ preferences of Section 2.2; see Leland (1999). An important issue in state-dependent utility is that the subjective probability distribution over states is not unique. For more on this, see Karni (1993).

References

Arrow, Kenneth J., "The Role of Securities in the Optimal Allocation of Risk Bearing," *Econometrie* (1953), as translated and reprinted in 1964, *Review of Economic Studies*, 31 (1964), 91–96.

Cook, Philip J. and Graham, Daniel A., "The Demand for Insurance and Protection: The Case of Irreplaceable Commodities," *Quarterly Journal of Economics*, 91 (1977), 143–156.

Copeland, Thomas E., Weston, J. Fred, and Shastri, Kuldeep, *Financial Theory and Corporate Policy*, 4th edition, Reading, MA: Addison-Wesley, 2004.

Debreu, Gerard, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, New Haven: Yale University Press, 1959.

Ingersoll, Jonathan E., *The Theory of Financial Decision Making*, New York: Rowman and Littlefield, 1987.

Karni, Edi, "Subjective Expected Utility Theory with State-Dependent Preferences," *Journal of Economic Theory*, 60 (1993), 428–438.

Leland, Hayne E., "Beyond Mean-Variance: Performance Measurement in a Nonsymmetrical World," *Financial Analysts Journal*, 55 (1999), 27–36.

Marshall, John M., "Gambles and the Shadow Price of Death," *American Economic Review*, 74 (1984), 73–86.

3

Comparative Statics of the Risk-Bearing Optimum

The elements of the decision problem under uncertainty – the individual's preferences, opportunities, and beliefs – were surveyed in Chapter 1. We distinguished between *terminal* choices, actions undertaken on the basis of given probability beliefs (covered in Part I of this volume), and *informational* choices, actions designed to improve one's knowledge of the world before a terminal decision has to be made (to be covered in Part II). Chapter 2 analyzed the individual's *risk-bearing optimum*, the best terminal action to take in the face of uncertainty.

We now want to explore how these optimizing decisions change in response to variations in the person's character or situation (his or her wealth, tastes for risk, the endowment of goods, the market prices faced, and so forth). Modeling the before-and-after effects of such "parametric" changes, without attending to the dynamics of the actual transition path from one solution to another, is called the *method of comparative statics*. This chapter is devoted to the comparative statics of the individual's risk-bearing optimum.

3.1 Measures of Risk Aversion

The individual's risk-bearing optimum depends critically upon his attitudes toward risk. And, since parametric changes generally involve positive or negative wealth effects, it will often be crucial to take into account how attitudes toward risk vary as a function of wealth.

As discussed in Chapter 2, in a regime of Complete Contingent Markets with two states of the world, the individual's wealth constraint is:

$$P_1 c_1 + P_2 c_2 = P_1 \bar{c}_1 + P_2 \bar{c}_2 \equiv \bar{W} \qquad (3.1.1)$$

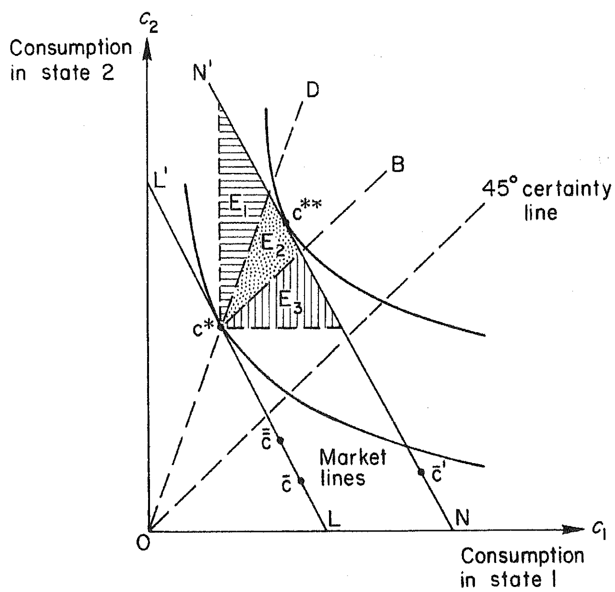

Figure 3.1. Wealth effects.

Endowed wealth, $\bar{W}$, represents the market value of the endowment vector $\bar{C} = (\bar{c}_1, \bar{c}_2)$. As shown in Figure 3.1, expected utility is maximized at the indifference-curve tangency C* along the original budget line LL′ (assuming an interior solution). As endowed wealth increases, the optimum position moves outward from C* along some *wealth expansion path* like C*B or C*D in the diagram.

Suppose that, after an increase in endowed wealth, the individual's new optimum lies northeast of the old optimum C* but below the line C*B (i.e., in region E_3). Since C*B is drawn parallel to the 45° line, all the points in E_3 lie closer to the 45° line than does C*. Thus, an individual whose wealth expansion path lies in this region reduces his *absolute consumption risk* (gap between c_1 and c_2) as his wealth increases. If instead (as shown in the diagram) his new optimum lies above the line C*B (in regions E_1 or E_2) his "tolerance" for absolute risk must be increasing with wealth. A solution along the dividing line C*B would represent constant tolerance for absolute risk. Or, putting it the other way, we can speak of increasing, decreasing, or constant *absolute risk aversion* as wealth increases.

These alternative responses to changes in wealth imply restrictions upon the shape of the individual's utility function $v(c)$. We can use these restrictions to construct a measure of the individual's absolute risk aversion. Consider a change in the absolute steepness of the indifference curve – the Marginal Rate of Substitution $M(c_1, c_2)$ – in moving from some arbitrary point (c_1, c_2) to a nearby point $(c_1 + dc_1, c_2 + dc_2)$. From Equation (2.1.2) the Marginal Rate of Substitution can be expressed as:

$$M(c_1, c_2) = \frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} \tag{3.1.2}$$

Taking the logarithm of both sides:

$$\ln M = \ln \pi_1 + \ln v'(c_1) - \ln \pi_2 - \ln v'(c_2)$$

The total differential of the expression is then:

$$d \ln M = \frac{dM}{M} = \frac{v''(c_1)}{v'(c_1)} dc_1 - \frac{v''(c_2)}{v'(c_2)} dc_2 \tag{3.1.3}$$

If c_1 and c_2 increase by the same absolute amount ($dc_1 = dc_2 = dx$), the (c_1, c_2) vector moves outward parallel to the 45° line. Then the proportionate change in $M(c_1, c_2)$ is:

$$\frac{dM}{M} = \left[\frac{v''(c_1)}{v'(c_1)} - \frac{v''(c_2)}{v'(c_2)}\right] dx \tag{3.1.4}$$

Suppose that, as depicted in Figure 3.1, an increase in wealth leads to a new optimum C** that lies in region E_2 and therefore is further from the 45° line. Since the Marginal Rate of Substitution is the same at C* and C**, it follows that $M(c_1, c_2)$ is necessarily lower where the new budget line NN′ intersects the line C*B parallel to the 45° line. If this holds everywhere, then, from (3.1.4):

$$c_1 < c_2 \Rightarrow \frac{dM}{M} = \left[\frac{v''(c_1)}{v'(c_1)} - \frac{v''(c_2)}{v'(c_2)}\right] dx < 0$$

Rearranging we obtain:

$$c_1 < c_2 \Rightarrow \frac{-v''(c_1)}{v'(c_1)} > \frac{-v''(c_2)}{v'(c_2)}$$

Of course, the condition is reversed below the 45° line where $c_1 > c_2$.

Thus, an individual displays decreasing aversion to absolute wealth risks if and only if $A(c)$ is a decreasing function, where:

$$A(c) \equiv \frac{-v''(c)}{v'(c)} \tag{3.1.5}$$

The function $A(c)$, which is evidently a property of the individual's utility function $v(c)$, measures the individual's *absolute risk aversion.*

Since the individual depicted in Figure 3.1 moves farther from the 45° line as his wealth grows, he exhibits Decreasing Absolute Risk Aversion (DARA). If instead he had Constant Absolute Risk Aversion (CARA), $A(c)$ would remain unchanged as wealth rises. In that case the wealth expansion path would be parallel to the 45° line (line C*B). Finally, if $A(c)$ rises with wealth, the individual exhibits Increasing Absolute Risk Aversion (IARA); the wealth expansion path would then converge toward the 45° line, and the new optimum would lie in region E_3.

A second useful measure of attitude toward risk is obtained by considering *proportional* rather than absolute changes in an individual's consumption levels. If c_1 and c_2 increase proportionately, the consumption vector moves outward along a ray out of the origin. Along such a ray:

$$c_2 = kc_1 \quad \text{and} \quad dc_2 = kdc_1$$

Eliminating k we obtain:

$$\frac{dc_2}{c_2} = \frac{dc_1}{c_1}$$

Rearranging terms in (3.1.3) and substituting, we can see that along the ray C*D in Figure 3.1 the Marginal Rate of Substitution changes in accordance with:

$$\frac{dM}{M} = \left[\frac{c_1 v''(c_1)}{v'(c_1)} - \frac{c_2 v''(c_2)}{v'(c_2)}\right]\frac{dc_1}{c_1} \tag{3.1.6}$$

As depicted, C**, the optimum at the higher wealth level, lies in region E_2 and so is closer to the 45° line than the ray C*D. To be the same at C* and C**, the absolute indifference-curve slope $M(c_1, c_2)$ must be increasing along the ray C*D. That is, from (3.1.6):

$$c_1 < c_2 \Rightarrow \frac{dM}{M} = \left[\frac{c_1 v''(c_1)}{v'(c_1)} - \frac{c_2 v''(c_2)}{v'(c_2)}\right]\frac{dc_1}{c_1} > 0$$

Rearranging, we obtain:

$$c_1 < c_2 \Rightarrow \frac{-c_1 v''(c_1)}{v'(c_1)} < \frac{-c_2 v''(c_2)}{v'(c_2)}$$

Let us define as measure of *relative risk aversion*:

$$R(c) \equiv \frac{-cv''(c)}{v'(c)} \tag{3.1.7}$$

Then an individual like the one depicted, who displays *increasing* aversion to proportional (or relative) risk as wealth grows, must be characterized by an $R(c)$ that is an increasing function. We say he displays Increasing Relative Risk Aversion (IRRA).

If the wealth expansion path is a ray out of the origin, so that the individual prefers to accept risks that are exactly proportionally larger as his wealth rises, his $R(c)$ is constant – he displays Constant Relative Risk Aversion (CRRA). Finally, if his tolerance for risk rises more than proportionally with wealth, then $R(c)$ declines with c and the individual is said to exhibit decreasing relative risk aversion (DRRA).

Both $A(c)$ and $R(c)$ are *local* measures. That is, they are defined in terms of small changes in wealth and consumption. In general, there is no reason why an individual should not exhibit increasing absolute or relative risk aversion over some consumption levels and decreasing risk aversion over others. However, in theoretical investigations it is common to make *global* assumptions about both measures of risk aversion.

Pratt (1964) has argued, as an empirical generalization, that individuals will be willing to bear greater *absolute* risk as wealth rises. This is very plausible. Every unit purchased of a risky asset buys a given absolute risk. Assuming that the state-1 yield of asset a is greater than its state-2 yield, the absolute risk per unit of a held is $Z_{a1} - Z_{a2}$. A rich individual, other things equal, should be willing to hold more of every kind of asset; acquiring more units of risky assets, he would inevitably accumulate a larger absolute consumption risk. Empirically less obvious is the contention by Arrow (1965) that individuals, as they become richer, will buy relatively more safety so as to reduce their *proportionate* risk. If both arguments are accepted, the typical wealth expansion path will lie in the region E_2 as depicted in Figure 3.1, rather than in E_1 or E_3. Individuals will be characterized by DARA and IRRA.

We now compare the preference maps of two individuals, one of whom is everywhere more risk averse. In Figure 3.2 the solid curve is an indifference curve for individual K. As depicted, this individual is indifferent between a certainty endowment point $\bar{C} = (\bar{c},\ \bar{c})$ along the 45° line and the gamble B. Moving away from the certainty line to the northwest, i.e., to consumption vectors with $c_2 > c_1$, his Marginal Rate of Substitution (steepness of the indifference curve) must evidently increase. Since the Marginal Rate of

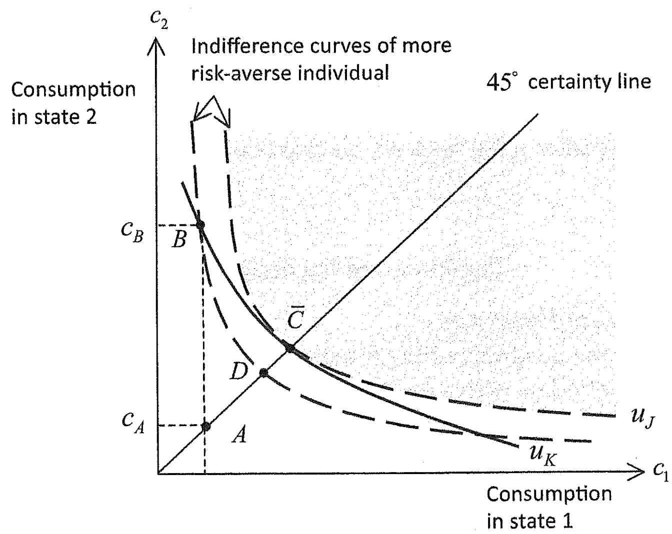

Figure 3.2. Acceptable gambles.

Substitution $M(c_1, c_2)$ is the same everywhere along the certainty line, the change $M(c_1, c_2)$ in moving from $\bar{C}$ to B can be expressed logarithmically as:

$$\begin{aligned}\ln M_B - \ln M_{\bar{C}} &= \ln M_B - \ln M_D \\ &= \int_{c_D}^{c_B} -\frac{v''(c_2)}{v'(c_2)} dc_2 \\ &= \int_{c_D}^{c_B} A_K(c_2) dc_2\end{aligned}$$

The proportional decline in $M(c_1, c_2)$ around the indifference curve therefore varies with the degree of absolute risk aversion. It follows immediately that if individual J has an everywhere greater degree of absolute risk aversion than individual K, so that $A_J(c) > A_K(c)$, the proportional decline in J's $M(c_1, c_2)$ is greater. Yet, if both have the same beliefs, we know from Section 2.1 and Equation (2.1.2) that along the 45° line the Marginal Rates of Substitution are equal to one another. Therefore, at any point B above the 45° line, the more risk-averse individual has a steeper indifference curve. Exactly the same logic reveals that below the 45° line the more risk-averse individual has a flatter indifference curve.

It follows that individual J's indifference curves must be "curvier" than K's (compare the dotted and the solid indifference curves in Figure 3.2).

So, starting from the same endowment, the set of gambles acceptable to J is strictly smaller than the acceptable set for K.

We have shown that greater $A(c)$ implies a smaller set of acceptable gambles. It is easy to verify that the converse is also true. That is, if J has an everywhere smaller acceptance set, his degree of absolute risk aversion $A_J(c)$ is greater.

Exercises and Excursions 3.1

1 *Well-Behaved Preferences*

(A) Does the quadratic utility function exhibit either DARA or IRRA?

(B) An individual with a utility function $v(c)$ such that $-v'(c)/v''(c) = \alpha + \beta c$ is said to exhibit linear risk tolerance. For what parameter values does such an individual exhibit DARA and IRRA?

2 *CARA and CRRA Preferences*

(A) Show that the utility function $v(c) = -e^{-Ac}$ exhibits CARA with coefficient of absolute risk aversion A. Show that it is the unique such function. That is, any other CARA utility function with constant coefficient of absolute risk aversion equal to A must be of the form $v(c) = -ae^{-Ac} + b$ where $a > 0$ and b are constants.

[HINT: To show uniqueness, integrate the function $-\frac{v''(c)}{v'(c)} = A$.]

(B) Show that the utility function (i) $v(c) = c^{1-R}$, $0 < R < 1$ exhibits CRRA with coefficient of relative risk aversion equal to R and (ii) $v(c) = \ln c$ exhibits CRRA with coefficient of relative risk aversion equal to 1.

3 *Preference for Positive Skewness*

Suppose two prospects $\tilde{c}_1$ and $\tilde{c}_2$ have the same mean and variance but $\tilde{c}_1$ has negative and $\tilde{c}_2$ positive skewness, that is:

$$E[(\tilde{c}_1 - \mu)^3] < 0 < E[(\tilde{c}_2 - \mu)^3]$$

(A) Does the typical lottery, offering just a few large prizes with a high probability of a small loss, exhibit positive skewness? Demonstrate.

(B) Ignoring moments higher than the third, use Taylor's expansion to show that, if $v'''(c)$ is positive, then positive skewness is indeed preferred.

(C) Show that decreasing absolute risk aversion is a sufficient condition for $v'''(c)$ to be positive.

4 *Absolute Risk Aversion and Concavity of the Utility Function*
Let $A_i(c)$ be the degree of absolute risk aversion corresponding to the twice-differentiable increasing functions $v_i(c)$, $i = \mathrm{J}, \mathrm{K}$.

(A) If individual J's utility function $v_J(c)$ can be written as an increasing twice-differentiable concave function of individual K's $v_K(c)$, that is:

$$v_J(c) = f(v_K(c)), \quad f'(\cdot) > 0, \; f''(\cdot) < 0$$

show that $A_J(c) > A_K(c)$.

(B) Since $v_J(c)$ is an increasing function, there is a one-to-one mapping $g\colon v \to v$ such that $g(v_K(c)) = v_J(c)$. Differentiate twice and rearrange to establish that:

$$g'(v_K(c)) = \frac{v_J'(c)}{v_K'(c)} \quad \text{and} \quad g''(v_K(c)) = \frac{-v_J'(c)}{v_K'(c)^2}[A_J(c) - A_K(c)]$$

Hence, establish the converse of (A).

5 *Small and Large Gambles*
An individual exhibits constant absolute risk aversion of degree A. This individual will *not* take the 50–50 gamble that gives a gain of \$110 and loss of \$100 (with equal probability).

(A) What is smallest value of A (to three decimal places) that is consistent with this behavior?

Now suppose that this individual is offered another 50–50 gamble with a gain of \$$G$ (greater than \$1,100) and a loss of \$1,000.

(B) Is there any value G that would make this individual accept this second gamble?

6 *The Risk Premium*
Formally, the "risk premium" associated with a risky prospect $\tilde{c}$ is the amount of income b that an individual is willing to give up in order to receive the expected value of $\tilde{c}$ with certainty. That is:

$$\mathrm{E}[v(\tilde{c})] = v(\bar{c} - b), \quad \text{where } \bar{c} = \mathrm{E}[\tilde{c}]$$

(A) If the risk is small so that third and higher-order terms can be neglected, apply Taylor's expansion to show that the risk premium is proportional to the degree of absolute risk aversion.

(B) Let b_0 be the initial risk premium and let b_1 be the risk premium when the individual's wealth rises by w. That is:

(i) $\mathrm{E}[v(\tilde{c})] = v(\bar{c} - b_0)$

(ii) $\mathrm{E}[v(w + \tilde{c})] = v(w + \bar{c} - b_1)$

If the risk is small, appeal to (A) to establish that, if absolute risk aversion is decreasing with wealth, then $b_0 > b_1$.

(C) *Show that this result holds for all risks as long as the degree of absolute risk aversion is decreasing with wealth.

[HINT: Define the new utility function $\bar{v}(c) = v(w + c)$. That is, we can think of the wealth effect as changing the utility function. Given the assumptions, explain why the utility function $\bar{v}(c)$ exhibits a lower degree of risk aversion than $v(c)$. Then appeal to Exercise 4 to establish that $v(w + \tilde{c}) = f(v(\tilde{c}))$, where f is a convex function.]

7 *Effect of an Uncertain Increase in Wealth on the Risk Premium*

Extending the analysis of Exercise 6, suppose two individuals face the same income risk $\tilde{c}$ but one has an additional uncertain endowment $\tilde{w}$. Let b_0 be the amount the first individual is willing to pay to replace the income risk with its expected value $\bar{c}$. Then b_0 satisfies (i) above. Similarly, let b_2 be the amount the second individual is willing to pay to replace the income risk with its expected value. That is:

(iii) $\mathrm{E}[v(\tilde{w} + \tilde{c})] = \mathrm{E}[v(\tilde{w} + \bar{c} - b_2)]$

Arguing by analogy with Exercise 6, it is tempting to think that, with decreasing absolute risk aversion, $b_0 > b_2$. However, consider the following example. There are three states and each is equally likely. The two risky prospects are:

$$\tilde{w} = (w, 0, 0) \quad \text{and} \quad \tilde{c} = (\bar{c}, \bar{c} + e, \bar{c} - e)$$

where $w > 0$ and $\bar{c} > e$.

(A) Write out Equations (i) and (iii) for this example and hence show that b_2 must satisfy:

$$[v(\bar{c} + w) - v(\bar{c} + w - b_2)] - [v(\bar{c}) - v(\bar{c} - b_2)]$$
$$= 3[v(\bar{c} - b_2)] - v(\bar{c} - b_0)]$$

* Starred questions or portions of questions may be somewhat more difficult.

(B) Explain why the left-hand side of this expression is negative and hence why $b_2 > b_0$. That is, the risk premium may *rise* as wealth increases stochastically regardless of whether or not absolute risk aversion is decreasing (see Machina, 1982).

(C) What is the intuition behind this result?

[HINT: You can also show that b_2 is a strictly increasing function of w. Then compare the effect of paying a risk premium on state-1 utility (i) when $w = 0$ and (ii) when w is large.]

3.2 Endowment and Price Effects

This section analyzes the effects of parametric changes in endowments and in prices upon the individual's risk-bearing optimum. We first take up the case where Complete Contingent Markets (CCM) are provided by a full set of tradable state claims (Section 3.2.1). Equivalent results can, of course, be obtained under Complete Asset Markets (CAM) – i.e., where the number of tradable assets with linearly independent return vectors equals the number of states. Section 3.2.2 then covers *incomplete* market regimes.

3.2.1 Complete Markets

In accordance with the analysis in Chapter 2, Section 1, under Complete Contingent Markets the individual chooses among state-claim bundles $(c_1, \ldots, c_S)$ so as to maximize expected utility $U = \sum_s \pi_s v(c_s)$ subject to the budget constraint:

$$\sum_{s=1}^{S} p_s c_s = \sum_{s=1}^{S} p_s \bar{c}_s \equiv \bar{W} \qquad (3.2.1)$$

Ignoring for expositional ease the possibility of a corner solution, the preferred position C^* is the one satisfying the budget constraint and the Fundamental Theorem of Risk Bearing:

$$\frac{\pi_1 v'(c_1)}{P_1} = \frac{\pi_2 v'(c_2)}{P_2} = \cdots = \frac{\pi_s v'(c_s)}{P_s} = \lambda \qquad (3.2.2)$$

where λ can be interpreted as the expected marginal utility of income. This condition can only be an optimum, of course, if risk aversion ($v''(c) < 0$) is postulated – else a corner solution would always be preferred.

In the two-state diagram of Figure 3.1, consider exogenous shifts in the individual's endowment vector, state-claim prices being held constant. The effect upon the risk-bearing optimum depends only upon whether or not the change in endowment alters endowed wealth $\bar{W}$ in (3.2.1). An endowment variation leaving $\bar{W}$ unchanged is illustrated in Figure 3.1 by a shift from $\bar{C}$ to $\bar{\bar{C}}$ along the same market line LL′. Such a change in the composition of the endowment does not in any way affect the position of the optimum vector C*. (On the other hand, this change will of necessity affect the scope of the *transactions* undertaken by the individual in order to attain the C* optimum.) The more interesting class of endowment shifts will be those in which $\bar{W}$ does change, so that the individual's optimum position must also be revised. (But, it is at least possible that the *transactions* he must undertake to attain his new optimum from his new endowment might remain unchanged.) For concreteness, we will speak in terms of *increases* in $\bar{W}$.

An increase in wealth at given prices must raise the optimum amount of contingent consumption claims held in at least one state t. Assuming risk aversion, λ must fall in (3.2.2). And, given the separable form of the expected-utility function, as in Equation (1.4.1), when any c_t rises in (3.2.2) then c_s must increase in each and every other state as well. Thus, in risk-bearing theory under the von Neumann–Morgenstern postulates, there are no "inferior-good" state claims; all wealth effects are necessarily "normal." Then the analysis of wealth expansion paths in the previous section can be applied directly. For any pair of states s and t the impact of an increase in wealth can be depicted essentially as in Figure 3.1. We need only let the axes be c_s and c_t, and interpret the budget lines LL′ and NN′ as indicating those combinations of state claims costing the same as the *optimal* purchase of state claims at the two wealth levels. It follows directly that under the assumption of decreasing (increasing) *absolute* risk aversion, the absolute difference between any pair of state claim holdings, $|c_t^* - c_s^*|$, rises (falls) with wealth. Similarly, under the assumption of decreasing (increasing) *relative* risk aversion, if $c_t > c_s$ then the ratio of expenditures $P_t c_t / P_s c_s$ rises (falls) with wealth.

We next consider the "pure substitution effect," the impact upon c_s of a *compensated* increase in the price P_s of one of the state claims. That is, we postulate an exogenous increase in P_s together with a simultaneous change in endowment such that expected utility (after the individual revises his state-claim holdings in accordance with the new price vector) is the same as before.

Suppose for concreteness that it is P_1, the price of claims to consumption in state 1, that rises. Under our standard assumption of state independence (so that $v'(c_s)$ is independent of consumption in any state $t \neq s$), and, since $P_2, \ldots, P_S$ are all unchanged, the Fundamental Theorem of Risk Bearing indicates that if λ rises $c_2, \ldots, c_s$ must fall and vice versa. That is, claims in all states other than state 1 move together. Since expected utility is required to remain constant, either c_1 falls and $c_2, \ldots, c_s$ all rise or the reverse. But, if c_1 and P_1 both were to rise, the marginal utility of income λ in Equation (3.2.3) would fall. Then, to maintain the equality, holdings of all other state claims would also rise. But this is inconsistent with constant utility. We have therefore established that the "pure substitution effect" of a price increase is negative and that all cross-effects are positive. In the language of traditional theory, state claims must be *net substitutes* in demand.

To determine the *uncompensated* effect of an increase in P_s on demand, note that, if $\bar{c}_s$ is the individual's state-s endowment, the Slutsky equation is:

$$\frac{\partial c_s}{\partial P_s} = \left.\frac{\partial c_s}{\partial P_s}\right|_{\text{comp}} - (c_s - \bar{c}_s)\frac{\partial c_s}{\partial W}$$

Since it has been shown that all wealth effects are normal, the two terms on the right-hand side of this expression are reinforcing as long as the individual is a net buyer of state-s claims. Informally, the increase in P_s makes a net buyer poorer, and so the income effect is negative. However, when we consider the effect of a rise in P_t upon the demand for state-s claims c_s, the substitution effect tending to increase c_s must be weighed against the income effect that tends to reduce c_s. So, state claims may be either *gross substitutes* or *gross complements* in demand.

Exercises and Excursions 3.2.1

1 *The Law of Demand for State Claims*

(A) Suppose an individual is a net buyer of consumption claims in state 1, i.e., $c_1^* > \bar{c}_1$. If the price of state-1 claims rises, show directly that the quantity of state-1 claims demanded must fall.

[HINT: Suppose the proposition is false. Then c_1^* does not fall and spending on state-1 claims must rise. Apply the Fundamental Theorem of Risk Bearing to show that $c_2^*, \ldots, c_s^*$ must also rise. Hence obtain a contradiction.]

(B) What if the individual is a net *seller* of state-1 claims?

2 *Elastic Own-Demand Curves and Gross Substitutes in Demand*

An individual begins with a fixed nominal wealth $\bar{W}$.

(A) Show that if, for each state s, the own-price elasticity of demand exceeds unity so that total spending on state-s claims falls with a rise in P_s, then all state claims are gross substitutes in demand.

[HINT: Use the Fundamental Theorem of Risk Bearing, and the fact that $P_s\, c_s$ declines as P_s rises, to establish that λ must decline.]

(B) If an individual is highly risk averse (so that indifference curves are essentially L-shaped) explain graphically why consumption in each state declines as P_s rises. In this case, does the own-price elasticity of demand exceed unity?

(C) Show that the own-price elasticity of demand exceeds unity if and only if relative risk aversion is less than unity.

[HINT: Show that consumption in state s and state 1 must satisfy:

$$\pi_s c_s v'(c_s) = (P_s c_s)\pi_1 v'(c_1)/P_1$$

Use (A) to establish that, as P_s rises, the right-hand side of this equation declines if and only if own-price elasticity exceeds unity. What happens to the left-hand side as c_s declines?]

3.2.2 Incomplete Markets

To illustrate trading in a regime of incomplete markets, in a world of $S > 2$ states consider an individual who must balance his portfolio between a riskless asset and a single risky asset. As before, we want to examine the effect of changes in endowments or in prices upon the individual's risk-bearing optimum.

Let the first asset, with price P_1^A, have the certain return z_1 while the second asset, with price P_2^A, pays off z_{2s} dollars in state s ($s = 1, 2, \ldots, S$). Equivalently, the return on asset 2 is a random variable $\tilde{z}_2$ with realizations $z_{21}, \ldots, z_{2S}$. Then, if an individual with utility function $v(\cdot)$ holds q_1 units of asset 1 and q_2 units of asset 2, and if he has no source of income in any state other than returns from asset holdings, his expected utility is:

$$U(q_1, q_2) = \mathrm{E}_s[v(q_1 z_1 + q_2 \tilde{z}_2)] = \sum_s \pi_s v(q_1 z_1 + q_2 z_{2s}) \quad (3.2.3)$$

As long as the individual is risk averse so that $v(c)$ is a concave function, it can be shown that his derived preferences *over assets* must be convex ("bowed toward the origin") as depicted in Figure 3.3.[1]

Asset demands must satisfy the budget constraint:

$$P_1^A q_1 + P_2^A q_2 = P_1^A \bar{q}_1 + P_2^A \bar{q}_2 \equiv \bar{W} \quad (3.2.4)$$

[1] See the first exercise at the end of this section.

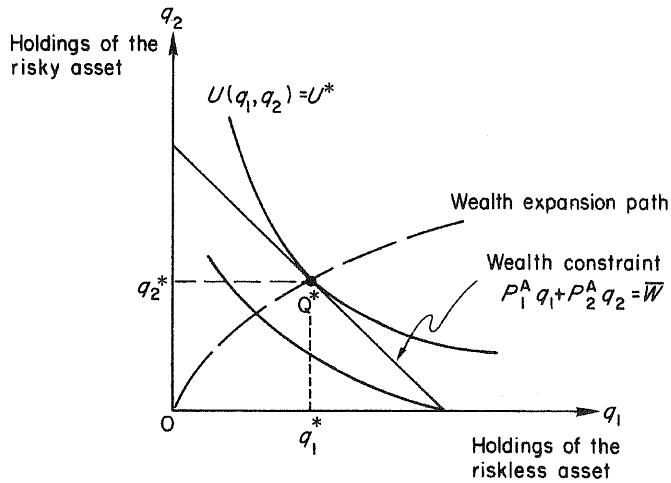

Figure 3.3. Optimal portfolio selection.

Using the budget constraint, we can substitute for q_1 in (3.2.3) so as to express expected utility as a function $\bar{U}(q_2)$ of q_2 only:

$$\bar{U}(q_2) = \mathrm{E}\left[v\left(\left(\frac{\bar{W}}{P_1^A} - \frac{P_2^A}{P_1^A}q_2\right)z_1 + q_2\tilde{z}_2\right)\right]$$
$$= \mathrm{E}\left[v\left(\frac{\bar{W}z_1}{P_1^A} - P_2^A q_2\left(\frac{\tilde{z}_2}{P_2^A} - \frac{z_1}{P_1^A}\right)\right)\right] \tag{3.2.5}$$

For the analysis that follows, it is more convenient to work with net asset yields per dollar invested (or, for short, simply the *yields*) rather than with the asset payoffs or returns.

For any asset a, the yield distribution is definitionally related to the payoff distribution by:

$$1 + \tilde{R}_a \equiv \frac{\tilde{z}_a}{P_a^A} \tag{3.2.6}$$

Substituting into (3.2.5) we obtain:

$$\bar{U}(q_2) = \mathrm{E}\left[v((1 + R_1)\bar{W} + (\tilde{R}_2 - R_1)P_2^A q_2)\right]$$

where R_1 (like z_1) is non-stochastic. The marginal increase in expected utility associated with an increased demand for the risky asset is then:

$$\bar{U}'(q_2) = P_2^A \mathrm{E}[(\tilde{R}_2 - R_1)v'(\tilde{c})] \tag{3.2.7}$$

The random variable $\tilde{c}$ represents contingent consumption over states:

$$\tilde{c} \equiv (1 + R_1)\bar{W} + (\tilde{R}_2 - R_1)P_2^A q_2 \tag{3.2.8}$$

In particular, at $q_2 = 0$ we have:

$$\bar{U}'(0) = P_2^A v'((1 + R_1)\bar{W})\mathrm{E}[\tilde{R}_2 - R_1]$$

This is positive if and only if the expected yield on the risky asset exceeds the yield on the riskless asset.

Thus, no matter how risk averse an individual is, he will always want to hold some amount of the risky asset if its expected yield is even slightly higher than the riskless yield. While this may at first seem puzzling, the explanation is simple. As long as the risk holding is sufficiently small, the individual's final consumption distribution is nearly riskless, and so marginal utilities are almost the same across states. His behavior toward a very small favorable gamble is therefore essentially that of a risk-neutral agent.

We now suppose that $\mathrm{E}[\tilde{R}_2] > R_1$ so that the optimal asset holding, q_2^*, is strictly positive. From (3.2.7) and (3.2.8), q_2^* satisfies the first-order condition:

$$\bar{U}'(q_2^*) = P_2^A \mathrm{E}[(\tilde{R}_2 - R_1)v'(\tilde{c}^*)] = 0 \tag{3.2.9}$$

What happens to demand for the risky asset as wealth increases? From the previous section, we would anticipate that an individual characterized by decreasing absolute risk aversion (DARA) should want to hold more of the risky asset. We now show that this intuition is correct.

The analysis proceeds by asking what happens to the marginal utility of q_2 as $\bar{W}$ increases. At the asset-holding optimum for the initial wealth level, we have seen, $U'(q_2^*) = 0$. If the effect of an increase in wealth is to raise this marginal utility, the new optimum will have to be at a higher level of q_2.

Differentiating (3.2.9) by $\bar{W}$ and making use of (3.2.8) we obtain:

$$\begin{aligned}\frac{d}{d\bar{W}}\bar{U}'(q_2) &= P_2^A(1 + R_1)\mathrm{E}[(\tilde{R}_2 - R_1)v''(\tilde{c}^*)] \\ &= -P_2^A(1 + R_1)\mathrm{E}[(\tilde{R}_2 - R_1)v'(\tilde{c}^*)A(\tilde{c}^*)] \end{aligned} \tag{3.2.10}$$

where $A(\tilde{c}) \equiv -v''(\tilde{c})/v'(\tilde{c})$ is the degree of absolute risk aversion.

If $A(c) = \bar{A}$ is constant, this can be rewritten as:

$$\frac{d}{dW}\bar{U}'(q_2) = -P_2^A(1 + R_1)\bar{A}\mathrm{E}[(\tilde{R}_2 - R_1)v'(\tilde{c}^*)]$$

From (3.2.9) the expectation is zero. Therefore, under Constant Absolute Risk Aversion (CARA), demand for the risky asset is independent of wealth.

Under the normal assumption of decreasing $A(c)$, the analysis is only a bit more complicated. If $R_2 = R_1$ then, from (3.2.8), $c = (1 + R_1)\bar{W}$ and so $A(c) = A((1 + R_1)\bar{W})$. If $R_2 > R_1$, c is larger and so $A(c) < A((1 + R_1)\bar{W})$. Then:

$$(R_2 - R_1)A(c) < (R_2 - R_1)A((1 + R_1)\bar{W})$$

If $R_2 < R_1$, c is smaller and so $A(c) > A((1 + R_1)\bar{W})$. Then again:

$$(R_2 - R_1)A(c) < (R_2 - R_1)A((1 + R_1)\bar{W})$$

Since this is true for all c, it follows from (3.2.10) that:

$$\begin{aligned}\frac{d}{d\bar{W}}\bar{U}'(q_2^*) &> -P_2^A(1 + R_1)\mathrm{E}[(\tilde{R}_2 - R_1)v'(\tilde{c}^*)A((1 + R_1)\bar{W})] \\ &= -P_2^A(1 + R_1)A((1 + R_1)\bar{W})\mathrm{E}[(\tilde{R}_2 - R_1)v'(\tilde{c}^*)]\end{aligned}$$

Again, from (3.2.9), the right-hand side of this inequality is zero. Therefore, at the initial optimum, an increase in wealth raises the expected marginal utility of investing in the risky asset and so raises demand for the asset.

Returning to Figure 3.3, it follows that the wealth expansion path in asset space is upward sloping.[2] As depicted, it bends forward so that, as wealth increases, there is a less-than-proportional increase in demand for the risky asset. You are asked to confirm in an exercise at the end of this section that this will be the case under Increasing Relative Risk Aversion (IARA).

In conclusion, in a regime of incomplete markets with a single risky asset, the wealth effect upon the demand for that asset will be positive, zero, or negative according as *absolute* risk aversion $A(c)$ is decreasing (DARA), constant (CARA), or increasing (IARA). So the uncompensated demand for the risky asset, in the region where the individual is a net buyer, will surely have negative slope under DARA or CARA but not necessarily so under IARA. As for the *riskless* asset, its demand (once again, in the region where the individual is a net buyer) must always have negative slope – since the wealth effect is surely positive. (A richer individual will always want to increase his contingent consumption in each and every state of the world.)

[2] It is tempting to generalize from this and conjecture that, with one riskless asset and several risky assets, total spending on the latter would rise with wealth. However, as Hart (1975) has shown, special cases can be constructed for which this is not the case. Despite this, there remains the presumption that wealth and total spending on risky assets will be positively related.

Exercises and Excursions 3.2.2

1 Concavity of the Derived Utility Function

Let $q = (q_1, \ldots, q_A)$ be an individual's holdings of A assets. In state s the return on each of these assets is $z_s = (z_{1s}, \ldots, z_{AS})$ so that the total state-s income is:

$$q \cdot z_s = \sum_{a=1}^{A} q_a z_{as}$$

Expected utility is then:

$$U(q) = \sum_{s=1}^{S} \pi_s v(q \cdot z_s)$$

where v is an increasing strictly concave function.

(A) Show that $U(q)$ is also strictly concave, that is, for any pair of vectors q^α, q^β:

$$U(\lambda q^\alpha + (1-\lambda)q^\beta) > U(q^\alpha) + (1-\lambda)U(q^\beta), \quad 0 < \lambda < 1$$

(B) Hence confirm that preferences are convex, as depicted in Figure 3.3.

2 Asset Demand with Constant Absolute Risk Aversion

Suppose $v(c) = \kappa_1 - \kappa_2 e^{-Ac}$, $\kappa_1, \kappa_2 > 0$. There are M assets, all of which are risky except asset 1.

(A) Write down the individual's optimization problem and then substitute for q_1, the demand for the riskless asset, using the wealth constraint.

(B) Write down the necessary conditions for an optimal portfolio and confirm that demands for risky assets are independent of initial wealth $\bar{W}$.

3 Demand for a Risky Asset under Increasing Relative Risk Aversion

Let κ_2 be the proportion of initial wealth invested in the risky asset in the portfolio problem described in this section.

(A) Obtain an expression for $U(\kappa_2)$, expected utility as a function of κ_2, and hence show that under constant relative risk aversion (CRRA) the optimal proportion κ_2^* is independent of wealth $(\partial U'(\kappa_2^*)/\partial \bar{W} = 0)$.

(B) Apply methods similar to those used in Section 3.2 to establish that, under increasing relative risk aversion (IRRA), κ_2^* declines with wealth.

(C) What occurs under DRRA?

4 *Demand for a Risky Asset with Different Attitudes towards, Risk (Pratt 1964)*

Suppose that individual J is everywhere more risk averse than K, so that (in accordance with an earlier exercise) J's utility function $v_J(c)$ is an increasing concave transformation of $v_K(c)$:

$$v_J(c) = f(v_K(c)), \quad f(\cdot) > 0, \quad f''(\cdot) < 0$$

(A) Show that, if both individuals face the portfolio-choice problem described in this section, and the prices of the riskless and risky assets are both unity, the optimal holding of the risky asset for J, q_2^J, satisfies:

$$\bar{U}_J'(q_2^J) = P_2^A \mathrm{E}[(\tilde{R}_2 - R_1) f'(v_\kappa(\tilde{c})) v_\kappa'(\tilde{c})] = 0$$

where $\tilde{c} = (1 + R_1)\bar{W} + (R_2 - R_1) P_2^A q_2^J$.

(B) Confirm that for each possible realization R_2:

$$(R_2 - R_1) f'(v_K(\tilde{c})) < (R_2 - R_1) f'(v_K(\bar{W}))$$

Hence show that $\bar{U}_K'(q_2^J) > 0$ and therefore that:

$$q_2^J < q_2^K$$

3.3 Changes in the Distribution of Asset Payoffs

The previous section examined the effects of parametric changes in wealth, or in the prices of state claims or of assets, upon the risk-bearing optimum of the individual. For example, we showed that, under a regime of Complete Contingent Markets (CCM), if the individual was previously at an interior optimum, then an increase in wealth would increase his holdings of each and every state claim. Owing to this positive wealth effect, the uncompensated demand curve for any contingent claim is negatively sloped in the region where the individual is a net buyer. An analogous conclusion evidently holds for Complete Asset Markets (CAM). But, with incomplete markets, the uncompensated demand for an asset on the part of a net buyer is unambiguously negatively sloped only if the individual is characterized by decreasing or constant absolute risk aversion (DARA or CARA).

This section represents a shift in point of view. Here the parametric changes impacting upon the individual take the form of shifts in the distribution of the contingent returns or payoffs z_{as} of some particular asset a. Let us reconsider the example of the previous section, in which an individual chooses to hold q_1 units of the riskless asset with payoff z_1 and q_2 units of the risky asset with state s return z_{2s} for $s = 1, \ldots, S$. For simplicity, assume that the endowment is entirely in units of the riskless asset. Then, with endowed wealth $\bar{W} = P_1^A \bar{q}_1$ the utility-maximizing portfolio choice (q_1 and q_2) is the solution to:

$$\underset{q_1, q_2}{\text{Max}} \left\{ U(q_1, q_2) \equiv \sum_{s=1}^{S} \pi_s v(q_1 z_1 + q_2 z_{2s}) \,\middle|\, P_1^A q_1 + P_2^A q_2 = \bar{W} \right\}$$

For simplicity, let $P_1^A = P_2^A = 1$ so that the budget constraint becomes $q_1 + q_2 = \bar{W}$. Then if the individual purchases q_2 units of the risky asset, his final state-s consumption is:

$$c_s = (\bar{W} - q_2)z_1 + q_2 z_{2s} = \bar{W} z_1 + q_2(z_{2s} - z_1) \qquad (3.3.1)$$

As a first illustration of a parametric change, suppose the return to the risky asset 2 declines in one state of the world but is otherwise unaltered. (In the market as a whole, such a shift would tend to change P_2^A, or more generally, the entire pattern of asset prices – a topic to be covered in Chapter 4. But in this chapter we are continuing to focus upon a single individual so that asset prices are assumed constant.) It is tempting to conclude that the individual would then respond by investing less in the risky asset. However, this intuitively appealing argument is not in general true!

Continuing to assume a single riskless and a single risky asset, suppose there are only two states ($S = 2$) as depicted in Figure 3.4. (This is therefore a situation of CAM.) Since the individual's endowment holding consists only of the riskless asset 1, in state-claim space the endowment point is $\bar{C} = (\bar{W} z_1, \bar{W} z_1)$ on the 45° certainty line. If the individual invests everything in the risky asset instead (so that $q_2 = \bar{W}$), his income claims are represented by the point $R = (\bar{W} z_{21}, \bar{W} z_{22})$. Before the postulated parametric change occurs, then, the set of feasible contingent incomes consists of all the weighted averages of $\bar{C}$ and R yielding non-negative consumptions in both states. Geometrically, this is the line LL′ through $\bar{C}$ and R, extending to the axes (since short sales are allowed within the first quadrant). Along this line the optimum position is the tangency at C*.

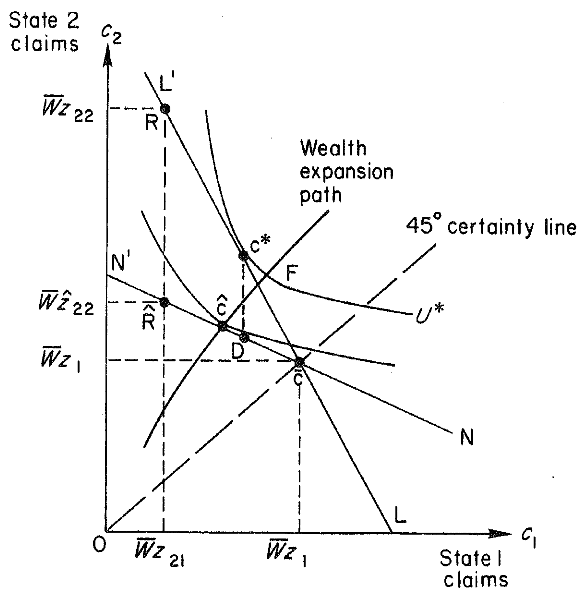

Figure 3.4. Demand for a risky asset.

As depicted, $z_{22} > z_{21}$. That is, the risky asset yields more in state 2 than in state 1. Now suppose that the state-2 yield of the risky asset declines, from z_{22} to $\hat{z}_{22}$. This is equivalent, from the individual's point of view, to a rise in the effective price P_2 of state-2 claims. On the basis of the discussion in the preceding Section 3.2.1, we know that for a net buyer of the risky asset there will be a *substitution effect* away from state-2 consumption and toward state-1 consumption, together with an *income effect* tending to reduce consumption in both states. So the result is unambiguous that at the new risk-bearing optimum there will be a reduction of consumption in state 2, i.e., that $\hat{c}_2 < c_2^*$. Nevertheless, it does not necessarily follow that there will be a reduction in purchases of the risky asset. Since each unit of asset 2 now yields fewer units of state-2 return ($\hat{z}_{22} < z_{22}$) than before, it *may* be the case that the individual would have to buy more units of asset 2 even to generate the reduced quantity of c_2 that he wants to consume.

In Figure 3.4, note that the postulated shift in the risky asset's payoff does not affect the endowment position $\bar{C}$ (which consists of holdings of the *riskless* asset only). So the feasible consumption vectors in Figure 3.4 are now the points on the flatter market line NN′ through $\bar{C}$ and $\hat{R}$. Also, it follows from (3.3.1) that, for any given portfolio (q_1, q_2), the shift in z_{22} leaves income in state 1 unchanged. Geometrically, the new income vector generated by any portfolio lies vertically below the old. In particular, $\hat{R}$ lies vertically below R and the state-income yield vector C* corresponding to the original asset-holding optimum becomes D.

Whether the individual buys more or less of the risky asset then hinges upon whether his indifference-curve tangency lies to the northwest or southeast of D along the line NN′. Specifically, if (as shown in the diagram) the new risk-bearing optimum $\hat{C}$ lies northwest of D, there will be *increased purchases* of asset 2 $(\hat{q}_2 > q_2^*)$ even though there is *decreased contingent consumption* in the event of state 2 occurring $(\hat{c}_2 < c_2^*)$. As an obvious corollary, whenever this occurs there will also be a reduction in state-1 consumption. Thus, the reduced state-2 yield of the risky asset "spills over" into reduced consumption in both states. On the other hand, should the new optimum $\hat{C}$ lie southeast of C*, there will be *reduced purchases* of asset 2 (as well as *reduced consumption* of state 2), hence increased holdings of asset 1 and increased consumption in state 1.

The more rapidly the slope of the indifference curve changes, the smaller is the substitution effect away from c_2-consumption. Hence the more likely it is that the wealth effect of the implicit increase in P_2 dominates, so that the individual's purchase of the risky asset increases. Referring back to Section 3.1 we see that the curvature of the indifference curve is greater the larger is the individual's aversion to risk. So the seemingly paradoxical result, that demand for the risky asset (after a decline of z_{22} to $\hat{z}_{22}$) can increase, is more likely if an individual exhibits a high degree of risk aversion. (From another point of view, however, this is not paradoxical at all. The shift from z_{22} to $\hat{z}_{22}$ has made the risky asset "less risky" – has reduced the gap between z_{22} and z_{21} – which has to some extent increased its attractiveness for highly risk-averse individuals.)

In conclusion, in a simplified regime of two tradable assets (one risky, the other riskless) and two states of the world, a reduction in one of the contingent payoffs z_{as} for the risky asset is equivalent, from the individual's point of view, to an increase in the price P_s of the corresponding state claim. It follows that at the new optimum the individual will reduce his contingent consumption c_s in that state of the world. But he will not necessarily reduce his portfolio holding of the risky asset. And, in particular, if z_{as} declines for the higher-yielding state, a highly risk-averse individual's optimal holding of the risky asset may actually increase – since that asset has in effect become "less risky."

Exercise 1 below proves a related proposition, that if z_{as} falls then the demand for the risky asset will decline if the degree of relative risk aversion, *R*, is not greater than unity.

Exercises and Excursions 3.3

1 State Returns and Relative Risk Aversion

Choosing units so that $P_1^A = P_2^A = 1$, a risk-averse individual is endowed with $\bar{W}$ units of a riskless asset 1 returning z_1 in each state. He can also make purchases of a risky asset 2 whose payoff is z_{2s} in state s. Initially his optimum holding of the risky asset is positive. Show that if the return z_{2s} on asset 2 rises in some state s, and if the individual's constant relative aversion to risk CRRA is no greater than unity, then his optimal holding of this asset will rise.

2 Parametric Change Lowering Mean and Raising Variance of Asset Payoff

An individual with an initial wealth of \$50 must choose a portfolio of two assets, both of which have a price of \$50. The first asset is riskless and pays off \$50 in each of the two possible states. The second returns z_{2s} in state s, for $s = 1, 2$. The probability of state 1 is π.

(A) If the individual splits his wealth equally between the two assets, confirm the correctness of the following table, where the risky asset returns may have the form of α, β, or γ.

(B) Suppose the individual has a utility function:

$$v(c) = -e^{-Ac}$$

where $A = \frac{\ln 4}{30}$ (and hence $e^{30A} = 4$). Confirm that the individual's preference ranking of the three risky assets is $\gamma > \alpha > \beta$.

	Risky asset returns (z_{21}, z_{22})	Probability of state 1	Final consumption (c_1, c_2)	$E(c)$	$\sigma^2(c)$
α	(20,80)	1/5	(35,65)	59	144
β	(38,98)	1/2	(44,74)	59	225
γ	(30,90)	1/3	(40,70)	60	200

(C) With preferences as given in (B) show that in each case the individual's optimal decision is to spend an equal amount on each of the two assets.

3.4 Stochastic Dominance

From the exercises at the end of Section 3.3 it is clear that, to derive strong qualitative predictions as to asset holdings in response to parametric changes in payoff distributions, we must introduce additional restrictions – either upon probability distributions or upon preferences. In the following section we describe some restrictions that do have general implications.

3.4.1 Comparison of Different Consumption Prospects

This section adopts a somewhat different approach to the risk-bearing decision. Instead of considering the specific effects of changes in wealth, in state-claim prices, in asset payoffs, etc., we ask under what general conditions it is possible to assert that one prospect or state-distributed consumption vector is preferred over another. We want to be able to answer this question by comparing the probability distributions of consumption alone, while calling only upon standard properties of individuals' preferences – to wit, positive marginal utility of income ($v'(c) > 0$) and risk aversion ($v''(c) < 0$). In this section it will be more convenient to deal with continuous distributions, equivalent to assuming a continuum rather than a finite or countably infinite number of states of the world.

Let $\tilde{c}_1$ and $\tilde{c}_2$ be two consumption prospects and suppose that an individual with utility function $v(c)$ prefers the former. That is:

$$E[v(\tilde{c}_1)] > E[v(\tilde{c}_2)] \tag{3.4.1}$$

We can write the two cumulative distribution functions as:

$$F(c) = \text{Prob}[\tilde{c}_1 \leq c]$$
$$G(c) = \text{Prob}[\tilde{c}_2 \leq c]$$

Let us assume that both $\tilde{c}_1$ and $\tilde{c}_2$ lie between the limits α and β and that both F and G are continuously differentiable. Then for the two distributions there are associated density functions, $F'(c)$ and $G'(c)$. We can rewrite (3.4.1) as:

$$\underset{F}{E}[v(\tilde{c})] \equiv \int_{\alpha}^{\beta} v(c)F'(c)dc > \int_{\alpha}^{\beta} v(c)G'(c)dc \equiv \underset{G}{E}[v(\tilde{c})] \tag{3.4.2}$$

In general, two individuals with different preferences will have different rankings of these two consumption prospects. However, in some cases it is possible to obtain an ordering that holds for all individuals regardless of their preferences (subject only to the standard properties of positive

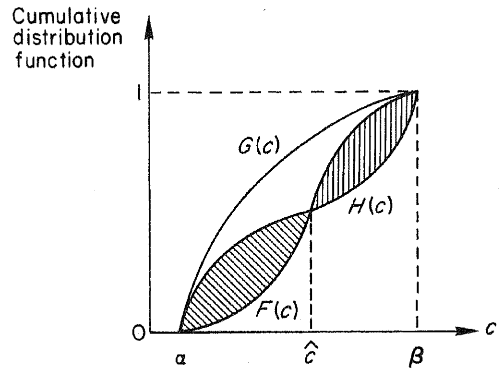

Figure 3.5. First- and second-order stochastic dominance.

marginal utility and risk aversion). In other words, we want to see how far we can get looking only at the probability distributions. When a choice between two prospects can be made using this information alone, it will be said that one distribution *stochastically dominates* the other.

Definition 3.1: First-order stochastic dominance

If, for all c, $F(c) \leq G(c)$ and the inequality is strict over some interval, the distribution F exhibits first-order stochastic dominance over G.

This definition leads immediately to:

Ranking Theorem I

For all increasing, piecewise differentiable functions $v(c)$, if F exhibits first-order stochastic dominance over G, then:

$$\underset{F}{\mathrm{E}}[v(\tilde{c})] > \underset{G}{\mathrm{E}}[v(\tilde{c})]$$

Consequently, if the prospect or distribution F is first-order stochastically dominant over G, then any individual with positive marginal utility of income will prefer F to G.

The property of the distribution functions leading to first-order dominance is evident from inspection of Figure 3.5. Here F and H are both first-order stochastically dominant over G but neither F nor H is first-order dominant over the other. Diagrammatically, the F and H curves both lie always below (and so also to the right of) G, but the F and H curves cross.

Following our usual practice, we will emphasize the intuitive meaning of this condition. (More rigorous statements are left as exercises.) First, compare the F and G curves. F being always below G means that, for each and every income level c° between α and β, the cumulative probability that

c is smaller than that income, that $c \leq c^{\circ}$, is greater for G than for F. Thus, no matter what level of income we look at between these limits, G always has a greater probability mass in the lower tail than does F. Alternatively, we could express this in terms of the income associated with any given probability level. For example, a lower-tail cumulative probability of, say, 0.5 occurs at a higher income for F than for G. In other words, the distribution F has a higher median (50th percentile) income than G. And, similarly, each and every percentile of the F distribution is at a greater income than the corresponding percentile of the G distribution. So we can confidently say (provided only that the marginal utility of income $v'(c)$ is always positive) that F will surely be preferred. We cannot make a similar comparison of F and H, however. Since F and H cross, comparisons of probability masses in the lower tail (or of income levels associated with any percentile of probability) will not always point the same way.

Only under quite stringent conditions will one distribution ever exhibit first-order stochastic dominance over another. So Ranking Theorem I is not very far reaching. This is not surprising, because only the first standard property of the utility function, that $v'(c) > 0$, has been exploited. A more powerful theorem, involving the concept of *second-order* stochastic dominance, also makes use of the risk-aversion property $-v''(c) < 0$.

Definition 3.2: Second-order stochastic dominance

If for all c:

$$\int_{-\infty}^{c} F(r)dr \leq \int_{-\infty}^{c} H(r)dr \tag{3.4.3}$$

with the inequality holding strictly over some part of the range, then the distribution F exhibits second-order stochastic dominance over H.

Geometrically, F is second-order dominant over H if, over every interval $[\alpha, c]$, the area under $F(c)$ is never greater (and sometimes smaller) than the corresponding area under $H(c)$. This is equivalent, of course, to the diagonally shaded area in Figure 3.5 being greater than the vertically shaded region.

Definition 3.2 leads directly to:

Ranking Theorem II

For all *increasing concave* twice-piecewise-differentiable functions $v(c)$, the concavity being strict somewhere, if F exhibits second-order stochastic dominance over H, then:

$$\underset{F}{\mathrm{E}}[v(c)] > \underset{H}{\mathrm{E}}[v(c)]$$

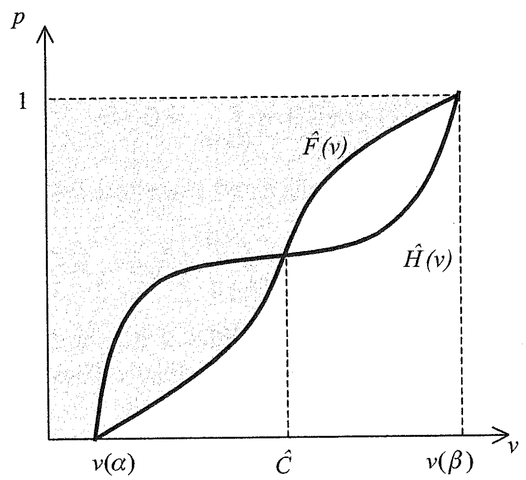

Figure 3.6. Cumulative distribution and expected value.

The intuitive interpretation of second-order stochastic dominance parallels the interpretation of first-order stochastic dominance. As a first step, it is useful to define $\hat{F}(v)$ to be the cumulative distribution function for *final utility* v, when c has the distribution $F(c)$. Note then that:

$$\hat{F}(v(c)) \equiv \text{Prob}\,[\tilde{v} \leq v(c)] = \text{Prob}\,[v(\tilde{c}) \leq v(c)] = \text{Prob}\,[\tilde{c} \leq c] = F(c)$$

That is, at the specific value c° located at any given percentile of the distribution of $\tilde{c}$, the corresponding $v(c^\circ)$ is at the same percentile of the distribution of v. Similarly, we define $\hat{H}(v)$ to be the cumulative distribution function for v when c has the distribution $H(c)$.

The key point to appreciate is that, as pictured by the dotted region in Figure 3.6, the area lying to the left of the cumulative distribution $\hat{F}(v)$ represents the expected value of v – that is, expected utility $U = \text{E}[v(c)]$ – under the distribution F. To see this, define $p = \hat{F}(v)$ to be the cumulative probability, so that:

$$dp = \hat{F}'(v)dv$$

The mean of $\hat{F}$ is, of course, defined by:

$$\text{E}\,[v] = \int_{v(\alpha)}^{v(\beta)} v\hat{F}'(v)dv = \int_0^1 v\,dp = \int_0^1 \hat{F}^{-1}(p)dp$$

Geometrically, this corresponds to finding the dotted area by integrating along the vertical rather than the horizontal axis in Figure 3.6. Equivalently, the expected value of v is the area of the rectangle less the area under $\hat{F}(v)$. That is:

$$E[v] = v(\beta) - \int_{v(\alpha)}^{v(\beta)} \hat{F}(v)dv$$

Then, to compare two distributions F and H we note that:

$$\underset{F}{\mathrm{E}}[v(\tilde{c})] - \underset{H}{\mathrm{E}}[v(\tilde{c})] = -\int_{v(\alpha)}^{v(\beta)} [\hat{F}(v) - \hat{H}(v)]dv$$

That is, the difference in the expected utilities of distributions F and H is just the difference in areas under the implied cumulative distribution functions $\hat{F}$ and $\hat{H}$.

Finally, we can rewrite this integral as:

$$\begin{aligned}\underset{F}{\mathrm{E}}[v(\tilde{c})] - \underset{H}{\mathrm{E}}[v(\tilde{c})] &= -\int_{\alpha}^{\beta} [\hat{F}(v(c)) - \hat{H}(v(c)]\frac{dv}{dc}dc \\ &= -\int_{\alpha}^{\beta} [F(c) - H(c)]v'(c)dc\end{aligned}$$

Returning now to Figure 3.5 and looking at the $F(c)$ and $H(c)$ probability distributions, remember that the condition for second-order stochastic dominance requires that the diagonally shaded area (representing the superiority of F over H at low values of c) exceed the vertically shaded area (representing the superiority of H over F at high values of c). But, as we have seen:

$$\begin{aligned}\underset{F}{\mathrm{E}}[v(\tilde{c})] - \underset{H}{\mathrm{E}}[v(\tilde{c})] = &\int_{\alpha}^{\hat{c}} [H(c) - F(c)]v'(c)dc \\ &- \int_{\hat{c}}^{\beta} [F(c) - H(c)]v'(c)dc\end{aligned}$$

As long as $v(c)$ is concave so that $v'(c)$ is declining, $v'(c) \geq v'(\hat{c})$ for $c \leq \hat{c}$ and $v'(c) \leq v'(\hat{c})$ for $c > \hat{c}$. It follows that:

$$\begin{aligned}\underset{F}{\mathrm{E}}[v(\tilde{c})] - \underset{H}{\mathrm{E}}[v(\tilde{c})] > &\int_{\alpha}^{\hat{c}} [H(c) - F(c)]v'(\hat{c})dc \\ &- \int_{\hat{c}}^{\beta} [F(c) - H(c)]v'(\hat{c})dc \\ = &\, v'(\hat{c})\int_{\alpha}^{\beta} [H(c) - F(c)]dc\end{aligned}$$

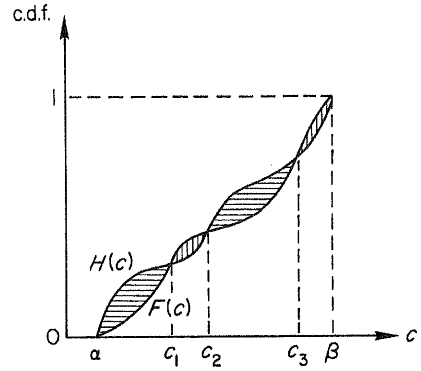

Figure 3.7. Multiple crossings.

From the definition of second-order stochastic dominance, the last integral is positive, and so expected utility is indeed higher under F.

With multiple crossings of F and H the argument is only slightly more complicated. Consider the case of three crossings as in Figure 3.7. Arguing exactly as above, the concavity of $v(c)$ implies that:

$$\int_{\alpha}^{c_2} (H - F)v'(c)dc > v'(c_1) \int_{\alpha}^{c_2} (H - F)dc \qquad (3.4.4)$$

and

$$\int_{c_2}^{\beta} (H - F)v'(c)dc > v'(c_3) \int_{c_2}^{\beta} (H - F)dc \qquad (3.4.5)$$

Second-order stochastic dominance implies that the integral on the right-hand side of (3.4.4) is positive. Therefore, from the concavity of $v(c)$:

$$\int_{\alpha}^{c_2} (H - F)v'(c)dc > v'(c_3) \int_{\alpha}^{c_2} (H - F)dc \qquad (3.4.6)$$

Adding (3.4.5) and (3.4.6) we have, at last:

$$\underset{F}{\mathrm{E}}[v(\tilde{c})] - \underset{H}{\mathrm{E}}[v(\tilde{c})] = \int_{\alpha}^{\beta} (H - F)v'(c)dc > v'(c_3) \int_{\alpha}^{\beta} (H - F)dc$$

Again, given second-order stochastic dominance, the last integral is positive and so F is the preferred distribution.

The limiting case where the two distributions have the same mean, but F exhibits second-order stochastic dominance over H, represents a formalization of the idea that one random variable can be more risky than another. This is illustrated in Figure 3.5 where the diagonally shaded region has the

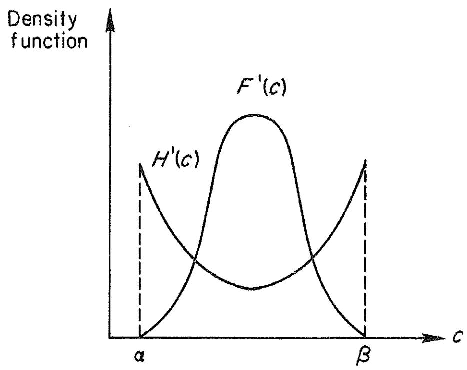

Figure 3.8. Density functions and spread.

same area as the vertically shaded region. Note that the slope of H is greater than the slope of F at both tails of the distribution, while F has a greater slope toward the middle.

From Figure 3.5 we can map the density functions $F'(c)$ and $H'(c)$. These are depicted in Figure 3.8. As already noted, $H'(c) > F'(c)$ toward the ends of the income distribution while $F'(c) > H'(c)$ toward the middle. Then H must have more probability weight in both tails than F. This case, in which probability weight is shifted toward the tails but in such a way that the old and new distributions cross only once, is often referred to as a *simple mean-preserving spread* (Rothschild and Stiglitz, 1971).

As is intuitively clear from Figure 3.8, in the special case where the distribution H represents a simple mean-preserving spread of F it must be that H has higher variance. However, a more powerful result (which is not limited to the single-crossing case) also holds: if F and H have the same mean but F exhibits second-order stochastic dominance over H, then H must have higher variance. This is a direct implication of the following proposition, which follows directly from Ranking Theorem II:

Ranking Theorem III

For all concave functions $v(c)$, the concavity being strict somewhere, if F and H have the same mean and F exhibits second-order stochastic dominance over H then:

$$\underset{F}{\mathrm{E}}[v(\tilde{c})] > \underset{H}{\mathrm{E}}[v(\tilde{c})]$$

Notice that we need not necessarily interpret $v(c)$ as a utility function here (or, for that matter, in the preceding Ranking Theorems). Specifically, we can choose here to let $v(c) = -(c-\mu)^2$, which is, of course, concave in the sense required. Then:

$$-\underset{F}{\mathrm{E}}\left[(\tilde{c}-\mu)^2\right] > -\underset{H}{\mathrm{E}}\left[(\tilde{c}-\mu)^2\right]$$

That is:

$$-\sigma_F^2 > -\sigma_H^2$$

So if F and H have the same mean μ and F is second-order stochastically dominant, then F has smaller variance.

Thinking now of $v(c)$ as a utility function, we have seen that (i) if F is second-order stochastically dominant over H, then F is strictly preferred, and (ii) if in addition F and H have the same mean, then F has smaller variance. But it does not in general follow that, if two distributions $\tilde{c}_1$ and $\tilde{c}_2$ have the same mean and $\tilde{c}_1$ has smaller variance, then $\tilde{c}_1$ is preferred, for $\tilde{c}_1$ might not be stochastically dominant.

Example:

$$\tilde{c}_1 = \begin{cases} 0.4, & \text{with probability } 1/2 \\ 2.1 & \text{with probability } 1/2 \end{cases}$$
$$\tilde{c}_2 = \begin{cases} 0.25, & \text{with probability } 1/9 \\ 1, & \text{with probability } 7/9 \\ 4, & \text{with probability } 1/9 \end{cases}$$

It is readily confirmed that $\mathrm{E}[\tilde{c}_1] = \mathrm{E}[\tilde{c}_2]$ and $\mathrm{var}[\tilde{c}_1] < \mathrm{var}[\tilde{c}_2]$. However, with $v(c) = \ln c$, expected utility is negative in the first case and zero in the second, that is, $\mathrm{E}[v(\tilde{c}_2)] > \mathrm{E}[v(\tilde{c}_1)]$ so that the second prospect is preferred. It is left to the reader to graph the cumulative distribution functions for $\tilde{c}_1$ and $\tilde{c}_2$ and hence to confirm that neither stochastically dominates the other. □

Exercises and Excursions 3.4.1

1 *First-Order Stochastic Dominance*

For any pair of distributions F and H and differentiable function $v(c)$, integrate by parts to establish that:

$$\underset{F}{\mathrm{E}}[v(\tilde{c})] - \underset{H}{\mathrm{E}}[v(\tilde{c})] = \int_\alpha^\beta v(c)[F'(c) - H'(c)]dc$$
$$= \int_\alpha^\beta v'(c)[H(c) - F(c)]dc$$

Hence establish Ranking Theorem I.

2 Second-Order Stochastic Dominance

Appealing to Exercise 1, and integrating by parts a second time, establish Ranking Theorem II.

3 Mean-Preserving Spreads

Use your answer to Question 2 to establish Ranking Theorem III.

4 Stochastic Dominance as a Necessary Condition

(A) In the text it was shown that, if $G(c) \geq F(c)$, then for any non-decreasing function $v(c)$:

$$\underset{F}{\mathrm{E}}[v(\tilde{c})] > \underset{G}{\mathrm{E}}[v(\tilde{c})]$$

By considering the example:

$$v_1(c) = \begin{cases} -1, & c < r \\ 0, & c \geq r \end{cases}$$

establish that, if the condition for first-order stochastic dominance does not hold, there are some non-decreasing utility functions for which the ranking is reversed. That is, for the entire class of non-decreasing utility functions to rank F over G (at least weakly), first-order stochastic dominance is a necessary condition.

(B) By considering the example:

$$v_2(c) = \begin{cases} c - r, & c < r \\ 0, & c \geq r \end{cases}$$

establish the necessity of the second-order stochastic dominance condition (3.4.3) for Ranking Theorem III to hold for all concave functions.

3.4.2 Responding to Increased Risk*

We conclude this chapter by asking how a change in the probability distribution of income that satisfies the conditions of second-order stochastic dominance affects decisions.

Let $c = c(x, \theta)$ be the consequence of taking action x when some exogenous variable takes on the value θ. Moreover, suppose that this exogenous

* Starred sections represent more difficult or specialized materials that can be omitted without significant loss of continuity.

variable is state dependent. Then, given the underlying beliefs about the likelihood of different states and the way θ varies across states, there is some derived distribution function for θ, $F_1(\theta)$. For expositional convenience let θ be distributed continuously. Then, one can write the expected utility of taking action x as:

$$U_1(x) = \int_{-\infty}^{\infty} v(c(x,\theta))F_1'(\theta)\, d\theta$$

A simple illustration is provided by the portfolio choice problem analyzed in Section 3.2. Suppose an individual invests x in a risky asset with gross yield of $1+\theta$ per dollar and his remaining wealth $W-x$ in a riskless asset returning 1 per dollar. His final income is:

$$c(x,\theta) = (W-x)1 + x(1+\theta) = W + x\theta \tag{3.4.7}$$

The question we wish to address is how the individual's portfolio decision is affected by a change in the distribution of the random variable θ.

Returning to the general formulation, suppose that, with distribution function F_1, expected utility is maximized by taking action x_1^*. That is, the rate at which expected utility changes with x:

$$U_1'(x) = \int_{-\infty}^{\infty} \frac{\partial v}{\partial x}(c(x,\theta))F_1'(\theta)d\theta$$

is zero at $x = x_1^*$.

Suppose next that there is a change in the way θ varies with the underlying state of nature. In particular, suppose that the old and new distribution functions for θ, F_1 and F_2, have the same mean and that F_1 exhibits second-order stochastic dominance over F_2.

From Ranking Theorem III it follows immediately that, if $\partial v/\partial x$ is a strictly concave function of θ, then:

$$U_1'(x) = \underset{F_1}{\mathrm{E}}\left[\frac{\partial v}{\partial x}(c(x,\theta))\right] > \underset{F_2}{\mathrm{E}}\left[\frac{\partial v}{\partial x}(c(x,\theta))\right] = U_2'(x)$$

In particular, this inequality holds at x_1^*. Therefore:

$$0 = U_1'(x_1^*) > U_2'(x_1^*)$$

It follows that the individual can increase his expected utility by choosing an action $x_2 < x_1^*$. Thus, assuming $U_2(x)$ has a unique turning point under the new (second-order dominated) probability distribution of returns on the risky asset, the individual will reduce his risky investment.

Note that the requirement that $\partial v/\partial x$ should be concave introduces a restriction on the third derivative of the utility function v. As we shall see in the exercises, this may be satisfied by imposing plausible restrictions on the way absolute and relative risk aversion vary with wealth.

Developing the arguments above only a little bit more yields the following result:

> *Optimal Response Theorem I* (Rothschild and Stiglitz 1971):
>
> Suppose that the distribution functions $F_1(\theta)$ and $F_2(\theta)$ have the same mean and F_1 exhibits second-order stochastic dominance over F_2. Let x_i^* be the solution to:
>
> $$\operatorname*{Max}_x U_i(x) = \int_{-\infty}^{\infty} v(c(x,\ \theta))F_i'(\theta)\, d\theta, \quad i = 1, 2$$
>
> Suppose further that x_i^* is the unique turning point of $U_i(x)$. Then if $\partial v/\partial x$ is a concave (convex) function of θ, x_1^* is greater than (less than) x_2^*.

While this proposition has been widely used in attempts to analyze the effects of mean-preserving increases in risk, results have been somewhat limited. The analysis of Section 3.3 suggests the reason. Mean-preserving increases in risk have both income and substitution effects, and these are often offsetting.

In an effort to overcome this problem, Diamond and Stiglitz suggested considering the effect of a change to a more risky distribution of returns that keeps expected utility constant. To be precise, suppose that the new distribution leaves expected utility constant at the old optimum x^* but the new distribution of *utility* is more risky in the sense of second-order stochastic dominance. The following theorem provides conditions to sign the effect of such a mean-preserving increase in risk.

> *Optimal Response Theorem II* (Diamond and Stiglitz, 1974)
>
> Suppose that $c \in [\alpha, \beta]$ has a continuously differentiable distribution function $F(c)$. Suppose, furthermore, that the solution x^* of the following problem is the unique turning point of $U_F(x)$:
>
> $$\operatorname*{Max}_x U_F(x) = \int_{\alpha}^{\beta} v(x, c)F'(c)dc, \quad \text{where } \frac{\partial v}{\partial c} > 0$$
>
> Then if the distribution shifts from F to G in such a way that, at $x = x^*$, expected utility is unchanged but the new distribution of utility is

more risky (in the sense of second-order stochastic dominance), the new optimum x^{**} is less than x^* if:

$$\frac{\partial^2}{\partial c \partial x} \ln \frac{\partial v}{\partial c} < 0$$

Moreover, if the last inequality is reversed, x^{**} exceeds x^*.

The derivation of this result is only a bit more complicated than that of Optimal Response Theorem I. The interested reader will find a sketch of the proof in the exercises at the end of this section. From these exercises it will become clear that significantly stronger results are possible using Optimal Response Theorem II.

Exercises and Excursions 3.4.2

1 *Optimal Responses to a Change in Risk*

(A) In Optimal Response Theorem I, suppose that the assumption that F_1 and F_2 have the same mean is replaced by the assumption that $\partial v/\partial x$ is an increasing function of θ. Show that the theorem continues to hold.

(B) What conclusions can be drawn if $\partial v/\partial x$ is a *convex* function of θ?

2 *Life-Cycle Saving with Future Income Uncertainty (Leland, 1968)*

(A) Show that a sufficient condition for $v'''(c) > 0$ is that absolute risk aversion, $A(c) \equiv -v''(c)/v'(c)$, is decreasing with wealth.

(B) An individual with current income I_0 and uncertain future income $\tilde{I}_1$ can earn $1 + r$ dollars on each dollar saved. His life-cycle utility is given by the intertemporally additive utility function:

$$v(c_0, c_1) = v_0(c_0) + v_1(c_1)$$

where c_0 is current consumption and c_1 is future consumption. Show that, if the distribution of future income becomes less favorable in the sense of second-order stochastic dominance and if $v_1'''(c_1) > 0$, the optimal level of savings increases.

3 *Portfolio Choice*

(A) Show that:

$$\mu x \frac{v''(\lambda + \mu x)}{v'(\lambda + \mu x)} = -R(\lambda + \mu x) + \lambda A(\lambda + \mu x)$$

where $A(c) = -v''(c)/v'(c)$ and $R(c) = -cv''(c)/v'(c)$.

(B) An individual with wealth $\bar{W}$ invests x in a risky asset with a return of $\tilde{z}_2$ and the rest of his wealth in a riskless asset with yield $z_1 = 1$.

If $A(c)$ is decreasing and $R(c)$ is less than unity and non-decreasing, apply Optimal Response Theorem I to establish the impact on x of a change in the distribution of risky returns that is strictly less favorable in the sense of second-order stochastic dominance.

(C) Analyze also the effect of a mean-utility-preserving increase in risk under the assumptions of decreasing absolute and increasing relative risk aversion.

4 *Mean-Utility-Preserving Increase in Risk*

(A) Under the hypothesis of Optimal Response Theorem II, let $c = \phi\,(x, v)$ be the inverse of the mapping $V = v(x, c)$, that is, $\phi\,(x, v) = v^{-1}\,(x, V)$. Furthermore, let $\hat{F}(V)$ be the implied distribution of V. Confirm that:

$$U_F'(x) = \int_{V_\alpha}^{V_\beta} \frac{\partial v}{\partial x}(x, \phi(x, V)\hat{F}'(V)dV$$

(B) Let x^* be the optimum under the distribution F. Write the corresponding expression for a new distribution G that has the property that, at $x = x^*$, $\hat{G}(V)$ is a mean-preserving spread of $\hat{F}(V)$.

(C) Let x^{**} be the optimum under the new distribution G. Appeal to Ranking Theorem II to show that x^{**} is less than x^* if $\partial v(x^*, \phi(x^*, V))/\partial x$ is a concave function of V.

(D) Define $y(V) \equiv \partial v(x, \phi(x, V))/\partial x$, i.e.:

$$y(v(x, c)) = \frac{\partial v}{\partial x}(x, c)$$

Differentiate by c and hence show that $y'(V)$ can be expressed as follows:

$$y'(v(x, c)) = \frac{\partial v}{\partial x} \ln \frac{\partial v}{\partial c}(x, c)$$

(E) Differentiate this expression again and hence establish Optimal Response Theorem II.

5 *Owner-Operated Firms Facing Demand Uncertainty (Sandmo, 1971)*

Each firm in an industry is owned and operated by a single agent whose best alternative is working elsewhere at wage w. In the production of q units of output, the cost to firm i of all other inputs, $C(q)$, is an increasing convex function, with $C(0) = 0$. Each owner must choose his output level q^* before knowing the final product price $\tilde{p}$. There is free entry into, and exit from, the industry.

(A) If owners are risk neutral, show that the equilibrium expected price denoted as $\bar{p}_n$ must satisfy:

(i) $\bar{p}_n = C'(q^*)$ (ii) $\bar{p}_n q^* - C(q^*) - w = 0$

(B) If owners are risk averse show, that the equilibrium expected price $\bar{p}_a$ exceeds $\bar{p}_n$.

(C) Suppose that initially there is no uncertainty so that the equilibrium price is $\bar{p}_n$ and output per firm is q^*. If prices become uncertain, apply Optimal Response Theorem II to establish that the output of firms remaining in the industry will decline if the following expression is a decreasing function of p:

$$\phi(q, p) = [pq - qC'(q)]\frac{v''(pq - C(q))}{v'(pq - C(q))}$$

(D) Show that, for all $q > 0$, $qC'(q) > C'(q)$. Hence establish that, under the assumptions of decreasing absolute and non-decreasing relative risk aversion, the equilibrium output per firm declines.

[HINT: Appeal to Optimal Response Theorem II and (A) of Exercise 3.]**

SUGGESTIONS FOR FURTHER READING: For an in-depth exposition of risk-aversion, including an extension to inter-temporal models, see Gollier (2001).

References

Arrow, Kenneth J., "The Theory of Risk Bearing," in *Aspects of the Theory of Risk Bearing*, Helsinki: Yrjö Jahnssonin Säätio, 1965. Reprinted in *Essays in the Theory of Risk Bearing*, Chicago: Markham, 1971.

Diamond, Peter A. and Stiglitz, Joseph E., "Increases in Risk and in Risk Aversion," *Journal of Economic Theory*, 8 (1974), 337–360.

Gollier, Christian, *The Economics of Risk and Time*, Cambridge, MA: MIT Press, 2001.

Hart, Oliver D., "Some Negative Results on the Existence of Comparative Statics Results in Portfolio Theory," *Review of Economic Studies*, 42 (1975), 615–621.

Leland, Hayne E., "Savings and Uncertainty: The Precautionary Demand for Saving," *Quarterly Journal of Economics*, 82 (1968), 465–473.

** End of starred section.

Machina, Mark J., "A Stronger Characterization of Declining Risk Aversion," *Econometrica*, 50 (1982), 1069–1080.

Pratt, John W., "Risk Aversion in the Small and in the Large," *Econometrica*, 32 (1964), 122–136.

Rothschild, Michael and Stiglitz, Joseph E., "Increasing Risk I: A Definition," *Journal of Economic Theory*, 2 (1970), 225–243 and "Increasing Risk II: Its Economic Consequences," *Journal of Economic Theory*, 3 (1971), 66–84.

Sandmo, Agnar, "On the Theory of the Competitive Firm under Price Uncertainty," *American Economic Review*, 61 (1971), 65–73.

4

Market Equilibrium under Uncertainty

We have so far considered only the decisions of the individual. In this chapter the level of analysis shifts to market interactions and the conditions of equilibrium. The *firm* will be introduced as an agency of individuals engaged in the process of production. We continue to deal only with *event uncertainty* under "perfect markets" – ruling out *market uncertainty* with its characteristic attendant phenomena of search and of trading at non-clearing prices. But account will be taken of possibly incomplete regimes of markets; i.e., we will not always assume that each distinct state claim is, directly or indirectly, tradable. In Section 4.3, we build on Section 2.2 to flesh out the capital asset pricing model, which is the cornerstone of asset pricing in financial economics.

4.1 Market Equilibrium in Pure Exchange

In the regime of Complete Contingent Markets (CCM), where claims to a generalized consumption good C under each and every state contingency are separately tradable, as shown in Chapter 2, Equation (2.1.6), the individual's optimum position can be expressed as the Fundamental Theorem of Risk Bearing:

$$\frac{\pi_1 v'(c_1)}{P_1} = \frac{\pi_2 v'(c_2)}{P_2} = \cdots = \frac{\pi_s v'(c_s)}{P_s} \tag{4.1.1}$$

(This form of the theorem is valid only for interior solutions.[1]) At this point, we call attention to the fact that, in principle at least, all of the following elements may differ among the various individuals $j = 1, \ldots, J$:

[1] Unless otherwise indicated, it will be assumed throughout that interior (and not corner) solutions apply for all economic agents.

the probability beliefs π_s^j, the consumption quantities c_s^j, and the utility functions $v_j(c^j)$. However, the prices P_s will be the same for all market participants.

In moving from individual optimization to market equilibrium under CCM, Equation (4.1.1) must hold for each and every market participant. It follows immediately that, for any two individuals *j* and *k*, and comparing state 1 with any other state *s*:

$$\frac{\pi_s^j v_j'(c_s^j)}{\pi_1^j v_j'(c_1^j)} = \frac{P_s}{P_1} = \frac{\pi_s^k v_k'(c_s^k)}{\pi_1^k v_k'(c_1^k)} \qquad (4.1.2)$$

Thus, for each and every individual (at an interior solution), the price ratio between any two state claims will equal the ratio of expected marginal utility of incomes in the two states. As an evident corollary, if in addition individuals *j* and *k* have the same beliefs, then for all *s*:

$$\frac{v_j'(c_s^j)}{v_k'(c_s^k)} = \xi \quad \text{(a constant)}$$

In words: for any individuals *j* and *k*, the ratio of *j*'s marginal utility of contingent income to *k*'s corresponding marginal utility is the same over all states.

The other conditions required for equilibrium represent market clearing. In equilibrium under pure exchange, for each and every traded state claim the sum of the desired holdings (demand quantities) must equal the sum of the endowment amounts (supply quantities):

$$\sum_{j=1}^{J} c_s^j = \sum_{j=1}^{J} \bar{c}_s^j, \quad \text{for } s = 1, \ldots, S$$

Example 4.1: In a world of two equally probable states ($\pi_1 = \pi_2 = 1/2$), suppose there are two equally numerous types of individuals: *j* and *k*. The *j* types have endowment $\bar{C}^j = (\bar{c}_1^j, \bar{c}_2^j) = (40, 40)$ while the *k* types have endowment $\bar{C}^k = (20, 140)$. The respective utility functions are $v_j = \ln c^j$ and $v_k = (c^k)^{1/2}$. Find the equilibrium price ratio and the optimum risky consumption vectors $\hat{C}^j$ and $\hat{C}^k$.

Answer: One way of solving the system is to consider the respective demands for c_2 claims as a function of the unknown P_2. For the type *j* individuals, the Fundamental Theorem of Risk Bearing can be expressed as:

$$\frac{0.5\left(1/c_1^j\right)}{P_1} = \frac{0.5\left(1/c_2^j\right)}{P_2}$$

or:

$$P_1 c_1^j = P_2 c_2^j$$

And the budget equation is:

$$P_1 c_1^j + P_2 c_2^j = 40P_1 + 40P_2$$

Setting $P_1 = 1$ as numeraire, the type j demand for state 2 claims becomes:

$$c_2^j = 40(1 + P_2)/2P_2$$

An analogous development for the type k individuals leads to:

$$c_2^k = (20 + 140P_2)/\left(P_2^2 + P_2\right)$$

Making use of the clearing condition that $c_2^j + c_2^k = 40 + 140 = 180$, it may be verified that the equilibrium price is $P_2 = 1/2$. The associated optimal consumption vectors are $\hat{C}^j = (30, 60)$ and $\hat{C}^k = (30, 120)$. □

Instead of contingent-claims trading, more generally there might be trading of *assets* ($a = 1, \ldots, A$) at prices P_a^A, where a unit of each asset represents a fixed bundle of state-claim payoffs z_{as}. The corresponding Risk-Bearing Theorem for Asset Markets is the individual optimum condition for holdings of assets q_a, applicable under regimes of Complete Asset Markets (CAM) and even for incomplete asset-market regimes:[2]

$$\frac{\sum_s \pi_s v'(c_s) z_{1s}}{P_1^A} = \frac{\sum_s \pi_s v'(c_s) z_{2s}}{P_2^A} = \cdots = \frac{\sum_s \pi_s v'(c_s) z_{As}}{P_A^A}, \qquad c_s = \sum_a q_a z_{as}$$

[2] As indicated in Chapter 2, an asset-market regime is "complete" if the set of available assets $a = 1, \ldots, A$ allows each individual to achieve the same contingent consumption vector as under CCM. A necessary, though not sufficient, condition for complete asset markets is $A \geq S$.

At an interior optimum, this equation will hold in asset-market equilibrium for each and every economic agent. The equilibrium price ratio between asset 1 and asset a will be such that, for any pair of individuals j and k:

$$\frac{\sum_s \pi_s^j v_j'(c_s^j) z_{as}}{\sum_s \pi_s^j v_j'(c_s^j) z_{1s}} = \frac{P_a^A}{P_1^A} = \frac{\sum_s \pi_s^k v_k'(c_s^k) z_{as}}{\sum_s \pi_s^k v_k'(c_s^k) z_{1s}}$$

Notice once again that not only the asset prices P_a^A are taken as given and thus the same for all individuals, but also the asset state payoffs z_{as}. That is, there is no disagreement among individuals about what each asset will return in each and every state. So, the only possible disagreement allowed for in this simple model concerns the *probabilities* of the different states.[3]

Finally, of course, in asset-market equilibrium there must also be market clearing:

$$\sum_{j=1}^{J} q_a^j = \sum_{j=1}^{J} \bar{q}_a^j, \quad \text{for } a = 1, \ldots, A$$

Example 4.2: Under the conditions of the previous example, imagine now that the same endowments are expressed as asset holdings $\bar{q}_a$. Thus, suppose the type j endowment consists of 40 units of asset 1 with state-return vector $(z_{11}, z_{12}) = (1, 1)$ while k's endowment consists of 20 units of asset 2 with return vector $(z_{21}, z_{22}) = (1, 7)$. Find the equilibrium asset price ratio P_1^A/P_2^A and the associated optimum asset holdings $\hat{q}_a$.

Answer: This is evidently a CAM regime. Using the Risk-Bearing Theorem for Asset Markets from Section 2.1.2 of Chapter 2, and since $c_s = \Sigma_a q_a z_{as}$, following the method of the previous example we could develop the parties' demands for one of the assets as a function of its unknown price. An easier analysis suffices here, however, since we know that under the CAM condition the same consumption vectors can be attained as under the CCM assumption of the previous example. We also know from the development in Chapter 2 that the asset prices P_a^A are related to the contingent-claim prices P_s by:

$$P_a^A = \sum_s z_{as} P_s$$

[3] This is, of course, a very drastic idealization of individuals' portfolio-holding choice situations in the real world.

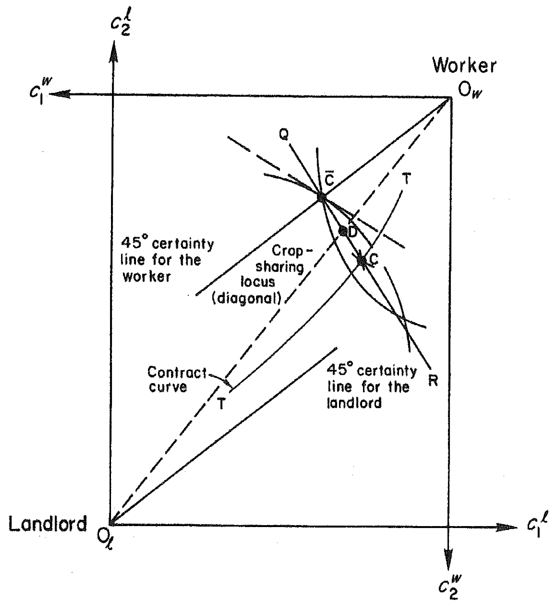

Figure 4.1. Risk sharing via state-claim trading.

Letting the state-claim price $P_1 = 1$ be the numeraire as before, it follows immediately that $P_1^A = 1(1) + 1(0.5) = 1.5$, while $P_2^A = 1(1) + 7(0.5) = 4.5$. And knowing the optimal $\hat{c}_2$ for each type of individual, the equations $c_s = \Sigma_a q_a z_{as}$ can be inverted, leading to the optimum asset holdings $\hat{Q}^j = (\hat{q}_1^j, \hat{q}_2^j) = (25, 5)$ while $\hat{Q}^k = (\hat{q}_1^k, \hat{q}_2^k) = (15, 15)$. □

4.1.1 Application to Share Cropping

In an agricultural situation, assume there are two decision makers: landlord l owns land but no labor, while worker w owns labor but no land. There are two states of the world: $s = 1$ (loss state, or "bad weather") versus $s = 2$ (non-loss state, or "good weather"). The respective agreed probabilities are π_1 and $\pi_2 \equiv 1 - \pi_1$. For the sake of the argument, assume that all productive decisions have been made so that the only choices remaining are how to share the contingent outputs and the associated risks.

Figure 4.1 is an Edgeworth box, on axes representing c_1 (income in the loss state) and c_2 (income in the non-loss state). Because of the difference in the social totals, the box is vertically elongated. The two parties' 45° certainty lines cannot coincide; it is impossible for *both* individuals to attain certainty positions, though either one could do so if the other were to bear all the risk. For individual h (where $h = l, w$), the absolute indifference-curve slope (the Marginal Rate of Substitution M^h) at any point in the Edgeworth box is given by $\pi_1 v_h'(c_1^h)/\pi_2 v_h'(c_2^h)$. Along the worker's 45° line, $M^w = \pi_1/\pi_2$ since $c_2^w = c_1^w$. But at any such point the landlord's $c_1^l < c_2^l$, hence $v_l'(c_1^l) > v_l'(c_2^l)$, so the landlord's indifference curves must all be steeper than the worker's along the latter's 45° line. Reasoning similarly for the landlord's 45° line, we see that the indifference curves must be shaped as shown in the diagram. It follows that the Contract Curve TT, connecting all the mutual-tangency points where the two parties' indifference-curve slopes are equal, must lie between the two 45° lines. This means that in equilibrium the parties will share the risk.

Assuming price-taking behavior – which would be applicable if there were a large number of competing individuals on both sides of the market – the equilibrium point would depend upon the endowment position. In the diagram, suppose this endowment is $\bar{C}$: the worker is initially receiving a fixed wage while the landlord is bearing all the risk. Under CCM there will then be trading in state claims c_1 and c_2, leading to the CCM equilibrium price ratio P_1/P_2. Specifically (4.1.2) takes the form:

$$\frac{\pi_1 v_w'\left(c_1^w\right)}{\pi_2 v_w'\left(c_2^w\right)} = \frac{P_1}{P_2} = \frac{\pi_1 v_l'\left(c_1^l\right)}{\pi_2 v_l'\left(c_2^l\right)}$$

The solution point is, of course, on the Contract Curve, as indicated by point C in the diagram. Note the following properties of the equilibrium position:

1 The parties have shared the risk, and will in fact do so regardless of the endowment position.
2 Since $c_1 < c_2$ for each party at equilibrium, it follows that $v'(c_1) > v'(c_2)$ for both l and w, and hence that $P_1/P_2 > \pi_1/\pi_2$. That is, the price of contingent income in the loss state is high *relative* to the corresponding probability. This is, of course, what we would expect: apart from the probability weighting factor, claims to income in an affluent state of the world should be cheap in comparison with claims to income in an impoverished state.

We digress now to make some remarks on contract structures. If the worker initially receives a contractually fixed wage placing him on his 45° line, we have seen that a certain amount of trading of contingent claims is necessary to achieve an efficient distribution of risk. The same holds if the landlord initially receives a fixed contractual rent placing her on her 45°

line. Since such trading is costly, we would expect to observe a tendency to avoid these extreme contractual forms. And, in fact, the worker and landlord functions are very commonly combined in owner-operated farms. An important element affecting the cost of trading is the problem of enforcement of contract. For instance, the landlord may find it difficult to control shirking by workers or the two parties might not be able to unambiguously identify which state of the world has occurred in order to distribute the contingent payoffs. The potential for trouble and disagreement on that score is all the greater since in practical situations the number of states S is large. (Consider how many states would have to be distinguished within the general category of "good weather.") One way of reducing the difficulty is a *share-cropping* arrangement in which the parties need only decide in what fixed proportions to divide the crop, whatever its size.

In Figure 4.1, the possible proportional divisions of the product would be represented by points along the main diagonal of the Edgeworth box (dashed line). In general, no such division could exactly reproduce the CCM solution along the Contract Curve TT in the diagram. Thus, Equation (4.1.2) would not be satisfied. But a point like D, on the main diagonal of Figure 4.1, may be a reasonably good approximation of the CCM solution at point C. (It would be possible to reproduce the *exact* state-claim solution by combining share cropping with side-payments; for example, if there are only two states as in the diagram, the landlord might receive x% of the crop less a side-payment of \$$y$ in each state. However, with more than two states, it would in general be necessary to have a different side-payment for each of $S-1$ distinct states, which would involve essentially the same high transaction costs as full state-claim trading.)

Alternatively, consider a CAM regime. In Figure 4.2, the endowment point $\bar{C}$ could be regarded as representing (i) the worker's initial holding $\bar{q}_a^w$ of a certainty asset a, i.e., an asset with payoffs $(Z_{a1}, Z_{a2}) = (1,1)$, which is reflected by the 45° slope of the line from 0_w to $\bar{C}$ or (ii) the landlord's initial holding $\bar{q}_b^l$ of a risky asset b, the ratio of whose returns (Z_{b1}, Z_{b2}) is reflected in the steeper slope of the line from 0_l to $\bar{C}$. We are free to choose units for each asset, so suppose that a unit of asset a is represented by the unit vector parallel to $0_w\bar{C}$ while a unit of asset b is represented by the unit vector parallel to $0_l\bar{C}$. Then the length of the line $0_l\bar{C}$ is the number of units of asset b initially held by the landlord.

Any exchange of assets by the landlord is a move back along the line $0_l\bar{C}$ and out along a line parallel to $0_w\bar{C}$. By exchanging assets in such quantities as to move to a point on the contract curve such as C^*, each party is satisfying the Risk-Bearing Theorem for Asset Markets, plus, of course, the

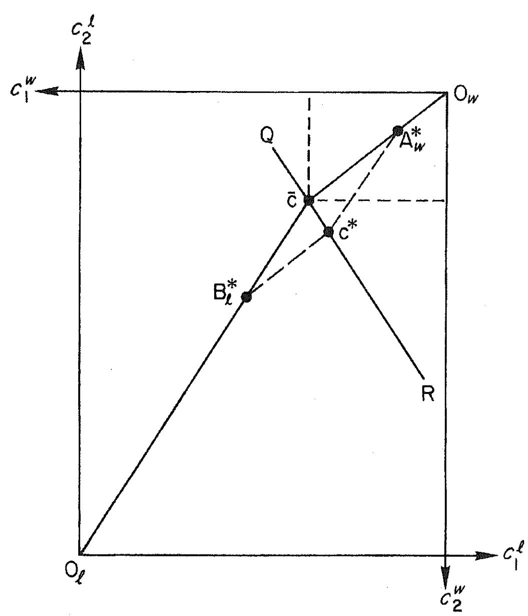

Figure 4.2. Risk sharing via asset trading.

market clearing guaranteed by the fixed size of the Edgeworth box. Once again, however, for $S > 2$ the ideal CCM solution at C* cannot in general be attained by trading in only two assets *a* and *b*.

As an interesting interpretation, we can think of this type of risk-sharing as the exchange of "equity shares" in the two parties' endowments. This interpretation will be developed further later in the chapter.

4.1.2 Application to Insurance

The Edgeworth box of Figure 4.1 can be given another interpretation: the risk-sharing that takes place there can be regarded as "mutual insurance." Indeed, all insurance is best thought of as mutual. Insurance companies, since they do not dispose of any resources other than those possessed by their owners, creditors, and policy holders, are only intermediaries in the risk-sharing process.

Once again, imagine "loss" and "non-loss" states of the world. For example, an earthquake might or might not occur. The Edgeworth box will again be vertically elongated, the social total of income being smaller in the loss state (state 1). From any endowment point like $\bar{C}$ in Figure 4.1, price-taking traders under CCM would arrive at a risk-sharing equilibrium like C* on the Contract Curve TT. As before, the absolute slope P_1/P_2 of the equilibrium market line QR exceeds the absolute slope of the dashed line representing the "fair" price ratio (equal to the probability ratio π_1/π_2). That is, claims to income in the less affluent state of the world command a relatively high price. This "social risk" helps explain why insurance is not offered at strictly fair (or "actuarial") terms.[4] But the influence of social risk depends upon a number of factors, as will be worked out in detail in what follows.

As an instructive special case, imagine there are two individuals j and k with equal initial gross incomes $\bar{c}$, each being subject to a fixed loss hazard L with the same probability p. Thus, each person faces the prospect $(\bar{c}, \bar{c} - L; 1 - p, p)$. Here, p is the probability of an independent private event. Four possible social states can be defined, according as loss is suffered by (1) neither party; (2) j only; (3) k only; or (4) both parties. The corresponding state probabilities are:

$$\pi_1 = (1 - p)^2$$
$$\pi_2 = p(1 - p) = \pi_3$$
$$\pi_4 = p^2$$

It is evident that, even under CCM, there is no possibility of risk-sharing in states 1 and 4. So the only trading that can take place will be exchanging state-2 for state-3 claims. From the symmetry of the situation, such exchange will occur in a 1:1 ratio, so that the equilibrium condition (4.1.1) takes the form, for each individual:

$$\frac{\pi_1 v'(\bar{c})}{P_1} = \frac{\pi_2 v'(\bar{c} - L/2)}{P_2} = \frac{\pi_3 v'(\bar{c} - L/2)}{P_3} = \frac{\pi_4 v'(\bar{c} - L)}{P_4}$$

where, of course, $P_4/\pi_4 > P_3/\pi_3 = P_2/\pi_2 > P_1/\pi_1$.

The equilibrium condition would not take quite so simple a form if inter-individual differences were permitted – as between the losses L^j and L^k, the loss probabilities p_j and p_k, initial incomes $\bar{c}^j$ and $\bar{c}^k$, and the utility functions $v^j(c^j)$ and $v^k(c^k)$. But it nevertheless remains true that, after trading, $c_1 > c_2$, $c_3 > c_4$ for each individual and that in equilibrium:

$$\frac{P_4}{\pi_4} > \frac{P_3}{\pi_3} \text{ and } \frac{P_2}{\pi_2} > \frac{P_1}{\pi_1}$$

4 In actual insurance practice, transaction costs place a "loading" upon the premiums offered purchasers of policies. In accordance with our assumption of perfect markets, transaction costs will be set aside here.

Any actual system of contractual insurance arrangements will only approximate the ideal results under CCM. One possible arrangement might be a mutual-insurance system with no stated premiums. Instead, policy holders would be proportionately assessed for the amounts required to match the aggregate of losses experienced. In our fully symmetrical example above, in state 1 there would be no loss and no assessment; in state 2 the assessment would be $L/2$ to each party, summing to the full L required to indemnify individual j; similarly in state 3, except that the indemnity would go to individual k; and in state 4, each party would be assessed L and indemnified L, so that no actual transfer of funds would take place. Thus, the assessment system would replicate the results of CCM. More generally, however – allowing for inter-individual differences in loss magnitudes L, loss probabilities p, endowments, etc. – such an assessment arrangement would diverge from results under CCM. But, if only the loss magnitudes varied among individuals, proportionate assessment would be similar to share cropping. An individual whose risk is x% of the total would be assessed x% of the loss *ex post*, so that the outcome would lie along the main diagonal of an Edgeworth box in four-dimensional space.

Coming closer to conventional insurance arrangements, standard practice would be to quote (say, for individual j) a fixed premium H^j to be paid into the pool regardless of which state obtains, while a fixed indemnity I^j will be receivable from the pool in either state 2 or state 4. Inability to provide for differential net payments in these two states, together with the corresponding failure in the case of individual k to distinguish state 3 from state 4, represents a serious incomplete-markets problem. Indeed, under mutual insurance the problem is an impossible one, since owing to social risk the totals of premiums paid in could not always match the totals of indemnities payable. For example, if any premiums at all are collected, should state 1 occur, there would be no losses to absorb them. But a zero premium would be absurd, providing no funds for the indemnity payments required in all other states of the world. In practice, this difficulty is avoided by having mutual-insurance pools take on a legal personality, e.g., via the corporate form. Then, premium levels H^j and H^k might be set, for example, to cover indemnities in state 4, the most adverse possibility. Should any other state actually come about, the corporation will show a "profit" that can be rebated back to its owners, the policy holders. (Alternatively, the corporation might engage in time averaging, reinvesting profits in good years to accumulate "reserves" to help cover losses in bad years and thus permit a lower level of premiums.) It will be evident that such a system is essentially equivalent to assessable premiums.

Social risk comes about whenever private risks are not perfectly offsetting. It is sometimes thought that the variability of the per-capita social risk is only a result of small numbers. If so, for pools with a sufficiently large membership N, mean income could be regarded as effectively constant. It follows that, for large N, insurance premiums would become practically actuarial (fair) – apart from transaction costs, of course.

Instead of a fixed loss L^j, assume more generally now that each individual faces a loss-probability distribution $f^j(\tilde{L}^j)$, and for simplicity suppose all the distributions are identical. Then the question is whether the per-capita loss $\tilde{\lambda} = (1/N)\Sigma_{j=1}^{N}\tilde{L}^j$ becomes approximately constant over states as N grows large. In accordance with the Law of Large Numbers, as N increases, the variance of $\tilde{\lambda}$ does decline, hence the error committed by ignoring social risk does diminish. Nevertheless, this error does *not* tend toward zero as N increases, unless indeed the separate risks are on average uncorrelated.[5]

Suppose that the individual L^j distributions all have the same mean μ and variance σ^2, and suppose, also, that the correlations between all pairs of risks equal some common r (which, of course, can be possible only for $r \geq 0$). That is, for any pair L^j, L^k:

$$E[(L^j - \mu)(L^k - \mu)] = r\sigma^2$$

The mean average loss $E[\tilde{\lambda}]$ is then just μ. The variance of the average loss is:

$$\begin{aligned}
\sigma_{\lambda}^2 &= E[(\tilde{\lambda} - \mu)^2] \\
&= \frac{1}{N^2}E\left[\left(\sum_{j=1}^{N}(L^j - \mu)\right)^2\right] \\
&= \frac{1}{N^2}\sum_{j=1}^{N}E\left[\sum_{k=1}^{N}(L^j - \mu)(L^k - \mu)\right] \\
&= \frac{1}{N^2}\sum_{j=1}^{N}[\sigma^2 + (N-1)r\sigma^2] \\
&= \sigma^2\left[\frac{1 + r(N-1)}{N}\right]
\end{aligned}$$

In the limit as N increases, the variance of per-capita loss approaches $r\sigma^2$, and thus always remains positive unless $r = 0$.

[5] See Markowitz (1959), p. 111.

We see, therefore, that social risk is not exclusively due to small numbers; it persists even with large numbers if risks are positively correlated. Somewhat offsetting this consideration is the possibility of time averaging via the accumulation of reserves. Doing so is to employ the Law of Large Numbers in a different dimension: the law tends to operate over time as well as over risks at any moment in time. If risks that are correlated at any point in time are serially uncorrelated over time, aggregated over a number of time periods, the variance of per-capita losses will diminish. (The power of the Law of Large Numbers over time will be weakened to the extent that positive serial correlation exists, that is, if high-social-loss states tend to be followed by similar high-loss states.)

Interpreting the main result of this section in terms of the language of portfolio theory, risks have a "diversifiable" element that can be eliminated by purchasing shares in many separate securities (equivalent to mutual insurance among large numbers of individuals), and an "undiversifiable" element due to the average correlation among risks. It follows then that a particular asset will be more valuable the smaller is the correlation of its returns over states of the world with the aggregate returns of all assets together – the variability of which is the source of undiversifiable risk.[6]

Social risk, therefore, provides two reasons why insurance prices may not be fair or actuarial: (i) if the number of risks in the insurance pool is small, the Law of Large Numbers cannot work very fully; (ii) if the separate risks are on average positively correlated, then even with large numbers the variance of the per-capita return does tend to diminish but does not approach zero. In either case there will still be relatively poor social states for which claims to income will command prices that are disproportionately high relative to the corresponding probabilities (with the reverse holding for relatively affluent social states).

In addition, other factors may help bring about non-actuarial terms of insurance: (1) as mentioned in footnote 4, insurance premiums are "loaded" in order to cover transaction costs, and (2) *adverse selection and moral hazard*, phenomena essentially due to information asymmetries between buyers and sellers, may tend to affect the terms of insurance transactions.

[6] In modern investment theory, the correlation of a particular security's return with that of the market as a whole – which represents the returns on all securities together – is measured by a "beta" parameter. Securities with low or, even better, negative betas tend to trade at relatively high prices. That is, investors are satisfied with relatively low expected rates of return on these assets, since they tend to provide generous returns in just those states of the world where aggregate incomes are low (and, therefore, marginal utilities are high). We analyze this topic in detail in Section 4.3.

Exercises and Excursions 4.1

1 *Complete versus Incomplete Asset-Market Equilibria*

(A) In a world of three equally probable states, with equally numerous individuals of types j and k, the endowments are $\bar{C}^j = (45, 45, 45)$ and $\bar{C}^k = (15, 67.5, 315)$. The utility functions are $v^j = \ln c^j$ and $v^k = (c^k)^{1/2}$. Verify that under CCM the equilibrium price ratios are $P_1{:}P_2{:}P_3 = 3{:}2{:}1$. Find the individual optimum positions.

(B) Suppose the same endowment positions are expressed in terms of asset holdings. Thus, j holds 45 units of asset a with state returns (1, 1, 1) while k holds 1 unit of asset b with state returns (15, 67.5, 315). Verify that the CCM equilibrium cannot be attained if the parties can exchange only assets a and b.

2 *Efficiency of Proportional Sharing*

In the landlord-worker problem, show that, if the two parties have common probability beliefs and identical utility functions $v(c^w)$ and $v(c^l)$ characterized by constant relative risk aversion R, then – for any finite number of states S – the CCM solution will lie along the main diagonal of the S-dimensional Edgeworth box. (Hence simple proportional sharing of the crop will be efficient.)

3 *Risk Sharing with μ, σ Preferences*

Suppose preferences are given by:

$$U^i = \mu(c^i) - \alpha^i \sigma^2(c^i), \quad i = w, l$$

The aggregate output in state s is y_s. The worker, individual w, is to be paid a fixed "wage" ω plus a share y of the residual $y_s - \omega$.

(A) Obtain expressions for $\mu(c^i)$ and $\sigma^2(c^i)$ in terms of ω, y, and the mean and variance of y.

(B) Write down a first-order condition for the Pareto-efficient choice of ω. Hence show that along the Pareto frontier $dU^w/dU^l = -1$.

(C) Hence, or otherwise, establish that the worker's efficient share of aggregate risk is:

$$\gamma^* = \frac{\alpha^l}{\alpha^l + \alpha^\omega}$$

(D) Is it surprising that this share is constant along the Pareto frontier?

(E) Would a similar result hold if there were M workers and N landlords?

4 *Insurance with Transaction Costs*

Suppose each individual faces the risk of a loss L, the different risks being independent (uncorrelated). Also, there are sufficiently large numbers that the per-capita risk is negligible and so insurance is offered on actuarily fair (i.e., zero profit) terms.

What would be the equilibrium insurance policy if, whenever a loss takes place, the insurance company incurs a transaction cost c?

5 *Complete Contingent Markets (CCM) with Constant Absolute Risk Aversion*

Suppose each of N individuals exhibits constant absolute risk aversion. All have the same probability beliefs. Under a CCM regime, let P_s denote the equilibrium price in state s ($s = 1, \ldots, S$).

(A) If individual i's degree of absolute risk aversion is A_i show that his optimum claims in states s and t must satisfy:

$$A_i(c_s^i - c_t^i) = \ln(\pi_s/\pi_t) - \ln(P_s/P_t)$$

(B) Hence obtain an expression for the logarithm of relative prices in terms of the average endowments in states s and t, $\bar{c}_s$ and $\bar{c}_t$.

(C) Let A^* be the harmonic mean of the degrees of absolute risk aversion, that is:

$$A^* = N\left[\sum_{i=1}^{N} \frac{1}{A_i}\right]^{-1}$$

Show that the difference between the price ratio P_s/P_t and the ratio of probabilities π_s/π_t is positive if and only if $\bar{c}_t > \bar{c}_s$.

(D) Discuss also the effect of changes in the distribution of endowments, and of an increase in A^*, upon P_s/P_t.

6 *Insurance Premiums with State-Dependent Utility*

Suppose that health risks are independently distributed for all individuals. Suppose, furthermore, that numbers are sufficiently large so that insurance against a deterioration in health is offered on actuarially fair terms.

(A) Suppose an individual has a utility function $v(c, h) = (ch)^{\frac{1}{2}}$, his health level h being either h_b or h_g (where $h_b < h_g$). Would this individual wish to buy insurance against bad health?

(B) Suppose bad health also reduces income by 50%. Would the individual now wish to buy insurance? If not necessarily, under what conditions would this be the case? Would the individual ever buy enough insurance to completely offset his income loss?

(C) Another individual has a utility function $\bar{v}(c, h) = \ln(ch)$. Confirm that there is a function $u(\cdot)$ such that:

$$\bar{v}(c, h) = u(v(c, h))$$

Hence draw a conclusion as to which individual is more risk averse.

(D) Repeat (A) and (B) with the new utility function.

(E) What can you say about an individual who is more risk averse than both these individuals but again has the same indifference map?

[HINT: You might refer back to Section 2.3 in Chapter 2 before attempting to answer this question.]

4.2 Production and Exchange

The previous section examined regimes of complete and incomplete markets in a pure-exchange economy with a single generalized consumption good *C*. We showed that the analysis of market equilibrium can be interpreted in terms of an S-dimensional Edgeworth box diagram with one axis for each state claim c_s. Just as goods are allocated efficiently in the traditional commodity-market equilibrium under certainty so, under uncertainty, a complete-market equilibrium (i.e., the CCM case where there are markets in all *S* states, or else the CAM case where an equivalent regime of asset markets exists) distributes social risk efficiently.

We will now generalize this conclusion. In Section 4.2.1 we show that, even in a world of production, and allowing for many commodities *G* as well as any number of states *S*, a complete-market equilibrium allocation is Pareto efficient. Of course, the assumption of complete markets is a strong one. Section 4.2.2 takes up production decisions in a special regime of incomplete markets called a "stock market economy." Conditions are derived under which shareholders unanimously agree upon value maximization as the objective of the firm.

4.2.1 Equilibrium with Production: Complete Markets

Suppose there is a single commodity (corn), a single firm, and two states of the world (rain or no rain). By varying the production process the firm

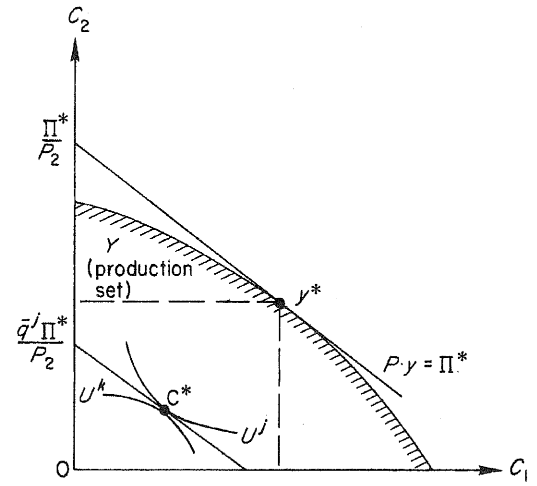

Figure 4.3. Equilibrium and efficiency.

chooses a state-distributed vector of production levels $y = (y_1, y_2)$. The set of possible production vectors or "production set" Y is illustrated in Figure 4.3.[7] We assume that this set is convex.

To illustrate, suppose that when rows of corn are planted close together there will be an especially big harvest *if* the weather is hot. However, if the weather is cool, a better yield is obtained by planting at lower densities. As a special case, let:

$$y_1(x) = 20x, \quad y_2(x) = 100x - 10x^2$$

be the state-dependent outputs associated with a crop density of x. Eliminating x yields the production frontier:

$$y_2 - 5y_1 + \frac{y_1^2}{40} = 0$$

The production set Y consists of the production vectors $y = (y_1, y_2)$ on or inside the production possibility frontier.

Given state-claim prices P_1 and P_2, the profit of the firm (since there are no purchased factors of production) is the revenue $P_1 y_1 + P_2 y_2$. Writing the price vector as $P = (P_1, P_2)$, a profit-maximizing price-taking firm chooses y^* and hence profit level Π to satisfy:

$$\Pi^* = P \cdot y^* \geq P \cdot y, \quad \text{for all } y \in Y$$

[7] In a more complete model, inputs would be purchased at $t = 1$ and output produced at $t = 2$, as in Exercise 3 at the end of this section.

(Here and henceforth, we will assume a unique solution exists.) The isoprofit line $P \cdot y = \Pi^*$ is also depicted in Figure 4.3 along with the profit-maximizing production vector y^*.

We now examine the consumer-shareholders' demands for state-contingent corn. With just two individuals j and k, we can illustrate via the Edgeworth box formed in Figure 4.3 with corners at the origin and at y^*. Suppose the two proportionate shareholdings in the firm are $\bar{q}_f^j$ and $\bar{q}_f^\kappa \equiv 1 - \bar{q}_f^j$.[8] With 0 as the origin for j, his budget constraint is:

$$P_2 c_1^j + P_2 c_2^j = \bar{q}_f^j \Pi^* \qquad (4.2.1)$$

This is, as depicted, parallel to $\Pi^* = P \cdot y$ but with vertical intercept $\bar{q}_f^j \Pi^* / P_2$. Then the risk-bearing optimum for individual j is point C* where $U^j(c_1^j, c_2^j)$ is maximized subject to the budget constraint (4.2.1).

Of course, this budget line can also be viewed as the budget line for individual k, using the point y^* as her origin. As depicted, the state-claim prices P_1 and P_2 are such that aggregate consumption equals aggregate production in each state – markets are cleared. Since neither individual can do any better without making the other worse off, the competitive equilibrium is Pareto efficient. And, specifically, price-taking behavior and profit maximization result in an allocation in which the Marginal Rate of Substitution of state-1 claims for state-2 claims, $M \equiv -dc_2/dc_1$, is, for each individual, equal to his or her marginal rate of productive transformation $-dy_2/dy_1$.

Note that, just as in the pure-exchange case, the analysis is formally equivalent to the traditional certainty model. However, whereas in the standard certainty model y is a vector of outputs of different commodities, y here becomes a state-contingent output vector. This suggests a way of demonstrating the efficiency of a CCM regime in a much more general setting. All we have to do is to show that the description of individual optimization and market clearing is formally equivalent to that in the traditional certainty model, where the efficiency of competitive equilibrium is a standard result.

As a first step, let us briefly review the traditional general-equilibrium model. Suppose the economy consists of I individuals indexed by i, F firms indexed by f, and G commodities indexed by g. Firm f chooses a vector of inputs and outputs $y^f = (y_1^f, \ldots, y_G^f)$ from the set Y^f of feasible vectors.

[8] In Chapter 2 and elsewhere, q_a^i, denoted the *number of units* of asset a held by individual i. Here, the total number of shares in firm f is defined as unity, so each individual's shareholding q_f^i will represent a fractional number of units.

We assume that a firm can always choose not to produce so that the zero vector is in Y^f. A positive y_g^f indicates that the firm produces more than it purchases of commodity g while a negative y_g^f indicates that the firm is a net purchaser of the commodity. With commodity prices $P = (P_1, \ldots, P_G)$, the firm chooses y_*^f to maximize profit,[9] that is:

$$y_*^f \text{ solves } \underset{y^f}{\text{Max}}\{P \cdot y^f | y^f \in Y^f\}$$

Since firm f can always choose not to produce, maximized profit Π^f is non-negative.

Each individual i, with utility function $U^i(c^i)$ where $c^i \equiv (c_1^i, \ldots, c_G^i)$, has some initial endowment of commodities $\omega^i \equiv (\omega_1^i, \ldots, \omega_G^i)$ and owns a proportion $\bar{q}_f^i$ of firm f. The individual then chooses c_*^i to maximize utility subject to the constraint that total expenditure on commodities does not exceed the value of his endowment plus profit shares. That is, c_*^i is the solution of:

$$\text{Max}\left\{U^i(c^i) \middle| P \cdot c^i \leq P \cdot \omega^i + \sum_f \bar{q}_f^i \Pi^f\right\}$$

For P to be a market equilibrium price vector, supply must equal demand in every market,[10] that is:

$$\sum_f y_*^f + \sum_i \omega^i = \sum_i c_*^i$$

And on the assumption that each individual, regardless of his consumption vector, always strictly prefers more of some commodity we know also – from the first theorem of welfare economics (Debreu, 1959) – that this market equilibrium is Pareto efficient.

We now seek to extend this result to include uncertainty. Instead of just G markets, one for each commodity, we introduce markets for each

[9] For example, the production set of the neoclassical firm producing Q units of output with capital and labor according to the production function $Q = \Phi(K, L)$ is:

$$Y^f = \{(Q, -K, -L) | Q \leq \Phi(K, L),\ K, L \geq 0\}$$

With prices $(P_1, P_2, P_3) = (p, r, w)$, the firm chooses $y^f \in Y^f$ to maximize:

$$P \cdot y^f = pQ + r(-K) + w(-L) = pQ - (rK + wL).$$

[10] As a more general statement (allowing also for corner solutions), supply must at least equal demand, and the price must be zero for any market in which there is excess supply.

commodity in each state of the world – $G \times S$ markets in all. The price P_{gs} is then the price of purchasing a unit of commodity g for delivery if and only if state s occurs.

Each firm makes a decision as to its purchases and sales in each state. For example, a firm producing commodity 1 using commodities 2 and 3 as inputs might have a state-dependent production function:

$$y_{1s}^f = \phi_s^f(-y_{2s}^f, -y_{3s}^f)$$

The firm then contracts to purchase contingent inputs and deliver contingent outputs in order to maximize its profit:

$$\Pi^f = (P_{1s}y_{1s}^f + P_{2s}y_{2s}^f + P_{3s}y_{3s}^f) = P \cdot y^f$$

In general, just as in the certainty case, firm f chooses y_*^f so that:

$$\Pi^f = P \cdot y_*^f \geq P \cdot y^f, \quad y^f \in Y^f$$

In the same way, individual i with endowment ω^i and utility function $v^i(c_s^i)$, where $c_s^i = (c_{1s}^i, \ldots, c_{Gs}^i)$, chooses his final consumption bundle to maximize expected utility:

$$U^i(c^i) = \sum_s \pi_s^i v^i(c_s^i)$$

This maximization is, of course, subject to the budget constraint:

$$P \cdot c^i \leq P \cdot \omega^i + \sum_f \bar{q}_f^i \Pi^f$$

Viewed in this way, it is clear that any conclusions about the certainty model must carry over. In particular the equilibrium allocation must lead to a Pareto-efficient allocation of risk bearing.

Several aspects of the equilibrium are worthy of note:

1 Under complete markets, the efficient allocation is achieved when firms simply maximize profit (net market value). Profit being deterministic rather than stochastic, there is no need to consider expected profit or to adjust for some concave function of profit representing owner risk aversion. The point is that, at the time a production decision is made, the firm can also complete all sales of its *contingent* outputs at the ruling state-claim prices. Net earnings or profit can then be handed

over to stockholders. Of course, actual input and output levels will be uncertain. However, the market equilibrium state-claim prices already provide the correct adjustments for the risk factor, so owners are best served when the chosen production vector maximizes net market value. It follows also that stock markets and stock trading have no special role. Indeed, no one has any incentive to trade his initial asset endowment except to make final consumption purchases.

2 Different consumers need not have the same beliefs about the likelihood of different states. The CCM equilibrium is efficient with respect to beliefs *actually held.*

3 All trading in this economy takes place prior to learning which state s has occurred. This raises the question as to whether any individual might wish to engage in posterior trading after the state of the world is revealed. To answer this question, suppose that all prior trading takes place on the anticipation that markets will *not* reopen after the state is revealed. Consumer i will then initially select his state-distributed consumption vector so that his Marginal Rate of Substitution of commodity 1 for commodity g in a particular state is equal to the price ratio:

$$\frac{\dfrac{\partial U^i}{\partial c_{1s}}}{\dfrac{\partial U^i}{\partial c_{gs}}} = \frac{\pi_s^i \dfrac{\partial v^i}{\partial c_{1s}}}{\pi_s^i \dfrac{\partial v^i}{\partial c_{gs}}} = \frac{P_{1s}}{P_{gs}} \tag{4.2.2}$$

Now suppose that the state is revealed to be s, and that, unexpectedly, markets do in fact reopen for posterior trading. If the state-s market-price *ratios* among the G commodities were to remain unchanged from the prior ratios of (4.2.2), individual i, now with utility function $v^i(c_s)$, will wish to trade so that his new Marginal Rate of Substitution of commodity 1 for commodity j will equal the unchanged price ratio, that is:

$$\frac{\dfrac{\partial v^i}{\partial c_{1s}}}{\dfrac{\partial v^i}{\partial c_{gs}}} = \frac{P_{1s}}{P_{gs}} \tag{4.2.3}$$

Comparing (4.2.2) and (4.2.3), it follows immediately that individual i will have no need to trade again. Thus, the prior-trading price ratios for the state-s commodity claims dictate a posterior equilibrium in which no retrading occurs.[11] This proposition will play an important role when we consider the topic of speculation in Chapter 6.

So far, we have considered only the *unanticipated* opportunity for retrading after the state is revealed. If the possibility of posterior trading is indeed *anticipated*, consumers must form beliefs about prices in future states ("future spot prices"). Our argument indicates that, as long as everyone believes that relative future spot prices in the state that actually occurs will be the same as relative prior contingent prices, there will be no gains to multiple rounds of trading. Moreover, such beliefs will be self-fulfilling – the market-clearing future spot price ratios will indeed equal the corresponding contingent price ratio.[12]

4 We have implicitly been assuming that production, consumption, and exchange all occur at a single date in time. This also is an expositional simplification that can easily be generalized. The same equation format for a CCM regime can allow for specifying the commodity, the state, and also the date. The price P_{gst} is then the price paid, in the current period, for commodity g to be delivered at time t in the eventuality that state s occurs. As in the one-period model, firm f chooses $y^f = (y^f_{111}, \ldots, y^f_{GST})$ from its production set Y^f to maximize $P \cdot y^f$ – which is the net present value of the production plan or, more simply, the *value of the firm* at today's prices. It should be noted that the firm's plan will, in general, be a contingent plan. That is, some farther future decisions may be contingent upon some still uncertain nearer-future events.

Exercises and Excursions 4.2.1

1 Exchange Equilibrium with Complete Markets

Consider an economy with two states. Every individual has the same utility function $v(c) = \ln(c)$ and believes that state 1 will occur with probability π.

(A) Show that the CCM equilibrium price ratio satisfies:

$$\frac{P_1}{P_2} = \frac{\pi}{1-\pi}\left(\frac{y_2}{y_1}\right)$$

where y_s is the aggregate endowment of claims in state s.

[11] It is left to the reader to confirm that no firm will wish to change its production plan in state s either.

[12] Beliefs may then be called "rational," as in the common but confusing term "rational expectations equilibrium" – a more accurate term would be "self-fulfilling beliefs equilibrium." In the absence of such concordant beliefs about future spot prices, those agents whose beliefs were incorrect will wish to re-enter the market. This in turn opens up opportunities for sophisticated traders to "speculate." We shall have more to say on this topic in Chapter 6.

(B) If the price of a riskless asset yielding 1 unit in each state is 1, show that the state-claim prices are:

$$P_1 = \frac{\pi y_2}{\pi y_2 + (1-\pi)y_1} \quad \text{and} \quad P_2 = \frac{(1-\pi)y_1}{\pi y_2 + (1-\pi)y_1}$$

(C) Suppose there are two types of asset in the economy. A unit of the riskless asset (asset 1) pays off 1 in each state and has market price $P_1^A = 1$. A unit of the risky asset (asset 2) returns $z_{21} = ½$ in state 1 and $z_{22} = 2$ in state 2. Aggregate supplies of the two assets are q_1 and q_2. If the two states are equally likely, show that the price of the risky asset is:

$$P_2^A = \frac{5q_1 + 4q_2}{4q_1 + 5q_2}$$

(D) Suppose initially there are no units of the risky asset. However, there is a technology that will create units of the risky asset at the cost of one unit of the riskless asset. There is free entry into the industry.

What will be the equilibrium price of the risky asset? What will be the equilibrium supply of the risky asset, expressed as a proportion of the equilibrium supply of the riskless asset?

2 *Complete-Market Equilibrium with Production*

Consider an economy in which a single firm produces a single commodity. There are two states of the world, state 1 and state 2. The n-th plant in the firm can produce any state-dependent output vector $y = (y_1, y_2)$ lying in the production set $Y^n = \{(y_1, y_2) | y_1^2 + y_2^2 \le 2\}$. There are two individuals in the economy, each of whom has a 50% share in the firm, and who behave as price takers.

(A) If there are two plants, confirm that the aggregate production set is $Y = \{(y_1, y_2) | y_1^2 + y_2^2 \le 8\}$. Hence, or otherwise, show that with state-claim prices $(P_1, P_2) = (1,1)$ the firm will produce an output vector $(y_1^* + y_2^*) = (2, 2)$.

(B) If individual 1 believes that state 1 will occur with certainty, explain why, *at the above prices*, his final consumption vector is $(c_1^1, c_2^1) = (2, 0)$.

(C) If the second individual believes that state 2 will occur with certainty, confirm that $P = (1,1)$ is the complete-market equilibrium price vector.

Newberry, David M., "Futures Markets, Hedging and Speculation," in *The New Palgrave Dictionary of Economics*, Steven N. Durlauf and Lawrence E. Blume (eds.), 2nd edition, vol. 3, pp. 521–526, New York: Palgrave MacMillan, 2008.

Novos, Ian E. and Waldman, Michael, "The Emergence of Copying Technologies: What Have We Learned?", *Contemporary Policy Issues*, 5 (1987), 34–43.

Radner, Roy, "Competitive Equilibrium under Uncertainty," *Econometrica*, 36 (1968), 31–58.

———, "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica*, 47 (1979), 655–678.

Shleifer, Andrei and Lawrence H. Summers, "The Noise Trader Approach to Finance," *Journal of Economic Perspectives*, 4(2) (1990), 19–33.

Smith, Vernon L., "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72 (1982), 923–955.

Verrecchia, Robert E., "Information Acquisition in a Noisy Rational Expectations Economy," *Econometrica*, 50 (1982), 1415–1430.

Working, Holbrook, "Futures Trading and Hedging," *American Economic Review*, 43 (1953), 314–343.

———, "New Concepts Concerning Futures Markets and Prices," *American Economic Review*, 52 (1962), 431–459.

7

Strategic Uncertainty and Equilibrium Concepts

For the most part, the analysis to this point has dealt with *event uncertainty*. Individuals were mainly uncertain about nature's choice of state of the world. In the following chapters the focus shifts to *strategic uncertainty*, where the best course of action for individual A depends upon individual B's choice, and vice versa. So the main risks that a person has to deal with concern the actions and reactions of others. A first step is the choice of an equilibrium concept for such an environment, which turns out to be a subtle and still controversial issue. As usual, our discussion will not attempt to address formal issues of existence or uniqueness of equilibrium. Our aim instead is to provide an intuitive interpretation of the key ideas.

7.1 Dominant Strategy

We begin with the Prisoners' Dilemma, a game that is easy to analyze. The story behind this game is as follows. Two accomplices have been arrested on suspicion of committing a major crime. The prosecutor does not have sufficient evidence to convict them of this crime. Without a confession from at least one of the two accomplices the prosecutor can only send them to prison for one year for the lesser charge of illegal possession of weapons. The two accomplices (or prisoners) are locked up in different cells and cannot communicate with each other. The prosecutor approaches each of them separately and says, "If you confess and your friend does not then I will drop all charges against you. On the other hand, if your friend confesses and you do not then you will do 10 years. If you both confess then I'll see to it that you get parole after 5 years in prison. Think hard about what you want to do and let me know tomorrow morning." Each prisoner knows that if both of them remain silent and do not confess then each will be sentenced

Table 7.1. *Prisoners' Dilemma*

		Prisoner k	
		Defect x_1^2	Cooperate x_2^2
Prisoner j	Defect x_1^1	−5, −5	0, −10
	Cooperate x_2^1	−10, 0	−1, −1

to only one year of prison on firearms possession charges. What do you think the prisoners' will do?

Assume that each prisoner acts purely in his own self-interest and seeks to minimize his time in prison. Let us call the two choices that each prisoner faces Defect (if he confesses) and Cooperate (if he does not confess).[1] The consequence (i.e., the number of years in prison) for each prisoner depends not only on his own choice but also on the choice of his accomplice. The consequences are shown in Table 7.1. If, for example, prisoner j Defects and prisoner k Cooperates, then j spends 0 years in prison and k spends 10 years. This is represented as (0, −10) in Table 7.1. (We represent payoffs of games so that higher numbers are better.)

First, consider this game from prisoner j's viewpoint. Suppose that prisoner k were to Defect; then prisoner j is sentenced to only 5 years if he Defects whereas he gets 10 years if he Cooperates. Suppose, instead, that prisoner k were to Cooperate; then prisoner j spends no time in prison if he Defects whereas he spends 1 year in prison if he Cooperates. Thus, Defect is the best course of action for prisoner j, *regardless* of prisoner k's choice of action. In the language of game theory, Defect is a *strictly dominant strategy* for prisoner j. A symmetric argument establishes that Defect is a strictly dominant strategy for prisoner k as well. For each prisoner, Cooperate is a *strictly dominated strategy* – it is strictly dominated by Defect.

Thus, each prisoner will Defect and end up spending 5 years in prison. If, instead, each prisoner had selected Cooperate then each would spend only 1 year in prison. However, a choice of Defect by both players is the only strategically stable outcome. To convince yourself that this is the case, suppose that the two prisoners can communicate with each other. They meet and agree to Cooperate so that each is sentenced to only 1 year in prison. After reaching this non-binding agreement they go back to their cells. What

[1] By not confessing, a prisoner Cooperates with his accomplice. Confessing betrays the accomplice; it is an act of defection.

Table 7.2. *Prisoners' Dilemma* (general payoffs) (e > f > g > h)

		Defect x_1^2	Cooperate x_2^2
Defect	x_1^1	g, g	e, h
Cooperate	x_2^1	h, e	f, f

do you think each prisoner will do when he meets the prosecutor the next morning? Each prisoner, even if he believes that his accomplice will stick to the agreement, has an incentive to deviate from Cooperate to Defect – better to spend no time rather than 1 year in prison. And if the prisoner believes that his accomplice may renege on their agreement and Defect, then that makes Cooperate all the more attractive.

There are other strategic situations which have a payoff structure similar to the Prisoners' Dilemma. Consider two firms that compete in a market with completely inelastic demand. Each firm chooses between a Low Price and a High Price. As demand is inelastic, the total quantity sold by the two does not change with price. Thus, each firm's profit is higher when both choose High Price (which corresponds to Cooperate) than if they both choose Low Price (which corresponds to Defect). However, if one chooses Low Price and the other High Price, then the former firm corners the market and makes a greater profit than if both firms had selected High Price.

Table 7.2 gives the general form of the Prisoners' Dilemma. The inequalities $e > f$ and $g > h$ on the payoffs ensure that Defect is a strictly dominant strategy. The inequality $f > g$ implies that payoff obtained when the two players choose (Defect, Defect) is Pareto dominated by the payoff under (Cooperate, Cooperate). We shall return to this game in Chapter 11, where we investigate whether repeated interactions between the same two players increases the possibility of cooperation.

In most games, players do not have a strictly dominant strategy, and it is less obvious how they should play. We turn to this issue in the remainder of this chapter.

7.2 Nash Equilibrium

In a *coordination game*, the parties' interests are somewhat parallel. A specific example known as "Stag Hunt" originates in a situation presented by the philosopher Jean-Jacques Rousseau. Two hunters can either hunt hare on their own or cooperate to hunt stag. A stag is more difficult to hunt and

Table 7.3. *Stag Hunt*

		Player k	
		Stag x_1^k	Hare x_2^k
Player j	Stag x_1^j	10, 10	0, 7
	Hare x_2^j	7, 0	6, 6

requires the combined efforts of the two; it also provides more meat for each hunter than a hare. This is reflected in the payoffs in Table 7.3, where if both players (hunters) j and k select Stag they each get a payoff of 10, whereas if each player chooses Hare, the payoff for each is 6. If player j chooses Stag and player k chooses Hare then j gets 0 – he does not succeed in hunting any game – whereas k gets 7.[2] Thus, the two parties both gain by coordinating their activities on a stag hunt. But a trust dilemma arises here. Even if the two hunters were to agree to hunt stag, can each trust the other not to stray from this decision? To quote from Rousseau's *A Discourse on Inequality* (1755):

> If it was a matter of hunting a deer, everyone well realized that he must remain faithfully at his post; but if a hare happened to pass within the reach of one of them, we cannot doubt that he would have gone off in pursuit of it without scruple and, having caught his own prey, he would have cared very little about having caused his companions to lose theirs.

The decision is further complicated by the fact that payoff from Hare (7 if the other player chooses Stag, 6 if the other player chooses Hare) is much less variable than the payoff from Stag (either 10 or 0).

Approaching this problem in terms of game theory, we can view each player as choosing an action without knowledge of the other player's action choice. In effect, the players can be thought of as choosing their actions simultaneously. We can depict this situation with a game tree in Figure 7.1. Player j is represented here as having the first move, but player k must make a decision without knowing j's choice. This is represented by means of the dashed line called the "information set" connecting the two decision nodes of player k. Player j moves first, and then player k gets to move. Player k does

[2] When a player is the only one hunting hare, his payoff is slightly higher than when both hunt hare (7 instead of 6).

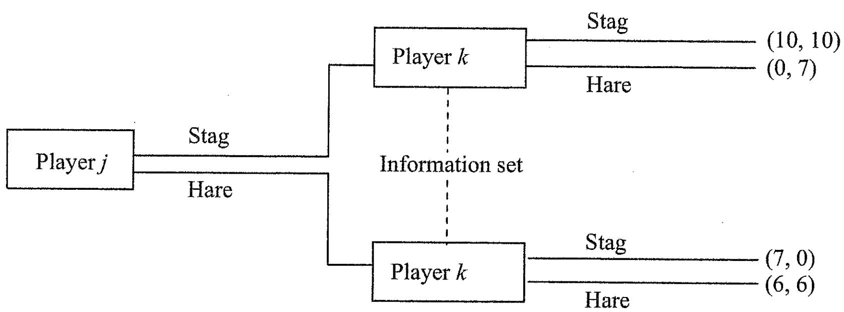

Figure 7.1. Game tree of Stag Hunt.

not know which of the two decision nodes in his information set has been reached; i.e., he does not know player j's action.[3]

The most frequently used solution concept for games is called the Nash non-cooperative solution, or *Nash equilibrium.*[4] The key idea is that there is an equilibrium when, given the strategies of all the other players, each single participant finds that his own strategy is (at least weakly) a best response to their choices. Thus the Nash equilibrium is a "no regret" equilibrium.

Our simple example is a two-player game in which each player simultaneously chooses one of two actions. With simultaneous play, each player is assumed to have made a hypothesis about the strategies of his opponents. His own strategy is then a best response to the others' hypothesized strategies. If, for each player, the chosen action coincides with what the other players have hypothesized about his strategy, a Nash equilibrium exists.[5]

In the Stag Hunt game, it is a best response to always match the action of the other player. If player k, say, were to choose Stag, the best response of player j is also Stag since it yields a payoff of 10 while Hare gives only 7 to player j. Similarly, if j were to choose Stag, k's best response is Stag. Hence, the action or strategy pair (Stag, Stag) is a Nash equilibrium in this game. In this equilibrium the parties achieve the mutually preferred outcome (10, 10).

[3] We could also have drawn this game tree with the decision nodes of the two players exchanged. That is, player k moves first and player j has two decision nodes within an information set. This game tree represents exactly the same strategic situation.

[4] Nash (1951). The Nash equilibrium is a generalization of a solution to the oligopoly problem that goes back to Cournot (1838).

[5] Note that the Nash equilibrium is *not* justified by appealing to some plausible dynamic process. Rather, it is a state of affairs in which, if it were somehow to come about, no party would unilaterally want to revise his action.

A similar argument establishes that (Hare, Hare), another mutual best response pair of actions, is also a Nash equilibrium. In this second Nash equilibrium, the parties achieve the inferior outcome (6, 6).

Unlike the Stag Hunt game, many games do not have a Nash equilibrium in simple actions or *pure strategies*. In such situations, a fruitful approach is to extend the range of choice beyond the pure strategies available to each player so as to consider *mixed strategies* as well; that is, the set of probabilistic combinations of the available pure strategies. If there are only two pure strategies, the complete set of pure and mixed strategies available to player i can be expressed (in analogy with the "prospect notation" of chapter 1) as:

$$\bar{X} = \left\{\left(x_1^i, x_2^i; \pi^i, 1-\pi^i\right) \middle| 0 \leq \pi^i \leq 1\right\}$$

where π^i = probability that player i chooses x_1^i.

More generally, if player i has a set of A^i feasible pure strategies $X^i = \{x_1^i, \ldots, x_{A^i}^i)$, then player i's complete set of strategies (the set of probability vectors over these pure strategies) can be expressed as:

$$\bar{X}^i = \left\{\left(x_1^i, \ldots, x_{A^i}^i; \pi_1^i, \ldots, \pi_{A^i}^i\right) \middle| 0 \leq \pi_a^i \leq 1 \text{ and } \sum_{a=1}^{A^i} \pi_a^i = 1\right\}$$

where π_a^i is the probability that player i chooses strategy π_a^i.

Returning to the Stag Hunt game tree, recall that there are two Nash equilibria in pure strategies. Either the players coordinate on strategy 1 (Stag) or on strategy 2 (Hare). In each case, either party acting alone can only lose by changing to a different action. But now there is also an equilibrium in mixed strategies. The following condition provides a technique for locating a Nash equilibrium in which at least one player uses a mixed strategy:

Suppose player i has chosen a mixed strategy. For his mixed strategy to be part of a Nash equilibrium, player i must then be indifferent – given the chosen strategies (pure or mixed) of the other players – among all of the pure strategies entering with non-zero probability into his own mixed strategy.

In the Stag Hunt game, suppose the players have chosen respective mixed strategies $(\pi^i, 1-\pi^i)$, $i = j, k$. We now ask when player j will be indifferent between the pure strategies 1 (Stag) and 2 (Hare). If he chooses Stag, his expected payoff is:

$$\pi^k(10) + (1-\pi^k)(0) = 10\pi^k$$

Table 7.4. *Chicken*

			Player k	
			Coward x_1^k	Hero x_2^k
Player j	Coward	x_1^j	4, 4	0, 8
	Hero	x_2^j	8, 0	−6, −6

If he chooses Hare, his gain is:

$$\pi^k(7) + (1 - \pi^k)(6) = 6 + \pi^k$$

Player j will be indifferent between Stag and Hare if and only if player k's probability mixture is $(\pi^k, 1 - \pi^k) = (2/3, 1/3)$. Given the symmetry of the game, player k's expected gains are equal when player j's mixture is $(2/3, 1/3)$ also. Evidently, this strategy pair is the only mixed-strategy Nash equilibrium.

Note that while the two pure strategy Nash equilibria here are *strong*, meaning that a player who unilaterally switches to any other strategy will end up actually worse off for having done so, the mixed strategy Nash equilibrium is *weak*. In fact, as follows directly from the condition stated above for finding the mixed strategy solution, if all other parties are playing in accordance with the mixed strategy Nash equilibrium then *any single player could equally well have chosen any of the pure strategies entering into his Nash equilibrium mixture* – or, indeed, any other mixture of them as well. More generally, a Nash equilibrium in pure strategies may be either strong or weak, but a Nash equilibrium in mixed strategies is always weak.

We now consider an alternative payoff environment, the famous game of Chicken[6] (Table 7.4), again under the assumption of simultaneous play. For instance, if j chooses Coward and k chooses Hero, the payoffs to j and k are 0 and 8 respectively.

In the Chicken game, there are once again two pure-strategy Nash equilibria, but in this case they are asymmetrical – at the off-diagonal cells (x_1^j, x_2^k) and(x_2^j, x_1^k). There is a mixed strategy Nash equilibrium as well. Given the specific payoffs of Table 7.4, the equilibrium mixed strategy is symmetrical: each player chooses strategy 1 (Coward) and strategy 2 (Hero) with probabilities 0.6 and 0.4, respectively. At the mixed strategy Nash

[6] In the biological literature, the game of Chicken is known as Hawk-Dove (Maynard Smith, 1976).

equilibrium each player's expected return is 2.4, intermediate between the Hero payoff of 8 and the Coward payoff of 0 at each of the pure-strategy Nash equilibria.

The "normal form" matrices of Tables 7.1, 7.3, and 7.4 describe the *payoff environments* of Prisoners' Dilemma, Stag Hunt, and Chicken, respectively. To represent other aspects of the game such as the *procedural rules* (for example, whether the players move simultaneously or in sequence, and if in sequence who moves first) and the *information or beliefs* that the different parties possess, one turns to the game tree or "extensive form."[7] In the next three sections of the chapter we will describe how these procedural and informational aspects of the problem affect possible solutions of a game. And, in particular, we will explore how they provide possible ways of separating more plausible Nash equilibria from those that are less plausible.

Exercises and Excursions 7.2

1 Tender Trap

Another example of a co-ordination game is Tender Trap (Hirshleifer, 1982). The Dvorak typewriter keyboard is, it has been claimed, ergonomically superior to the currently standard "Qwerty" arrangement. But having settled on the current standard keyboard, largely by historical accident, now manufacturers are supposedly reluctant to produce Dvorak keyboards so long as almost all typists are trained on Qwerty, while typists do not want to train on Dvorak when almost all keyboards are Qwerty.[8] Even the inferior keyboard as a matched choice is superior to failing to coordinate at all.

		Player k (typist)	
		Dvorak x_1^k	Qwerty x_2^k
Player j (manufacturer)	Dvorak x_1^j	10, 10	4, 4
	Qwerty x_2^j	4, 4	6, 6

The above table shows the payoffs to the two players – manufacturer and typist – in Tender Trap. Draw the game tree corresponding to this payoff

[7] Figure 7.1 is an example of a game tree.

[8] Liebowitz and Margolis (1990) claim that this story is mythical and that the Dvorak keyboard is not superior to Qwerty.

table. Find all pure strategy Nash equilibria. Is there a mixed strategy Nash equilibrium in this game?

Tender Trap illustrates the binding force of convention (of having an agreed rule) even allowing for the possibility that the convention is not ideal. We tacitly agree upon many conventions to order our daily lives – rules of the road, rules of language, rules of courtesy. Although better rules might well have been arrived at, it is hard to change a settled convention.

7.3 Subgame-Perfect Equilibrium

While the Nash equilibrium concept remains at least a preliminary guide, frequently there are multiple Nash equilibria. Not only do multiple Nash equilibria create difficulties when it comes to prediction, they also pose problems for the theory itself. A condition of equilibrium is that each player's choice be a best response to the strategy of his opponents. But, if Nash equilibrium is not unique, how will a player know whether a given strategy choice on his part is a best reply when the opponents may be choosing among several different Nash equilibria strategies? Owing to such problems, much effort has gone into "refining" the Nash equilibrium concept (Selten, 1965, Selten, 1975; Myerson, 1978; Kreps and Wilson, 1982; Kohlberg and Mertens, 1986; Grossman and Perry, 1986; Cho and Kreps, 1987). Two widely accepted refinements of Nash equilibrium will be examined in this and the next section.

Consider the following game. One firm, the "entrant," moves first by deciding whether or not to invade a market now occupied solely by an "incumbent" firm. If she chooses to enter, the entrant will quote a price lower than that previously ruling. The incumbent must respond in one of two ways. He can (i) match the entrant's price and hence share the market, or (ii) quote a price still lower than hers so as to drive out the new competitor. In the latter case, the incumbent's profits are further reduced while the entrant suffers a loss.

Figure 7.2 depicts the game tree of this game; Table 7.5 depicts the same game in tabular or normal form. In these diagrams, the first number is the incumbent's payoff and the second is the entrant's payoff; thus, if the entrant stays out the incumbent's payoff is 6 and the entrant's payoff is 0. Looking at the table, there are two Nash equilibria in pure strategies,[9] indicated by the asterisks. (1) If the incumbent is going to choose Undercut, the entrant's best response is to choose Out. And if she chooses Out, the

[9] Since this is a sequential-move game, we need not consider mixed strategies.

Table 7.5. *Entry game*

		Player 2 (entrant)	
		Enter	Out
Player 1 (incumbent)	Match	2, 2*	6, 0
	Undercut	0, −1	6, 0*

incumbent loses nothing by being prepared to Undercut. (2) On the other hand, if the entrant chooses Enter, the incumbent's best response is Match. And, given the choice Match, the entrant is indeed better off choosing Enter.

But is the first equilibrium really plausible? In other words, once entry has taken place, will the incumbent carry out this threat or intention to Undercut her price? In terms of the decision sequence or game tree, the entrant might reason as follows: "Once I have chosen Enter, the incumbent will be better off choosing Match. Undercut is an empty threat. I therefore am better off choosing Enter."

Formally, a *subgame* of a game tree starts at a single node that is not in an information set with other nodes. The entry game has the simple subgame depicted in Figure 7.3. Given that fact, instead of requiring only that strategies be best replies for the original game, it seems reasonable to impose the additional condition that the relevant parts of each player's overall strategy be a best response in any subgame as well. In other words, the chosen strategy should not only be rational in the Nash equilibrium "best response" sense but in addition should not involve the player in an irrational choice among available options at any later (decision) node, even nodes that may not be reached if the Nash equilibrium is played. Whenever the Nash equilibrium strategies are rational at any *subgame* starting from a node in a tree, the equilibrium is said to be *perfect*, or, more precisely, *subgame perfect* (Selten, 1965).

The Nash equilibrium (Out, Undercut) involves a suboptimal choice by the incumbent in the event that the entrant deviates from the Nash

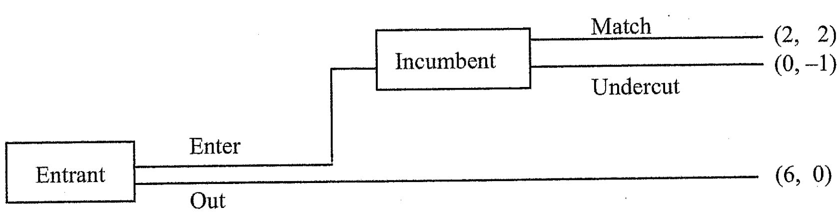

Figure 7.2. Game tree of entry.

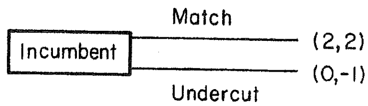

Figure 7.3. Subgame for incumbent.

equilibrium and chooses to Enter. In other words, this Nash equilibrium is supported by a threat by the incumbent to Undercut upon entry, a threat that the incumbent is unlikely to carry out as it is not in his self-interest to do so. The requirement of subgame perfection eliminates Nash equilibria (such as this one) that are supported by non-credible threats.

As another illustration of subgame-perfect equilibrium, consider an auction conducted under the following rules. The auctioneer starts the bidding at $1,000 and will make raises in steps of $1,000. The n bidders draw numbers out of a hat. The buyer drawing number 1 has the first opportunity to accept or reject the initial asking price. If buyer 1 rejects, he is out of the auction and buyer 2 has a chance to bid $1,000. If buyer 1 accepts, the asking price is raised by $1,000 and the auctioneer moves to buyer 2 who then must decide whether to accept at $2,000 or reject (and hence exit). The auction continues until the asking price is rejected by all buyers, in which case the last acceptance becomes the actual sale.

Suppose there the two buyers are bidding for a diamond tiara. Alex, who drew the number 1, values the tiara at $3,500. Bev, who drew the number 2, values the tiara at $2,500.[10] The "sensible" solution is for Alex to accept the opening price of $1,000 while Bev accepts the next asking price of $2,000. Alex then bids $3,000 and wins the tiara. However, there are other Nash equilibria. Consider the following alternative strategy pair as a solution:

Alex's strategy: Reject the initial price.
Bev's strategy: Accept an asking price if and only if it is less than $5,000.

The strategy described for Bev seems rather weird, since it raises the possibility that she could end up paying for the tiara more than her valuation of $2,500. But let us follow the logic of the proposed solution. Given that Alex does reject immediately, Bev will take the tiara for $1,000, getting her maximum net payoff of $2,500 – $1,000 = $1,500. So this strategy for Bev is indeed a best response to Alex. Now consider whether Alex's strategy is a best response to Bev's. If Alex rejects he ends up with nothing. But if he

[10] We are implicitly assuming that the seller does not have full knowledge of the buyers' reservation prices. For, if he did, rather than hold an auction he would simply announce an asking price close to $3,500.

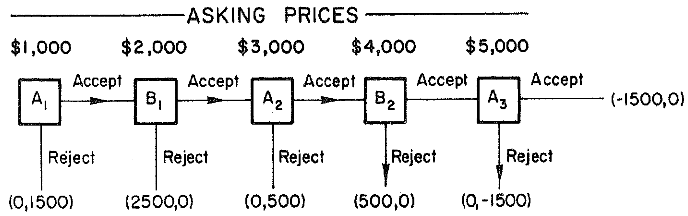

Figure 7.4. Open auction with alternating bids.

accepts, that is, if he follows some strategy *other* than the one considered here, for example, bidding up to his own valuation, the bidding sequence might go as follows:

Alex bids $1,000
Bev bids $2,000
Alex bids $3,000
Bev bids $4,000 (!) and gets the tiara.

No matter what specific strategy Alex chooses (he might, for example, set himself an upper limit of $1,000 or $3,000 or $5,000 or . . .), he will either end up with nothing or, worse, end up paying more for the tiara than it is worth to him. To reject the initial asking price is therefore indeed a best response to Bev's strategy. It follows that the pair of proposed strategies is also a Nash equilibrium.

Just as in the entry game, we can eliminate this "implausible" Nash equilibrium by requiring that the equilibrium be subgame perfect. The tree or extensive form of the game is depicted in Figure 7.4, with initial node A_1. There are four subgames beginning at B_1, A_2, B_2, and A_3. Each is easily analyzed. Starting with the last subgame beginning at A_3, Alex loses $1,500 by accepting once Bev has bid $4,000, so his best response is to reject. This is denoted by the arrow pointing down from A_3.

Next, consider the subgame originating at B_2. If Alex had bid $3,000, Bev's payoff from accepting at $4,000 would be −$1,500 (= $2,500 − $4,000) since, as we have just argued, Alex will reject at his next opportunity. Her optimal move at B_2 is therefore to reject, which would violate the strategy under consideration. It follows that the implausible Nash equilibrium associated with that strategy pair is not subgame perfect.

To confirm that the "sensible" intuitive Nash equilibrium is subgame perfect, consider the subgame with initial node A_2. From our previous argument, if Alex accepts the asking price of $3,000, Bev will reject the asking price of \$4,000, and so Alex's net payoff is \$500 (= \$3,500 − \$3,000). It follows that his optimal strategy is to accept. This is depicted by an arrow pointing across from the node A_2.

Now consider the subgame with initial node B_1. Bev is outbid if she accepts and gets nothing if she rejects. Therefore to accept the asking price of \$2,000 is a best response. Given this, Alex's optimal strategy in the opening round is to accept also. We conclude therefore that the "sensible" Nash equilibrium (in which Alex will bid up to \$3,000 and Bev up to \$2,000) is subgame perfect.

Unfortunately, this is not the end of the story. From node B_1, Bev is indifferent between accepting and rejecting. Rejecting \$2,000 is therefore also a best response. It follows that there is a second subgame-perfect equilibrium in which Alex accepts the starting offer of \$1,000 and Bev then drops out of the bidding. Nor is this an entirely implausible outcome. Intuitively, Bev may note that Alex always has an incentive to outbid her, and so she may well decide not to bother going through the exercise of pushing up the price on the seller's behalf. However, if there is any chance at all that Alex will not continue bidding, Bev is strictly better off staying in and accepting the asking price of \$2,000.

This suggests a further approach to "refining" the Nash equilibrium concept. Starting with some game *G*, one might perturb the payoffs and consider what happens as the perturbation approaches zero. A Nash equilibrium for the original game *G* that is the limit of Nash equilibria in the perturbed game is surely more credible than if this were not the case. For example, in the previous bidding game suppose that Alex's valuation is \$3,500 with probability $1 - \pi_A$ and \$1,500 with probability π_A while Bev's valuation is \$2,500 with probability $1 - \pi_B$ and \$500 with probability π_B. (Here each person's probability distribution is known to the opponent, but only the individual knows his or her own actual realization.) If Alex accepts the initial price of \$1,000, he will take the tiara at that price with probability π_B, since Bev will not bid if her valuation is \$500. If Bev has a valuation of \$2,500, she will accept at \$2,000 and win with probability π_A, since Alex will bid \$3,000 only if his valuation is \$3,500. Of the two subgame-perfect equilibria for the bidding game, the equilibrium in which the price is bid up to \$3,000 is therefore more credible than the one in which Bev rejects the asking price of \$2,000.

A second approach, also due to Selten (1975), introduces "noisy" strategies. Suppose an individual who intends to select some strategy x_a from his set of feasible pure strategies $(x_1, x_2, \ldots, x_A)$ unintentionally plays some other strategy x_b with probability $\pi_b > 0$, where $\sum_{b \neq a} \pi_b = \varepsilon$ and ε is small. Then an opponent may want to choose her strategy in the light of this "tremble" possibility.

For our alternating-bid auction, the possibility of such trembles may induce each buyer to stay in the bidding until the asking price exceeds his or her reservation price. Again, the reason should be clear. As long as there is a chance that an opponent will make a mistake and drop out, a buyer is better off accepting any asking price below his or her reservation price, since there is a positive probability of winning.

We explore this idea more systematically in the next section.

Exercises and Excursions 7.3

1 Entry Game with Two Types of Entrant
Suppose that with probability π the entrant, if she decides to enter, signs a short-term contract with a supplier. If so the payoffs in the game are exactly as in Table 7.5. With probability $1 - \pi$ the entrant, if she decides to enter, signs a long-term contract. In that case payoffs are:

ENTRY GAME WITH A LONG-TERM CONTRACT

		Player 2 (entrant)	
		Enter	Out
Player 1 (incumbent)	Match	3, 1	6, 0
	Undercut	−1, −2	6, 0

The incumbent does not know whether the entrant has signed a short-term or long-term contract.

(A) If $\pi = 1$, we have seen that there are two Nash equilibria, one of which is subgame perfect. Show that the conclusion is the same if $\pi = 0$.
(B) There are two subgame-perfect equilibria if $0 < \pi < 1$. Explain.
(C) Suppose that the entrant "trembles" as she chooses her strategy so that there is a small probability that she will stray from her pure Nash equilibrium strategy. What will be the outcome of such a game?
(D) Does your answer change if the incumbent also "trembles" with small probability?

2 Second-Price Sealed-Bid Auction
Alex has a valuation of $3,500, Bev a valuation of $2,500. Bids must be submitted in hundreds of dollars. Valuations are common knowledge to the two buyers but are unknown to the seller.

Each buyer makes a sealed bid, without knowledge of the opponent's bid. The high bidder is the winner and pays the *second highest* bid.

(A) Explain why bids by Alex and Bev of $3,500 and $2,500, respectively, are Nash equilibrium bids. (Since there are no subgames the solution is also subgame perfect.)
(B) Explain why bids of $0 by Alex and $10,000 by Bev are also Nash equilibrium bids. Are there other equilibria as well?
(C) Appeal to arguments along the line of those at the end of the section to conclude that the Nash equilibrium of (A) is more credible than any other equilibrium.

7.4 Further Refinements

As has been seen, the additional requirement of subgame perfectness can reduce the number of Nash equilibria. But subgame perfectness is applicable only to games in which players move one at a time (and where these moves are public information). And, even when there is a strong subgame structure, there may be multiple subgame-perfect equilibria, some of which seem more credible than others. It would be desirable therefore to find other criteria for ruling out certain of the implausible Nash equilibria.

Consider the following modification of the entry game analyzed in the previous section. The rules are as before except that now the entrant as first mover can enter in two different ways. These affect the outcome if the seller tries to Match. If the entry is Mild, the entrant accepts a price match by the incumbent, whereas, if the entry is Tough, the entrant fights with a further price cut. The incumbent must choose his strategy knowing whether entry has occurred but without knowing whether the entrant has chosen Mild or Tough.

Payoffs in this game are given in Table 7.6 (note that if the entrant chooses Mild, the payoffs are exactly as in the example of Table 7.5). The tree or extensive form of the game is depicted in Figure 7.5. As before, the nodes connected by the dashed line (the information set) indicate that the player at that point must choose without knowing which branch of the tree he or she is on.

Table 7.6. *Entry game with tough entry-I*

		Player 2 (entrant)		
		Mild entry	Tough entry	Out
Player 1 (incumbent)	Match	2, 2	1, 1	6, 0
	Undercut	0, −1	0, 0	6, 0

There are two Nash equilibria in pure strategies: (1) the entrant stays Out and the incumbent chooses Undercut; (2) the entrant chooses Mild Entry and the incumbent chooses Match.

Since there are no subgames of this game, we need to employ some other refinement of Nash equilibrium to rule out one of these equilibria. Consider how the incumbent fares against each possible entry strategy. With Mild entry, the incumbent has a payoff of 2 if he matches and 0 if he undercuts. If the entry strategy is Tough, again the incumbent is strictly better off choosing Match. Therefore, regardless of the type of entry, the incumbent is strictly better off choosing Match.

Undercut is a *weakly dominated strategy*. Once entry has occurred, Undercut yields the incumbent a lower payoff than Match. And of course, if the entrant remains Out, it makes no difference what the incumbent would have done. It thus seems reasonable to conclude that the incumbent will

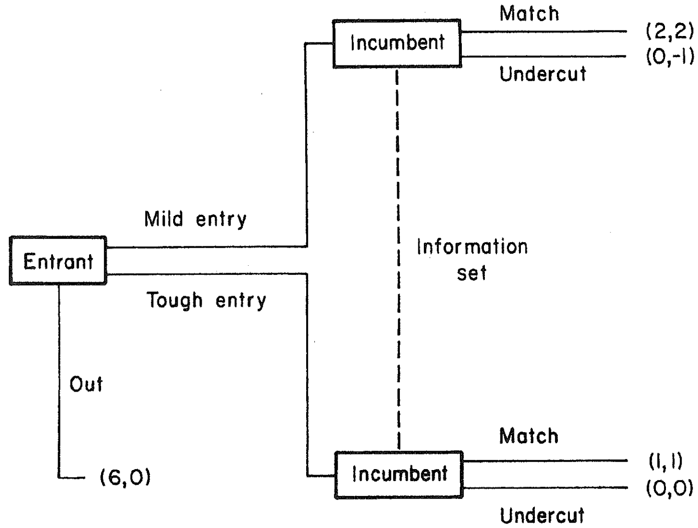

Figure 7.5. Entry game with two entry strategies.

Table 7.7. *Entry game with tough entry-II*

		Player 2 (entrant)		
		Mild entry	Tough entry	Out
Player 1 (incumbent)	Match	2, 2	−1, 1	6, 0
	Undercut	0, −1	0, −1	6, 0

respond to entry with Match. This eliminates the Nash equilibrium in which the entrant stays out because of the (empty) threat of undercutting.

Another very similar example is given in Table 7.7. Here the payoffs in the second column of Table 7.6 (when the entrant chooses Tough) have been changed. Now it is no longer the case that Undercut is a weakly dominated strategy. However, consider the choices of the entrant. If the incumbent matches, the entrant's payoff is higher if she chooses Mild. If the incumbent undercuts, the entrant's payoff is the same whether she chooses Mild or Tough. Tough entry is therefore a weakly dominated strategy for the entrant. By eliminating such a strategy the game is reduced to the original entry game analyzed in the previous section (Table 7.5).

Eliminating (weakly) dominated strategies[11] is a relatively uncontroversial further refinement. However, only in rare cases is the dominance criterion applicable. Consider next the three-player game depicted in tree form in Figure 7.6. Each player chooses either Up or Down. The two nodes for player 3 are connected, indicating that this is an information set. That is, player 3 must select his action without knowing whether it was player 1 or player 2 who made the previous move. The payoffs to the players are specified at the terminal nodes: for instance, if players 1 and 3 each chooses Up, then players 1, 2, and 3 get −1, 0, and 1, respectively.

If player 1 chooses Up, player 3's best response is Down. Player 1 then ends up with a payoff of 3. Since player 1 is certain to have a lower payoff if he chooses Down, this is a Nash equilibrium.

But what if player 1 and player 2 both choose Down? This will occur if player 1 is a pessimist and thinks that player 3 will choose Up. Player 1 therefore chooses Down. If player 2 is also a pessimist he too will choose Down rather than Up, out of fear that player 3 will choose Down.

Note that this outcome occurs because player 1 thinks that player 3 will choose Up while player 2 thinks that player 3 will choose Down. That is, the players have mutually inconsistent beliefs.

[11] A somewhat stronger refinement is the *successive* elimination of weakly dominated strategies.

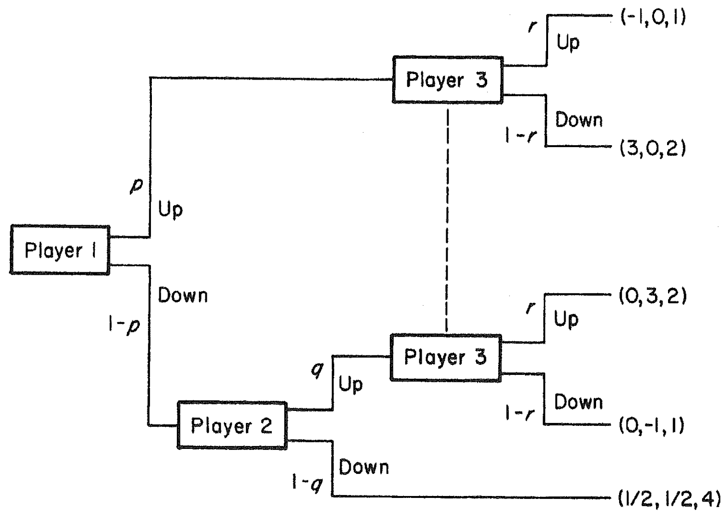

Figure 7.6. Playing against an unknown opponent.

One way of overcoming this problem has been suggested by Kreps and Wilson (1982). Take any strategy vector $s = (s_1, \ldots, s_n)$ for the n players. Consider, for each agent, a completely mixed strategy[12] that is close to s. Since all nodes are reached with completely mixed strategies, Bayes' Theorem can be applied in a straightforward manner to compute beliefs of each information set. Beliefs are then *consistent* if they are the limiting beliefs as the mixed strategies approach s.

For our example, suppose players 1 and 2 choose Down. Then consider this as the limit of mixed strategies in which player 1 chooses Up with small probability p, player 2 chooses Up with small probability q, and player 3 chooses Up with probability r. The probability that player 3 is then called on to play is $p + (1-p)q$ and the conditional probability that it was player 1 who chose Up is:

Prob(1 chose Up | 3's information set is reached)

$$= \frac{p}{p + (1-p)q} = \frac{(p/q)}{(p/q) + 1 - p} \qquad (7.4.1)$$

In the limit, as $p \to 0$ this approaches:

$$\frac{p/q}{1 + p/q}$$

[12] A completely mixed strategy for a player is a mixed strategy in which every pure strategy of the player is selected with non-zero probability.

If q goes toward zero faster than p, this ratio approaches 1 in the limit. On the other hand, if p goes to zero more quickly than q, this ratio approaches zero. Indeed, any conditional belief is consistent. To see this, set $p/q = \pi/(1-\pi)$ where $0 < \pi < 1$. Substituting into (7.4.1) and taking the limit we obtain:

$$\lim_{p,q\to 0} \text{Prob}(1 \text{ chose Up} \mid 3\text{'s information set is reached}) = \pi$$

Thus, for this example, consistency imposes no restrictions upon the beliefs of agent 3. However, it still has predictive power since it imposes the restriction that players 1 and 2 must agree about agent 3's beliefs.

For the tree in Figure 7.6, the payoffs of players 1 and 2 for the top pair of terminal nodes are just the mirror image of those for the second pair of terminal nodes. If $\pi > \frac{1}{2}$ so that player 3 is more likely to be at the upper node, his best response is Down. If $\pi < \frac{1}{2}$ player 3's best response is Up. Finally, if $\pi = \frac{1}{2}$, then player 3 is indifferent and so willing to play a mixed strategy.

We can now establish that, for any consistent beliefs, the terminal node $(\frac{1}{2}, \frac{1}{2}, 4)$ will not be reached in equilibrium. For if $\pi > \frac{1}{2}$ player 3 chooses Down and so player 1's best response is Up. And if $\pi < \frac{1}{2}$ player 3's best response is Up, in which case player 2's best response is Up. Finally, if $\pi = \frac{1}{2}$ and player 3 adopts the mixed strategy of choosing Up with probability r, the payoff of player 1 is $3 - 4r$ if he chooses Up. Moreover, if player 1 chooses Down and player 2 chooses Up, player 2 has an expected payoff of $4r - 1$.

For all possible values of r, the larger of these two payoffs is at least 1. It follows that, for all r, either player 1 or player 2 or both have a best response of Up. The terminal node of $(\frac{1}{2}, \frac{1}{2}, 4)$ is therefore never reached if beliefs are consistent.

We now formalize the concept of consistency.

Definition: Consistent Beliefs

Let $\{(s_t^1, \ldots, s_t^n) | t = 1, 2, \ldots\}$ be a sequence of completely mixed strategies which converges to $(\bar{s}^1, \ldots, \bar{s}^n)$. Let $\{(H_t^1, \ldots, H_t^n) | t = 1, 2 \ldots\}$ be the corresponding beliefs of the n agents at each node of the tree induced by the completely mixed strategies (via Bayes' Theorem). If $(\bar{H}^1, \ldots, \bar{H}^n)$ is the limit of this sequence then these beliefs are consistent with strategy $(\bar{s}^1, \ldots, \bar{s}^n)$.

Having characterized consistent beliefs at every node and information set in the tree, it is then a straightforward matter to work backwards through the tree and determine whether strategies are best responses. If so, the strategies are said to be sequential (that is, sequentially consistent).

Definition: Sequential (Nash) Equilibrium

Let $(\bar{H}^1, \ldots, \bar{H}^n)$ be beliefs that are consistent with the Nash equilibrium strategy $(\bar{s}^1, \ldots, \bar{s}^n)$. If, moving sequentially through the entire tree, the Nash equilibrium strategies are best responses under these beliefs, then the equilibrium is sequential.

The sequential equilibrium concept is a modification of an earlier approach due to Selten (1975). He too begins by considering a sequence of completely mixed strategies that converges to a Nash equilibrium. However, in contrast with the definition of a sequential equilibrium, the beliefs induced by the completely mixed strategies and not just the limit of these beliefs are a part of the definition.

The basic idea is to start with a Nash equilibrium strategy for the n players $(\bar{s}^1, \ldots, \bar{s}^n)$ and then ask whether, for each i, $\bar{s}_i$ is still a best response if each opponent trembles when he tries to play his equilibrium strategy and instead plays each of his other feasible strategies with a small positive probability.

If such a set of trembles can be found, then the equilibrium is said to be trembling-hand perfect. Formalizing our earlier intuitive discussion, we have the following definition:

Definition: Trembling-Hand Perfect Equilibrium

Let $\{(s_t^1, \ldots, s_t^n) | t = 1, 2, \ldots\}$ be a sequence of completely mixed strategies converging to $(\bar{s}^1, \ldots, \bar{s}^n)$ and let $\{(H_t^1, \ldots, H_t^n) | t = 1, 2, \ldots\}$ be the beliefs of the n agents induced by the completely mixed strategies. If for each i and all t sufficiently large, the best response by agent i, given beliefs H_t^i, is $\bar{s}^i$, then $(\bar{s}^1, \ldots, \bar{s}^n)$ is trembling-hand perfect.

While the requirements of a trembling-hand perfect equilibrium are mildly stronger, it is only in rather special cases that a sequential equilibrium is not also trembling-hand perfect.

One of these is the simple bidding game of Exercise 2 at the end of Section 7.3. There Alex had a valuation of $3,500 and Bev a valuation of $2,500. It is a weakly dominant strategy for Alex to remain in the bidding as long as the asking price is less than $3,500. Therefore the belief that Alex will be willing to stay in the bidding beyond Bev's valuation is sequentially rational. Thus it is a sequential equilibrium for Bev to reject the initial asking price of $1,000 and drop out of the bidding. However, if there is a positive probability that Alex will tremble and drop out of the bidding at a price below $2,500, Bev is strictly better off staying in the bidding. So, dropping out immediately is not a trembling-hand perfect equilibrium strategy for Bev.

At a conceptual level, most people find the trembling-hand analogy quite straightforward. In contrast, consistency as a limit of beliefs seems abstract and difficult to grasp. Given this, it is natural to employ the trembling-hand perfect equilibrium when this proves to be relatively uncomplicated. However, for sophisticated applications it can be significantly easier to check whether there exist sequentially consistent beliefs for a Nash equilibrium.

Exercises and Excursions 7.4

1 Sequential Equilibria
Consider the game depicted in Figure 7.6.

(A) Show that (Up, Down, Down) is a sequential equilibrium of this game.
(B) Is (Down, Up, Up) also a sequential equilibrium?
(C) Are there any sequential equilibria in which player 3 adopts a mixed strategy?

2 Elimination of a Family of Nash Equilibria
Suppose that the payoffs for the game depicted in Figure 7.6 are modified as follows:

	Payoff vector
If player 1 chooses Up and 3 chooses Up	(3, 3, 2)
If player 1 chooses Up and 3 chooses Down	(0, 0, 0)
If player 1 chooses Down, 2 chooses Up and 3 Up	(4, 4, 0)
If player 1 chooses Down, 2 chooses Up and 3 Down	(0, 0, 1)
If players 1 and 2 both choose Down	(1, 2, 1)

Let p be the probability that player 1 chooses Up, q be the probability that player 2 chooses Up, and r be the probability that player 3 chooses Up.

(A) Show that for $q \leq 2/3$ and $r \leq 1/3$, two Nash equilibria are (Up, q, Up) and (Down, Down, r).
(B) Show that only one of these two classes of equilibria meets the condition for sequential equilibrium.

3 Trembling-Hand Perfect Equilibrium

(A) Explain why a trembling-hand perfect equilibrium is a sequential equilibrium.

(B) If the payoff matrix of a simultaneous-move game is as shown below, draw the tree of the corresponding sequential-move game, in which player 1 moves first and player 2 must respond without knowing what player 1 has chosen.

		Player 2	
	($\alpha > 1$)	l	r
Player 1	L	1, 1	$1, \alpha$
	R	2, 0	$-1, -1$

(C) Confirm that there are two Nash equilibria. Then let ε_1 be the probability that player 1 takes an out-of-equilibrium action. For each equilibrium confirm that, as long as ε_1 is small, player 2's best response is unaffected.

(D) Suppose, in addition, that player 2 makes an out-of-equilibrium move with probability ε_2. Show that neither player's best response is affected as long as ε_1 and ε_2 are sufficiently small. That is, the equilibria are trembling-hand perfect.

(E) Suppose instead that $\alpha = 1$. Show that there is only one trembling-hand perfect equilibrium. However, both the Nash equilibria identified in (C) are sequential.

COMMENT: This example illustrates that it is only for very specific parameter values of the normal-form payoff matrix that a sequential equilibrium is not also trembling-hand perfect.

4 Open Bidding with Different Valuations

Section 7.3 took up an example in which Alex and Bev made sequential bids for a tiara. For this game there are two subgame-perfect equilibria.

(A) Explain why both are sequential equilibria.

(B) Show that only one is trembling-hand perfect.

(C) Would any small change in the parameters of the model change your answer to (B)?

(D) Try to reconcile your answer with the comment at the end of the previous question.

Table 7.8. *To fight or not to fight*

		Player b (Bev)	
	$v^a, v^b \in \{1, -4\}$	Aggressive	Passive
Player a (Alex)	Aggressive	v^a, v^b	6, 0
	Passive	0, 6	3, 3

7.5 Games with Private Information

While the "refinements" discussed in the previous sections succeed in excluding some of the implausible Nash equilibria, important difficulties remain. This is especially the case when players have private information but *their actions may signal their type.* Further refinements can be sensibly applied in such games of private information.

Consider the following example. Each of two players simultaneously chooses Aggressive or Passive. If both choose Aggressive so that a fight ensues, for either player the payoff is 1 if he is naturally mean and −4 if he is naturally kind. A player knows whether or not he is himself a mean type of individual, but this information is private.

The normal form of the game is depicted in Table 7.8. If player a (Alex) thinks that player b (Bev) is likely to play aggressively, his best response is to choose Aggressive if $v^a = 1$ and Passive if $v^a = -4$. Since the game is symmetric, the same is true for Bev. On the other hand, if Alex thinks that Bev is likely to choose Passive, then his best response is Aggressive, regardless of his private information.

But what will Alex think about Bev? And what will Alex think Bev will think about Alex? And what will Alex think that Bev will think Alex will think about Bev?. . . . And so on.

Economic theorists have, almost exclusively, chosen to rely on a resolution of this puzzle proposed by Harsanyi (1967–68). Suppose that the uncertain payoffs v^a and v^b are draws from some joint distribution. Moreover, and this is critical, suppose that this joint distribution is *common knowledge.*[13] That is, each player knows the joint distribution, each player knows that the other knows the joint distribution, each player knows that the other knows the joint distribution, and so on. Then each player is able to utilize this information to compute a best response.

[13] The concept of common knowledge was discussed in Section 5.3.

For our example, there are four possible payoff pairs when both players choose to be aggressive. Suppose the probability of each payoff-pair is $\frac{1}{4}$. With private information, a complete description of a player's strategy is a description of his strategy (possibly mixed) for each possible private message. We now confirm that it is a Nash equilibrium for each player to choose Aggressive if his or her parameter value is positive and to choose Passive if it is negative.

Suppose Bev behaves in this manner. Given our assumption that each payoff vector is equally likely, there is a probability of 0.5 that Bev will choose Aggressive. If Alex's valuation is v^a, his expected payoff to Aggressive is $(\frac{1}{2})v^a + (\frac{1}{2})6$ while his expected payoff to Passive is $(\frac{1}{2})0 + (\frac{1}{2})3$. The net advantage of Aggressive is therefore $\frac{1}{2}(v^a + 3)$. This is positive if $v^a = 1$ and is negative if $v^a = -4$. Therefore Alex's best response is to behave as proposed. Given the symmetry of the example, it follows that the proposed strategy is a Nash equilibrium.

To reiterate, in a game with private information, a strategy is a description of a player's action (or probabilistic mix of actions) for each possible private information state. As long as the underlying distribution of private informational messages is common knowledge, each player can compute his expected payoff against a particular strategy of his opponent. Equilibrium strategies are then strategies that are best responses, just as in the earlier discussion of Nash equilibrium with no private information. Because of the importance of the common-knowledge assumption, economists sometimes acknowledge the distinction by referring to the equilibrium as a *Bayesian Nash equilibrium.*

As a second example, let us consider an advertising game with private information. In this game, there is an equilibrium in which a seller of high-quality products can signal that fact by costly advertising.

Suppose that a manufacturer is about to introduce a new product that will be of either superior or mediocre quality. These define two "types" of firm. If the product is superior, optimal use by consumers is High. If it is mediocre, optimal use is Low. High rates of consumption generate high revenue and profit for the firm. Before distribution of the new product the firm chooses either Zero advertising, Z, Radio advertising, R, or more expensive TV advertising, T.

It is common knowledge that the odds of a superior product are only 1 to 4. Consumers are not able to observe product quality until it has been used for some time. They do, however, observe the advertising decision of the firm. The firm observes product quality before taking an advertising decision.

Table 7.9. *Advertising game*

			Consumer's choice		
			Low	High	
Type of manufacturer	Mediocre (prob. = 0.8)	T	−3, 1	1, 0	ε
		R	−1, 1	3, 0	ε
		Z	4, 1	8, 0	$1 - 2\varepsilon$
	Superior (prob. = 0.2)	T	−3, 2	5, 4	ε
		R	−1, 2	7, 4	ε
		Z	4, 2	12, 4	$1 - 2\varepsilon$

The payoff matrix for this game is given in Table 7.9. Note that, in switching from Zero advertising to Radio advertising, the manufacturer's payoff declines by 5. This reflects the cost of the advertising. TV advertising costs 7 so there is a further decline of 2 in the manufacturer's payoff if he switches from Radio to TV.

One Nash equilibrium of this game is for consumers to choose a Low rate of consumption and for both types of manufacturers to choose Zero advertising. This is readily confirmed from Table 7.9. With consumers choosing Low, there is no incentive for a manufacturer to incur any advertising costs.

This equilibrium is also sequential (and trembling-hand perfect). To see this, consider the mixed strategies of the two types given by the final column in Table 7.9. Since both types choose Radio advertising with probability ε, the conditional probability that an advertiser is mediocre is equal to the prior probability, that is, 0.8. The expected payoff to consuming at a high rate is therefore $(0.8)(0) + (0.2)(4) = 0.8$, while the expected payoff to consuming at a low rate is $(0.8)(1) + (0.2)(2) = 1.2$. Given such beliefs, consumers will choose the low rate.

Exactly the same argument holds for TV advertising. Therefore the belief that both types of manufacturers will choose Zero advertising is consistent. This "pooling" equilibrium in which the different types are not differentiated is not the only equilibrium, however. There is a second "separating" equilibrium in which a superior manufacturer signals his product's quality via advertising.

Suppose consumers believe that superior manufacturers will choose TV while mediocre advertisers will choose Radio or Zero advertising. From Table 7.9, given such beliefs, the best response to TV is a high rate of use, and the best response to Radio or Zero advertising is a low rate of use. Finally,

Table 7.10. *Equilibrium payoffs in the advertising game*

		Mediocre manufacturer	Superior manufacturer	Consumer
E_1:	Neither type advertises	4	4	1.2
E_2:	Superior chooses TV	4	5	1.6
E_3:	Superior chooses Radio	4	7	1.6

given such choices by consumers, a superior manufacturer has a payoff of 5 if he chooses TV, a payoff of −1 if he chooses Radio, and a payoff of 4 if he chooses Zero advertising. His best response is therefore to advertise on TV. On the other hand, a mediocre manufacturer has a payoff of 4 without advertising and a payoff of 1 if he advertises on TV. It follows that the proposed strategies are Nash equilibrium strategies. Arguing almost exactly as above, it may be confirmed that the consumers' beliefs are consistent. Therefore the Nash equilibrium is also sequential.

But this is far from the end of the story. Suppose consumer beliefs are different, and instead they believe that any manufacturer who advertises on either Radio or TV is of high quality while a manufacturer who does no advertising is of mediocre quality. From Table 7.9, a mediocre manufacturer is still better off not advertising while a superior manufacturer will choose Radio advertising. We therefore have a third Nash (and sequential) equilibrium.

It is instructive to compare the payoffs in the different equilibria, as summarized in Table 7.10. Note that no player is made worse off and at least one is made better off in moving from the first to the second and then to the third equilibrium. In particular, a superior manufacturer has a strong incentive to try to convince players to play the third equilibrium. We shall now argue that such a player, if he is allowed to communicate, can plausibly talk his way out of the other two equilibria. More precisely, we begin by proposing a (sequential) Nash equilibrium and then ask whether the equilibrium beliefs are likely to survive if players can communicate.

Suppose, for example, that the proposed equilibrium has the superior-quality manufacturer choosing TV and the mediocre manufacturer choosing Zero advertising.

The superior firm might send the following message to consumers: "I am a superior firm but I am going to advertise on Radio rather than Television. You should believe me and choose a high rate of usage since a mediocre firm would be worse off if it were to choose Radio and you were to make the same response."

Looking at Table 7.9, we see that this message is correct. With a high rate of usage, the short-run gains for a mediocre firm are offset by the cost of advertising and so profit is 3 – which is less than the equilibrium profit of 4. However, the superior firm is clearly better off. Therefore if consumers recognize that the argument is correct, the Nash equilibrium with TV advertising fails the communication test.

More formally, let $\Theta = \{\theta_1, \ldots, \theta_n\}$ be the set of possible types of player. We will describe an equilibrium as being *weakly communication proof* if no message of the following type is credible:[14]

> I am taking an out-of-equilibrium action and sending you the true message that my type is $\theta_i \in \Theta$, and you should believe me. For if you do and respond optimally, I will be better off while any other type of player mimicking me would end up worse off.

This communication test hinges upon the availability of an out-of-equilibrium action that would be in the interests of only one type of player. The following stronger test allows for an out-of-equilibrium action that would be in the interests of a subset of the possible types of player. We describe an equilibrium as being *strongly communication proof* if no message of the following type is credible:

> I am taking an out-of-equilibrium action and sending you the true message that my type is in B, a subset of Θ, and you should believe me. For if you do so and respond optimally (using prior beliefs about types), any type in B would be better off while any other type of seller attempting to do the same would end up worse off.

It may be that none of the Nash equilibria survive this strong communication test. Thus, at least so far as we now can tell, game-theoretic methodology will sometimes fail to generate a credible equilibrium. While perhaps regrettable this is, we believe, hardly surprising. In general, games with private information have informational externalities. One player's return yields information about his type and, by inference, information about other players who choose different actions. In the presence of such externalities it would be much more surprising if there were a universal existence theorem for credible equilibria.

Nevertheless, even without equilibrium explanations for every imaginable situation, the Nash equilibrium concept with its refinements has proved to be fruitful for analyzing a wide range of strategic interactions, as will be illustrated further in the chapters to come.

[14] This is what Cho and Kreps (1987) refer to rather obliquely as the "intuitive criterion."

Exercises and Excursions 7.5

1 Bayesian Nash Equilibrium with Correlated Beliefs

In the first example considered in Section 7.5, the probability that an opponent has a positive payoff when both are aggressive is 1/2 regardless of a player's type. That is, types are independent. Suppose instead that $v^a = 1$ and $v^b = 1$ with probability $\frac{1}{4}$ while $v^a = 1$ and $v^b = -4$ is β and $v^a = -4$ and $v^b = 1$ is β. That is, the joint probability matrix is symmetric.

(A) Show that types are positively correlated if and only if $\beta < \frac{1}{4}$.
(B) If $\beta \geq \frac{9}{40}$ holds, show that it is a Bayesian Nash equilibrium for a player to choose Aggressive when v^a is positive and Passive when v^a is negative.
(C) Show also that if $\beta < \frac{9}{40}$ then this is no longer a Bayesian Nash equilibrium.
(D) What is the Bayesian Nash equilibrium in this case?

7.6 Evolutionary Equilibrium

The approach outlined in the preceding sections is not useful in exploring the long-run equilibria of games that are played repeatedly in large anonymous populations subject to natural selection. In such cases what can be said about the outcome and what equilibria are likely to emerge? One way of resolving this problem is to introduce evolutionary or natural-selection considerations, as proposed by the biologist John Maynard Smith (1976, 1982). Imagine a large uniform population of organisms that randomly encounter one another in pairwise interactions, with payoffs for each single encounter given by some game matrix. Then, owing to the force of natural selection, over the generations a strategy yielding above-average return will gradually come to be used by larger and larger fractions of the population while strategies with below-average returns will shrink in representation. Among economic players, *imitation* may replace or supplement natural selection, with somewhat similar results (Alchian, 1950; Winter, 1964). If the dynamic evolutionary process leads to a population whose members are fractionally distributed over a set of strategies – or, as a special case, all of whom are following some single strategy – then that distribution is called an *evolutionary equilibrium*, provided that the evolutionary process works to maintain and restore the distribution in the face of all sufficiently small arbitrary displacements of the population proportions ("shocks").

To begin with, consider only pure strategies. Let us assume a symmetrical game (so that the row and column players could be interchanged without affecting the matrix of payoffs).[15] Denote as $V(x_a \mid x_b)$ the expected payoff to an organism playing x_a in an environment where everyone else in the population is playing x_b. Maynard Smith defined what he termed an "evolutionarily stable strategy" (ESS) as follows. Strategy x_a is an ESS if either of these two conditions hold:

(i) $V(x_a \mid x_a) > V(x_b \mid x_a)$
(ii) Or, $V(x_a \mid x_a) = V(x_b \mid x_a)$ and $V(x_a \mid x_b) > V(x_b \mid x_b)$

for any $b \neq a$

The first condition corresponds essentially to the strategy pair (x_a, x_a) being a *strong* Nash equilibrium. If when everyone else is playing x_a, any single player finds that x_a is strictly better for him than any other strategy, then x_a is an ESS. The second condition says that, even if (x_a, x_a) is only a *weak* Nash equilibrium, with x_a yielding the same expected payoff as another strategy x_b, that strategy pair can still be an ESS provided that x_a can defeat any other strategy x_b when the population consists almost entirely of players of x_b. In effect, the first condition says that the home team prevails when it can beat any intruder. The second says that the home team can still win out even if it only ties some intruders, provided it can beat any such intruder on the latter's own home field.

Satisfying the conditions for an ESS does not necessarily suffice for evolutionary equilibrium, however. In the first place, the ESS definition above was pitched in terms of a single evolutionarily stable *strategy* – whereas, more generally, evolutionary stability is a characteristic of a population distribution over a set of strategies. If strategies *a*, *b*, and *c* are being played within a certain population in proportions p_a, p_b, and p_c, respectively, that distribution may or may not be evolutionarily stable – with no implications one way or the other as to whether any of the three component strategies is an ESS standing alone.

It is important not to confuse a *mixed population* with a *uniform population playing a mixed strategy*. A population distribution in the proportions p_a, p_b, and p_c over the pure strategies *a*, *b*, and *c* does not in general have the same stability properties as a population uniformly playing the corresponding probability mixture of the same three strategies. (Except that when there are only two pure strategies, the stability properties are indeed equivalent.)[16]

[15] This amounts to assuming a homogeneous population in which everyone is of one single type.
[16] Maynard Smith (1982), pp. 184–186.

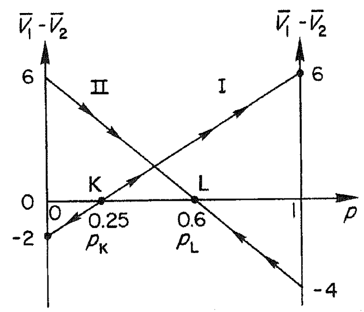

Figure 7.7. Evolutionary equilibria in Tender Trap and Chicken.

Second and more fundamentally, since an evolutionary equilibrium is the stable terminus of a natural-selection process over the generations, it is characterized not only by the payoff elements entering into the definition of the ESS but also by the dynamic formula governing the change in population proportions in response to yield differentials in any generation (Taylor and Jonker, 1978; Zeeman, 1981; Friedman, 1991; Hirshleifer and Martinez Coll, 1988). It may be, for example, that, even if the payoffs remain unchanged, increasing the sensitivity of the dynamic response formula can lead to explosive cycling instead of the damped behavior of a system consistent with ultimate stability.[17]

In what follows, we will generally be employing the evolutionary equilibrium terminology, although in most of the simple cases dealt with the ESS definition originally proposed by Maynard Smith suffices to locate the equilibrium.

Along the lines of the analysis in Hirshleifer (1982), the essentials of the relation between Nash equilibrium and evolutionary equilibrium are pictured in Figure 7.7. On the horizontal axis is plotted p, the proportion of the population playing the first or "more cooperative" strategy in the games Chicken and Tender Trap considered in Section 7.2.[18]

[17] Consider the discrete-generation dynamic formula:

$$\Delta p_a = \kappa p_a(\bar{V}_a - V), \quad \text{for all strategies } a = 1, \ldots, A$$

Here p_a is the proportion of the population playing strategy a (where, of course, $\sum_a p_a = 1$), $\bar{V}_a$ is the mean payoff received by a player of strategy a, V is the average mean yield for the population as a whole, and κ is a parameter representing the sensitivity of the dynamic process. Then, if there is an interior Nash equilibrium consisting of a population distributed over a set of pure strategies, whether or not that Nash equilibrium is also an evolutionary equilibrium depends upon κ being sufficiently small (Hirshleifer and Martinez Coll, 1988, pp. 387–390).

[18] Recall that "Dvorak" and "Coward" are the more cooperative strategies in (the simultaneous move version of) Tender Trap and in Chicken, respectively.

On the assumption that each individual will be randomly encountering other members of the population in a one-in-one interaction, let $\bar{V}_1$ denote the average payoff, as a function of p, to an individual choosing the more cooperative strategy 1. Similarly, let $\bar{V}_2$ be the expected payoff to the less cooperative strategy 2 as a function of p. Thus:

$$\begin{aligned} \bar{V}_1 &\equiv pV(x_1|x_1) + (1-p)V(x_1|x_2) \\ \bar{V}_2 &\equiv pV(x_2|x_1) + (1-p)V(x_2|x_2) \end{aligned} \tag{7.6.1}$$

Evidently, the second strategy will be more successful, and thus over the generations will be naturally selected over strategy 1, whenever $\bar{V}_1 - \bar{V}_2 < 0$. From (7.6.1):

$$\bar{V}_1 - \bar{V}_2 = p[V(x_1|x_1) - V(x_2|x_1)] + (1-p)[V(x_1|x_2) - V(x_2|x_2)]$$

If strategy x_1 is a strong Nash equilibrium so that the first bracketed term is strictly positive, then $\bar{V}_1 - \bar{V}_2$ is necessarily positive when p approaches unity. Moreover, even if the payoffs were such that the first bracket is zero, $\bar{V}_1 - \bar{V}_2$ is still positive if the second bracket is positive. Hence sufficient conditions for an evolutionary equilibrium are indeed those given by (i) and (ii) in the definition of ESS above.[19]

Figure 7.7 indicates that two qualitatively different types of situations associated with the payoff environments of Tender Trap (line I) and Chicken (line II). Line I is positively sloped, owing to the fact that in Tender Trap it is more profitable always to conform to what the great majority of the other players are doing; line II is negatively sloped since in the environment summarized by the Chicken payoff matrix it is more advantageous to do the contrary.

For the Tender Trap example of Exercise 7.2.1, $p = 0$ corresponds to the mutually less profitable "Qwerty" Nash equilibrium at (x_2, x_2); $p = 1$ similarly corresponds to the more profitable "Dvorak" Nash equilibrium at (x_1, x_1); and finally the crossover point K at the population proportions $(p_K, 1 - p_K) = (0.25, 0.75)$ corresponds to the mixed Nash equilibrium. As can be seen, for $p > p_K$, the difference $\bar{V}_1 - \bar{V}_2$ is positive. Then, as indicated by the arrows, the proportion adopting the first strategy will grow over the generations, eventually achieving the extreme at $p = 1$. For any initial proportion $p < p_K$, on the other hand, the evolutionary process over time will go the other way, terminating at $p = 0$. Thus, the mixed-strategy Nash

[19] This argument can also be used to show that even when there are more than two pure strategies (i) any strong Nash equilibrium is an evolutionary equilibrium and (ii) a weak Nash equilibrium might not be an evolutionary equilibrium.

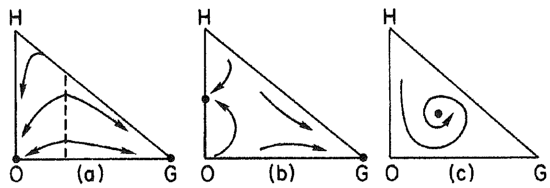

Figure 7.8. Examples of possible evolutionary equilibrium patterns when $A = 3$.

equilibrium[20] represented by point K is not an evolutionary equilibrium (is not "evolutionarily stable"); only the pure-strategy Nash equilibria at $p = 0$ and at $p = 1$ are evolutionary equilibria.

For the Chicken payoff matrix of Table 7.4, there is a mixed Nash equilibrium at the crossover point L, where $(p_L, 1 - p_L) = (0.6, 0.4)$. Since in Chicken it is the *less* prevalent strategy that has the advantage, as indicated by the arrows along line II, the evolutionary progression is always away from the extremes and toward the interior solution at point L. Thus, for Chicken only, the mixed Nash equilibria is an evolutionary equilibrium.

A number of new possibilities emerge when we consider strategy sets where the number of pure strategies, A, exceeds 2.[21] Figure 7.8 is a suggestive illustration for $A = 3$. *Any* point within each triangle represents a distribution of the population over the three strategy options. The horizontal coordinate represents the proportion p_1 playing strategy 1, the vertical coordinate the proportion p_2 playing strategy 2, while the remaining proportion p_3 is measured by the horizontal (or, equivalently the vertical) distance from the hypotenuse. Thus, the origin is the point where $p_3 = 1$. The arrows show convergence possibilities for a number of different possible cases. Without going into the specific conditions here,[22] it is evident that there may be evolutionary equilibria (as indicated by the heavy dots in the diagrams) at various combinations of: (i) one or more vertices; (ii) an edge and a vertex; or (iii) in the interior.[23] A vertex evolutionary equilibrium corresponds to the situation where only a single strategy is represented in equilibrium; an edge evolutionary equilibrium corresponds to more than

[20] Since in both Tender Trap and Chicken only two pure strategies are involved, for purposes of analyzing evolutionary stability we can, as indicated above, deal with a uniform population playing a mixed strategy as if it were a mixed population playing the corresponding distribution of pure strategies.

[21] Here we will always be thinking of *mixed populations*, each member of which plays some pure strategy, rather than *uniform populations playing a mixed strategy*.

[22] On this, see Hirshleifer and Martinez Coll (1988, pp. 379–380).

[23] Some of the possibilities are illustrated in the exercises below.

one, but not all, strategies being represented; an interior evolutionary equilibrium indicates that all strategies are represented. Or, finally, there may be no evolutionary equilibrium at all.

For $A \geq 3$, there are several other interesting implications. First, it may be that each of two or more strategies can defeat all others, but are tied against one another. In such cases, no single one of the winning strategies will meet the conditions for an evolutionary equilibrium, since it is not stable with regard to displacements shifting its proportionate representation as against other members of the winning group. Yet, the group as a whole represents a kind of *evolutionary equilibrium region*, since any given starting point of the evolutionary progression will always be attracted to some terminus along the edge connecting those strategies. Second, it is also possible to have a different kind of attractive region: a closed "limit cycle" toward which the population proportions spiral from within or without – but where, along the curve itself, the proportions cycle perpetually.

Exercises and Excursions 7.6

1 Nash Equilibrium and Evolutionary Equilibrium

Identify the pure-strategy and mixed-strategy Nash equilibria, and the evolutionary equilibria as well, of the following payoff matrices. Show that 1 has evolutionary equilibria at all three vertices; 2 has an evolutionary equilibrium only along an edge; 3 has only an interior evolutionary equilibrium; and 4 has no evolutionary equilibrium at all.

1	*a*	*b*	*c*	2	*a*	*b*	*c*
a	8, 8	3, 2	1, 7	*a*	3, 3	4, 4	2, 2
b	2, 3	5, 5	4, 0	*b*	4, 4	3, 3	2, 2
c	7, 1	0, 4	5, 5	*c*	2, 2	2, 2	1, 1
3	***a***	***b***	***c***	**4**	***a***	***b***	***c***
a	1, 1	2, 2	3, 3	*a*	3, 3	3, 3	2, 1
b	2, 2	1, 1	2, 3	*b*	3, 3	3, 3	2, 2
c	3, 3	3, 2	1, 1	*c*	1, 2	2, 2	1, 1

2 Evolutionary Dynamics

This exercise is designed to illustrate the dynamics corresponding to the two vertex evolutionary equilibria shown in the first diagram of Figure 7.8.

For the payoff matrix below, let (p_1, p_2, p_3) be the population proportions using strategies *a*, *b*, and *c*, respectively.

	a	*b*	*c*
a	2, 2	0, 0	0, 0
b	0, 0	1, 1	0, 2
c	0, 0	2, 0	1, 1

(A) Using the discrete-generation dynamic formula given in footnote 17, show that the population proportions evolve according to:

$$\Delta p_1 = \kappa p_1(2p_1 - V)$$
$$\Delta p_2 = \kappa p_2(p_2 - V)$$
$$\Delta p_3 = \kappa p_3(2p_2 + p_3 - V)$$

(B) Since these changes sum to zero, show that:
$V = 2p_1^2 + p_2^2 + p_3(2p_2 + p_3^2) = 2p_1^2 + (p_2 + p_3)^2 = 2p_1^2 + (1 - p_1)^2$

(C) Hence show that $\Delta p_1 > 0$ if and only if $1 > p_1 > \frac{1}{3}$.

(D) Show that V exceeds $\frac{2}{3}$ for all p_1. Hence explain why $\Delta p_2 < 0$ for all $p_2 < \frac{2}{3}$.

(E) Show that if $p_1 < \frac{1}{3}$ and $p_3 \approx 0$, Δp_2 is strictly positive.

(F) Use these results to explain why the population proportions will evolve as depicted in Figure 7.8a (as long as κ is sufficiently small).

3 *Nash Equilibrium versus Evolutionary Equilibrium, and Evolutionary Equilibrium Region*

(A) For the payoff matrix below, show that there is a weak Nash equilibrium at the (c, c) strategy pair along the main diagonal (corresponding to the *c*-vertex of the triangle). However, show that this Nash equilibrium cannot be an evolutionary equilibrium.

[HINT: If $p_c = 1 - \varepsilon$, are there any p_a, p_b population fractions, summing to ε, for which strategy *a* and/or strategy *b* has higher payoff than strategy *c*?]

	a	*b*	*c*
a	3, 3	3, 3	1, 4
b	3, 3	3, 3	2, 2
c	4, 1	2, 2	2, 2

(B) Show that strategy c will be dominated, even as ε approaches unity, by mixtures of a and b in which the latter has more than 50% representation. Accordingly, identify a range along the a–b edge of the triangle that represents an evolutionary equilibrium region.

4 Interacting Populations

The analysis in the text postulated interactions within a single homogeneous population. But sometimes we want to consider interactions between members of two distinct populations: e.g., males versus females, buyers versus sellers, predators versus prey.

(A) The payoff matrix below is associated with the game known as Battle of the Sexes. What are the three Nash equilibria of this game?

(B) In considering evolutionary equilibria of games with interacting populations, we seek stable vectors of population proportions ($p_1, p_2, \ldots, p_A; q_1, q_2, \ldots, q_A$) – where the p's represent the proportions of the first population and the q's are proportions of the second population, both distributed over the A available pure strategies. Find the evolutionary equilibria of the Battle of the Sexes game, if any.

		Player k	
		x_1^k	x_2^k
Player j	x_1^j	10, 8	4, 4
	x_2^j	4, 4	8, 10

(C) In the Chicken game, described earlier in the chapter as taking place within a single homogeneous population, there were three Nash equilibria of which only the single "interior" Nash equilibrium (representing a mixed-strategy or mixed-proportions solution) was an evolutionary equilibrium. Suppose now that the interacting players come from different populations. Do the results differ?

5 The "War of Attrition" (Maynard Smith, 1976; Riley, 1979)

Two animals are competing for a single prize (item of food). As long as they engage in a merely ritualistic (non-damaging) struggle, neither wins the prize and each incurs an opportunity cost of c per unit of time. If player

2 withdraws at time t, the payoffs are:

$$U_1 = V - ct$$
$$U_2 = -ct$$

(A) Explain why no player will, in equilibrium, choose any particular stopping time $\hat{t}$ with positive probability.

(B) Suppose player 2 adopts the continuously mixed strategy of dropping out by time t with cumulative probability $G(t)$. Write the expected payoff to player 1.

(C) Hence, or otherwise, show that the Nash equilibrium strategy of this game is:
$G(t) = 1 - \exp^{-ct/V}$

(D) Explain why it is only a matter of notational convenience to choose a unit of measurement so that $c = V$.

(E) *Under this assumption, show that the expected payoff to playing the Nash equilibrium strategy against a mutant strategy $H(t)$ can be expressed as:

$$V(v|\mu) = \int_0^\infty G'(t)\left[H(t) - t + \int_0^t H(x)dx\right]dt$$
$$= \int_0^\infty [G'(t)(H(t) - t) + H(t)(1 - G(t))]dt$$

[HINT: Integrate $\int_0^\infty G'(t) \int_0^t H(x)\, dxdt$ by parts, using the fact that $G'(t) = -\, d(1 - G(t))/dt$.]

(F) *Show that $V(\mu \mid \mu) - V(v \mid \mu)$ reaches a maximum at $\mu = v$, that is, the Nash equilibrium is also an evolutionary equilibrium.

[HINT: Write the Euler condition and then use the fact that $G'(t) + G(t) - 1 = 0$.]

SUGGESTIONS FOR FURTHER READING: There are several first-rate books on game theory. Friedman (1986) and Gibbons (1992) are excellent introductions to game theory. For more in-depth treatments, the reader should consult Fudenberg and Tirole (1991), Osborne and Rubinstein (1994), or Myerson (1997). For evolutionary game theory, the subject of Section 7.6, see Weibull (1997).

* Starred questions or portions of questions may be somewhat more difficult.

References

Alchian, Armen A., "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy*, 58 (1950), 211–221.

Cho, In-Koo and Kreps, David M., "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102 (1987), 179–221.

Cournot, Augustin, *Recherches sur les Principes Mathématiques de la Théorie des Richesses*, Paris: Hachette, 1838.

Friedman, Daniel, "Evolutionary Games in Economics," *Econometrica*, 59 (1991), 637–666. Reprinted in Eric Maskin, ed., *Recent Developments in Game Theory*, a volume in the series *The International Library of Critical Writings in Economics*, Northampton, MA: Edward Elgar, 1999.

Friedman, James W., *Game Theory with Applications to Economics*, New York: Oxford University Press, 1986.

Fudenberg, Drew, and Jean Tirole, *Game Theory*, Cambridge, MA: MIT Press, 1991.

Gibbons, Robert, *Game Theory for Applied Economists*, Princeton: Princeton University Press, 1992.

Grossman, Sanford and Perry, Motty, "Perfect Sequential Equilibrium," *Journal of Economic Theory*, 39 (1986), 120–154.

Harsanyi, John G., "Games with Incomplete Information Played by 'Bayesian' Players," *Management Science*, 14 (1967–68), 159–182, 320–334, 486–502.

Hirshleifer, Jack, "Evolutionary Models in Economics and Law: Cooperation versus Conflict Strategies," *Research in Law and Economics*, 4 (1982), 1–60.

Hirshleifer, Jack and Coll, Juan Carlos Martinez, "What Strategies Can Support the Evolutionary Emergence of Cooperation?" *Journal of Conflict Resolution*, 32 (1988), 367–398.

Kohlberg, Elon and Mertens, Jean-François, "On the Strategic Stability of Equilibria," *Econometrica*, 54 (1986), 1003–1038.

Kreps, David M. and Wilson, Robert, "Sequential Equilibrium," *Econometrica*, 50 (1982), 863–894.

Liebowitz, Stan J. and Margolis, Stephen E., "The Fable of the Keys," *Journal of Law and Economics*, 33 (1990), 1–25.

Maynard Smith, John, "Evolution and the Theory of Games," *American Scientist*, 64 (1976), 41–45.

———, *Evolution and the Theory of Games*, Cambridge: Cambridge University Press, 1982.

Myerson, Roger, "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 7 (1978), 73–80.

———, *Game Theory: The Analysis of Conflict*, Cambridge, MA: Harvard University Press, 1997.

Nash, John F., Jr., "Non-Cooperative Games," *Annals of Mathematics*, 54 (1951), 286–295.

Osborne, Martin J. and Ariel Rubinstein, *A Course in Game Theory*, Cambridge, MA: MIT Press, 1994.

Riley, John G., "Evolutionary Equilibrium Strategies," *Journal of Theoretical Biology*, 76 (1979), 109–123.

Selten, Reinhard, "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrägetragheit," *Zeitschrift für die gesamte Staatswissenschaft*, 121 (1965), 301–324 and 667–689.

———, "A Reexamination of the Perfectness Concept for Equilibrium Concepts in Extensive Games," *International Journal of Game Theory*, 4 (1975), 25–55.

Taylor, Peter D. and Jonker, Leo B., "Evolutionarily Stable Strategies and Game Dynamics," *Mathematical Biosciences*, 40 (1978), 145–156.

Weibull, Jorgen W., *Evolutionary Game Theory*, Cambridge, MA: MIT Press, 1997.

Winter, Sidney G., "Economic 'Natural Selection' and the Theory of the Firm," *Yale Economic Essays*, 4 (1964), 225–272.

Zeeman, Erik C., "Dynamics of the Evolution of Animal Conflicts," *Journal of Theoretical Biology*, 89 (1981), 249–270.