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

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