

1 Introduction

Life is fraught with uncertainty. The benefits of human capital, such as schooling, on the job experience, and children, are unpredictable. Personal health is another cause of great uncertainty. Insurance can only be purchased against some traumatic personal events such as the serious disability from injury, or the death of household breadwinners, but not for all, such as marital breakdown, which is often accompanied by economic hardship and child neglect. Homeowners bear the risk of changes in their neighborhood property values and that often contribute significantly to changes in their total lifetime wealth. Entrepreneurs and small businessmen typically assume a lot of business risk. The timing of death itself is random. This chapter analyzes rational choice under uncertainty.

Our point of departure is the simple hypothesis, that in games for a single player, individuals maximize the expected value of their wealth. Several examples in the next section illustrate how this assumption can be used as a benchmark for calculating how information should be used and its value to the decision maker. Expected wealth maximization is a useful assumption to make in some situations, but seems inappropriate for others. In Section 3 we demonstrate how to test the expected wealth maximization hypothesis using experimental methods. We define the certainty equivalent as the minimal (maximal) amount a person would accept (pay) to avoid a lottery. The certainty equivalent of a wealth maximizer is the expected value of the lottery, or its actuarial value; the certainty equivalent of a risk seeker is more than the expected value, while a risk avoider's certainty equivalent is less. Although more wealth is preferred to less, for many people, the certainty equivalent of a lottery with monetary payoffs is not equal to its actuarial value.

Expected utility maximization generalizes expected wealth maximization to account for different attitudes towards risk. Expected utility maximizers assign a utility to each possible outcome that might occur, and sum (or integrate) the probability weighted utilities from the payoffs. To predict the value of information to an expected utility maximizer, we need to know her utility function. Fortunately, testing the expected wealth maximization hypothesis accomplishes the dual purpose of revealing the utility function of an expected utility maximizer. The experiments explained in Section 3 not only reveal a subject's attitudes towards risk, but also provide a way of recovering her utility function.

Sections 4 and 5 continue our analysis of uncertainty under expected utility maximization. The expected utility hypothesis is widely used to assess attitudes towards risk. We define several parametric forms that are commonly used to model whether players are risk averse, risk loving or risk neutral (expected value maximizers), and describe experiments that help to identify which of these forms can be ascribed to experimental subjects. The approach of assuming that players are risk averse utility maximizers is at the heart of almost all study of insurance policies and

financial portfolio management. The examples on insurance, investment in risky securities and portfolio choice in financial markets in Section 5 illustrate why risk attitudes guide decision making on these issues. We show that a risk averse expected utility maximizer fully insures her property at actuarially fair rates, but invests some of her assets in risky securities if their expected rate of return exceeds the interest rate.

The two remaining topics we discuss in this chapter are devoted to testing and then relaxing the hypothesis that players are expected value maximizers. Whether a person is an expected utility maximizer or not depends on whether they know the laws of probability, and obey the three axioms of rational behavior, completeness, transitivity and independence. In section 6 we define these axioms, and provide several tests for investigating whether experimental subjects obey them, along with the laws of probability.

Predicting individual behavior under uncertainty is feasible but more arduous if people are not expected utility maximizers. One approach is to entertain a more general set of preferences than expected utility theory permits, and extract more detail in experimental testing sessions about subjects' preferences over risks than what expected utility theory demands. In this way we could develop a theory that relaxes the independence axiom but still provides us with predictions about how individuals acquire, use and value new sources of information. An alternative approach is to limit our analysis and the range of our predictions to situations where subjects might rank lotteries the same way even though their attitudes towards uncertainty are heterogeneous. In Section 7 we take up the last topic in this chapter by defining several notions of dominance between probability distributions, and testing whether subjects form the same rankings.

2 Expected Wealth Maximization

A first approach to choice under uncertainty is to maximize the expected present value of wealth. This involves attaching a monetary value to each possible outcome, defining the expected wealth from a decision as the probability weighted sum of the monetary values of the possible outcomes, and the implementing the decision that yields the highest expected wealth. For example, in nonunionized and unregulated corporations with a board of directors responsive to shareholders whose individual portfolios are well diversified, properly incentivized managers might maximize the expected value of the firm. Indeed expected wealth maximization is not only the simplest way to intelligently aggregate across uncertain outcomes and impute the value of information; it is quite realistic to believe many businesses do so. The following two examples show how much information to seek when expected wealth maximization is the objective.

2.1 Foreign investment

A multinational deliberates over the prospect of acquiring a plant in previously unexplored territory. Having already undertaken some study, it must choose between

making a final decision now, or deferring until more information is gathered and processed. The cost of acquiring the plant is \$8 million. Without the benefit of the additional study, the company believes there is a probability of 3/4 that the investment will generate sales revenue with net present value \$16 million, but with probability 1/4 the net present value of sales will be only \$4 million. If a final decision was made now, the company would make the investment, because its net present of \$12 million (3/4 of \$16 million) plus \$1 million (1/4 of 4) minus \$8 million is \$5 million, positive. The decision tree for this problem is depicted in Figure 4.1.




Figure 4.1

Direct foreign investment

Extending the example a little, we might suppose the value of the company increases by $S > 0$ if the venture is successful and by $F < 0$ if it fails. Also denote the probabilities of success and failure by P_S and P_F respectively. Assuming there are no opportunities to gather further information that might affect the probabilities a wealth maximizing company would acquire the plant if and only if

$$P_S S + P_F F > 0$$

Now suppose the company can commission consultants to provide more information about the plant's potential. The consultants provide different kinds of information. One of them is a technical expert with local knowledge who can provide information about the on plant maintenance and supply costs. Reliable suppliers of raw materials are critical to the success of the project. The project cannot be successful otherwise, and of those which fail, a half are due to a breakdown in the supply chain. For a fee of \$1 million the technical consultant will write a report of one letter Y (suppliers are reliable) or N (they are not) that determines whether a reliable supply chain can be established.

The other expertise of the other consultant is in marketing. She has a strong record identifying successful projects, correctly identifying profitable projects eight times out of nine. When the project is weak, however, he still attaches probability of one sixth

that it will succeed. Her report also costs \$1 million to obtain, and would postpone the project by a year. This delay would reduce the value of a successful project by \$2 million, but not affect the value of the project in the event of failure. Again the report consists of a letter, either G , meaning marketing prospects are good, or B , meaning they are bad.

The question facing the company is whether it should commission either or both reports before taking a final decision. First we consider how the company would process the information contained in the technical report about the supply chain. The decision tree for this problem is displayed in figure 4.2.

More generally, suppose factors that doom the project can be identified beforehand with probability $P_{N|F}$. In the numerical specialization $P_{N|F} = 0.5$. Upon receiving the report N , the company immediately terminates the project because $P_{F|N} = 1$. Noting $P_{F|Y} = 1 - P_{S|Y}$, the company's value from receiving Y is

$$P_{S|Y}S + P_{F|Y}F = S + P_{F|Y}(F - S)$$

The two conditional probability $P_{F|Y}$ can be computed from the data given in this problem. In this application $P_{Y|S} = 1$, so

$$P_{F|Y} = \frac{P_{Y|F}P_F}{P_Y} = \frac{P_{Y|F}P_F}{P_S + P_{Y|F}P_F}$$

We conclude that the project should be undertaken if and only if

$$S + \frac{P_{Y|F}P_F}{P_S + P_{Y|F}P_F}(F - S) > 0$$

This still leaves the question whether the company should commission the technical report unanswered. The value of the company from commissioning a technical study and then taking the appropriate action is

$$\begin{aligned} P_Y[S + P_{F|Y}(F - S)] &= P_Y S + P_F P_{Y|F}(F - S) \\ &= (P_S + P_{Y|F}P_F)S + P_F P_{Y|F}(F - S) \\ &= P_S S + P_F P_{Y|F}F \end{aligned}$$

so the value of the information is $P_F P_{N|F} F$.

The value of the marketing report can be analyzed in a similar way. As in the previous example, the information is only valuable if it is decisive. This means:

$$P_{S|G}S + P_{F|G}F > 0 > P_{S|B}S + P_{F|B}F$$

In this case

$$P_{S|G} = \frac{P_{G|S}P_S}{P_G} = \frac{P_{G|S}P_S}{P_{G|F}P_F + P_{G|S}P_S}$$

and

$$P_{S|B} = \frac{P_{B|S}P_S}{P_B} = \frac{P_{B|S}P_S}{P_{B|F}P_F + P_{B|S}P_S}$$

If the report is B , then the value of the acquisition is

$$P_{S|B}P_S + P_{F|B}P_F = \frac{P_{B|S}P_{SS} + P_{B|F}P_{FF}}{P_{B|F}P_F + P_{B|S}P_S}$$

The value of the firm with the marketing report is

$$P_G(P_{S|G}P_S + P_{F|G}P_F) = P_{G|S}P_{SS} + P_{F|G}P_{FF}$$

Therefore the value of obtaining the marketing report is

$$P_{G|S}P_{SS} + P_{F|G}P_{FF} - P_{SS} - P_{FF} = (P_{G|F} - 1)P_{FF} + (P_{G|S} - 1)P_{SS}$$

The first expression is positive, the second negative.

worth to use the marketing report, and asseUndertaking the study would yield more precise odds. Otherwise the company will incur losses of \$be profitableAlthough the factory's condition is determined before the firm takes its first decision, since it is not revealed until after a decision has been taken the solution to the extensive form in Figure 4.20 is identical to the solution to the extensive form in 4.2.

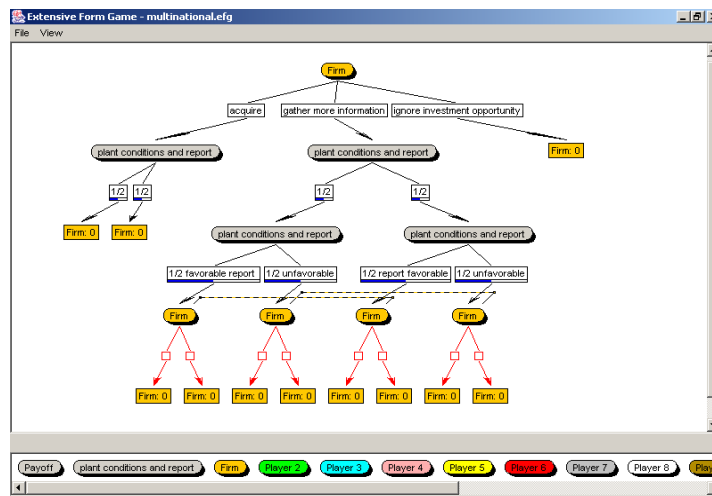



Figure 4.2

Direct foreign investment redrawn

Since the moves of uncertainty are adjacent it is useful to consolidate them. This is undertaken in Figure 4.22.



The expected value of not inspecting and buying is found by solving the subproblem

depicted in 4.22.

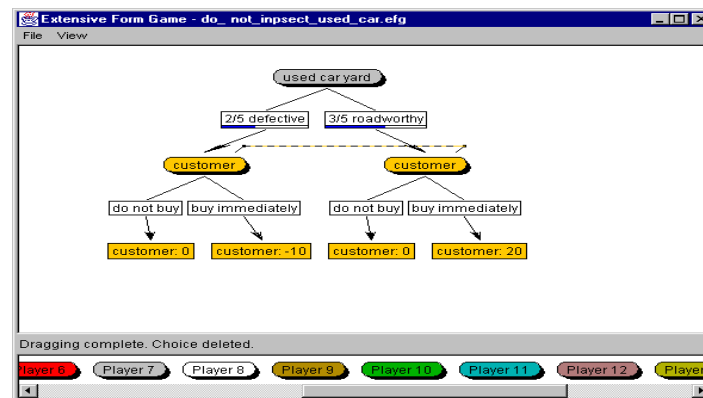


Figure 5.2

Calculating the of not inspecting

The value from not inspecting is the maximum of the value from rejecting the car (nothing) versus buying it immediately ($12 - 4 = 8$).

Inspecting the car

Focusing on the opportunity of making the purchase decision after undertaking an inspection for 2 monetary units, we form the sub-game (?):

Figure 5.3

Calculating the value of an inspection

Simplifying the role of uncertainty

Figure 5.4

Simplifying the role of uncertainty

The expected value of inspecting the car first:

There are four strategies to consider:

1. Buy the car regardless of the outcome of the inspection, yielding an expected value of 6.
2. Do not buy the car regardless of the outcome, yielding an expected value of -2 .
3. Buy the car if it passes the inspection test, but not if it fails, yielding an expected value of 10.
4. Buy the car if it fails the inspection test, but not if it passes, yielding an expected value of -6.67 .

The solution to the used car problem

Notice that the first two strategies do not exploit the conditioning information that the test provides, and thus simply add the cost of the test, 2, to the calculation of the buy/sell strategies. This illustrates the fairly obvious rule that information should not be purchased unless future decisions will condition on its content. Notice that 10, the expected value of the third strategy of the third strategy exceeds all the others. For

example unconditionally buying the car yields 8.

2.2 Filling a vacancy

Acquiring information by search rather than by purchasing it. There are time costs involved. Consider a department within an organization that makes offers to a succession of job candidates. We suppose the value of filling the position to the department is a positive number denoted x , and that the department can offer the job to at most to N candidates, but that if an offer is rejected there is some probability β that the department will lose the position. No bargaining takes place between the department and any candidate: the department simply makes one offer to the current candidate which is either accepted or rejected. We denote the department's n^{th} offer by w_n . If the n^{th} candidate accepts the department's offer then the net benefit to the department is $x - w_n$. Each job candidate $n \in \{1, \dots, N\}$ has a reservation wage which we denote by v_n , meaning that he will accept any wage offer of at least v_n and reject every wage offer below it. The department does not observe v_n , but believes the probability distribution of v_n is independently and identically distributed with probability distribution function $F(v)$ over all candidates $n \in \{1, \dots, N\}$.

In the last period there is no reason to offer more than \bar{v} since every candidate would accept \bar{v} . Setting $N = 1$ the expected value to the department from making an offer of w_1 is the probability of acceptance $F(w_1)$ multiplied by the net value conditional on acceptance ($x - w_1$). Assuming an interior solution, the unconstrained first order condition to this problem is:

$$F(w_1) = (x - w_1)F'(w_1)$$

Generalizing this problem to all finite positive integers N and $n \in \{1, \dots, N-1\}$, let V_n denote the value of the firm when $N - n$ job candidates have rejected their respective offers and the firm follows an optimal wage offer policy for the n remaining candidates. Setting $V_0 = 0$, the firm's value follows the recursion:

$$V_n = \max_{w_n} \{(x - w_n)F(w_n) + [1 - F(w_n)]\beta V_{n-1}\}$$

Substituting V_0 and V_1 in the equation above, noting $V_1 \geq 0$, and comparing it to the same equation when V_1 and V_2 , we deduce $V_2 \geq V_1$, and can show by the induction principles that $V_n \geq V_{n-1}$ for all $n \leq N$. The longer the list of prospective job candidates, the greater the value of the firm. Differentiating the maximand with respect to w_n we obtain the first order condition:

$$(x - w_n^o - \beta V_{n-1})F'(w_n^o) - F(w_n^o) = 0$$

If the stationary point solves the problem it must satisfy the second order condition

$$SOC(w_n^o, V_{n-1}) \equiv (x - w_n^o - \beta V_{n-1})F''(w_n^o) - 2F'(w_n^o) < 0$$

Dropping the subscripts and superscripts, and totally differentiating with respect to V and w we obtain in this case:

$$\frac{dw}{dV} = \frac{\beta F'(w_n^o)}{SOC(w_n^o, V_{n-1})} < 0$$

The more candidates remaining, the lower the offer, wage offers rising as the number of prospective candidates shrink. Thus with each successive rejection, the expected value of the firm falls for two reasons. The probability that the vacancy will not be filled at all increases, and if someone is hired, it will be on less favorable terms to the firm. If the opportunities to hire are unlimited and N is infinite, the value function has the recursive structure

$$V_\infty = F(w_\infty^o)(x - w_\infty^o) + [1 - F(w_\infty^o)]\beta V_\infty$$

Solving for V_∞ in terms of w_∞^o yields

$$V_\infty = \frac{F(w_\infty^o)}{1 - \beta[1 - F(w_\infty^o)]}(x - w_\infty^o)$$

and substituting this equation back into the first order equation yields an implicit solution for w_∞^o :

$$w_\infty^o = x - \frac{F(w_\infty^o)}{F'(w_\infty^o)} \left(1 + \frac{\beta}{1 - \beta} F(w_\infty^o) \right)$$

As an illustrative example if v_n is uniformly distributed on $[\underline{v}, \bar{v}]$, then $\underline{v} \leq v_n \leq \bar{v}$ with $F(v) = (v - \underline{v})/(\bar{v} - \underline{v})$, then for $N = 1$ the department chooses w_1 to maximize:

$$(x - w_1)F(w_1) = (x - w_1)(w_1 - \underline{v})/(\bar{v} - \underline{v})$$

In the unconstrained maximization problem there is a stationary point at

$$w_1^o = (x + \underline{v})/2$$

Noting that the division is assured of hiring the candidate if it offers the candidate \bar{v} , the optimal offer is therefore \bar{v} if $2x > 2\bar{v} - \underline{v}$ and w_1^o otherwise. In this example we assume that $x > \underline{v}$; otherwise the department would have no interest in hiring anyone. For $\underline{v} = 0$ and $\bar{v} = 1$ the interior solution applies with $w_1^o = 1/2$. In this special case the first order condition for $N > 1$ simplifies to

$$x - w_n^o - \beta V_{n-1} = w_n^o$$

and the recursive equation for the value function is

$$V_n = (x - w_n^o)w_n^o + (1 - w_n^o)\beta V_{n-1}$$

Substituting into the equation for V_n and V_{n-1} from the first order condition we thus obtain the solution for w_n^o in recursive form

$$w_{n+1}^o = \frac{(1 - \beta)x + 2\beta w_n^o - \beta(w_n^o)^2}{2}$$

Finally the formula for derived for w_∞^o implies:

$$w_\infty^o = x - w_\infty^o \left(1 + \frac{\beta}{1 - \beta} w_\infty^o \right)$$

This is a quadratic in w_∞^o with the unique positive root:

$$w_{\infty}^o = \frac{1-\beta}{\beta} \left[-1 + \sqrt{1 + \frac{\beta}{1-\beta}} \right]$$

3 Certainty Equivalence

Maximizing expected wealth is a useful assumption to start with, especially when thinking about the objectives of a publicly traded corporation. Shareholders typically hold a small stake in each company, and thus use the law of large numbers to reduce their exposure to risk. In addition they can hold safer assets, such as bonds, if they choose. Consequently those shareholders with higher risk tolerance hold riskier portfolios, so the premium demanded for holding them is modest. But there is plenty of causal evidence against wealth maximization. The returns from (non-tradable) human capital are high relative to (tradable) physical capital. Homeowners (and drivers) partially insure their houses (and cars) at actuarially unfair rates. Individuals insure their health treatment costs at actuarially unfair rates. Entrepreneurs seek financial partners notwithstanding costs of the moral hazard and hidden information. Is wealth maximization reasonable assumption in these situations too? This section devises a test of the wealth maximization hypothesis, and explains why .

3.1 Lotteries

A first approach to testing expected value maximization is to determine how much people value choices with uncertain outcomes. Specifically, how much are they willing to pay to avoid a risk, or equivalently, how much are they willing to pay to gamble? To investigate this question, we consider a lottery where an experimental subject randomly draws one prize from a finite number of potential prizes, knowing the probability of drawing any given prize in advance. Suppose there are K possible prizes, whose values are denoted by x_1 through x_K , and K probabilities, denoted by p_1 through p_K where $0 \leq p_k \leq 1$ for all $k \in \{1, 2, \dots, K\}$ and

$$\sum_{k=1}^K p_k = 1$$

Summarizing, if a person plays the lottery, outcome $k \in \{1, 2, \dots, K\}$ occurs with probability p_k , and in that event she receives a prize of x_k .

How much a subject $n \in \{1, \dots, N\}$ is willing to pay for a lottery ticket reveals much about her attitude towards risk. We define the maximal amount she is willing to pay for a ticket as her reservation value, or certainty equivalent, and denote it by v_n . Because the value of her prize from participating in the lottery invariably lies between the worst and the best of all possible outcomes, we assume

$$\min\{x_1, \dots, x_L\} \leq v_n \leq \max\{x_1, \dots, x_L\}$$

Inducing an experimental subject to truthfully announce her certainty equivalent is a subtle matter. For example if a player is awarded a lottery ticket at the price she bids, she has an incentive to bid less than her certainty equivalent. Bidding less than her valuation yields a profit, but bidding more yields a loss.

To induce experimental subjects to report their reservation values truthfully, we adapt the second price auction mechanism to lottery ticket purchases for one player. To begin the experimental session, the N participating subjects are shown the same lottery, defined by the L prizes with values (x_1, \dots, x_K) and their associated probabilities (p_1, \dots, p_K) . Then each subject $n \in \{1, 2, \dots, N\}$ is asked to place a bid we denote by b_n . After the bids have been recorded, the moderator draws a random number denoted c_n , which is continuously distributed on the closed interval covering (x_1, \dots, x_K) . Whether the subjects know probability distribution $F(c_n)$ or not, and whether c_n is independently distributed across n , is immaterial; however c_n must be drawn independently of the bid b_n . If $b_n \geq c_n$, the subject receives a ticket to the lottery and randomly draws her prize x_n according to the probability distribution (p_1, \dots, p_K) . If $b_n < c_n$ the subject is denied a lottery ticket, and she receives no payoff. Therefore the expected value of the payoff to the n^{th} subject is

$$E[1\{b_n \geq c_n\}(v_n - c_n)] = \int_{-\infty}^{b_n} (v_n - c) dF(c)$$

3.2 Optimal bidding

Suppose the certainty equivalent of the n^{th} subject for lottery L is the value v_n . What is her optimal bid? If n bids less than her certainty equivalent, that is $b_n < v_n$, then she gains $v_n - c_n$ whenever $b_n \geq c_n$, and zero otherwise. Now compare the outcomes of this bidding strategy with the alternative of bidding v_n . There are three possibilities to consider. If $b_n \geq c_n$ or $v_n \leq c_n$, then the bidding strategies yield the same outcome, $v_n - c_n$ and 0 respectively. However if $b_n < c_n < b_n + \Delta$, then bidding b_n yields 0, but bidding v_n nets a gain of $v_n - c_n$. Consequently bidding v_n yields a payoff as least as high as bidding less than v_n for all random draws of c_n , and a strictly higher payoff for some draws of c_n .

We have just proved that subjects should not bid less than their certainty equivalent in this lottery. Now suppose subject n bids more than her certainty equivalent. If $b_n > v_n$, then the n^{th} subject gains $v_n - c_n$ if $c_n \leq v_n$ but loses $c_n - v_n$ if $v_n < c_n \leq b_n$. The only other possibility is that $b_n < c_n$ in which case she pays nothing but does not receive a lottery ticket either. By way of comparison, bidding v_n avoids the loss that occurs when $v_n < c_n \leq b_n$. Again we see that v_n is always at least as beneficial as submitting a higher bid, and sometimes is strictly more beneficial. Therefore v_n is the optimal bid, so in this lottery each subject fully reveals her certainty equivalent through her bid.

3.3 Are bidders expected wealth maximizers?

To test whether experimental subjects are expected wealth maximizers or not, we graph their bids on a sequence of lotteries $L^{(1)}$ through $L^{(J)}$, graphing their bids against the expected values of the lottery. We seek answers to two questions. do subjects have the same attitudes towards risk? In this experiment subjects are asked to how much they are willing to pay for a lottery called L when they know its

probabilities. Call that number b for bid. We then draw a random number n from a probability distribution which lies between (that has support on) 0 and 100. If $n \leq b$, then you pay n in exchange for the lottery L , and receive the payoff from playing the lottery. If $b < n$, then you neither pay nor receive anything. The first lottery pays \$100 half the time and \$0 half the time. The second lottery pays \$100 one eleventh of the time, \$90 one eleventh of the time, . . . , and \$0 eleventh tenth of the time. The third lottery pays \$100 one tenth of the time, \$90 one tenth of the time, . . . , and \$10 one tenth of the time, so always pays out something

Here we report an experiment to see whether we can reject the hypothesis that there is a 45 degree straight line from origin joining all the certainty equivalents. We run a regression with intercept and slope coefficients and conduct an F test.

4 Expected Utility

A more general objective function than expected wealth maximization is expected utility maximization. An expected utility maximizers attaches a utility to all the lottery prizes and sums up the weighted average to evaluate the uncertain prospect. In the case of two lotteries $L_i = (p_{i1}, \dots, p_{iK})$ and $L_j = (p_{j1}, \dots, p_{jK})$ let (u_1, \dots, u_K) denote the vector of utilities associated with the respective prizes. Then L_i is at least as good as L_j if and only if

$$\sum_{k=1}^K p_{ik} u_k \geq \sum_{k=1}^K p_{jk} u_k$$

Further specializing, suppose w_k is the wealth received when the k^{th} outcome in the lottery occurs. Setting $u_k = w_k$ proves that expected wealth maximization is an application of expected utility maximization, but the converse is false. That is an expected utility maximizer might rank the two lotteries differently than an expected wealth maximizer. To see this, suppose $K = 2$ and $(w_1, w_2, w_3) = (1, 3, 7)$. Then a wealth maximizer prefers $L_i \equiv (\frac{1}{2}, 0, \frac{1}{2})$ to $L_j \equiv (0, 1, 0)$ because:

$$\sum_{k=1}^K p_{ik} w_k = 4 > 3 = \sum_{k=1}^K p_{jk} w_k$$

However an expected utility maximizer with log utility prefers L_j to L_i because upon setting $u_k = \log(w_k)$ we see that:

$$2 \sum_{k=1}^K p_{ik} u_k = \log(1) + \log(7) < \log(9) = 2 \log(3) = 2 \sum_{k=1}^K p_{jk} u_k$$

4.1 Constructing a utility function

Expected utility and certainty equivalence. Let $u(x)$ denote the utility from a realization of the random variable x . Assume your preferences obey the expected utility hypothesis. Then your expected utility from playing a lottery F is $EF[u(x)]$, where E denotes the expectations operator. It follows that your certainty equivalent for F is the value vF which solves $u(vF) = EF[u(x)]$

Constructing a utility function that obeys the expected utility hypothesis. We can use a sequence of experiments to construct a utility function. Experiment with different

probabilities. Graph out the results. Construct a utility function based on experiments by first determining a fit for some lotteries. We assign a value of \underline{x} to the lower value and \bar{x} to the upper value. Then we define a sequence of lotteries $L(\underline{x}, \bar{x}, p)$ by a two outcome lottery, \underline{x} and \bar{x} , attaching probability p to the lower outcome, and graph the results.

4.2 Attitudes towards risk

Expected utility and attitudes towards risk . In this case we can characterize you attitude towards risk quite simply. If you are risk neutral, then $vF = EF[x]$ for all F , and hence we can write $u(x) = x$. If you are a risk lover, meaning $vF > EF[x]$ for all F , then $u(x)$ is convex. If you are a risk avoider, meaning $vF < EF[x]$ for all F , then $u(x)$ is concave.

An expected utility maximizer is said to exhibit risk aversion if her certainty equivalent for a lottery is less than its expected value. In the case of an expected utility maximizer if w is a random variable denoting her wealth, and $u(w)$ is her utility function, a risk averse person prefers the expected value of a random wealth over the random variable itself:

$$u(E[w]) > u(c) = E[u(w)]$$

The inequality holds if and only if $u(w)$ is a concave function, that is $u(w)$ satisfies the inequality

$$\lambda u(w_1) + (1 - \lambda)u(w_2) < u(\lambda w_1 + (1 - \lambda)w_2)$$

for any wealth pair (w_1, w_2) and weight $\lambda \in (0, 1)$. We can see this point from Figure 3.1. The horizontal measures wealth, the outcome of a lottery which in this example lies between w_1 and w_2 . We suppose the expected value of




Figure 3. 1
Certainty Equivalence and Concavity




Figure 3.1: A Strictly Concave Utility Function

how risk averse are experimental subjects and how can experiments determine/measure it?

An alternative to assuming investors exhibit quadratic risk aversion is that their utility function takes the form

$$u(c) = c^\eta$$

for some positive constant η . Note that if $\eta > 1$, then $u(c)$ is convex increasing and the investor is risk loving, if $\eta = 1$ then the investor is risk neutral and maximizes expected value, and if $\eta < 1$ the utility function is concave increasing implying risk aversion. One can also show that as η converges to zero, $u(c)$ converges to $\log(c)$. Because of its flexibility and parsimony, the relative risk aversion utility function is widely used to approximate attitudes towards risk.

A person is said to have constant relative risk aversion if her utility function takes the form

$$u(c) = c^\eta$$

The key property of this utility function is that the ratio $u''(c)/u'(c)$ is proportional to c . A person has absolute risk aversion if her utility function takes the form

$$u(c) = -\exp(-\delta c)$$

In this case $u''(c)/u'(c)$ is a constant.

4.3 Measuring risk aversion

Parameterizing the utility function

consider the problem of choosing γ to minimize

$$\sum_{n=1}^N \sum_{j=1}^J (x_j^\gamma - u_{nj})^2$$

where u_{nj} is the certainty equivalent for a lottery with mean x_j . We might minimize over everyone or minimize over each person separately and then test whether they are the same.

5 Expected Utility Maximization

Expected utility is applied in many finance and insurance applications

5.1 The value of information reconsidered

Acquiring information reduces the uncertainty, and hence can be viewed as partial insurance. This remarks suggests that a person's attitude towards risk affects their willingness to pay for information. In particular risk averse players might be more willing to pay for information than risk seekers. The following example explores this conjecture.

Occupational Hazard

Another example

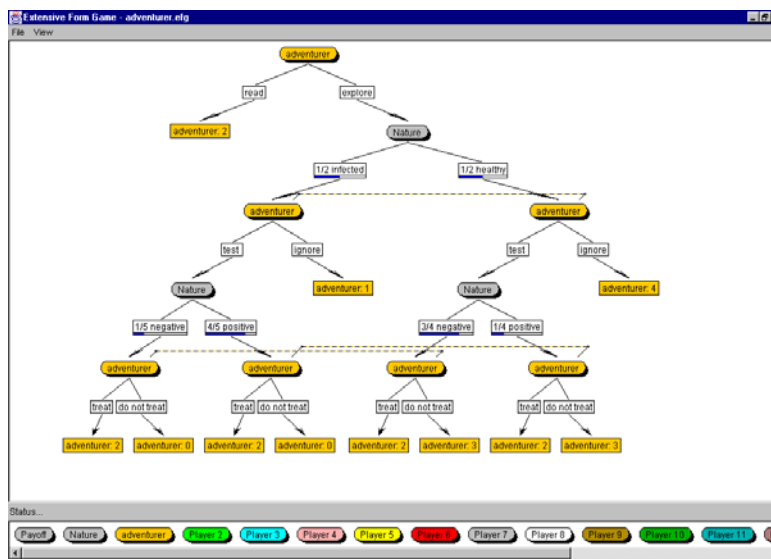


Figure 5.5
Occupational Hazard

A final application to uncertainty. In Figure 3.1 an investor chooses between a stock that has an uncertain return and a bond. After the choice is made the return on the stock is revealed. Alternatively, consider figure 3.2

Figure 4.18
The Value of Information

Inspecting Figure 3.10, you should notice that nature does not reveal the true value of the car until after the consumer has made a decision to purchase it or not.

Figure 4.19
Computing the Value of Information

Having redrawn the decision tree in this way, it is now straightforward to apply the first

rule of strategy to calculate the value of having the car inspected.

5.2 Insurance

Insurance is a pervasive feature of the consumer durable goods and housing sectors, but there are other reasons why cars and houses are insured. Mortgage banks insist that secured credit is backed by insurance. In the event of a loss, the creditor is thus assured that the loan will be repaid. This prevents risk lovers from betting their house against fire or against damage to their car incurred from an accident. The reason why the mortgage or creditor might insist on insurance is not related to their attitude towards risk. After all they can diversify over large number of customers. It has to do with the fact that as first claimant on the collateral, these companies are protecting themselves against a h. Note that if everyone handled the asset in the same way, the lender should not be concerned: it simply incorporates the losses into the interest rate. The problem is that some people are safer and more careful than others, and the loan company resists specializing in those kinds of risks, parceling the job out to an insurance agency. In the end much insurance might be at least as much a response to moral hazard issues that arise in collateralized loans as it is to the demand by risk averse demanders seeking to insure themselves against hardship brought on by accidents.

How much insurance a person buys is depends on his attitude towards risk and the premium he pays. As we suggested in the previous slide, the outcomes of lotteries you own might be negatively correlated with assets you can buy. Therefore insurance decisions are part and parcel of personal asset management. Optimal insurance. A risk averse person fully insures himself against a calamity if actuarially fair insurance is available. This follows from the fact that for risk averters the certainty equivalent of a gamble is greater than its expected value. By the same argument we have used on portfolio choices, a risk averter will not fully insure himself if the rates are not actuarially fair. He will retain ownership over a portion of the lottery.

A common reason given for why consumers buy housing, car and life insurance is that they or their designated beneficiaries are risk averse. The discussion above questions for house and car insurance but seems less contraversial for life insurance. To illustrate this point suppose that a driver begins with wealth w and faces a gamble in which she might lose d from damage in the event of a accident which occurs with probability p . To offset this potential loss an insurance company offers her the opportunity to reimburse her $q \in [0, d]$ for her loss, for a premium of qf . Suppose the driver has an expected utility maximizer with a twice differentiable, increasing utility function $u(w)$ defined on her wealth w . How much insurance should she buy? A risk neutral driver fully insures his vehicle if the rate better than actuarially fair but buys no insurance if it is unfactorable to him. Finally a risk seeker does not buy any insurance sold at actuarilly fair rates, but can be induced to buy it at rates that are favorable to him. Suppose the driver is risk averse. If she only partially insures herself with $q < d$ units, her expected utility is:

$$\begin{aligned}
& pu(w - qf) + (1 - p)u(w - qf - d + q) \\
&= u(w - qf) + (1 - p)[u(w - qf - d + q) - u(w - qf)] \\
&< u(w - qf)
\end{aligned}$$

But we see from the top line that a utility of $u(w - qf)$ can be obtained if insurance is actuarially fair, which means $f = p$ per unit, and the driver purchases full insurance, setting $q = d$. It follows that a risk averse driver would fully insure his vehicle if the premium is actuarially fair. However if the premium was less favorable to the driver, the discussion in the next example below demonstrates that even a risk averse driver would not fully insure his vehicle, preferring to take on some risk. In this case the risky asset, a partially uninsured vehicle, offers a higher expected return than the safe full insurance alternative.

In the following experiment, suppose $u(w) = -\exp(aw)$ and f and $p = 0.1$.

5.3 Pension fund

Consider a worker planning retirement who allocates w , the amount of wealth to be invested for future consumption, between buying shares in a pension fund with a random return denoted by π , and saving at a constant interest rate denoted by r . Alternatively we might like to think of the proportion a pension fund allocates to bonds and the amount allocated to stocks. Denoting the amount of his wealth deposited in his savings account by s , the amount of wealth consumed in retirement is then

$$c = s(1 + r) + (w - s)\pi$$

We assume the worker seeks to maximize his expected utility in retirement

$$E[u(s + sr + w\pi - s\pi)]$$

by choosing $s \in [0, w]$, where $u(c)$ is a real valued, twice differentiable, increasing function in retirement consumption, and the expectation is taken over the random variable π .

A risk averse worker would not allocate any of his wealth into the fund if the expected return on pension funds was less than the return on savings. Historically, the average return from investing in pension funds has exceeded the return on savings. Accordingly we shall assume that $1 + r < E[\pi]$. Under this assumption, a risk neutral or risk seeking worker would not allocate any wealth to his savings account. A weaker result holds for a risk averse worker: he allocates a strictly positive amount to the pension fund. Placing all his wealth in a saving account yields a retirement utility of $u(s + sr)$ to the worker, and the change in utility from allocating a marginal amount from his savings account to a pension fund would be positive since

$$E[u'(s + sr + w\pi - s\pi)(1 + r - \pi)|s = 1] = u'(s + sr)\{(1 + r) - E[\pi]\} > 0$$

Therefore regardless of his risk preferences, it cannot be optimal for a worker to concentrate all his wealth in savings if the return on savings is less than the expected return on the pension fund.

We can also establish the condition under which a risk averse worker puts ignores his savings account. Note first that if the workder is risk averse, then $u(c)$ is concave, second that the expectations operator over the random variable π preserves concavity, and third, that the savings deposit is chosen on the convex set $[0, w]$. Therefore the worker maximizes a concave function on a convex set. The Kuhn-Tucker theroem implies that is a unique local maximum defines the optimal choice of s . We complete the proof by finding sufficient conditions for a local maxmium at the point $s = 0$, where all wealth is concentreated in the pension fund. There is a local maximum at this no savings choice if taking a dollar in from the pension fund and placing it in a savings account lead to a decline in the worker's utility. Thus the necessary and sufficient condition for specialization in the pension fund is

$$(1+r)E[u'(w\pi)(1+r)] < E[u'(w\pi)\pi]$$

where $u'(c)$ denotes the first derivative of $u(c)$ with respect to c .

Otherwise an interior solution obtains for a risk averse worker, characterized by the first order condition for the optimal savings s^o . It states that the expected marginal utility from adding another dollar to either account is equal:

$$(1+r)E[u'(s^o + s^o r + w\pi - s^o \pi)] = E[u'(s^o + s^o r + w\pi - s^o \pi)\pi]$$

In an experiment we normalized wealth to unity, that is $w = 1$, let the utility function take the function form $u(c) = c^\gamma$ for some $\gamma > 0$, set the interest rate r to zero and assumed π is uniformly distributed between $\underline{\pi}$ and $\underline{\pi} + 1$ for some real number $\underline{\pi}$. Then

$$\begin{aligned} E[u'(s^o + s^o r + w\pi - s^o \pi)] &= \int_{\underline{\pi}}^{\underline{\pi}+1} \gamma(s^o + \pi - s^o \pi)^{\gamma-1} d\pi \\ &= \left[\frac{(s^o + \pi - s^o \pi)^\gamma}{1 - s^o} \right]_{\underline{\pi}}^{\underline{\pi}+1} \end{aligned}$$

$$\begin{aligned} E[u'(s^o + s^o r + w\pi - s^o \pi)\pi] &= \int_{\underline{\pi}}^{\underline{\pi}+1} \gamma\pi(s^o + \pi - s^o \pi)^{\gamma-1} d\pi \\ &= \left[\gamma(s^o + \pi - s^o \pi)^{\gamma-1} \right]_{\underline{\pi}}^{\underline{\pi}+1} - \int_{\underline{\pi}}^{\underline{\pi}+1} \gamma(s^o + \pi - s^o \pi)^{\gamma-1} d\pi \\ &= \left[\gamma(s^o + \pi - s^o \pi)^{\gamma-1} \right]_{\underline{\pi}}^{\underline{\pi}+1} - E[u'(s^o + \pi - s^o \pi)] \end{aligned}$$

Equating the two expressions we obtain an equation that relates γ , the relative coefficient of risk aversion, with $\underline{\pi}$, the minimal rate of return on the pension fund, to s^o , the optimal share of wealth deposited in a savings account. Figure 3.5 illustrates the optimal allocation of wealth between savings and the pension fund in this specialization. The horizontal axis indicates different values of $\underline{\pi}$, the vertical axis measures γ , and the isoquants, and isoregions indicate the optimal values of s^o for $(\underline{\pi}, \gamma)$ coordinate pairs.

Overlaying the table are the results of an experiment in which

6 Testing the Expected Utility Hypothesis

The examples above illustrate how the assumption of expected utility helps to rationalize choices made under uncertainty. Indeed the use of expected utility analysis to study attitudes towards risk is well established in economics and finance. Nevertheless the hypothesis is questioned, partly because experimental evidence does not fully support the axioms that justify it. There are essentially three parts to this question. First, do decision makers understand the laws of probability? Second, do players act as if only the outcome of uncertain events matter, or have they preferences over how a lottery is conducted? And finally, do players treat outcomes as mutually distinct events? In this section we review its theoretical basis of the expected utility hypothesis and elaborate on the experimental methodology used to test them.

6.1 Probability law

In other words are people cognizant of logical implications of the facts that no probability is negative, and that the probabilities over all the elements in a partitioned outcome space sum to one.

Ambiguity defining the lotteries

We don't always know the probabilities of the different outcomes, and that can affect the choices we make. However the fact that the subjective probabilities that rational experimental subjects form over the outcomes over the outcomes must sum to one generates some testable restrictions on their behavior. The following experiment shows that these restrictions are not always satisfied in laboratory sessions.

Ellsberg paradox

6.2 Simple and compound lotteries

A compound lottery is formed from $J + 1$ lotteries. The first lottery simply determines which one of the remaining lotteries is played. We also interpret linear combinations as lotteries in which the probabilities of the compound lottery is found by taking the weighted probabilities of the original lotteries. Thus the lottery

$$\phi L_i + (1 - \phi)L_j \equiv (\phi\pi_{1i} + (1 - \phi)\pi_{1j}, \dots, \phi\pi_{Ki} + (1 - \phi)\pi_{Kj})$$

ascribes a probability of $\phi\pi_{ki} + (1 - \phi)\pi_{kj}$ to winning the prize p_K for each $k \in \{1, \dots, K\}$. Are compound lotteries treated the same way as reduced lotteries? We can test to see if subjects switch their preferred lottery, depending on whether they are certain they have the choice or not. This test directly compares compound with simple lotteries.

The previous slides define simple lotteries. A compound lottery is defined by forming a lottery over several other lotteries. We might consider K lotteries denoted by L_k where $k = 1, 2, \dots, K$. The probability of lottery L_k occurring is given by q_k . The probability of outcome l occurring is then:

$$p_{l1}q_1 + p_{l2}q_2 + \dots + p_{lK}q_K$$

where p_{kl} is the probability that lottery k yields outcome l .

A reduced lottery

For example if the probability that you will be neglectful is 0.4, the probability of you being lazy is 0.3 and the probability of building family capital is 0.3, then you are facing compound lottery of how you behave, and then how that affects household decision-making for the summer. A reduced lottery can be formed by calculating the odds of each outcome occurring from playing the compound lottery.

Are compound and reduced lotteries fundamentally different?

It is useful to know whether people are indifferent between playing in reduced lotteries and the compound lotteries which generated them. We consider the following choices over the lotteries, which seek to reveal whether subjects inherently prefer one or the other type. Test of expected utility (we could use the discrete form here)

Problem 1: The decision maker chooses between two simple lotteries: Option A: $(q,0)$ and option B: $(0,1)$.

Problem 2: The decision maker faces a lottery in which, with probability $(1-r)$, she receives x_3 and, with probability r , she faces a subsequent choice between two options, each of which is a simple prospect: Option A: $(q,0)$ and option B: $(0,1)$.

The decision maker faces a lottery in which, with probability $(1-r)$, she receives x_3 and, with probability r , she received one of the options listed below, each of which is a simple prospect. She is required to choose which option to receive before the initial lottery is resolved. Option A: $(q,0)$ and option B: $(0,1)$. Test of expected utility continues:

Problem 4: The decision maker faces a choice between two compound lotteries:

Option A: First stage gives x_3 with probability $(1-r)$ and the simple prospect $(q,0)$ with probability r ;

Option B: First stage gives x_3 with probability $(1-r)$ and the simple prospect $(0,1)$ with probability r .

Problem 5: The decision-maker chooses between two simple lotteries: Option A: $(rq, 0)$, Option B: $(0,r)$

6.3 Independence axiom

In addition to knowing the laws of probability and caring only about consequences of play, expected utility maximizers also adhere to the independence axiom. The independence axiom states that preference over lotteries, or more generally probability distributions over uncertain outcomes, should not be reversed when an independent event is introduced that makes playing either lottery equally less likely.

We motivate the independence axiom with the following example. An employee in a large firm spends her free weekends outside when the weather is fine, but inside when it rains. The Wednesday evening forecast for the upcoming weekend put the odds of fine to rainy weather at four to one. When these odds fall to three to two on Thursday evening the employee is disappointed. In other words she prefers lottery $L_1 \equiv (\frac{4}{5}, \frac{1}{5})$ to lottery $L_2 \equiv (\frac{3}{5}, \frac{2}{5})$, where the first element in $L_i \equiv (p_{i1}, p_{i2})$ is the probability of being outdoors on the upcoming weekend, and p_{i2} is the probability of

being indoors. Appealing to the notation we developed in the previous chapter, $L_1 \succeq L_2$. We illustrate the independence axiom by comparing this scenario with another hypothetical situation facing the same person. After seeing the Wednesday evening forecaste of L_1 , next morning her inflexible and unapproachable boss announces that she might be required to work on an office project (inside) over the (entire) weekend. The independence axiom is violated if employee is relieved when she sees Thursday evening's revised forecaste of L_2 , doubling the likelihood of weekend rain. An expected utility maximizer who prefers L_1 to L_2 is disappointed by Thursday's forecaste regardless of the chance of being randomly assigned to office work over the weekend, because whether she prefers clear skies to rain does not depend on how likely the weather matters.

The dual combination of uncertainty about the weather conditions, and being assigned office work or not, can be modeled as a composite lottery. Let $L_3 \equiv (0, 1)$ denote the degenerate lottery of being indoors, a certainty if the boss requires the employee to work over the weekend. Supposing θ is the probability of having a free weekend, then the composite lottery facing the employee on Thursday afternoon is $\theta L_1 + (1 - \theta)L_3$, whereas the composite lottery facing her on Friday morning before the boss resolves some of the uncertainty is $\theta L_2 + (1 - \theta)L_3$. Together, the premise that only consequences matter, and that the independence axiom holds, imply that if $L_1 \succeq L_2$ then for all θ in the unit interval:

$$\begin{aligned} \left(\frac{4\theta}{5}, \frac{5-4\theta}{5} \right) &= \theta \left(\frac{4}{5}, \frac{1}{5} \right) + (1 - \theta) \left(\frac{0}{5}, \frac{5}{5} \right) \\ &= \theta L_1 + (1 - \theta)L_3 \\ &\succeq \theta L_2 + (1 - \theta)L_3 \\ &= \theta \left(\frac{3}{5}, \frac{2}{5} \right) + (1 - \theta) \left(\frac{0}{5}, \frac{5}{5} \right) \\ &= \left(\frac{3\theta}{5}, \frac{5-3\theta}{5} \right) \end{aligned}$$

As this inequality suggests, the independence axiom imposes strong refutable predictions on behavior under uncertainty.

More generally consider lotteries defined on the space of K (finite) outcomes, with $L = (p_1, \dots, p_K)$, where $0 \leq p_k \leq 1$ for all $k \in \{1, \dots, K\}$, and $p_1 + \dots + p_K = 1$. The preference ordering of the player satisfies the independence axiom if and only if the following condition holds. For all lottery triplets (L_h, L_i, L_j) and any probability $\phi \in [0, 1]$, if $L_i \succeq L_j$ then

$$\phi L_i + (1 - \phi)L_h \succeq \phi L_j + (1 - \phi)L_h$$

This axiom says that in order to rank the two compound lotteries $\phi L_i + (1 - \phi)L_h$ and $\phi L_j + (1 - \phi)L_h$, it suffices to know how the two simple lotteries L_i and L_j are ordered. Augmenting two lotteries, L_i and L_j by introducing a new lottery L_h that down weights each of the original lotteries to ϕ does not affect the original ranking.

Intuitively appealing, yet questioned by some, the independence axiom yields

powerful refutable predictions. To illustrate these, consider the position of a manager who has just played a major role in restoring the profitability of his operating unit. One of three outcomes will occur. He will either be promoted to manage a larger operating unit, or receive a bonus and be asked to remain in charge of the unit to see if it can become even more profitable, or his efforts will be ignored. The probability of these mutually exclusive events occurring are respectively denoted by p_1 , p_2 , and p_3 , so we may define a lottery over these three events by the triplet (p_1, p_2, p_3) . Regarding his preferences over lotteries on these events, the manager certainly prefers a promotion to a bonus, and being ignored is the worst outcome of all. We interpret the symbolic expression $L_i \succ L_j$, defined as the two events $L_i \succeq L_j$ but $L_j \not\succeq L_i$, to mean that the lottery L_i is strictly preferred to L_j . Then in terms of the newly defined notation, the manager's preferences can be expressed as $(1, 0, 0) \succ (0, 1, 0) \succ (0, 0, 1)$.

Figure 4.1 depicts lotteries defined on the three outcomes. The horizontal axis measures the probability of receiving a bonus, while the vertical axis measures the probability of being promoted. Thus the point $(\frac{1}{2}, \frac{1}{4})$ depicts the lottery $L_1 = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, while the origin $(0, 0)$ represents being ignored for sure. The figure also graphs several indifference curves for the manager: every pair of points L_i and L_j on the same indifference curve satisfies both $L_i \succeq L_j$ and $L_j \succeq L_i$. To explain why they are parallel, for any linear combination ϕ of L_i and L_j note:

$$L_h \equiv \phi L_i + (1 - \phi)L_j \succeq \phi L_i + (1 - \phi)L_i = L_i$$

and vice versa, the preference inequality following from the independence axiom. Hence L_h is on the same indifference curve as L_i and L_j . Because the origin represents the least desirable outcome the lines have a negative slope, and since $(0, 1)$ is less desirable than $(1, 0)$, their slope is less than one in absolute value.

Denoting the slope of the indifference lines by α , an expected utility function for the manager is $p_2 - \alpha p_1$. In this way Figure 1 depicts how utility functions can be constructed for people who obey the independence axiom. More generally, if a rational person obeys the independence axiom then we can construct a utility function to represent his preferences that is linear in the probability weights. This means that in the finite case where there are K outcomes a person's utility function can be modeled as:

$$\sum_{k=1}^K p_k u_k$$

or more generally as

$$\int_X u(x) dF(x)$$

where x is an element in the outcome space X , the real valued mapping $u(x)$ is the utility value placed on outcome $x \in X$, and $F(x)$ is a probability distribution function for the random variable.

Testing the expected utility theorem

There are two tests of the independence axiom, and by implication, the expected utility theorem: We can test whether the indifference curves over simple lotteries from parallel lines or not

7 Dominance

Deterministic

If alternative A dominates Alternative B deterministically, then we know for sure that we will always be better off with alternative A than with alternative B (completeness). This is the weakest of axioms above transitivity and completeness.

The concept of Stochastic Dominance

In this section we show that the postulates of Expected Utility lead to a definition of two alternative concepts of dominance which are weaker and thus of wider application than the concept of state-by-state dominance. These are of interest because they circumscribe the situations in which rankings among risky prospects are preference-free, i.e., can be defined independently of the specific trade-offs (between return, risk and other characteristics of probability distributions) represented by an agent's utility function.

7.1 First order stochastic dominance

Consider two different probability distributions $F(x)$ and $G(x)$. That is $F(x) \geq G(x)$ for some real number x . We say that F first-order stochastically dominates G if and only if $F(x) \geq G(x)$ for all x . This is formally equivalent to saying that if x is a random variable drawn from $G(x)$, if y is a random variable that only takes on positive values, and $F(z)$ is the probability distribution function for $z = x + y$, then F first-order stochastically dominates G .

First-order stochastic dominance and expected utility

Now consider a person who obeys the expected utility hypothesis, obtaining expected utility $EF[u(x)]$ from playing lottery F . Also suppose $u(x)$ is increasing in x . Then we can prove that if F first-order stochastically dominates G , then he prefers F to G , that is $vG \leq vF$.

An example

Notice that the the third lottery first-order stochastically dominates the second.

States of nature	1	2	3
Probabilities	.4	.4	.2
Investment Z_1	10	100	100
Investment Z_2	10	100	2000
	$EZ_1 = 64, \sigma_{z_1} = 44$		
	$EZ_2 = 444, \sigma_{z_2} = 779$		

Table 2.1: Sample Investment Alternatives




Figure 2.8: First Order Stochastic Dominance for the Investment case

Link this to the decision tree examples:

Definition: Let $F_A(\tilde{x})$ and $F_B(\tilde{x})$, respectively, represent the cumulative distribution functions of two random variables (cash payoffs) that, without loss of generality assume values in the interval $[a, b]$. We say that $F_A(\tilde{x})$ first order stochastically dominates (FSD) $F_B(\tilde{x})$ if and only if $F_A(\tilde{x}) \leq F_B(\tilde{x})$ for all $x \in [a, b]$.




Figure 2.9: First Order Stochastic Dominance: A more General Representation

Theorem: Let $F_A(\tilde{x}), F_B(\tilde{x})$, be two cumulative probability distribution for random payoffs $\tilde{x} \in [a, b]$. Then $F_A(\tilde{x})$ FSD $F_B(\tilde{x})$ if and only if $E_A U(\tilde{x}) \geq E_B U(\tilde{x})$ for all non decreasing utility functions $U(\cdot)$.

If alternative A has first order stochastic dominance relative to alternative B, then you will always have a higher probability of getting a better deal with alternative A than with B. Remember the main assumption: $U' > 0$ (you prefer more money to less).

Implications of First Order Stochastic Dominance

- FSD also implies mean dominance and geometric mean dominance.
- If A has first order stochastic dominance with respect to B, then the mean of the distribution of A is higher than the mean of the B distribution.
- FSD implies Second order stochastic dominance, implies third order...ect.

Decision Tree example

7.2 Second order stochastic dominance

Second-order stochastic dominance

Consider two different probability distributions $F(x)$ and $G(y)$ with the same mean. That is $EF[x] = EG[x]$. We say that F second-order stochastically dominates G if and only for all t : This is saying that if x is a random variable drawn from $G(x)$, if y is a random variable with mean 0, and that probability distribution function for $z = x + y$ is $F(z)$, then F second-order stochastically dominates G.

Second-order stochastic dominance and expected utility

Assume a person follows the expected utility hypothesis, and thus obtains expected utility $EF[u(x)]$ from playing a lottery F. We now assume $u(x)$ is concave increasing in x , implying that the person is risk averse. Suppose $EF[x] = EG[x]$. Then a risk averse person prefers F to G, that is $vG \leq vF$, if F second-order stochastically dominates G.

Example: we should use our lotteries here

Investment 3		Investment 4	
Payoff	Prob.	Payoff	Prob.
4	0.25	1	0.33
5	0.50	6	0.33
12	0.25	8	0.33

Table 2.2: Two independent Investments




Figure 2.9: Second Order Stochastic Dominance Illustrated

Definition: Let $F_A(\tilde{x})$ and $F_B(\tilde{x})$, be two cumulative probability distributions for random payoffs in the interval $[a, b]$. We say that $F_A(\tilde{x})$ second order stochastically dominates (SSD) $F_B(\tilde{x})$ if and only if for any x :

$$\int_{-\infty}^x [F_B(t) - F_A(t)] dt \geq 0$$

(with strict inequality for some meaningful interval of value of t).

Theorem: Let $F_A(\tilde{x})$, $F_B(\tilde{x})$, be two cumulative probability distribution for random payoffs $\tilde{x} \in [a, b]$. Then $F_A(\tilde{x})$ SSD $F_B(\tilde{x})$ if and only if $E_A U(\tilde{x}) \geq E_B U(\tilde{x})$ for all non decreasing and concave $U(\cdot)$.

- If A has a SSD with respect to B then if you are a rule person (i.e. $U' > 0$) + you are risk averse, you will always choose alternative A to B.

Main assumption for SSD: If you are a rule person + you are risk averse : i.e. $U' > 0$, $U'' < 0$.

More or less risky \cong mean preserving spread

$$\tilde{x}_B = \tilde{x}_A + \tilde{Z}$$




Figure 2.10: Mean preserving spread

Theorem: Let $F_A(), F_B(),$ be two distribution functions defined on the same state space with identical means. Then the following statement is equivalent:

- $F_A(\tilde{x})$ SSD $F_B(\tilde{x})$

- $F_B(\tilde{x})$ is a mean preserving spread of $F_A(\tilde{x})$ in the sense of $\tilde{x}_B = \tilde{x}_A + \tilde{Z}$.

8 Summary

Behaving strategically means that the individual in question acquires and processes information in a purposeful manner with a view to furthering her own ends, recognizing those ends may conflict with the goals of other people, and taking account of how they react to her decisions. This chapter abstracted from the strategic aspects to amplify the premise that strategic play is based on rational choice under uncertainty.

We started out with a very simple hypothesis about behavior under uncertainty, that in games for a single player, individuals maximize the expected value of their wealth. Computing how expected wealth maximizers should use their information and its value to them, amounts to calculating the expected value from following different routes. The examples we showed that

Although expected wealth maximization is a useful assumption to make in some situations, it seems inappropriate for others. A generalization of expected wealth maximization is expected utility maximization: although more wealth is preferred to less, people might not be indifferent between gambling some of their wealth, versus avoiding the gamble and taking its expected value instead. In this case their utility is not linear in wealth, but might possibly be represented by utility function, nonlinear but increasing in wealth. If the function is concave they are risk averse, and prefer the

certain expected value of a gamble for sure over the random payoff implied by the gamble. If the function is convex, they accept all actuarially fair bets. Then we provided some experiments which determined whether players are expected value maximizers or not. The experiments also reveal the subject's attitudes towards risk, and provide a way of recovering their utility function if they are expected utility maximizers. Once we know their utility function, we can readily adapt the methods devised for valuing information. By definition, an expected wealth maximizer is risk neutral, but an expected utility maximizer might be risk averse or risk loving. The examples on insurance and financial investing illustrated how easy it is to work with expected utility formulations of preferences.

Whether a person is an expected utility maximizer or not depends on whether they know the laws of probability, and obey the three axioms that define rational behavior. To check whether laid out on the three axioms, complete and transitive preferences, plus the independence axiom, that justify expected utility maximization. We showed We also provided several tests of the expected utility hypothesis for investigating whether experimental subjects obey the three axioms, or even the laws of probability.

Finally we argued that all is not lost if people are not expected utility maximizers. In this case we might expect subjects to show preferences towards payoff probability distributions that first order stochastically dominate. Moreover risk averse players prefer distributions that also second order dominate others.

The advantage of assuming a person is an expected value maximizer is that you only need to know the probability distribution in order to know where preferences will be. On the other hand