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 very 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.

Section 3 demonstrates how to test the expected wealth maximization hypothesis using experimental methods. There 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 expected wealth maximization is a useful assumption to make in some situations, it seems inappropriate for others. 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. The extra generality comes with an additional burden, for the decision maker's attitude towards risk partly determines the value she places information. If she is an expected utility maximizer, then each possible outcome that might occur is assigned a utility, and the expected utility is found by summing (or integrating) the weighted utilities from the payoffs, where the weights are just the respective probabilities of occurrences. 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 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 securites and portfolio choice in financial markets in Section 5 illustrate why risk attitudes guide decisionmaking on these issues. We show that a risk averse expected utilty maximizer fully insures her property at acturially fair rates, but invests some of her assets in risky securites 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 indpendence. 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 operationalize a theory that relaxes the independence axiom but still provides us with predictions about how indiviudals 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 chpater by defining several notions of dominance between probability distributions, and testing whether subjects form the same rankings.

Expected Wealth Maximization

A first approach to choice under uncertainty is to attach a monetary value to each possible outcome, and implement decisions that maximize the expected wealth from the outcomes, where the expected wealth from a decision plan is defined as the weighted sum of the monetary values of the outcomes, and the weight placed upon each outcome is the probability of its occurence under that decision plan. Under the hypothesis of wealth maximization, the optimal acquisition and use of information can be calculated in a straightforward way. We consider three examples to illustrate the applicability of this approach.

Foreign Investment

Figure 3.1 depicts the extensive form of a game in which a multinational is deliberating over the prospect of acquiring a plant in new, 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.

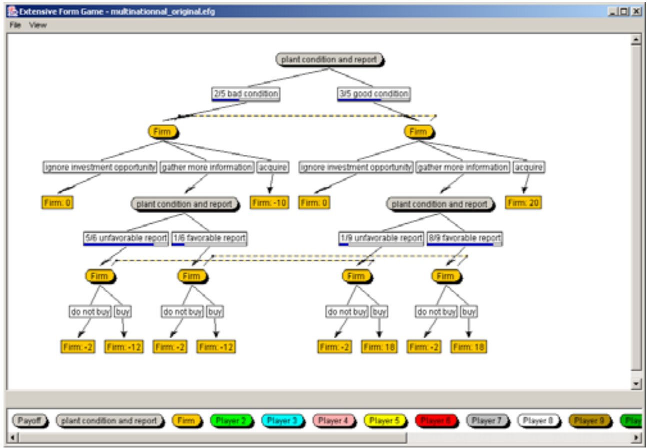

Figure 3.1
Direct foreign investment

Although 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 from in 3.2.

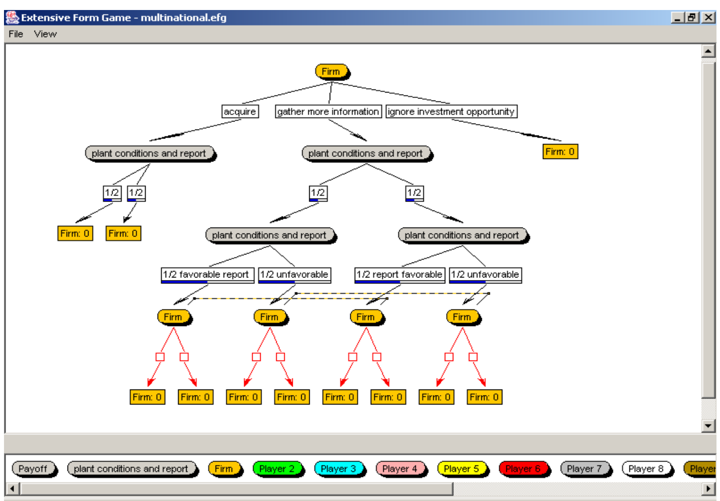

Figure 3.2
Direct foreign investment redrawn

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

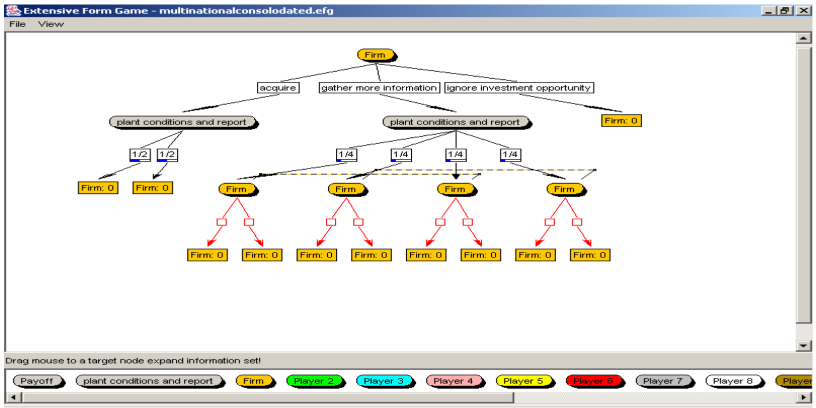

The expected value of not inspecting and buying is found by solving the subproblem depicted in 4.22.

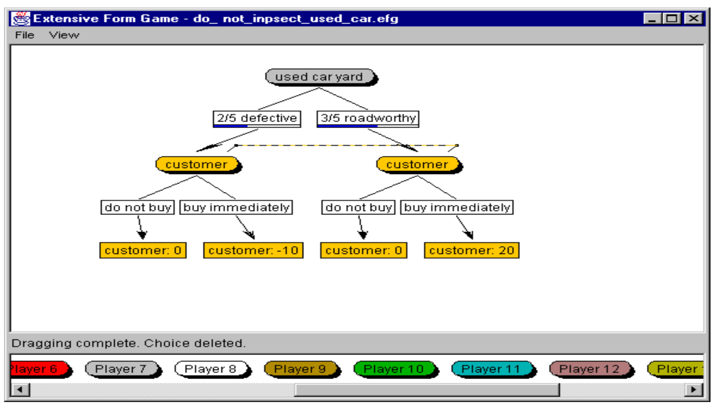

Figure 5.2
Calculating the of not inspecting

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.

Product Testing

In this final example we consider a pharmaceutical company which cannot market a drug unless it has passed guidelines set out by the Federal Drug Administration. we suppose there are two tests that must be passed before marketing is permitted.

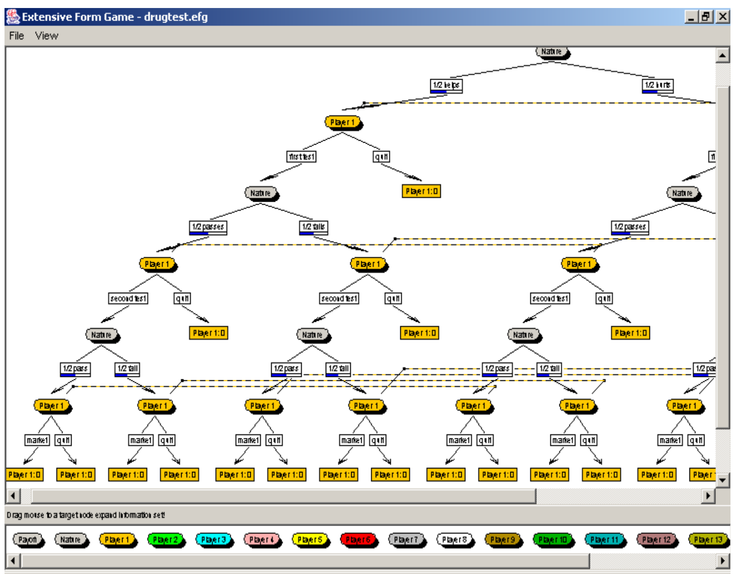

Figure 4.
Product Testing

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 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 dpeartment is $x - w_n$. Each job candidate $n \in \{1, \ldots, N\}$ 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, identically and unifromly distribributed over all candidates $n \in \{1, \ldots, N\}$ with support $[\underline{v}, \bar{v}]$. Thus the probability that $v_n \leq w$ is $(w - \underline{v})/(\bar{v} - \underline{v})$. We assume that $x > \underline{v}$; otherwise the department would have no interest in going through the hiring process.

First consider the case $N = 1$. There is no reason to offer more than $\bar{v}$ since every candidate would accept $\bar{v}$ expected value to the department from making an offer of w_1 is the probability of acceptance multiplied by the net value conditional on acceptance. Accordingly the department chooses w_1 to maximize:

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

subject to the constraint that $w_1 \leq \bar{v}$. From the first order condition to this problem, we derive the interior solution:

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

Noting that the division is assured of hiring the candidate if it offers the candidate 1, the optimal offers is therefore

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

Now consider a two period extension of the optimization problem. In the last period the value of the game is V_1, so in the preceding period, the firm sets r to maximize

$$V_2 = \max_{w_2} \{(x - w_2)F(w_2) + [1 - F(w_2)]\beta V_1\}$$

More generally

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

Differentiating with respect to w_n we obtain

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

Comparing the first order condition for w_2^o with w_1^o we see that the extra term on the left side, βV_{n-1}, enters negatively.

Returning to the uniform distribution case, so the first order condition simplifies to

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

which implies

$$\beta V_{n-1} = x - 2w_n^o$$

But

$$\begin{aligned} V_n &= (x - w_n^o)F(w_n^o) + [1 - F(w_n^o)]\beta V_{n-1} \\ &= (x - w_n^o)w_n^o + [1 - w_n^o]\beta V_{n-1} \\ &= (x - w_n^o)w_n^o + [1 - w_n^o](x - 2w_n^o) \end{aligned}$$

Substituting into the equation for V_{t+1} we thus obtain

$$\begin{aligned} 2r_{t+1} - 1 &= (1 - r_t)r_t + r_t(2r_t - 1) \\ 2r_{t+1} &= r_t^2 + 1 \end{aligned}$$

which is a first order difference equation with solution

Experiment. We consider several sequential offer games. Consider a game of sequential offers in which there is a probability that the game might end after any given period if a suitable candidate does not accept. Compare your results with a game in which continues for exactly N periods

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 planty 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 welath maximization hypothesis, and explains why .

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 L possible prizes, whose values are denoted by x_1 through x_L, and L probabilities, denoted by p_1 through p_L where $0 \leq p_l \leq 1$ for all $l \in \{1, 2, \ldots, L\}$ and

$$\sum_{l=1}^{L} p_l = 1$$

Summarizing, if a person plays the lottery, outcome $l \in \{1, 2, \ldots, L\}$ occurs with probability p_l, and in that event she receives a prize of x_l. For expository convenience we rank the value of the prizes in ascending order. Thus $x_1 = \min\{x_1, \ldots, x_L\}$ and $x_L = \max\{x_1, \ldots, x_L\}$.

How much a subject 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. Since the value of her prize from participating in the lottery invariably lies between x_1 and x_L, we assume $x_1 \leq v_n \leq x_L$. That is, we impose the principle of weak dominance upon preference orderings, a topic discussed more fully in the last section of this chapter, and also Chapter 8.

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 a strong incentive to bid less than her certainty equivalent. To induce experimental subjects to report their reservation values truthfully, we adapt the second price auction mechanism, analyzed in Chapter 15, 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 $(x_1, \ldots, x_L)$ and their associated probabilities $(p_1, \ldots, p_L)$. Then each subject $n \in \{1, 2, \ldots, N\}$ 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 joining $\min\{x_1, \ldots, x_L\}$ to $\max\{x_1, \ldots, x_L\}$. 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 distirbtuion $(p_1, \ldots, p_L)$. If $b_n < c_n$ the subject is denied a lottery ticket, and she receives no payoff.

In this experiment you are asked to how much you are willing to pay for a lottery called L when you 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 £ 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

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 outocmes of this bidding strategy with the alternative of bidding v_n. There are three possibilites 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 $c_n.$ In other words bidding v_n weakly dominates submitting a lower bid.

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 possiblity 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 weakly dominates submitting a higher bid. Therefore v_n is the optimal bid, so in this lottery each subject fully reveals her certainty equivalent through their bid.

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?

Are bidders expected wealth maximizers?

Here we report an experiment to see whether we can reject the hypotheis 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.

The Expected Utility Hypothesis

Under the expected utility hypoethesis a utility is attached to all the lottery prizes $(p_1, \ldots p_K),$ and Let $(u_1, \ldots u_K)$ denote the vector of utilities that. Then $L_i \succeq L_j$ if and only if

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

The key implication of this remarkable theorem is that if the independence axiom is satisfied then we can represent preferences over uncertainty by re-scaling payoffs through a utility function and then taking the expectation, in other words using expected utility instead of expected value as the objective function.

Consider a single argument such as a numeraire payoff, as we have done in representing the payoff in the extensive form representation, and suppose that independence axiom applies. The we can continue as before just by re-scaling

At the same time, we are bound to acknowledge that to the extent that the independence axiom is invalid the prediction based on expected utility theory are discredited.

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.

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

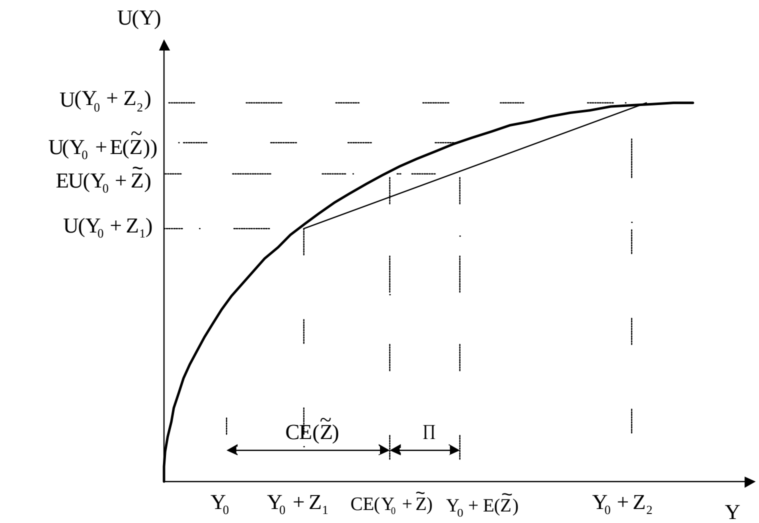

Figure 3. 1
Certainty Equivalence and Concavity

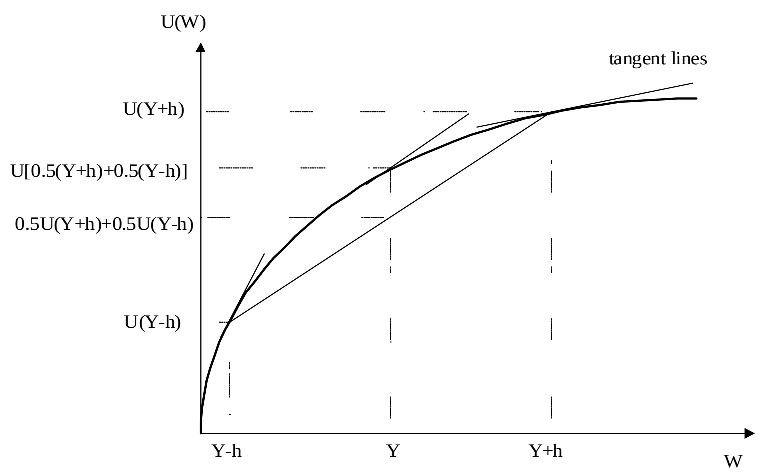

Figure 3.1: A Strictly Concave Utility Function

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.

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.

Expected Utility Maximization

Expected utility is applied in many finance and insurance applications

Insurance

Insurance is a pervasive feature of the consumer durable goods and housing sectors, but there are other reasons why are 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 unfacorable 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 prurchases 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, prefering 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 $u(w) = -\exp(\alpha w)$ and f and $p = 0.1.$

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 $\pi,$ 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 exceed 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 opitmal 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 ignnores 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

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

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

Equating the two expressions we obtain an equation that relates $\gamma,$ 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 $\gamma,$ 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

Portfolio choice

Another area that draws heavily on theories about how people deal with uncertainty is asset pricing. This section shows how the theory of expected utility can be applied to a portfolio choice problems. In this game with nature, we assume there are K securities that have uncertain payoffs as well as a security, such as cash or a bond, that has a certain payoff. At the beginning of the game, a player is endowed with nonnegative amounts of the risky securities, summarized by the vector $(\overline{q}_1, \ldots, \overline{q}_K),$ plus some quantity of the risk free security, $\overline{q}_0.$ She chooses how many securities of each type to buy subject subject to an overall budget constraint that total spending on securities she buys is less than or equal to her total wealth endowment plus the revenue generated from the sale of the securities she sells. Write p_k for the (strictly positive) price of the k^{th} security and q_k for the held. We normalize the return of the risk free security setting $p_0 = 1,$ meaning the price of all the risky securities are stated in terms of the normalized security, and impose the assumption that q_k is nonnegative, thus ruling out short sales. Therefore the budget constraint may be expressed as

$$\sum_{k=0}^{K} p_k(q_k - \overline{q}_k) \leq 0$$

At the end of the game, each risky security realizes a return of π_k and the player receives $\pi_k q_k$ from her holdings of that security. The riskless security yields a payoff of $\pi_0.$ Consequently at the end of the game she receives in total:

$$\sum_{k=0}^{K} \pi_k q_k$$

The vector of returns $(\pi_1, \ldots, \pi_K)$ is not know by the player before she makes her portfolio choice, but she does know the joint probability distribution for this multivariate random variable. Letting μ_k denote the mean of π_k, an expected wealth maximizer would maximize:

$$E\left\{\sum_{k=0}^{K} \pi_k q_k\right\} = \sum_{k=0}^{K} \mu_k q_k$$

By inspection we can see that if all the mean returns are distinct, then the optimal portfolio is achieved by specializing in the security with highest mean return. The fact that that investors in the stock market do not typically specialize to such an extent has lead analysts to conclude that investors are risk averse who hold diversified portfolios to limit their risk. Accordingly we now suppose that the player maximizes the expected value of $u(c)$, a concave increasing utility functionin im her consumption c, defined as:

$$c = \sum_{k=0}^{K} \pi_k q_k$$

She picks the vector $(q_1, \ldots, q_K)$ subject to the wealth of her initial endowment and the $K+1$ nonegativity constraints $q_k \geq 0$. The Lagrangian for this problem is

$$E\left[u\left(\sum_{k=0}^{K} \pi_k q_k\right)\right] + \lambda \sum_{k=0}^{K} p_k(q_k - \overline{q}_k) + \sum_{k=0}^{K} \lambda_k q_k$$

The first order condition for this problem is

$$E[u'(c)\pi_k] + \lambda p_k + \lambda_k = 0$$

with $K+2$ complementary slackness conditions

$$\lambda_k q_k = 0$$

$$\lambda \sum_{j=0}^{K} p_j\left(q_j - \overline{q}_j\right) = 0$$

It is easy to show that because $u(c)$ is increasing and there is a risk free asset, the budget constraint will be met with equality. Thus λ is strictly positive and can be interpreted as the marginal utility of wealth. Supposing the first and second securities form part of the investor's portfolio, the complementary slackness conditions imply

$$\lambda_k = \lambda_j$$

and

$$E\left[u'(c)\left(\frac{\pi_k}{p_k} - \frac{\pi_j}{p_j}\right)\right] = 0$$

This is the fundamental equation that characterizes the choice of alternative securities that are liquidated, or at least reassessed, at the same time. The expected payoffs weighted by the marginal utility of consumption for each outcome are equated across those securities that are actually purchased. Suppose, for example, one security only pays off in two out of a possible ten outcomes; the more the investor buys of that security, the lower the marginal utility she obtains, and hence the lower the weight that is placed on the payoffs in those two states, both for that security, and any other security that pays off in those same outcomes.

To operationalize this framework as an experiment, we now suppose that the investor's utility function is quadratic:

$$u(c) = \alpha_1 c - \alpha_2 c^2$$

where α_1 and α_2 are both positive. We also assume that given the endowment the maximum consumption achieved in any state is bounded above by $\alpha_1/2\alpha_2$, an assumption we make to avoid worrying about the entirely unrealistic possibility that consumers prefer less wealth to more. In this case for all interior $k \in \{1,\ldots,K\}$

$$E[c(\pi_k - \pi_0)] = 0$$

Define the return on the individual portfolio as

$$\pi_m = \frac{c}{w + \sum_{k=1}^{K} \bar{q}_k}$$

so that

$$E[\pi_m(\pi_k - \pi_0)] = 0$$

or

$$E[\pi_m \pi_k] = E[\pi_0]E[\pi_m]$$

Subtracting $E[\pi_m]E[\pi_k]$ from both sides of this equation yields

$$E[\pi_m \pi_k] - E[\pi_m]E[\pi_k] = -E[\pi_k]E[\pi_m] + E[\pi_0]E[\pi_m]$$

Applying the definition of a covariance to the left side we see that

$$cov(\pi_m, \pi_k) = -E[\pi_k - \pi_0]E[\pi_m]$$

This condition is not only true for each of the individual securities, but also hold for the market index itself, a fact that can be derived directly, or by including the market index as one of the K risky securities in the original definition of the problem. Consequently

$$var(\pi_m) = -E[\pi_m - \pi_0]E[\pi_m]$$

Define the coefficient β_k as the ratio of the covariance of the stock return with the portfolio return to the variance of the portfolio return

$$\beta_k = \frac{cov(\pi_m, \pi_j)}{var(\pi_m)}$$

Combining the two equations to eliminate $E[\pi_m]$ we obtain

$$E[\pi_k - \pi_0] = \beta_k E[\pi_m - \pi_0]$$

In words the excess return on a stock over the risk free rate is proportional in expectation to the excess return of the individuals portfolio, where the coefficient of proportionality is β_k as defined above. This equation holds for every individual in the market who holds quadratic preferences, regardless of the preferences of the other players.

If a person only cared about the mean and variance of his total portfolio, he would only choose amongst those portfolios, which for a given mean, minimized the variance. In teh experimetn we should show how much we lose by only focusing on the traded securities and ignoring nontraded securities

The value of information reconsidered

Acquiring information reduces the uncertainty, and hence can be viewed as partial insurance. This remarks suggests that aperson's attitude towards risk affects their willingness to pay for infomation. 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

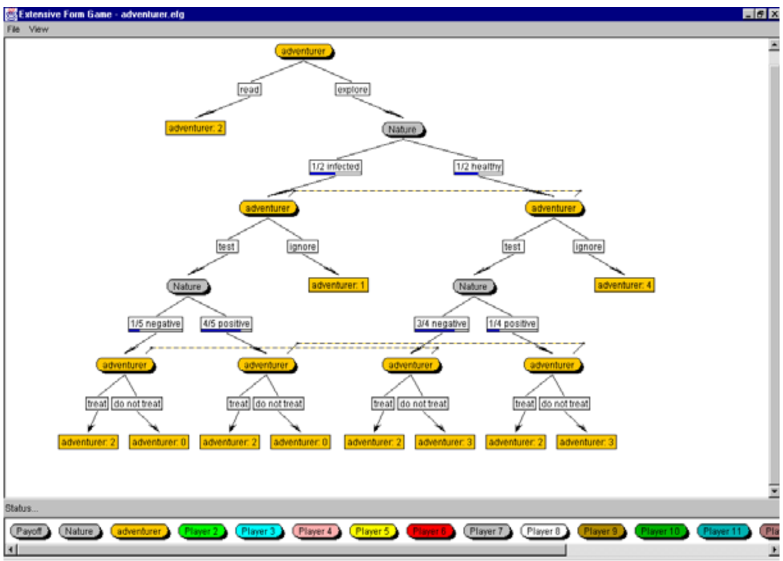

Figure 5.5
Occupational Hazard

A final application to uncertainty. In Figure 3.1 an investor chooses between a stock that as 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 now straightforward to apply he first rule of strategy to calculate the value of having the car inspected.

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.

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

Simple and compound lotteries

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 Lk where k = 1,2, . . . ,K. The probability of lottery Lk occurring is given by qk. The probability of outcome l occurring is then:

$$p1l\ q1 + p2l\ q2 + . . . + pLl\ qK$$

where pkl 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).

Problem2: The decision maker faces a lottery in which, with probability (1-r), she receives x3 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 x3 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 x3 with probability (1-r) and the simple prospect (q,0) with probability r;

Option B: First stage gives x3 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)

Independence axiom

The basis for expected utility theory is the independence axiom. We motivate the axiom with an example that captures its essence.

Suppose an employee in a small firm is planning his next weekend’s activities. In this case there is no preparation involved in either weekend activity. The weather is a little unpredictable but will not be unbearable. There are no other relevant factors. Reflecting on this choice set Friday evening, he decides to tend his outside garden rather than read inside. However just as he makes his decision, the boss phones home to say there is an even chance that he might need to work over the weekend. The probability of this occurring is due to a client’s unusual request and in particular entirely unrelated to the weekend weather. The independence axiom implies that the employee will tend his garden if he is not called into work next day. In other words the fact that the phone call that has halved the chance the employee spends the weekend at home does not affect the priorities he had already determined.

An abstract way to formulate the independence axiom is in terms of lotteries. We suppose there are K possible outcomes, events or prizes denoted $(p_1, \ldots p_K)$ and the j^{th} lottery, called $L_j \equiv (\pi_{1j}, \ldots \pi_{Kj})$, ascribes a probability of π_{jk} to winning the prize p_K

for each $k \in \{1,\ldots,K\}$. Thus $\pi_{jk} \geq 0$ and

$$\sum_{k=1}^{K} \pi_{jk} = 1$$

This describes a simple lottery. To define the independence axiom, we not only need to characterize a simple lottery but also a compound 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

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

ascribes a probability of $\phi\pi_{ki} + (1-\phi)\pi_{kj}$ to winning the prize p_K for each $k \in \{1,\ldots,K\}$.

The example above can be used to illustrate this notation. In the example suppose the probability of poor weather is $\frac{1}{10}$, the probability of a client making an unusual request is $\frac{7}{10}$, the probability that the boy will have the same choice to make on the following day is $\frac{3}{10}$ and the probability of a short queue at the station is $\frac{4}{10}$. In this example there are three prizes, an ice cream, p_1, a front row seat at the top of a double decker, p_2, and nothing p_3. Therefore lotteries in this example are defined over coordinate pairs (π_1, π_2, π_3). The lottery from taking a bus can then be written as $\left(\frac{1}{10}, 0, \frac{9}{10}\right)$ the lottery accepted by walking is $\left(0, \frac{7}{10}\right)$, while the lottery implied by taking teh tube is $\left(0, \frac{4}{10}\right)$. Finally the compound lotteries the boy faces are

$$\begin{bmatrix} \frac{1}{10} & 0 \\ 0 & \frac{4}{10} \end{bmatrix} = \frac{3}{10}\left(\frac{1}{10}, 0\right) + \frac{7}{10}\left(0, \frac{4}{10}\right) = \left(\frac{3}{100}, \frac{28}{100}\right)$$

for walking, and $\left(0, \frac{49}{100}\right)$ for taking the bus. The independence axiom asserts in this case that if $\left(\frac{1}{10}, 0\right) \succeq \left(0, \frac{7}{10}\right)$ then $\left(\frac{3}{100}, \frac{28}{100}\right) \succeq \left(0, \frac{49}{100}\right)$.

A player ranks lotteries in order of preference, and we use the symbol $\succeq$ to denote the ordering. Thus $L_i \succeq L_j$ means that the i^{th} lottery is at least as good as the j^{th}. Loosely speaking L_i, for example, might be the probability distribution for the weather when gardening next weekend, and L_j might denote the probability distribution of the quality of the book that could be read next weekend. (This can be made rigorous by carefully defining exactly what is meant by the events $(p_1, \ldots p_K)$ but there is no need to venture into such detail here.)

The preference ordering of the player satisfies the independence axiom if and only is 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, 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 $\pi_1,$ $\pi_2,$ and $\pi_3,$ so we may define a lottery over these three events by the triplet $(\pi_1, \pi_2, \pi_3).$ We consider his preferences over lotteries on these events. He 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 preferences can be expressed as $(1,0,0) \succ (0,1,0) \succ (0,0,1).$

Figures 4.1 and 4.2 show how to graphically represent $(\pi_1, \pi_2, \pi_3).$ In the first figure

The independence axiom is a sensible premise if you believe there is no fundamental difference between a compound lottery and its reduced lottery. It states the following: Consider any three lotteries, denoted by L1, L2, and L3, plus any number z in the [0,1] interval. Suppose L1 is preferred to L2. Then the simple lottery [z L1 +(1-z) L3] is preferred to [z L2 +(1-z) L3].

Illustrating the independence axiom

Expected utility theorem

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. In other words the independence axiom implies that a person's utility function can be modeled as:

p1u(x1) + p2u(x2) + . . . + pLu(xL)

or more generally as

EF[u(x)]

where F is a lottery or probability distribution over x and EF is the expectations operator.

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

Dominance

Deterministic

If alternative A dominates Alternative B deterministically, then we know for sure that we will always be better off with alternative A then with alternative B (completness). This is the wekest 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.

First order stochastic dominance

Consider two different probability distributions F(x) and G(x). That is F(x) ¹ G(x) for some real number x. We say that F first-order stochastically dominates G if and only F(x) £ 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 £ 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

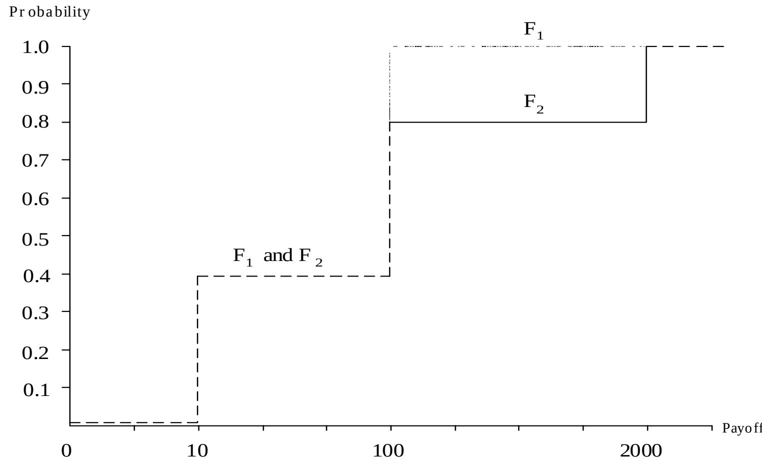

Figure 2.8: First Order Stochastic Dominance for the Investment case

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 $F_A(\tilde{x}) \le F_B(\tilde{x})$for all x $\in [a,b]$.

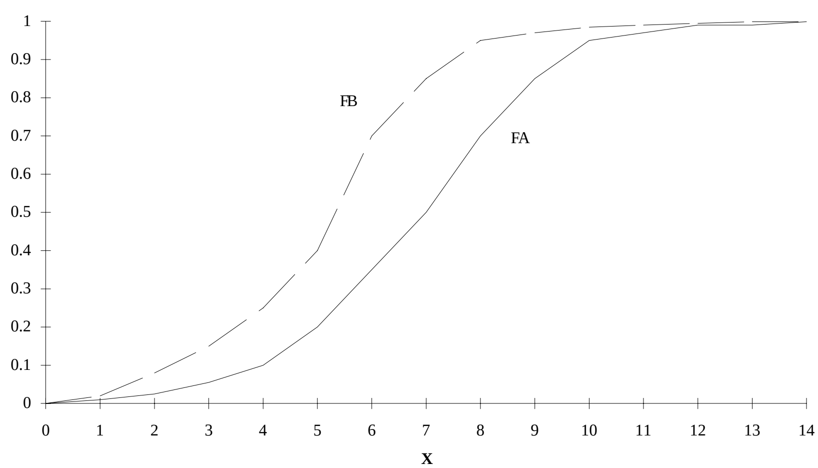

Figure 2.9: First Order Stochastic Dominance: A more General Representation

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 $E_A U(\tilde{x}) \ge E_A U(\tilde{x})$ for all non decreasing utility functions U( ).

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 stocastic 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 ......

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 £ 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

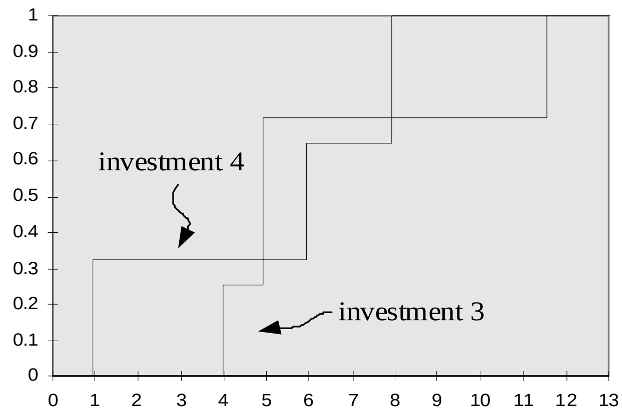

Figure 2.9: Second Order Stochastic Domince Illustrated

Definition: Let $F_A(\tilde{x})$ and $F_B(\tilde{x})$, be two cumulative probabilty 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 $E_A U(\tilde{x}) \geq E_A U(\tilde{x})$ for all non decreasing and concave U( ).

- 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 ≅ mean preserving spread

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

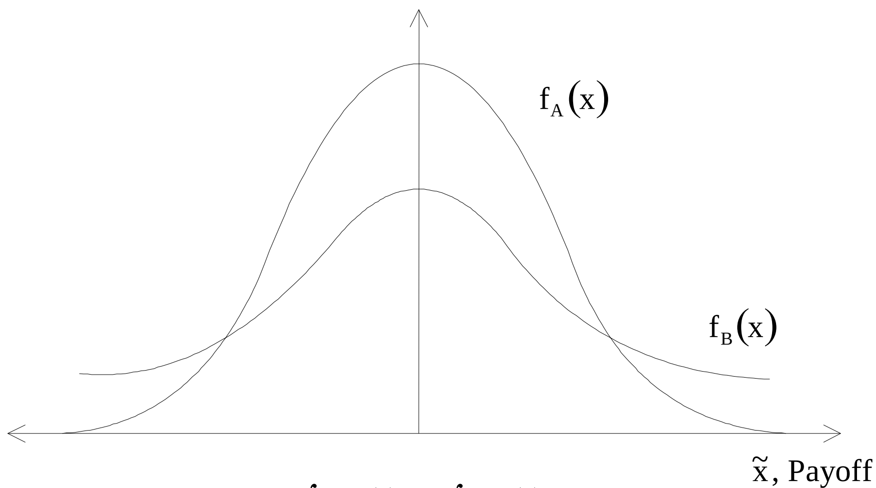

$$\mu = \int x \, f_A(x) dx = \int x \, f_B(x) dx$$

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

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 differnt 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 utitiy is not linear in weatlh, but might possibily be represented by utlity function, nonlinear but increasingin 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 fucntion is convex, they accept all acturially 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 utlity function, we can readily adpat the methods deviced for valuing infomration. By definition, an expected welath maximizer is risk neutral, but an expected utility maximizer might be risk averse or risk loving. The examples on insurance and finanicial 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 the three axioms that define rational behavior. to check whether laid out on the three axioms, complete and transitive preferences, plus the indpendence axiom, that justify expected utilty 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 maxmiaers. In this case we might expect subjects to show preferences towards payoff probability distirubtions that first order stochasiticaly dominate o. Morevoer risk averse plyayers prefer distributions that also second order dominate others.

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