

Notes on Uncertainty and Expected Utility

Ted Bergstrom, UCSB Economics 210A

November 16, 2016

1 Introduction

Expected utility theory has a remarkably long history, predating Adam Smith by a generation and marginal utility theory by about a century.¹ In 1738, Daniel Bernoulli wrote:

“Somehow a very poor fellow obtains a lottery ticket that will yield with equal probability either nothing or twenty thousand ducats.² Will this man evaluate his chance of winning at ten thousand ducats? Would he not be ill-advised to sell this lottery ticket for nine thousand ducats? To me it seems that the answer is in the negative. On the other hand I am inclined to believe that a rich man would be ill-advised to refuse to buy the lottery ticket for nine thousand ducats.”

“... the determination of the value of an item must not be based on its price, but rather on the utility it yields. The price of the item is dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate. Thus there is no doubt that a gain of one thousand ducats is more significant to a pauper than to a rich man though both gain the same amount.”

¹Even when Cournot (1838), Gossen (1858) developed recognizably utility theory in the mid nineteenth century, they gained little attention for at least another half century.

²A ducat contained 3.5 grams of gold. I believe that at the time, a ducat was worth about 1/2 of an English pound sterling. English laborers earned about 20 pounds per year. So annual earnings for a poor fellow might be 40 ducats. If this is the case, 20,000 ducats would be 500 times the annual earnings of a poor fellow. TB

“If the utility of each possible profit expectation is multiplied by the number of ways in which it can occur, and we then divide the sum of these products by the total number of possible cases, a mean utility [moral expectation] will be obtained, and the profit which corresponds to this utility will equal the value of the risk in question.”

Another remarkable feature of Bernoulli’s discussion is his focus on human capital and wealth rather than just current income. Bernoulli proposes that the utility function used to evaluate gambles should be a function of one’s wealth, and not just current income flows.

Bernoulli’s suggests a form for the utility function stated in terms of a differential equation. In particular, he proposes that marginal utility is inversely proportional to wealth. Thus we have

$$\frac{du(W)}{dW} = \frac{a}{W}.$$

for some constant a . We can solve this differential equation to find the function u . In particular, we must have

$$u(W) = a \ln W + b$$

for some constants a and b .

Bernoulli points out that with this utility function, people will be risk averse. Moreover, one’s willingness to pay for the poor fellow’s lottery ticket depends on one’s initial wealth.

To see the power of Bernoulli’s model, let us try some calculations:

Exercise:A) Calculate the wealth of the poorest fellow with logarithmic utility who would buy this ticket for 9000 ducats.

Let X be the fellow’s wealth. He is just indifferent between buying the ticket for \$9,000 or not buying it if

$$a \ln X + b = \frac{1}{2} (a \ln (X + 20,000 - 9,000) + b) + \frac{1}{2} (a \ln (X - 9,000) + b).$$

This expression simplifies to

$$\ln X = \frac{1}{2} \ln (X + 11,000) + \frac{1}{2} \ln (X - 9,000).$$

Exponentiating both sides we have

$$X = \sqrt{X + 11,000}\sqrt{X - 9,000}$$

Squaring both sides we have

$$X^2 = X^2 + 2,000X - 99,000,000.$$

Solving for X , we have $X = 49,500$ ducats.

Notice that the answer is independent of the parameters a and b . This is no surprise since risk behavior is independent of affine transformations of the Bernoulli utility function. So the answer is the same as if the Bernoulli utility function were just $\ln X$.

B) Calculate the wealth of the richest fellow who would be willing to sell this ticket for 9,000 ducats. Is the answer the same as the answer to A? If not, can you explain why not?

Suppose that this guy has X safely invested in addition to the ticket. Then if he sells the ticket, his wealth will be $X + 9,000$ with certainty. If he keeps the ticket, his wealth is equally likely to be $X + 20,000$ and X . He will be indifferent about selling the ticket if

$$\ln(X + 9,000) = \frac{1}{2} \ln(X + 20,000) + \frac{1}{2} \ln X$$

Exponentiating this and squaring it, we have

$$(X + 9,000)^2 = X^2 + 20,000X.$$

Simplifying this gives us

$$81,000,000 = 2,000X$$

and therefore $X = 40,500$. Notice that if he sells the ticket, his wealth with certainty will be the same as that of the fellow in Part A.

C) If Bernoulli's poor fellow has wealth of 25 ducats, calculate the lowest price at which he would be willing to sell the lottery ticket.

If he sells at price P , his wealth will be $P + 25$ with certainty. If he does not sell, his wealth will be 20,025 with probability 1/2 and 25 with probability 1/2. He is indifferent about selling if

$$\ln(P + 25) = \frac{1}{2} \ln 20025 + \frac{1}{2} \ln 25.$$

Exponentiate and square both sides to find

$$(P + 25)^2 = 20025 \times 25.$$

So $P + 25 = \sqrt{20025 \times 25} = 225$ and $P = 687.55$.

Since the expected value of the ticket is 10,000 ducats, he is willing to sell for less than 7% of the expected value. Does this suggest how venture capitalists might get rich?

1.1 The Saint Petersburg Paradox

Daniel Bernoulli's cousin, Nickolaus (1687-1759) posed a problem which has come to be known as the St. Petersburg paradox.

“Peter tosses a coin and continues to do so until it should land ”heads” when it comes to the ground. He agrees to give Paul one ducat if he gets ‘heads’ on the very first throw, two ducats if he gets it on the second, four if on the third, eight if on the fourth, and so on, so that with each additional throw the number of ducats he must pay is doubled. Suppose we seek to determine the value of Paul’s expectation.

My aforementioned cousin discussed this problem in a letter to me asking for my opinion. Although the standard calculation shows that the value of Paul’s expectation is infinitely great, it has, he said, to be admitted that any fairly reasonable man would sell his chance, with great pleasure, for twenty ducats.”

–passage from Daniel Bernoulli’s paper

The expected value of winning this lottery is

$$\sum_{t=1}^{\infty} 2^{t-1} \frac{1}{2^t} = \infty.$$

Daniel Bernoulli suggests that the actual value of this lottery to a person with wealth x will be

$$\sum_{t=1}^{\infty} \ln(x + 2^{t-1}) \frac{1}{2^t}.$$

This series can be shown to converge to a finite quantity.

It has been pointed out that Bernoulli's solution would not work for a lottery in which the price for a run of t heads is e^{2^t} . In this case, the expected utility of someone whose only asset is this this lottery would be

$$\sum_{t=1}^{\infty} \ln e^{2^t} \frac{1}{2^t} = \sum_{t=1}^{\infty} \frac{2^t}{2^t} = \infty$$

This difficulty could be conquered by assuming that the Bernoullian utility function is bounded from above. (Note that a function can be always increasing and still bounded.) If $u(x) \leq b$ for all x , then

$$\sum_{t=1}^{\infty} u(x) \frac{1}{2^{t-1}} \leq b \sum_{t=1}^{\infty} \frac{1}{2^{t-1}} = b.$$

2 Constructing Expected Utility Functions

2.1 Contingent Commodity Approach

Let us follow Bernoulli's suggestion of dividing the possible events into a large number of equally likely events and considering preferences over lotteries that assign a specific prize to each possible event. Let the events be e_1, \dots, e_n where the probability of each event is $1/n$. Let $x(e_i)$ be the vector of commodities that one would receive contingent on event i happening. These will be called contingent commodities. Then a vector $x = (x(e_1), \dots, x(e_n))$ specifies the prizes that you would win in each possible outcome. Let us suppose that a decision maker has continuous, transitive preferences over these events and that these preferences are additively separable. Then preferences of this consumer could be represented by a utility function of the form

$$U(x) = \sum_{i=1}^n u_i(x(e_i)).$$

The separability assumption requires that how you rank two alternative lotteries that differ only in what happens in one compound event does not depend on what would happen if the outcome is not in this compound event.

We assume that the events are preference neutral in the sense that one has no intrinsic interest in which event happens, but only in the prize associated with the event. (This is like outcomes on a toss of a coin or roll of a die and not like health outcomes.) Since the events are equally likely, we also assume

that you are indifferent if you swap the outcomes of any two events. Then there u_i 's could all be the same function and utility could be represented by a utility of the form

$$U(x) = \frac{1}{n} \sum_{i=1}^n u(x(e_i)).$$

The function $u(\cdot)$ is known as the von Neumann-Morgenstern utility function or expected utility function.

Events with (rational) unequal probabilities can be broken up into equal size pieces. So the utility of bundle x with probability k/n and bundle y with probability $1 - k/n$ is

$$\frac{k}{n}u(x) + \frac{n-k}{n}u(y) = pu(x) + (1-p)u(y).$$

Since additively separable representations are unique up to affine transformations, the von Neumann-Morgenstern utilities are also uniquely determined up to affine transformations.

2.2 Another Way to Construct Expected Utility

Assume that preferences are continuous, transitive and complete over lotteries. We will show how to construct a utility function over a set of outcomes whose “utilities” range between a best outcome and a worst one. These might not be the best and worst imaginable, but just the top and bottom of the range of desirability of things that you are able and willing to think about.

Assign utilities to the best outcome and the worst outcome so that $U(best) = 100$ and $U(worst) = 0$. Let U of any lottery be the probability mix of best and worst outcomes that is indifferent to that lottery.

Define a simple gamble to consist of a finite number n of possible prize bundles, $a_1 \dots, a_n$ such that with probability p_i you get prize bundle i .

Axioms: Assume that preferences on gambles are

- 1) Complete
- 2) Transitive $a_1 \succ a_2, \dots, \succ a_n$
- 3) Continuous. For any gamble g , there is some p such that $g \sim pa_n + (1-p)a_1$.
- 4) Monotonic: If $a > b$, then getting best bundle with probability a and worst with probability $(1-a)$ is preferred to best with probability b and worst with probability $(1-b)$.

5) Substitution If for each outcome, you are indifferent between the prize you get in lottery g and the prize you get in lottery h , then you are indifferent between g and h .

6) Reduction to Simple gambles: You care only about the probabilities and the prizes and not on the way that gambles might be compounded.

With these axioms, preferences on gambles are representable by a von Neumann-Morgenstern utility.

The utility of outcome x is determined as follows: If getting x for sure is indifferent to a gamble in which you get *best* with probability p and *worst* with probability $1 - p$, then

$$u(x) = pU(\textit{best}) + (1 - p)U(\textit{worst}).$$

(Lecturer's aside: Constructing your own utility under this procedure would be psychologically taxing and probably unreliable if *best* were truly fantastic and *worst* almost unimaginably terrible.)

Let us show that utility constructed in this way is unique up to affine transformations: Suppose that u and v are two alternative von Neumann-Morgenstern utility functions that represent the same preferences over lotteries. For any x such that $\textit{worst} \succ x \succ \textit{best}$, there is some p such that having x for sure is indifferent to having *best* with probability p and *worst* with probability $1 - p$. Then since u and v both represent the same preferences, it must be that

$$u(x) = pu(\textit{best}) + (1 - p)u(\textit{worst})$$

and

$$v(x) = pv(\textit{best}) + (1 - p)v(\textit{worst})$$

with the same p in either case. Rearranging terms, we find that

$$\frac{u(x) - u(\textit{worst})}{u(\textit{best}) - u(\textit{worst})} = \frac{v(x) - v(\textit{worst})}{v(\textit{best}) - v(\textit{worst})} = p.$$

But this implies that

$$u(x) - u(\textit{worst}) = \frac{u(\textit{best}) - u(\textit{worst})}{v(\textit{best}) - v(\textit{worst})} (v(x) - v(\textit{worst}))$$

and hence

$$u(x) = Av(x) + B$$

where

$$A = \frac{u(best) - u(worst)}{v(best) - v(worst)}$$

and

$$B = u(worst) - Av(worst).$$

We see that A and B are constants and so u is an affine transformation of v .

2.2.1 Cardinal measurement of temperature: an analogy

We have shown how to define von Neumann-Morgenstern utility by equivalence to a probability mix of a “best” and a “worst” prospect. It is instructive to think of an analogous procedure for measuring water temperature, where we use physical mixtures rather than probability mixtures. Let us consider temperatures between the freezing point and the boiling point of water. Let us assign a “Celsius temperature” of zero to the freezing point and 100 to the boiling point. Suppose that we can always tell when two pots of water are at the same temperature. Take any pot x of water and find the proportions of $p(x)$ of boiling water and $1 - p(x)$ of just-about-to-freeze water that you would need to get a pot of water of the same temperature as your original. Then define the Celsius temperature of the water in pot x to be

$$C(x) = p(x)100 + (1 - p(x))0.$$

Alternatively, you could define a Fahrenheit temperature where boiling water is 212 degrees and freezing water is 32 degrees. Then the Fahrenheit temperature of the water in your pot is

$$F(x) = p(x)212 + (1 - p(x))32.$$

For any pot of water, the proportions of boiling and freezing water needed to match the temperature of your original pot is the same regardless of the temperatures that you assign to boiling and freezing water. How are the Celsius and Fahrenheit temperatures related? If we rearrange Equation 2.2.1, we find that

$$p(x) = \frac{C(x)}{100}$$

and if we rearrange Equation 2.2.1, we see that that

$$p(x) = \frac{F(x) - 32}{180}.$$

From Equations 2.2.1 and fp it follows that

$$\frac{C(x)}{100} = \frac{F(x) - 32}{180}$$

or equivalently,

$$C(x) = \frac{5}{9}(F(x) - 32) = \frac{5}{9}F(x) - \frac{160}{9}.$$

Thus temperatures in Fahrenheit and Centigrade are affine transformations of each other.

The analogy to our second method of constructing expected utility is close. Let the temperature of water at the boiling point correspond to $u(\textit{best})$ and the temperature of water at the freezing point correspond to $u(\textit{worst})$. We measure the temperature of a specified object as 100 times the fraction p of boiling water in a mix of boiling and freezing water that is exactly as warm as the object. We measure a person's Bernoulli utility of an outcome x as the probability p such that this person is indifferent between having the outcome x for sure and having a probability of p of getting the best outcome and $1 - p$ as the worst outcome.

You will notice that the method we have discussed for measuring temperature only works for temperatures between the freezing point and the boiling point of water. To cardinalize more extreme temperatures, we would need to consider the behavior of other chemical elements with higher and lower boiling and freezing point. Alternatively, we could work with the expansion of columns of mercury or other substances.

Similar problems arise with the construction of utility functions based on a best and worst outcome. Really horrible outcomes or really wonderful outcomes may be so far from ordinary experience that one is a poor judge of comparing lotteries between such extremes. Perhaps the most dramatic comparison of this type is the so-called "Pascal's wager" with which religious priesthods attempt to gain power over ordinary people. "He who believes what I say and does what I say will enjoy unimaginable bliss for all eternity. He who does not will experience an infinity of torture." Let there be a utility cost C of following the priest's instructions. Suppose you think that the chance that the priest is telling the truth is p . Then your expected utility if you believe the priest and do what he says is $p \times \infty + (1 - p)(X - C) = \infty$. If on the other hand you choose not to do what the priest says, is

$p \times (-\infty) + (1 - p)X = -\infty$. If $p > 0$, then no matter how small it is, you will prefer to believe the priest and follow his instructions.³

Other ordinal measures of temperature are possible. For example, we could measure the temperature of everything that is warmer than freezing water by the square of its Celsius temperature and everything colder than freezing water by minus the square of its Celsius temperature. This measure would always assign higher numbers to warmer things. You could still use this scale with a mercury thermometer. (Of course the temperature numbers painted on the side of your thermometer would not be equally spaced.) But this measure would be inconvenient if you wanted to calculate the temperature of mixtures.

Exercise: (Mainly for entertainment...not to be handed in.)

A.) According to Wikipedia “A now somewhat obsolete scale ... was created by R A F de Réaumur (1683-1757) a French scientist. He used the freezing point of water as his zero mark, and put the boiling point at 80 degrees. This scale was widely used in the 18th and 19th centuries, especially in France, in scientific communities.” How is the temperature in the Réaumur scale related to the temperature in degrees Fahrenheit?

B) Scientists have determined that the coldest possible temperature is -273.15 degrees Celsius. Temperature in degrees Kelvin assigns a temperature of 0 to this coldest temperature and assigns a difference of 100 degrees between the freezing point of water and the boiling point of water (at sea level). How is temperature in degrees Kelvin converted to degrees Fahrenheit?

C) Another temperature scale, named after the Scottish engineer W.J.M. Rankine, is the Rankine scale. It assigns a number of zero to the coldest possible temperature, but a Rankine degree is equal to one degree Fahrenheit. What are the temperatures of the freezing point and boiling point of water in the Rankine scale?

D) Suppose that you measured temperatures between the freezing point and the boiling point of water by the formula

$$S(x) = \frac{1}{100}C(x)^2$$

where $C(x)$ is the temperature in degrees Celsius. If pot of water x is the same temperature as a mixture with proportions $p(x)$ of water that is boiling

³This argument was proposed by philosopher and mathematician, Blaise Pascal, in *Pensees*, 1670

and $(1 - p(x))$ of water that is just about to freeze? How would you calculate the temperature of this mixture?

3 Risk Aversion and Preferences over Money Gambles

We could conduct this discussion by assuming that there is only one commodity. But we can do better, and to do so is instructive.

If there is more than one commodity, the natural way to handle preferences over money gambles is to employ the notion of indirect utility. Suppose that there are n ordinary commodities and that a von Neumann-Morgenstern utility function is defined over lotteries in which the prizes are n -commodity bundles. Let us consider a consumer with von Neumann-Morgenstern utility function $u(x)$. Let us consider lotteries in which the consumer's income is randomly determined, but the price vector p is the same for all outcomes of the lottery. Let

$$v(p, y) = \max\{u(x) | px \leq y\}$$

be this consumer's indirect utility function. Since u is determined up to an affine transformation, so is v . (You should be able to show this.)

If a consumer plans to maximize her utility subject to whatever her budget turns out to be, then her preferences over money lotteries will be represented by expected indirect utility. Not only does this setup allow us to analyze preferences over gambles with random income where prices are held constant, but it also enables us to describe preferences over gambles in which prices as well as incomes are random.

Example 1: There are two commodities, 1 and 2, and the von Neumann-Morgenstern utility function is

$$u(x_1, x_2) = (x_1 + 2x_2^{1/2})^{1/2}.$$

Suppose that $p = (p_1, p_2) = (1, p_2)$. To find indirect von Neumann-Morgenstern utility, $v(p, y)$, we find the Marshallian demand function for someone with income y and prices $(1, p_2)$. Setting marginal rate of substitution equal to the price ratio, we find that

$$x_2(1, p_2, y) = p_2^{-2}.$$

At an interior solution, it must be that

$$x_1(1, p_2, y) = y - p_2 x_2(1, p_2, y) = y - \frac{1}{p_2}.$$

Thus there is an interior solution if and only if $y > \frac{1}{p_2}$. If $y \leq \frac{1}{p_2}$, then $x_1(1, p_2, y) = 0$ and $x_2(1, p_2, y) = \frac{y}{p_2}$. Then for $y > \frac{1}{p_2}$, we have

$$v(1, p_2, y) = y - \frac{1}{p_2} + 2\frac{1}{p_2} = y + \frac{1}{p_2}.$$

Then indirect von Neumann-Morgenstern utility is

$$v(1, p_2, y)^{1/2} = \left(y + \frac{1}{p_2} \right)^{1/2}.$$

Example 2: There are n commodities and the von Neumann-Morgenstern utility function is homogeneous of degree k . Then $x(p, y) = yx(p, 1)$ and so

$$v(p, y) = u(yx(p, 1)) = u(x(p, 1))y^k.$$

Thus when p is held constant, $v(p, y)$ is just a constant times y^k .

Example 3: A special case of Example 2.

$$u(x_1, x_2) = (x_1^a + x_2^a)^b.$$

The indirect utility function for a CES utility function with parameters a and b is

$$v(p_1, p_2, y) = (p_1^r + p_2^r)^{-b/r} y^{ab}$$

where $r = a/(a - 1)$. See Jehle and Reny Example 1.2, pp 31-32 for the case where $b = 1/a$.

A consumer is said to be risk averse with respect to a money gamble if she prefers to receive the expected value of the gamble with certainty rather than to have the gamble. If she prefers the gamble, she is said to be risk loving. A consumer who is risk averse with respect to all money gambles is said to be (just plain) risk averse.

If the expected utility function $v(p, y)$ is concave in y , then a consumer is risk averse. If it is convex in y , she is risk loving. A consumer for whom $v(p, y)$ is neither concave nor convex will be risk averse for some gambles and risk loving for other gambles. Draw some pictures to illustrate this.

3.1 Measures of Risk Aversion

Degree of absolute risk aversion (at y) is defined to be

$$R_a(y) = \frac{-u''(y)}{u'(y)}$$

Things to Notice:

- 1) This measure does not change with affine transformations of u .
- 2) Suppose that absolute risk aversion is a constant, $a > 0$, then it must be that we can write the von Neumann Morgenstern utility function as

$$u(y) = 1 - e^{-ay}.$$

This function is called the CARA function.

Proof of this claim: To find the utility function, we need to solve the differential equation

$$\frac{-u''(y)}{u'(y)} = a.$$

This is equivalent to the equation

$$\frac{d \ln u'(y)}{dy} = -a$$

which implies that

$$\ln(u'(y)) = -ay + b$$

for some b . Exponentiating both sides gives us

$$u'(y) = e^b e^{-ay}.$$

Integrating both sides of this equation, we have

$$u(y) = c - \frac{1}{a} e^b e^{-ay}$$

for some constants b and c . Since von Neumann Morgenstern utility functions are unique up to increasing affine transformations, we could choose the constant c to be any constant and b to be any positive constant. Economists have found it convenient to choose $c = 1$ and $e^b/a = 1$. With these choices, we can write the utility as $1 - e^{-ay}$.

- 3) Risk aversion is positive for risk averters, negative for risk lovers.

4) If $v(y) = h(u(y))$ where h is increasing and concave, then v represents more risk averse preferences than u .

Proof: $v'(y) = h'(u(y))u'(y) > 0$ and $v''(y) = h''(u(y))u'(y)^2 + h'(u(y))u''(y)$.

So

$$\frac{v''(y)}{v'(y)} = \frac{h''(y)}{h'(y)}u'(y) + \frac{u''(y)}{u'(y)} < \frac{u''(y)}{u'(y)}.$$

Conversely, if v is more risk averse than u , then $v(y) = h(u(y))$ where h is a concave function.

5) We say there is decreasing absolute risk aversion DARA if for all y , the degree of absolute risk aversion is a decreasing function of y .

6) A consumer with decreasing absolute risk aversion who is faced with a choice of portfolios containing one risky asset and one safe asset, will choose to purchase a greater amount of the risky asset as he gets wealthier. (In other words, the risky asset is a “normal good” for a consumer with DARA.) See proof in Jehle and Reny.

Risky asset gives payoff $1 + r_i$ in event i per dollar invested in it. Initial wealth is w units of safe asset. If he spends β dollars on the risky asset, his wealth in event i will be $w - \beta + \beta(1 + r_i) = w + \beta r_i$. His expected utility will be $\sum_i (u(w + \beta r_i))$. At an optimum, it must be that

$$\frac{d}{d\beta} \sum_i (u(w + \beta r_i)) = \sum_i u'_i(w + \beta r_i) r_i = 0.$$

How does wealth affect optimal β ? Differentiate both sides with respect to w . We have

$$\sum_i u''_i(w + \beta r_i) r_i + \sum_i u''_i(w + \beta r_i) r_i^2 \frac{d\beta}{dw} = 0.$$

Rearrange terms to find

$$\frac{d\beta}{dw} = - \frac{\sum_i u''_i(w + \beta r_i) r_i}{\sum_i u''_i(w + \beta r_i) r_i^2}.$$

The denominator of this expression must be negative. If there is DARA, the numerator is also negative. See Jehle and Reny for the details.

3.1.1 Relative risk aversion

Degree of relative risk aversion at y is defined as

$$R_r(y) = \frac{-yu''(y)}{u'(y)}.$$

Things to Notice:

1) If there is constant relative risk aversion, then there must be decreasing absolute risk aversion. $R_r(y) = yR_a(y)$ so $R_a(y) = R_r(y)/y$. If $R_r(y) > 0$ is constant, then $R_a(y)$ is clearly decreasing in y .

2) There is constant relative risk aversion CRRA if and only if $u(y)$ is an affine transformation of $u(y) = \frac{1}{a}y^a$ for some constant $a \neq 0$ or an affine $u(y) = \ln y$.

3) Suppose that preferences over lotteries are homothetic and representable by an expected utility function over income. Then by Bergson's theorem, they belong to the constant relative risk aversion family.

4) If there are two assets, one risky, one safe, then if there is DRRA, the wealthier you get, the larger fraction of your wealth you will keep in the risky asset. If there is CRRA, that fraction will be constant.

3.1.2 Certainty equivalent utility function

Another way to represent preferences on money lotteries is by their "certainty equivalents". How much money would I need to have with certainty to be exactly as well off as if I had this lottery.

The risk premium for a gamble is defined as the difference between the expected return of the gamble and the certainty equivalent of the gamble.

Draw a graph with two possible events, contingent commodities. Points in the graph are indifferent to some sure thing found on the 45 degree line. Certainty equivalent is the projection of this line onto either axis.

For example, suppose that $u(y) = \ln(y)$ for all $y > 0$. What is the certainty equivalent of a gamble in which the toss of a fair coin determines whether you will have $y = y_0$ or $y = W > y_0$? To answer this question, solve the equation $u(C) = 1/2u(y_0) + 1/2u(W)$, which is equivalent to

$$\ln C = 1/2 \ln y_0 + 1/2 \ln W$$

which is equivalent to

$$C = \sqrt{W} \sqrt{y_0}.$$

The risk premium of this gamble is then

$$P = W/2 + y_0/2 - \sqrt{W} \sqrt{y_0}.$$

For example if $W = 100$ and $y_0 = 1$, the risk premium is $55.5 - 10 = 45.5$

Suppose that like Bernoulli's poor fellow, you found a lottery ticket that would pay either 0 or W with probability $1/2$ for each outcome and that your initial wealth was y_0 , so that the gamble you face is that your wealth will be y_0 with probability $1/2$ and $y_0 + W$ with probability $1/2$. What is the certainty equivalent of your position? Solve

$$\ln C = 1/2 \ln y_0 + 1/2 \ln (W + y_0).$$

Then

$$C = \sqrt{y_0} \sqrt{W + y_0}.$$

The risk premium is

$$P = W/2 - \sqrt{y_0} \sqrt{W + y_0}.$$

Another example: Suppose that $u(y) = \sqrt{y}$. What is the certainty equivalent of a gamble in which the toss of a fair coin determines whether you will $y = 0$ or $y = W$? Solve the equation $\sqrt{C} = 1/2\sqrt{0} + 1/2\sqrt{W}$ which gives you $C = W/4$. The risk premium of this gamble is

$$P = W/2 - W/4 = W/4.$$

Example: Suppose that $u(y) = \sqrt{y}$. What is the certainty equivalent of a gamble in which a fair coin is tossed repeatedly until tails appear. If heads show up the first t times it is tossed and then comes up tails, your wealth will be 2^t .

Solve the equation:

$$\sqrt{CE} = \sum_{t=1}^{\infty} \frac{1}{2^t} 2^{t/2} = \sum_{t=1}^{\infty} \left(2^{-1/2}\right)^t.$$

Recall that if $-1 < a < 1$, the sum of a geometric series $\sum_{t=1}^{\infty} a^t$ is $\frac{a}{1-a}$. The above series is a geometric series in which $a = \frac{1}{\sqrt{2}}$. Therefore our equation becomes

$$\sqrt{CE} = \frac{2^{-1/2}}{1 - 2^{-1/2}} = \frac{1}{\sqrt{2} - 1}$$

Squaring both sides, we find $CE = \frac{1}{(\sqrt{2}-1)^2} \sim 5.83$.

4 Subjective Probability

So far we have treated probabilities as if they are objectively known. But in many economic applications, there is not a clear “objective” probability that would be recognized and agreed to by all reasonable observers. Instead, individuals have their own, “subjective” notions of the likelihood of events. We can, however, construct a coherent theory of “subjective probability” that seems to correspond to the way that a systematic rational decision-maker would organize his thoughts about uncertain events.

One theory of subjective probabilities works as follows. Suppose that an individual is an expected utility maximizer over events with known probability p . Consider two outcomes g and b , such that having g for sure is preferred to having b for sure. Now consider an event E and define a lottery $L(E)$ that gives you prize g if E occurs and prize b if E does not occur. Let us define the decision-maker’s *subjective probability* of the event E to be the probability p such that the decision-maker is indifferent between the lottery $L(E)$ and a lottery in which he gets prize g with probability p and prize b with probability $1 - p$.⁴ This procedure works well, in cases where the decision-maker does not care about the “events” themselves, but only with the “prizes” associated with the events. This seems appropriate for evaluating outcomes of the spin of a roulette wheel, but without some careful interpretation is not appropriate if the events themselves affect preferences. This is known as *state dependent utility*. Examples of events that may directly affect preferences are: “It will rain today” or “I will have a serious illness today” or more drastically, “I will be eaten by a bear or run over by an SUV. ”

An example illustrates what can go wrong with the previous method of finding subjective probability when utility is state dependent. Suppose that you are interested in the probability that it will rain on the day of an outdoor sporting event which you are interested in attending. A ticket to this event costs \$36. Bundle $g = (1, 64)$ consists of 1 ticket to the game and \$64 in spending money. Bundle $b = (0, 100)$ has no ticket to the game and \$100 in spending money. If it rains on the day of the game, you will throw away the ticket and not go to the game. Although you realize that there is a chance that you may have to throw away the ticket because of rain, you still prefer bundle g to bundle b .

Lets try to apply the method previously discussed to determine your

⁴If the decision-maker has continuous preferences, there will be exactly one p .

subjective probability $p(E)$ of the event E that it will rain on the day of the game. Construct the lottery $L(E)$ that gives you the bundle g if it rains on the day of the game and b if it does not rain on the game. You like bundle g better than bundle b . Now ask “For what value of p am I indifferent between the lottery $L(E)$ and a simple lottery in which I get bundle g with probability p and b with probability $(1 - p)$. ” Why is this not a reasonable way to find out how likely you think it is to rain? What do we mean by how likely you think it is to rain?

The lottery $L(E)$ is one in which if it rains, you will pay \$36 and get a (useless) ticket to the game and if it does not rain you will have \$100 and no ticket to the game. If you believed that the probability of rain is $p > 0$, you would always prefer any simple lottery between g and b to the lottery $L(E)$.

A more satisfactory way to deal with subjective probability for state dependent events, is to make sure that the description of the “prizes” over which you have preferences includes a full description of the preference-relevant part of the state. So for example, suppose we describe the prize not only by how many tickets and how much money you have, but also by whether it is raining or not. We would now describe the prize by the three variables (R, T, X) where R is 1 or 0 depending on whether it rains on the day of the game, T is 0 or 1 depending on whether you have a ticket, and X is money available for other stuff. Now take two fully specified outcomes G and B such that G is preferred to B . For example, let $G = (1, 0, 100)$ and $B = (0, 0, 90)$. We know that you prefer G to B . Let $L(E)$ be the lottery that gives you prize G if it rains and prize B if it does not rain. We then look for the probability p and construct the simple lottery that gives you prize G with probability p and prize B if it does not rain. Let us find the probability p such that you find that this simple lottery is indifferent to $L(E)$ and define this p to be your subjective of p . Now, unlike the earlier example, we see that the lottery $L(E)$ and the simple lottery lead to the same two outcomes. Either it rains and you stay home and have \$100 or it does not rain and you stay home and have \$90.

4.1 Proper scoring rules

A mechanism that gives agents an incentive to truthfully report their subjective probability of an event is known as a *proper scoring rule*. How do you get the weatherman to truthfully report his subjective probability that it will rain tomorrow? There are many possible payoff functions that will give

a risk-neutral weatherman an incentive to report these probabilities truthfully. Two commonly discussed examples are the Brier Score and the logarithmic score.

The Brier score works as follows. Suppose that the weatherman has subjective probability π that it will rain tomorrow. If he announces that his forecasted probability of rain is f , then if it rains, his Brier score will be $1 - (1 - f)^2$. If it doesn't rain, his Brier score will be $1 - f^2$. The weatherman's expected payoff if he announces probability f is then

$$1 - (\pi((1 - f)^2) + (1 - \pi)f^2).$$

The logarithmic scoring rule will pay the weatherman $C - \ln f$ if it rains and $C - \ln(1 - f)$ if it doesn't rain. In this case, the weatherman's expected payoff is

$$C - (\pi \ln f + (1 - \pi) \ln(1 - f)).$$

In either case, the weatherman chooses f to maximize his payoff. Simple calculus shows that in both cases he maximizes his payoff by selecting $f = \pi$.

5 Value of Life

How much would I have to pay you to play a game of Russian roulette? If your answer is "No amount of money would make me do that", must I assume that you have "infinite disutility" for dying?

Do you ever cross busy streets or drive a compact car in fast freeway traffic? If you do so, you are taking an avoidable risk of being killed. How can a willingness to take small risks for small gains be consistent with unwillingness to play Russian roulette for any price?

There is a simple, and I believe fairly satisfactory answer within the standard expected utility model. Let us start with a single period version. Suppose that there is just one consumption good and your preferences can be represented by a utility function $U(\pi, c)$ where c is your consumption contingent on survival and with probability $1 - \pi$ you will die at the beginning of the period and with probability π you will survive. Suppose that the function $u(\cdot)$ is continuously differentiable and strictly increasing, but that there is a bound b such that $u(c) \leq b$ for all c .

Suppose that in the initial situation you have survival probability $\bar{\pi}$ and consumption \bar{c} . In order to study tradeoffs that you would be willing to take,

we would like to consider the indifference curve specified by $\{(\pi, c) | \pi u(c) = \bar{\pi} u(\bar{c})\}$. Since $u(c) \leq b$ for all c , it must be that along this indifference curve,

$$\pi \geq \frac{\bar{\pi} u(\bar{c})}{b}.$$

So, for example if

$$\frac{\bar{\pi} u(\bar{c})}{b} > 5/6,$$

you would never voluntarily play Russian roulette. (Draw a picture of the indifference curve.)

On the other hand, since u is differentiable, there is a well-defined and finite marginal rate of substitution between survival probability and wealth. Someone with this type of utility function is willing to undertake small risks for a finite price.

This is the foundation of the “Value of expected lives” method in benefit cost analysis. Many public safety projects, like road improvements, are likely to save lives, but at an economic cost. Public authorities must decide whether they are “worth it.” How can we avoid falling into the trap that “Since human lives are priceless, every such project is worthwhile.” Most such projects result in a tiny improvement in survival probability for each of a large number of people. Each of the people who is made a little safer has some finite willingness to pay for this additional safety. Let V_i be the marginal rate of substitution of person i between survival probability and money. Let $\Delta\pi_i$ be the gain in i ’s survival probability from the project. Then i ’s willingness to pay for the increased survival project resulting from the project is $V_i \Delta\pi_i$. Suppose for example that V_i is the same for all i . Then total willingness to pay for the project is $V \sum_i \Delta\pi_i$, which is V times the expected number of lives saved.

6 Applications to Insurance

6.1 Insurance against financial loss

Suppose that a consumer is a risk averse expected utility maximizer with von Neumann-Morgenstern utility function $\pi_i u(x_i)$ where x_i is the number of dollars he has to spend in event i . If he does not buy insurance, his income will be y_i in event i , where y_i may vary from one event to another. Suppose that he is able to buy actuarially fair insurance. What does he do?

Actuarially fair means that the cost of \$1 contingent on event i is π_i . So he seeks to maximize

$$\pi_i u(x_i)$$

subject to

$$\sum_{i=1}^n \pi_i (x_i - y_i) = 0.$$

Writing the Lagrangean for this problem and taking derivatives, we see that at a constrained maximum it must be that for some λ , $u'(x_i) = \lambda$ for all i . Since the consumer is assumed to be risk averse, we know that $u''(x) < 0$ for all x , so the only solution must be one in which consumption is the same amount, \bar{x} in every event. The budget constraint tells us that

$$\sum_{i=1}^n \bar{x} = \bar{x} = \sum_{i=1}^n \pi_i y_i.$$

His insurance policy is one in which he receives $\bar{x} - y_i$ when $y_i < \bar{x}$ and he pays $y_i - \bar{x}$ when $y_i > \bar{x}$.

Exercise: Work out the special case when y_i takes on only two values \bar{y} and $\bar{y} - L$.

6.2 Medical Insurance

6.2.1 Example 1

Let us assume that the probabilities of getting various diseases are known to everyone from medical experience and that everyone concerned agrees about these probabilities.

Suppose that someone is sure to get exactly one of n different diseases. His attitude toward risks, conditional on being healthy is described by a strictly concave von Neumann-Morgenstern utility function $u(x)$. He will get disease i with probability π_i . If he gets disease i , he can go to the hospital for repairs. The hospital will fix him up as good as new at a cost of c_i for disease i .

Suppose that he has fixed wealth, $W > \sum \pi_i c_i$ and that he can buy any amount of insurance at actuarially prices, for each of the various diseases. What will he buy?

Suppose that some diseases are so expensive, that he can not afford to buy insurance that will pay for curing them. How do we handle that?

6.2.2 Example 2

Consider someone whose income is independent of his health. Suppose that he is faced with the possibility of several illnesses, none of which is curable. His expected utility function is $\pi_i u_i(x_i)$ where π_i is the probability that he gets disease i and x_i is the amount of money available to him contingent on his getting disease i where

$$u_i(x_i) = u(x_i) - a_i.$$

Show that this person will buy no health insurance.

6.2.3 Example 3

Consider the same situation, but this time suppose that his expected utility function is $\pi_i u_i(x_i)$ where

$$u_i(x_i) = \frac{1}{\alpha} (k_i x_i)^\alpha$$

where $\alpha < 1$. In this case, we say that a disease is more severe, the lower is k_i . When will this person buy insurance that pays him if he has severe illness and where he pays the insurance company if he is well? When will he buy insurance such that the insurance pays him if he is well and he pays the insurance company if he has severe illness.