

Expected Utility Analysis

TeX file: *expectedutility* — Daniel A. Graham, September 8, 2011

1 Preferences for Probabilities

We now turn to characterizing preferences for lotteries with the goal of identifying a utility function to represent these preferences. A difficulty is that even simple lotteries have too many components to illustrate in two or even three dimensional graphs. Consider, for example, the *Allais Lotteries*¹ described below

Choice A

Lottery 1: (\$1 million; 1)

Lottery 2: (\$0, \$1 million, \$5 million; 0.01, 0.89, 0.10)


Choice B

Lottery 3: (\$0, \$1 million; 0.89, 0.11)

Lottery 4: (\$0, \$5 million; 0.90, 0.10)

These lotteries involve as many as 3 different prizes with corresponding probabilities. Even recognizing that the probabilities must add up to 1 and thus that there will only be two independent choices of probabilities, means that we still need 5 axes — 1 for each of the three prize amounts and 2 more for the probabilities — to illustrate the set of such lotteries.

The solution, as usual, is to illustrate cross sections. In the Allais Lotteries, for example, there are actually only 3 distinct prize amounts: \$0, \$1 million and \$5 million. By restricting attention to lotteries that involve just these prizes, we need only to deal with two-dimensions to graph the probabilities. Let p_1 denote the probability of winning \$1 million and p_2 denote the probability of winning \$5 million and the probability of winning \$0 must then be $1 - p_1 - p_2$. This is illustrated below:



¹These lotteries were used in a famous experiment by the French Economist Allais who observed that the modal (most common) choices were to choose Lottery 1 in Choice A and Lottery 4 in Choice B.

It is important to realize that *every* lottery involving only the prizes \$0, \$1 million and \$5 million corresponds to some point in this yellow triangle and vice versa. The corners correspond to the degenerate lotteries in which only one prize is possible, the edges correspond to lotteries in which only two prizes are possible (have positive probabilities) and the interior of the triangle corresponds to lotteries in which all three prizes are possible.

Compound lotteries of lotteries involving only these prizes also correspond to points in the triangle. Consider, for example, the *simple lotteries*


Lottery A: (\$0, \$1 million; 1/3, 2/3)

Lottery B: (\$0, \$5 million; 2/3, 1/3)

and consider a *compound lottery* which pays Lottery A with probability 1/2 and Lottery B with probability 1/2, i.e., a fair coin will be tossed and if it comes up heads then the person will win Lottery A and if it comes up tails then the person will win Lottery B. Note that this compound lottery is equivalent to the following simple lottery:

Lottery C: (\$0, \$1 million, \$5 million; 1/2, 1/3, 1/6)

since the ultimate probabilities of winning \$0, \$1 million and \$5 million, respectively, are $1/2 \times 1/3 + 1/2 \times 2/3 = 1/2$, $1/2 \times 2/3 = 1/3$ and $1/2 \times 1/3 = 1/6$. These lotteries are illustrated below.



Note that Lottery C, which corresponds to a 1/2 chance of winning Lottery A and a 1/2 chance of winning Lottery B, lies exactly half way along the line segment connecting B and A. This is, of course, no accident — the coordinates of the compound lottery are just the weighted average of the coordinates of the simple lotteries with the probability of receiving the simple lotteries as the weights. We will thus use the shorthand notation:

$$pA + (1 - p)B$$

to denote a compound lottery that pays Lottery A with probability p and Lottery B with probability $1 - p$. Geometrically, this compound lottery lies a fraction p of the way along the line segment from B to A.

How will indifference curves appear in the triangle diagram? To address this issue we introduce the following short-hand notation:

Strictly preferred: \succ

At least as desirable: \succsim

Indifferent to: \sim

Using this notation we would, for example, express the fact that the lottery A is at least as desirable as the lottery B by simply writing $A \succsim B$.

Next we make the same assumptions about preferences that we made under consumption theory — that they are complete, reflexive and transitive — and add two *new* assumptions or axioms:

1.1 The Continuity Axiom

The *Continuity Axiom* states that if A , B and C are *any* three lotteries with $A \succ B$ and $B \succ C$, then there is a $p \in (0, 1)$ such that $pA + (1 - p)C \sim B$. In other words, there is a compound lottery involving A (the best) and C (the worst) that the individual will regard as indifferent to B (the intermediate). Geometrically, the indifference curve through B (the intermediate) must cross the line segment connecting A (the best) and C (the worst).


Returning to the Allais Lotteries, the Continuity Axiom asserts that somewhere on the vertical axis — the line connecting the certainty of winning \$5 million with the certainty of winning \$0 — there is a lottery that the person would regard as indifferent to the certainty of winning \$1 million.

Where this point lies is, of course, a matter of personal taste for the individual. We might identify the point by offering the choice between the following two lotteries

Lottery M: (\$1 million; 1)

Lottery I: (\$0, \$5 million; $1 - p$, p)

where we begin with a small value for p and increase it until the person is just indifferent. Having done this we have the two indifferent lotteries, I and M , illustrated below:



The rest of the indifference curve connecting M and I is the subject of the next axiom.

1.2 The Independence Axiom


The *Independence Axiom* states that if A , B and C are *any* three lotteries with $p \in [0, 1]$, then

$$pA + (1 - p)C \succsim pB + (1 - p)C$$

if and only if


$$A \succsim B$$

This is illustrated below. The two spinners are identical save for the labels, i.e., spinning either wheel entails the same probability, p , of the wheel stopping with the pointer in the yellow region. The only difference between the lotteries is that Lottery 1 pays A and Lottery 2 pays B if the wheel stops with the pointer in the yellow. Both lotteries pay C otherwise.



The *Independence Axiom* asserts that the individual will rank the compound lotteries in the same order as the simple lotteries A and B , i.e., Lottery 1 will be preferred to Lottery 2 if and only if A is preferred to B . Thus the ranking of the compound lotteries depends only upon the ranking of A and B and is *independent* of lottery C which is common to both..

In terms of the Allais Lotteries, we know that since the individual is indifferent between M and I , the Independence Axiom implies that the individual must also be indifferent between $N = pM + (1 - p)I$ and $I = pI + (1 - p)I$. [In terms of the axiom, we are letting $A = M$, $B = I$ and $C = I$.] But N lies along the line segment connecting I with M a fraction p of the way toward M . Since this must be true for any choice of $p \in [0, 1]$, it follows that the indifference curve connecting I and M must be a *straight line*!




The Independence Axiom not only implies that indifference curves are straight lines but also that they are parallel. To see this, consider the following two compound lotteries:

$$M' = pM + (1 - p)C$$

$$I' = pI + (1 - p)C$$

illustrated below:



The Independence Axiom implies that since the individual is indifferent between I and M , the individual must also be indifferent between I' and M' . A second application of the independence axiom then implies that the indifference curve connecting I' and M' must also be a straight line. But the fact that the same p is used in both of the compound lotteries means, by a similar triangles argument, that the indifference line $I'M'$ must be parallel to the indifference line IM . Since this argument works for any choice of $p \in [0, 1]$, every parallel to IM must be an indifference line.

A similar argument replacing C with D establishes that every parallel below IM must be an indifference line as well.

1.3 Constructing the Utility Function

Since the Independence Axiom implies that indifference curves are linear and parallel, one might suspect that a utility function could be found which is linear in the probabilities. This is indeed correct. Since the prizes are fixed, we need to find constants corresponding to the three fixed prizes. Let's call these constants the "utilities" of the respective prizes and use the labels u_0 , u_1 and u_2 . The "utility" of a lottery p_1, p_2 will then be

$$\begin{aligned} U(p_1, p_2) &\equiv (1 - p_1 - p_2)u_0 + p_1u_1 + p_2u_2 \\ &= u_0 + p_1(u_1 - u_0) + p_2(u_2 - u_0) \end{aligned}$$

The task is to choose u_0 , u_1 and u_2 so that the indifference curves for this utility function should have the proper slope. This is an easy task. Setting utility equal to a constant and solving for p_2 (the variable plotted on the vertical axis) gives the following equation for an indifference curve:

$$p_2 = -\frac{u_1 - u_0}{u_2 - u_0}p_1 + \text{constant} \quad (1)$$

Now if we know the actual slope of the person's indifference curves is equal to s , then it is a simple matter to choose the utilities of the prizes so that

$$-\frac{u_1 - u_0}{u_2 - u_0} = s$$

Not only can this always be done, it can be done more than one way. If u_0 , u_1 and u_2 work then so will $au_0 + b$, $au_1 + b$ and $au_2 + b$ for *any* positive a and *any* b whatever. This follows from two facts. In the first place, the indifference curves have the same slope:

$$\begin{aligned} -\frac{au_1 + b - au_0 - b}{au_2 + b - au_0 - b} &= -\frac{a(u_1 - u_0)}{a(u_2 - u_0)} \\ &= -\frac{u_1 - u_0}{u_2 - u_0} \end{aligned}$$

And in the second place, the fact that a is positive means that higher indifference curves lie in the same direction.

Multiplying by a positive constant and adding another constant is called a *positive affine transformation*. The utilities of prizes that we have identified are unique up to such a transformation. Note in particular that other positive monotonic transformations will *not* work, e.g., squaring the utilities of the prizes will give a different slope.

The fact that the utilities of the prizes are unique only to a positive affine transformation means that we have two “degrees of freedom” in choosing the utilities, e.g., we could arbitrarily choose $u_0 = 0$ and $u_2 = 1$ and then solve for u_1 .² Suppose, for example, that when we ask our individual to choose between a certainty of \$1 million and the lottery (\$0, \$5 million; $1 - p$, p), the individual is exactly indifferent when $p = 0.4$. Then we can determine u_1 from the requirement that $u_1 = 0.6u_0 + 0.4u_2$. Substituting $u_0 = 0$ and $u_2 = 1$ gives $u_1 = 0.4$.

Assigning a utility to an arbitrary prize proceeds analogously. Suppose we want to compute the utility of, say, \$7 million. We then must experimentally determine what value of p makes the individual indifferent between a certainty of \$5 million and the lottery (\$0, \$7 million; $1 - p$, p). If $p = 0.8$ is the critical value then the unknown utility u_7 must satisfy $u_2 = 0.2u_0 + 0.8u_7$. Substituting $u_0 = 0$ and $u_2 = 1$, we obtain $1 = 0.8u_7$ or $u_7 = 1.25$. The complete assignment of utilities to all possible prizes is called the individual's *von Neumann utility function*.

Viewed the other way around, if we are given the von Neumann utility function, e.g., $u(x) = \ln(x)$, and we wish to predict which of two possible lotteries an individual would choose,

$$X = (\$1, \$100; 0.8, 0.2)$$

$$Y = (\$10, \$1090; 0.99, 0.01)$$

then we simply compute the associated expected utilities:

$$U(X) = 0.8 \ln(1) + 0.2 \ln(100) \approx 0.92$$

$$U(Y) = 0.99 \ln(10) + 0.01 \ln(1090) \approx 2.35$$

²This choice of “origin and scale” is analogous to calling the temperature at which water freezes 0, and the temperature at which water boils 100. Note, in this regard, that the Fahrenheit and Centigrade temperature scales are affine transformations of one another.

and predict that the individual would choose Y since it has the larger expected utility.

Taking stock, we have seen that the Independence Axiom implies that indifference curves are parallel lines in probability space. This, in turn, implies that utilities can be found for the prizes so that the one lottery will have a larger expected utility than another if and only if it is strictly preferred.

1.4 Attitudes Toward Risk


The family of indifference lines is not the only family of parallel lines in probability space. The set of lotteries which have the same expected value is also a straight line and lotteries associated with a different expected value lie along a parallel. To see this consider the lottery $L = (v_0, v_1, v_2; p_0, p_1, p_2)$ where $v_0 < v_1 < v_2$ and, of course, $p_0 + p_1 + p_2 = 1$ with $p_i \geq 0$, $i = 0, 1, 2$. The expected value of this lottery is just $E = p_0 v_0 + p_1 v_1 + p_2 v_2$. Setting this equal to a constant, substituting $p_0 = 1 - p_1 - p_2$ and solving for p_2 gives

$$p_2 = -\frac{v_1 - v_0}{v_2 - v_0} p_1 + \text{another constant} \quad (2)$$

In the context of the Allais Lotteries this reduces to


$$p_2 = -\frac{1}{5} p_1 + \text{constant}$$

The set of lotteries with the same expected value as (\$1 million; 1) is illustrated below:



Other “constant expected value” lines are parallel. Now consider parallel arrow drawn to the left from A and note that, by construction, the expected value is unchanged by this movement, i.e., it is a mean preserving movement. Note also that the movement reduces the probability of receiving \$1 million and increases the probabilities of winning \$0 and \$5 million. The movement, in short, is a *mean preserving spread*. Any movement to the left along any “constant expected value” line is, similarly, a mean preserving spread.

The *IM* indifference curve derived earlier is illustrated together with the blue “constant expected value” line below.




Note that movements to the left along the “constant expected value” line entails movements to *lower* indifference lines. This follows directly from the fact that the indifference lines are steeper than the “constant expected value” lines. When this is the case, the person is said to be *risk averse*.

While the slope of the constant expected value lines are fixed by the size of the prizes, the slope of the indifference lines is a matter of taste. If the indifference lines had the same slope as the constant expected value lines, then the person would be indifferent to mean preserving spreads and would be called *risk neutral*. If the slope of the indifference lines were flatter than the slope of the constant expected value lines then mean preserving spreads would move the person to *higher* indifference lines and the person would be called *risk loving*.

1.5 The Shape of the Utility Function

Consider a lottery involving three prizes, $v_0 < v_1 < v_2$ and suppose that $u_0 = u(v_0)$, $u_2 = u(v_2)$ and $u_2 > u_1$. Three possible values for the utility, $u(v_1)$, of the intermediate prize illustrated below.



Consider u_1^N first. This value is selected so that (v_1, u_1^N) will lie *along* the line segment connecting (v_0, u_0) with (v_2, u_2) . As such, it must satisfy

$$\frac{u_1^N - u_0}{u_2 - u_0} = \frac{v_1 - v_0}{v_2 - v_0}$$

From Equation 1, the left side of this expression is the negative of the slope of an iso utility curve and from Equation 2, the right side is the negative of the slope of an iso expected value line. When (v_1, u_1) lies on the line segment, these slopes must be equal and the individual must, therefore, be *risk neutral* for lotteries involving these three prizes.

Now consider u_1^A which is selected so that (v_1, u_1^A) lies *above* the line segment connecting (v_0, u_0) with (v_2, u_2) . Here the left side of

$$\frac{u_1^A - u_0}{u_2 - u_0} > \frac{v_1 - v_0}{v_2 - v_0}$$


must be greater than the right side, the iso-utility line must then be steeper (more negative) than the iso expected value line and the individual must be *risk averse* for lotteries involving these three prizes.

Consider finally u_1^L which is selected so that (v_1, u_1^L) lies *below* the line segment connecting (v_0, u_0) with (v_2, u_2) . Here the left side of

$$\frac{u_1^L - u_0}{u_2 - u_0} < \frac{v_1 - v_0}{v_2 - v_0}$$

must be less than the right, the iso utility line must be flatter (less negative) than the iso expected value line and the individual must be *risk loving* for lotteries involving these three prizes.

If we now suppose that the graph of the utility function is concave over the interval $[v_0, v_2]$, as illustrated by the red curve below, then the individual will be risk averse for lotteries involving *any* set of prizes selected from this interval.



Alternatively, if the graph of the utility function is linear, as illustrated by the blue curve, then the individual will be risk neutral for lotteries involving prizes selected from this interval. Finally, if the graph of the utility function is convex, as illustrated by the green curve below, then the individual will be risk loving for lotteries involving prizes selected from this interval.