

PROBABILITY MODELS FOR ECONOMIC DECISIONS

Chapter 3: Utility Theory with Constant Risk Tolerance

[excerpt]

3.4. Note on foundations of utility theory

John Von Neumann and Oskar Morgenstern showed in 1947 how simple rationality assumptions could imply that any decision-makers' risk preferences should be consistent with utility theory. In this section, we discuss a simplified version of their argument, which justifies our use of utility theory.

To illustrate the logic of the argument, let us begin with an example. Suppose that we want to study a decision-maker's preferences among gambles that involve possible prizes between (say) \$0 and \$5000. We might start by asking him whether he would prefer to get \$1000 for sure or a lottery ticket that would pay \$5000 with probability 0.20 and would pay \$0 otherwise. If the decision-maker were risk-neutral, he would express indifference between these two alternatives (as they both have expected monetary value \$1000), but suppose that our decision-maker expresses a clear preference for the sure \$1000 over the lottery ticket. Next, we might ask him whether he would prefer the sure \$1000 or the lottery ticket that pays either \$5000 or \$0 if its probability of paying \$5000 were increased to 0.30. Suppose that the decision-maker now responded that he would be willing to give up a sure \$1000 for such a lottery ticket. So somewhere between 0.20 and 0.30 there should be some number p such that this decision-maker would be indifferent between getting \$1000 for sure and getting a lottery ticket that pays \$5000 with probability p and pays \$0 with probability $1-p$. Suppose we ask the decision-maker to think about such lotteries and tell us what probability p would make him just indifferent, and he tells us (perhaps after some long thought) that he would be just indifferent between \$1000 for sure and a lottery ticket that pays \$5000 with probability 0.27 and pays \$0 otherwise.

Next, we might ask this decision-maker a similar question about \$2000 for sure: "For what probability p would you be indifferent between getting \$2000 for sure and a lottery ticket that would pay \$5000 with probability p and \$0 otherwise?" Suppose that, in answer to this

question, the decision-maker says that he would be just indifferent between \$2000 for sure and a lottery ticket that pays \$5000 with probability 0.50 and pays \$0 otherwise.

Now consider the following two gambles

Gamble 1 pays \$2000 with probability 0.5 and pays \$1000 otherwise.

Gamble 2 pays \$5000 with probability 0.4 and pays \$0 otherwise.

Given only the above information about this decision-maker, should we be able to predict which of these two gambles he should prefer?

The answer to this question is Yes. If his preferences are logically consistent, he should prefer Gamble 2 here. Let me explain why. Suppose we offered to exchange Gamble 1 for the following two-stage gamble: At stage 1 we toss a fair coin, and if the coin comes up Heads (which has probability 0.5) then we give him a lottery ticket that pays \$5000 with probability 0.50 and pays \$0 otherwise, but if the coin comes up Tails then we give him a lottery ticket that pays \$5000 with probability 0.27 and pays \$0 otherwise. He already told us that he would be indifferent between \$2000 for sure and the first of these lottery tickets, and he would be indifferent between \$1000 for sure and the second of these lottery tickets. So he should be just indifferent between Gamble 1 and this two-stage gamble. But in this two-stage gamble, his ultimate prize is either \$5000 or \$0, and his probability of getting \$5000 is

$$0.5 \cdot 0.50 + 0.5 \cdot 0.27 = 0.385$$

(This calculation could be done with a probability tree, as in Section 1.6 of Chapter 1.) Gamble 2 also pays either \$5000 or \$0, and its probability of paying \$5000 is 0.4, which is greater than 0.385. So the decision-maker should rationally prefer Gamble 2 over the two-stage gamble, which he should think is just as good as Gamble 1. Thus he should prefer Gamble 2 over Gamble 1 if he is logically consistent. Figure 3.4 shows a spreadsheet illustrating this argument.

	A	B	C	D	E	F	G	H
1	The derivation of Utility Theory begins by choosing High and Low Prizes.							
2	5000	High Prize\$		0	Low Prize\$			
3	For other prizes \$X in between these prizes, we ask the Decision Maker:							
4	What probability p would make \$X as good as a lottery that pays							
5	the High Prize with probability p, the Low Prize with probability 1-p?							
6	We can then make a table of the DM's responses.							
7								
8	Given	DM assesses		So the DM considers the Prize\$ as good as				
9	Prize\$	p=P(High\$):		this equivalent High-Low lottery:				
10	0	0		0				
11	920	0.25		0				
12	1000	0.27		0				
13	2000	0.5		5000				
14	3300	0.75		5000				
15	5000	1	(rand)	5000				
16			0.321601					
17	Now consider any other lotteries.							
18								
19	LOTTERY 1							
20	Prize\$	Prob'y		P(Get High\$ in H-L-equiv)				
21	1000	0.5		0.27				
22	2000	0.5		0.5				
23	Simulations of Lottery 1 and an equivalent High-Low lottery:							
24	Lottery 1		H-L-equiv lottery	P(Get High\$ in H-L-equiv)				
25	2000			5000	0.385			
26								
27	LOTTERY 2							
28	Prize\$	Prob'y		P(Get High\$ in H-L-equiv)				
29	0	0.6		0				
30	5000	0.4		1				
31	Simulations of Lottery 2 and an equivalent High-Low lottery:							
32	Lottery 2		H-L-equiv lottery	P(Get High\$ in H-L-equiv)				
33	0			0	0.4			
34								
35	The decision-maker, if rational, should prefer Lottery 1 over Lottery 2							
36	if the value of cell E28 is greater than the value of cell E37.							
37	Now let the DM's assessed equivalent probabilities of the high prize							
38	in cells B10:B15 be called "Utilities" of the prizes in A12:A17.							
39	Then the decision maker, if rational should always prefer the lottery							
40	that has the higher Expected Utility (in cell E25 or cell E33)							
41	FORMULAS FROM RANGE A1:H40							
42	C15.	=RAND()						
43	D10.	=IF(\$C\$15<B10,\$A\$2,\$D\$2)			D10 copied to D10:D15			
44	D21.	=VLOOKUP(A21,\$A\$10:\$B\$15,2,0)			D21 copied to D21:D22,D29:D30			
45	A25.	=DISCRINV(RAND(),A21:A22,B21:B22)						
46	D25.	=VLOOKUP(A25,\$A\$10:\$D\$15,4,0)						
47	E25.	=SUMPRODUCT(B21:B22,D21:D22)			A25:E25 copied to A33:E33			

Figure 3.4. A spreadsheet illustrating the derivation of utility theory.

This argument can be generalized to show that that logically consistent decision-makers should always satisfy utility theory. To formulate this general argument, we first need to introduce some basic notation. Let us consider a decision-maker who may have to choose among a variety of possible monetary lotteries or gambles. A gamble may be denoted here by a random variable (in boldface) that represents the unknown monetary value that the decision-maker would be paid from this gamble, if he or she accepted it. Given such a decision-maker and given any two gambles **X** and **Y**, we may write

$$\mathbf{X} \succ \mathbf{Y}$$

to denote proposition that "the decision-maker would strictly prefer the gamble **X** over the gamble **Y**." Similarly, we may write

$$\mathbf{X} \sim \mathbf{Y}$$

to denote the proposition that "the decision-maker would be indifferent between the gambles **X** and **Y**."

Utility theory says that each decision-maker's preferences among all possible monetary gambles can be represented by some utility function $U(\bullet)$ such that, for any gambles **X** and **Y**,

$$\mathbf{X} \succ \mathbf{Y} \text{ when } E(U(\mathbf{X})) > E(U(\mathbf{Y})),$$

and

$$\mathbf{X} \sim \mathbf{Y} \text{ when } E(U(\mathbf{X})) = E(U(\mathbf{Y})).$$

Thus, according to utility theory, once we know a decision-maker's utility function, then we can predict the decision-maker's preferences over all possible monetary gambles. According to utility theory, then, different people may have different preferences for taking and avoiding risks because they have different utility functions.

To keep the argument simple, let us only consider lotteries where the outcome will be selected from some finite set of monetary prizes, which we may denote by

$$\{W_1, W_2, \dots, W_n\}.$$

Let W_1 denote the best of these possible outcomes, and let W_n denote the worst of these possible outcomes. Now, for any other possible prize W_i in this set, let us ask the decision-maker:

"If you were comparing the alternatives of (1) getting W_i dollars for sure, and (2) a binary

lottery in which you will get either the best prize W_1 or the worst prize W_n , then what probability of getting the best prize W_1 in this binary lottery would make you just indifferent between these two alternatives."

Obviously, if the probability of the best prize were 1 then the binary lottery would be better; but if the probability of the best prize were 0 then the sure W_1 would be better. So there should exist some probability of getting the best prize such that the decision-maker would be just indifferent between these two alternatives.

Given the decision-maker's answer to this question, let Z_i denote such a random variable that will be equal to either W_1 or W_n with the probability distribution that the decision-maker considers just as good as W_i for sure. That is, let Z_i be such that

$$P(Z_i=W_1) + P(Z_i=W_n) = 1, \text{ and } Z_i \sim W_i.$$

Notice that we must have

$$P(Z_n=W_1) = 0,$$

because the worst prize W_n would be worse than any lottery that gave any positive probability of the best prize. Similarly, we must have

$$P(Z_1=W_n) = 1.$$

Now consider any two general lotteries X and Y that have outcomes drawn from this finite set of possible prizes $\{W_1, W_2, \dots, W_n\}$. Since the decision-maker is indifferent between each W_i and the corresponding lottery Z_i , he should not care if we substitute a ticket to the lottery Z_i wherever he would have gotten the prize W_i . Repeating this substitution argument for each i from 1 to n , we conclude that the lottery X should be just as good, for this decision-maker, as a two-stage lottery in which the first stage gives a ticket to a lottery Z_1 or Z_2 or ... or Z_n respectively with the probabilities $P(X=W_1)$, $P(X=W_2)$, ..., $P(X=W_n)$. In this two-stage lottery, the final outcome will be either the best prize W_1 or the worst prize W_n , and the probability of getting the best prize W_1 is

$$P(X=W_1)*P(Z_1=W_1) + P(X=W_2)*P(Z_2=W_1) + \dots + P(X=W_n)*P(Z_n=W_1)$$

So the lottery X should be just as good, for this decision-maker, as a lottery that pays either the best prize W_1 or the worst prize W_n , where the probability of the best prize is

$$\sum_{i=1}^n P(X=W_i)*P(Z_i=W_1).$$

By a similar argument, the decision-maker should be indifferent lottery \mathbf{Y} should be just as good, for this decision-maker, as a lottery that pays either the best prize W_1 or the worst prize W_n , where the probability of the best prize is

$$\sum_{i=1}^n P(\mathbf{Y}=W_i) * P(\mathbf{Z}_i=W_1).$$

Among these two binary lotteries where the only possible outcomes are W_1 or W_n , the better one is obviously the one with the higher probability of the best prize W_1 . So the decision-maker should think that \mathbf{X} is better than \mathbf{Y} if

$$\sum_{i=1}^n P(\mathbf{X}=W_i) * P(\mathbf{Z}_i=W_1) > \sum_{i=1}^n P(\mathbf{Y}=W_i) * P(\mathbf{Z}_i=W_1).$$

The trick now is to let $U(W_i) = P(\mathbf{Z}_i=W_1)$. Then we have just shown that

$$\mathbf{X} \succ \mathbf{Y} \text{ whenever } \sum_{i=1}^n P(\mathbf{X}=W_i) * U(W_i) > \sum_{i=1}^n P(\mathbf{Y}=W_i) * U(W_i).$$

That is, the better gamble should always be the one with the higher expected utility. Thus, we have found a way of assigning utility values to prizes such that the decision-maker, if he is logically consistent, should always prefer the lottery with the higher expected utility.

You may notice that the utility numbers $U(X_i)$ that are generated by this argument are all between 0 and 1. However, there can also be valid utility functions in which the utility numbers are negative, or bigger than 1. In fact, we can multiply a utility function by any positive constant and add any other constant, and we will still have an equally valid utility function:

Fact 1. Suppose that the decision-maker's preferences over all monetary gambles can be represented by a utility function $U(\bullet)$. Suppose that another function $V(\bullet)$ is an increasing linear transformation of $U(\bullet)$. That is, suppose that there exist two numbers A and B such that $A > 0$ and

$$V(x) = A * U(x) + B, \text{ for all } x.$$

Then $V(\bullet)$ is also a valid utility function for representing this decision-maker's preferences over all monetary gambles.

Proof of Fact 1. For any two gambles \mathbf{X} and \mathbf{Y} , if $E(U(\mathbf{X})) > E(U(\mathbf{Y}))$ then

$$\begin{aligned} E(V(\mathbf{X})) &= E(A * U(\mathbf{X}) + B) = A * E(U(\mathbf{X})) + B \\ &> A * E(U(\mathbf{Y})) + B = E(A * U(\mathbf{Y}) + B) = E(V(\mathbf{Y})). \end{aligned}$$

So comparing expectations of V and U both yield the same preference ordering over gambles.