Ambiguity Aversion

Mark Dean

Lecture Notes for Spring 2015 Behavioral Economics - Brown University

1 Subjective Expected Utility

So far, we have been considering the 'roulette wheel' world of objective probabilities: our decision maker has been choosing between lotteries, which give a know probability of winning each of the available prizes. We are now going to think about the 'horse race' world of subjective probabilities (which we will also call uncertainty or ambiguity). Remember the key difference: when talking about a roulette wheel all (or at least most) people would agree on the probability of each outcome: if there are 38 slots in the wheel, the probability of each slot is $\frac{1}{38}$. We can therefore treat the probabilities of lotteries based on roulette wheels as observable (i.e. we, as the researcher, know what they are). In contrast, we may be much less confident about the probabilities associated with a horse race. Just because there are 3 horses does not mean that each has a $\frac{1}{3}$ chance of winning. In fact, reasonable people may assign very different probabilities to outcomes of the same horse race. Therefore we cannot treat these probabilities as externally observable. If they exist anywhere, they exist in the head of the decision maker, in the same way that utilities exist in the head of the decision maker. We therefore need another approach.

In order to capture what it is that the decision maker is uncertain about, we will introduce the concept of *states of the world*. These summarize the possible outcomes that could occur. If we think of the simple case of a race between 3 horses (Archibald, Byron and Cumberbach) then there are 6 possible states of the world, related to the different orders in which the horses can finish (excluding ties)

State	Ordering	Payoff Act f	Payoff Act g
1	A, B ,C	\$10	-\$10
2	A, C, B	\$10	-\$10
3	B, A, C	-\$10	-\$10
4	B, C, A	-\$10	-\$10
5	C, A, B	-\$10	\$20
6	C, B, A	-\$10	\$20

We then think of a decision maker choosing between *acts*. An act is defined by the outcome it provides in each state of the world. The following are two examples of acts:

- Act f: A \$10 even money bet that Archibald will win
- Act g: A \$10 bet at odds of 2 to 1 that Cumberbach will win

The above table shows what each of these acts will pay in each state of the world. Act f pays \$10 if Archibald wins, and -\$10 otherwise. Act g pays \$20 if Cumberbach wins, and -\$10 otherwise?

If you had to make a choice between act f and act g, how would you do it? The subjective expected utility model assumes the following procedure

1. Figure out the probability you would associate with each state of the world
2. Figure out the utility you would gain from each prize
3. Figure out the expected utility of each act according to those probabilities and utilities
4. Choose the act with the highest utility

So, in the above problem, the DM would choose act f over g if

$$\begin{aligned} & [\pi(ABC)+\pi(ACB)]\,u(10)+[\pi(BAC)+\pi(BCA)]\,u(-10)+[\pi(CBA)+\pi(CAB)]\,u(-10) \\ \geq\ & [\pi(ABC)+\pi(ACB)]\,u(-10)+[\pi(BAC)+\pi(BCA)]\,u(-10)+[\pi(CBA)+\pi(CAB)]\,u(20) \end{aligned}$$

If we rearrange this expression and (for the moment) assume that utility is linear, we get

$$\frac{[\pi(ABC)+\pi(ACB)]}{[\pi(CBA)+\pi(CAB)]}\geq\frac{3}{2}$$

In other words, you will choose f over g if you think the probability of Archibald winning is at least $\frac{3}{2}$ higher than the probability of Cumberbach winning.

We say that a DM's preferences over acts have a Subjective Expected Utility (SEU) representation if we can find some utility function and some probabilities that are consistent with their choices

Definition 1 *Let X be a set of prizes, Ω be a (finite) set of states of the world and F be the resulting set of acts (i.e. F is the set of all functions $f : \Omega \to X$). We say that preferences $\succeq$ on the set of acts F has a subjective expected utility representation if there exists a utility function $u : X \to \mathbb{R}$ and probability function $\pi : \Omega \to [0,1]$ such that $\sum_{\omega\in\Omega} \pi(\omega) = 1$ and*

$$\begin{aligned} f &\succeq g \\ &\Leftrightarrow \sum_{\omega\in\Omega} \pi(\omega) u\left(f(\omega)\right) \geq \sum_{\omega\in\Omega} \pi(\omega) u\left(g(\omega)\right) \end{aligned}$$

Some things to note about the subjective utility model:

- In the case of objective world, we only had to identify one thing from the decision maker's choices - the utility function. Everything else was externally observable. Now we have to identify two things: their utility function and their beliefs
- As in the case of choice under objective risk, we might be interested in understanding what behavior is consistent with a subjective utility representation and what is not. In other words we would like a set of axioms such that preferences have an SEU representation if and only if they satisfy those axioms. Sadly, doing so is beyond the scope of this course, but you should be aware that this has been done. In fact it has been done twice. First by Savage[1] and later (using a trick to make the process a lot simpler) by Anscombe and Aumann[2]
- You should be interested in whether the utility numbers and probability assessments from an SEU representation are unique: i.e. can we find multiple sets of utility numbers and probability assessments that match the same preferences. We showed in the last lecture that, for objective risk, utilities were unique up to a strictly positive affine transformation (i.e. if

[1] Savage, Leonard J. 1954. The Foundations of Statistics. New York, Wiley.

[2] Anscombe, F. J.; Aumann, R. J. A Definition of Subjective Probability. The Annals of Mathematical Statistics 34 (1963), no. 1, 199–205. doi:10.1214/aoms/1177704255.

some utility function u represented preferences, then v would represent the same preferences if and only if $v(x) = au(x) + b$ for some $a > 0$ and $b \in \mathbb{R}$). The same is true here. Helpfully, it turns out that the probability assessments are unique - i.e. there is only one set of probabilities that will explain any preferences. This is good, in the sense that we can think of choices as completely revealing the DM's beliefs.

2 The Ellsberg Paradox

Unfortunately, while the SEU model is a simple, neat and intuitively plausible model, it has problems as a predictive model of choice. In the same way that Allais came up with a thought experiment that demonstrated robust violations of Expected Utility theory for choice over objective risk, so Daniel Ellsberg came up with counter examples for choice under uncertainty. One such example we demonstrated in class:

- Fill a bag with 20 red and 20 black tokens (we will call this the 'risky' bag). Offer your subject the opportunity to place a \$10 bet on the color of their choice (i.e., if they choose 'red', then they receive \$10 if a red token is drawn from the bag and \$0 otherwise). Then elicit the amount x such that the subject is indifferent between playing the gamble and receiving \$$x$ for sure.

- Repeat the above experiment, but provide the subject with no information about the number of red and black tokens (we will call this the 'uncertain' bag). Then elicit the amount y such that the subject is indifferent between playing the gamble and receiving \$$y$ for sure.

Typically, in this experiment, we find that x is much higher than y - in other words subjects prefer betting on the 'risky' bag than on the 'uncertain' bag. This seems like a very natural way to behave - after all, we know a lot more about the risky bag than we do about the uncertain bag. However, it is not in line with SEU theory. Why is that? Well let's first think about the risky urn. Here, we know that there is an equal number of red and black tokens, so whichever color you bet on, the probability of winning is 50%. Thus (assuming the utility of \$0 is zero), the value of betting on the risky bag is

$$0.5u(\$10)$$

What about the uncertain bag? Well in order to apply SEU we need to figure out what are the states of the world. Here we have two relevant states: that the token drawn from the uncertain bag will be red, or that it will be black. We can therefore represent the acts of betting on red (r) and betting on black (b) as follows

State	r	b
red	10	0
black	0	10

If our decision maker is an SEU type person, then they act as if they assign a probability $\pi(red)$ to the token being red, and $\pi(black) = 1 - \pi(red)$ to it being black. Thus, the value of betting on red is

$$\pi(red)u(\$10)$$

and betting on black is

$$(1 - \pi(red))\, u(\$10)$$

Now notice that in the experiment, the decision maker gets to *choose* which color to bet on. Thus, the value of betting on the bag is

$$\max\{\pi(red)u(\$10).\,(1 - \pi(red))\, u(\$10)\}$$

The key thing here is, *whatever the decision maker's beliefs, the value of betting on the uncertain urn cannot be lower than the value of betting on the risky urn.* If $\pi(red) \geq 0.5$, then $\pi(red)u(\$10) \geq 0.5u(\$10)$. If $\pi(red) \leq 0.5$ then $(1 - \pi(red)) \geq 0.5$ and so $\pi(black)u(\$10) \geq 0.5u(\$10)$. The probability distribution in the risky urn is the worst possible beliefs to have given that you are allowed to pick what color to bet on. Thus, the observation that x is greater than y (sometimes called 'ambiguity aversion') cannot be squared with SEU

3 Maxmin Expected Utility

There have been many attempts to adjust the SEU model to allow for ambiguity aversion. One of the most popular is the Maxmin Expected Utility model introduced by Gilboa and Schmeidler.[3]

[3] Gilboa, Itzhak & Schmeidler, David, 1989. "Maxmin expected utility with non-unique prior," Journal of Mathematical Economics, Elsevier, vol. 18(2), pages 141-153, April.

This model also has a nice interpretation: imagine that you were of a pessimistic frame of mind, and that you thought the world was out to get you: in particular you though that the sneaky experimenter could observe whether you bet on red or black, then alter the number of red and black tokens in the bag in order to screw you over. If they saw you bet on red, they could increase the number of black tokens, and if they saw you bet on black the could increase the number of red tokens. This is basically the Maxmin Expected Utility model. Rather than having a single probability distribution over states of the world, the maxmin expected utility model assumes that the decision maker believes that there are a *set* of possible probability distributions out there, and that nature will pick the worst of those given the act that they have chosen

Definition 2 *Let X be a set of prizes, Ω be a (finite) set of states of the world and F be the resulting set of acts (i.e. F is the set of all functions $f : \Omega \to X$). We say that preferences $\succeq$ on the set of acts F has a Maxmin expected utility representation if there exists a utility function $u : X \to \mathbb{R}$ and convex set of probability functions Π and*

$$\begin{aligned} f &\succeq g \\ &\Leftrightarrow \min_{\pi \in \Pi} \sum_{\omega \in \Omega} \pi(\omega) f(\omega) \geq \min_{\pi \in \Pi} \sum_{\omega \in \Omega} \pi(\omega) g(\omega) \end{aligned}$$

The maxmin expected utility model can explain Ellsberg paradox type behavior. Say that the decision maker believed that there were a set of possible probabilities Π of a red token being drawn from the uncertain bag between 0.25 and 0.75 (with associated probabilities of a black token being drawn between 0.25 and 0.75). Then the value of betting on red in the uncertain back would be

$$\min_{\pi \in \Pi} \pi(red) u(\$10) = 0.25 u(\$10)$$

whereas the value of betting on black would be

$$\min_{\pi \in \Pi} \pi(black) u(\$10) = 0.25 u(\$10)$$

Thus the maximum value of betting on the uncertain bag would be $0.25u(\$10)$, less than the $0.5u(\$10)$ one would get from betting on the risky urn.

4 Ambiguity Aversion and No-Trade Prices

Ambiguity aversion has a number of implications for behavior in (for example) finance and insurance. One example, illustrated by Dow and Werlang,[4] is the existence of a 'no trade' region in asset prices.

Imagine that there is a financial asset that pays \$10 if a company is a success, and \$0 otherwise. The price of the asset is p. As an investor, you are can buy 1 unit of this asset, or you can short sell 1 unit of the asset. If you buy the asset you pay p and receive \$10 if the company is a success. If you short sell the asset, then you have receive p for sure, but have to pay \$10 if the company does well.

What would an SEU guy choose to do? Well they would figure out their probability of the company doing well - let's call this $\pi(good)$. The value of buying the asset is therefore

$$\pi(good)u(10) - u(p)$$

while the utility of short selling the asset is

$$u(p) - \pi(good)u(10)$$

Thus, (assuming that utility is linear, for convenience), the decision maker will buy if $p \leq 10\pi(good)$, and will short sell if $p \geq 10\pi(good)$

The key point here, is that, at *any* price the decision maker will be prepared to trade. For $p \leq 10\pi(good)$ they would like to buy, while for $p \geq 10\pi(good)$ they would like to sell (if $p = 10\pi(good)$ then they would be happy to do either.)

Now compare this to the behavior of a maxmin expected utility guy. Remember that such a person has a range of possible probabilities of the firm doing well. Let's say that $\pi^*(good)$ is the highest probability they will consider for the firm to be good, and $\pi_*(good)$ the lowest probability they will consider. Now, what happens if they buy the asset? In this case, they will assume the worst, and assign probability $\pi_*(good)$ to the firm doing well. They will therefore buy the asset if

$$\pi_*(good)u(10) - u(p) \geq 0$$

[4] Dow, James & Werlang, Sergio Ribeiro da Costa, 1992. "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," Econometrica, Econometric Society, vol. 60(1), pages 197-204, January.

If they sell the asset, they will again assume the worst, and assign the probability $\pi^*(good)$ to the firm doing well. They are therefore prepared to short sell the asset if

$$u(p) - \pi^*(good)u(10) \geq 0$$

Therefore, if $\pi_*(good) < \pi^*(good)$, then the decision maker will *not* be prepared to trade at any price. Again, assuming linear utility, if the price p falls in the range

$$10\pi^*(good) > p > 10\pi_*(good)$$

Then they will not be prepared to buy or sell. The reason is that their assessment of how well the firm will do depends on whether they have bought or sold the financial asset: if they buy the asset then the will assume the firm will do poorly. If they short the asset they will assume it will do well. for prices in that range, they would prefer to do neither.