Juliane Kaden – October, 24th, 2003.

Beyond Expected Utility Theory

Expected Utility Theory (EUT) has long been the prime building block of individual choice theory. This can be attributed to its simplifying and normative appeal which enabled economists to model individual choice in the presence of risk.

It was first proposed by Daniel Bernouilli (1738) in response to the so-called St. Petersburg puzzle which describes a gamble where a coin is flipped repeatedly until a head is produced. By entering the game, agents receive a payoff off, say, \$2 to the power of n where n is the number of the throw producing the first head. Although it is straightforward to realize that the expected monetary payoff is infinite, experiments showed that most prople would onlz be prepared to pay a relatively little amount to enter. Bernouilli suggested a discrepancy between the value an individual assigns to a gamble and its expected monetary value, i.e. individuals place a subjective value ("utility") on monetary values and the final expected payoff of the gamble is the expectation of these utilities. While Bernouilli's solution to the Paradox relies on a cardinal utility scale, von Neumann and Morgenstern (1947) derived the EU hypothesis from a set of preference axioms based on the assumption of individual rationality. In order to evaluate the relevance and accuracy of EU predictions comparing it to empirical finding and alternative choice theories, I will briefly present these axioms and EU derivation.

Choice under risk is modelled as following:

Note, to begin with, that risk in this respect has to be distinguished from simple uncertainty in so far as individuals are able to assign known and well-defined probabilities to all future prospects, i.e. they have a well-defined choice set based on clear probability distributions concerning the likelihood of each prospect. This is in contrast to choice under uncertainty where individuals cannot assign precise probabilities (Starmer, 2000).

1) Define a set of certain prospects (outcomes), X and assume that a certain prospect can always be characterized by specifying the values of a finite number of variables.

2) Define a set of risky prospects (lotteries), Y. The set of simple risky prospects is the set of probability distributions over X. Composite risky prospects consist of components which may themselves be risky prospects. A risky prospect is formally described as: $B=(A1, p1; A2, p2; \ldots; An, pn)$, where pi is the probability of Ai occurring, $pi \geq 0$ for all I, and $\Sigma pi=1$. Thus, the elements of X are to be understood as an exhaustive and mutually exclusive list of possible consequences which may follow from a particular course of action.

3) Assume that the decision maker has preferences over the set Y, $\geq$, i.e. preferences are defined over prospects. All consequences and probabilities are known to the agent, i.e. the agent faces a situation of risk (in contrast to uncertainty).

4) Consider now the following axioms of choice under risk:

A1 Ordering. The preference relation $\geq$ is a complete and transitive binary relation over Y.

A2 Independence. $A > B$ iff $(A, p; C, 1\text{-}p) > (B, p; C, 1\text{-}p)$, for $p \in (0,1)$

A3 Continuity. If A>B>C, then there is some p $\in(0,1)$ such that (A, p; C, 1-p) B.

5. Together the axioms of ordering and continuity imply that preferences over prospects can be represented by a function V(.) which assigns a real-valued index to each prospect. The function V(.) is a preference representation in the sense that $V(q)\geq V(r) \Leftrightarrow q\geq r$: an individual will choose the prospect q over the prospect r if, and only if, the value assigned to q by V(.) is no less than that assigned to r.

This way we arrive at well-defined preferences and their funcional relation simply by imposing minimal restrictions on the precise form of those preferences. On the contrary, the indepence axiom imposes strong restrictions of the form of preferences. It gives the standard theory most of its empirical content and as we will see later, it is violated in many instances. Alternatives to EUT will relax this assumption.

The Expected Utility Theorem says that if all three axioms hold, preferences can be represented by: $V(q)=\Sigma pi.u(xi)$ where q is any prospect, and u(.) is a "utility" function defined on the set of consequences and there will always be a utility function representing $\geq$ which satisfies the expected utility property.

EUT provides a very simple way of combining probabilities and consequences into a single "measure of value" and is therefore intuitively suitable for modelling decisions under risk. Furthermore, it has a number of appealing properties. Firstly, it is monotonic. Monotonicity can be defined as follows: Let x1, …, xn be consequences ordered from worst (x1)to best (xn). We may say that one prospect q=[p(q1),…,p(qn)] first-order stochastically dominates another prospect r=[p(r1),…,p(rn)] if for all I=1,…,n: $\Sigma p(qj)\geq\Sigma p(rj)$ with a strict inequality for at least one i. Monotonicity is the property that stochastically dominating prospects are preferred to prospects which they dominate. It is commonly acknowledged that any satisfactory theory– descriptive or normative – should embody monotonicity.

Secondly, the shape of the utility function has a behavioural interpretation. Concavity (Convexity) of U(.) implies risk averse (prone) behaviour; an agent with a concave utility function will always prefer a certain amount x to any risky prospect with expected value equal to x - $[u(E(A))\geq U(A)]$ where $E(A)=\Sigma pW$. The certainty equivalent of a lottery is the amount ce such that the decision maker (DM) is indifferent between the lottery and ce with certainty, i.e. U(A)=u(ce). The risk premium π is the difference between the expected value of a lottery and its certainty equivalent, i.e. π=E(A)-ce.

EUT as a normative device has very powerful implications for the analysis of choice under and therefore also for theoretical as well as real-world applications, such as game theory, insurance and stock-market pricing. However, does EUT really fit the empirical evidence? What are the descriptive merits of the theory and does it usefully represents actual choice behaviour? There is a large body of empirical evidence that suggests that is doesn't and that on the contrary, real world choice behaviour exhibits some serious violations of the standard theory. In the following I will briefly present the development up to the present of alternatives to EUT and empircal testing of actual choice.

It seems that violations of EUT fall under two broad headings which also spurred the theoretical development of alternative theories: firstly, violations which could possibly be explained in terms of some "conventional" theory of preferences, i.e. relying on well-defined preferences but integrating violations of the independence axiom. Secondly, violations which could only be accounted for in terms of "non-conventional" theories which challenge the assumption of well-defined preferences determining choice.

"Conventional theories" generally have the following features in common: preferences are represented by a function V(.) defined over individual prospects; the function satisfies ordering and continuity; and while V(.) is designed to permit observed systematic violations of the independence axiom, the principle of monotonicity is retained.

There is wide evidence that the independence axiom is frequently violated. Maurice Allais (1953) postulated two very common examples of such phenomena which have also played a particularly important role in shaping theoretical developments in non-EU theory. These are the so-called common consequence and common ratio effects. The common-consequence effect names the following situation: assume you have choices between pairs of prospects of the following form: $s^*=(y,p;c,1-p)$ and $r^*=(q,p;c,1-p)$, where $q=(x,\lambda;0,1-\lambda)$ and $0<\lambda<1$. The payoffs c, x and y are nonnegative (usually monetary) consequences such that $x>y$. Both prospects give outcome c with probablility 1-p (i.e. the "common consequence"). By the independence axiom choices between s* and r* should be independent of the value of c. However, numerous studies have found substantial influence of the value of c.

The common ratio effect can be thought of as follows: suppose agents have to choose between getting s**=($3000, 1) or entering a gamble r**=($4000, 0.8; 0, 0.2). Make up you mind!

Then they are asked what they would do if they had to choose between a 25% chance of getting $3000 or a 20% chance of getting $4000. Most people would go for sure in the first place and take the $3000, but opt for the 20% chance of the $4000 in the second place. However, this violates EUT because preferences should not depend on the value of p, yet empirical evidence reveals a tendency for people to switch their choice from s** to r** as p falls.

Graphically violations of preference can be easily shown with the help of the Marshak-Machina triangle. EU preferences in the triangle are parallel, linear, upward-sloping indifference curves. More risk aversion implies steeper curves. The angles of the triangle represent the probabilities of high, medium and low prize, i.e. each point on the triangle defines probabilities associated to a specific lottery.

EU in the Marshak-Machina triangle

The Allais Paradox I and II

Regarding the numerous experiments in laboratory and field conditions, there are three main stylized facts about the triangle:

1) EUT is most often violated near the boundaries (boundary effects), particularly if the agent is made certain of getting at least something by one of the lotteries (quasi-certainty boundary effect).
2) There is left edge fanning in, i.e. the agent becomes more risk-loving as he moves up the left edge of the triangle.
3) Indifference curves appear to be neither parallel nor linear (non-linearity).

There are some Generalized Expected Utility Theories (GEU) ("conventional") that fit these facts reasonably well and provide definite improvements over EU while at the same time being methodologically sound. These are for example the Generalized Disappointment model, Cumulative Prospect Theory, Concave RDEU and "Prospect Theory" (Camerer's version). By giving suitable decision probability weights, these models yield s-shaped utility functions.

s-shaped U(.) – Cumulative Prospect Theory

Evidence for the convexity in losses is unclear, but one dominant feature is that the utility function is steeper in losses than in gains, people are more averse to losses than they love gains (loss aversion). Furthermore, they seem to make their decisions according to some reference point (Cumulative Prospect Theory, Camerer, 1995). One example of loss aversion is for example the so-called endowment effect: if people are randomly given something (e.g. a prize, chocolate: Bateman et al., 1997), then this will tend to be worth more to them than if it had not been given.

Experimental evidence is often criticized on the following lines: (i) laboratory conditions would not reflect real world decision making environments,(ii) rewards could only inadequately create enough motivation to participate (i.e. giving no rewards could distort outcomes because people do not participate in the same way as in real situations; giving too high rewards could do as well; financial rewards may not really trigger the same intrinsic motivation etc.),(iii) repetition or non-repitition and the adequacy of capturing learning in a suitable framework. However, field work in many real world situations has replicated most of the laboratory outcomes. Most of the EUT violation appear furthermore relatively immune to changes in financial incentives. Sometimes (eg. Endowment effects, framing effects, probability updating etc.) violations can be reduced but most of the time they cannot be eliminated.

In spite of GEU being an improvement to the standard EU theory, they are still not able to capture all violations that systematically appear in experiments (e.g. intransitivities/instable preferences, incompleteness (incomplete choice sets), irreflexivitiy (endowment effect), framing effects, mishandling of probability information etc.). There is massive evidence for bounded rationality, i.e. agents behave neither completely rational nor completely stupid and display degrees of rationality on a larger band. Rubinstein (1998) notes that bounded rationality could possible explain (at least partly) for ambiguity aversion (evidence in the Ellsberg Paradox (Mukerji at al. 2002)(agents preferring certain outcomes to lotteries for which the exact probabilities are unknown), preference reversal, inability to optimise and framing effects. If preferences are unclear, this may mean that agents use framing to get more information and to infer their own preferences, i.e. framing may affect the way agents perceive a situation and categorise it. Categorisation and perception are again limited by the individual computational abilities. Although bounded rationality can indeed possibly explain for a lot of issues, but it is also a lot more difficult to handle it methodologically in order to develop well-specified, prescriptive and above descriptive models of choice.

The conventional theories are more or less preoccupied with finding some underlying feature of preferences that could explain violation of indepence while retaining well-defined preference functions. However, numerous studies suggest that violations in fact run a lot deeper. Conventional theory implicitely assumes procedure invariance (preferences over prospects are independent of the method used to elicit them) and description invariance (preferences over prospects are purely a function of the probability distributions of consequences implied by prospects and do not depend on how these given distributions are described). A common failure of procedure invariance is preference reversal.
Also there is wide evidence that small changes in the presentation (i.e. the “framing”) of prospects can significantly alter final choice (“Asian Disease” case, Tversky, Kahneman (1986)). Kahneman and Tversky’s original version of prospect theory (1979) (“unconventional”) allowed for a kink in the origin and did not rule out violations of monotonicity. Prospect Theory assumes two phases of the choice process: “editing” and “evaluation”. The editing stage is a cognitive process which consists of a preliminary analysis and simplification of the decision problem – agents essentially try to reduce the complexity of the problem and code prospects relative to a reference point (status quo), they put features of the problem into categorisations. Evaluation then

simply consists in choosing the highest valued prospect. Although this seems a lot a rational maximising choice process (i.e. has in principle a lot in common with the standard EU theory), it allows for a large part of possible violations. There are several points in which those can potentially come in before any EU-maximisation process becomes relevant: first of all, how is the analysis done? What determines the perception of the problem? How do agents categorise? Kahneman and Tversky (1979) already suggest that the reference point formation is essentially "dependent on the formulation of the offered prospects and the expectations of the decision maker", this allows for framing effects, omission of information etc..

Conventional theory has great difficulty in handling intransitivities (unstable preferences), framing effects and mishandling of probability information as these significantly question the basic assumptions of well-behaved preferences on which rational choice relies, also violations do not occur in a predictable way. This suggests that bounded rationality together with myopic agents (i.e. agents who put relatively more value on the near future then events which are further away, they are therefore mainly concerned over the short-run and are likely to exclude information about later events when making their choices) could explain these phenomena to a large extent.

Can we really assume that all agents have complete preference orderings for all economically relevant choices? There is large evidence of cyclical choices (e.g. A>B, B>C, C>A) according to some systematic pattern. Cyclical patterns are largely connected with so-called event-splitting effects (ESE) [p: p1, p2, p3 = 1/3 each; $: a>b (say, a=5, b=3) ESE: L1>L2 but L3<L4] (Starmer and Sugden, 1993; Humphrey, 1995).

	probability		
lottery	p1	p2	p3
L1	a	0	0
L2	0	b	b

	probability	
lottery	p1	p2 + p3
L3	a	0
L4	0	b

People obviously like to have as many positive amounts listed as possible. Further studies (Kahneman, Snell, 1990; (Kahneman, Varen, 1992) also show that only very few agents can correctly predict their own future tastes, the correlation between actual and predicted preference shifts tends close to zero. Psychologists (Liebenstein, Slovic, 1971) also found that people tend to make preference reversals when confronted with choices between $-bets and probability-bets (p-bets). Suppose you are offered the choice between a \$-bet with a lower probability of winning a higher sum of money and a p-bet with a higher probability of winning a lower sum of money. If you are now asked to first, choose between the two and second, to state your certainty equivalents for the two bets, what would you do? Most people tend to prefer the \$-bet in choice 1 but then assign a higher certainty equivalent to choice 2.

Preference reversal has been replicated many times (Grether and Plott, 1979), Lichtenstein and Slovic (1973) even replicated in a field study on a Las Vegas Casino. Experience and incentives may reduce preference reversals, but are unlikely to fully eliminate them (Knez, Smith, 1987).

All this suggests that agents are far from having clear, complete preferences. This implies the possiblity that nonconventional theories are probably even better and more useful in accounting for the empirical evidence than GEU.

Other instances where non-conventional approaches are needed to explain for the evidence are base-rate neglect and the conjunction fallacy. Base-Rate neglect is best illustrated by the "taxicab" problem (Tversky, Kahneman (1982): Imagine "a cab is involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following:

a) 85% of the cabs in the city are Green and 15% are Blue
b) A witness identified the cab as Blue. The court tested the reliability of the witness…and concluded that the witness correctly identified each one of the two colors 80% of the time and failed 20% of the time. What is the probability that the cab involved in the accident was Blue rather than Green?"

Bayesian updating suggests the use of the following: B=Cab is Blue; G=Cab is green; C_B=Cue that cab is blue

$P(B/ C_B)=P(C_B/B)P(B) / P(C_B)$ (Bayes' Rule)

i.e. $P(B/ C_B)= .8x.15 / .8x.15+.2x.85 \approx .41$

The modal answer of the experiment coincided with the 80% answer of the witnesses. The base rates were completely ignored, suggesting that people give immediate, personal information more weight than older information. Although base-rates are not always neglected, they remain substantially underweighed.

The conjunction fallacy violates the monotonicity axiom of probability, i.e. it is related to base-rate neglect but is a more basic violation. A fraction of people tends to rate the conjunction of two events more likely to occur than one of the events alone (e.g. Zizzo et al, 2000; Stolarz-Fantino at al, 2003).

Obviously, the decision making process is a much more complicated process than the standard EU theory suggests. Decision-makers take the context of their problem into account, they are sensitive to different representations of one and the same problem, i.e. in order to choose an action the DM must have a representation of the decision problem according to one or more dimensions. This process of perception is called categorisation of the decision problem. Regard for example a decision over lotteries. In order for the DM to make the final choice, he or she typically needs a specified number of outcomes, the probability p and an expected $-outcome for each outcome, this is the standard theory. However, apparently, there are also "irrelevant" aspects of the problem that may matter (the way of phrasing, the group of possible choices, the split-up etc.).

Rubinstein (1988) presents a model that focuses on one particular of such features: similarity. His model can explain the Allais paradoxes. For example:

Choice 1:

A1 (100.000,1)

B1 (500.000, 0.1; 100.000; 0.89)

Choice 2:

A2 (100.000, 0.11)

B2 (500.000, 0.1)

A1>B1 but A2<B2

Look at choice 2: 0.11 and 0.1 are almost the same, but 500.000 is definitely dissimilar from and much larger than 100.000, so one would choose B2>A2. For choice 1, A1 and B1 are dissimilar in probabilities, as you will get something for certain and people as we know tend to favour certainty (assuming quasi-certainty boundary effect), i.e. A1>B1. Rubinstein assumes a Boolean function (use of dummy variable, True=1, False=0) over the dimensions of the decision problem, which are in this case the monetary outcome $ and the probability p, i.e. p-similarity=1 and $-similarity=1. Relevant cases for the theory are when p-similarity=0 and $-similarity=0, which means that the decision problem is inconclusive and the DM uses some other unspecified dimension to choose between the lotteries. The other, in this respect more interesting, case is when p-similarity+$-similarity=1, because then the DM treats the similar dimensions as the same between the lotteries and decided on grounds of the dissimilar dimension (ex: B2>A2).

Other models which departed from the base model Rubinstein proposed play with the sequence of judgements in the decision problem (Leland, 1994), e.g. agents first maximize according to EUT and only when the appeal to preference is uninformative, agents will edit the lottery according to similarity or some other criterion. The similarity function is Boolean as well. The theory's outcome is basically the same, but the sequence assumption goes against the evidence.

Buschena and Zilberman (1999) question the Boolean similarity expression as too rigid and introduce similarity degrees. They find a 0.71 negative correlation between (dis)similarity rating and likelihood of choosing a risky lottery. They addressed too questions: first of all the degree of similarity, secondly, the possibly relevant other features of the decision problem. Similarity increases risk-taking and the noise error (heteroskedastic, i.e. increasing with the degree of similarity). Comparing the predictive power of their similarity model, GEU and noisy EU models, they find that the similarity model appears to be the single best model presently available. Furthermore, they find evidence against betweenness (Camerer, Ho, 1994), indifference curves are most likely not straight line in the triangle.

Is this the final point in research? Have we found satisfactory explanations and models for real individual choice behaviour? No, we are still a long way off from explaining everything we observe in experiments and field work. However, the similarity model being a nonconventional approach to individual choice, yields most promising results. Its relative performance in explaining for the evidence we find, is at least as good as GEU models if not better. Is this the end for EU models? From a descriptive point of view we are likely to be better of to recognise its limitations. This is especially important because EUT is so embedded in modelling individual choice which is again widely used for real –world purposes (insurance markets, capital markets, game theory & strategy etc.). Neglecting the empirical findings, the insufficiency of EUT and the availability of alternative models with higher explanatory power could possibly lead to the wrong conclusions. However, the EU framework can still be used for normative analysis.