Utility Functions

We've seen some claims about the requirements of rational belief, one key ingredient in rational choice. Let me say a few more things about the other ingredient: values, represented by utilities.

Earlier we represented the value an outcome had for an agent with a *cardinal utility*. Continuing to set aside, briefly, exactly what utility *is*, we'll continue to assume that something about an agents values or well-being fixes information about cardinal values that attach to outcomes. We can represent this, just like probability, with a ***utility function*** U.

We'll assume:

> ***Completeness***: U is defined for all events.

So, in a sense, individuals have "settled" views about the utility any outcome. For now, think of this as an idealizing assumption. Note that since we know that utilities are numeric values,

> ***Transitivity***: If U(a) > U(b) and U(b)>U(c), then U(a)>U(c).

This might constitute an assumption that the agent is rational (to this extent).

Maximizing Expected Utility

First let's develop some ways of comparing the payoffs you get from uncertain outcomes. I give you two options:

> (A) I'll flip a coin and give you $1000 if heads, $1 if tails.
> (B) I'll flip a coin and give you $2 if heads, $3 if tails.

If we ask "How valuable are these choices?" there's an intuitive sense in which the first is more valuable. Here's a ways of saying how by using what's call the *expected payoff* of each choice. Sum over the product of the likelihood of events with their payoffs. In this case:

A:	.5 x ($1000) + .5 x ($1) =	$500.5
B:	.5 x ($2) + .5 x ($3) =	$ 2.5

$500.5 is a lot more than $2.5. So, in an important sense, you stand to get more money on the first choice. Note, *neither expected payoff is a payoff you would ever get*. It's an amount chosen to represent a kind of "average" payoff given the uncertainty in outcome.

Maybe money isn't very valuable, or good (according to an agent). Maybe not all money is equally valuable or good. But *utility* represents what's valuable in choice (according to an agent) by stipulation. So if we want to figure out which actions are rational to perform under uncertainty, we should take 'average' utilities.

Now, in the above example our choice didn't affect the likelihood that the coin would land heads or tails. So when we "averaged" we could just take the probabilities as given. But we can't always do that. Sometimes, as we've seen, our actions *influence* the likelihood of the outcomes we're interested in. Recall our classic dilemma

	Pass	*Fail*
Study	10	0
Party	15	5

Studying *influences* whether or not you pass. How much? We now know (abstractly) how to quantify this kind of influence! We can appeal to the likelihood of passing *conditional* upon studying.

So we basically have all the tools we need to give an "expected utility" to individual choices that is completely general. Just two more definitions needed...

> States e_1, e_2, e_3,… *exhaustive* if (e_1 & e_2 & e_3 &…) is guaranteed to be true.
> States e_1, e_2, e_3,… *disjoint* if no two events in the list can jointly be true.

We're going to use exhaustivity to ensure we count up "all" the utility. We use disjointness to ensure we don't "double count" utility. Here's our notion:

> Let "A" be the state of performing action A, and E = {e_1, e_2, e_3,…} be an mutually incompatible set of states such that {not A, e_1, e_2, e_3,…} is exhaustive. Let P be a probability function, U be a utility function. Then the ***expected utility of performing action*** A, EU(A) is

$$EU(A) = P(e_1|A)U(e_1\&A) + P(e_2|A)U(e_2\&A) + P(e_3|A)U(e_3\&A) \ldots$$

For those in the know, the right hand side can also be written more succinctly...

$$\sum_{e_i \in E} P(e_i|A)U(e_i\&A)$$

This allows us to formulate our "super rule" of subjective, instrumental rational choice.

> ***Rule of Maximizing Expected Utility:*** Choose any action which maximizes expected utility.

The rule forms the core basis of decision theory. It's the go-to standard rule for subjective, instrumental rational choice. Before we see why, let's see it in action.

Example 1
You're haven't studied for our midterm at all, so you've got about a 20% of passing. Studying should quadruple that. You could also go out partying instead...Utilities are as in the previous table. What should you do to maximize expected utility?

Example 2
You want to drive to the store to get some food. There's a standing .0001% chance you'll get into a serious accident on the way and die. Utilities are as on handout 2.

Intuitively the rule gives the right answers here. Let's note some general features of maximizing utility.

(A) No immediate problems of actions and outcomes.

We saw with the study/partying case that obvious rules like dominance could fail if our actions influenced the outcomes of choice. This forced us to look for decision tables with outcomes which “factored out” the influence of choice, which can be difficult. Now there's no problem. Expected Utility has built into it mechanisms for tracking the influence of actions on outcomes.

(B) No inexhaustive or conflicting rules.

When we looked at basic choice strategies on handout 1, we saw that many of the most plausible choice strategies only yielded definitive verdicts in special cases. The more controversial choice strategies conflicted with each other. Here neither problem arises. We have an all-purpose rule that (allegedly) applies to all choices. It's a coherent rule, so it can't “conflict” with itself. And it seems to have a very firm intuitive grounding.

(C) Some nice structural features.

(i) Maximizing expected utility *subsumes* strong dominance when outcomes are independent of choices (and the actions themselves aren't intrinsically valuable). That's just what we'd hope for.

(ii) Maximizing expected utility gives transitive recommendations. If you should chose A over B, and B over C, then you should chose A over C.

(iii) Maximizing expected utility satisfied the “independence of irrelevant alternatives”. If you should do A when your options are A, B, and C, then you should do A when your options are A and C.