Simplified Preference Elicitation

Greg Hines
Agents, Interaction, and Complexity Research
School of Electronics and Computer Science or Other
University of Southampton

Preference elicitation

Help people make decisions in complex settings

Examples:

- Whether to install solar panels
  - Revenue from panels is dependent on amount of sunlight; a wet summer may result in low revenue
  - If the solar panels are bought using a loan, a loss risk averse homeowner would be especially sensitive to months where the solar panel revenue does not cover the loan payment

- Whether to buy an electric car
  - An electric car has a higher upfront cost than an equivalent petrol or diesel car
  - The benefit of an electric car is a long term savings in refilling cost and CO_2 emissions
  - Are the future benefits worth the upfront costs?

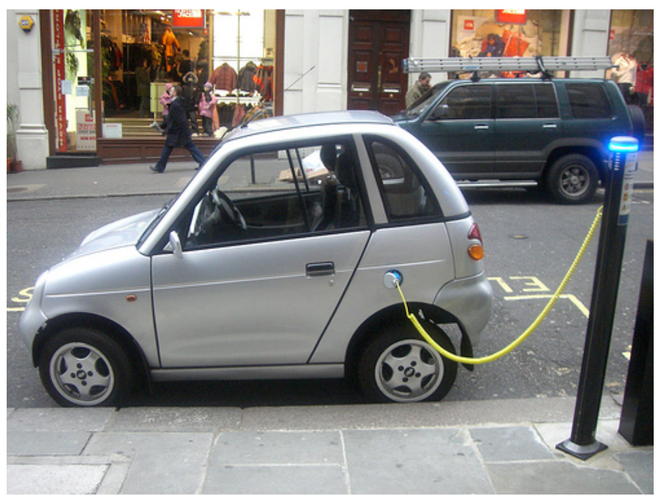

Standard Approach in Preference Elicitation

Standard Model

- Each decision, *d,* is represented by a probability distribution over a set of outcomes.
- Each outcome is represented by a range of possible utility values, e.g. $u^{min} < u(x) < u^{max}(x)$, and a probability distribution, $P(u)$, over that range
- We recommend the decision which maximizes the expected expected utility (EEU) (Chajewska, Koller and Parr, 2000)

$$EEU(d,P) = \int_u EU(d,u)P(u)\,du$$

Improving Recommendations

To improve the EEU, the standard approach to refine the range of utility values is to ask the user a standard gamble query (SGQ):

"Would you prefer a guaranteed outcome *x* or a gamble where, with probability *p*, you will receive the best possible outcome, and otherwise, you will receive the worst possible outcome?"

If the user prefers the guaranteed outcome, then $u(x) > p$, and otherwise, $u(x) < p$.

Unfortunately, SGQs pose a high cognitive burden on the user.

Thus, we have developed a new approach based on modeling users with the utility function (Tversky and Kahneman, 1992)

$$u(x) = \begin{cases} x^r & if\ x \geq 0 \\ -s|x|^r & otherwise \end{cases}$$

where r measures how risk averse the user is and s measures how loss averse they are.

We then ask the user,

"Would you accept a gamble with a probability p of winning $x, and otherwise losing $y?"

A Graphical Interpretation

An indifference curve (IC) gives all values of (*r,s*) for which the user would be indifferent towards the proposed gamble.

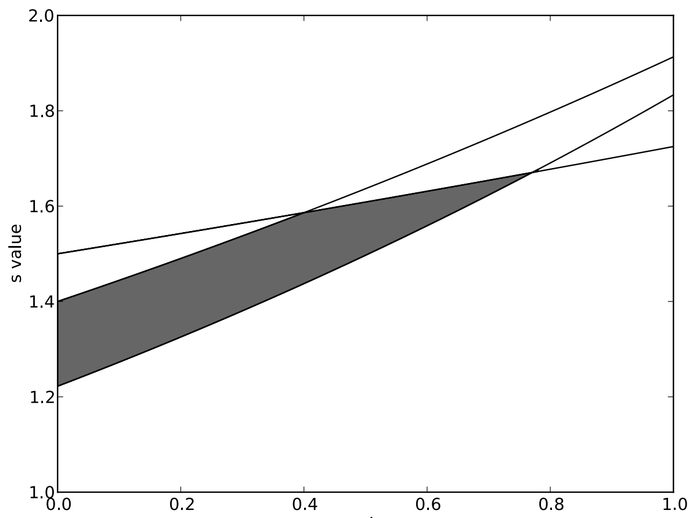

If the user accepts the gamble, we reject all (*r,s*) above the IC and otherwise, we reject all (*r,s*) below the IC. We then calculate the EEU assuming all remaining (*r,s*) values are equally likely to be correct.

Future Work

Human experiments – possibly with Amazon Mechanical Turk

Research into optimizing query parameters