Towards a Theory of Decision–Making without Paradoxes

Roman V. Belavkin (r.belavkin@mdx.ac.uk)

School of Computing Science,
Middlesex University, London NW4 4BT, UK

8 April 2006

MOTIVATION

- The *expected utility* theory leads to many paradoxes.
- Data suggests that humans and animals often violate principles of the rational choice (Allais, 1953; Ellsberg, 1961; Myers, Fort, Katz, & Suydam, 1963; Tversky & Kahneman, 1981).
- Many AI systems and cognitive architectures (e.g. ACT–R, Anderson & Lebiere, 1998) use the $E\{u\}$.
- Noise seems to play an important role optimising the behaviour (Belavkin & Ritter, 2003)

THE EXPECTED UTILITY THEORY

The classical decision–making theory is due to Pascal and Fermat, Bernoulli (1738/1954), von Neumann and Morgenstern (1944), Savage (1954) and Anscombe and Aumann (1963).

1. Represent preferences by some *utility* function $u : X \to \mathbb{R}$

$$x \succ y \iff u(x) > u(y)\,,$$

2. Under uncertainty, the *expected utilities* ($E\{u\}$) are considered (due to Pascal and Fermat):

$$p \succ q \iff \sum_{z \in Z} p(z)\, u(z) > \sum_{z \in Z} q(z)\, u(z)\,,$$

where Z is a set of prizes, P a set of probability measures.

DECISION MAKING IN ACT–R

In ACT–R (Anderson & Lebiere, 1998), the choice between several alternative decisions (i.e. rules) is implemented by the conflict resolution mechanism. A rule with the highest *utility* is selected: $i = \arg\max U_i$, where

$$U_i = P_i G - C_i + \text{noise}(s)$$

rule's properties :

- P_i – probability of success
- C_i – cost (e.g. time)

global parameters (constants) :

- G – goal value
- s – controls noise variance σ^2

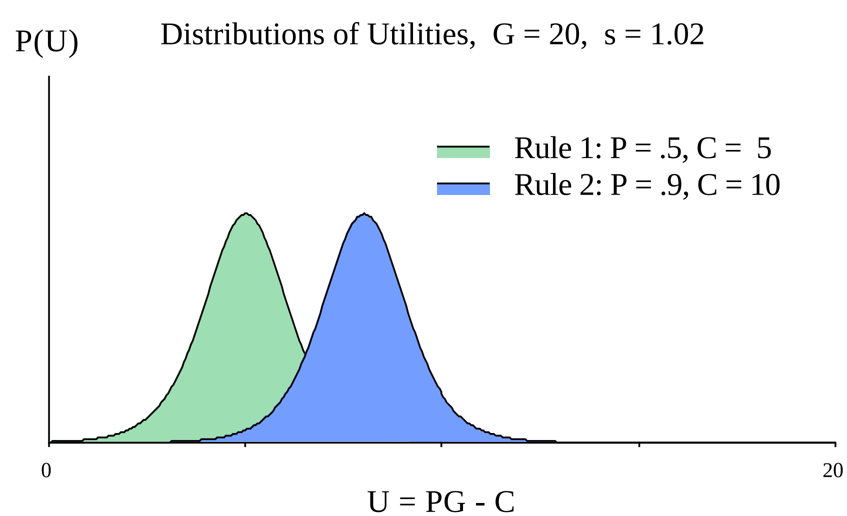

ACT–R AND EXPECTED UTILITY

- For each decision, two outcomes: Success $\vee$ Failure
- Let $U^s = U($Success$)$ and $U^f = U($Failure$)$. Then

$$\begin{aligned} E\{U\} &= P^s U^s + P^f U^f \\ &= P^s U^s + (1 - P^s) U^f \\ &= P^s (U^s - U^f) + U^f \end{aligned}$$

- If $G = U^s - U^f$ and $U^f = -C$, then $E\{U\} = PG - C$
- ACT–R uses the expected utility and therefore is prone to all the paradoxes.

THE RATIONAL DONKEY PARADOX

?

↙ ↘

Haystack A **Haystack B**

- $\max EU$ fails if there is no unique $\max$ (use a roulette wheel).
- Human subjects and animals always retain some degree of randomness (e.g. Myers et al., 1963).
- ACT–R uses noise (:egs) to model this.

THE ALLAIS PARADOX

Due to Allais (1953). Consider two lotteries A and B

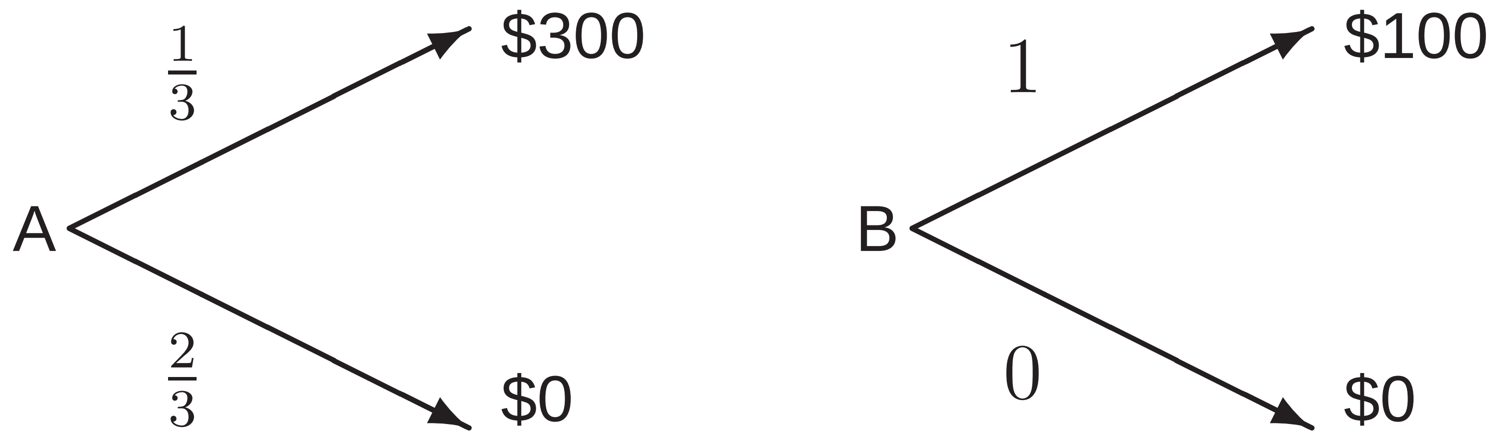

$$\frac{1}{3} \cdot \$300 + \frac{2}{3} \cdot \$0 = \$100 \qquad 1 \cdot \$100 + 0 \cdot \$0 = \$100$$

About 80% of subjects prefer $A \prec B$.

THE ALLAIS PARADOX (LOSSES)

When the gains are changed to losses, the preferences reverse

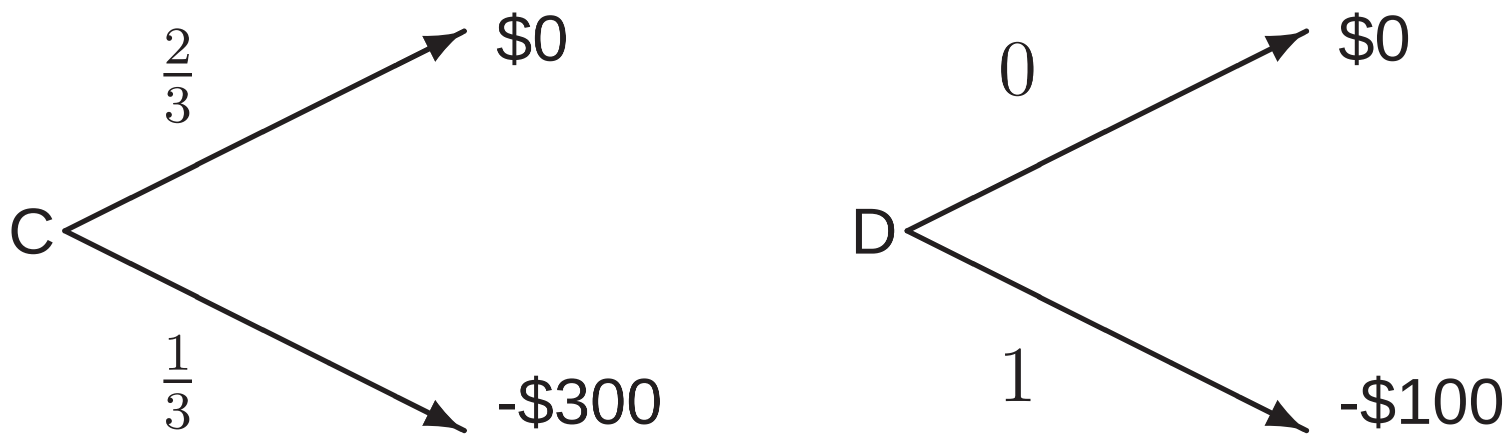

$$\frac{2}{3} \cdot 0 - \frac{1}{3} \cdot \$300 = -\$100 \qquad 0 \cdot \$ - 1 \cdot \$100 = -\$100$$

About 80% of subjects express preference $C \succ D$

Confirmed in many studies (e.g. Tversky & Kahneman, 1981)

Professional traders behave this way too (List & Haigh, 2005).

THE ELLSBERG PARADOX

Due to Ellsberg (1961). Consider two lotteries A and B, and probabilities of outcomes for A are given

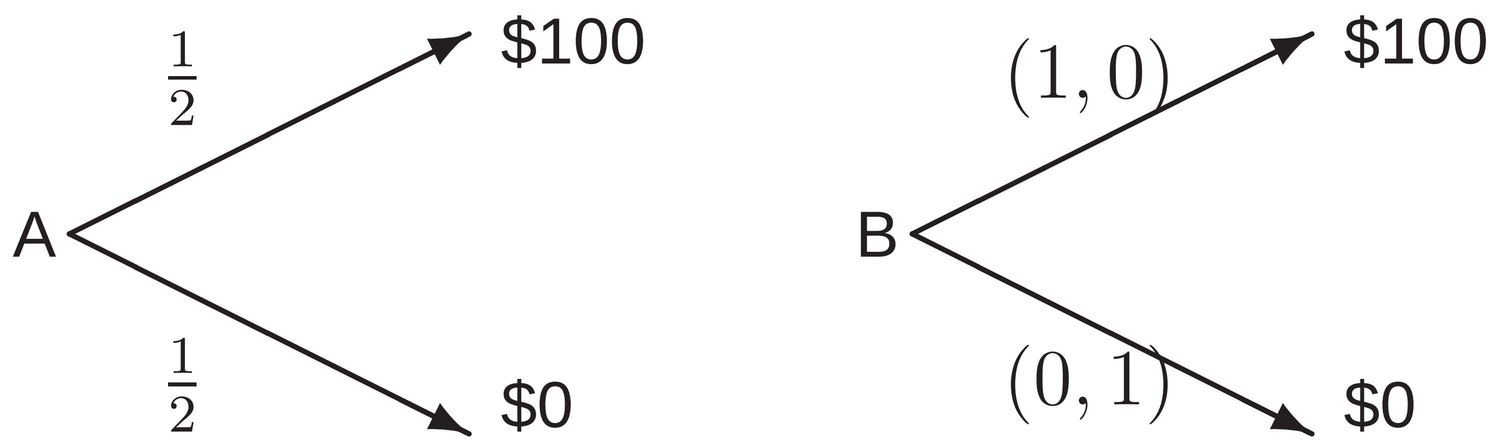

$EU(A) = \$50$ $\qquad$ $EU(B) = \$50$

$$A \succ B$$

ISSUES TO CONSIDER

Decision–making under uncertainty is estimation of utilities (sampling or sensing) and then choosing based on the highest estimate.

- Many paradoxes occur when $E\{u\}$ is used to estimate future utility based on some $p(u)$.
- Is $E\{\}$ the optimal estimator of utility?
- Are the lottery problems good examples of estimation (regression) problems?
- Should we use **subsymbolic** or **symbolic** mechanisms to build models of the paradoxes (e.g. quantitative vs qualitative)?

WHAT IS THE BEST ESTIMATION OF UTILITY?

x unobservable random (e.g. future utility)

y observable (e.g. past utilities)

- Estimation of x through y is finding some regression function $x \approx g(y)$

$$x = g(y) + C(x, y)$$

- If $C(x, y) = (x - y)^2$, then optimal $g(y) = E\{x \mid y\}$
- If $C(x, y) = 1 - \delta_x^y$ (i.e. success if $y = x$, failure otherwise), then optimal $g(y) = \arg_x \max p(y \mid x) \equiv \max L(x, y)$ (maximum likelihood estimate).

MAX LIKELIHOOD vs. EXPECTED VALUE

- Often (e.g. for non–Gaussian) $\arg\max p(y \mid x) \neq E\{x \mid y\}$

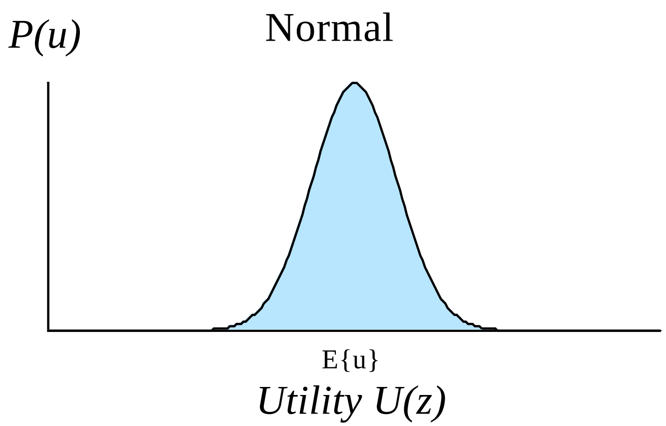

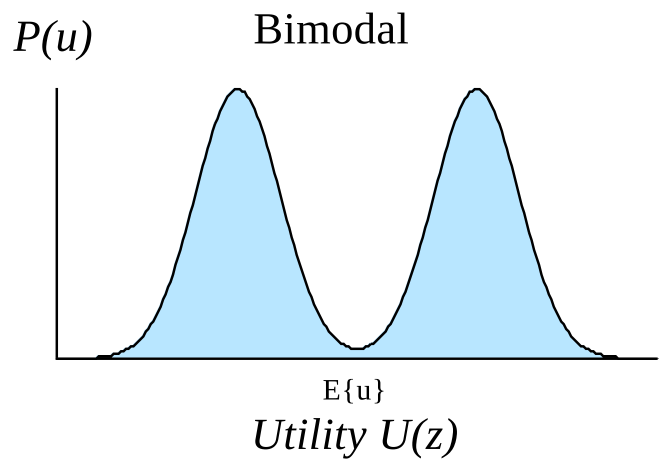

- Indeed, the MLEs of lotteries A and B are $\$0 < \$100$.
- Similarly, the MLEs of lotteries C and D are $\$0 > -\$100$

$$A \prec B\,, \qquad C \succ D$$

EXPLORATION vs EXPLOITATION

The quality of estimation depends on information about the utility in $P(u)$. What is the best sampling strategy?

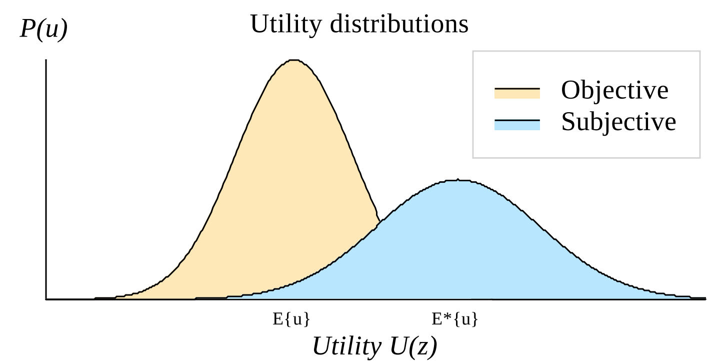

- Exploration with maximum information

$$I(u) = -\log p(u) \sim \frac{1}{p(u)}$$

- This contradicts exploitation

$$\arg\max p(u)$$

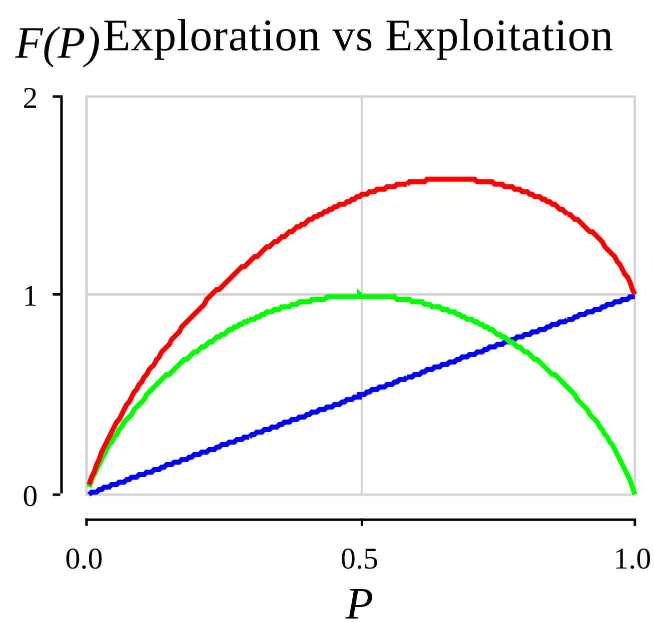

RANDOM ESTIMATION

- Instead of $x \approx \sum_y p(y)y$ or MLE, we can use $p(x \mid y)$ to draw random estimates of x (i.e. Monte–Carlo simulation).
- If $F(y)$ is the distribution function for $p(y)$ (PDF), then sampling can be done using the inverse PDF method:

$$x \approx F^{-1}(p)\,, \quad \text{where } p \in (0,1)$$

- Asymptotically, this estimation is similar to both MLE and $E\{x\}$.
- Given $P(u)$, decisions can be made based on the largest random estimates of utility.

RANDOM UTILITY IN ACT–R

Each rule i has a history of successes and failures $P_i(\text{Outcome})$. For a set of conflicting rules, the following scheme is used to generate random utilities U_i

$$\begin{aligned} P_i(\text{Outcome}) &\to \text{Success} \vee \text{Failure} \\ U_i &= U_i^s \vee U_i^f \\ &= G + U_i^f \vee U_i^f \\ &= G - C_i \vee -C_i \end{aligned}$$

where C_i is the cost. We can also use Gamma noise

$$U_i = G - \text{Gamma}(\theta_i) \vee -\text{Gamma}(\theta_i)$$

INVERSE PDF (A and B)

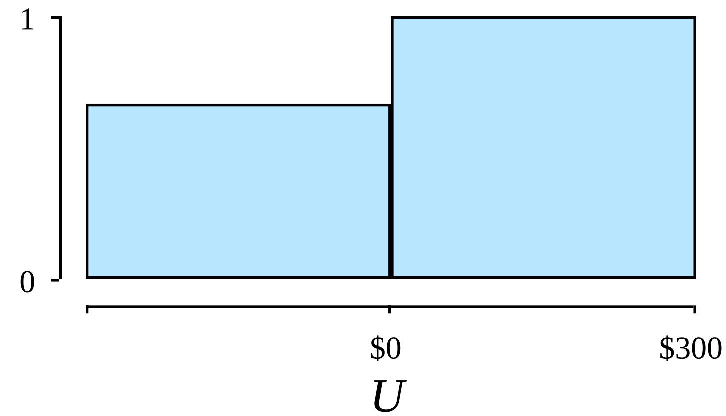

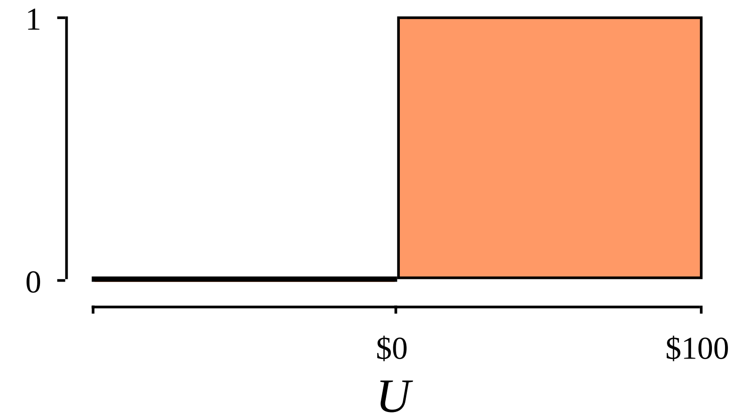

Utility $= F^{-1}(P)$

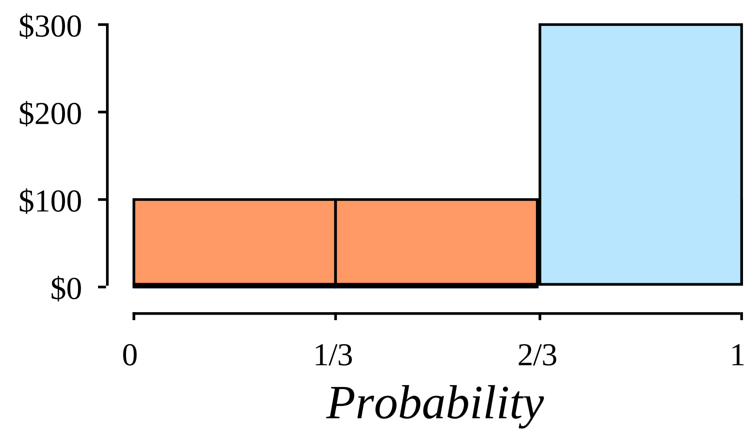

$RU_A < RU_B$ 2 out of 3 times, which supports experimental evidence

$$A \prec B$$

INVERSE PDF (C and D)

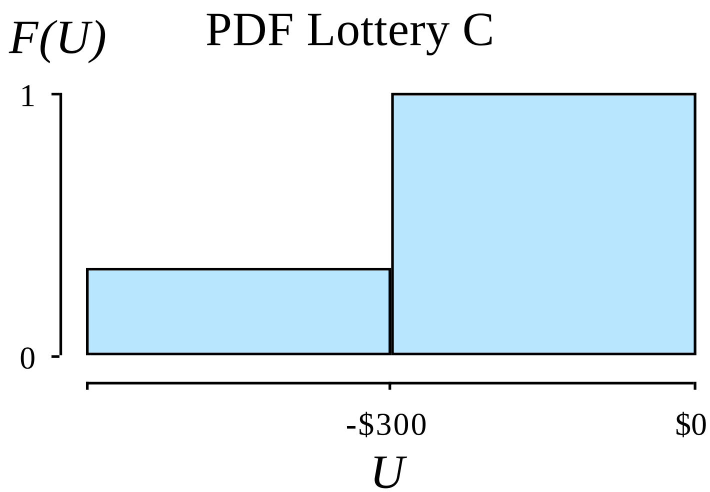

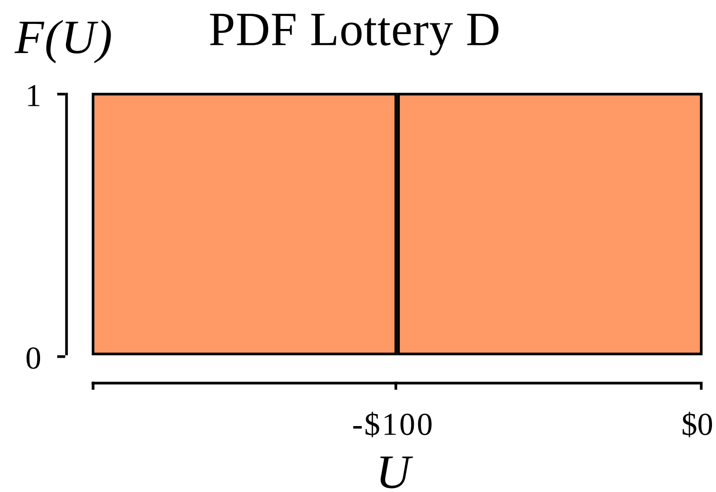

Utility $= F^{-1}(P)$

$RU_C > RU_D$ 2 out of 3 times. Again, corresponds to experimental results

$$C \succ D$$

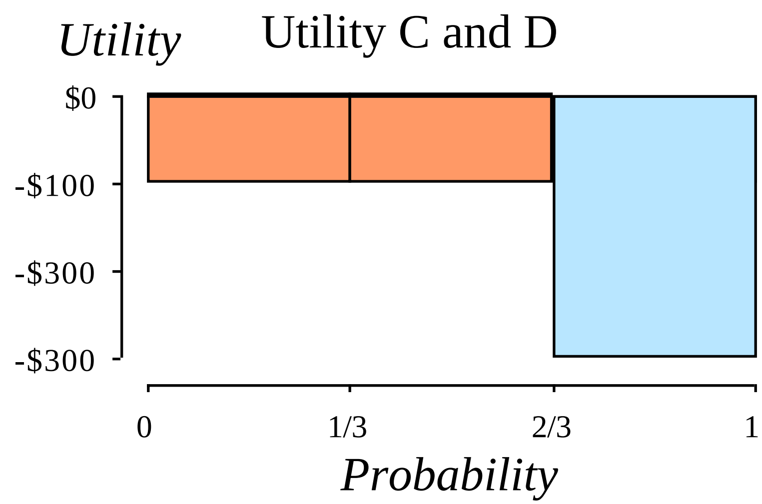

MODEL RESULTS (THE ALLAIS PARADOX)

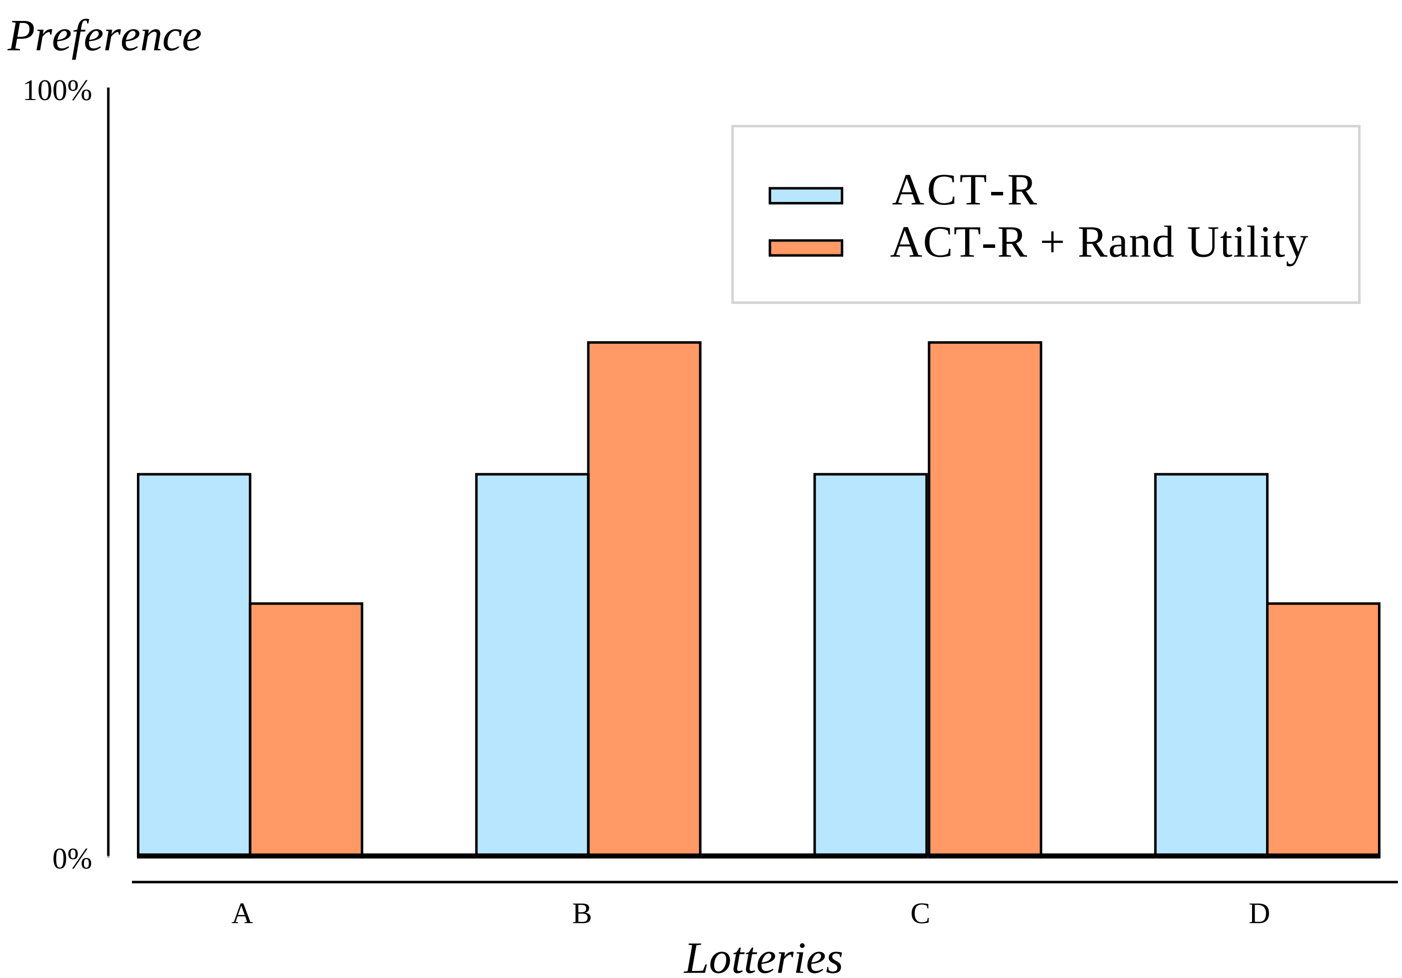

THE EFFECT OF PROBABILITY

$P(A)$	A	B	$A \prec B$	
1/3	$600	$200	28%	(Tversky & Kahneman, 1981)
1/4	$1000	$240	16%	(Tversky & Kahneman, 1981)
1/4	$1000	$240	38%	(List & Haigh, 2005)

- In the lottery task, P of uncertain prize does not seem to have consistent effect on % of subjects preferring it.
- Probabilities are given, no sampling allowed.
- Could qualitative decision–making be used to model the task symbolically?

THE LOGIC OF CHOICE

$\sim$ indifference (any can be chosen)

$\succ$ preference (the preferred is chosen)

Object A		Object B
attribute 1	$\succ$	attribute 1
attribute 2	$\prec$	attribute 2
$\vdots$		$\vdots$
attribute n	$\sim$	attribute n

Combining preferences

$$\succ \text{ and } \prec \quad = \quad \sim \quad = \quad \succ \text{ or } \prec$$

Combination of $\succ$ or $\sim = 3/4$ chance of choosing A.

QUALITATIVE CHOICE MODEL (SYMBOLIC)

In ACT–R, can be implemented at least in two ways

- Using parallel rules for each attribute

```
(p A or B, attribute 1 A ≻ B ==> choose A)
                    ⋮
(p A or B, attribute n A ≺ B ==> choose B)
```

- Using OAV triplets (e.g. `A gain better`) and rules such as

```
(p A or B, =oav A better ==> choose A)
                ⋮
(p A or B, =oav B better ==> choose B)
```

CHOOSING LOTTERIES QUALITATIVELY

Attribute:	A		B	C		D
U^s	\$300	$\succ$	\$100	\$0	$\sim$	\$0
P^s	1/3	$\prec$	1	1/3	$\succ$	0
U^f	\$0	$\sim$	\$0	-\$300	$\prec$	-\$100
P^f	2/3	$\prec$	0	2/3	$\succ$	1
Union		$\prec$			$\succ$	

Moreover, the chance of choosing A is

$$P(A) = \frac{1}{4} \times \left(1 + 0 + \frac{1}{2} + 0\right) = \frac{3}{8} \approx 38\%$$

OTHER OBSERVATIONS

- Can model the Ellsberg paradox: If one prefers certainty, then $A \succ B$ follows.
- Symbolic model can be improved to take into account other effects of choosing (e.g. how many attributes are considered, how long does it take to choose. etc).
- How to encode real values, such as $P = 0.1, 0.2$? Both small or one larger than another? Can explain the violations of the **independence axiom** (Allais, 1953).

CONCLUSIONS

- The $E\{u\}$ theory does not provide the optimal decision–making strategy (Belavkin, 2005).
- The MLE and the random utility estimation of utility can explain some data contradicting the $E\{u\}$ theory.
- Qualitative reasoning be used to make choice, and symbolic models can also explain the data.
- Subsymbolic mechanisms may be better for modelling tasks where some statistics has to be learnt (e.g. trials and errors).
- Symbolic models may also (and perhaps better) represent the decision–making in the lottery task.

THE ORIGINS OF THE EXPECTED UTILITY THEORY

- Blaise Pascal and Fermat used $E\{\}$ to solve several problems (e.g. rolling a dice, etc).
- Pascal also proposed to use $E\{u\}$ to argue that a rational agent should believe in God (yet, there are some people who are atheists).
- Because there is no prior $P(\text{God})$, the max. likelihood or the random estimation of utility may explain this fact.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'École americaine. *Econometrica*, *21*, 503–546.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum.

Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, *34*, 199–205.

Belavkin, R. V. (2005, December). Acting irrationally to improve performance in stochastic worlds. In M. Bramer, F. Coenen, & T. Allen (Eds.), *Proceedings of AI–2005, the 25th SGAI International Conference on Innovative Techniques and*

Applications of Artificial Intelligence (Vol. XXII, pp. 305–316). Cambridge: Springer. (ISBN 1-84628-225-X)

Belavkin, R. V., & Ritter, F. E. (2003, April). The use of entropy for analysis and control of cognitive models. In F. Detje, D. Dörner, & H. Schaub (Eds.), *Proceedings of the Fifth International Conference on Cognitive Modelling* (pp. 21–26). Bamberg, Germany: Universitäts–Verlag Bamberg. (ISBN 3-933463-15-7)

Bernoulli, D. (1954). Commentarii acad [English translation]. *Econometrica*, *22*, 23–36. (Reprinted from *Scientiarum Imperialis Petropolitanae*, 1738, *5*, 175–192)

Ellsberg, D. (1961, November). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, *75*(4), 643–669.

List, J. A., & Haigh, M. S. (2005). A simple test of expected utility theory using professional traders. *PNAS*, *102*(3), 945–948.

Myers, J. L., Fort, J. G., Katz, L., & Suydam, M. M. (1963). Differential monetary gains and losses and event probability in a two–choice situation. *Journal of Experimental Psychology*, *77*, 453–359.

Neumann, J. von, & Morgenstern, O. (1944). *Theory of games and economic behavior* (first ed.). Princeton, NJ: Princeton University Press.

Savage, L. (1954). *The foundations of statistics*. New York: John Wiley & Sons.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.