Modelling the Paradoxes of Decision–Making

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

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

8 July 2005

OVERVIEW

1. Expected utility and ACT–R
2. The Rational donkey paradox
3. Noise and dynamic variance
4. The Allais paradox
5. The random utility solution
6. The Ellsberg paradox
7. Future work

DECISION MAKING

- Classical decision–making theory is due to von Neumann and Morgenstern (1944), Savage (1954) and Anscombe and Aumann (1963).

- Despite the differences in treating the *uncertainty*, the main idea is that of a *utility* and its *expected value* (the EU), and the choice made by maximising EU

$$\text{Decision}(i) = \arg\,\max \sum_{i=1}^{n} P_i U_i$$

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

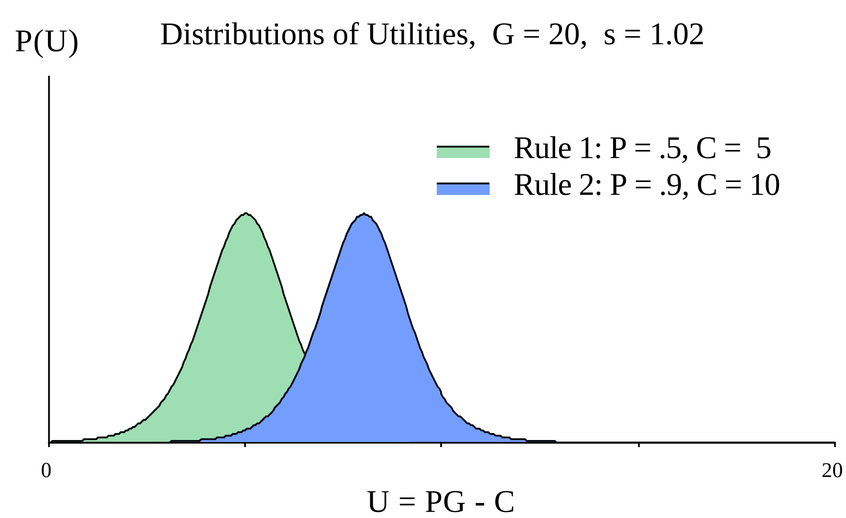

ACT–R AND EXPECTED UTILITY

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

$$?$$

$$\swarrow \qquad \searrow$$

$\boxed{\textbf{Haystack A}}$ $\qquad$ $\boxed{\textbf{Haystack B}}$

- $\max EU$ theory fails when there is no unique $\max$.
- ACT–R uses noise (`:egs`) which ensures this does not happen
- How large should be noise variance?
- There are other paradoxes related to $\max EU$.

DYNAMIC EXPECTED GAIN NOISE

- Dynamic noise variance has been discussed recently (e.g. Taatgen, 2001; Belavkin, 2001)
- Entropy–based method to control `:egs` was proposed in Belavkin and Ritter (2003)

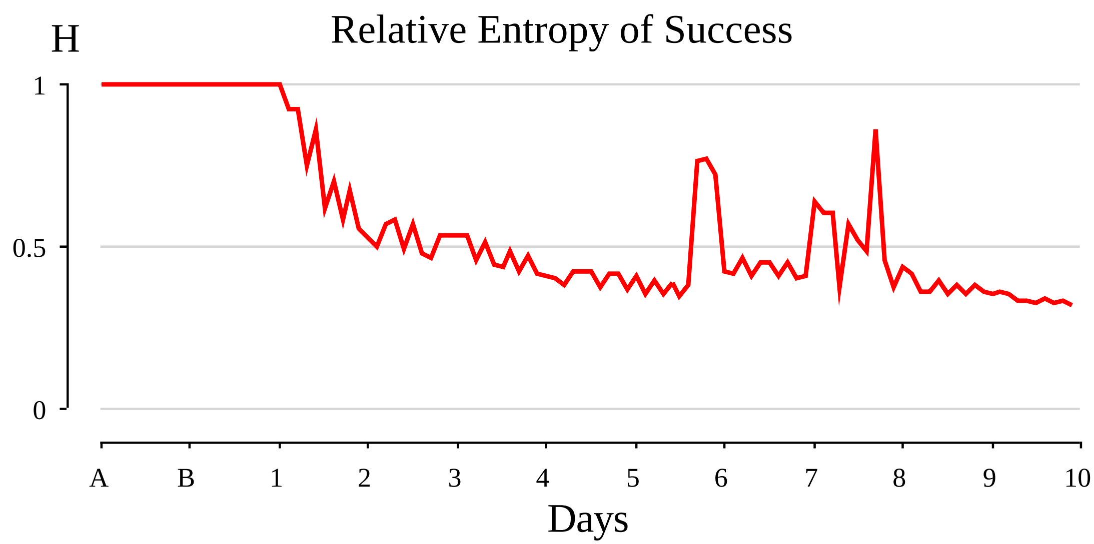

GAMMA NOISE (OPTIMIST)

- The probability distributions of utilities can be used directly to control the variance (Monte–Carlo).
- The time component of the cost can be estimated using Poisson distribution $p = 1 - e^{-1/\theta}$ (Belavkin, 2003)

$$U_i = P_i G - \text{Gamma}(\theta_i)\,, \quad \text{where } \theta = \frac{\text{Efforts}}{\text{Successes}}$$

- The OPTIMIST overlay (Belavkin & Ritter, 2004) for ACT–R is available at

http://www.cs.mdx.ac.uk/staffpages/rvb/

THE ALLAIS PARADOX (GAINS)

Due to Allais (1953). Also studied by Tversky and Kahneman (1974) in many interpretations. Consider two lotteries A and B

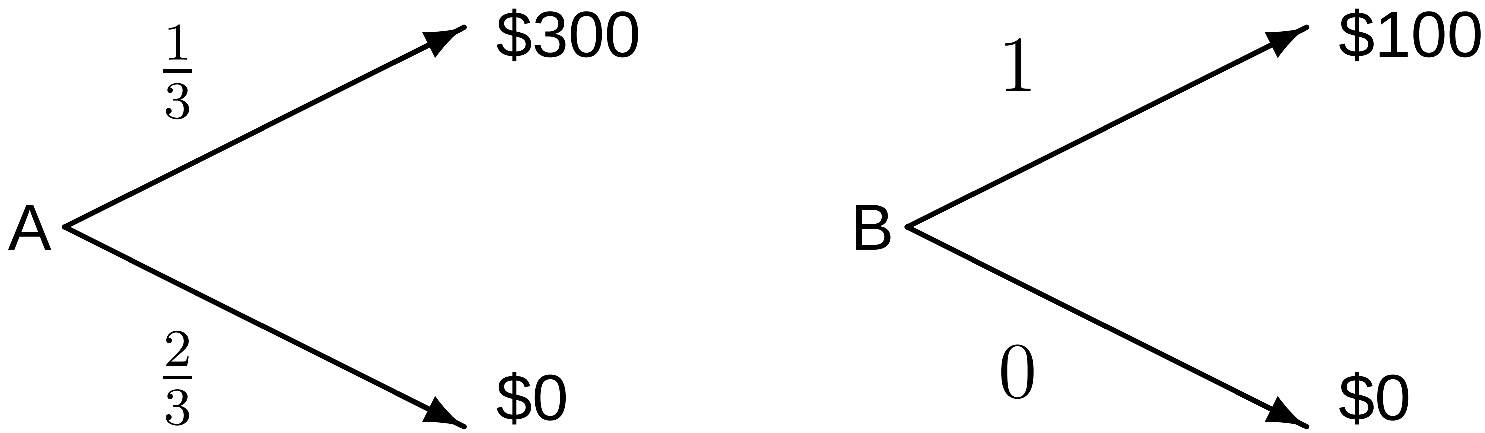

$$\frac{1}{3} \cdot \$300 + \frac{2}{3} \cdot \$0 = \$100 \qquad 1 \cdot \$100 + 0 \cdot \$0 = \$100$$

About 80% of subjects express preference $A \prec B$

THE ALLAIS PARADOX (LOSSES)

When the gains are changed to losses, the preferences reverse

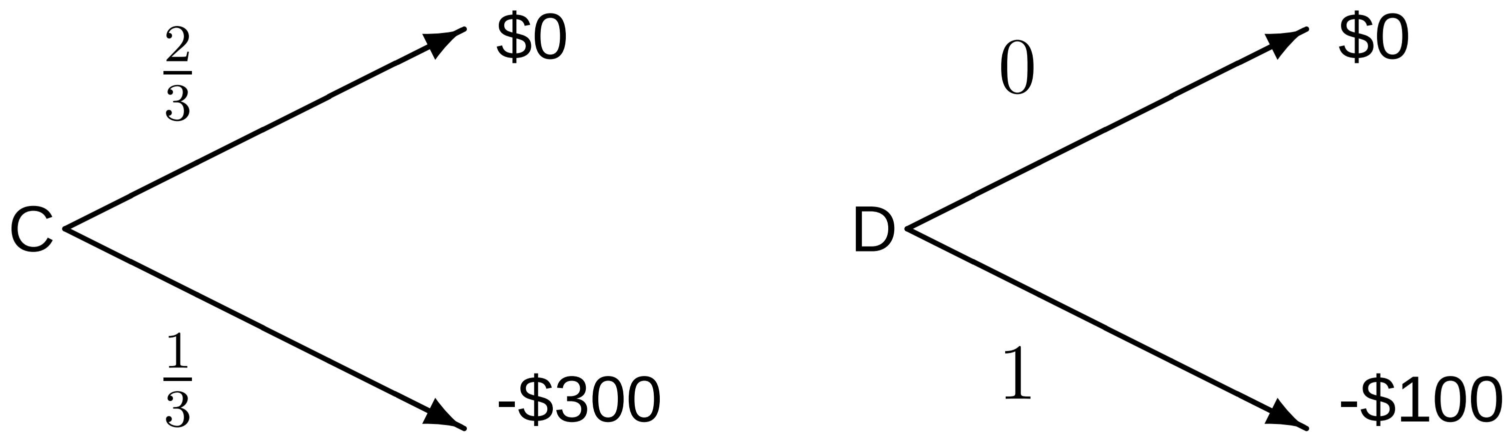

$$\frac{2}{3} \cdot 0 - \frac{1}{3} \cdot \$300 = -\$100 \qquad 0 \cdot \$ - 1 \cdot \$100 = -\$100$$

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

FRAMING OF DECISIONS

- Tversky and Kahneman (1974) suggested *decision framing* theory of using a function $\pi(P)$ of the probability.
- In ACT–R, one suggests to use G as the 'framing' global parameter

Lottery A and B $\frac{1}{3} \cdot G - \$0 \quad \prec \quad 1 \cdot G - \$0$

Lottery C and D $\frac{2}{3} \cdot G - \$0 \quad \succ \quad 1 \cdot G - \$100$

- However, the above formulae are incorrect as C should also be relative to goal value G. The correct formula is $C = G - U^s$
- Note also that not 100% of subjects preferred as above.

RANDOM UTILITY

For each decision i, the outcome is sampled from its distribution $\boldsymbol{P}(\text{Outcome} \mid i)$ conditional to rule i. The utility of this outcome is called *random utility* RU_i

$$\text{Decision}\, i = \arg\max_i RU_i \;, \;\text{ where } RU_i \leftarrow \boldsymbol{P}(\text{Outcome} \mid i)$$

Here $\boldsymbol{P}(\cdot \mid i)$ is probability distribution of successes and failures for a given rule, and RU_i is the utility of each outcome.

Sampling can be implemented using the inverse PDF method

$$\text{Outcome} = F^{-1}(P)\,, \quad \text{where } P \in (0,1)$$

RANDOM UTILITY vs $\max EU$

- Tested on agents with Bayesian learning of Markov Decision models (i.e. transitional probability tables $\boldsymbol{P}_{ij}^{k}$).
- The random utility agents are as good as the $\max EU$ agent, and often outperformed them 2:1 (Belavkin, 2005)

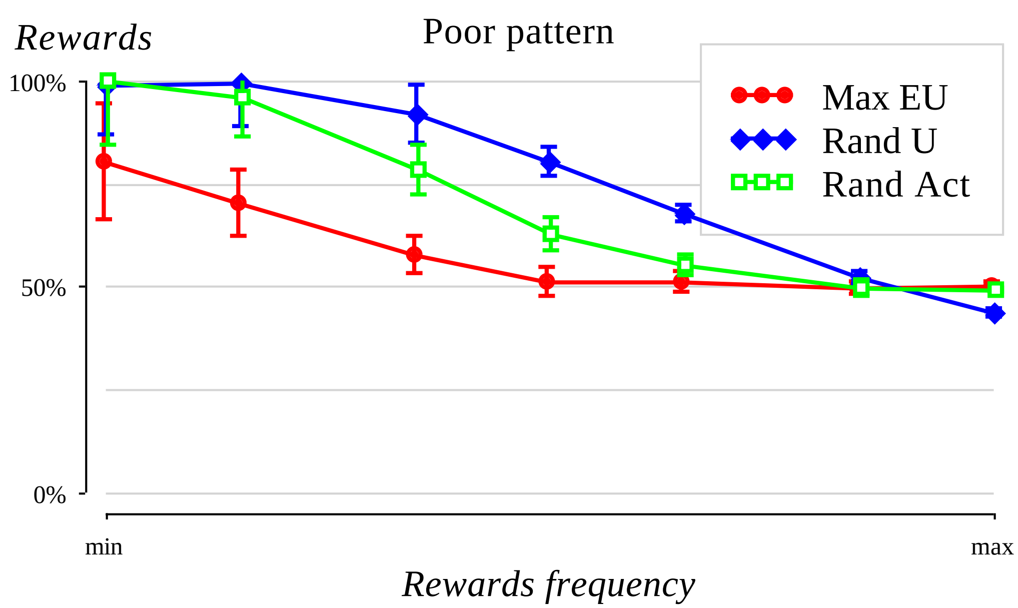

INVERSE PDF (A and B)

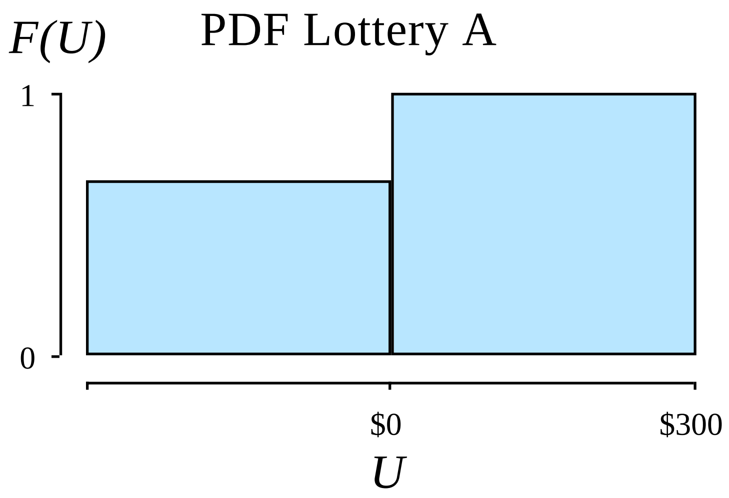

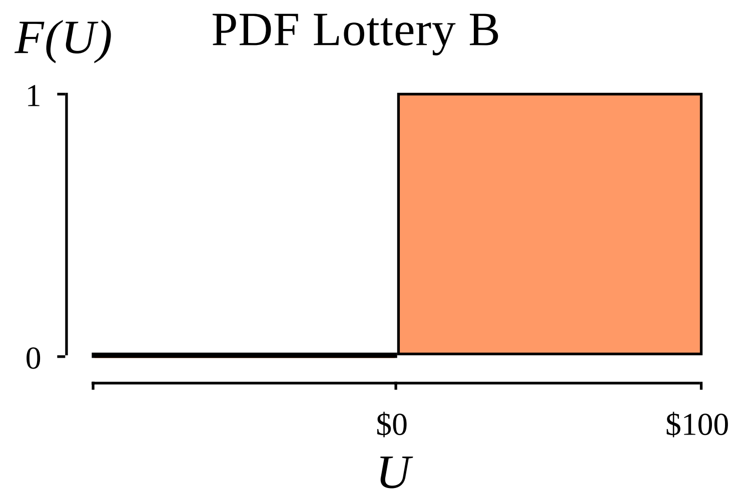

$$\text{Utility} = F^{-1}(P)$$

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

$$A \prec B$$

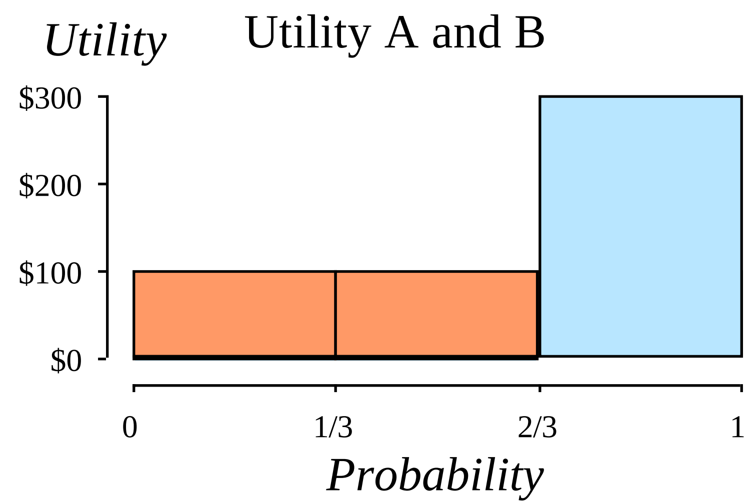

INVERSE PDF (C and D)

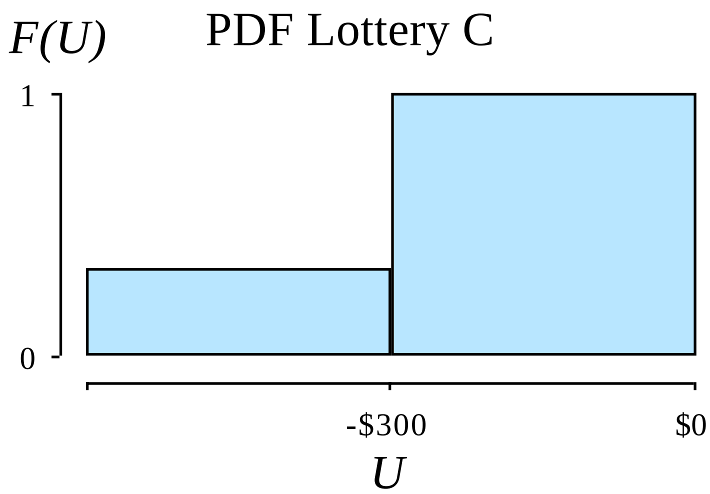

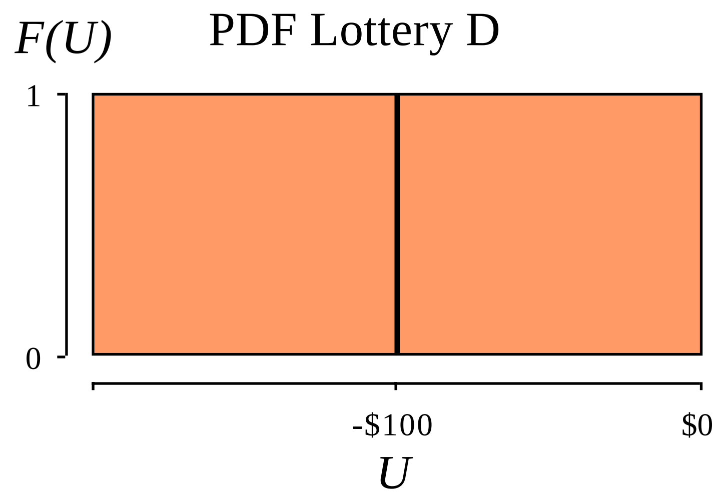

Utility $= F^{-1}(P)$

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

$$C \succ D$$

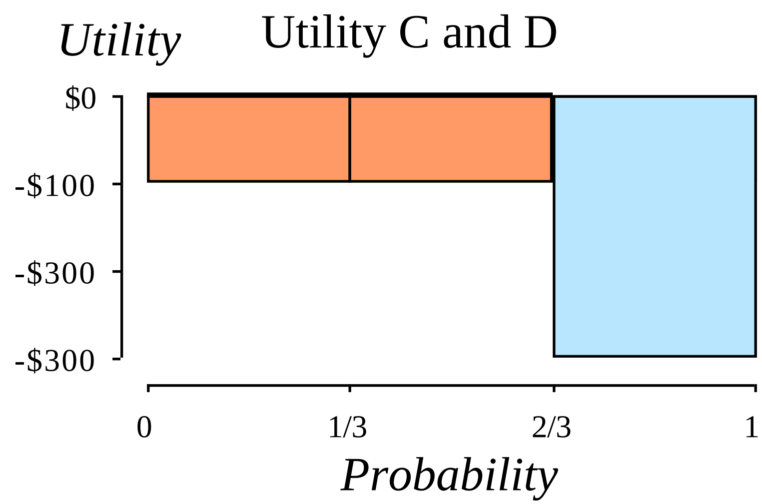

RANDOM UTILITY IN ACT–R

Each rule i has a history of successes and failures $\boldsymbol{P}(\text{Outcome} \mid i)$.

For a set of conflicting rules, the following scheme is used to generate random utilities RU_i

$$
\begin{aligned}
\boldsymbol{P}(\text{Outcome} \mid i) &\longrightarrow \text{Success} \vee \text{Failure} \\
RU_i &= U_i^s \vee U_i^f \\
&= G + U_i^f \vee U_i^f \\
&= G - C_i \vee -C_i
\end{aligned}
$$

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

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

PROPERTIES OF RANDOM UTILITY

- The expected value of random utility

$$\begin{aligned} E\{RU_i\} &= P_i(G - C_i) - (1 - P_i)C_i \\ &= P_i G - C_i \end{aligned}$$

- Allows to model the Allais paradox
- The use of Gamma noise implements the features of the OPTIMIST conflict resolution: Rule specific and dynamic noise variance $\sigma^2 = \theta^2$.

MODEL RESULTS (THE ALLAIS PARADOX)

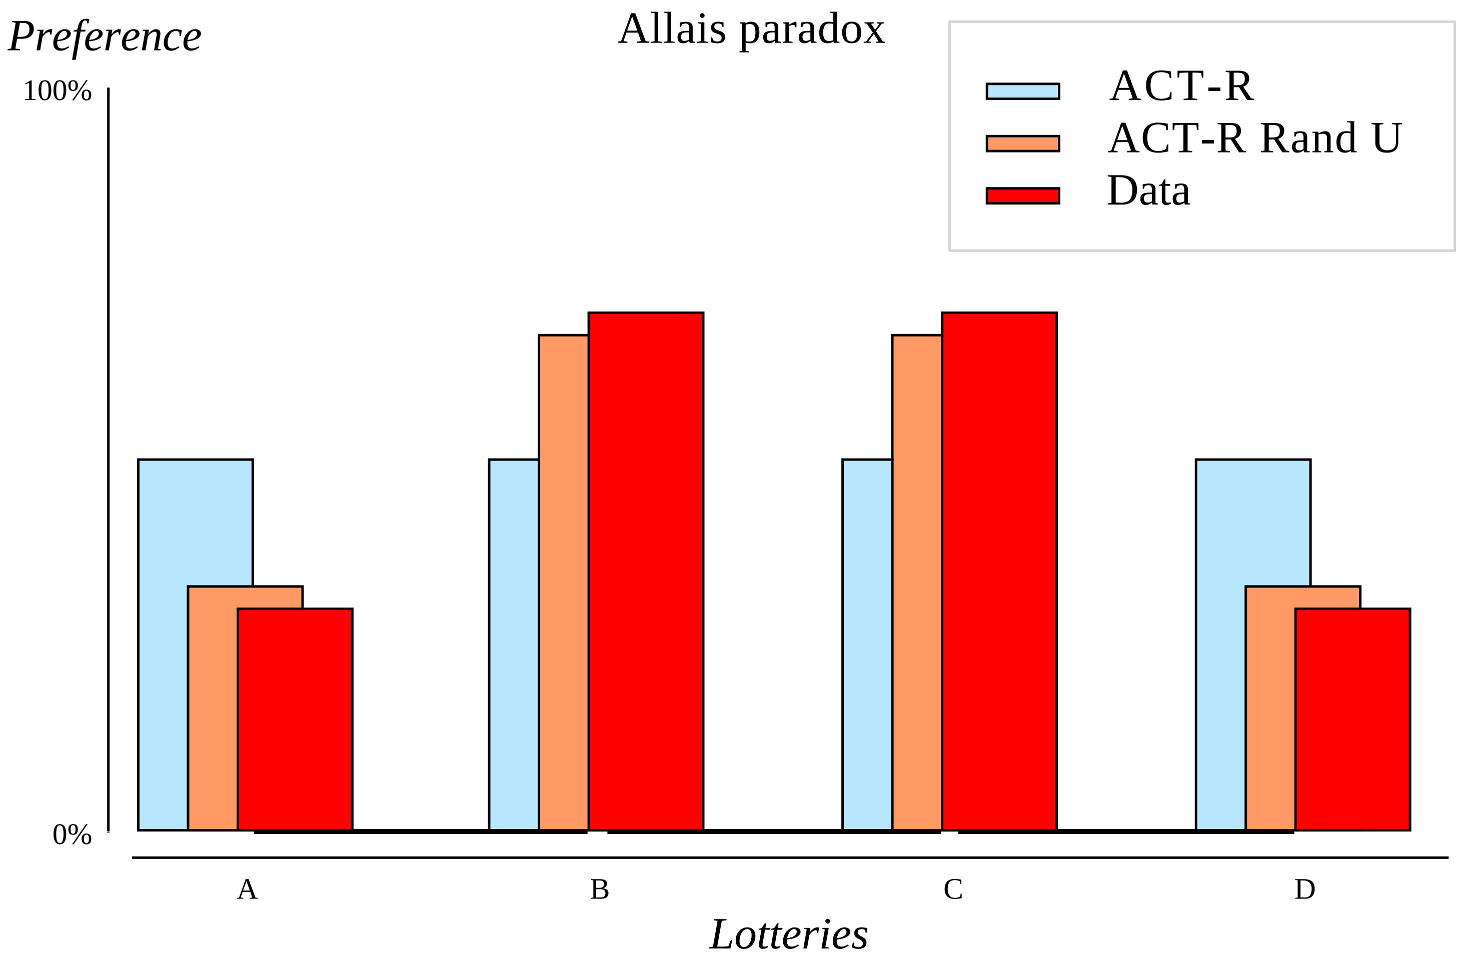

THE ELLSBERG PARADOX

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

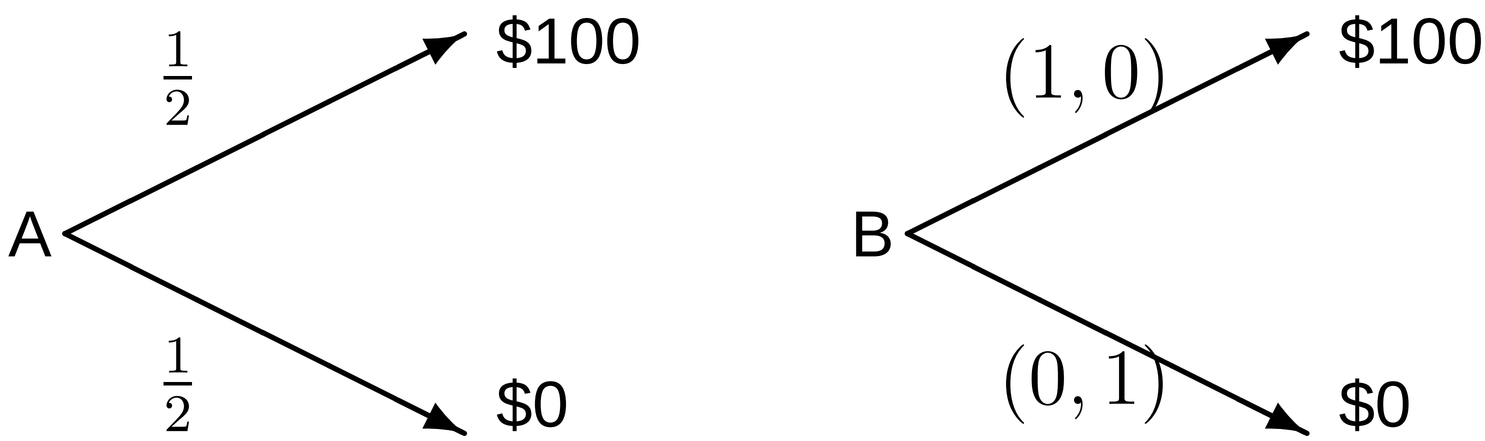

$EU(A) = \$50$ $EU(B) = \$50$

$$A \succ B$$

UNCERTAINTY OF INFORMATION

$$PU^{s} + (1 - P)U^{f}$$

Although the expected utilities are the same, the procedures involved in choosing are clearly different

$$\frac{1}{2} \cdot \$100 + \frac{1}{2} \cdot \$0 \quad \neq \quad \left\{ \begin{array}{c} \frac{1}{100} \cdot \$100 + \frac{99}{100} \cdot \$0 \\ \vdots \\ \frac{99}{100} \cdot \$100 + \frac{1}{100} \cdot \$0 \end{array} \right.$$

Using random utility would involve drawing two samples in lottery B (one for P and one for U) while only one sample is needed for A.
Thus, lottery A may be perceived as closer to the goal and less risky.

CONCLUSIONS

- The Expected utility theory is probably not a good model of the decision–making in the brain.
- Cognitive architectures and ACT–R need to consider the paradoxes arising from the $\max EU$ principle.
- The random utility method has been suggested as a cost–effective solution to the problem.
- The role of uncertainty in decision–making is not well understood (e.g. Ellsberg, 1961).
- Sub–symbolic mechanism may not be the best way to model tasks, where probabilities are given as instructions.

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. (2001, March). The role of emotion in problem solving. In C. Johnson (Ed.), *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing* (pp. 49–57). Heslington, York, England: AISB. (ISBN 1-902956-19-7)

Belavkin, R. V. (2003). *On emotion, learning and uncertainty: A cognitive modelling approach*. PhD Thesis, The University of Nottingham, Nottingham, UK.

Belavkin, R. V. (2005). *Acting irrationally to improve performance in stochastic worlds.* (Submitted to *The Twenty–fifth SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*)

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)

Belavkin, R. V., & Ritter, F. E. (2004). Optimist: A new conflict resolution algorithm for ACT–R. In *Proceedings of the Sixth International Conference on Cognitive Modelling* (pp. 40–45). Mahwah, NJ: Lawrence Erlbaum. (ISBN 0-8058-5426-6)

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

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.

Taatgen, N. A. (2001, July). Production compilation. In *Eights annual post–graduate summer school*. Retrieved from `act-r.psy.cmu.edu`.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.