

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, United Kingdom

Abstract

Human subjects often violate the rational decision-making theory, which is based on the notion of expected utility and axioms of choice (Neumann & Morgenstern, 1944; Savage, 1954). The counterexamples, suggested by Allais (1953) and Ellsberg (1961), deserve special attention because they point at our lack of understanding of how humans make decisions. The paradoxes of decision-making are particularly important for the ACT-R theory which currently relies on expected utility. This paper presents two alternative methods: A random prediction method that uses subsymbolic computations and a method that uses symbolic reasoning for qualitative decision-making. Both methods are tested on ACT-R models of the paradoxes, and the advantages of each method are discussed.

Introduction

Decision-making under uncertainty (i.e. when the outcomes of decisions are not certain) is an extremely important subject in economics, artificial intelligence, cognitive science and psychology. The mathematical theory of choice has been influenced largely by the works of Bernoulli (1738/1954), von Neumann and Morgenstern (1944), Savage (1954) and Anscombe and Aumann (1963). The central idea of this theory is to express preferences of an agent by some *utility function* $u : X \rightarrow \mathbf{R}$ such that

$$x \succ y \iff u(x) > u(y), \quad (1)$$

where $\succ$ is a *strict preference* relation — a binary relation that satisfies among other the antisymmetric and transitive properties: $x \succ y \Rightarrow x \not\succ y$ and $x \succ y \succ z \Rightarrow x \succ z$. Note that strict preference is an idealisation and a mathematical model of real human behaviour (indeed, strict preference of apples over oranges implies one would never choose an orange).

Under uncertainty, one defines probability measures P on some set of prizes Z , and condition (1) is usually replaced by

$$p \succ q \iff \sum_{z \in Z} p(z) u(z) > \sum_{z \in Z} q(z) u(z), \quad (2)$$

where $p, q \in P$ are two probability distributions corresponding to different $x \in X$ (elements of the choice set), and the sums on the right represent the *expected utilities* $E\{u\}$. Although different approaches have been considered in treating sets Z and P (i.e. a set of prizes or acts, objective or subjective probabilities), the theories are usually translated into the following recommendation for a decision-maker

$$\text{Decision}(x) = \arg \max_{x \in X} E\{u\} \quad (3)$$

Thus, a rational agent should make decisions that maximise the expected utility. This principle, known as the $\max E\{u\}$ principle, has been used successfully in many applications for economics, artificial intelligence and cognitive science.

Despite the successes, however, soon after its emergence, the theory of rational decision-making has been strongly criticised by some psychologists and economists. Perhaps, the most obvious problem is that the $\max E\{u\}$ principle fails to suggest the choice when expected utilities of alternatives are equal. This situation is sometimes referred to as the *rational donkey* paradox (i.e. when a donkey is placed between two identical haystacks). Therefore, some additional mechanism must be involved in choosing, such as a roulette wheel. Interestingly, it has been noticed experimentally that human subjects always express some degree of randomness in their choice behaviour (Myers, Fort, Katz, & Suydam, 1963). Cognitive architectures, such as ACT-R (Anderson & Lebiere, 1998), have to use noise in the utility in order to model this ‘imperfect’ property of choice. Moreover, several studies have demonstrated recently that this noisy or ‘irrational’ component of decision-making may, in fact, play an important function optimising the behaviour in stochastic environments (Belavkin & Ritter, 2003).

Another famous and powerful counter example to the $\max E\{u\}$ principle has been suggested by Allais (1953), in which he presented subjects with a choice of different lotteries and asked them which they preferred to play. One example of such a choice of lotteries is described below and shown on Figure 1:

A: 1/3 chance of winning \$300 or 2/3 of not winning anything;

B: A sure win of \$100.

One can easily check that both lotteries have equal expected utilities (\$100 exactly). Thus, there should be no preference according to the $\max E\{u\}$ principle. However, most of the subjects (about 70%) prefer *B* over *A* demonstrating risk-averse behaviour. Interestingly, when the problem is presented with gains replaced by losses (i.e. losing money instead of winning), then the preferences of subjects also revert, and a risk-taking behaviour is observed. Indeed, consider the example below (see Figure 2):

C: 1/3 chance of losing \$300 or 2/3 of not losing anything;

D: A sure loss of \$100.




Figure 1: When choice involves gains, then for majority of subjects $A \prec B$ (risk averse).




Figure 2: When choice involves losses, the preferences revert to $C \succ D$ (risk taking).

Once again, the expected utilities of C and D are equal. However, in many similar studies, the majority of subjects preferred $C \succ D$ (i.e. take the risk).

Even more intriguing, Allais demonstrated that participants often switched their preferences after the probability distributions have been multiplied by some constant factor, which is in violation of the so-called *independence* axiom of the von Neumann and Morgenstern theory. This phenomenon is known as the Allais paradox, and it has been tested by many researchers using not only college students as subjects, but even professional traders (List & Haigh, 2005).

There have been several theories attempting to accommodate the inconsistency of human choice with the rational decision-making theory. One of them is the *prospect* theory due to (Tversky & Kahneman, 1981), which suggests that a function $\pi(p)$ on a probability should be used rather than the probability itself. Such adjustments can explain some violations of the independence axiom, but there is no account for the stochastic nature of choice in the prospect theory. Moreover, replacing probability by a function still suggests that the decision-makers involve some sort of averaging operators (like expected value). There is experimental evidence, however, in favour of the idea that humans often do not use the average.

An important illustration is the paradox, suggested by Ellsberg (1961). One version of this paradox can be explained as follows. You are asked to draw a ball from one of two urns labelled A and B. Each urn contains 100 balls some of which are black and the rest are white. You have been told that urn A contains exactly 50 white balls, but it is not known how many white balls are in urn B. Before you draw a ball, you have to select the colour. If you draw a ball of the same colour you have named, you win \$100. Which urn will you choose to draw the ball from? According to the expected utility, one should be indifferent between A and B. The majority of subjects, however, prefer urn A with known probabilities, even though the ‘average’ probability of each colour in urn B is the same as for urn A (i.e. 1/2) (see Figure 3). Thus, subjects prefer more certain information, and at least in this problem they do not average the odds.

It is clear from the discussion above that the expected utility theory fails to provide a good model of human decision-




Figure 3: A variation on the Ellsberg paradox. Probabilities of lottery B are unknown. The majority prefers $A \succ B$.

making behaviour. This problem is particularly important for cognitive science since many models and underlying architectures rely heavily on the $\max E\{u\}$ principle. In this paper, two alternative methods will be discussed. The next section will be dedicated to the theory of random decision-making that uses Monte-Carlo techniques to generate predictions. This method can explain some of the data associated with the Allais paradox. The following section will consider data from different studies and discuss whether subsymbolic or symbolic mechanisms should be used to model the choices subjects make for lotteries communicated through instruction. A simple symbolic model will be presented. This model will use symbolic representations to encode the lotteries and use qualitative reasoning to make decisions. The issues related to encoding the decision-making problems will also be considered. The paper will conclude by a discussion of the results and the possible directions for future research.

Subsymbolic Model

In this section, subsymbolic mechanisms for decision-making will be considered. This is because ACT-R, the cognitive architecture used by the author to implement the models, employs subsymbolic computations to make decisions (i.e. the conflict resolution). However, this algorithm will be modified in order to accommodate new theories. The method presented here is based on the Markov decision process theory and implements Monte-Carlo technique to make decisions randomly and sample the distributions simultaneously.

Rational and Irrational Components of Choice

One can see from the discussion in previous section that most of the paradoxes occur when the $\max E\{u\}$ principle is used to make decisions. Indeed, the main objective of a rational agent to maximise the utility (1) has been replaced by maximisation of its expected value (2). It is known from the theory of optimal parameter estimation (as well as regression) that the expected value function produces optimal estimator when quadratic error measures a mismatch (the cost). However, the outcomes of many decision-making problems can only be described in terms of successes or failures, and the precise distance to the best solution is not known. When the mismatch is represented by a binary function (i.e. a δ -function), then the optimal estimator is the maximum of probability. For Gaussian distributions, which historically have been considered more often (and due to the maximum entropy principle), the expected value also corresponds to the maximum of probability. More recently the assumption of non-Gaussianity have proven to be very fruitful (e.g. the Independent Components Analysis, non-linear filtering). In general, for non-Gaussian distributions, the expected value $E\{x\}$ does not necessarily identify the most probable $x \in X$.




Figure 4: Two very different distributions with equal $E\{u\}$.

Indeed, both distributions on Figure 4 have the same expected values, but the expected value of a bi-modal distribution on the right corresponds to x between the two maximums. In fact, it is quite possible that $E\{x\} \notin X$. For example, lottery A on Figure 1 presents set of prizes $Z = \{\$0, \$300\}$, and one can easily check that $E\{z\} = \$100 \notin Z$. One can see that decisions made according to $\max E\{u\}$ will be in general different than for $\max P(u)$.

Apart from expected and most probable values, the distributions may have very different other characteristics, such as variances σ^2 , and the $\max E\{u\}$ method does not take them into account. Moreover, the preference relation in (2) assumes that distributions p and q are accurate. The distributions, however, can be seen as subjective approximations of some objective distribution, information about which can only be received through sampling. In this setting, the subjective expected value $E^*\{u\}$ (or $\max P^*(u)$) is only an approximation of the objective, which may be quite different (see Figure 5). If the choice of decisions is made according




Figure 5: Subjective approximation of objective distribution.

to $\max P^*(u)$ (or $\max E^*\{u\}$), then the sampling of objective P will not be very efficient because $\max P$ corresponds to the minimum of information $I = -\log P \sim \frac{1}{P}$ (this is, in fact, the least informative sampling strategy). Samples from parts of distribution other than $\max P$ can allow for a faster updating of the distributions. Therefore, a decision-maker under uncertainty should optimise two conflicting objectives:

1. Maximise information about the utility ($\min P$).
2. Maximise the utility ($\max P$).

These two objectives represent the problem of balancing exploration and exploitation.

A method for simultaneous sampling and decision-making using a random process (i.e. Monte-Carlo style) has been suggested and evaluated using decision-theoretic agents in stochastic environments (Belavkin, 2005). This method uses

random predictions of utilities, which are generated using their prior distributions $P(u)$. This type of sampling is optimal for the maximum of probability: The most probable u corresponds to $\max P(u)$ (can be different from $E\{u\}$). Furthermore, because each sample can be slightly off $\max P$ (i.e. suboptimal or ‘irrational’), the method implements a better sampling strategy. Note, however, that the expected value of these random predictions is $E\{u\}$. Moreover, because they are distributed according to $P(u)$, all other characteristics, such as variance of u , are used.

The performance of agents using the random (i.e. ‘irrational’) method was compared with agents using the classical $\max E\{u\}$ method. These experiments demonstrated that the random agents can significantly outperform the $\max E\{e\}$ when distributions of utility are uncertain and non-Gaussian (Belavkin, 2005). The next section of this paper will present a modification to the ACT-R rule selection mechanism that incorporates the Monte-Carlo technique.

Monte-Carlo Rule Selection for ACT-R

It has been mentioned earlier that the ACT-R cognitive architecture (Anderson & Lebiere, 1998) uses noise in the utility equation to account for the stochastic or ‘irrational’ properties of human choice behaviour. Thus, ACT-R also implements some form of random rule selection mechanism. However, this mechanism assumes Gaussian distribution of utility. Indeed, the choice between several alternative decisions (i.e. rules) in ACT-R is implemented by the subsymbolic conflict resolution mechanism: A rule with the highest utility U_i is selected, where

$$U_i = P_i G - C_i + \text{noise}(\sigma^2) \quad (4)$$

Here, G is called the goal value, P_i is the probability that the goal will be achieved if the rule fires, and C_i is the cost associated with evaluating the rule. Gaussian noise of zero mean and variance σ^2 corrupts the utilities, which allows for modelling many psychological experiments. The noise allows also ACT-R to choose between alternatives with identical prospects (i.e. equal expected utilities). However, the rational component of the utility ($P_i G - C_i$) is based on the expected utility. Indeed, once a rule is selected, there are two possible outcomes: Success (goal achieved) or failure (otherwise). Let U^s be the utility of success and U^f the utility of failure. The probability of success is P_i , and $1 - P_i$ is the probability of failure. The expected utility is

$$E\{U_i\} = P_i U^s + (1 - P_i) U^f = P_i (U^s - U^f) + U^f$$

If we denote $U^s - U^f$ as G (goal value) and U^f as $-C$ (cost), then the above equation will be identical to (4). Note that $U^s = U^f$ in ACT-R terms means that the goal value is zero.

One can see that ACT-R, as many other applications, relies on the $\max E\{u\}$ principle, and, therefore, the conflict resolution mechanism of ACT-R alone cannot model the data from the paradoxes discussed earlier. The use of noise randomises the utility, but the distributions are assumed to be Gaussian, and it cannot capture properties of more unusual distributions, such as used in lotteries on Figures 1 and 3. A modification of the conflict resolution using Monte-Carlo technique has been implemented in the following way.

The ACT-R architecture considers only two possible outcomes of each decision: Success or Failure. The statistics about these events is stored for each rule i in its probability of success P_i . This probability can be used to generate random predictions about the outcomes using the inverse probability distribution function (PDF) method:

$$\text{Success} \vee \text{Failure} = F^{-1}(p), \quad p \in [0, 1],$$

where F is the PDF (cumulative) for P_i , and p is a uniform random number on $[0, 1]$. The utilities of these two outcomes in ACT-R notation are

$$G - C_i \vee -C_i$$

That is, the system expects either to pay the cost and gain some goal value, in case of success, or just to pay the cost and gain nothing, in case of failure. These random predictions are generated for each rule, and the rule with the maximum random prediction is selected. In addition, the costs can be randomised using Gamma noise, as described by Belavkin and Ritter (2004). This should exclude the possibility of two rules in the conflict set having equal utilities.

The method, described above, can be used to explain some of the results associated with the Allais paradox. Indeed, Figure 6 shows PDFs for lotteries A and B on Figure 1. Note that




Figure 6: Distribution functions for lotteries on Figure 1

for lottery A, we should win nothing two out of three times, while in lottery B we always win £100. Thus, according to the random utility method, two out of three times the random utility of lottery A will be smaller than that of lottery B (Figure 7). This confirms that approximately 70% of subjects prefer $B \succ A$.




Figure 7: Random utilities for Lotteries A and B. There is a 2/3 chance that $B \succ A$.

One can easily verify in a similar manner that for lotteries C and D the reverse preference holds $C \succ D$. Indeed, two out of three times the random utility of lottery C will be greater than that of lottery D.

The random utility method using inverse PDF has been implemented as an overlay for the ACT-R version 5. A simple model of a two choice task has been implemented to test the

subsymbolic mechanism. In this model, a choice of two lotteries is represented by a chunk that has two alternatives as slot values. Two rules are competing for the goal each selecting a different alternative. The information about the gains and probabilities of each lottery is encoded in the subsymbolic form for each rule. For example, for lottery A and B (Figure 1), the two rules has the following settings:

$$\begin{aligned} P_A &= 1/3, & C_A &= \$0, & P_A G - C_A &= \$100 \\ P_B &= 1, & C_B &= \$200, & P_B G - C_B &= \$100 \end{aligned}$$

For lottery C and D (Figure 2), the following settings were used

$$\begin{aligned} P_C &= 2/3, & C_C &= \$300 & P_C G - C_C &= -\$100 \\ P_D &= 0, & C_D &= \$100 & P_D G - C_D &= -\$100 \end{aligned}$$

In both cases, the goal value was set to \$300 representing the difference between the utilities of success and failure. One can see from above that in both cases the standard expected utilities ($PG - C$) of ACT-R of conflicting rules are equal, and, therefore, the model should express no preference. When the same model is run using the random utility method, the preferences should be $A \prec B$ and $C \succ D$. The results of both models are shown on Figure 8 confirming the prediction.




Figure 8: Preferences for lotteries A, B, C and D of the ACT-R model using the $PG - C$ and random methods for conflict resolution.

One can see that the random utility method changes the performance of the model quite dramatically. It is more sensitive to the characteristics of the utility distributions, which can be quite different despite having equal expected values. The random utility method is also more dynamic: Unlike the noise in ACT-R, which has static variance σ^2 , the utility distributions are constantly changing as a result of learning, and controlling the uncertainty allowing for a better adaptation of the decision-making mechanism. However, there are reasons for which the subsymbolic model, described above, cannot be considered as the model of the Allais paradox. These reasons will be explained in the next section, in which another model based on symbolic computation will be presented.

Symbolic Model

One should question the validity of using subsymbolic mechanisms to model the task of choosing between lotteries, such as on Figure 1. The main purpose of subsymbolic computations is to capture the statistical properties of the learning

mechanisms in the brain. This statistical (or Bayesian) learning assumes some sampling procedures (e.g. trial and error) of distributions that evolve over time. Quite clearly, the tasks in the lottery problem are very different. Indeed, the information about the odds and monetary gains is communicated through instruction in a symbolic form, not through trial and error. Moreover, as has been discussed earlier and tested elsewhere (Belavkin, 2005), the main advantage of the random utility method is because it implements a better sampling strategy of distributions, which are uncertain. In the lottery task, on the other hand, the distributions are given to the participants, and they cannot perform several trials.

The Effect of the Probability

Furthermore, the inverse PDF method suggests that the proportion of people that chooses a lottery with higher monetary gain should depend only on the odds of this lottery to be successful, and not on the amount of the gain itself. Indeed, by looking at Figure 7, one can see that the results will be the same, if lottery A offers a gain of \$200 or \$300. The only parameter that matters there is $P_A = 1/3$ — the probability of success in A.

The analysis of several studies on choice between lotteries fails to support the idea of such a clear dependence between the probability and the proportion of people opting for a risky lottery. Tversky and Kahneman report results of several experiments involving choices of lotteries similar to A and B. In one example, the probability of winning \$600 (against a sure win of \$200) is $1/3$. They report 28% of participants opting for this risky option. In another experiment, with $P_A = 1/4$ (win \$1000 against a sure win of \$240) only 16% chose A. However, List and Haigh report results of their experiment with identical parameters, but with 38% of subjects choosing the risky option. The comparison of data from other similar experiments, where the expected utilities of both alternatives are equal and one option having a sure win, indicates against strong dependence between the probability and the proportion of people choosing the risky option. This may be because subsymbolic mechanisms are not playing a significant role in such tasks, and that symbolic reasoning may yield a better model. Note that the average of those choosing lottery A calculated from several studies is 27%.

To address these issues, a symbolic model has been implemented using the ACT-R production system. This model implements simple logical strategies and implements qualitative analysis of the decision-making problem. The model is described below.

The Logic of Choice

The symbolic model of the Allais paradox uses chunks and production rules of ACT-R to encode symbolic representations of lotteries and reasoning required to make choice. The reasoning implements qualitative analysis of quantities, such as monetary gains and probabilities of the alternatives, and the decisions are made according to the following preference relations:

$\sim$ indifference (any can be chosen)

$\succ$ preference (e.g. if $x \succ y$, then x is chosen)

If more than one pair of relations is considered, then a union is assumed (i.e. logical or). This can be proven using the following equality

$$\succ \wedge \prec = \sim = \succ \vee \prec$$

In probabilistic terms, $\sim$ means that there is equal chance of any of the two objects to be chosen. The union operation also can be used to estimate the probability when more than one preference relation holds. For example, $\sim$ or $\succ$ yields probability $3/4$ that the first object will be chosen. This disjunctive logic allows for a very simple production system using simple rules that compare the properties in parallel and assign the preference relations. Any of these rules can fire, and, therefore, the probability that an object is selected depends on how many of its properties are preferred.

The choice of two objects (e.g. the lotteries) is represented in the model by a chunk of special type `choice` that has pairs of slots for each property that can be different. For example, slots `name1` and `name2` hold the names of each alternative lottery:

```
(ab ISA choice name1 A name2 B ...)
```

All rules in the model can apply until slot `chosen` of the goal chunk is not `nil` and holds the name of a chosen object.

Each property of the two objects is compared, and if the value of one object is preferred to a value of another, then this object is preferred. If the values are equal, then indifference holds. Numerical values can be compared either directly (e.g. if $x > y$, then $x \succ y$) or using a set of categories such as `{large, medium, small}`. The issue of converting quantitative information into qualitative will be discussed later.

Encoding the Lotteries

One can see from Figures 1 or 2 that each lottery, apart from its name, can be described by the following four properties: Gain (U^s), the probability of gain (P^s), loss (U^f) and the probability of loss (P^f). A preference relation is inferred by production rules comparing any of these four properties. For example, if $U_A^s > U_B^s$, then $A \succ B$. If the values are equal, then indifference holds, and any lottery can be selected.

It is easy to see that if one lottery has a number of advantages (i.e. several properties are preferred), then there is a higher chance that it will be selected, because there are more rules in the conflict set that will choose it. The preference is less obvious when different properties have different preferences. For example, one lottery may have higher gain than another, but with lower probability, such as on Figure 1. Recall that these examples were designed in such a way that they have equal expected utilities. However, qualitative analysis can yield different results. The table below shows values of each property for lotteries A, B, C and D with the preference relations assigned to each pair:

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

The bottom row shows cumulative preferences between the lotteries that are computed by a union of preferences on all properties. One can see that $A \prec B$ and $C \succ B$. Moreover, it is also possible to compute the probability of these preferences. For example, A is preferred to B with probability $3/8 \approx 38\%$, which is confirmed by the ACT-R model. Interestingly, this is exactly the proportion that List and Haigh report in their study. Of course, this may be just a coincidence, and, as has been mentioned earlier, other authors report slightly different results (e.g. 28% and 16% in Tversky & Kahneman, 1981). However, the symbolic model may explain why there is no clear dependence between the probability of a gain in one lottery and the proportion of people that choose it: The number of properties that are different may play a greater role, because each pair of properties is evaluated separately.

Another important issue to consider is the way the values of properties are encoded. Indeed, the discussion above uses numerical values. However, one may speculate that the numbers have to be converted into categories, such as {large, medium, small}. This creates additional ambiguity of how, for example, small probabilities 0,01 and 0,02 are categorised: As both small and equal or one greater than another? This ambiguity may explain the violations of the independence axiom noticed by Allais.

Finally, the symbolic model suggests a very elegant explanation of the Ellsberg paradox. Indeed, one can see from Figure 3 that both lotteries have identical gains and losses. The only different property is how certain is the information: The probabilities of lottery B are not known. If a decision-maker prefers certainty, then $A \succ B$ follows.

Discussion

Decision-making is a complex and very important process that is involved in almost every other aspect of cognition. This is why the study of the paradoxes unexplained by the theory is so important for cognitive science. These paradoxes have been mainly associated with the expected utility theory, pioneered by Bernoulli and developed by von Neumann and Morgenstern. In this paper, two other theories have been considered: The random (or irrational) and the qualitative decision-making methods. The former is based on the theory of Markov-decision processes and Monte-Carlo statistical estimation. This theory can bear significant advantages over the $\max E\{u\}$ method in stochastic environments with unknown and non-Gaussian utility. The method can be used as a model of subsymbolic learning processes in the brain, that is when learning involves accumulation of some statistics through trial and error (i.e. unsupervised).

The latter theory is based on qualitative analysis and Boolean logic. Qualitative decision-making allows one to make decisions using logical inferences without any quantitative computations. This method can be used to model tasks where information is encoded symbolically and communicated through instruction.

It has been shown both analytically and using models that these methods enable one to simulate and explain decision-making data that do not follow the expected utility theory. However, there is still a long way before we fully understand how these different mechanisms interact and work together.

ACT-R is a cognitive architecture that can use both types of computations. This is why it has been chosen to build and test the models. However, the current algorithms in ACT-R use the traditional theory that is prone to paradoxes. The new conflict resolution algorithm is publicly available on the author's webpage. The methods, outlined in this paper, have been tested on simple models of the famous paradoxes of decision-making, but a wider testing on a variety of models is desirable. The author believes this is the way towards a better theory of human decision-making without the paradoxes.

References

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'École américaine. *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.
- Belavkin, R. V., & Ritter, F. E. (2003, April). The use of entropy for analysis and control of cognitive models. In F. De-tje, D. Dörner, & H. Schaub (Eds.), *Proceedings of the Fifth International Conference on Cognitive Modelling* (pp. 21–26). Bamberg, Germany: Universitäts-Verlag Bamberg.
- 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.
- 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.