J Risk Uncertainty (2006) 33:155–164
DOI 10.1007/s11166-006-0390-3

Learning in the Allais paradox

Gijs van de Kuilen · Peter P. Wakker*

© Springer Science + Business Media, LLC 2006

Abstract Whereas both the Allais paradox, the first empirical challenge of the classical rationality assumptions, and learning have been the focus of many experimental investigations, no experimental study exists today into learning in the pure context of the Allais paradox. This paper presents such a study. We find that choices converge to expected utility maximization if subjects are given the opportunity to learn by both thought and experience, but less so when they learn by thought only. To the extent that genuine preferences should be measured with proper learning and incentives, our study gives the first pure demonstration that irrationalities such as in the Allais paradox are less pronounced than often thought.

Keywords Learning · Rational choice · Allais paradox · Nonexpected utility

JEL Classification D81 · D83

The first empirical challenge of the classical rationality assumption in economics was made by Allais (1953), through his famous choice paradox for decision under risk. It showed that the majority of people systematically violate expected utility, the model commonly accepted as the hallmark of rationality for choice under risk in those days. Allais' paradox initiated a flurry of empirical demonstrations that more and more basic rationality assumptions are violated empirically (Camerer, 1995). To model these findings, theories were developed that explicitly depart from rational choice and only pretend to be descriptive, such as prospect theory (Kahneman and Tversky, 1979). The surprising achievement of these theories about irrational behavior was that, contrary to earlier beliefs, they could maintain enough regularity to be analytically tractable and to yield empirical predictions.

*Paul Anand and Stefan Trautmann made helpful comments.

G. van de Kuilen · P. P. Wakker (✉)
CREED, Department of Economics, University of Amsterdam, The Netherlands & Econometric Institute, Erasmus University, Rotterdam, The Netherlands
e-mail: Wakker@FEW.EUR.NL

A reaction to this literature, and a defense of the classical economic assumption of rationality, came from experimental economics. The original empirical evidence against expected utility can be challenged on several grounds. Most of the original evidence came from one-shot decision making experiments. It is likely that subjects never faced the considered decisions before, and their acts may have been based on simple misunderstandings rather than on irrationalities in genuine preference. The Allais paradox may arise, for example, because of misperception and unfamiliarity with probability, an effect that will be reduced or will even disappear after learning. In addition, the early experiments often used hypothetical choice, so that subjects may not have been well motivated to reveal their true preferences.

The standards of experiments have been greatly improved by the modern principles of experimental economics, with an emphasis on well-controlled experiments (Smith, 1982). Thus, Plott's discovered preference hypothesis asserts that "individuals have a consistent set of preferences over states, but such preferences only become known to the individual with thought and experience" (Myagkov and Plott, 1997, p. 821). Binmore (1999) took an extreme position and compared testing an economic theory in a laboratory environment without sufficient learning opportunities, without adequate incentives, and without feedback, with performing a chemistry experiment using dirty test tubes. The importance of proper learning and incentives, such as occurring in markets, to avoid irrationalities due to basic misunderstandings and lack of motivation, has been demonstrated by many studies in many domains (Brookshire and Coursey, 1987; Camerer and Ho, 1999; Cox and Grether, 1996; Erev and Roth, 1998; Evans, 1997; List, 2004; Loomes, Starmer and Sugden, 2003; Myagkov and Plott, 1997).

Whereas both the Allais paradox and learning have, thus, been the focus of numerous experimental investigations (see also Camerer, 2003; Conlisk, 1989; Einav, 2005; Hertwig and Ortmann, 2001, Section 3; Offerman, Potters, and Sonnemans, 2002; Starmer, 2000), there exists no experimental study into learning in the context of purely the Allais paradox today. Given the historical importance of the Allais paradox, and its continued importance as a benchmark for risk models today, this absence is remarkable. There have been some studies of the Allais paradox with repeated choice and learning, described in Section 3, but these always involved other more complex phenomena. Hence, a pure test of the individual rationality tested by the Allais paradox has not yet been obtained under learning. This paper presents such a pure test.

Myagkov and Plott (1997) distinguished between learning through mere thought and learning through thought and experience. To separate the effects of learning and of experience, we gave feedback to our subjects in one treatment, so that learning is generated by both thought and experience, and no feedback in another treatment, so that learning results merely from thought and not from experience. Other than that, our setup was chosen to be as simple and transparent as possible, so as to avoid all strategic interactions, framing effects, income effects, or other effects beyond (violations of) independence. Cubitt, Starmer, and Sugden (2001) recommended the use of setups as simple as possible to test individual rationality (p. 391). We will find that thought and experience together generate a drift towards expected utility, but thought alone does not. Myagkov and Plott (1997) predicted a drift to rationality:

> Thus, when individuals are first given questions, they are characterized by a type of confusion. As they begin to formulate decisions in this state they are influenced by "frames" in much the way that prospect theory asserts. As an understanding of the context evolves, the manifestation of the underlying preferences becomes more clearly observable in the data and decisions approach those predicted by the classical theory of choice and preference. (p. 821)

In general, to what extent the initial preferences with biases are important and to what extent those that result after learning, will depend on context and application and is a topic for future investigations. The importance of this topic is agreed upon by experimental economists and behavioral economists alike. For example, Tversky and Kahneman (1986) wrote:

> Indeed, incentives sometimes improve the quality of decisions, experienced decision makers often do better than novices, and the forces of arbitrage and competition can nullify some effects of error and illusion. Whether these factors ensure rational choices in any particular situation is an empirical issue, to be settled by observation, not by supposition. (p. S273)

1 The common ratio effect

We will focus on the *common ratio* version of the Allais paradox. Consider a pair of prospects $S = (p{:}s)$ and $R = (0.8p{:}r)$ with $0 \leq p \leq 1$. Prospect S (safe) yields a prize of €s with probability p and nothing otherwise and prospect R (risky) yields a prize of €r with probability $0.8p$ and nothing otherwise. If a decision maker faces a choice between S and R, *expected utility* entails that the decision maker applies the following decision rule:

$$S \succcurlyeq R \Leftrightarrow p\mathrm{U}(s) \geq 0.8p\mathrm{U}(r) \Leftrightarrow \mathrm{U}(s) \geq 0.8\mathrm{U}(r),$$

where $\mathrm{U}(\cdot)$ is the *utility function* (with $\mathrm{U}(0) = 0$) and $\succcurlyeq$ denotes a weak preference for prospect S over prospect R. Thus, expected utility theory predicts that preferences between prospect pairs of this type are independent of p. However, numerous laboratory experiments have shown that individual preferences are affected by the value of p (Camerer, 1995). In particular, preferences often switch from S to R if p changes from 1 to an intermediate value. This phenomenon contradicts expected utility and is known as the *common ratio effect*.

The common ratio effect is accommodated, for instance, by prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). If decision makers are asked to choose between prospects S and R, prospect theory assumes the following decision rule:

$$S \succcurlyeq R \Leftrightarrow w(p)\mathrm{U}(s) \geq w(0.8p)\mathrm{U}(r),$$

where $w(\cdot)$ is a probability transformation function that transforms the objective probabilities into subjective decision weights. This probability transformation function is usually assumed to be inverse S-shaped, which implies that individuals overweight extreme outcomes relative to moderate outcomes (Tversky and Kahneman, 1992). Thus, for $p = 1$ prospect S is more attractive relative to prospect R than for intermediate p, explaining the common ratio effect. Other popular theories of decision under risk that can explain the common ratio effect include rank-dependent utility (Quiggin, 1981) and regret theory (Loomes and Sugden, 1982).

2 Existing evidence on the Allais paradox with repeated choice

The findings presently available in the literature do not give clear suggestions about what will happen in a pure individual learning experiment for the Allais paradox, the topic of our study. In MacCrimmon (1968) and Slovic and Tversky (1974), subjects were first asked to choose between Allais paradox-type prospects. Then the subjects were given prepared normative arguments for and against the independence axiom and were asked to reconsider their choices. This did not lead individual preferences to adhere to the independence axiom of expected utility. In these experiments, subjects made decisions without real incentives though, and opportunities to learn from experience did not exist.

Bone, Hey and Suckling (1999) also provided evidence against expected utility. In their experiment, subjects were asked to register three common-ratio type choices without feedback and were found to violate expected utility more often in the third round than in the first round. However, in the second round, subjects were allowed to communicate and were asked to make joint decisions with another participant, so that effects of social interaction came into play.

Results from an individual choice experiment by Keren and Wagenaar (1987) showed, on the other hand, that individual choice behavior tends to converge to expected value maximization (i.e., expected utility with a linear utility function) if decision makers are asked which prospect from a common ratio pair they prefer to play 10 times in a row. Here subjects made hypothetical choices and individual learning opportunities were absent. Under repeated decisions with repeated (hypothetical) payments, income effects play a role and expected value is dictated merely by stochastic dominance and the law of large numbers, which confounds the test of the Allais paradox or independence. To avoid these confounds, the random lottery incentive system is commonly used in experiments on individual choice under risk today, and we will use it too.

Barron and Erev (2003) asked some participants to choose between prospects $S =$ (1:\$0.0075) and $R =$ (0.8:\$0.01), while others were asked to choose between prospects $S =$ (0.25:\$0.0075) and $R =$ (0.2:\$0.01), for 400 rounds. After each decision, participants received the prize of the chosen prospect. Results indicated a tendency for preferences to converge towards expected value maximization over time. Subjects were not informed about the true probabilities, so that attitudes towards unknown probabilities came in. In addition, there were repeated payments, so that income effects and the law of large numbers distorted the test of independence as they did in Keren and Wagenaar (1987). Remarkable in this study is how much expected value, the obvious model for a sum total of 400 independent choices, was violated.

Humphrey (2006) investigated the effect of learning on violations of coalescing, monotonicity, and a common-consequence effect that was more complex than the Allais paradox (with no certain option available). Choices were not elicited directly, but were derived indirectly from certainty-equivalent matching questions. In the learning treatment, subjects were shown 10 drawings from the relevant prospects before deciding. These drawings were manipulated so as to be representative. As in our treatment with feedback below, subjects could infer about the probability distribution this way, but in Humphrey's experiment no outcomes were associated with the learning events so that they did not provide experience in this sense. Mixed results were found, with some violations reduced but others enlarged.

Among others, Carbone and Hey (2000), Harless and Camerer (1994), Loomes, Moffat, and Sugden (2002), and Schmidt and Neugebauer (2006) experimentally analyzed the role of errors in decision making. Typically, subjects were first asked to register their preferences between a large number of different prospect pairs two or three times. A random mechanism then determined the prize of the randomly selected prospect to be played out for real at the end of the experiment. Thus, subjects did not receive feedback from their earlier decisions and only made a particular pairwise choice two or three times, limiting the opportunity to learn. These studies mostly seemed to support expected utility theory plus a particular error term.

<table>
<tr><th></th><th>1</th><th>2</th><th>3</th><th>4</th><th>5</th><th>6</th><th>7</th><th>8</th><th>9</th><th>10</th><th>11</th><th>12</th><th>13</th><th>14</th><th>15</th><th>16</th><th>17</th><th>18</th><th>19</th><th>20</th></tr>
<tr><td>Prospect A</td><td colspan="20">€s</td></tr>
<tr><td>Prospect B</td><td colspan="16">€r</td><td colspan="4">€0</td></tr>
</table>

<table>
<tr><th></th><th>1</th><th>2</th><th>3</th><th>4</th><th>5</th><th>6</th><th>7</th><th>8</th><th>9</th><th>10</th><th>11</th><th>12</th><th>13</th><th>14</th><th>15</th><th>16</th><th>17</th><th>18</th><th>19</th><th>20</th></tr>
<tr><td>Prospect C</td><td colspan="5">€s</td><td colspan="15">€0</td></tr>
<tr><td>Prospect D</td><td colspan="16">€0</td><td colspan="4">€r</td></tr>
</table>

Fig. 1 The presentation of the prospect pairs

3 Experimental design

3.1 Subjects

Our experiment was conducted at the Hogeschool of Amsterdam. Participants consisted of undergraduate and postgraduate students from a wide range of disciplines. We approached the students ourselves in a cantina, and asked if they wanted to participate in an individual decision making experiment. In total, 52 subjects took part, divided equally over two treatments as well as over both sexes. The experiment was purely individual and subjects made choices under the direct supervision of the experimenter.

3.2 Stimuli

Each subject was asked to choose between two pairs of prospects $S = (p{:}s)$ and $R = (0.8p{:}r)$ in fifteen rounds. In one prospect pair, which we will call *non-reduced*, p was 1. In the other, *reduced*, prospect pair, p was 0.25. Thus, the subject registered one choice between reduced prospects S and R and one choice between non-reduced prospects S and R in each round. The prospect pairs were presented in a *matrix display*, as shown in Figure 1.

All prospects yielded a possible prize, depending on a roll with a twenty-sided die. For example, prospect C in Figure 1 yielded a prize of €s if the roll of a twenty-sided die was 1, 2, 3, 4, or 5 and nothing otherwise. We decided to use the above juxtaposition of prizes with the r-payment in D disjoint from the s-payment in C. There is evidence that, primarily due to regret effects, this juxtaposition yields more violations of expected utility than other juxtapositions (Loomes, 1988; Starmer and Sugden, 1989; Harless, 1992).

The probabilities of winning in each prospect were the same over all rounds. The prizes of prospects S and R varied across the rounds. Table 1 displays the fifteen combinations of s and r that were used.

The reduced and non-reduced prospect pairs as well as the prospects themselves were presented in random order between subjects to avoid order effects.

3.3 Procedure

At the beginning of the experiment, each subject received instructions (see Appendix) and was first asked to register preferences between four non-common ratio prospect pairs that served to familiarize them with the procedure. Then they were asked to register their preferences between the fifteen common-ratio prospect pairs. To determine the final payoff of participants,

Table 1 The fifteen (s, r) combinations used

s	6.00	6.25	6.50	7.00	7.50	8.00	8.50	9.00	9.50	10.00	10.50	11.00	11.50	12.00	12.50
r	7.75	8.50	9.00	9.25	9.50	11.00	11.25	11.50	13.00	13.50	14.00	15.00	15.50	15.75	16.00

a random lottery incentive system was used. That is, at the end of the experiment, one choice pair was selected at random and the prize of the selected prospect was paid out in cash. On average, subjects earned about €6.50 while the total experiment lasted approximately fifteen minutes.

The amount of feedback that participants received differed between the two treatments. Under the treatment *with feedback*, subjects were asked to roll a twenty-sided die to determine the prize of the chosen prospect after each choice. Even if subjects indicated preference for a prospect that yielded a particular prize for certain we asked them to roll the twenty-sided die, although some did not seem to pay much attention to the result. They were then asked to write this prize down on their decision sheet, before continuing to the next choice. Thus, after each decision, subjects immediately received feedback and, therefore, directly experienced what the resolution of uncertainty was. Under the treatment *without feedback*, participants only received feedback at the end of the experiment when the real payoff, resulting from the randomly chosen choice situation, of the particular participant had to be determined. They, therefore, did not experience what the consequences of other decisions were.

4 Results

Figure 2 depicts individual choice behavior under the two treatments over time. The difference in violations of expected utility between the two treatments in the first round reflects

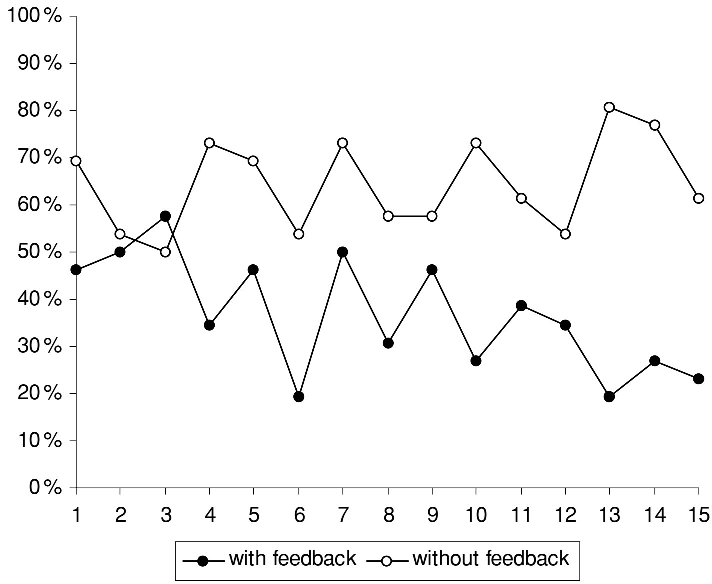

Fig. 2 Percentage of expected utility violations over rounds

Table 2 Estimates of regression

β_0	β_1	β_F	β_{NF}
0.597 (0.053)	−0.087 (0.074)	−0.018* (0.006)	0.006 (0.006)

Notes: Standard errors in parentheses.
*Significant at the 5% level.

between-subjects randomness ($t_{50} = 1.70$, $p = 0.10$). The following analyses will concern within-subject differences.

We first consider the individual choice behavior under the treatment with feedback. Here, 46.15% of the participants violated expected utility theory in the first round of the experiment. These subjects either preferred prospect *S* over *R* in the non-reduced prospect pair combined with a preference for prospect *R* over *S* in the reduced prospect pair, or vice versa. In the final round of the experiment only 23.08% of the individual choices violated the independence axiom. This increase in expected utility consistency is significant, $t_{25} = 1.81$, $p = 0.04$. Moreover, in the first round 26.9% of the participants maximized expected value, which is significantly lower than 58.0% in the final round of the experiment, $t_{25} = -3.33, p = 0.001$. Finally, there was a significant negative linear trend in the percentage of expected utility violations over rounds, $t_{25} = 3.15, p = 0.005$. Thus, individual choice behavior converged to the descriptive predictions of expected utility under the treatment with feedback.

We next consider the treatment without feedback. Here, 69.2% of the individual choices violated the predictions of expected utility theory in the first round of the experiment. In the fifteenth round, the amount of expected utility violations dropped insignificantly to 61.5%, $t_{25} = 0.46$, $p = 0.3232$. The expected value violation rate in the first round also did not differ significantly from that in the final round, $t_{25} = -0.30$, $p = 0.38$. We neither found significant evidence for the existence of a negative linear trend in the amount of expected utility violations over rounds, $t_{25} = 0.98$, $p = 0.34$. Thus, convergence of individual behavior to expected utility predictions was not found under the treatment without feedback.

Without feedback, the average violation rate over all subjects and over all rounds ($M = 0.64$, $SD = 0.04$) was higher than that of subjects with feedback ($M = 0.36$, $SD = 0.05$). Analysis of variance indicates that this difference is significant, $F(1, 50) = 19.65, MSE = 0.05$.

We also estimated the following simple regression model using OLS:

$$\pi = \beta_0 + \beta_1 D + \beta_F\, D\, Round + \beta_{NF}(1 - D) Round + \varepsilon,$$

where π is the percentage of expected utility violations, D is a treatment dummy which equals 1 for the with-feedback treatment and 0 for the without-feedback treatment, and *Round* is the number of the round. Table 2 gives the results.

Only β_F is statistically significant, which is consistent with our finding that convergence to rationality over rounds is only present in the with-feedback treatment. Here, the percentage of expected utility violations is estimated to drop by 1.8% per round.

5 Discussion and conclusion

Holt (1986) formulated a potential theoretical problem for the random-lottery incentive system, if subjects interpret this system as one grand overall lottery. Subsequent studies showed that this problem does not occur empirically (Cubitt, Starmer, and Sugden, 1998; Starmer and Sugden, 1991). This real-incentive system has become the almost exclusively used one in individual choice experiments today. Its main features are that it avoids income and house money effects.

The violations of expected utility found in the first round agree with the common findings in the field (Camerer, 1995; Starmer, 2000). If subjects were given the opportunity to learn by both thought and experience, the number of expected utility violations dropped significantly over time. Subjects seemed to learn to maximize expected value, which is in line with the findings of Keren and Wagenaar (1987) and Barron and Erev (2003). A possible explanation is that probability transformation is reduced because of learning. With repetition and feedback, decision makers learn not only about the prize of the chosen prospects, but also about the prize that the non-chosen prospects would have yielded. When a decision maker prefers prospect *S* in the non-reduced prospect pair because of subjective probability distortion, he experiences that the possible prize of prospect *R* was higher 80% of the time. This could induce the decision maker to assess probabilities better, decreasing the amount of expected utility violations over rounds.

Results from the treatment without feedback support the above explanation. Under this treatment, convergence of individual preferences to the descriptive predictions of expected utility was not found. Clearly, subjects are unlikely to learn to assess probabilities better if they are not able to learn about the prize of the chosen prospects or the prize of the non-chosen prospects. This was predicted by Cubitt, Starmer, and Sugden (2001): "...what is repeated must include not only the act of decision, but also the resolution of any uncertainty and the experience of the resulting outcome" (pp. 393–394).

This paper has given a pure experimental demonstration that learning can reduce violations of expected utility. Our experiment avoided distortions due to other factors beyond individual risk attitude. Thus, to the extent that genuine preferences can be revealed only after proper learning and with proper real incentives, this paper gives support for a better descriptive validity of expected utility than suggested by earlier experimental studies of pure individual decisions under risk.

Appendix A. Experimental instructions

The following experimental instructions have been translated from Dutch into English.

Welcome to this experiment. If you have any questions while reading these instructions, feel free to ask the assistant of this experiment. The experiment consists of 2 practice rounds followed by 15 real rounds. Every round consists of 2 parts. In each part of each round you can earn a prize (in euros). At the end of the experiment you will randomly select 1 of the 15 real rounds by rolling a 20-sided die (in case you then roll a 16, 17, 18, 19 or 20, we will ask you to re-roll the die). Thereafter you will randomly select the first (the roll is even) or the second (the roll is odd) part of this real round by again rolling a die. *Only the prize of the selected part of the selected real round will be paid to you.* The prizes of the lotteries in this experiment range from €0 to €16. It is thus possible that by rolling the die at the end of the experiment, you select a lottery with a prize of €0, and thus no euros will be paid to you. It is also possible that by rolling the die at the end of the experiment, you will select a lottery with a prize of €16, which will then be paid out to you. On average, the prize per participant is about €6. At the beginning of each round you receive a sheet on which you can write down your decisions. We will now explain the filling in of such a sheet on the basis of the example-sheet that has already been handed out to you.

First, you see the number of the current round at the top of each decision sheet. In Part 1 of each round, we ask you to make a choice between two lotteries, named Lottery A and

Lottery B. Both lotteries yield a particular prize that depends on the result of your roll with a die. The rolling of this die using a cup takes place after you have made a choice between both lotteries. If you choose Lottery A on the example-sheet and the result of the roll of the 20-sided die is between 1 and 4, the prize of Lottery A is equal to €4. However, if the result of the roll of the die is between 5 and 20, the prize of Lottery A is equal to €0. If you choose Lottery B on the example sheet and the result of the roll of the 20-sided die is between 1 and 12, the prize of Lottery B is equal to €0. However, if the result of the roll of the die is between 13 and 20, the prize of Lottery B is equal to €3. After you have made a choice between Lottery A and Lottery B by encircling either A or B on the decision sheet, we ask you to roll the 20-sided die using the cup once and encircle the result of the roll on the decision sheet. As mentioned before, the result of the roll determines the prize of the lottery that you have chosen. After you have written down this prize on the decision sheet, Part 1 has ended and Part 2 begins.

The second part of each round is almost identical to the first part. We again first ask you to choose between two lotteries, this time named Lottery C and D, by encircling either C or D on your decision sheet. Then we ask you to roll the 20-sided die using the cup once and encircle the result of the roll on the decision sheet. The result of the roll again determines the price of the lottery that you have chosen, which you will write down on your decision sheet. After filling in this price, the next round begins.

There are no right or wrong answers during this experiment. It exclusively concerns your own preferences. In those we are interested. At each part of each round it is best to encircle the lottery that you prefer. Surely, that part of that round can be selected at the end of the experiment, and then you will get the prize of the lottery you have encircled. It is therefore best for you to encircle your preferred lottery in each part. If you have no questions at this point, the first of the 2 practice rounds will start now. Good luck!

References

Allais, Maurice. (1953). "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Américaine," *Econometrica* 21, 503–546.

Barron, Greg and Ido Erev. (2003). "Small Feedback-Based Decisions and Their Limited Correspondence to Description-Based Decisions," *Journal of Behavioral Decision Making* 16, 215–233.

Binmore, Kenneth G. (1999). "Why Experiments in Economics?" *Economic Journal* 109, F16–F24.

Bone, John, John Hey and John Suckling. (1999). "Are Groups More (or Less) Consistent than Individuals?" *Journal of Risk and Uncertainty* 18, 63–81.

Brookshire, David S. and Don L. Coursey. (1987). "Measuring the Value of a Public Good: An Empirical Comparison of Elicitation Procedures," *American Economic Review* 77, 554–566.

Camerer, Colin F. (1995). "Individual Decision Making." In John H. Kagel and Alvin E. Roth (eds), *Handbook of Experimental Economics* 587–703. Princeton, NJ: Princeton University Press.

Camerer, Colin F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction* (Roundtable Series in Behaviorial Economics). Princeton, NJ: Princeton University Press.

Camerer, Colin F. and Teck-Hua Ho. (1999). "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica* 67, 827–874.

Carbone, Enrica and John D. Hey. (2000). "Which Error Story Is Best?" *Journal of Risk and Uncertainty* 20, 161–176.

Conlisk, John. (1989). "Three Variants on the Allais Example," *American Economic Review* 79, 392–407.

Cox, James C. and David M. Grether. (1996). "The Preference Reversal Phenomenon: Response Mode, Markets and Incentives," *Economic Theory* 7, 381–405.

Cubitt, Robin P., Chris Starmer and Robert Sugden. (1998). "On the Validity of the Random Lottery Incentive System," *Experimental Economics* 1, 115–131.

Cubitt, Robin P., Chris Starmer and Robert Sugden. (2001). "Discovered Preferences and the Experimental Evidence of Violations of Expected Utility Theory," *Journal of Economic Methodology* 8, 385–414.

Einav, Liran. (2005). "Informational Asymmetries and Observational Learning in Search," *Journal of Risk and Uncertainty* 30, 241−259.

Erev, Ido and Alvin E. Roth. (1998). "Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review* 88, 848–881.

Evans, Dorla A. (1997). "The Role of Markets in Reducing Expected Utility Violations," *Journal of Political Economy* 105, 622–636.

Harless, David W. (1992). "Actions versus Prospects: The Effect of Problem Representation on Regret," *American Economic Review* 82, 634–649.

Harless, David W. and Colin F. Camerer. (1994). "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica* 62, 1251–1289.

Hertwig, Ralf and Andreas Ortmann. (2001). "Experimental Practices In Economics: A Challenge for Psychologists?" *Behavioral and Brain Sciences* 24, 383–403.

Holt, Charles A. (1986). "Preference Reversals and the Independence Axiom," *American Economic Review* 76, 508–513.

Humphrey, Stephen J. (2006). "Does Learning Diminish Violations of Independence, Coalescing and Monotonicity?" *Theory and Decision* 61, 93–128.

Kahneman, Daniel and Amos Tversky. (1979). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263–291.

Keren, Gideon B. and Willem A. Wagenaar. (1987). "Violation of Utility Theory in Unique and Repeated Gambles," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13, 29–38.

List, John A. (2004). "Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace," *Econometrica* 72, 615–625.

Loomes, Graham. (1988). "Further Evidence of the Impact of Regret and Disappointment in Choice under Uncertainty," *Economica* 55, 47–62.

Loomes, Graham, Peter G. Moffat and Robert Sugden. (2002). "A Microeconometric Test of Alternative Stochastic Theories of Risky Choice," *Journal of Risk and Uncertainty* 24, 103–130.

Loomes, Graham, Chris Starmer and Robert Sugden. (2003). "Do Anomalies Disappear in Repeated Markets?" *The Economic Journal* 113, C153–C166.

Loomes, Graham and Robert Sugden. (1982). "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal* 92, 805–824.

MacCrimmon, Kenneth R. (1968). "Descriptive and Normative Implications of the Decision-Theory Postulates." In Karl H. Borch and Jan Mossin (eds), *Risk and Uncertainty* 3–23, St. Martin's Press, New York.

Myagkov, Mikhail G. and Charles R. Plott. (1997). "Exchange Economies and Loss Exposure: Experiments Exploring Prospect Theory and Competitive Equilibria in Market Environments," *American Economic Review* 87, 801–828.

Offerman, Theo, Jan Potters and Joep Sonnemans. (2002). "Imitation and Belief Learning in an Oligopoly Experiment," *Review of Economic Studies* 69, 973–997.

Quiggin, John. (1981). "Risk Perception and Risk Aversion among Australian Farmers," *Australian Journal of Agricultural Economics* 25, 160–169.

Schmidt, Ulrich and Tibor Neugebauer. (2006). "Testing Expected Utility in the Presence of Errors". *Economic Journal*, forthcoming.

Slovic, Paul and Amos Tversky. (1974). "Who Accepts Savage's Axiom?" *Behavioral Science* 19, 368–373.

Smith, Vernon L. (1982). "Microeconomic Systems as an Experimental Science," *American Economic Review* 72, 923–955.

Starmer, Chris. (2000). "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature* 38, 332–382.

Starmer, Chris and Robert Sugden. (1989). "Probability and Juxtaposition Effects: An Experimental Investigation of the Common Ratio Effect," *Journal of Risk and Uncertainty* 2, 159–178.

Starmer, Chris and Robert Sugden. (1991). "Does the Random-Lottery Incentive System Elicit True Preferences? An Experimental Investigation," *American Economic Review* 81, 971–978.

Tversky, Amos and Daniel Kahneman. (1986). "Rational Choice and the Framing of Decisions," *Journal of Business* 59, S251–S278.

Tversky, Amos and Daniel Kahneman. (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5, 297–323.