Int J Game Theory (2015) 44:113–133
DOI 10.1007/s00182-014-0421-2

On loss aversion, level-1 reasoning, and betting

Ido Erev · Sharon Gilat-Yihyie · Davide Marchiori · Doron Sonsino

Accepted: 20 March 2014 / Published online: 1 April 2014

Abstract Previous research suggests that human reaction to risky opportunities reflects two contradicting biases: "loss aversion", and "limited level of reasoning" that leads to overconfidence. Rejection of attractive gambles is explained by loss aversion, while counterproductive risk seeking is attributed to limited level of reasoning. The current research highlights a shortcoming of this popular (but often implicit) "contradicting biases" assertion. Studies of "negative-sum betting games" reveal high rate of counterproductive betting even when limited level of reasoning and loss aversion imply no betting. The results reflect two reasons for the high betting rate: initial tendency to participate and slow learning. Under certain conditions, the observed betting rate was higher than the rate predicted under random choice even after 250 trials with immediate feedback. These results can be captured with a model that assumes a tendency to select strategies that have led to good outcomes in a small set of similar past experiences, and allows for an initial framing effect.

Keywords Loss aversion · Level-1 reasoning · Samuelson's Colleague · Acquiring a company problem · Market for lemons

I. Erev (✉)
Faculty of Industrial Engineering and Management, Technion, 32000 Haifa , Israel
e-mail: erev@tx.technion.ac.il

S. Gilat-Yihyie
Department of Psychology, Western Galilee College, 24121 Acre , Israel

D. Marchiori
Strategic Organization Design Unit, University of Southern Denmark, 5230 Odense M , Denmark

D. Sonsino
The School of Business, The College of Management Academic Studies, 75190 Rishon Lezion , Israel

JEL Classification C63 · C73 · D03 · D82 · D83

1 Introduction

Experimental studies of choice between the status quo and riskier alternatives reveal contradicting deviations from maximization. Whereas most studies show a strong tendency to prefer the status quo even when this behavior impairs expected returns (see Samuelson and Zeckhauser 1988), some studies document a tendency to prefer counterproductive risky options over the status quo. A clear indication for these contradicting tendencies comes from comparison of the two classical problems presented in Table 1.

Investigation of "Samuelson's Colleague problem" (hereafter "colleague problem") reveals a strong tendency to prefer the status quo. For example, only 43 % of the participants in Redelmeier and Tversky (1992) were willing to play the hypothetical gamble, even though its expected value is much higher (+750) than the status quo. A similar pattern was documented in many other studies (e.g., Tom et al. 2007; Wedell and Böckenholt 1994).

Studies of the "Acquiring a company problem" (hereafter "company problem") reveal the opposite tendency. Most subjects bid significant amounts, even though the expected value of a positive bid is negative. Samuelson and Bazerman (1985) observed a mean bid of 58 in a one-shot play of the problem. Analyses of repeated play of the acquiring the company problem document limited learning. Ball et al. (1991) found that the mean bid in a 20-trial experiment with immediate feedback is 57. Grosskopf et al. (2007) and Bereby-Meyer and Grosskopf (2008) prove the robustness of this pattern in 100 repetitions.

The leading explanation of the tendency to prefer the status quo in the colleague and similar problems is the loss aversion hypothesis (Kahneman and Tversky 1979).

Table 1 "Samuelson's Colleague" and the "acquiring a company" problems

Problem	Samuelson's Colleague problem (Samuelson 1963; Redelmeier and Tversky's, 1992 variant)	The acquiring a company problem (Samuelson and Bazerman 1985)
Instructions	Imagine that you have the opportunity to play a gamble that offers a 50 % chance to win \$2,000 and a 50 % chance to lose \$500. Would you play the gamble?	The value of company C (in the market and to the current owner) is between 0 and 100 (uniformly distributed). You do not know the exact value, but know that if you acquire it, the value will increase by 50 %. To acquire the company you should submit a bid. What is your bid?
Expected value	The expected value of the gamble is \$750	The expected value of a bid of x ($0 \leq x < 100$) is $-x/4$
Main results	Proportion of play: 43 %. A preference for a counterproductive status quo.	Mean bid: 58. A preference for a counterproductive change of the status quo.

According to this explanation, losses loom larger than gains. That is, the possible loss of 500 looms larger than a possible gain of 2,000. This hypothesis provides an elegant explanation for a wide set of economic puzzles, including the endowment effect (Knetsch and Sinden 1984), the default effect (Johnson and Goldstein 2003), the disposition effect (Shefrin and Statman 1985), and under-investment in the stock market (Benartzi and Thaler 1995).

The common explanations of the opposite bias assume that people feel that they are smarter than others. The high bids in the company problem are explained with the assertion that participants tend to simplify their decision task by ignoring the fact that the seller can use her private information (Carroll et al. 1988; Tor and Bazerman 2003). This explanation can be described as "level-1" reasoning (see Stahl and Wilson 1995; Nagel 1995): Selecting the best option under the assertion that all the other agents are level-0 (behave randomly in the current case). Camerer et al. (2004) present an elegant generalization of this idea within a "cognitive hierarchy" theory. A similar "overconfidence" explanation was offered to capture overtrading in the stock market: Traders feel that they know more than the others (see Odean 1999).

Although both explanations are elegant, their coexistence can impair their usefulness. To predict behavior in new situations, it is necessary to understand the conditions that trigger the contradicting tendencies. The current paper tries to address this problem by comparing two distinct abstractions of the relation between loss aversion and level-1 reasoning. The first abstraction assumes a two-stage decision process (see related idea in Thaler et al. 1997, and Fox and Tversky 1998). At the first stage, agents use level-1 (or similar) reasoning to estimate the payoff distributions associated with the possible alternatives; in the company problem, for example, subjects estimate the payoff distributions under the assumption that sellers behave randomly. At the second stage, agents choose in accordance to prospect theory and exhibit loss aversion. The difference between the colleague and the company problems emerges, under this "2-stage" hypothesis, as a result of the fact that the colleague problem does not require an estimation of the payoff distribution (since the distributions are given). Thus, the results in this problem reveal the net effect of loss aversion.

The second abstraction assumes that loss aversion and level-1 reasoning are only two of the possible strategies considered by decision makers (see related ideas in Stahl 1996, and in Costa-Gomes et al. 2001). The probability of using these strategies depends on the framing of the task, and on the decision makers' experience. Certain framings trigger loss aversion, other framings trigger level-1 reasoning, and it is also possible that some framings trigger different behaviors. Thus, it is possible that the tendency to take counterproductive risk in the company problem and similar situations indicates that the framings of these tasks trigger risk taking rather than loss aversion and/or level-1 reasoning. This "framing hypothesis" is supported by two lines of recent research. One line demonstrates that initial behavior in variants of the colleague problem is highly sensitive to the framing of the task. For example, Ert and Erev (2008) found that abstract framing (choice between "0 for sure" and a gamble that pays "2,000 or −500 with equal probability") increases the rate of risky choices to 78 %. Another line of research reveals a high rate of risky choices in market entry games even when the incentive structure is unknown and the feedback after each choice is complete (or that the information conveyed from entering and staying out is the same). For

Table 2 Three variants of a betting game

State	Information sets		Payoffs from betting	
	Player 1	Player 2	Player 1	Player 2
A	A or B	A	+40	– 40
B	A or B	B or C	−30	+30
C	C or D	B or C	+20	– 20
D	C or D	D	−10	+ 10

Nature moves first and selects one of four equiprobable states A, B, C, or D. Each player receives a signal that provides some information about the realized state and has to choose if she wants to enter the zero-sum bet. The decisions are made simultaneously, and the bet is played if and only if both players decide to enter. The information sets and the payoffs from betting are identical in all three games. The games differ with respect to the outside payoff: The payoff when the bet is not played. This payoff is 0, 4, or 6 in, respectively, games Betting 0, Betting 4, and Betting 6

example, Erev et al. (2010a) gave subjects the following instructions: "This experiment includes several games. In each game you will be matched to interact with three other participants, for several trials. At each trial, each participant will be asked to choose between two options: 'stay out' or 'enter'. Your payoff in each trial will depend on your choice, the state of nature, and on the choices of the other participants (such that the more people enter, the less is the payoff from entry)." The observed entry rate at the first trial of the first game played by each player was 71 %. That is, most subjects preferred to enter to this risky game with unknown payoff rule.

The framing hypothesis implies that it should be possible to find situations in which behavior deviates from the predictions of both loss aversion and level-1 reasoning. Experiment 1 tests this prediction.

2 Experiment 1

Experiment 1 studies three variants of the two-person betting game presented in Table 2. At each iteration (round, trial) of the repeated game, nature selects a state (A, B, C, or D with equal probability), and each agent receives a signal (partial information) concerning the realized state. Specifically, depending on the drawn state, Player 1 receives one of the two signals "the state of nature is A or B" or "the state of nature is C or D," whereas Player 2 receives one the three signals "the state of nature is A," "the state of nature is B or C," or "the state of nature is D." Based on her signal, each agent has to decide between "Enter" and "Stay out." If both players enter, a "bet is played" and the payoffs are those reported on the right hand side of Table 2. If one or both players choose to stay out, the payoff for both players is the "outside outcome." Outside outcomes change across the three variants, and are 0, 4, and 6 in games "Betting 0," "Betting 4," and "Betting 6," respectively.

Notice that, from a normative point of view, Table 2 betting games are similar to the company problem. Iterated elimination of dominated strategies implies no betting in all three games (see Milgrom and Stokey 1982; Sebenius and Geanakoplos 1983).[1]

[1] Sonsino (1995, 1998) and Perets and Sonsino (1999) generalize the result to noisy strategic environments.

To clarify this prediction, note that Player 2 should never enter at Set A (staying out is the dominant strategy), and for that reason Player 1 should never enter at Set "A or B." If Player 1 does not enter at "A or B," Player 2 cannot benefit from entry at "B or C," and this implies that Player 1 cannot benefit from entering at "C or D."

Level-1 reasoning (assuming the other player randomly decides whether to enter or stay out) implies very different predictions: high entry rates (100 % in all information sets but "A") in Games Betting 0 and Betting 4, and no betting in Betting 6. The difference in predictions for the three games is a result of the fact that the expected payoff from betting, when uncertainty is resolved at level-1 (in information sets "A or B," "C or D," and "C or B"), is 5. Thus, level-1 agents will decide to enter if and only if the outside outcome is smaller than 5.

Previous experimental studies of betting games (see Sonsino et al. 2001; Søvik 2009; Rogers et al. 2009) focused on games where level-1 reasoning implies high betting rates, as in Betting 0 and Betting 4. The results of these studies can be elegantly approximated by the cognitive hierarchy generalization of the level-1 idea (Camerer et al. 2004), or by stochastic variants of this model that can be described as applications of truncated quantal response equilibrium (see Rogers et al. 2009). In the present context, however, the generalizations are not required: Like the basic Level-1 concept, all these models predict high betting rates in Games 0 and 4, and low betting rate in Game 6.

The two-stage hypothesis predicts no betting in Game 6, and allows for the possibility of betting in Games 0 and 4. Behavior in these games is likely to depend on the magnitude of loss aversion. For example, when Player 2 receives the signal "C or D" in Game Betting 0, she is predicted to estimate the payoff distributions as:

Enter: -10 with probability 0.25; $+20$ with probability 0.25; 0 otherwise
Stay out: 0 with certainty.

Under the linear variant of prospect theory used by Thaler et al. (1997), entering in this set is predicted if the subjective value of $+20$ is larger than the absolute subjective value of -10. This condition is met when the loss aversion parameter is smaller than two.

Assuming that the exact value of the outside options (e.g., the difference between 4 and 6) does not affect the framing of the task, the framing hypothesis predicts that the initial differences in behavior across the three games will not be large. Thus, this hypothesis allows for the possibility of high initial betting rates in Game Betting 6 as well.

2.1 Procedure

Fifty-two pairs of Technion students participated in the experiments. Twelve pairs were assigned to Game 0, twelve to Game 4, and 28 to Game 6.[2] Each pair played one

[2] We run more subjects in Game 6 because (under the level-1 assertion) the other games are similar to each other and to the games studied in previous research. Game 6 is different, and provides the clearest discrimination between the current hypotheses.

Table 3 Initial entry and betting rates in Experiment 1

Game	Very first trial (excluding sets A and D)	AB	CD	A	BC	D	Both players enter in the first play of each state
0	.76	.92	.83	.08	.50	.92	.48
4	.82	.67	1.00	.08	.92	1.00	.56
6	.85	.75	.82	.28	.78	.93	.46
Mean	.82	.77	.90	.19	.75	.94	.49

Average individual entry rates in the very first trial and in the first trial of each information condition

of the games repeatedly for 250 trials. The interaction between players was conducted via personal computers using the z-Tree software (Fischbacher 2007).

Instructions included a complete description of the payoff matrices. Subjects were told that, at every trial of the experiment, they would receive an information signal describing the state of nature, and would then have to decide whether to enter the bet or stay out. If at least one of them decides to stay out, both of them would receive the outside payoff of 0, 4, or 6 depending on the game they were assigned to. If both decide to enter, payoffs will be determined by the state of nature. The (translated) instructions are reported in Appendix 1. Instructions were provided to subjects in printed copy and explained verbally by the experimenters.

Each participant started the experiment with an initial endowment of 3,000 points (equivalent to 30 Shekels ≈9 US $). At every trial of the game, subjects could gain or lose points according to their realized payoffs. In particular, if one agent (or both) refuses the proposed bet at some trial, then both of them get the outside outcome (0, 4, or 6) points for that trial. The final payoff ranged from 28 to 46 Shekels.

Each trial started with the draw of one of the four possible states of nature (A, B, C, or D with equal probabilities). The selected state determined the information set (signal) presented to the subjects. For example, when the state was A, Player 1 saw the signal "A or B" and Player 2 saw the signal "A." After having observed the signal, players were asked whether they want to enter the game. After both players made their choices, they received feedback including: their choice, their opponent's choice, their payoff, their forgone payoff (the payoff for the choice they did not make), and their accumulated earning.

2.2 Results

The average individual entry rates at the very first trial (t = 1), and at the first trial in each information set, are presented on the left hand side of Table 3. The results reveal high entry rates, and almost no difference among the three conditions. For example, the entry rate in the very first trial over the subjects who did not face the trivial information sets (A and D) is significantly higher than 50 % in all three games. The exact values (and number of observations plus significance level in a binomial test) are 76 % (n = 21, $p = .03$), 82 % (n = 17, $p = .01$), and 85 % (n = 40, $p < .001$) in Games 0, 4, and 6, respectively. The difference among the three games is not significant.

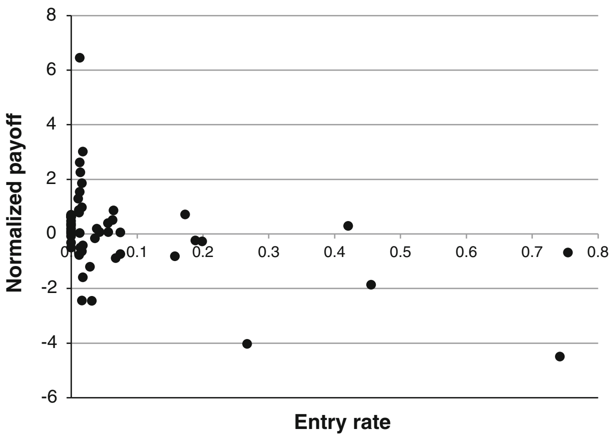

Fig. 1 Average normalized payoff (i.e., mean payoff minus the outside option) of Player 2 as a function of the entry rate in set A (52 observations)

The high entry rates in betting game 6 violate the predictions of the two-stage hypothesis. The two components of the two-stage hypothesis, level-1 reasoning and loss aversion imply no entry in all information sets but D. The results show the opposite: More entries than predicted by random choice.

Comparison of initial entry rates in the trivial information sets (A and D) reveals a surprising pattern. The results show more deviations from the dominant strategy in set A (19 % average entry rate), where the deviation can lead to a large loss of 40 points, compared to deviations in set D (6 % stay-out decisions), where this deviation has a smaller effect of losing no more than 10 points (10 points minus the outside payoff). The significant difference in deviation rates ($t_{51} = 2.44$, $p < .05$) suggests that the initial deviations cannot be easily explained as random noise (where errors should be symmetric) or by using a logistic response model (where the error rate decreases in the cost of error).

Under one explanation, the high initial entry rate in set A reflects the attempt to confuse the other agent (i.e., to make him or her believe that attaining the maximal payoff of +40 is possible). In order to evaluate this explanation we plot, in Fig. 1, the average normalized payoff of Player 2 as a function of the entry rate in set A. Each dot in this curve presents one of the 52 Players 2 (across the 3 conditions). Recall that entry in this set is a dominated action, and can be very costly (leading to a loss of 40 points). Nevertheless, the results suggest that some entries can be effective. The normalized payoff of the 29 Players 2 that entered in Set A at least once but in no more than 10 % of the trials is higher (Mean = 0.40, std = 1.77) than the normalized payoff of the 14 Players 2 that never entered in Set A (Mean = 0.26, std = 0.39). More entries at A appear counterproductive, but the sample is small. The normalized payoff of the 9 agents that entered in more than 10 % of the trials is negative (Mean = −1.28, std = 1.85). Thus, it is possible that at least some of the entries at set A reflect sophisticated behavior. A second possible contributor to the entry in Set A is overgeneralization

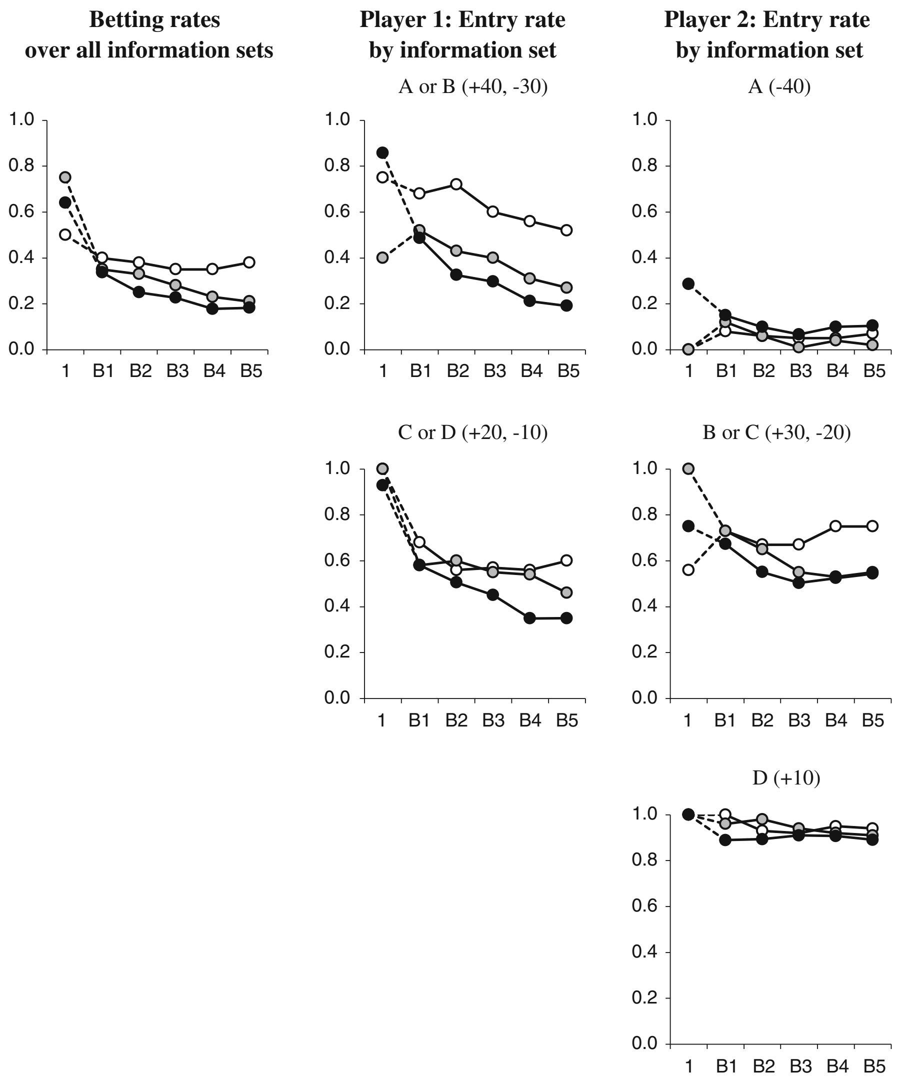

Fig. 2 Betting rates (proportion of trials in which both players enter) and entry rates by game in the first trial and in five blocks of 50 trials. The *hollow circles* are Game 0, the *gray circles* are Game 4, and the *black circles* are Game 6

and/or failure to carefully read the instructions. It is possible that some subjects came to the lab thinking that they should first try to enter and then think again.

Figure 2 presents the observed entry rates in five blocks of 50 trials and in the very first trial (t = 1). The results show faster learning to stop betting in Games 4 and 6 than in Game 0. Whereas the betting rates at the first block of 50 trials are very similar, the betting rates at the last block are 38 % in Game 0, 21 % in Games 4, and 18 % in Game 6. The differences between Game 0 and each of the other two games are significant ($t_{22} = 1.98$, $p < .05$, for Game 4, and $t_{38} = 2.81$, $p < .05$, for Game 6).

The difference between Games 4 and 6 is insignificant ($t_{38} = 0.50$, $p = .31$). These comparisons suggest that effect of relative payoffs on betting rates slowly emerges through limited learning.

Individual differences Table 4 presents the entry rates and the average payoff of each of the 52 dyads in the last block of 50 trials. The results reveal that 34 out of the 52 Players 1 earned less than their outside option, and only 6 earned more than their outside option. The pattern for Players 2 is more balanced: 21 lost, and 19 benefited from entering. Twelve pairs (2 pairs in Game 0, 2 in Game 4, and 8 in Game 6) never played the bet in the last 50 trials and always got the outside payoff.

The boredom hypothesis One explanation of the high entry rates is that they reflect the fact that staying out, and getting the same payoff in all trials, is very boring. Thus, the high entry rate can be rationalized claiming that subjects maximize a utility function that decreases with boredom. This explanation was evaluated in Sonsino et al. (2001). One of the conditions in that paper examined a variant of Game Betting 0 in which the outside option was a play of a gamble with three possible outcomes. The expected value of the gamble was equal to the original outside option, and the variance was similar to the variance of payoffs under mutual entry. The change did not decrease the high entry rates.

3 Experiment 2

The framing hypothesis implies that the high initial entry rates observed in Experiment 1 are triggered by certain features in the description of the game. The similarity of current results to the results Erev et al. (2010b) market entry experiments suggests that the critical factor can be the framing of decision as the choice between "entering" an interaction or "staying outside." Experiment 2 was designed to evaluate this hypothesis. It studies Game Betting 6 under two limited information conditions: "Enter" and "Risk."

In Condition Enter, agents were told (see instructions in Appendix 2) that at each trial they would be asked if they want to participate in a betting game. They were informed that the payoff would depend on the state of nature. It was explained that the payoff for playing (entering) can be positive or negative, and the payoff from not playing (staying out) would be 6 points.

The instructions in condition risk (Appendix 2) were identical to the instructions in Condition Enter, except for changing the names of the two alternatives. The decision to enter was named "risky participation and data collection," whereas the decision to stay out was changed to "safe participation and data collection."

3.1 Procedure

The procedure was identical to the one in Experiment 1 with one exception: as explained above, the instructions did not provide a description of the payoff rule. The participants were not even informed that the experiment involves an interaction with another person. The feedback was limited to their decisions and outcomes. Twelve pairs of subjects participated in each condition.

Table 4 Entry rates and average payoffs in the last block of 50 trials by Dyad

	Player 1's information sets		Player 2's information sets			Payer 1's avg. payoff	Payer 2's avg. payoff
	A or B	C or D	A	B or C	D		
Game Betting 0							
Pair 1	0.78	0.63	0.00	0.68	1.00	−1.80	1.80
Pair 2	1.00	0.00	0.00	0.90	0.93	−6.00	6.00
Pair 3	0.71	0.77	0.00	0.93	1.00	−4.00	4.00
Pair 4	0.15	1.00	0.06	0.78	1.00	1.00	−1.00
Pair 5	0.13	0.88	0.00	0.82	1.00	−0.40	0.40
Pair 6	1.00	0.10	0.00	0.10	1.00	0.00	0.00
Pair 7	0.50	0.41	0.00	0.00	0.56	−0.40	0.40
Pair 8	0.00	0.60	0.09	0.87	1.00	0.00	0.00
Pair 9	0.15	1.00	0.08	0.56	1.00	−0.20	0.20
Pair 10	0.55	0.11	0.06	0.74	0.80	−1.80	1.80
Pair 11	0.88	0.72	0.57	0.84	1.00	1.40	−1.40
Pair 12	0.38	1.00	0.00	0.90	1.00	2.00	−2.00
Mean	0.52	0.60	0.07	0.68	0.94	−0.85	0.85
Game Betting 4							
Pair 1	0.15	0.67	0.07	0.31	1.00	2.44	3.64
Pair 2	0.35	0.38	0.20	0.67	1.00	−1.16	7.24
Pair 3	0.17	0.56	0.00	0.45	1.00	2.12	4.12
Pair 4	0.33	0.96	0.00	0.83	1.00	1.00	2.20
Pair 5	0.20	0.48	0.00	0.80	1.00	2.88	2.88
Pair 6	0.00	0.00	0.00	0.50	1.00	4.00	4.00
Pair 7	0.64	0.72	0.00	0.43	1.00	3.52	2.72
Pair 8	0.04	0.08	0.00	0.86	1.00	3.96	3.56
Pair 9	0.32	0.32	0.00	0.79	0.92	3.04	3.04
Pair 10	0.00	0.00	0.00	0.00	0.00	4.00	4.00
Pair 11	0.00	0.63	0.00	0.44	1.00	1.84	4.24
Pair 12	1.00	0.72	0.00	0.05	1.00	1.52	4.72
Mean	0.27	0.46	0.02	0.51	0.91	2.43	3.86
Game Betting 6							
Pair 1	0.00	0.36	0.00	0.43	0.89	6.20	4.60
Pair 2	0.48	0.36	0.00	0.31	1.00	2.52	7.32
Pair 3	0.00	0.08	0.00	0.23	1.00	6.28	5.48
Pair 4	0.44	0.36	0.00	0.40	1.00	2.72	7.12
Pair 5	0.00	1.00	0.00	0.55	1.00	4.12	3.32
Pair 6	0.00	0.13	0.00	0.00	1.00	5.04	6.24
Pair 7	0.00	0.78	0.00	0.64	1.00	3.88	4.28
Pair 8	0.59	0.61	0.81	0.77	0.42	8.44	−0.76
Pair 9	0.00	0.00	0.00	0.00	0.00	6.00	6.00

Table 4 continued

	Player 1's information sets		Player 2's information sets			Payer 1's avg. payoff	Payer 2's avg. payoff
	A or B	C or D	A	B or C	D		
Pair 10	0.00	0.00	0.00	0.00	1.00	6.00	6.00
Pair 11	0.16	0.56	0.00	0.00	0.94	3.12	6.72
Pair 12	0.00	0.31	0.00	0.42	0.94	5.60	5.20
Pair 13	0.04	0.00	0.00	0.00	1.00	6.00	6.00
Pair 14	0.00	0.10	0.00	0.05	0.92	5.36	6.16
Pair 15	0.05	0.35	0.13	0.17	0.83	5.28	5.28
Pair 16	0.62	0.67	0.31	1.00	1.00	−1.32	7.08
Pair 17	0.42	0.83	0.00	1.00	1.00	1.88	3.88
Pair 18	0.23	0.25	0.00	1.00	1.00	1.28	8.08
Pair 19	0.04	0.77	0.07	1.00	1.00	5.00	2.20
Pair 20	0.61	0.82	0.00	0.97	1.00	−0.24	5.76
Pair 21	0.00	0.00	1.00	1.00	1.00	6.00	6.00
Pair 22	0.00	0.00	0.00	0.00	0.00	6.00	6.00
Pair 23	0.68	0.48	0.00	0.31	1.00	4.20	5.40
Pair 24	0.52	0.52	0.42	1.00	1.00	−0.24	6.96
Pair 25	0.00	0.00	0.00	1.00	1.00	6.00	6.00
Pair 26	0.46	0.45	0.00	1.00	1.00	3.16	5.96
Pair 27	0.00	0.00	0.08	1.00	1.00	6.00	6.00
Pair 28	0.00	0.00	0.11	1.00	1.00	6.00	6.00
Mean	0.19	0.35	0.10	0.54	0.89	4.30	5.51

Table 5 Initial entry and betting rates in Experiment 2

Condition	Very first trial	AB	CD	A	BC	D	Betting (both enter in first play of each state)
Enter	.95	.92	.75	.83	.92	.75	.56
Risk	.55	.75	.42	.58	.50	.33	.37

The average individual entry rates at the very first trial, overall and by information set

3.2 Results

Table 5 presents the initial entry rates in Experiment 2. The mean entry rate in the very first trial is 95 % in Condition Enter, and 55 % in condition risk. The difference between the two conditions is highly significant ($t_{46} = 4.08$, $p < .001$).

Note that the high entry rate in Condition Enter appears to contradict the basic concepts of risk aversion, loss aversion, and ambiguity aversion: Most participants

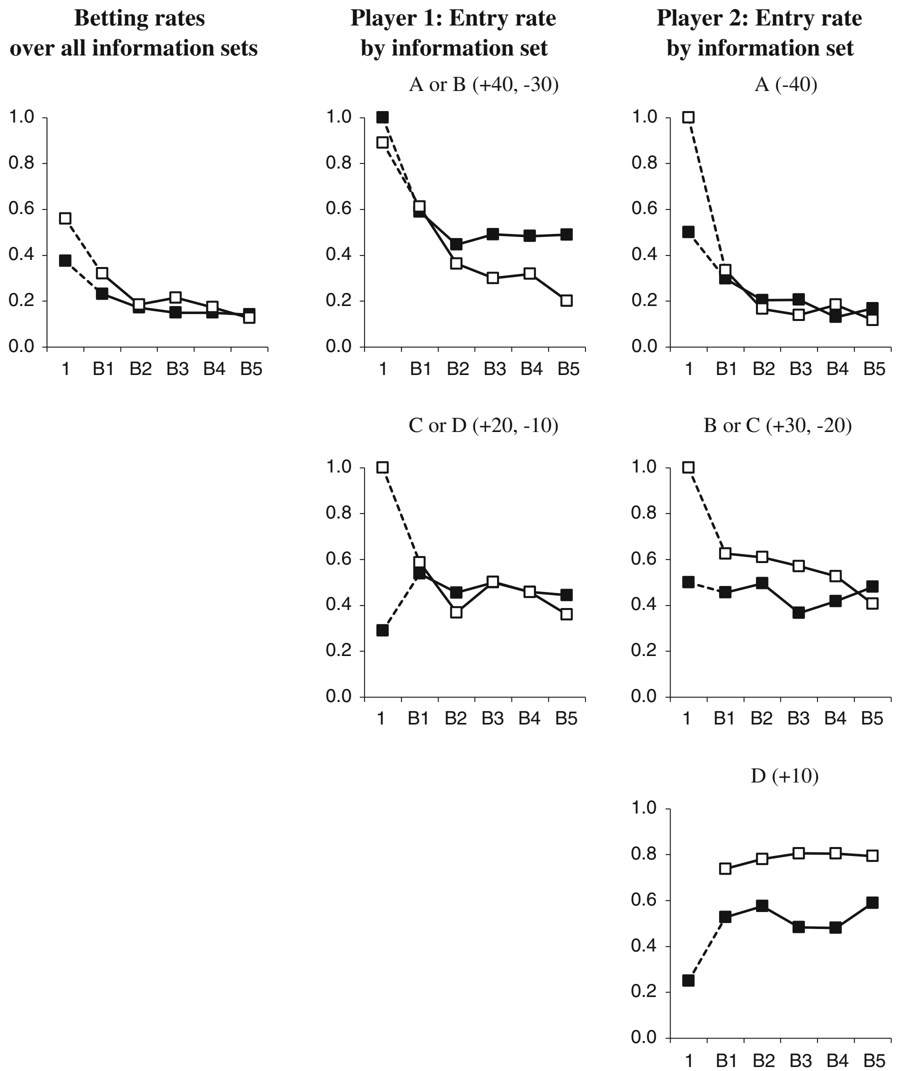

Fig. 3 Betting rates (proportion of trials in which both players enter) and entry rates by information set in the first trial and in five blocks of 50 trials. The *hollow squares* are condition Enter and the *filled squares* are condition risk. The missing first trial point in Set D reflects the fact that Set D did not occur in the very first trial in Condition Enter

(95 %) preferred an option that could lead to unspecified gain or loss over a sure gain.[3]

The learning curves presented in Fig. 3 show that the large initial difference in betting rates between the two conditions was eliminated by experience. Start-

[3] In addition, the results show higher entry rates in sets that include the letter A. Whereas we have not predicted this pattern in advance, it appears to be consistent with the current overgeneralization story: in many natural settings the signal A implies higher expected outcomes than the signal D.

Table 6 Entry rates and average payoffs in the last block of 50 trials by Dyad

	Player 1's information sets		Player 2's information sets			Payer 1's avg. payoff	Payer 2's avg. payoff
	A or B	C or D	A	B or C	D		
Condition Enter							
Pair 1	0.31	0.29	0.20	0.89	0.59	6.52	3.32
Pair 2	0.18	0.64	0.00	0.21	1.00	3.24	5.64
Pair 3	0.13	0.22	0.00	1.00	1.00	6.36	3.96
Pair 4	0.00	0.00	0.00	0.13	1.00	6.00	6.00
Pair 5	0.23	0.71	0.15	0.04	0.85	4.40	5.20
Pair 6	0.08	0.04	0.17	0.13	1.00	4.96	6.56
Pair 7	0.09	0.21	0.00	0.00	0.80	5.04	6.24
Pair 8	0.50	0.77	0.13	0.40	0.60	5.56	3.56
Pair 9	0.31	1.00	0.00	0.44	1.00	1.80	6.60
Pair 10	0.00	0.00	0.11	1.00	0.75	6.00	6.00
Pair 11	0.58	0.42	0.22	0.13	0.27	4.60	6.20
Pair 12	0.00	0.00	0.43	0.50	0.67	6.00	6.00
Mean	0.20	0.36	0.12	0.41	0.79	5.04	5.44
Condition risk							
Pair 1	1.00	0.89	0.00	0.00	0.00	6.00	6.00
Pair 2	1.00	1.00	0.46	0.20	0.33	7.20	1.20
Pair 3	0.04	0.17	0.06	0.91	1.00	5.32	5.72
Pair 4	1.00	1.00	0.00	0.00	0.00	6.00	6.00
Pair 5	0.03	0.14	0.00	0.82	1.00	5.52	5.52
Pair 6	0.80	0.63	0.00	0.40	0.00	5.96	4.36
Pair 7	0.69	0.13	0.42	0.63	0.50	4.52	5.32
Pair 8	0.77	0.67	0.92	0.52	0.63	9.60	-3.60
Pair 9	0.22	0.30	0.10	0.66	1.00	4.84	5.24
Pair 10	0.00	0.00	0.00	0.27	0.63	6.00	6.00
Pair 11	0.28	0.20	0.06	0.36	1.00	4.88	5.68
Pair 12	0.05	0.21	0.00	1.00	1.00	5.16	5.16
Mean	0.49	0.44	0.17	0.48	0.59	5.92	4.38

ing from the second block, the results reveal low and similar betting rates in both conditions.

Individual differences Table 6 presents the last block results for each of the 24 pairs. Comparison of these results to the results in Game betting 6 in Table 4 suggests that the lack of information reduced the advantage of Player 2, and reduced the between-pair variance. In addition, the data reveals strong negative correlation between entry at AB and BC ($r = -.72$, $p < .01$) in condition risk. This pattern can explain the high entry rate in Set AB in condition risk: most entries in this set are by Players 1 that interacted with Players 2 who rarely entered at BC.

4 The joint effect of instructions, framing and experience

Comparison of the learning curves in the five conditions reveals two interesting patterns. First, there is a large difference between the effect of the incentive structure (the outside payoffs in Experiment 1), and the effect of the framing (the labels in Experiment 2). The different outside payoffs had little effect on initial behavior (an initial entry rate of about 80 % in all three games), but had a large effect on learning. The results show almost no reduction in betting in Game 0, and a clear reduction in Game 6. The betting rates in the last block were 38 and 18 %, respectively, in Games 0 and 6. Analysis of the framing manipulations of Experiment 2, however, reveals an opposite pattern: A large effect on initial entry rates (95 vs. 55 % at the very first trial), but almost no effect after 50 repetitions. The betting rate in the last block was around 14 % in the two conditions.

A second interesting pattern emerges in comparison of the two experiments. The complete description of the game payoffs in Experiment 1 had a large effect on behavior in information sets A and D, but this effect did not speed up learning in the other information sets. For example, the betting rate at the last block of Game 6 with information (18 %) was not lower than the respective rate for the minimal information conditions of Experiment 2 (14 and 13 %).

In order to clarify the implications of these observations we tried to capture them with a simple abstraction of the joint effect of framing and experience. This quantification, referred to as the "instructions & sampling" model includes the following three assumptions.

A1. Two sources of information When facing the information set at trial $t + 1$, players consider two sources of information: first, the experimental instructions and, second, all their previous experiences with the relevant information set along the experiment. The number of previous experiences with a given set is denoted as $t_{set.}$

A2. Instructions-based evidence The information implied by the instructions is described by two parameters: s, which represents the weight of the instructions in terms of the number of "evidence," and $0 < q < 1$, which represents the proportion of the evidence that favors entering.

A3. Reliance on small samples The decision at trial $t+1$ is based on $t_{set} + s$ evidences: The t_{set} past experiences with the relevant information set, and the *s* instructions-based evidences. The decision is a best reply to a random sample of *k* evidences from this set of $t_{set} + s$ evidences. The value of *k* is a free parameter.

When the sample of k evidences includes a decisive instruction-based evidence (positive or negative), this evidence determines the decision at trial $t+1$ independently of the experiences during the first t trials. That is, the members of the sample that belong to the t_{set} experiences affect the final choice only if the instructions-based evidences lead to indifference and/or when the sample does not include instructions-based evidence. The effect of the experiences increases during the experiment because the probability of a decisive instructions-based evidence decreases with time. When t_{set} is much larger than s, most samples do not include instructions-based evidence.

Thus, the probability that the t_{set} experiences will determine the choice at trial t + 1 increases with t.[4]

Notice that the current model can be described as simplified Bayesian inference, and as an example of a case-based decision process (Gilboa and Schmeidler 1995). The final assumption is implied by the payoff sampling equilibrium concept proposed by Osborne and Rubinstein (1998). The descriptive potential of the sampling concept in simple games is demonstrated in Selten and Chmura (2008), as well as in two choice prediction competitions (Erev et al. 2010a, b; and see review in Erev and Haruvy, in press).

4.1 Quantitative fit

The current model has three parameters: s–the strength of the instructions, q–the proportion of instructions-based evidence that favors entry, and k–the sample size. We estimated q to fit the initial entry rates in each information set (as in Roth and Erev 1995), and estimated s and k to fit the observed learning curves.

The estimation was based on a grid search procedure. We derived the model predictions under different values of s and k in the range $s = 1$ to 20, and $k = 1$ to 10, and selected the values that minimize the mean squared deviation between the observed and predicted betting rate curves (in five blocks of 50 trials each). Each simulation started with the estimation of the value of q in the different information sets rates given the assumed value of k. The estimated value is the one that yields the observed initial entry rate under the assumed value of k. For example, when the initial entry rate is 0.82, and $k = 3$, the estimated value of q is .74 (with this value, a sample of 3 favors entering in 82 % of the cases).

Best fit was obtained with the parameters $s = 8$ and $k = 3$. Figure 4 presents the betting rates in virtual replications of the two experiments with simulated agents that behave in accordance with the current model with these parameters.

The results reveal that the instructions and sampling model provides sufficient conditions to the joint effect of framing and experience: It captures the large initial framing effect, the fact that experience decreases the framing effect, the high sensitivity to outside options, and the failure to converge to rational behavior. In addition, the model suggests a large difference between the effect of framing and experience. Whereas we had to assume situation specific framing parameters (the value of q varied between the different conditions and information sets), it was not necessary to assume situation specific learning parameters. Moreover, the current abstraction of the effect of experience is a simplified variant of the abstraction that won two recent choice prediction competitions (Erev et al. 2010a, b). It seems that the effect of experience is relatively general and can often be approximated

[4] Under one justification of this assumption, the numerical value of the instructions-based evidence is much larger than the numerical value of the new evidence (between −40and + 40). For example, if the value is +1,000 or −1,000 the existence of single instructions-based evidence in the sample masks the effect of at least 25 new experiences. The probability of the event "at least one instructions-based evidence in the sample after t_{set} experiences with a particular information set" is $1 - (t_{set}/[s + t_{set}])^k$.

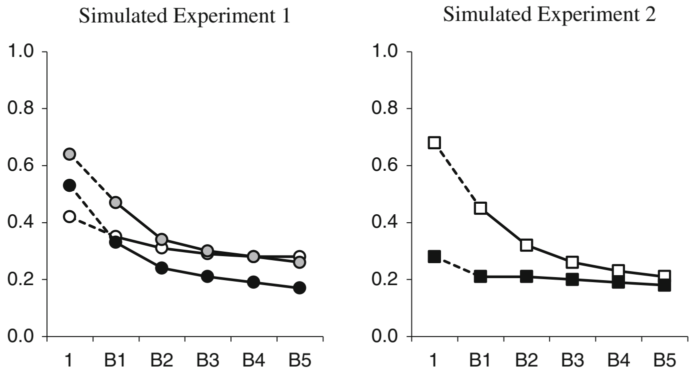

Fig. 4 Simulated betting rates in the first trial, and in five blocks of 50 trials. Experiment 1: the *hollow circles* are Game 0, the *gray circles* are Game 4, and the *black circles* are Game 6. Experiment 2: the *hollow squares* refer to Condition Enter, the *filled squares* to condition risk

with the assertion that people tend to rely on small samples of similar past experiences.

5 General discussion

Previous research explains the tendency to reject attractive risky offers in the colleague and similar problems as a reflection of loss aversion, and explains the opposite bias in the company acquisition game and similar problems as a reflection of level-1 reasoning. The current attempt to clarify the joint effect of the two tendencies, leads to surprising findings. It turns out that behavior in betting game contradicts both assertions. In certain settings (Game Betting 6), people tend to bet even when loss aversion and level-1 reasoning imply no betting (and no betting is predicted under rational choice, too).

We believe that the current results, and the contradicting deviations from rational choice that motivated this study, are best explained with the assertion that the likelihood of the behavior predicted by loss aversion and by level-1 reasoning is sensitive to the framing of the task. The framings used in the current study did not trigger these tendencies. Rather, the results suggest that the framing of the risky option as "entry to a game" triggers deviations from rational choice that can contradict both loss aversion and level-1 reasoning. In addition, the results show that experience of 250 trials with immediate feedback decreases the effect of the framing, but still does not imply convergence to rational behavior.

5.1 Values, rationality, and experience

To clarify the wider implications of the current results, it is constructive to distinguish between three interpretations of experimental evidence of suboptimal behavior. The most popular interpretation assumes that experimental results shed light on the subjective values and beliefs that drive behavior. Thus, the results can be used to refine the rational model. For example, indications for "loss aversion" and for "other regarding preferences" are assumed to reflect the shape of the subjective utility function, and indications for "level-1" reasoning are taken as reflections of the nature of subjective beliefs.

Under an "initial errors" interpretation, deviations from optimality reflect initial errors that are made by inexperienced agents. Repeated experiences with the appropriate incentive structure are expected to eliminate these mistakes. This interpretation is implicit in mainstream research that ignores contradicting experimental evidence.

Finally, under the "learning" approach, deviations from optimal behavior occur since agents can never be sure that they know the correct incentive structure. Thus, they rely on past experiences in similar situations, and for that reason exhibit high sensitivity to the framing of the task. This interpretation implies that experience affects choice behavior (as it can reduce overgeneralizations), but does not imply convergence to rational choice. The effect of experience is assumed to depend on the interaction between the incentive structure and the basic properties of human learning.

We believe that the current results highlight important shortcomings of the first two interpretations, and demonstrate the potential value of the learning approach. The shortcoming of the "subjective values and beliefs" approach is demonstrated by the fact that the current results are in the opposite direction of three of the best-known deviations from the rational model suggested by previous research (loss aversion, level-1 reasoning and ambiguity aversion). The shortcoming of the "initial errors" approach is clarified by the observation that the learning process can be very slow. In some cases, like in Game Betting 0, 250 trials with immediate feedback are not enough to move behavior to the "rational side" of random choice. That is, the deviation from random choice (betting rate of 38 % in the last block) is in the opposite direction of the prediction of the rational model. The potential of the learning approach is reflected in the similarity of the learning processes documented here and in previous research: The main results can be approximated assuming that participants rely on small samples of past experiences.

5.2 Methodological implications and open questions

It is important to emphasize that the current results do not question the existence of loss aversion and/or level-1 reasoning. That is, the results do not question the view that these phenomena provide useful summary of human behavior in important sets of situations. The main contribution of the current analysis concerns the interpretation of these phenomena, and the methodology that should be used to clarify their implications. The possibility that these phenomena are best described as "strategies that were found to be effective in certain settings" (and are sometimes overgeneralized), leads to very different implications than the common working assumption that these phenomena are reflections of stable subjective values and/or beliefs. Specifically, the current learning interpretation suggests that careful study of one of these phenomena in certain settings may not shed much light on behavior in other settings. We believe that the most important open questions involve the marking of boundary conditions for the emergence of different phenomena.

6 Summary

The current results challenge the view that human reaction to risky offers is driven by a tradeoff between loss aversion and limited level of reasoning that leads to overconfidence. Our participants chose to accept negative sum bets even when limited level of reasoning implies that these bets should be avoided. These results can be explained with the assertion that loss aversion and level-1 reasoning are only two of many possible framing effects. Certain framings trigger loss aversion and/or level-1 reasoning, and other framings trigger very different behaviors. Specifically, the naming of the risky option as "Enter" and the safe option as "Stay out" was found to trigger deviations from rational choice in the opposite direction of loss aversion and level-1 reasoning. In addition, the results show that experience decreases the effect of the framing, but does not insure convergence to rational choice. The joint effect of framing and experience can be captured with a model that assumes reliance on small samples of evidence.

Acknowledgments This paper was supported by a grant from the Israel Science Foundation. The paper benefitted for constructive comments from the participants of the "Theory, Decision, and Applications" meeting held in Paris in June 2011. All authors contributed equally.

Appendix 1: instructions of "Game Betting 6"

You are given the possibility to participate in a betting-game as follows:

The game has two players: One and Two
The bet has four possible outcomes represented by letters: A, B, C, and D
The probability of each outcome is 1/4

The payoff to both players will be determined by the following rule:

If the realized state is A, Player 2 will pay Player 1 40 points
If the realized state is B, Player 1 will pay Player 2 30 points
If the realized state is C, Player 2 will pay Player 1 20 points
If the realized state is D, Player 1 will pay Player 2 10 points

You have to decide whether you would want to participate in the bet without knowing the actual realized state. Yet, the bet's organizers will give you some partial information (signal) concerning the realized state before your decision.

Player 1 will receive the signal **"A or B"** when the realized state is A or B. He will get the signal **"C or D"** when the realized state is C or D.
Player 2 will receive the signal **A** when the realized state is A, the signal **"B or C"** when the realized state is B or C, and the signal **D** when the realized state is D.

The structure of the bet is summarized in the next figure:

	A	B	C	D
Player 1	+40	-30	+20	-10
Player 2	-40	+30	-20	+10

The bets' payoffs will be actually played only if both players (1 and 2) will decide to accept the bet. If one of the players will decide to reject, the bet will not take place and each player will get a payoff of +6 points.

The experiment includes 250 rounds. In each round of the experiment, you have to decide whether you would like to take the bet. If you decided to accept the bet, you should press **Enter**. Otherwise, you should press **Stay Out**.

After each round of the experiment, you will observe a feedback window disclosing the following information:

Your response
The other player's response
The realized state
Your payoff for that round
Your foregone payoff for that round
Your total accumulated score

To exit the feedback window and move to the next trial, you should press **Continue**.
You start the experiment with 3,000 points.
Your payoff will be determined by your final cumulative score at the end of the experiment. The value of each point is 1 agora (0.01 Shekels).

Appendix 2: instructions of Experiment 2

Condition Enter

The current experiment includes 250 rounds. In each round, you are given the possibility to participate in a game that has four possible states represented by letters: A, B, C, and D. The probability of each state is 1/4.

You have to decide whether you want to participate in the game without knowing the actual realized state. Yet, the bet's organizers will give you some partial information (signal) concerning the realized state before your decision. The potential gain or loss from participating in the bet is predetermined by a payoff rule for each realized state.

At each trial you have to decide between "Enter" and "Stay out." If you decide to stay out you will get a fixed payoff of 6 points. If you decide to enter, your payoff can differ from 6 points.

After each round of the experiment, you will observe a feedback window disclosing the following information:

Your response
Your payoff for that round
Your foregone payoff for that round
Your total accumulated score

You start the experiment with 3,000 points.
Your payoff will be determined by your final cumulative score at the end of the experiment. The value of each point is 1 agora (0.01 Shekels).

Condition risk

The instructions in condition risk were identical to the instructions in Condition Enter with the exception that third paragraph (in italic) read as follows:

At each trial you have to decide between “safe participation and data collection” and “risky participation and data collection.” If you decide to participate and collect data safely, you will get a fixed payoff of 6 points. If you decide to participate riskily, your payoff can differ from 6 points.

References

Ball SB, Bazerman MH, Carroll JS (1991) An evaluation of learning in the bilateral winner’s curse. Organ Behav Hum Decis Process 48:1–22

Benartzi S, Thaler RH (1995) Myopic loss aversion and the equity premium puzzle. Quart J Econ 110:73–92

Bereby-Meyer Y, Grosskopf B (2008) Overcoming the winner’s curse: an adaptive learning perspective. J Behav Decis Mak 21:15–27

Camerer CF, Ho TH, Chong JK (2004) A cognitive hierarchy model of games. Q J Econ 119:861–898

Carroll JS, Bazerman MH, Maury R (1988) Negotiation cognition: a descriptive approach to negotiators’ understanding of their opponents. Organ Behav Hum Decis Process 41:352–370

Costa-Gomes M, Crawford VP, Broseta B (2001) Cognition and behavior in normal-form games: an experimental study. Econometrica 69:1193–1235

Erev I, Ert E, Roth AE (2010a) A choice prediction competition for market entry games: an introduction. Games 1:117–136

Erev I, Ert E, Roth AE, Haruvy H, Herzog SM, Hau R, Hertwig R, Stewart T, West R, Lebiere C (2010b) A choice prediction competition: choices from experience and from description. J Behav Decis Mak 23:15–47

Erev I, Haruvy H (in press) Learning and the economics of small decisions. In: Kagel JH, Roth AE (Eds) The handbook of experimental economics. Princeton University Press, Princeton

Ert E, Erev I (2008) The rejection of attractive gambles, loss aversion, and the lemon avoidance heuristic. J Econ Psychol 29:715–723

Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments. Exp Econ 10:171–178

Fox CR, Tversky A (1998) A belief-based account of decision under uncertainty. Manag Sci 44:879–895

Gilboa I, Schmeidler D (1995) Case-based decision theory. Q J Econ 110:605–639

Grosskopf B, Bereby-Meyer Y, Bazerman MH (2007) On the robustness of the winner’s curse phenomenon. Theory Decis 63:389–418

Johnson EJ, Goldstein D (2003) Do defaults save lives? Science 302:1338–1339

Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. Econometrica 47:263–291

Knetsch JL, Sinden JA (1984) Willingness to pay and compensation demanded: experimental evidence of an unexpected disparity in measures of value. Q J Econ 99:507–521

Milgrom P, Stokey N (1982) Information, trade and common knowledge. J Econ Theory 26:17–27

Nagel R (1995) Unraveling in guessing games: an experimental study. Am Econ Rev 85:1313–1326

Odean T (1999) Do investors trade too much? Am Econ Rev 89:1279–1298

Osborne MJ, Rubinstein A (1998) Games with procedurally rational players. Am Econ Rev 88:834–847

Perets H, Sonsino D (1999) On preplay negotiations and zero-sum betting. Int Game Theor Rev 1:192–196

Redelmeier DA, Tversky A (1992) On the framing of multiple prospects. Psychol Sci 3:191–193

Rogers BW, Palfrey TR, Camerer CF (2009) Heterogeneous quantal response equilibrium and cognitive hierarchies. J Econ Theory 144:1440–1467

Roth AE, Erev I (1995) Learning in extensive form games: experimental data and simple dynamic models in the intermediate term. Games Econ Behav 8:164–212

Samuelson PA (1963) Risk and uncertainty: a fallacy of large numbers. Scientia 98:108–113

Samuelson W, Bazerman MH (1985) Negotiation under the winner’s curse. In: Smith VL (ed) Research in experimental economics. Jai Press, Greenwich, pp 105–137

Samuelson W, Zeckhauser R (1988) Status quo bias in decision making. J Risk Uncertain 1:7–59

Sebenius JK, Geanakoplos J (1983) Don't bet on it: contingent agreements with asymmetric information. J Am Stat Assoc 78:424–426

Selten R, Chmura T (2008) Stationary concepts for experimental 2×2 games. Am Econ Rev 98:938–966

Shefrin H, Statman M (1985) The disposition to sell winners too early and ride losers too long: theory and evidence. J Finance 40:777–790

Sonsino D, Erev I, Gilat S (2001) On rationality, learning and zero-sum betting: an experimental study of the no-betting conjecture. Working Paper, http://ie.technion.ac.il/Home/Users/erev/Sonsino_Erev_Gilat_2001.pdf. Accessed 30 March 2014

Sonsino D (1995) "Impossibility of speculation" theorems with noisy information. Games Econ Behav 8:406–423

Sonsino D (1998) Sebenius and geanakoplos model with noise. Int J Game Theory 27:111–130

Søvik Y (2009) Strength of dominance and depths of reasoning: an experimental study. J Econ Behav Organ 70:196–205

Stahl DO (1996) Boundedly rational rule learning in a guessing game. Games Econ Behav 16:303–330

Stahl DO, Wilson PW (1995) On players' models of other players: theory and experimental evidence. Games Econ Behav 10:218–254

Thaler RH, Tversky A, Kahneman D, Schwartz A (1997) The effect of myopia and loss aversion on risk taking: an experimental test. Q J Econ 112:647–661

Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision making under risk. Science 315:515–518

Tor A, Bazerman MH (2003) Focusing failures in competitive environments: explaining decision errors in the monty hall game, the acquiring a company problem, and multiparty ultimatums. J Behav Decis Mak 16:353–374

Wedell DH, Böckenholt U (1994) Contemplating single versus multiple encounters of a risky prospect. Am J Psychol 107:499–518