Effects of Sequential Context on Judgments and Decisions in the Prisoner's Dilemma Game

Ivaylo Vlaev (ivaylo.vlaev@psy.ox.ac.uk)
Department of Experimental Psychology, University of Oxford, Oxford, OX1 3UD, UK

Nick Chater (nick.chater@warwick.ac.uk)
Department of Psychology, University of Warwick, Coventry, CV47AL, UK

Abstract

Existing models of interactive game-theoretic decision making typically assume that only the attributes of the game need be considered when reaching a decision, i.e., these theories assume that the utility of a strategy is determined by the utility of the outcomes of the game, and transforms of the probabilities of each outcome. The strategic decisions are assumed to be based on these utilities. The two experiments presented here provide strong evidence against these assumptions. We investigated choice and predictions about the choices of other players in Prisoners Dilemma game. The cooperativeness of the games in each condition was varied and the results demonstrate that the average cooperation rate and the predicted cooperation of the other player in each game strongly depended on the cooperativeness of the preceding games, which suggests that games are not considered independently. It is proposed that people have poor notions of absolute cooperativeness, risk, and utility, and instead make their judgments and decisions in relative terms. The proposed accounts for these results are based on existing psychophysical and cognitive theories of perception and judgment of magnitude information.

Introduction

Understanding how people predict others' behavior and make choices on the basis of these prediction and the available opportunities and rewards is a central question for psychology; but also how people trade off risk and return when interacting with other people is a central issue for economics because the foundations of economic theory are rooted in models of interactive decision-making, in which the outcome of a situation depends also on the decisions of others agents (the latter known as game theory). In economics and part of psychology the starting point for investigating how people make decisions under uncertainty has not been empirical data on human behavior. Instead, the starting point has been a normative theory of decision making, expected utility theory and game theory (first axiomatized by von Neumann & Morgenstern, 1947) which specifies how people ought to make decisions, and which plays a key role in theories of rational choice (for a review see Shafir & LeBoeuf, 2002). The assumption has then been that people make rational decisions as the theory prescribes, that is, expected utility theory and game theory can be viewed as a descriptive, as well as a normative, theories of human behavior. At the core of expected utility theory is the assumption that people make choices that maximize their utility, and they value a risky option or a strategy by the expected utility (in a probabilistic sense of expectation) that it will provide. In game theory, which is based on expected utility theory, the utility from a strategy is judged on the basis of the sum of each individual outcomes weighted by their respective probability (e.g., the agents beliefs about the other players distribution of choices). In these models the basic assumption is that each prospect or game is considered separately and the resulting choice should be based only on the attributes of the particular prospect or game.

Stewart et al. (in press) have already described the phenomenon of "prospect relativity": that the perceived value of a risky prospect (e.g., "p chance of x") is relative to other prospects with which it is presented. This is counter to expected utility theory, according to which the perceived value of each prospect should be dependent only on its own attributes. Stewart et al. suggest that this phenomenon arises in the representation of the magnitudes that define the prospects, and suggest that the phenomenon has a common origin with related effects in the perception of sensory magnitudes (Garner, 1954; Lockhead, 1995, Laming, 1997). Importantly, Stewart et al (in press) found that only the simultaneously considered choice options affect the decisions about risky prospects, without finding evidence for sequential effects. Birnbaum (1992), however, demonstrated that skewing the distribution of certainty equivalents (CE) offered for simple gambles, whilst holding the range constant, influenced the selection of a CE. When the CE options were positively skewed (i.e., more small values) gambles were overvalued compared to the negatively skewed context, consistent with range-frequency theory (Parducci, 1974).

None of the existing studies however have investigated whether these context effects also hold during choice under uncertainty in the context of interactive decision making (modeled in economics as game playing). It would of course be difficult investigate the effects of the immediately available games in the context of interactive game playing because games are always played in a sequence. Therefore we decided to focus on testing the sequential effects. In particular, this study investigates the possibility that the previous context influences decisions in prisoner's dilemma game, in which players have to decide whether to cooperate or not in order to obtain certain payoffs and also to predict their co-player's decisions. The context was assumed to be the cooperativeness of the previously played games in the sequence while the dependent variables that this context was expected to affect were the cooperation rate of the participants and their estimate of the probability that the other player will cooperate.

The cooperation index scale

The table below illustrates the structure of the Prisoner's Dilemma Game, which is the subject of vast literature in economics, behavioral decision making, and cognitive and social psychology.

		Player 2	
		Cooperate	*Defect*
Player 1	*Cooperate*	C, C	S, T
	Defect	T, S	D, D

The game is defined by the chain of inequalities T >C >D >S, where C is the payoff if both cooperate (C), D is the payoff if both defect (D), T is the payoff if one defects and the other cooperates (and its called the temptation payoff), S is the payoff if one cooperates and the other defects (and this payoff is also called the sucker payoff). For example, this inequality could be represented by the numbers 20 > 10 > 8 >0.

In order to manipulate the cooperativeness of the games in each session we used a measure developed by Rapoport (1965) who investigated whether certain structural properties of the game will affect people's propensity to cooperate. He derived a cooperation index for predicting the probability that people will cooperate defined by the ratio:

K = (C - D) / (T - S)

This index varies from 0.1 to 0.9 (for simplicity denoted in the rest of the text with the integers 1 to 9). A game with index 1 could defined by the inequality 20 > 10 > 8 >0, and is a very non-cooperative one because it is characterized by high temptation to defect – offering a potential increase from 10 to 20 points and a low potential loss if both defect – from 10 to 8 points; while a game with index 9, which could be defined by the inequality 20 > 19 > 1 > 0, is very cooperative one because it is characterized with a low relative gain from defection – from 19 to 20 points, and a high potential loss if both defect – from 19 points to 1 point. In a seminal study Rapoport (1965) demonstrated a linear relationship between the cooperation index and the cooperation rate, i.e., people tend to cooperate more when playing games with a higher index.

In the two experiments presented here various groups played in conditions with different distribution of the games along the cooperation index in order to test whether participants' cooperation rate and their predicted cooperation of the other player in each game strongly depend on the cooperativeness of the preceding games. Thus, the present research departs fundamentally from previous work in game theoretic decision making by trying to model the highly flexible and contextually variable way in which people represent magnitudes such as payoffs and probabilities, rather than assuming that these can be represented on 'absolute' internal psychological scales.

Experiment 1

In this experiment, we manipulated the skewing of the distribution of the values of the cooperativeness index positive or negative in each session, whilst holding its range constant. There was also a control condition in which the values were equally distributed along the whole scale. If the participant were not affected by the context provided by the distribution of the cooperativeness of the games in each condition, then the responses to any given game in the skewed conditions would be expected to be similar to the control condition.

Method

Participants were 96 students recruited from the University of Oxford student population via a mailing list of people who have asked to be contacted participated in this experiment. 48 took part in the control condition, 24 participated in the first context condition, and 24 in the second context condition. Each session took no more than 60 min, and they were paid in cash at the end of it depending on their performance. Participants were paid a £2 fixed fee, and could have gained up to £7 in total, with an estimated average of £6.

Design There were three conditions in this experiment. The first aimed to test how people behave when they play the whole range of the cooperation index from 1 to 9. In the second condition the participants played games with index 1, 2, 3, 4, 5, 9, while in the third condition they played games with index 1, 5, 6, 7, 8, 9. In addition, there were three basic version for each condition of the experiment.

• Hypothetical play, in which the participants had to make decisions and judgments without real interaction. In this setting they made judgments and decisions by just imagining what would they would do if playing against a real opponent. This condition aimed to test whether people would exhibit the same context effects as in the interactive version of the game, which would indicate that these effects are due to the perceptual attributes of the games rather to the dynamics of the particular interaction in the group.

• Interactive play against anonymous player from the group. In this condition, the participants were informed that on each round of the game they would play against a randomly selected player from a group of between four to eight people, and therefore it is impossible to infer the strategy of the other player from the history of the game. This random matching was intended to control for possible learning affects of group dynamics during the play (e.g., a particular group might settle into particularly cooperative or uncooperative modes thus biasing our data).

• Play against the computer. In this condition the computer was pre-programmed to respond with a strategy reflecting the values of the cooperation index of each game, which means that when for example the computer plays games with index .5 the computer cooperates in 50% of the rounds this game is played. This condition was designed to test whether the context effects can overpower the effects of the automated response consistent with the value of the cooperation index, which would lead participants' responses to more in line with the values of the index.

Each subject participated in only one experimental design (hypothetical play, interactive play, or play against the computer) and in only one condition of the experiment, i.e., there were 8 participants in each design-condition pairing.

Procedure. The participants were informed that they would play 56 rounds of the game and on each round of the game they would play against a randomly selected player from their group (usually between four and eight people were in each group). This random matching makes it impossible to infer the strategy of the other player from the history of the game and thus controls for possible learning affects during the play. Thus we also aimed to prevent people from building a model of their opponent, which is another significant contextual factor that has been shown to affect strategic behavior (e.g., Pruitt and Kimmel, 1977 for a review).

Each condition consisted of a sequence of rounds of prisoner's dilemma game in which players make their choices simultaneously. The games were presented in a random order and the payoffs differed on each round depending on the condition of the experiment. In order to control for the effects related to the absolute magnitude of the received payoff from each round, the payoffs of each game were multiplied by the factors 1, 4, 7 and 10, so there were four versions of each game index in terms of the absolute magnitudes of the payoffs.

On each round of the game the participants were presented with a matrix of the game on the computer screen and they had to make a judgment and a decision. We used the abstract label 1 to denote the cooperative response and 2 for the uncooperative one in order not to prime certain social values in the group, which might induce certain strategies that could additionally bias the results. The judgment was to state how probable they think it is that the other will play 1 in this game. In order to make an estimation they had to move the slider on the screen, using the pointer of your mouse, to the position between 0 and 100%, which reflected their subjective prediction of the probability (likelihood) that the other player will choose to play '1' in the current round. They were awarded points for the accuracy of these predictions. Then the participants had to choose their decision strategy (1 or 2), and after both players in each pair have made their decisions the round ends and they were informed on the screen about the decision made by the other player, and their received payoffs from the game and from the accuracy of their predictions. The participants were paid for their participation in cash according to their performance. At the end of the experiment the accumulated score in points was transferred into cash according to an exchange rate, i.e., the experiment was conducted incentive compatibly.

Results

The results from the control condition are presented in Figure 1. In general, participants demonstrated linearly increasing cooperation rate and predicted cooperation as the index value increased, which replicated Rapoport's (1965) finding that mean cooperation increase as the index increases.

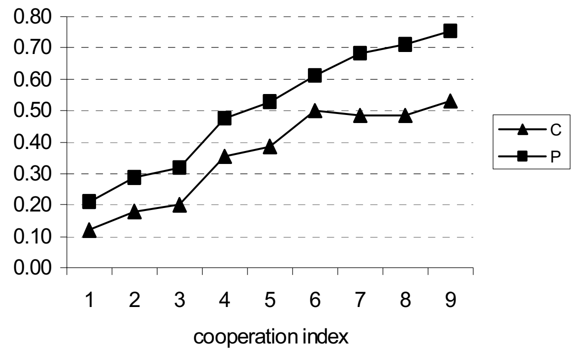

Figure 1: Mean cooperation rates (C) and mean predictions (P) for each value of the cooperation

The mean cooperation rates for the two context conditions of the interactive design are shown in Figure 2, and for the mean prediction in Figure 3.

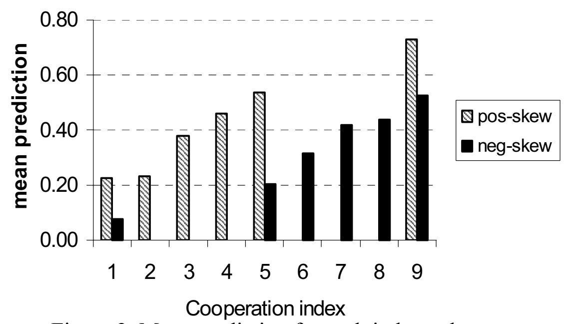

Figure 2: Mean prediction for each index value in the two contextual conditions with interactive

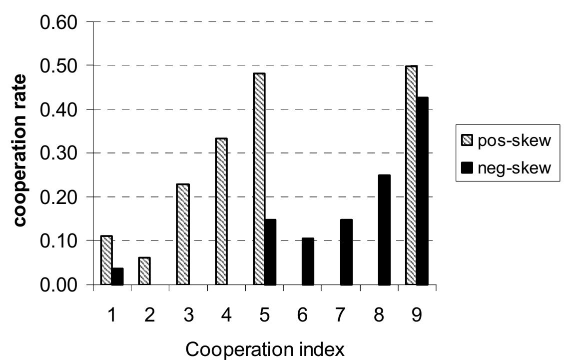

Figure 3: Mean cooperation rate for each index value in the two contextual conditions with interactive play

These results demonstrate that games of index 5 that have a higher rank in the positively skewed distribution (condition 1) had a significantly higher cooperation rate and higher mean predictions compared with the games with a lower rank (condition 2). Another interesting result is the almost identical pattern between the results for games with index 1 to 5 in the positive skew and the results for the games with index 5 to 9 in the negative skew, which do not increase as in the control condition. The same pattern appeared in the version involving only hypothetical play. This result indicates that human judgments and decisions are relative to the other values in the context rather than being represented on an absolute scale.

The mean cooperation rate and mean prediction values for games with index 5 in the positively and negatively skewed distributions for each of the three designs are shown in the Table 1 below.

Table 1. Summary results and statistics for game 5.

Design	Task	Pos	Neg	t	p
Hypothetical	Cooperation	0.60	0.35	2.28	.021
	Prediction	0.60	0.51	2.10	.049
Interactive	Cooperation	0.48	0.17	3.19	.002
	Prediction	0.54	0.24	4.67	.0001
Computer	Cooperation	0.24	0.27	.42	.671
	Prediction	0.54	0.53	.10	.915

The shaded cells indicate the conditions between which there was a significant difference while the last two columns present the *t* tests of the difference between the two context conditions. In four out of six comparisons, during the hypothetical and interactive play, the games with a higher rank in the distribution (condition 1) had higher significantly higher cooperation rate and higher mean predictions compared with the games with a lower rank (condition 2), which indicates that the interaction was not essential to produce these effects. In the condition of play against the computer there is no difference between the groups suggesting that people are very sensitive to the automated response that enables them to form an absolute judgment of the cooperativeness of the games and the context effects are not strong enough to counter this effect.

Discussion

The cooperation rate and the predicted cooperation were strongly influenced by the preceding games in the sequence, and participants' behavior in games with the same level of cooperativeness differed significantly between the two context groups. In particular, the games with index 5 had higher values in the first group in which these games had a higher rank (fifth) compared with the games having the same index in the second group but being second in rank. The same effects were observed when people made only hypothetical decisions without real interaction.

These results clearly demonstrate that the rank of a game in the distribution (stimulus set) will significantly affect participants' perception of the cooperativeness of that game and they will be more likely to cooperate and will predict higher cooperation when the game has a higher rank. There are well-known theories of context effect in perceptual judgment and general models of decision making under risk, which explicitly model and empirically demonstrate the effects of the ranking of the stimuli in the distribution like for example Parducci's range-frequency theory (1974) and Tversky & Kahneman's prospect theory (1992).

Experiment 2

The first experiment showed that games are judged relative to previously played games. In this experiment we decided to test whether these context effects are emerging because people compare and contrast the current stimulus (game) only with the previous one, or also whether these effects can appear on a larger scale when there is (implicit) comparison between two separate sequences of games. For this study we decided to use only two types of games, very cooperative ones with index 9 and very uncooperative ones with index 1, and we tested whether order in which people played these games was affecting their choices and predictions.

In order to further accentuate the difference between the two game types we decided to design the uncooperative games with negative payoffs. Thus we expected to provoke bigger perceptual and reinforcement dissociation between the two games (although their strategic structure is identical), which could further enforce the contrast between them. There were three order conditions in this study. One in which the two game types were mixed and appeared in random order in the sequence and thus were contrasted on a trial by trial basis; second in which the cooperative games were played first and then the uncooperative ones; and third condition where the uncooperative games were played first and then the cooperative ones.

Method

Participants 20 participants took part in each condition of this study (60 in total). All participants were paid £3 plus performance related winnings of up to £3.

Design The cooperative games were similar to the ones in the first experiment. In order to make the magnitudes of the payoffs in the two game types comparable and of equal importance for the players, we decided that the low cooperation game should offer similar absolute amounts as the cooperative game, but the payoffs will be negative. The following matrix game is an example of an uncooperative prisoners dilemma game with index 1, which offers negative payoffs.

		Other *1*	Other *2*
You	*1*	-5, -5	-11, 0
	2	0, -11	-6, -6

The payoffs of each game were multiplied by the factors 1, 4, 7 and 10, as Experiment 1, so there were four versions of each game index in terms of the absolute magnitudes of the payoffs.

Procedure The experiment was done interactively in groups containing 4 to 8 participants and the procedure was the same as the one in the first experiment. The only difference is that the participants played 96 rounds in the session. In the conditions with separate sequence for the positive and the negative games they played each game type separately for 48 rounds, while in the mixed condition the two game types were presented randomly in the sequence.

Results

In order to see how each variable changed over time we took three types of running average over time. The first was over the whole sequence of trials up to the current trial, i.e., averaged over the number of trials so far, the second also running average over the last 5 trials, and the third one was over the last 10 trials. Since the three measures produced almost identical results here we report only the data from the running average estimated over the whole sequence of trials. We calculated separate running average time series for the positive and the negative games, in order to estimate the trends for each game type independently of each other (otherwise, if we take the running average to include both games, then the effects will be confounded between the positive and the negative games).

Cooperation rate Figure 4 presents evidence that in the negative-positive condition the time series results for the running average of the cooperation rate change abruptly in the second half of the session after the participants started to play the positive games (first playing the negative ones), jumping instantly from just below .2 to above .5. By contrast, in the positive negative condition the cooperation rate for the positive games (played first) is always below .4 and in the second half, when people start playing the negative games drops down to the level of the other negative-positive condition, which is below .2. This is clear evidence of order effects on the behavior in the positive games.

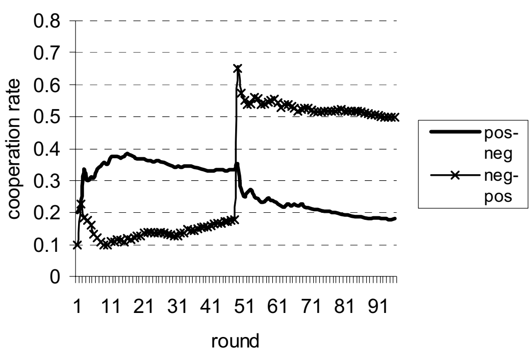

Figure 4: Running average of the cooperation rate in the two conditions with separate sequence for the positive and the negative games

When we compare these results with the outcome from the condition with mixed games presented in Figure 5, it immediately evident that the score for the negative games is very similar and again just below .2, while the cooperation rate for the positive games reaches almost .8 (80%) and stabilizes. It appears that in the mixed condition the contrast was even stronger than in the negative-positive condition (exceeding it with more than 20%), which was possibly due to the more immediate effect of the negative games in the mixed condition. Another characteristic feature of the mixed condition is that the cooperation is constant over the whole session, while in the other two conditions it drops down as in most experimental findings with prisoner's dilemma game.

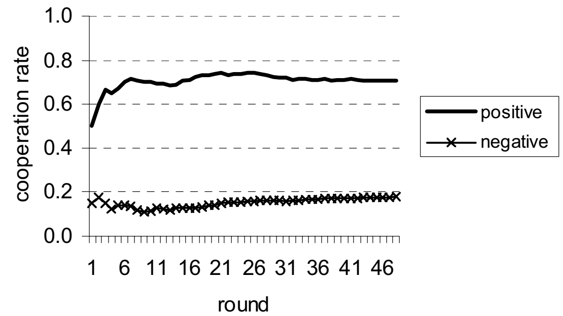

Figure 5: Running average of the cooperation rate in the condition with mixed games

The cooperation rates were averaged for every subject separately over each game index and were analyzed in a repeated-measures analysis of variance with game type (positive vs. negative) as a within-subjects factor and the experimental condition (*pos-neg* vs. *neg-pos* vs. *mixed*) as a between-subjects factor. The interaction between the game type and the experimental condition was found to be statistically significant, $F(2,57)=37.77$, $p<.001$. Pair-wise comparison post-hoc test was used between the experimental groups for each game type. We found a significant difference in the cooperation rate between the each of the sequence conditions and the mixed condition for the positive games ($F(2,57)=23.34$, $p<.001$) without however finding significant difference between the two sequence conditions. For the negative games the difference was significant between all groups ($F(2, 57)=4.53$, $p=.01$) except between the *neg-pos* condition and the *mixed* condition. The main effect of the game type was also significant, $F(1,57)=36.09$, $p<0.001$, i.e., the participants cooperated more in the positive games.

Prediction For this dependent variable, the results for two sequential conditions were very similar to the ones for the cooperation rate (Figure 6). Again, in the neg-pos condition there was a sharp increase with almost 40% in mean prediction immediately after the start of the second half of the session when participants started to play the positive games, which is again above the level for the positive games in the other condition. At the same time the level for the negative games in the pos-neg condition decreases sharply with more than 20% to the level for this game in the neg-pos condition.

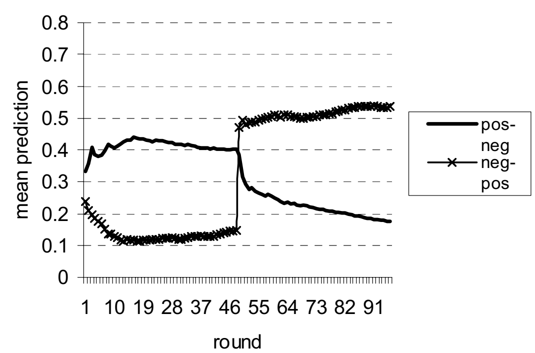

Figure 6: Running average of the cooperation rate in the two conditions with separate sequences for the positive and the negative games

The interesting result however was found in the mixed condition shown in Figure 7, in which the prediction levels for both game types increased steadily during the whole session, reaching .8 for the positive games (as opposed to .4 and .5 in the other conditions), and .6 for the negative games (as opposed to .1 and .2 in the sequential conditions). These results are very strong evidence for the powerful effects of the context provided by the other games in the sequence.

The prediction levels were averaged for every participant separately over each game index, and were analyzed in a repeated-measures analysis of variance with game type (positive vs. negative) as a within-subjects factor and the experimental condition (*pos-neg* vs. *neg-pos* vs. *mixed*) as a between-subjects factor. The interaction between the game type and the experimental condition was statistically significant, $F(2, 57)=4.32$, $p = 0.018$. In order to compare the experimental groups for each game type, a post-hoc test was used. We found a significant difference in the mean prediction between the three experimental groups for positive games ($F(2,57)=35.41$, $p<.001$), and for the negative games ($F(2,57)=12.08$, $p<.001$). The difference between the *pos-neg* condition and the *mixed* condition was not significant. The main effect of the game type was also significant, $F(1,57)=36.09$, $p < .001$. Participants made significantly higher predictions in the positive games.

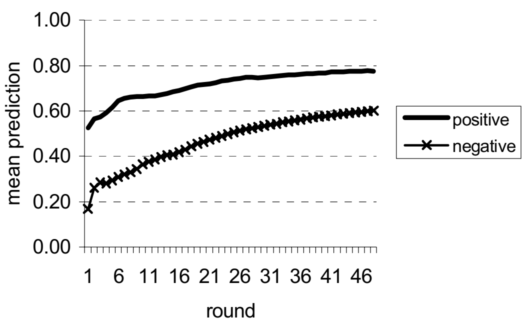

Figure 7: Running average of the mean prediction in the condition with mixed games

Discussion

The results clearly show sequential context effects on judgment and choice in strategic game playing, which were significant in all conditions, but particularly increased in the mixed condition where the contrasting between the two game types interleaved on a trial by trial basis. When the games were divided in two separate sequences then the negative games in the first half of the session produced powerful contrast effect with the positive games in the second half---the latter is considered even more cooperative than normally. There was also a significant contrast effect in the opposite direction (if a change is made from a cooperative game to a non-cooperative one, but the effect was stronger when the transition was from negative to positive games than from positive to negative. This asymmetry may arise because of the way our incentive scheme works. We pay participants according to their accumulated payoff in the whole session. Thus when the negative game was played first, participants initially lose points; and when the positive games started they might have immediately seized the opportunity to compensate that loss by looking for mutual cooperation (otherwise they risk ending up with no payment if their final payoff was negative). When there was transition from positive to negative games, however, there was probably less incentive to shift immediately to non-cooperation because the accumulated positive payoff was not perceived to be at risk.

General Discussion

Our study investigates sequential context effects on strategic choice in game playing in combination with elicitation of players' beliefs about her opponent's decisions. We demonstrated that the attributes of the previous games influence the judgments and decisions in the current game. This result implies that games are not considered independently of the previously played ones. Thus the present research demonstrates a new and large anomaly for expected utility theory and game theory in decision-making under uncertainty. Specifically, the reported results seem to indicate that people do not possess a well-defined notion of the utility of a strategy and the cooperativeness of a game, and hence they do not view such utilities and games in terms of expected utility. Instead, people's perceived utility for a strategy appears highly context sensitive and it depends on the other games being played.

References

Birnbaum, M. H. (1992). Violations of monotonicity and context effects in choice-based certainty equivalents. *Psychological Science, 3,* 310-314.

Garner, W. R. (1954). Context effects and the validity of loudness scales. *Journal of Experimental Psychology, 48,* 218-224.

Laming, D. R. J. (1997). *The measurement of sensation.* London: Oxford University Press.

Lockhead, G. (1995). Psychophysical scaling method reveal and measure context effects. *Behavioral and Brain Sciences. Vol 18(3)*: 607-612.

Parducci, A. (1974). Contextual effects: A range-frequency analysis. In L. Carterette and M. P. Friedman (Eds.), *Handbook of Perception* (Vol. II). New York: Academic Press.

Pruitt, D. G. & Kimmel, M. J. (1977). Twenty years of experimental gaming: critique, synthesis, and suggestions for the future. Annual Review of Psychology, 28:363-92.

Rapoport, A. (1965). Prisoner's dilemma: a study in conflict and cooperation. Ann Arbor: University of Michigan Press.

Shafir, E., & LeBoeuf, R. A. (2002). Rationality. *Annual Review of Psychology, 53,* 491-517.

Stewart et al. (in press). Prospect Relativity: How Choice Options Influence Decision Under Risk. *Journal of Experimental Psychology: General.*

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*, 297-323.

von Neumann, M., & Morgenstern, O. (1947). *Theory of games and economic behavior. (2nd ed.).* Princeton, NJ: Princeton University Press.