

Context effects in games: Local versus global sequential effects on choice in the prisoner's dilemma game

Ivo Vlaev* and Nick Chater

Department of Psychology

University College London

Abstract

We report an experiment exploring sequential context effects on strategy choices in one-shot Prisoner's Dilemma (PD) game. Rapoport and Chamah (1965) have shown that some PDs are cooperative and lead to high cooperation rate, whereas others are uncooperative. Participants played very cooperative and very uncooperative games, against anonymous partners. The order in which these games were played affected their cooperation rate by producing perceptual contrast, which appeared only between the trials, but not between two separate sequences of games. These findings suggest that people may not have stable perceptions of absolute cooperativeness. Instead, they judge the cooperativeness of each fresh game only in relation to the previous game. The observed effects suggest that the principles underlying judgments about highly abstract magnitudes such as cooperativeness may be similar to principles governing the perception of sensory magnitudes.

Keywords: decision making; cooperation; prisoner's dilemma; context effects.

1 Introduction

Most applications of standard normative models, whether of individual or strategic decision making, make the basic assumption that each risky prospect or game is considered separately and the resulting choice should be based only on the attributes of the particular prospect or game (Fudenberg & Tirole, 1991; Kreps, 1990). The validity of this, and related, independence assumptions has been challenged extensively in the past especially in the context of individual decision making under risk. Thus, Allais (1953) first demonstrated behavior violating the independence axiom of expected utility theory (here the independence is between mutually exclusive possible outcomes). Later on, regret theory (Loomes & Sugden, 1982) showed how regret can modify the utility of an outcome that results from a particular choice depending on the outcomes that would have resulted from other choices in the choice set. Recent psychological theories of individual decision making have also

been developed, in which prospects are judged in relation to one another, such as the stochastic difference model (González-Vallejo, 2002), multialternative decision field theory (Roe, Busemeyer, & Townsend, 2001), the componential-context model (Tversky & Simonson, 1993) and decision-by-sampling (Stewart, Chater, & Brown, 2006). These theories all have in common the idea the mere presence of an option in a choice set may change the way another option is judged; or, more broadly, that preferences are constructed afresh in the light of the salient options in each situation or the recent past. Thus, preference is *constructed* rather than *revealed* (see Slovic, 1995).

In an attempt to ground this constructivist idea onto some fundamental properties of the perceptual system, Stewart, Chater, Stott, and Reimers (2003) describe a phenomenon called *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. Similar effects were also found in financial (saving and investment) decision making under risk (Vlaev, Chater, & Stewart, 2007a; 2007b). These prospect relativity effects are counter to expected utility theory, the basic normative principle for individual choice, which assumes that the perceived value of each prospect should be dependent only on its own attributes. Stewart et al. (2003) suggested that this phenomenon arises because of way in which the magnitudes that define the prospects are de-

*This work was conducted while Ivo Vlaev was a doctoral student at the Department of Experimental Psychology, University of Oxford. We thank Neil Stewart for his comments and Daniel Zizzo for providing the experimental economics laboratory at the Department of Economics, University of Oxford. We would also like to thank two anonymous referees and especially the editor of this journal, Jonathan Baron, for working out the details of the regression analyses and other invaluable suggestions. Addressed: Ivo Vlaev, Department of Psychology, University College London, 26 Bedford Way, London, WC1H 0AP, United Kingdom. Email: i.vlaev@ucl.ac.uk, n.chater@ucl.ac.uk

fined, and that the phenomenon has a common origin with similar psychophysical effects on perception of sensory magnitudes like brightness and loudness (Garner, 1954; Laming, 1997; Lockhead, 1995).

1.1 Sequential vs. simultaneous context

Stewart et al. (2003) also found that only the simultaneously considered choice options affect the decisions about risky prospects, without finding evidence for sequential effects. Birnbaum (1992), however, demonstrated such sequential effects by showing that skewing the distribution of certainty equivalents offered for simple gambles, whilst holding the range constant, influenced the selection of a certainty equivalent. (In making a certainty equivalent judgment, participants select from a set of options the amount of money for certain that is worth the same to them as a single chance to play the prospect.) When the certainty equivalent options were positively skewed (i.e., more small values), gambles were overvalued compared to the negatively skewed context, consistent with range-frequency theory (Parducci, 1965, 1974).

There has been a wealth of publications on various other types of sequential effects on decision making. In repeated games, Rapoport, Stein, Parco, and Nicholas (2003) showed the effect of the outcomes of whole series of previous games. Knez and Camerer (2000) showed that cooperation in social dilemmas can be increased by preceding play of coordination games. Thus they demonstrated that cooperation in one game spills over into cooperation in Prisoners' Dilemma game. Mellers, Schwartz, Ho, and Ritov (1997) report the effect of series of gains vs. series of losses vs. a mixture. Others have exhaustively analyzed the dynamics of learning over series of games (Erev & Roth, 1998; Camerer & Ho, 1999). Thaler, Tversky, Kahneman, and Schwartz (1997) have clearly shown the impact of *trends* (series of gains vs. series of losses) in investors' decision-making. For example, Loewenstein and Thaler (1989), Loewenstein and Prelec (1992), and Frederick, Loewenstein, and O'Donoghue (2002) have carried out extensive research on choice over time, showing large and complex effects of short-term and long-term time-contexts. Other studies have also shown that people frequently cooperate in scenarios which appear to have a one-shot Prisoners' Dilemma structure (see, e.g., Dawes & Thaler, 1988).

However, none of these studies on history dependence (apart from Birnbaum, 1992, and Stewart et al., 2003) used theoretical accounts based on some theory of the fundamental perceptual mechanisms. Also, few existing studies have investigated whether perceptual context effects like prospect relativity also hold during choice under uncertainty in the context of interactive (strategic) deci-

sion making. Therefore, we focus here on testing only (perceptual types of) sequential effects in game playing (which are typically played one at a time and hence sequential in nature). We were particularly interested in investigating whether the sequential context effects in games are similar to the context effects observed in perceptual judgment tasks like magnitude estimation (Stewart, Brown, & Chater, 2002; Ward & Lockhead, 1970, 1971).

In previous research, we have already found that the attributes of the previously seen games influenced the decisions in the current game (Vlaev & Chater, 2006). In particular, we investigated choice and predictions about the choices of other players in PD game during one-shot plays with different, anonymous opponents on each trial. Participants played a sequence of games with varying degrees of cooperativeness, as measured by Rapoport and Chammah's (1965) Cooperation Index (described below). 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 range and the skew of the distribution of the cooperativeness of the preceding games. That is, we found that the representation of the cooperativeness of the current game depends on the distribution of cooperation indices of games that the player has previously played. In particular, people tend to contrast the current game with the other games in the sequence depending on the position of the current game in the range, and also the rank order of the previous games on the cooperation index scale.

1.2 Local vs. global context

In the study reported here, we again investigated how the cooperativeness of previous games influenced decisions in PD game, but this time we compared the impact of *global* vs. *local* sequential context. That is, suppose that a person plays a long sequence of highly uncooperative games. During this sequence, players are likely to defect. How will they react on suddenly encountering more cooperative games? If past history is the main factor determining behavior, they might be expected to continue defect and to expect defection of their opponents. A past history of uncooperative games may have created "cynicism" in the face of a more cooperative game. However, if cooperativeness is represented in the same way as sensory magnitudes, the opposite possibility arises — that the new more cooperative games are viewed as "especially" cooperative, in contrast to the prior uncooperative games. Hence, players might be expected to cooperate more. Thus, instead of cynicism, the player is "grateful for small mercies," after a history of uncooperative games.

		Player 2	
		<i>Cooperate</i>	<i>Defect</i>
Player 1	<i>Cooperate</i>	C,C	S,T
	<i>Defect</i>	T,S	D,D

Figure 1: Model of the Prisoner's Dilemma Game.

Our previous work (Vlaev & Chater, 2006) has shown that such contrast effects can occur, when games of different levels of cooperativeness are mixed together. But can the effect of contrast (i.e., viewing a moderately cooperative game as more cooperative, given a history of less cooperative games) overcome the cynicism that might be expected to follow from a history of negative feedback? A further question concerns how the effects of perceptual contrast affect an entire sequence of future games; our previous work focussed on the impact of context on a single game. If effects of past games are highly local, then although we might observe context effects when there is a switch between, say, less and more cooperative games, this effect will rapidly disappear. If, though, context effects are long-lasting, then the cooperativeness of previous games may have a substantial influence on a long sequence of future play. Note that a substantial influence here constitutes a global contrast effect that is larger than the local contrast effect, or, at least a statistically significant effect on a whole sequence of future play (i.e., in terms of cooperation rate).

With psychophysical stimuli, local effects appear strongest — indeed a good deal of performance in absolute judgment tasks can be explained with reference to only the previous two stimuli as shown recently by Stewart, Brown and Chater (2005). In particular, previous psychophysical models assume that identification is achieved using long-term representations of absolute magnitudes. Stewart, Brown and Chater (2005) propose an alternative relative judgment model (RJM) in which the elemental perceptual units are representations of the differences between current and previous stimuli. These differences are used to respond without using long-term representations of absolute magnitudes. The logic of RJM is based on two main assumptions. The first assumption is that judgment is relative and not absolute, which is based on abundant evidence that long-term representation of attributes such as pitch and loudness may be very poor (see Stewart, Brown, & Chater, 2005, for an extensive review of this evidence). Models that use long-term representations of absolute magnitudes of stimulus values (as either exemplars, anchors, or criteria) do not capture the sequential effects adequately. In the RJM, judgment is instead relative to the immediately preceding stimulus. The second assumption is that the locus of the limit in performance

observed in magnitude estimation tasks when the range of the stimuli is increased, or when stimuli are presented in a larger set versus in isolation, is not perceptual but judgmental. In particular, models of absolute identification, which assume that the locus of the limit in information transmitted is perceptual, fail to predict (or require modification to predict) that channel capacity will remain severely limited even for very large stimulus spacings. In the RJM, the limit in channel capacity is not perceptual. Stewart, Brown, and Chater (2005) show that using difference information optimally within a limited capacity provides an account of the limit in information transmitted. In summary, in the assumed absence of stable, long-term absolute magnitudes, the representation of the difference between the stimulus on the current trial and the stimulus on the preceding trial is used in conjunction with the feedback from the previous trial to produce a response.

This new theoretical framework has yet to be explored in the domain of strategic games. If the principles underlying judgments about abstract magnitudes such as “cooperativeness” are similar to principles governing the perception of sensory magnitudes (as demonstrated by Vlaev & Chater, 2006), then it is very likely that the predictions of the relative judgment model will also hold for strategic decision making in games. In particular, we do not expect to find significant long lasting (global) effects between whole sequences of games (i.e., a lack of lasting effects of stimuli across longer time periods). Instead we expect to find only local context effects between current and previous games in a sequence (conforming to the RJM proposed by Stewart, Brown & Chater, 2005).

In this experiment, similarly to the Vlaev and Chater (2006) study, the stimulus context was assumed to be the cooperativeness of the previously played games in the sequence, while the dependent variable was the participants' cooperation rate. Games' cooperativeness was expected to be judged similarly to other perceptual attributes such as pitch and loudness.

1.3 The cooperation index scale

Figure 1 illustrates the structure of PD, which is the subject of a vast literature in economics, behavioral decision making, and cognitive and social psychology. The game is defined by the chain of inequalities $T > C > D > S$, where C is the payoff if both cooperate, D is the payoff if both defect, T is the payoff if one player defects and the other cooperates (often called the *temptation* payoff), S is the payoff if one player cooperates and the other defects (and this payoff is often called the *sucker* payoff).

In order to manipulate the cooperativeness of the games in each session we used a measure developed by Rapoport and Chammah (1965), who investigated

		Other	
		Cooperate	Defect
You	Cooperate	10,10	0,11
	Defect	11,0	1,1

Figure 2: Game matrix of a Prisoner’s Dilemma game with a cooperation index .8.

		Other	
		Cooperate	Defect
You	Cooperate	-5,-5	-11,0
	Defect	0,-11	-6,-6

Figure 3: Game matrix of a Prisoner’s Dilemma game with a cooperation index .1.

whether certain structural properties of the game will affect people’s propensity to cooperate. They derived a *cooperation index* (CI) for predicting the probability that people will cooperate, defined by the ratio: $(C - D)/(T - S)$.

In principle, the CI ranges from 0 to 1, where CIs close to 1 characterise games in which cooperation is probable; and CIs close to 0 characterise games in which defection is probable. For example, Figure 2 shows a game with a high cooperation index for which $T = 11 > C = 10 > D = 1 > S = 0$ (see Figure 2). Thus, CI is $(10-1)/11-0 = 9/11$, which is approximately .8. By contrast, Figure 3 shows an uncooperative game, in which $T = 0 > C = -5 > D = -6 > S = -11$. Thus, CI is $(-5-(-6))/(-5-(-11)) = 1/11$, which is approximately .1. In a seminal experimental study, Rapoport and Chammah (1965) demonstrated a roughly linear relationship between the cooperation index and the cooperation rate: people tended to cooperate more when playing games with a higher index.

In the experiment presented here, we used these two levels of the CI in order to test whether participants’ cooperation rate 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 cooperativeness (and payoffs and probabilities in this respect), rather than assuming that these attributes can be represented on stable internal psychological scales.

2 Experiment

In this experiment, we decided to test whether the sequential context effects caused by the previous games emerge

Table 1: Prisoner’s Dilemma games used in the study.

Game’s CI	Outcomes			
	CC	ST	TS	DD
.1	-5,-5	-11, 0	0,-11	-6,-6
.1 (x4)	-20,-20	-44, 0	0,-44	-24,-24
.1 (x7)	-35,-35	-77, 0	0,-77	-42,-42
.1 (x10)	-50,-50	-110, 0	0,-110	-60,-60
.8	10, 10	0, 11	11, 0	1, 1
.8 (x4)	40, 40	0, 44	44, 0	4, 4
.8 (x7)	70, 70	0, 77	77, 0	7, 7
.8 (x10)	100, 100	0, 110	110, 0	10, 10

Note. The payoffs in each cell of the table are as indicated in the cells of the game matrix shown in Figure 1.

because people compare and contrast the current game only with the previous one, or whether these effects can also appear on a larger scale when there is (implicit) comparison between two separate sequences of games. For this study, there were only two types of games, very cooperative ones with index .8 and very uncooperative ones with index .1, and we tested whether the order in which these games were played affected people’s choices and predictions.

In order to further accentuate the difference between the two game types, and hence to maximize the impact of context effects, the uncooperative games were given negative payoffs. Thus we expected to provoke stronger perceptual dissociation between the cooperative and uncooperative (positive and negative) games, which could further enforce the contrast between them (although their strategic structure is identical). There were three order conditions in this study. For one group, the cooperative (positive) games were played first and then the uncooperative (negative) ones; for the other, the uncooperative (negative) games were played first and then the cooperative (positive) ones. These two conditions were testing whether the perceptual context effects appear globally between two whole sequences of rounds. A third group of participants played a random mixture of the two game types where the perceptual context effects were expected to appear locally between two neighboring rounds, i.e., on a round-by-round basis as predicted by the relative judgment model (Stewart, Brown, & Chater, 2005). In the following sections, these ordering conditions are denoted in terms of the cooperativeness of the games in the sequence as High-Low, Low-High, and Mixed condition, respectively.

2.1 Method

2.1.1 Participants

Twenty participants took part in each of the three conditions (groups) of this study (so there were 60 participants in total) recruited from the University of Oxford student population via the experimental economics research group mailing list of people who have asked to be contacted participated in this experiment. All participants were paid £3 plus performance related winnings of up to £3.

2.1.2 Design

There were two types of PD game in this experiment – very cooperative game with cooperation index .8 and very uncooperative game with cooperation index .1. Figure 2 presents a game matrix of a very cooperative game with index .8.

In order to make the payoffs to be of comparable magnitudes in the two game types and therefore of equal importance for the players, we decided that the uncooperative game (offering negative payoffs) should offer similar absolute amounts as the cooperative game. Figure 3 presents the matrix of the uncooperative PD game with index .1, which offers negative payoffs (i.e., the strategic structure of the game remains unchanged).

In order to control for the effects related to the absolute magnitude of the received payoff from each round, the initial payoffs of each game, which were between 0 and 20, were multiplied by factors of 1, 4, 7 and 10, so finally there were four versions of each game index in terms of the magnitudes of the payoffs. Table 1 presents the eight games used in this study. Each game type (CI) is presented as a separate row in the table, and each cell presents the outcome payoffs for the two players. The cells and the payoffs were organized according to the cells of the abstract game matrix shown in Figure 1.

The cooperative games involved only positive payoffs and the uncooperative games included negative payoffs in order to make the distinction between the games very explicit, and thus to maximize any contrast effects that might affect participants' choices. In this respect, our goal was to demonstrate with this experiment the very existence and the power of the effect, which from a normative point of view should not exist. In other words, we aimed to demonstrate a *game type* effects without distinguishing those two factors (cooperativeness vs. payoff sign). Of course, other versions of the design could have cooperative games with negative payoffs and non-cooperative ones with positive payoffs, or all games to be either positive or all negative, and it would be interesting to explore these versions in future research. Consideration of these further cases, however, should not change

the nature of the argument we make here about the perceptual context sensitivity in strategic decision making that should, according to normative principles, be purely dependent on the current trial.

2.1.3 Procedure

The experiment was conducted interactively in separate groups for each condition. In the conditions with separate sequence for the positive and the negative games, the participants played each game type separately for 48 rounds (i.e., 96 rounds in total). In the Mixed condition, the game types were presented in a random order, which was the same for everyone. The participants were informed that they would play 96 rounds of the game and on each round of the game they would play against a randomly selected player from their group. This random matching aimed to make it impossible to infer the strategy of the other player from the history of the game, and thus to control for possible learning affects during the play. Thus we also aimed to prevent people from learning a model of their opponent, which is another significant contextual factor that has been shown to affect strategic behavior (see Pruitt & Kimmel, 1977, for a review).

Each condition consisted of a sequence of rounds of PD game in which players make their choices simultaneously. On each round of the game the participants were presented with a matrix of the game on the computer screen and they had to choose their decision strategy (1 or 2). 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. After both players in each pair have made their decisions the round ended and they were informed on the screen about the decision made by the other player, and their received payoffs from the game. Thus, 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.

2.2 Results

The cooperation rates were averaged for every participant separately over each game type in every condition and these averaged results are presented in Figure 4. In the Low-High condition the average cooperation rate changed in the second half of the session after the participants started to play the positive games (first playing the negative ones), jumping from 18% to 50%. In the High-Low condition, the cooperation rate for the positive games (played first) is 33% and drops down to 18% in the second half when people start to play the negative games.




Figure 4: Cooperation rate for each game type in the three experimental conditions. (Error bars are standard error of the mean.)

In the Mixed condition, the cooperation rate was 71% for the positive games and 18% for the negative ones. This result indicates that in all conditions the participants reacted to the change from negative to positive games and vice versa. The average cooperation in the positive games in the Mixed condition was significantly higher than the cooperation in the positive games in both the High-Low condition, $t(38) = 4.65$, $p < .0001$, and the Low-High condition, $t(38) = 2.31$, $p = .0264$. This result is evidence for the power of the contrast effect when the games are mixed between trials and hence the comparison is on trial by trial basis.

The cooperation rates 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 (*High-Low* vs. *Low-High*) as a between-subjects factor. There was a significant effect of the within-subject factor *game type* (indicating that there was significant change in cooperation from positive to negative games and vice versa), $F(1,38) = 16.5$, $p < .0001$. However, the interaction between the game type and the experimental condition (the between-subject factor) was not statistically significant, $F(1,38) = 2.12$, $p = .154$. If there is a significant global context effect between the two sequences of rounds (i.e., *Low-High* vs. *High-Low* condition), then there should be significant interaction effect due to effect of the condition on cooperation in the two game types. Note that a significant contrast effect on the positive games, which is implied by the higher cooperation in the positive games in the *Low-High*

condition, would also result in a significant interaction effect. The lack of statistical effect, however, suggests that, although people were sensitive to the cooperativeness of the two game types, their cooperativeness was not affected by the change from positive to negative and negative to positive games respectively. In other words, this is an indicator of a lack of global context effect.

When we added the Mixed condition as a third between-subject factor, then there was significant effect of both the game type, $F(1,57) = 55.5$, $p < .0001$, and the interaction between the game type and the experimental condition, $F(2,57) = 5.91$, $p = .0046$. This can only be explained by a very strong local contrast effect on the positive games in the Mixed condition (the negative games have the same cooperation level as in the other conditions), which is due to the local trial-to-trial context effect between the games in this condition. We also tested Mixed condition versus a combination of the other two conditions (i.e., “blocked”). Thus, instead of having three levels of the factor condition (H-L, L-H, and Mixed) we run an analysis with only two levels (H-L&L-H vs. Mixed). Again, there was significant effect of the interaction between the game type (H vs. L) and the experimental condition, $F(1,58) = 9.22$, $p = .0036$, which corroborated our initial conclusion.

If context effects exist and are only transitory, one should also observe a lack of global effect by analyzing the time series of choices. To test this hypothesis, we run two extra regression analyses of *trends* — one in the Mixed condition and second in the two other conditions (*Low-High* vs. *High-Low*).

2.2.1 Mixed condition trends

We looked directly at sequential effects in the Mixed condition by regressing, for each subject, the response to each round on the type of the previous round, the one before that, and so on. Finding greater coefficients for more recent rounds would support our conclusion. For example, round $n-3$ was expected to have little effect, because its effect would be diluted by rounds $n-2$ and $n-1$. We run a separate regression for each subject, and then averaged the coefficients across all participants in this condition. We used a one-sample t -test to test whether the average coefficients (for each factor) differ from zero (indicating the average size of the effect). However, the coefficients at t , $t-1$, and $t-2$ were not significantly higher than zero (inclusion of past choices as predictors did not change the results). This result suggests that there is no context effect within the mixed condition, while our model predicts negative (contrast) effect. However, another interpretation is that, because of the contrast between the two game types, the subjects quickly started to coordinate their responses in the first few rounds (in this condition). Note that PD can be seen as a coordination game. Any change in the environment that changes the ease with which coordination can occur can affect cooperation frequencies. It seems likely that coordination was easier in the Mixed condition. For example, the simple rule “cooperate in the positive payoff game, defect otherwise” could easily emerge. This rule cannot easily emerge in cases where only one type of game is repeatedly played. Thus, differences in ease of learning to coordinate between treatments could explain lack of effect of previous rounds on the current choice in the Mixed condition (while such coordination can still be caused by the local contrast in the first rounds of the session).

This “coordination” hypothesis would imply a positive interaction between the current game and trials. That is, the current game (at t) would have more effect as the subjects learned. To answer this question, we split the data into two halves (rounds 1–48, and rounds 49–96) and regressed the current choice only on the current game in each half. Higher coefficient for the current game in the second half would indicate that the subjects were learning to conditionalise more on the game (which imply less contrast effect of the previous game). Indeed, the average coefficients for current game in the first half was 9.98, which is lower than in the second half (18.4) by the factor of two and this difference was marginally significant, $t(19) = 2.00$, $p = .0604$. This result suggests that players learn to coordinate as the game progresses, which washes out the contrast effect with the previous game. This result could also explain the sustained strong contrast between the two game-types observed in this condition. (In this respect, we never claimed that there are no other effects

going on in this complex strategic interaction.)

The trend analysis in the Mixed condition also demonstrated that our experiment has the power to detect differences that we did not find within the conditions (conversely, the failure to find significant contrast effect here is not due to lack of statistical power, since the experiment did detect the context effect at the aggregate level of the results). Therefore, this trend analysis cannot reject our hypothesis that subjects are evaluating the current payoffs (games) by comparing them to the context set up by very recent payoffs. The next analysis aimed to support this hypothesis.

2.2.2 Low-High vs. High-Low trends

If context effects are only transitory, we should see this effect by looking at the time series of choices in these two conditions too. In particular, a difference between Low-High and High-Low would exist, but only at early rounds after the conditions have switched within subjects. To test this hypothesis, we fit a model to each subject in the Low-High and High-Low conditions respectively with the following components:

- a*: general linear trend over the 96 trials to capture the general effect of repetition;
- b*: dummy variable for condition, High vs. Low (allowing that the condition does matter, although apparently it doesn't);
- c*: dummy variable for the first 2 rounds in each block of 48 (after the transition) in order to capture the initial learning period;
- d*: dummy variable for the first 6 rounds in each block of 48 rounds after the transition in order to capture immediate context effects (in each condition).

The issue was whether *b* (global context effect) and *d* (local context effect) depend on the order of the conditions (High, Low). The idea here was to remove the effect of general changes over time by modeling those changes within each block of 48 trials (but removing the first few trials in that block). That is, if there is only a local context effect, then *d* in Low-High should be significantly higher than in High-Low, but no significant difference should be observed between *b* across the two conditions.

Indeed, we observed such pattern: $b(\text{High-Low}) = 1.34$ and $B(\text{Low-High}) = 2.91$, but this difference was not statistically significant, $t(38) = 0.47$, $p = .6383$; however, $d(\text{High-Low}) = -5.93$ was significantly higher than $d(\text{Low-High}) = 2.78$, $t(38) = 2.47$, $p = .0182$. (The other two parameters, *a* and *c*, did not differ significantly across the two conditions.)

These results indicate that the transition effect is significant only in the very first few rounds (i.e., showing a temporary contrast effect in shifting from high to low

or low to high), while there is no global effect of context (the context effect seems to disappear after the first six rounds). This is in line with the relative judgment model (Stewart, Brown, & Chater, 2005), in which effects from previous rounds $t-2$ and $t-3$ are allowed to have some residual effect, but to a much lesser degree. (Note that currently this is an intensely debated issue in the absolute identification literature: see Stewart, 2007.) This regression evidence also brings additional support for our hypothesis about the contextual locality of the effects, which also supported by the aggregate level analysis and the weak learning trends in the two control conditions.

3 General discussion

The results clearly show sequential local context effects on choice in strategic game playing, which coordinated responses and created a long-lasting increase in cooperation in the Mixed condition where the contrast between the two game types interleaved on a trial by trial basis. When the games were divided in two separate sequences, the negative games in the first half of the session did not produce a long-lasting contrast effect with the positive games in the second half. As a consequence, the latter were no more cooperative than normal. Also the cooperative games in the first half of the session did not produce any contrast effect on the uncooperative games in the second half. Thus, we demonstrated that sequential context effects arise, given variations in cooperativeness of PD games, only as local context effects, because people compare the current game with the previous game and quickly learn to coordinate their behaviour depending on the game played; but there are no significant global contrast effects, when the comparison is between entire sequences of games. Another important result was the fact that the average cooperation in the positive games in the Mixed condition was significantly higher than the positive games in both the High-Low condition and the Low-High condition. This result could be explained by the presence of a very powerful contrast (between the two game types) operating on a trial by trial basis, which enabled the participants to quickly learn cooperate in the positive (cooperative) games.

The average cooperation in the negative games in all three conditions was around 18%, indicating that the context affected only the positive games. This asymmetry of the context effect suggests that there is some default minimum cooperation level which cannot be reduced. This result is an indicator of a *floor effect* (i.e., cooperation cannot go lower than certain level), which is most likely due to the documented existence of so-called unconditional cooperators, who tend to cooperate no matter what (as Kurzban & Houser, 2001, seemed to find). Therefore

the stable ~18% cooperation rate across all conditions is not sufficient to rule out the significant context effect observed in the Mixed condition.

3.1 Theoretical accounts

Our results demonstrate strong evidence that people (explicitly or implicitly) compare the games between the trials, but the data do not indicate a long term representation of previous game sequences. This result is in line with the predictions of the (psychophysical) relative judgment model (Stewart, Brown, & Chater, 2005), which postulates that judgments are based the differences between current and previous stimuli without using long-term representations of absolute magnitudes. Thus, we also demonstrated that variation in local context would have a bigger effect on judgment than global context, which is a main prediction of the model, and has not been tested before in the context of strategic interaction (Vlaev & Chater, 2006, do not differentiate between such local and global context effects).

In order to provide an account of the results presented in this article, we could assume that people are unable to make reliable judgements of absolute cooperativeness, because they do not have direct access to internal representation and information about absolute magnitudes in the environment. Instead, they always need some reference standard either retrieved from memory, or existing in the current environment, with which to compare the current game. Such reference standard in our study is provided by the previous games that have different cooperativeness. This explanation is further motivated by a new theory of decision making developed by Stewart, Chater, and Brown (2006), in which, in contrast with traditional models, there are no underlying psychoeconomic scales. Instead, consistent with some interpretations of psychophysical data (e.g., Laming, 1997), this model assumes that an attribute's subjective value is based only on its rank within a small sample of attribute values drawn from memory or from the current decision context.

Another plausible account of the contrast effects caused by the games in the sequence could be some relativistic theory of human judgement like the range-frequency theory proposed by Parducci (1965, 1974), which claims that the subjective value given to an attribute is a function of its position within the overall range of attributes, and its rank. Therefore, this model implies that games' attributes (e.g., cooperativeness) are judged purely in relation to one another and their subjective value is independent of their absolute value (defined by the cooperation index in the case of PD games).

Vlaev and Chater (2006) already demonstrated a context effects in games, which could be explained by such relativistic theoretical frameworks. The results presented

in this article elaborate these findings by showing for first time that the source of the context effect is a comparison of the current game with the immediately preceding games as opposite to comparing the current game with some amalgam of the games experienced during the whole session (the short-lived contrast effect in the Low-High condition supports this conclusion). Our study was specifically motivated by the relative judgment model (Stewart, Brown, & Chater, 2005), which not only assumes that people have no absolute access to psychophysical magnitudes, but that they can only make *ordinal* judgments concerning the size of “jumps” between pairs of magnitudes (this model explains a wide range of absolute magnitude experiments, including trial-to-trial context effects). Note that other existing psychophysical models for a wide range of perceptual phenomena like the information transmission limits, bowed serial position effects, and sequential effects, assume that identification is achieved using long-term representations of absolute magnitudes. The relative judgment model (Stewart, Brown, & Chater, 2005) accounts for these phenomena without using long-term representations of absolute magnitudes.

Our results show that principles underlying judgments about cooperativeness may be based on principles postulated by the relative judgment model, instead of principles underlying other models like the range-frequency theory (Parducci, 1965, 1974) which assumes long-term representations. Note that models like the range-frequency theory were used to account for the context effects on games demonstrated by Vlaev and Chater (2006). However, this article shows that Vlaev and Chater’s (2006) results could also be due to psychological processes that do not utilise long-term representations of cooperativeness magnitudes because such magnitudes either are unavailable or, for some reason, are unused in judgment.

3.2 Concluding remarks

Our results indicate, first, that the notion of cooperativeness of a game, although not presented directly to participants and apparently highly abstract, nonetheless exhibits context effects, just as do perceptual magnitudes such as loudness and brightness. Second, these context effects are relatively short-term. This pattern is consistent with the assumption that the cooperativeness of a game is assessed not in absolute terms, but in comparison with the cooperativeness of a small number of the most recent games that have been played. This study also adds to previous results on perceptual context effects on games shown by Vlaev and Chater (2006) by revealing the local nature of the effect, which does not utilise long-term representations. The effects shown in this article are best ex-

plained by the relative judgment model (Stewart, Brown, & Chater, 2005).

Our results also imply that games are not considered independently of the previously played game. Thus, the present research demonstrates a new and large anomaly for normative rational theories (like game theory) of decision-making under strategic uncertainty. Therefore, the present study demonstrates the extent to which basic aspects of human magnitude representation play an important role in strategic decision making. Our results also suggest that any descriptive theory of choice will be incomplete without taking into account this fundamental aspect of human cognition.

References

- Allais, M. (1953). Le comportement de l’homme rationnel devant la risque: critique des postulats et axiomes de l’école Américaine. *Econometrica*, *21*, 503–546.
- Birnbaum, M. H. (1992). Violations of monotonicity and contextual effects in choice-based certainty equivalents. *Psychological Science*, *3*, 310–314.
- Camerer, C., & Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*, 827–874.
- Dawes, R., & Thaler, R. (1988). Cooperation. *Journal of Economic Perspectives*, *2*, 187–197.
- Erev, I., & Roth, A. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique mixed strategy equilibria. *American Economic Review*, *88*, 848–881.
- Frederick, S., Loewenstein, G., & O’Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, *XL*, 351–401.
- Fudenberg, D. & Tirole J. (1991). *Game Theory*. Cambridge, MA: MIT Press.
- Garner, W. R. (1954). Context effects and the validity of loudness scales. *Journal of Experimental Psychology*, *48*, 218–224.
- González-Vallejo, C. (2002). Making trade-offs: A probabilistic and context-sensitive model of choice behavior. *Psychological Review*, *109*, 137–154.
- Knez, M., & Camerer, C. (2000). Increasing cooperation in social dilemmas through the precedent of efficiency in coordination games. *Organizational Behavior and Human Decision Processes*, *82*, 194–216.
- Kreps, D. (1990). *A course in microeconomic theory*. Princeton, NJ: Princeton University Press.
- Kurzban, R., & Houser, B. J. (2001). Individual differences in cooperation in a circular public goods game. *European Journal of Social Psychology*, *15*, 37–52.
- Laming, D. R. J. (1997). *The measurement of sensation*. London: Oxford University Press.

- Lockhead, G. (1995). Psychophysical scaling methods reveal and measure context effects. *Behavioral and Brain Sciences*, 18, 607–612.
- Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics*, 107, 573–597.
- Loewenstein, G., & Thaler, R. H. (1989). Intertemporal choice. *Journal of Economic Perspectives*, 3, 181–193.
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92, 805–824.
- Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: emotional reactions to the outcomes of risky options. *Psychological Science*, 8, 423–429.
- Parducci, A. (1965). Category judgment: A range-frequency theory. *Psychological Review*, 72, 407–418.
- Parducci, A. (1974). Contextual effects: A range-frequency analysis. In L. Carterette and M. P. Friedman (Eds.), *Handbook of Perception* (Vol. II, pp. 127–141). 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, 363–392.
- Rapoport, A., & Chammah, A. (1965). *Prisoner's dilemma: a study in conflict and cooperation*. Ann Arbor: University of Michigan Press.
- Rapoport, A., Stein, W. E., Parco, J. E., & Nicholas, T. E. (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior*, 43, 239–265.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108, 370–392.
- Slovic, P. (1995). The construction of preference. *American Psychologist*, 50, 364–371.
- Stewart, N. (2007). Absolute identification is relative: A reply to Brown, Marley, and Lacouture (2007). *Psychological Review*, 114, 533–538.
- Stewart, N., Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 3–11.
- Stewart, N., Brown, G. D. A., & Chater, N. (2005). Absolute identification by relative judgment. *Psychological Review*, 112, 881–911.
- Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53, 1–26.
- Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: How choice options influence decision under risk. *Journal of Experimental Psychology: General*, 132, 23–46.
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: an experimental test. *Quarterly Journal of Economics*, 112, 647–661.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, 39, 1179–1189.
- Vlaev, I., & Chater, N. (2006). Game relativity: How context influences strategic decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 131–149.
- Vlaev, I., Chater, N., & Stewart, N. (2007a). Financial prospect relativity: Context effects in financial decision making under risk. *Journal of Behavioral Decision Making*, 20, 273–304.
- Vlaev, I., Chater, N., & Stewart, N. (2007b). Relativistic financial decisions: Context effects on retirement saving and investment risk preferences. *Judgment and Decision Making*, 2, 292–311.
- Ward, L. M., & Lockhead, G. R. (1970). Sequential effect and memory in category judgment. *Journal of Experimental Psychology*, 84, 27–34.
- Ward, L. M., & Lockhead, G. R. (1971). Response system processes in absolute judgment. *Perception & Psychophysics*, 9, 73–78.