האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

DETECTING CHANGE IN PARTNER'S
PREFERENCES

By

JUDITH AVRAHAMI and YAAKOV KAREEV

Discussion Paper # 557   July 2010

מרכז לחקר הרציונליות

CENTER FOR THE STUDY
OF RATIONALITY

DETECTING CHANGE IN PARTNER'S PREFERENCES

Judith Avrahami and Yaakov Kareev[1]

ABSTRACT

Studies of the detection of change have commonly been concerned with individuals inspecting a system or a process, whose characteristics were fully determined by the researcher. We, instead, study the detection of change in the preferences – and hence the behavior – of others with whom an individual interacts. More specifically, we study situations in which one's benefits are the result of the joint actions of one and one's partner when at times the preferred combination is the same for both and at times it is not. In other words, what we change is the payoffs associated with the different combinations of interactive choices and then look at choice behavior following such a change. We find that players are extremely quick to respond to a change in the preferences of their counterparts. This responsiveness can be explained by the players' impulsive reaction to regret – if one was due – at their most recent decision.

Keywords: Regret Driven Behavior; Strategic Interaction; Detection of Change; Stag Hunt; Snowdrift

INTRODUCTION

The world around us is constantly changing and the sooner we detect a change the faster we shall be ready to face the new situation. This is particularly true of changes in the preferences of people with whom we interact, of people whose choices, combined with our own affect our wellbeing. When certain combinations of actions are preferred by all matters are simple: One could just choose the action that would benefit all. Romantic interactions and many business and other social interactions, are examples of such a case. Matters become more complex when combinations of actions that are good for one are not so for the other. Obviously, if both acted to

---

[1] School of Education and The Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem, Israel

promote only their own good, the outcome of their joint actions would be bad for both. We shall denote the former case as one of *aligned* preferences and the latter as *misaligned.*

More complicated still is the fact that life neither involves only interactions with those whose preferences are aligned with one's own, nor only interactions with those whose preferences are misaligned, but periods of one and periods of the other. The tastes of romantic partners may change such that joint actions that brought them pleasure do so no more or the investment preferences of business partners may differ such that with some, certain activities are profitable to all and with others they may be in conflict. Here we focus on situations in which the preferences of one's counterpart change from being aligned to misaligned and vice versa and ask if, how and how fast people detect such changes and change their behavior accordingly.

A large part of the research concerned with the detection of change has been conducted from a prescriptive point of view, with mathematicians, statisticians, computer scientists, engineers, and psychologists proposing mathematical analyses, computer algorithms and heuristics designed to optimize the detection of change (e.g., Rapoport et al. 1979; for an overview, see Nikiforov 2001). The size of the changes, their frequency, and the variability (noise) in the generating process, are all important parameters determining the amount of data that should be used to establish both a base line and an estimate of the current state.

A descriptive approach to the detection of change can be found in the work of economists (e.g., Massey and Wu 2005), ecologists and psychologists interested in foraging behavior (e.g., Kacelnik et al. 1987; Shettleworth et al. 1988), and psychologists focusing on learning and on judgment and decision-making (e.g., Barry and Pitz 1979; Brown and Steyvers 2005 2009; Brown et al. 2007; Gallistel et al. 2001). One important finding of these studies is that organisms – humans and other animals alike – are fast at detecting changes (but see Brown and Steyvers 2005, for a case of slower detection). In fact, the observed reaction to changes is typically so fast that researchers in the tradition of learning theory point out the challenge that such fast detection poses to classical reinforcement-learning models (Gallistel et al. 2001). Another important result to come out of these studies is that of system neglect (Massey and Wu 2005; Gabaix et al. 2006): Although people are sensitive to change and capable of detecting it, they give too much weight to incoming data (the strength of the evidence), and not enough weight to characteristics of the systems generating

these data (the weight of the evidence). This result is in line with earlier findings by Griffin and Tversky (1992).

All the studies cited above focus on the detection of change by an individual observing the behavior of a mechanism or a process, whose outcomes are independent of the observer's actions. The focus of the current study is on the detection of change in social interactions, in which a person's rewards are determined by the joint action of that person and his or her counterpart. The question is, thus, how people would adjust their behavior following a change in the preferences of their counterparts.

Behavior in situations in which preferences are aligned and situations in which they are misaligned has also been studied before. The former as the assurance problem (Sen 1967) in a game called "The Stag Hunt" (e.g., Battalio et al. 2001), and the latter in a number of games like "The Snowdrift" (e.g., Hauert and Doebeli 2004; Kümmerli et al. 2007), "The Game of Chicken" (Rapoport and Chammah 1966; Bornstein et al. 1997; Bornstein and Gilula 2003), or "The Battle of the Sexes" (Cooper et al. 1989). The former – The Stag Hunt game – is based on Rousseau's metaphor of hunters agreeing to collaborate in hunting a stag but each may be tempted to abandon their post and go after a passing hare. A stag is worth much to every one of the hunters that share it but can only be hunted if all are committed to hunting a stag. If they were not all there, those who still are would return empty-handed. Each could hunt a hare on his own, hence this option is less risky; this option is, however, also less profitable because a hare is worth less than a share of a stag. In the game, usually played in pairs, if both players choose what is best for each of them, namely, a 'stag', both gain most but the choice involves the risk that the other would fail to act accordingly. From a game-theoretic perspective, this game has two pure strategy equilibria: one with both players always choosing to go for the stag and the other with both always choosing to go for the hare. From neither of these states can any of the players profitably deviate.

In the three games of misaligned preferences the combination of actions that is best for one is always bad for the other. For example, the Snowdrift game represents a situation in which two drivers are caught in a snowdrift blocking their way home. If none gets out of their car to shovel the snow – both may perish. If both do get out to shovel they would share the unpleasantness of the chore but get home safely and fast. Best for each of them, of course, is to stay in their car while the other does the shoveling. This game too has two pure-strategy equilibria – with either one or the

other shoveling – and also a mixed-strategy equilibrium – which, although less profitable, on average, is the only equitable one. The three games mentioned differ in what happens if both actors give up their preferred action. In the Snowdrift, both are better off when both choose to shovel than when they both choose to stay in their car; in the Battle of the Sexes or in the game of Chicken they are as bad off in both cases. Figure 1 presents a version of the payoff matrices of The Stag Hunt (a) and the Snowdrift (b) games.

	B (stag)	Y (hare)
B (stag)	8,8	0,4
Y (hare)	4,0	2,2

1a

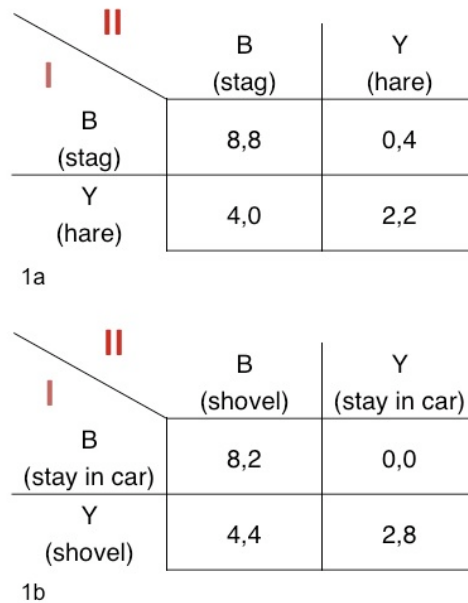|	B (shovel)	Y (stay in car)
B (stay in car)	8,2	0,0
Y (shovel)	4,4	2,8

1b

Figure 1. Payoff matrices for the Stag Hunt game (a) and for the Snowdrift game (b). The specific values are those used in the study.

All these games have been studied both theoretically and experimentally both in a single move and in repeated choice scenarios and both for pairs of players and for populations of players. The results of all earlier studies describe how people act in each of these situations separately. We study how people adapt to a transition between the two situations: How efficiently they make use of the situation of aligned preferences; how fast they realize that the preferences of their counterparts have become misaligned; and how fast they detect that those preferences have turned back to being aligned.

To find out, we designed a game in which two groups of players had the same set of payoffs but these were associated with different joint actions. In periods that modeled interaction with counterparts whose preferences were aligned, players were paired with members of their own group; the joint payoff matrix then resembled that of the Stag Hunt game. In periods that modeled interaction with counterparts whose

preferences were misaligned, players were paired with members of the other group; the joint payoff matrix then resembled that of the Snowdrift game.

Players could, of course, treat the situation as one of incomplete information (Aumann et al. 1995; Harsanyi 1967) and expect their counterpart's behavior to be a mixture of that of the two types. In that case, no change in behavior is expected following a change in a partner's type. However, since change in the counterpart's type occurred periodically – and players knew it to be so – we expected players to attempt to detect those changes and adjust their behavior accordingly.

To detect changes fast one ought to be reactive to one's most recent history. Taking into account a lengthy history of past occurrences would postpone one's adjustment to change. The highest level of adaptation to a changing situation would result from best responding to the outcome of the most recent action though, of course, at the cost of reacting too fast to noise. To find out if players were indeed reactive to the outcome of their most recent action we tested how well we could predict the chosen action on the basis of only the very last action and its outcome (see Avrahami and Kareev in press, for a similar analysis of repeated choices, albeit in stable environments). If the most recent action and its outcome were found to be a good predictor of players' following action we could see how each outcome affected the players upcoming choice. The study was designed in such a way that the value of an actual outcome (i.e., the size of the reward obtained) did not necessarily correspond to the difference between its value and that of the counterfactual one – the outcome that would have been obtained had an alternative choice been made. The design could thus reveal what players were most concerned about: actual gains only (in line with pure reinforcement-learning models, e.g., Roth and Erev 1995) or a comparison of actual and counterfactual gains (e.g., Avrahami et al. 2005; Grosskopf et al. 2006; Hart 2005; Hart and Mas-Colell 2000; Selten et al. 2005; Yechiam and Busemeyer 2006).

METHOD

The choices players faced were symbolized by colors: On each round players had to choose between blue and yellow. In order not to confuse players about their own preferences, the preferences of each player stayed stable throughout the game. Half of them preferred a joint choice of blue ('type 1' players) and the other half preferred a joint choice of yellow ('type 2' players). A change in one's counterpart's

preferences was modeled by changing the 'type' of the counterparts with whom one interacted: In periods of aligned preferences, players were paired with ones of the same 'type' (i.e., preferring the same joint color) and in periods of misaligned preferences they were paired with players of the other 'type' (i.e., preferring a different joint color). As a result, in periods of aligned preferences they were playing the Stag Hunt game and in periods of misaligned preferences they were playing the Snowdrift.

The instructions informed players of the existence of two types and presented the payoff of the joint choices for each type. Payoffs were 8, 4, 2, or zero points corresponding to the values in Figure 1. Note that a player's payoffs for the various combinations of choices were the same for both types. It is the alignment between one's payoffs and that of one's counterpart – in the different joint choices – that differed between the aligned and the misaligned situations. The instructions also made it clear that every player's own type would be assigned at the outset and remain unchanged throughout the game. On every round every player would be randomly paired with another: In certain periods this other would be of the player's own type and in others of the other type.

The game was computerized and played in sessions of 12 participants in the Ratiolab laboratory at the Psychology department of the Hebrew University. Participants were rewarded, on each round, in line with their value for the combination of their and their counterpart's choices. The conversion rate was 1 New Israeli Shekels (NIS, about 25 US cents) for every 15 points gained in the game.

Pretest (with no change)

To acquire a baseline for behavior with a counterpart of a different type (namely, a counterpart whose preferences were misaligned with one's own), two sessions were conducted in which players always played with others of the different type. Thus six players of one type were each randomly matched, on every round, with one of the six players of the other type. These sessions lasted for 60 rounds.[2]

---

[2] Earlier pretests in which players played continuously with same-type others indicated that they quickly converged to always playing their preferred option. We therefore did not see the need (and could hardly afford) to run same-type, no-change sessions.

Main experiment (with change in preference-alignment)

The main experiment involved periodic changes between aligned and misaligned pairings. For generality, the experiment was conducted under three different change conditions; these differed with respect to the overall rate of change and the predictability of a change. A change occurred at about every 15 rounds, every 20 rounds, and every 30 rounds. In the sessions with a change occurring at about every 15 rounds and in those with a change occurring at about every 30 rounds a random number, drawn from a uniform distribution in the range of -2 to +2, was added to 15 or 30 to make the change less predictable. In the sessions with a change occurring at about every 20 rounds a random number drawn from a uniform distribution in the range of -5 to +5 was added, Two sessions were conducted with each change condition, making 6 experimental sessions in all. Each experimental session lasted for 120 rounds.

Participants

Twelve students of the Hebrew University, recruited by email from the Ratiolab's database of participants, took part in each session. There were thus 24 participants in the pretest and 72 participants in the main experiment. Participants earned, on average, 38 NIS.

RESULTS

The first question was whether players detected the changes and reacted to them in their choice of action. As a dependent measure we used the proportion with which the players chose the color that corresponded to 'stag' in the aligned periods or to 'stay in the car' in the misaligned periods. This was the color that could have earned them 8 points (but also zero points, depending on their counterpart's choice). For brevity we shall call this the 'stag' choice. The average proportion of choosing 'stag' was 0.82 in periods of aligned preferences and 0.59 in periods of misaligned preferences. This difference indicates that players reacted to the situation they were in. The latter value is very close to the average proportion of choosing 'stag' in the pretest, in which preferences were always misaligned (0.61).

Finding that players detected the changes and reacted to them, we turned to explore the dynamics of that detection. How fast did behavior change following a transition from aligned to misaligned preferences? How fast did it change following

the opposite transition? Figure 2 presents the average proportion of choosing 'stag' on each round, together with the preference status line (aligned/misaligned).[3] It is clear from the figure that behavior changed almost immediately following a transition from aligned to misaligned preferences and also fast, but more slowly, following the opposite transition. Indeed, comparing the proportion of 'stag' choices on the last round of aligned and the second round of misaligned preferences (i.e., the first round on which the change could have been reacted to) reveals already a highly significant difference ($t(203) = 5.74, p < .001$). For the opposite transition a significant difference is evident one round later, on the third round, namely the second after a change could have been reacted to ($t(155) = -2.42, p = .017$). Neither the frequency of a change nor the level of noise in its predictability resulted in a significant difference in the speed of detecting a change.
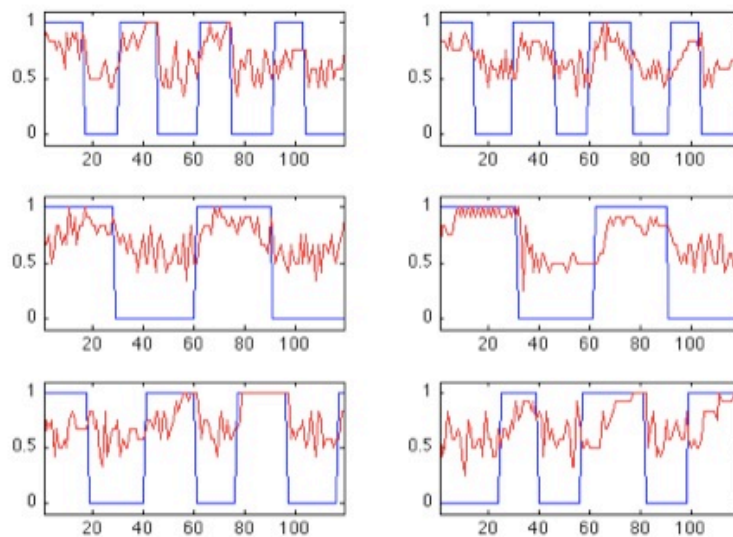

Figure 2. The average proportion of choosing the 'stag' color on every round, separately for each session. The additional, saw-like, line represents periods of aligned and misaligned preferences (1 and 0, respectively).

As to the players' gains, these were, on average, 6.18 points in periods of aligned preferences and 3.39 in misaligned periods. The latter value closely resembles that in the constantly misaligned pretest (3.13 points).

---

[3] Since change occurred in different rounds in the different sessions (due to the insertion of noise to the transition process) the dynamics of the game are presented separately for each session.

With the results demonstrating a quick response to change in the preferences of their counterparts, it is quite obvious that players reacted to the outcomes of their very recent choices. It was therefore interesting to find out how and how strongly they reacted to those outcomes. We remind the reader that there were four possible outcomes, two for each of the two available options. Following a choice of one's 'stag' color one could earn either 8 points – if the counterpart chose the same color – or zero – if the colors mismatched. Following a choice of the other color one could earn either 4 points – if the colors of both players mismatched – or 2 points – if the colors matched. As was already mentioned above, with regard to having chosen 'stag', one of the outcomes, that of not earning anything, may have given rise both to disappointment with the no-gain and to regretting one's own choice: Choosing otherwise would have resulted in a gain of 2 points. With regard to having chosen non-stag, earning 2 points may also have given rise to disappointment: had the counterpart chosen a mismatching color one would have gained 4 points instead. At the same time, in earning 2 points there is no reason for regretting one's choice because choosing the other color would have resulted in no gain at all. It is the choice of non-stag that resulted in earning of 4 points, that is most interesting: Although hardly a cause for disappointment (4 points are the second best level of gain and had the counterpart chosen differently the player would have gained less, namely, 2) it is clearly a cause for regret: choosing otherwise would have resulted in a gain of 8 points. How did players react to each of those outcomes? Which outcomes made them more likely to stay with their previous choice and which made them more likely to switch to the other? Did these patterns correspond to the emotional reaction likely to have emerged? Was it actual gain or regret for forgone payoffs that drove their choice behavior?

To find out, we used the four outcomes as predictors of the players' upcoming choices. More specifically, decisions were classified as either *repetitions* of the previous choice or *switches* from the previous choice and these two values were regressed on the four possible outcomes of the previous choice. The analysis was conducted separately for each participant. The goodness of prediction was high, as expressed in the average R^2 of the regression analysis, which was 0.31 (median $p <$

.001 or, perhaps more informative, only 5 out of the 72 players had a p-value higher than 0.05).[4]

The weights assigned by this analysis to the different predictors can be interpreted as the tendency to repeat a choice following either of its outcomes. Figure 3 presents the tendency to repeat each of the two options – 'stag' or 'non-stag' – following its two possible outcomes – the one that may and the one that may not have given rise to regret.
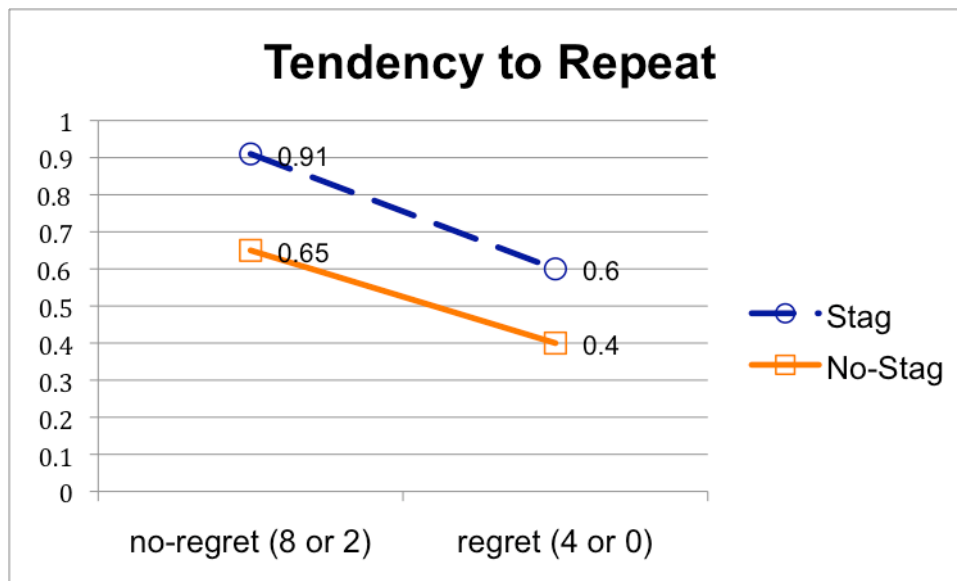

Figure 3. The tendency to repeat a choice following each of the outcomes. The higher line represents reaction to the two outcomes of having chosen the 'stag' color (gaining 8 or 0) and the lower represents reaction to the two outcomes of having chosen the non-'stag' color (gaining 2 or 4). Note that the tendency to repeat following a gain of 2 is higher than that following a gain of 4.

The figure shows that there was an overall higher tendency to repeat a choice of the 'stag' color (which reflects the overall higher proportion of the 'stag' choices, reported above). Most relevant to our thesis, the figure clearly shows that the tendency to repeat the previous choice was high when there was no room for regret and lower when there was. An analysis of variance carried out on the tendencies to repeat revealed these two main effects to be significant and no significant interaction between them. For the tendency to repeat the 'stag' versus the 'non-stag' color $F(1,69) = 46.85$, $p < .001$, $\eta_p^2 = .40$, for the tendency to repeat a choice following

_____

[4] Two players, who chose only 'stag' throughout the game, were dropped from the regression analysis.

regret versus no-regret $F(1,69) = 101.57$, $p < .001$, $\eta_p^2 = .59$, and for the interaction $F(1,69) = 1.83$, $p = .18$, $\eta_p^2 = .03$. It is particularly striking that the tendency to repeat a choice after the relatively high gain of 4 points was so much lower than the tendency to repeat a choice after earning 2 points. Only the effect of regret can explain that.

CONCLUSION

How quickly, then, do people react to a change in their partner's preferences? The results reported above indicate that they react very quickly. In addition, and as was to be expected, they respond faster to a change from aligned to misaligned preferences. This is because as long as preferences are aligned selfishness and care for the other converge and the general atmosphere is that of cooperation. Once preferences become misaligned, however, what used to be an act of cooperation is so no more. Now, cooperation would mean giving up on one's more lucrative rewards, being ready to make do with less. This, in turn, comes at the risk of being taken advantage of hence a mixed strategy may emerge. With partners mixing their choices, the signal of change in preferences – from being misaligned to being aligned – is be more difficult to detect.

Still, the results show that detection is very quick in both cases; so quick, in fact, as to rule out the assumption that the burden of accumulated history has an effect on decisions. In that, the results observed complement those of earlier studies involving individual judgments and decision making (Brown and Steyvers 2009; Gallistel et al. 2001; Shettleworth et al. 1988), and add support to the idea that changes in the environment are detected fast thanks to a limit on the past experience considered. We go further and propose a simple mechanism – an immediate reactance to the experience of regret – which renders navigation through repeated decisions in changing environments automatic and efficient. As we have shown elsewhere, for behavior in stable environments (Avrahami and Kareev in press), such reactance could well be completely impulsive: it occurs unawares and does not diminish with experience. This is not to mean that players do not consciously deliberate, calculate, and test their decisions. All it says is that there is a residual impulsive reactance, most strongly to the feeling of regret for not having decided otherwise, that guarantees

close tracking of the physical and behavioral contour of the environment – whether stable or changing – within which one operates.

The results of this study are also in line with the finding of system neglect (Massey and Wu 2005). In fact, the impulsive reaction to the most recent outcome can be regarded as the utmost neglect of the system, namely, of any prior information about the prevailing situation. Impulsive reaction thus reflects a deep suspicion of such information – either given or learnt through experience – resulting in heavy reliance on recent experience with the latter guaranteeing fastest adaptation to environmental changes.

ACKNOWLEDGMENT

REFERENCES

Aumann, R. J., Maschler, M., & Stearns, R. E. (1995). Repeated games with incomplete information. Cambridge, MA: MIT Press.

Avrahami, J., Güth, W., & Kareev, Y. (2005). Games of competition in a stochastic environment. Theory and Decision, 59, 255-294.

Avrahami, J., & Kareev, Y. (in press). The role of impulses in shaping decisions. Journal of behavioral decision making.

Barry, D. M., & Pitz, G. F. (1979). Detection of change in nonstationary , random sequences. Organizational Behavior and Human Performance, 24, 111-125.

Battalio, R., Samuelson, L., & Van Huyck, J. (2001). Optimization incentives and coordination failure in laboratory Stag Hunt games. Econometrica, 69, 749-764.

Bornstein, G., Budescu, D., & Zamir, S. (1997). Cooperation in intergroup, n-person, and two-person games of Chicken. The Journal of Conflict Resolution, 41, 384-406.

Bornstein, G. & Gilula, Z. (2003). Between-group communication and conflict resolution in assurance and chicken games. Journal of Conflict Resolution, 47, 326–339.

Brown, S. D., & Steyvers, M. (2005). The dynamics of experimentally induced criterion shifts. Journal of Experimental Psychology: Learning Memory, and Cognition, 31, 587-599.

Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. Cognitive Psychology, 48, 49-67.

Brown, s. D., Steyvers, M., & Hemmer, P. (2007). Modeling experimentally induced strategy shifts. Psychological Science, 20, 40-46.

Cooper, R., DeJong, D., Forsythe, R., & Ross, T. (1989). Communication in the Battle of the Sexes game: Some experimental results. The Rand Journal of Economics, 20, 568–587.

Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. American Economic Review, 96, 1043-1068.

Gallistel, C. R., Mark, T. A., King, A. P., & Latham, P. E. (2001). The rat approximates an ideal detector of changes in rates of reward: Implications for the law of effect. Journal of Experimental Psychology: Animal Behavior Processes, 27, 354-372.

Griffin, D., & Tversky A. (1992). The weighing of evidence and the determinants of confidence. Cognitive Psychology, 24, 411-435.

Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. International Journal of Game Theory, 34, 285-302.

Hauert, C., Doebeli, M. (2004). Spatial structure often inhibits the evolution of cooperation in the snowdrift game. Nature, 428, 643-646.

Harsanyi, J. C. (1967). Games with incomplete information played by "Bayesian" players, I-III. Part I. the basic model. Management Science, 14, 159-182.

Hart, S. (2005). Adaptive heuristics. Econometrica, 73, 1401–1430.

Hart, S., & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. Econometrica, 68, 1127–1150.

Kacelnik, A., Krebs, J., R., & Ens, B. (1987). Foraging in a changing environment: An experiment with starlings (Sturnus Vulgaris). In M. L. Commons, A. Kacelnik, & S, J. Shettleworth (Eds.). Quantitative analysis of Behavior, Vol VI: Foraging (pp. 63-88). Hillsdale: NJ: Lawrence Erlbaum Associates.

Kümmerli, R., Colliard, C., Fiechter, N., Petitpierre, B., Russier, F, & Keller, L. (2007). Human cooperation in social dilemmas: comparing the Snowdrift game

with the Prisoner's Dilemma. Proceedings of the Royal Society, B, 274, 2965-2970.

Massey, C., & Wu, G. (2005). Detecting regime shifts: The causes of under- and overreaction. Management Science, 51, 932-947.

Nikiforov, I. V. (2001). A simple change detection scheme. Signal Processing, 81, 149-172.

Rapoport, A., & Chammah, A. M. (1966). The game of Chicken. American Behavioral Scientist, 10, 10-28.

Rapoport, A., Stein, W. E., & Burkheimer, G. J. (1979). Response models for detection of change. Dordrecht: D. Reidel.

Roth, A.E., & Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior, 8, 164-212.

Selten, R., Abbink, K., & Cox. R. (2005). Learning direction theory and the winner's curse. Experimental Economics, 8, 5–20.

Sen, A. (1967). Isolation, assurance and the social rate of discount. Quarterly Journal of Economics, 80, 112-124.

Shettleworth, S. J., Krebs, J. R., Stephens, D. W., & Gibbon, J. (1988). Tracking a fluctuating environment: A study of sampling. Animal Behavior, 36, 87-105.

Yechiam, E., & Busemeyer, J. R. (2006). The effect of forgone payoffs on underweighting small probability events. Journal of Behavioral Decision Making, 19, 1-16.