University of Dundee

Is positive confirmation relevant for game theory?: an experimental test

Jones, Martin K.

Publication date:
2003

Is Positive Confirmation relevant for game
theory?: An Experimental Test

Martin K. Jones

Is Positive Confirmation relevant for game theory?:

An Experimental Test.

Martin K Jones*

Department of Economic Studies, University of Dundee, Dundee. DD1 4HN tel: (01382) 345164. e-mail:

m.k.jones@dundee.ac.uk fax: (01382) 344691

Abstract

Previous experimental work has found that Positive Confirmation is a significant factor in learning. However, to date, there has been no attempt to test this phenomenon in a game-theoretic setting. An experiment is undertaken which tests positive confirmation in a game against nature. It is found that positive confirmation does occur in this setting and results in behaviour that cannot be explained by current theories of choice. This suggests that current theories are incomplete and new approaches need to be developed to take account of positive confirmation.

During the last few years, tests of learning in game theory have become a growth industry in experimental economics. Most of these experiments have taken the form of subjects playing each other in repeated games and fitting different models to the generated data. This has produced many useful results and an increasing fit of models to the data (see Camerer (2003)). There are, however, many potential improvements that could be made to current models and theories. It will be argued in this paper that one fruitful direction would be to take insights on learning from the psychology literature and try to incorporate them into experimental tests of learning in game theory.

The aim of this paper is to test for the existence of *positive confirmation* in a tightly controlled series of games. Positive Confirmation is the tendency, when testing a belief, to search for and use evidence that is implied by that belief rather than evidence that is not implied by it[1]. To date, there have been two experimental studies made of this bias in the economics literature (Jones & Sugden 2001, Jones 2003) although there is a large literature on positive confirmation in the psychology literature (e.g. McKenzie 2004; Manktelow & Over 1993; Oaksford & Chater 1994; Cheng & Holyoak 1989, Wason 1960). The two experiments in the economics literature have focussed on positive confirmation as a cause of bias from expected utility when making decisions.

From these two experiments three conclusions were derived. First, positive confirmation exists and it causes a bias from expected utility maximisation when choosing information. Second, this bias occurs both in cases where the information has informative value but is non- expected utility maximising and also where the information is completely valueless. In the latter case, positive confirming choices of valueless information were shown to happen even when they incurred a cost for the subject. Finally it was shown that,

as well as positive confirmation in the selection of information, there is also positive confirmation in the *use* of information. This means that where information is interpreted as confirming a belief, this increases subjects' confidence in the truth of the belief and their tendency to act on it. This occurs even if, from a Bayesian point of view, that information has no value[2].

In this paper it is proposed to give a test of the existence of positive confirmation in a game theoretic setting rather than in the specialised environments of the previous experiments. We are interested in whether subjects exhibit Positive Confirming behaviour which is not predicted by Bayesian learning theory within this setting. Bayesianism has been selected as the control in the experiment because it is a standard of rationality and also allows for a test of more complicated structures than is allowed by the rigid formats of other learning theories. The experiment therefore is not constructed to test other theories such as fictitious play (Fudenburg and Levine 1998), reinforcement learning (Borgers and Sarin 1997) and Experience- Weighted Attraction (Camerer & Ho 1999). However there will be an interest in how people behave when they fall short of Bayesian rationality and what this means for theorising about learning.

To a certain extent this has already been done in that previous experiments have shown that subjects will tend to use information that is valueless. In that such information is ignored in the structure of game theory and by learning theories this already demonstrates the incompleteness of learning ideas in game theory. However, the aim here is to demonstrate an effect of positive confirmation in a game theoretic setting that is not accounted for by Bayesianism or bounded rationality even when all available information has value.

There is no well-specified theory of positive confirmation as there is for Bayesianism and the myopic learning theories so the experiment will not test a specific model against the data[3]. Instead the experiment will involve a test of the *mechanisms* of Bayesianism to see whether these mechanisms fully explain the results of the experiment. This means that the experiment will only involve a few repetitions of the game used. This focus on the individual stages of the game means that there is no attempt to look at the effects of large numbers of repetitions or asymptotic results. These are not at issue in this paper and are not analyzed.

The implications of previous work on positive confirmation for game theory would seem to be fairly straightforward. If a strategy chosen by an individual is "successful" (i.e. it gets good feedback/ payoffs) then it will be more likely to be chosen again by that person. However this is problematic because there are several potential confounding factors which are partly the result of the testing method used. We have not got a quantitative model to comparatively test with other models so instead we have to rely on the *directions* of strategy choices made by the subjects. Given the qualitative focus of the experiment, these confounding factors become inevitable. The problem is that positive confirmation in this context is hard to distinguish observationally from Bayesian updating and expected utility maximization without controls on the type of behaviour allowed or not allowed within the experiment.

One such problem comes from the payoffs. If in a game there is a comparatively higher payoff as a result of a subject using a strategy, this will increase the likelihood of that strategy being used again. This could either be the result of subjects maximising utility or Positive Confirmation on a "successful" strategy. Therefore, with differing payoffs, the

strategy with the highest payoff would always be chosen and Positive Confirmation would not be qualitatively distinguishable from Bayesianism or other theories.

Another potential confounding factor is the influence of the opposing subject. If a subject updates in a Bayesian format then the opposing subject's play history is used to form a distribution over strategies. This is in turn used to formulate a best reply by maximizing expected utility. If one strategy is consistently a best reply to the opponent's play then this will observationally look like Positive Confirmation. This is because the "best reply" is a successful strategy that will be chosen again and again by the subject. Again, this means that it is hard to distinguish between Bayesian updating and Positive Confirmation simply using the direction of choices.

Any successful experimental test of positive confirmation will have to take these confounding factors into account and control or test for them. Furthermore, any such experiment should be comparable with previous experiments in allowing for tests of positive confirmation. In particular, it should also allow us to distinguish between positive confirmation in the search for and use of information.

The first element of this controlled experiment is to have subjects play a repeated *game against nature*. This effectively removes the effects of playing against another subject and makes the distribution of strategies played by the opposing side (i.e. "nature") stationary. Furthermore, the experimenter can have far greater control over the prior probabilities and information given to the subjects as well as being able to follow in far greater detail the updating processes of the subjects[4].

The second element is to split apart the choosing of evidence and the selection of strategies. This allows a distinction to be made between positive confirmation in the choice

of evidence and in the use of evidence. This means that playing a strategy does not provide information as feedback on it is not provided until the end of the experiment. Instead the subjects receive information through items of evidence known as *combinations*. These combinations give information to the subjects about the strategies played by nature and so give the subject expected values for choosing each strategy. Combinations generally give information about several strategies at once.

This split effectively excludes a direct test of most of the "myopic" learning theories as these assume that strategies are updated when they are played whereas here playing the strategy is different from gaining information about it. This is something that is intrinsic to the experiment. However, this split also opens up the possibility of testing a plausible, alternative "bounded rationality" explanation of behaviour in the experiment. This, briefly, is the idea that a subject only focuses on updating one strategy at a time and ignores evidence for other strategies.

This needs to be tested because such a boundedly rational subject would be a problem when testing for Positive Confirmation. If such subjects find this one strategy "successful" then they will continue to use this strategy in a manner which would be identical to that of a Positive Confirmer. However, this would be the result of reinforcing a strategy while ignoring evidence for other strategies. This behaviour is not the same as positive confirmation where evidence, once obtained, is not usually ignored (although it may be misused). The tactic used in this experiment will be to *assume* that all available information is used but to *test* to see whether subjects are instead boundedly rational in focussing on the information for one strategy at a time.

This split between searching for and using evidence, while it causes complications, is crucial in that it allows for "backwards compatibility" with previous experiments which have tested for positive confirmation. This in turn means that there is more room for comparison and replication of results across experiments. Also, in maintaining the distinction between positive confirmation in searching for evidence (i.e. choosing combinations) and positive confirmation in the use of evidence (i.e. choosing strategies), it allows for an investigation of the links and differences between the two.

The third element in the experiment is to make this game against nature, at each stage, a pure strategic form matching game with equal expected payoffs. This is done by making the prior probabilities of each of nature's strategies equal and having the payoffs for correctly matching nature's strategy the same over all possible strategies. When the probabilities are updated in an objective Bayesian fashion then they eliminate precisely half of the strategies at every round. This means that the Bayesian probability of the eliminated strategies is zero while the positive probabilities are equally loaded onto the remaining strategies. The result of this is that the expected payoff at each stage of the game for the surviving strategies is equal.

These three elements together mean that at each stage of the game the subject has a choice between a set of equally weighted strategies. This naturally leaves the problem of what to choose where one is indifferent between all available strategies. In such a situation we will follow Harsanyi and Selten's (1988) notion of *symmetry invariance*. Informally this states that if strategies are identical save for their labelling or ordering then they should be chosen with equal probability. So, in this case, given that each strategy is equally likely to be chosen by nature and that successful matching results in the same payoff irrespective of

strategy, this means that each subject's strategy should be equally likely to be chosen. It follows that any deviation from this will be a good indicator for the effect of positive confirmation.

2. Experimental Design

The subjects in the experiment[5] were presented with a set of sixteen rules of which one was selected at random as the *hidden* rule. This random selection of a rule acted the role of "nature" in the game. The objective for the subjects was to find out which rule was the "hidden" rule. This was done by testing evidence presented in *rounds of evidence* and then choosing one of the rules. Each round of evidence consisted of a choice between two pairs of combinations of symbols which allowed the subject to increase her knowledge of the hidden rule. At the start and finish of the experiment, as well as in between the rounds of evidence, the subject was asked a question; in each case to state which rule she thought was the hidden rule. There were five questions interleaved with four rounds of evidence.

These questions can be seen as the subjects' play in the game against nature. The subject matched with nature when they answered a question by stating the hidden rule. The selection of combinations allowed the subjects to increase their knowledge of this rule so that each stage game involved matching with Nature over a smaller set of possible strategies.

The rules selected were related to symbols, each of which had four major *attributes* relating to their: a) Shape, b) Colour, c) Size, d) Possession of a border. Denote each attribute by $\aleph_i$ where subscript i is an index to which attribute is being referred. In each

case there were two possibilities or *aspects* for each attribute (e.g. the shape could either be a cross or an octagon. For a full list of rules see Appendix A). An aspect x_i of attribute $\aleph_i$ will be referred to as either α_i or β_i so that $x_i \in \{\alpha_i, \beta_i\}$. Each rule was a simple declarative statement and related to a layout where there were two such symbols; one denoted the *Left Symbol* and the other the *Right Symbol* .Together this pair of symbols comprised a combination. Rules were declarations that an aspect x_i did or did not hold for one or other of these symbols.

Denote the set of rules as **R**. The hidden rule was picked out at random without the knowledge of the subject so the probability of any one rule being picked out was one-sixteenth. Denote this hidden rule $H \in \mathbf{R}$. It should be noted that the rules are symmetrical in the sense that there are equal chances of H being about either aspect in each attribute. It is also equally likely that H would be about either the left symbol or the right symbol. These facts were known to the subject.

In each round of evidence in the experiment the subject was presented with a choice of two combinations in order to test for the rule H. Each symbol in a combination had one aspect from each of the four attributes. Suppose X_1 represents the collection of attributes for the left symbol and Y_1 represents the collection of attributes for the right symbol, both in the first round of evidence. The right symbol was constructed so that it shared at least one and at most three attribute aspects with the left symbol so, if $X_1 = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ then one possibility for the other symbol is $Y_1 = (\alpha_1, \beta_2, \beta_3, \beta_4)$. (Another possibility would be $(\beta_1, \alpha_2, \alpha_3, \beta_4)$).

The second combination offered in the round of evidence also had two symbols. However, in this case, the attribute aspects were inverted. Denote the corresponding left and right symbols for the second combination in the first round of evidence as X'_1 and Y'_1. In this case $X'_1 = (\beta_1, \beta_2, \beta_3, \beta_4)$ and $Y'_1 = (\beta_1, \alpha_2, \alpha_3, \alpha_4)$. This symmetry between the two combinations was important in the construction of the experiment since it meant that each attribute aspect, for each symbol position, appears just once in each round of evidence. This corresponds with the fact that each aspect is mentioned in just one rule in **R**. Together X_1 and Y_1 form the combination $\mathbf{Z_1}$ while X_1' and Y_1' form the combination $\mathbf{Z'_1}$ in the first round of evidence.

The evidence varied from round to round to enable the subject to acquire more information about the rule H. The general aim in the experiment was to have every symbol in the experiment vary in at least one attribute aspect from any other symbol. In order to maintain the symmetry between the four rounds of evidence it was decided that symbols in the same position and combination should have different aspects for two of the attributes from each other. So, if the attribute aspects for the first round of evidence were as above then, for the second round, one possibility would be: $X_2 = (\alpha_1, \alpha_2, \beta_3, \beta_4)$, $Y_2 = (\alpha_1, \beta_2, \alpha_3, \alpha_4)$, $X'_2 = (\beta_1, \beta_2, \alpha_3, \alpha_4)$, $Y'_2 = (\beta_1, \alpha_2, \beta_3, \beta_4)$. This allowed for each round having the same informative content as all the other rounds which means that the rounds of evidence could be presented in any order.

Once the subject had chosen one of the two combinations then she was told whether that combination was *allowed* by the hidden rule H or not. A combination $\mathbf{Z_j}$ (where j is the round of evidence) is said to be allowed by any rule $R \in \mathbf{R}$ if the attribute

aspect x_i mentioned and indicated in the symbol by R, is identical with the attribute aspect in the symbol in $\mathbf{Z_j}$. If a combination is allowed or disallowed by H then this communicates information about H to the subjects

Incentives were distributed on the following basis. At the end of the experiment the rule H was revealed to the subject and then the rule was compared to the subject's answers to each of the five questions. A die was rolled by the experimenter and the number on the die used to select one of the five questions. If the answer to that question was correct then the subject won a money prize of £20. Otherwise they won nothing. If a six was rolled then the subject was given an automatic "win" of £20.

4. Hypotheses

In what follows we will have two separate definitions of Positive Confirmation. *Positive Confirmation in the choice of evidence* is when a subject chooses a combination which is allowed by the rule given by the subject in the preceding question. This is equivalent to the "positive confirmation in the search for evidence" mentioned above. *Positive Confirmation in the choice of rule* is when the subject, having Positively Confirmed in the choice of evidence and been told that the combination is allowed by the hidden rule, then goes on to repeat the same rule in the following question. This is equivalent to "positive confirmation in the use of evidence" as mentioned above

The experiment is structured so as to allow for the updating of probabilities in the following manner. Examination of the structure of the combinations shows that in each round of evidence half the rules are eliminated so that in the final questions there is only

one possible rule left. It follows that in each question one can calculate the *objective*

Bayesian probability of choosing the correct rule. Hence, the prior probability of any one

rule being H is known at the start to be 1/16, after the first round of evidence for those rules

not eliminated it is 1/8. In the third question it is 1/4, in the fourth 1/2 and, in the fifth

question, the rule should be known.

A crucial problem which has to be assessed is how would a person choose amongst

these equally- rewarding rules each time they have to make a choice? The basic solution

has already been outlined above. One should assume, as a null hypothesis, that subjects are

rational in the sense that they obey symmetry invariance and assign equal probabilities to

each strategy. However, this is insufficient for this experiment since, as has been pointed

out by Schelling (1960) as well as by Sugden (1995) and Bacharach (1993), it is possible

for subjects, when faced with matching games, to rationally select a strategy by focussing

on the labelling (or "concepts") of the game.

This creates a problem of control in the experiment because it is possible that a person

playing the game could choose a strategy based on the labels (called "aspects" here) rather

than on positive confirmation. The answer to this is to follow Casajus' (2000) notion of

framed strategic forms. A frame strategic form is a strategic form game which includes the

labelling of the game within its definition. This means that any solution to the game must

take account of the labelling of the games.

The main concept of relevance here is Casjus' extension of symmetry invariance to

framed strategic forms. According to this a framed strategic form is symmetry invariant if it

is symmetry invariant in strategies but also if there is an isomorphism into itself of the

labels and attributes as well (Casajus 2000). This means that swapping the labels around

between two strategies will not have any effect on the choice between the two. Again, this implies that symmetric strategies are assigned equal probabilities.

An examination of the rules in this experiment will show that the design of the experiment is highly symmetric in this sense. There is no bias towards either symbol in each combination, there are equal numbers of rules covering each attribute (i.e. four each for colour, size, shape and border) while there are two aspects for each attribute. This means that there is no label (aspect) which is distinguished by different numbers of rules and so they are symmetrical and hence symmetric invariant.

It follows that one would expect rule choices to be randomly distributed amongst the possible rules with the expected numbers choosing each rule being roughly equal. From this, the principal null hypothesis in this experiment is that subjects choose rules at each stage of the game randomly, in line with the Bayesian probability of a particular rule being the hidden rule. Define the set $\mathbf{R_a} \subset \mathbf{R}$ as the set of rules in question "a" which have not been eliminated by objective Bayesian updating in the previous rounds of evidence. Suppose $\psi_a(.)$ is a function defined from the set $\mathbf{R_a}$ to the interval $[0,1]$ where subscript "a" is an index of the question. This can be defined as a *decision probability* over the set of rules $\mathbf{R_a}$. Suppose that $R_i \in \mathbf{R_a}$ is a rule that has been chosen by a given subject in the previous question. In the questions for $a = 2,3,4^6$ then $\psi_a(R_i) = $ 1/8, 1/4, 1/2 for $R_i \in \mathbf{R_a}$ (assuming that R_i is not eliminated at any stage). The alternative hypothesis in this case would be that of Positive Confirmation i.e. that $\psi_a(R_i) > $ 1/8, 1/4, 1/2 for $R_i \in \mathbf{R_a}$ as $a = 2,3,4$.

As it stands this null hypothesis does not strictly test Positive Confirmation in the choice of rules as it simply tests for significant amounts of repetition of rules. However, as

we mentioned previously, Positive Confirmation is used in a far more sophisticated way than this. Any reasonable test should only include those who are positive confirmers in their choice of evidence and whose selected combination is allowed by the hidden rule. The results for this hypothesis therefore will focus on this subsection of the sample

The secondary null hypothesis relates to Positive Confirmation in the choice of evidence. Here we will also derive our null hypothesis from Casajus' extension of symmetry invariance. For this the choice of evidence is important to the subject in that it gives information which is valuable in choosing which strategy to pick. In fact each combination in each round will eliminate precisely half of the rules on offer so they can be seen as being equally valuable and so symmetrically invariant in the sense of Harsanyi and Selten.

However, as was laid out in the second section in this paper, the combinations in each round of evidence are counterbalanced with each other in terms of attributes so that neither combination stands out against each other. The effects of this counter- balancing can be seen in Table 1. From this it can be seen that Casajus' extension of symmetric invariance applies here as well so that, in each round of evidence, a combination should be equally likely to be chosen. It follows that any systematic divergence from this will not be the result of the labelling of the combinations.

Suppose we define a function $\pi(.)$ as a decision probability defined from the set $\{Z_b, Z_b'\}$ of combinations to the interval $[0,1]$. $\pi(Z_b)$ or $\pi(Z_b')$ are the probabilities that a subject chooses combination Z_b or Z_b' in the bth round of evidence when attempting to find rule H. Suppose also that the combination Z_b is allowed by the rule chosen in the previous

question. In this case the null hypothesis would be $\pi (\mathbf{Z_b}) = \pi (\mathbf{Z_b'})$ where b=1,2,3,4 . The alternative hypothesis is that of Positive Confirmation in the choice of evidence. In this case one would expect that a positive confirmer would tend to choose combination $\mathbf{Z_b}$ over combination $\mathbf{Z_b'}$. It follows that the alternative hypothesis in this case would be $\pi (\mathbf{Z_b}) > \pi (\mathbf{Z_b'})$ where b=1,2,3,4.

One problem inherent in this experiment is the fact that it requires some effort at calculation (although not an implausible amount) and so subjects may make choices of rules which are not consistent with the information received[7]. "Consistency" in this case will simply refer to choosing rules which have not yet been eliminated by the evidence. As an example, after the first round of evidence there will be eight rules out of the original sixteen which would be eliminated by the chosen combination. If, in the second question, a subject chooses one of these eight rules then they are being inconsistent. Otherwise they are consistent.

This notion of consistency allows us to formulate a test of whether the subjects are Positive Confirming or are boundedly rational subjects who are focussing on the information for one strategy at a time. The first fact to notice is that these one- strategy subjects would not have any information about the likelihoods of rules apart from the one they had been focussing on. So, for example, it would be plausible for a subject to have the rule that she chose rejected by a combination and then choose another rule which has (logically) been rejected by previous combinations but has not been focussed on by the subject. The second fact to notice is that combinations in subsequent rounds of evidence may not again reject a rule that has been logically rejected by a previous combination.

From these two facts one can see that, if a person was simply focussing on information for one rule at a time, then the choices of rule may indeed look like Positive Confirmation in that one would see repetitions of successful rules. However, these repetitions of rules would not be guided by whether the subjects were being consistent or not as they would not be able (by assumption) to keep track of their own consistency. The repetition of rules by subjects would, therefore, be independent of the consistency of the subject.

It follows that a significant difference from independence would indicate that the subjects would not simply be focussing on information for one rule at a time. A bias towards inconsistency would suggest that positive confirmation would tie in with irrationality and that consistent subjects would not tend to be Positive Confirmers. By contrast, if there was a bias towards consistency then this would tie in with previous results that showed Positive Confirmation to be a "reasonable" tactic to use (See Jones & Sugden 2001).

4. Experimental Design: Details

The experiment was carried out at the University of Dundee in the year 2002. Subjects were recruited by e-mail on campus and came from a wide range of course programmes across all years. The 87 subjects took part in groups of up to 12 at a time.

The experiment was carried out using a questionnaire and followed the general structure in section 3. Before starting the main experiment, subjects were given full instructions about the nature of the tasks involved, how questions and rounds of evidence were to be

answered and how the random lottery device worked. These instructions were given orally with some visual aids that illustrated the workings of the experiment. This was followed by an example of two questions with a round of evidence in between- although using different symbols for the combinations and a different set of rules from those used in the experiment itself. Subjects worked through this example with the help of further oral instructions. In composing the instructions care was taken not to suggest that there was a right way to do the task or to suggest that any strategy was preferable to any other.

After this each subject answered two multiple-choice questions which were designed to test understanding of the task and scoring system. In the first question they were tested on whether they understood the meaning of the idea that a combination was allowed by a rule. In the second question the subject was asked to imagine that they had completed the experiment and had successfully stated the correct rule three times. They were then asked what chance they would have of getting the money prize of £20. In general, the answers to these test questions indicated a high level of understanding with the overwhelming majority getting both answers correct.

There were four rounds of evidence in the experiment, each of which consisted of a choice between two combinations. Which round of evidence appeared in which order in the experiment was randomly varied so as to prevent order effects. Table 1 gives the base rounds before the order was randomised. In table 1 column 1 gives the round of evidence while column 2 gives the combination of symbols to be chosen in each round. For a given round of evidence, Columns 3 and 4 give the attributes of each symbol in each combination. Examination of these attributes will reveal that they correspond with the framework set out in section 2.

The experiment proceeded as follows. The subjects were given all possible rules in the experiment on a sheet of paper, together with their main questionnaire. Each of these rules was also was written on a separate piece of paper and placed in an envelope. These sixteen envelopes were then shuffled and one was dealt out for each person in the experiment. The subjects were told which their envelope was and that their task in the experiment was to discover which rule was in this envelope.

After the instructions the experiment started and the subjects went through each question and round of evidence simultaneously. First of all the first question was asked and subjects had to write down which rule they thought was in their envelope. Next they were faced with their first round of evidence. Subjects chose a combination from the two available by ticking the box next to the one they wished to choose. Once this had been done they were told to halt. The experimenter went around the subjects with their envelopes and told them whether the combination they had chosen was allowed or not by the rule in their envelope. The subjects were then given a chance to record this. After this the subjects went on to the second question and then the second round of evidence and this was repeated through the remaining questions and rounds of evidence. After the fifth question the experiment finished. The experimenters then went around each of the subjects and rolled a die to determine whether the subject had won any money and the money was then paid to the subjects.

5. Results

Table 2 gives an initial analysis of the sample, concentrating on the questions and whether they were answered in a manner that corresponds to Bayesian reasoning. This

gives an overview of the level of Bayesian rationality in the sample. There are two ways of measuring this. One is through looking at the proportions who answered the question correctly (i.e. matched with nature). The second column gives the expected (Bayesian) proportions of correct answers for each question while the third column gives the actual numbers of correct answers. Using the binomial test one can see that, up to the third time of asking the question there is no significant deviation from the expected proportion. However, after this point there are significant differences between expectation and the numbers answering correctly.

The second way of measuring the Bayesian rationality of subjects is to look at the consistency of choices, which are given in the fourth column. This column gives the number of consistent choices for each question. For the first question there has been no evidence received so all responses are consistent. After that it can be seen that consistency is initially quite high but drops sharply until by the fourth question the majority (just) are inconsistent. In the fifth question there is a slight revival, suggesting that the subjects made a special effort for the final question. This drop in consistency suggests that there was an accumulation of mistakes that diminished the number of correct answers in the later stages.

The main results are given in table 3 that shows the evidence for positive confirmation both in choice of evidence and in choice of rules. The first column gives the numbers of the rounds of evidence and the corresponding question numbers on either side of the round of evidence. Round 1, for example, is that between the first and the second questions. The second column gives the total numbers of subjects who chose combinations that were allowed by the rule which they had chosen in the previous question. As can be seen the numbers here show overwhelming evidence of positive confirmation in the choice

of evidence in all rounds and a violation of symmetry invariance in both strategies and labelling.

One important question is how closely linked are Positive Confirmation in the choice of evidence and the repetition of rules if combinations are allowed by the hidden rule[8]. Since it is assumed that the latter will only occur if the combination chosen is allowed by the hidden rule then the test sample is restricted to those cases. The comparison takes the form of a χ^2 statistic on a cross tabulation of the occurrence of Positive Confirmation in the choice of evidence and rule repetition. In general (apart from round 3) there is a positive significant relationship between the two. It is not certain why there is no relationship in round 3, although it is possible that the lack of consistency in question 4 may have caused the relationship to break down.

One interesting result not reported in the table is the behaviour of those Positive Confirmers in the choice of evidence when a combination was *not* allowed by the hidden rule. In general, very few subjects then went on to repeat the rule from the previous question in the next question (2 subjects in round 1 and 1 subject in each of rounds 2,3 and 4). This suggests that, when faced with disconfirming evidence, positive confirmers did tend to use it properly and reject the rule being tested.

The fourth column acts as a guide for the next four columns and simply states whether the same rule is repeated in the question before as in the question after the round of evidence i.e. whether there is positive confirmation in the choice of rules. In columns 4 to 7 the sample is cut back to include only those who were positive confirmers in their choice of evidence and whose chosen combination was allowed by the hidden rule.

Column six gives the total numbers of consistent *and* inconsistent subjects while column seven only looks at those who were consistent. Column five gives the test proportions for columns six and seven. As mentioned earlier, these are based on the idea that, once a subject has updated their beliefs in light of the evidence, they will assign equal probabilities to the rules in line with symmetry invariance. In column six we can see that this means that for rounds 1 and 2 there are significantly higher numbers of people who positive confirm than one would expect from the test proportions. It can also be seen that there are substantial numbers who Positively Confirm in the fourth round. However, round 3 seems to be anomalous in that the number of Positive Confirmers in rule choice is significantly *lower* than the expected proportion.

However, it will be argued that this anomaly is an artefact of the lack of consistency of a large part of the population. This fits in with previous evidence that suggests that Positive Confirmation is a heuristic used by largely consistent (if not fully rational) people. If we look at column seven where all inconsistent subjects have been eliminated from the sample we can see that in all rounds the proportion of Positive Confirmers has substantially increased. In round 3 the proportions are roughly as expected implying that high levels of inconsistency influenced the anomalous negative result. Looking at the other rounds we see that there are also significant numbers of Positive Confirmers in the first two rounds. In the fourth round the restriction to consistent subjects is associated with a large increase in the proportion who positively confirm.

The eighth column tests for the independence of Positive Confirmation in the choice of rules from consistency. As explained earlier, this not only explains the interaction in the preceding paragraph but also tests for Positive Confirmation against bounded

rationality in focussing on one rule. It will be seen that for the second, third and fourth rounds there are significant deviations from independence, suggesting that the Positive Confirming behaviour is not the result of bounded rationality based on information for just one strategy. This also confirms the interaction observed in the previous columns.

6. Discussion

The aim of this paper was to test for the existence of the Positive Confirmation bias in a controlled game- theoretic setting. This allowed for a test of Positive Confirmation by demonstrating systematic behaviour that was not explained by Bayesian updating. It was decided that the best way to test for this was through a controlled test of the mechanisms of learning using a game against nature with highly symmetric payoffs, updating and evidence. In addition, the game was adapted so that Positive Confirmation in the search for evidence could be tested separately from Positive Confirmation in the choice of rules.

It was found that Positive Confirmation exists in both cases. When subjects were searching for evidence this was particularly convincing as there was overwhelming evidence for this phenomenon as against the extended notion of symmetry invariance. There was also strong evidence for Positive Confirmation in the choice of rules although this has to be qualified by the failure of one round of evidence to demonstrate this effect. It was also discovered that subjects who Positively Confirmed in the choice of evidence were more likely to Positively Confirm in the choice of rules indicating that there is a link between the two.

In general, the results of this paper do seem to corroborate previous experiments on the subject. It was found, as before, that Positive Confirmation exists and that it is not just the result of a lack of understanding by the subjects. Updating with Positive Confirmation strategies is carried out in a logical manner with a consistent pattern of reasoning. Being told whether a combination is allowed or not by the hidden rule *does* have an effect on the choices made. It is rare for a rejected rule to be repeated.

One striking and novel aspect of this is the role of consistency in whether subjects Positively Confirmed or not. It may be thought that subjects who were less consistent would be more likely to Positively Confirm. However this is not the case. Instead those who were consistent were *more* likely to Positively Confirm than those who were not. In fact it was the high level of *inconsistency* in round 3 that made sure that Positive Confirmation in the choice of rules was statistically significant in the wrong direction. It was also this fact that identifies Positive Confirmation behaviour as distinct from boundedly rational behaviour that focuses on information for one strategy at a time.

This also adds something to the question of why people Positively Confirm. It seems that, in many respects, Positive Confirmers have much in common with Bayesian updaters (although they have been shown to violate expected utility in some circumstances). This would help to explain why Bayesianism as a more effective method of updating does not replace it. People Positively Confirm because, quite often, it *does* lead to the right answer although it can also lead to subjects making suboptimal decisions. Positive Confirmation however is not obviously wrong and to subjects, in a world prone to random shocks, any failure of Positive Confirmation to achieve the correct answer may not result in its abandonment but may instead be attributed to other causes.

The success of this experiment lies in showing that Positive Confirmation is a systematic effect demonstrated by subjects in a game against nature. It also raises questions about other theories of learning such as fictitious play, reinforcement learning and experience-weighted attraction. By construction the myopic theories of learning cannot account for the results here as they only apply to situations where one just gathers information from the strategy one has chosen. However, even if we were in the latter situation, these theories do not adequately model Positive Confirmation in the choice of strategies. It follows that the next task will be to create and test a full- scale model of Positive Confirmation against the current myopic theories of learning to test for their empirical adequacy.

Bibliography

Bacharach M. (1993) "Variable Universe games" in *Frontiers of Game Theory* edited by Binmore K., Kirman A. & Tani P. MIT Press, Cambridge Massachusetts

Borgers T. & Sarin R. (1997) "Learning through reinforcement and replicator dynamics" *Journal of Economic Theory* 77,1-14

Camerer C. (2003) *Behavioral Game Theory: Experiments in Strategic Interaction* Priceton University Press, Princeton New Jersey

Camerer C. & Ho T. (1999) "Experience- weighted attraction learning in normal- for games" *Econometrica* 67, 827-74

Cheng P.W. & Holyoak K.J. (1989) "On the natural selection of reasoning theories" *Cognition* 31, 187- 276

Cosajus A. (2000) "Focal Points in Framed Strategic Forms" *Games and Economic Behavior* 32, 263-291

Festinger L. (1957) *A Theory of Cognitive Dissonance*. Stanford CA; Stanford University Press

Fischhoff B. & Beyth- Marom R. (1983) Hypothesis Evaluation from a Bayesian Perspective" *Psychological Review* 90, 239- 260

Fudenberg D. & Levine D.K. (1998) *The Theory of Learning in Games.* Cambridge Mass. London MIT Press

Harrison G. & Rutstrom E. (2001) "Doing it both ways- experimental practice and heuristic context" Behavioral and Brain Sciences 24, 413-414

Harsanyi J.C. & Selten R. (1988) *A General Theory of Equilibrium Selection in Games* Cambridge MA London England MIT Press

Jones M.K.. (2003) "Positive Confirmation in Rational and Irrational Learning" Dundee Discussion Paper in Economics no. 150

Jones M.K. & Sugden R. (2001) "Positive Confirmation bias in the Acquisition of Information" *Theory and Decision*, 50, 59- 99

Levine M. (1966) "Hypothesis behavior by humans during discrimination learning" *Journal of Experimental Psychology*, 71, 331- 338

McKenzie, C. (2004). Hypothesis testing and evaluation. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 200-219). Oxford: Blackwell.

Manktelow K.I. & Over D.E. (1993) (eds) *Rationality: Psychological and Philosophical Perspectives*. London: Routledge.

Oaksford M. & Chater N. (1994) "A Rational analysis of the Selection Task as Optimal Data Selection" *Psychological Review,* 101, 608-631

Rabin M. & Schrag J.L. (1999) "First Impressions Matter: A model of Confirmatory Bias" *Quarterly Journal of Economics*, 114, 37-82

Schelling T. (1960) "The Strategy of Conflict" Harvard University Press, Cambridge, Massachusetts

Sugden R. (1995) "A Theory of Focal Points" Economic Journal 105, 533-550

Wason P.C. (1960) "On the failure to eliminate hypotheses in a conceptual task" *Quarterly Journal of Experimental Psychology*, 12, 129 - 140

1) Left symbol is black

2) Right symbol is black

3) Left symbol is white

4) Right symbol is white

5) Left symbol is large

6) Right symbol is large

7) Left symbol is small

8) Right symbol is small

9) Left symbol is a cross

10) Right symbol is an cross

11) Left symbol is an octagon

12) Right symbol is an octagon

13) Left symbol has a border

14) Right symbol has a border

15) Left symbol does not have a border

16) Right symbol does not have a border

During this part of the experiment we ask you to test one of the two combinations below.

Tick one of the two boxes on the right to indicate which combination you would like to test:

Combination 1:

Combination 2:

Now please wait for the experimenter to tell you whether the combination you have chosen is allowed by the rule in the envelope

The combination chosen is (allowed/ is not allowed) by the rule in the envelope. *(Delete as appropriate)*

Table 1: Attributes used in Symbols in the experiment

Round of Evidence	Combination	Left Symbol	Right Symbol
1	1	Cross, Large, Black, No Border	Octagon, Large, White, Border
	2	Octagon, Small, White, Border	Cross, Small, Black, No Border
2	1	Cross, Small, Black, Border	Octagon, Small, White No Border
	2	Octagon, Large, White, No Border	Cross, Large, Black, Border
3	1	Octagon, Large, Black, Border	Cross, Large, White, No Border
	2	Cross, Small, White, No Border	Octagon, Small, Black, Border
4	1	Octagon, Small, Black, No Border	Octagon, Large, Black, No Border
	2	Cross, Large, White, Border	Cross, Small, White, Border

Table 2 - Numbers of correct answers and consistency for each question

Question	Test Proportion	Total (87)	Consistency
1	0.0625	6	N/A
2	0.125	16	77
3	0.25	18	57
4	0.5	22[a]	43
5	N/A	46	46

[a]Significant at the 5% level of significance

Table 3 –Positive Confirmation in the use of evidence and in choosing rules

Round (Questions)	Total PC Evidence	PC Evid v Repetition	Same Rule?	Test Prop/ Consistent	Total Allowed	Consistent Allowed	Consistent vs PC Rule
1 (1,2)	73	5.973[a]	Yes	0.125	13[a]	13[a]	2.543
			No		23	19	
2 (2,3)	66	13.72[a]	Yes	0.25	13[a]	12[a]	5.407[a]
			No		22	12	
3 (3,4)	64	3.459	Yes	0.5	10[a]	8	5.1[a]
			No		24	9	
4 (4,5)	65	32.29[a]	Yes	N/A	22	17	6.348[a]
			No		12	4	

[a] Significant at the 5% level of significance

[1] Positive Confirmation should be distinguished from the *confirmatory bias* of Rabin & Schrag (1999) as well as *Cognitive Dissonance* (Festinger 1957). Positive Confirmation is a bias in the *search* for data while the latter two are concerned with how the data is changed either by an innate information processing bias or by one's own preferences.

[2] All of this casts doubt on the interpretation by Fischoff and Beyth- Marom (1983), Oaksford and Chater (1994) and McKenzie (2004) (amongst others) that much of positive confirmation can be explained by Bayesian updating.

[3] Unlike what is done in many of the experiments covered in Camerer (2003). A quantitative model would also involve extra assumptions which would be jointly tested with positive confirmation and so reduce the effectiveness of the test.

[4] In principle, using games against nature should have no effect on "unsophisticated" Bayesian updating or indeed most myopic learning theories used in the literature (evolutionary and social learning models aside) as these theories do not take into account the opposing subject's payoffs or beliefs.

[5] The design (although not the interpretation) of the experiment is similar to the rule discovery task put forward by Levine (1966). The two are similar in many ways although it has been changed considerably to allow for controls in the form of incentives and prior probabilities.

[6] Note that a=1 is not included as there is no prospect of testing positive confirmation in the first question while a =5 is nonsensical in terms of the alternative hypothesis.

[7] It may be suggested that such inconsistency could be eliminated by repetition. However this is not followed here as there are potentially useful linkages between consistency and positive confirmation. Another reason concerns doubts about the efficacy of repetition. As Harrison and Rutstrom (2001) point out there may be different heuristics at work in one shot as opposed to those in repeated games.

[8] This is *not* the same as Positive Confirmation in the choice of rule. The latter *assumes* the combination selected is selected through Positive Confirmation in the choice of evidence.