



University of Dundee

Dundee Discussion Papers in Economics 150

Jones, Martin

Publication date:
2003

[Link to publication in Discovery Research Portal](#)

Citation for published version (APA):

Jones, M. (2003). *Dundee Discussion Papers in Economics 150: Positive confirmation in rational and irrational learning*. (Dundee Discussion Papers in Economics; No. 150). University of Dundee.

General rights

Copyright and moral rights for the publications made accessible in Discovery Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



University
of
Dundee

Dundee Discussion Papers in Economics

Positive Confirmation in Rational and Irrational Learning

Martin Jones

Department of
Economic Studies,
University of Dundee,
Dundee.
DD1 4HN

Working Paper
No. 150
September 2003
ISSN:1473-236X

Positive Confirmation in Rational and Irrational Learning¹

Martin Jones²

Abstract

An experiment is reported which tests for positive confirmation bias in a setting in which a person can make decisions about which evidence to select to test a rule. The experiment reveals strong evidence of positive confirmation bias and a tendency not to choose the Bayesian- optimal selection of information. There is also evidence that the information, once selected, is used in a pattern of reasoning which, while sub- optimal, is internally coherent. These results are consistent with and expand on previous results on the subject.

¹ The author would like to acknowledge the financial support of the Economic and Social Research Council (Award no.R000223606) and the assistance of Qing Lu and Mara Violato in helping with the experiment.

² Department of Economic Studies, University of Dundee, Dundee. DD1 4HN.

1 Introduction

Over the past twenty years there has been a gradual increase in interest in the theory of learning in economics. While economists once assumed that agents could be modelled as rational optimisers, now there are attempts to explain how boundedly rational individuals actually learn and whether this learning process converges on optimising behaviour. (c.f. Borgers T 1996 and Cubitt & Sugden 1998) In general the theoretical results have been mixed (c.f. Fudenberg & Levine 1998 for an in- depth look at the subject). Empirically, as well as testing specific theories of learning (see Erev & Roth 1998) there has been some empirical research into learning biases (e.g. on the representativeness heuristic- Grether 1980, 1992).

One bias which has not received much interest from economists is the *positive confirmation bias*. This is a tendency, when testing an existing belief, to search for information which could confirm that belief rather than information which could disconfirm it. When assessing information in a framework of rational learning both types of evidence are relevant. There is a bias if more than reasonable effort is devoted to searching for confirming information. If positive confirmation is indeed a significant part of the processes used by humans in learning then we can expect it to have a significant impact on the decisions made by economic agents when learning.

To date there has only been one experimental study made of this bias in the economics literature (Jones & Sugden 2001), although there is a large literature on positive confirmation in the psychology literature (e.g. Manktelow & Over 1993; Oaksford & Chater 1994; Cheng & Holyoak 1989). Jones and Sugden's paper adapted an experiment commonly used within psychology- the Wason Selection Task- by placing it within a Bayesian decision theory framework where the costs, benefits and prior probabilities of acquiring information were made explicit.

There were three principal conclusions derived from that experiment. One was that there was indeed such a thing as positive confirmation. While irrationality was not as pronounced in that experiment as similar experiments in the psychology literature, there was a general tendency towards positive confirmation in the data. Second, it was found that significant numbers of people would choose positive confirming, but irrelevant,

information even at a cost to themselves. Finally a new form of positive confirmation was discovered. This is positive confirmation in the use of information, where information which is interpreted as confirming a belief increased subjects' confidence in the truth of the belief even if, from a Bayesian point of view, that information had no value.

The Jones and Sugden experiment concentrated on demonstrating that positive confirmation could result in obviously irrational behaviour. For this reason, in the experiment, there was only a choice between relevant and rational behaviour on one hand and irrelevant and irrational behaviour on the other. However, this split only allows for one particular extreme category of irrationality caused by positive confirmation. The aim of the experiment in this paper is to test positive confirmation where information may be relevant, but not rational.

This broadens the scope of the positive confirmation bias from a relatively narrow situation, where positive confirmation results in irrelevant information being used, to a situation where information is relevant, but the choice of that information does not maximise expected utility. It follows that, as with the Jones and Sugden paper, the experiment here reveals a pattern of information-gathering which contravenes the fundamental principles of Bayesian decision theory.

2. Rule Discovery Experiments

The experiment which is reported in this paper is related to a group of experiments known as *rule discovery* experiments. A classic example of this is Wason's (1960) Numbers Game experiment. In this experiment the experimenter asked each subject to guess a hidden rule about triads of numbers (e.g. "Three numbers in ascending order of magnitude"). The subjects were then given one triad ("2,4,6") and were told that this followed the hidden rule. They were then told to make up their own triads of numbers to test the rule. For each triad they gave they had to write down the rule they thought they were testing. Then they were told whether the triad followed the rule or not. When the subject was happy that they knew the rule then they announced it and were told if they were correct or not.

The results of these experiments were interesting in that the original rule ("Three numbers in ascending order of magnitude") seemed to be too general for the subjects to easily discover. Most subjects did not manage to discover it first time round. More important was the process used. When subjects had formulated a rule, they tended to issue triads which were consistent with this rule rather than inconsistent with it. This suggested that they were suffering from positive confirmation bias. Subjects seemed to test a rule by enumerating examples of that rule rather than by attempting to find disconfirming evidence.

The positive confirming aspect of this was highlighted by Tweney et al. (1980) who, instead of putting the experiment in terms of testing for one rule, reframed the experiment as a test of two mutually exclusive and exhaustive rules (called "DAX" and "MED"). The subjects had to find out the rule "DAX" but were told whether a piece of evidence obeyed the rule "DAX" or "MED" rather than just whether it was in "DAX" or not. This was found to increase the number of disconfirmations because the evidence was not specified in terms of whether it did or didn't obey the rule "DAX".

There have been attempts to expand the scope of these rule discovery experiments by creating simple scientific environments which one can explore and attempt to find the rules by which they operate (e.g. Mynatt, Doherty and Tweney 1978). In the cases which have been carried out the experimenters have claimed a pronounced tendency towards positive confirmation.

There have been many variations on the rule discovery tasks to date. Some of them have attempted to vary the experimental conditions to test the robustness of positive confirmation. An example of this is in Gorman & Gorman (1984) where the experimenters gave the subjects advice about how to test rules, with one group of subjects being given advice biased towards confirmation and the other being given advice biased towards disconfirmation. In general, it was found that this was effective in increasing the amount of disconfirming behaviour and also resulted in more people finding the hidden rule.

However, they are all problematic, from the point of view of experimental economics because these experiments are not decision- making tasks in the economics sense. Subjects are simply told to discover the rule. There is no structure of payoffs, prior

probabilities or costs of acquiring information. This is a problem because, as Klayman and Ha (1987) point out, different prior probabilities for a rule could result in positive confirmation being used quite rationally to test for a rule. To be more precise, for a highly specific rule which is a subset of most other possible rules it may be reasonable, for some prior probabilities, to search for evidence which, if true, would disconfirm the rule. However, for a very general rule, this would not hold if most other possible rules are subsets of this rule. In this case positive confirmation would be rational. Any decision theory version of a rule discovery task would need to specify the prior probabilities of all possible rules.

Experiments purporting to show positive confirmation have also been criticised by Evans (1972) on the grounds that the subjects could be suffering from a *matching bias* rather than positive confirmation. The matching bias is a perceptive bias resulting from an inability to understand the problems as set in experiments. As a result, subjects in rule discovery tasks simply choose those pieces of information which correspond most closely with the rule being tested. There is no testing of the rule- simply an automatic choice of information which corresponds most closely to the rule being tested. This means that the evidence of Positive Confirmation which has been accumulated to date is simply evidence of subjects not properly understanding the task. This paper, amongst other things, aims to test this hypothesis.

3. Experimental Design: Principles

The basic aim of this experiment is to test for positive confirmation in a setting where individuals make information acquisition decisions with real financial consequences. In the experiment subjects are given a set of rules and "combinations". One of these rules is selected out as a "test rule" and the subject has to test this rule by using one of the possible combinations in order to find the unknown rule which is also one of these rules. This means that, while the experiment is related to the rule discovery tasks, it differs in that rather than have the subjects generate their own rules and "triads", they are given a set of rules and combinations.

In this experiment the subjects are asked to imagine that the rules are about balls in a bag. There are ten such balls in each bag and they may be either red or blue. Rules are descriptive of the numbers of balls in the bag such as:

"There must be at least 2 Red Balls"

In this experiment the rules concentrate on the minimum number of red balls in the bag. Suppose that there is a set \mathbf{V} of such rules. The set \mathbf{V} is "nested" in the sense that all members of \mathbf{V} are strictly implied by or imply other members. For example:

If the statement "There must be at least 2 red balls" is in the set \mathbf{V} then it is implied by "There must be at least 3 red balls" which is also in \mathbf{V} .

If we define rules which are implied by the majority of rules in \mathbf{V} as *outer* rules and rules which imply the majority of rules as *inner* rules then we can define the *middle* rule as the rule which implies and is implied by equal numbers of rules.

Such an imaginary bag, as suggested by the rules, would be tested by looking at its contents. An example of the contents of such a bag is known as a *combination*. A combination simply states the number of red balls and blue balls in such an imaginary bag e.g. "8 Blue; 2 Red". A Combination x is said to be *allowed* by a rule $V_i \in \mathbf{V}$ if V_i implies the existence of x .

For each $V_i \in \mathbf{V}$ there is a set of combinations Q_i which are allowed by V_i . Since the rules are assumed to be distinct, all $Q_i \in \mathbf{Q}$ are (strict) subsets of each other. N is the number of V_i in \mathbf{V} and is constant across tasks and subjects. It is useful to define subsets of Q_i known as *bands*. A band $B_i \in \mathbf{B}$ is a subset of $Q_i \in \mathbf{Q}$ but is not a subset of $Q_j \in \mathbf{Q}$ where Q_j is the largest strict subset of Q_i in \mathbf{Q} . Roughly, bands are the set of combinations which are allowed by a rule V_i but not by other rules which imply rule V_i .

Bands are *outside* a rule if its combinations are not allowed by the rule. Bands are *inside* a rule if its combinations are allowed by the rule.

Suppose that there is an unknown rule $U \in \mathbf{V}$ which is the rule which the subjects are trying to find. U is selected from \mathbf{V} by a random process, independently for each task with a probability $1/N$. In addition the experimenter chooses another rule $H \in \mathbf{V}$ which is

to be tested by the subjects. The aim of the experiment is for each subject to discover whether $U = H$ or not. In order to help them they are given the choice of a combination and, once they have chosen a combination, are told whether or not this combination is allowed by the rule U . Subjects are then told whether the combination they have chosen is allowed by the rule U or not.

Once the subjects have been given the information from the combination, then they have to make a judgement as to whether test rule H is the same as the unknown rule U . When the judgement has been made on all the tasks the subject is then rewarded according to whether the judgement is correct and whether $U=H$ or not.

In this experiment there are no costs involved in acquiring information. However, the rewards differ according to the judgement made. The subject starts with an endowment of points and then gains points according to the accuracy of her judgement. If it is the case that $U=H$ and the judgement made is that $U \neq H$ or if it is the case that $U \neq H$ and the judgement is made that $U=H$ then the subject is incorrect and gains no additional points. If it is the case that $U \neq H$ and the judgement is made that $U \neq H$ then the subject gains r additional points. If it is the case that $U=H$ and the judgement is made that $U=H$ then the subject gains $(N-1)r$ additional points. This counterbalances the advantage of saying that $U \neq H$ because of the a priori equal likelihood of each rule being chosen.

At the end of the experiment one of the tasks in the experiment is picked out to be played for real using a die. After this, the number of points won by the subject in that task is calculated. The subject enters a lottery in which the probability of winning a money prize is proportional to the total number of points credited to her. This means that, if the subject is rational in the sense of expected utility (and the money prize is preferred to nothing), then she will seek to maximise the expected number of points scored in each task.

This binary lottery system has been widely used in experimental economics as a means of inducing risk- neutral preferences. However, Selten, Sadrieh and Abbink (1999) have found evidence that, in fact, subjects are at least as risk averse with respect to payoffs that are denominated in terms of lottery tickets as they are with respect to payoffs that are denominated in money. In this case we cannot assume that subjects in the experiment are risk- neutral with respect to points. However, as will be explained in the next section, this

experiment does not rely on risk neutrality in the formulation of its null hypotheses. This means that a failure of risk neutrality would not ruin the tests for positive confirmation bias.

4. Bayesian Analysis

It is not the objective of this experiment to test whether the subjects are rational Bayesians. The use of the prior probabilities, the payments etc. are to achieve experimental control over the experiment rather than because the subject is expected to behave in a Bayesian fashion. Instead the aim is to test for a particular systematic bias, namely positive confirmation bias. For this reason the null hypotheses which we will give in the next section allow for the widest possible range of behaviour which is not consistent with positive confirmation. This includes Bayesianism but is not identical to it.

However, it is of interest to investigate what the optimal behaviour for a Bayesian subject would be in this situation. In this section we make the assumption that subjects do behave like Bayesian optimisers to the extent of behaving correctly under the binary lottery system.

It can be shown (see Appendix 1) that the payoff from testing a combination from the bands inside H is P_{inside} where:

$$(1) \quad P_{\text{inside}} = r / N [1 + (j - 1) / (N - 1)]$$

And the payoff from testing a combination from the bands outside H is P_{outside} where:

$$(2) \quad P_{\text{outside}} = r / N [1 + (N - j + 1) / (N - 1)]$$

Where $j \in \mathbf{J}$ is an index of bands going from the innermost band (i.e. the set of combinations implied by *all* the rules) outwards. It can be seen that increasing j increases the payoff from testing inside H and decreases the payoff from testing outside H. It follows that the optimal bands for picking combinations for testing H are the two neighbouring the rule H.

It is useful to define the concept of combinations or bands being *immediately* inside or outside a rule. A band $B_i \in \mathbf{B}$ is immediately inside rule $V_i \in \mathbf{V}$ iff when $Q_i \in \mathbf{Q}$ is the set of combinations allowed by V_i then the combinations in B_i are not allowed by any of the rules which imply V_i . Likewise a band $B_i \in \mathbf{B}$ is immediately outside rule $V_i \in \mathbf{V}$ iff when $Q_j \in \mathbf{Q}$ is the set of combinations allowed by all rules outside V_i then B_i is that subset of Q_j which is not allowed by V_i .

Suppose that rules are also labelled by indices from the set \mathbf{J} . If $h \in \mathbf{J}$ is the index for rule H then the optimal band to choose is $j = h + 1$ if testing a combination outside H and $j = h$ if testing a combination inside H .

Substituting into formulae (1) and (2) we get formulae (3) and (4) which show respectively the payoffs for testing for combinations in bands immediately inside and outside a rule:

$$(3) \quad P_{\text{Inside}} = r / N [1 + (h-1) / (N - 1)]$$

$$(4) \quad P_{\text{Outside}} = r / N [1 + (N - h) / (N - 1)]$$

If a rule H is an inner rule then it can be seen that $P_{\text{Inside}} < P_{\text{Outside}}$ so that it is rational to choose a combination in the band immediately outside H . If a rule H is an outer rule then it can be seen that $P_{\text{Inside}} > P_{\text{Outside}}$ so it is rational to choose a combination in the band immediately inside H . If H is a middle rule then the two payoffs are equal and one can rationally choose a combination in either band.

For those bands which are not immediately inside and outside the rule H then it can be seen through formulae (1) and (2) that, as the bands get further away from H then the payoff drops so that it becomes less rational to choose combinations in bands further away from H . The worst bands in this respect are the outermost and innermost bands. In these cases choosing a combination gives no additional information and so does nothing to increase the expected payoff.

An examination of the formulae allows us to see that the variation in payoffs is a result of variation in the number of rules excluded or included from the set of possible rules when a combination is chosen. This property will now be described using a more general notation which will be used to motivate the (more general) null hypotheses in the next section.

Define two sets \mathbf{W} and \mathbf{Z} as *distinguishing sets* where $[\mathbf{W}, \mathbf{Z}]$ is a partition of \mathbf{V} . The combination θ partitions the set \mathbf{V} into $[\mathbf{W}, \mathbf{Z}]$ where \mathbf{W} includes all those rules V_i which allow θ and \mathbf{Z} includes all those rules which do not allow θ .

In general, given that all rules are initially equally weighted in terms of expected utility, a rational subject would want to increase the number of rules which are eliminated by the combination selected, as this would decrease the number of possibilities for U in the final choice. So, supposing that \mathbf{W} includes the rule H , then the aim of the subject would be to reduce the number of rules in \mathbf{W} and (hence) maximise the number of rules in \mathbf{Z} . This would imply a policy of choosing in the immediate outside band for inside rules and choosing in the immediate inside band for outside rules. This conclusion is identical to that reached for the payoffs, but is done in far more general terms.

However, as was mentioned beforehand, it is not intended that this paper should be a test of Bayesianism. The latter analysis simply sets out a framework for choice which can be applied to more general situations.

Null Hypotheses in Experiment

i) Choice of bands

This experiment is cast in terms of stochastic choice, for given values of the parameters N , r , H and a given set of rules. A function $\pi(\cdot)$ is defined from the set \mathbf{B} to the interval $[0,1]$. For each $B_i \in \mathbf{B}$, $\pi(B_i)$ is defined as a decision probability. It is the probability that the subject chooses a combination in band B_i when attempting to find a combination to test rule H . Stochastic variation in choice is to be interpreted as resulting from imprecision or errors in individuals' preferences or beliefs (Loomes and Sugden 1995).

By allowing that some of subjects' behaviour is partly random, in testing this paper will look for patterns of behaviour that cannot be explained by randomness.

Suppose that in task S a subject picks combination θ in band $G \in \mathbf{B}$. This also creates two distinguishing sets $[E, F]$ for rules, partitioning \mathbf{V} . Suppose that in task S' a subject picks combination θ' in band $G' \in \mathbf{B}$. θ' also creates distinguishing sets $[E', F']$ for rules also partitioning \mathbf{V} . Further suppose that E' has the same number of rules as F while F' has the same number of rules as E . In this case we refer to the bands G and G' as being *symmetric*.

Assume that H is the rule to be tested in S while H' is the rule to be tested in S'. H and H' are also described as symmetric when the band immediately outside H is symmetric to the band immediately inside H' and the band immediately inside H is symmetric to the band immediately outside H' . Two tasks S, S' which have their test rules H, H' as symmetric are known as symmetric tasks. If H and H' are symmetric then combinations θ and θ' selected from symmetric bands G and G' will eliminate exactly the same number of rules and will have exactly the same chance of eliminating H or H' respectively. Because of the identical numbers of rules eliminated by θ and θ' we shall say that θ provides the same information in S as θ' does in S' i.e. that they are *isomorphic* to each other. It follows that G and G' are also isomorphic to each other since all combinations in the two bands are isomorphic to each other. This means that for all symmetric G, G' then a person making a decision should have $\pi(G) = \pi(G')$. This condition will be described as *symmetric neutrality*³.

When testing the results of the experiment it will be necessary to test for combinations in more than one symmetric band. If $X \subset \mathbf{J}$ is the set of indices for bands $A_i \in \mathbf{B}$ which have been selected for testing in task S and $X' \subset \mathbf{J}$ is the set of indices for symmetric bands $A_i' \in \mathbf{B}$ in task S' then the null hypothesis with aggregated bands is: $\sum_{i \in X} \pi(A_i) = \sum_{i \in X'} \pi(A_i')$

If we assume that decision bands $A_1, A_2, \dots \in \mathbf{B}$ are bands inside the test rule H in task S then we can construct an alternative hypothesis for positive confirmation. If positive

³ Symmetric neutrality is satisfied by any stochastic theory of choice (c.f. Loomes & Sugden 1995) where the "core" of the choice model is a theory of dominance respecting preferences over lotteries (e.g. expected

confirmation exists then one would expect there to be a larger number of people choosing in bands inside the test rule H rather than outside. In the symmetric task S', the symmetric bands $A_1', A_2' \dots \in \mathbf{B}$ would be outside the test rule H' so we would expect there to be fewer choices in those bands. For individual bands therefore the alternative hypothesis would be $\pi(A_i) > \pi(A_i')$ while for groups of bands the alternative hypothesis would be $\sum_{i \in X} \pi(A_i) > \sum_{i \in X'} \pi(A_i')$

In the experiment there are tasks which are identical except for differences in the combinations which could be selected out of a given band. Suppose that D and D' are two such tasks with the bands which are tested being $K \in \mathbf{B}$ and $K' \in \mathbf{B}$ where K and K' are the same bands in D and D'. However, since the combinations are in the same bands then it follows that there is no informational difference between picking a combination from K and a combination from K'. There is only a difference in labelling and the two questions are isomorphic with each other. Given the definitions above, for individual bands K and K', the null hypothesis is $\pi(K) = \pi(K')$. Likewise, if $K_1, K_2, K_3 \dots \in \mathbf{B}$ are bands in D and $K_1', K_2', K_3' \dots \in \mathbf{B}$ are the same bands in D' then we can define the null hypothesis as $\sum_{i \in X} \pi(K_i) = \sum_{i \in X} \pi(K_i')$ where $X \in \mathbf{J}$ is the index set of bands selected for testing.

In this case there is no "direction" in which we may expect any differences between the bands in the two tasks. For this reason the alternative hypotheses are those of inequality i.e. $\pi(K) \neq \pi(K')$ for bands K and K' and $\sum_{i \in X} \pi(K_i) \neq \sum_{i \in X} \pi(K_i')$ for a collection of bands.

ii) Judgements about the test rule

As well as examining what evidence a subject chooses when testing H, it is also interesting to find out whether there is a bias in whether subjects judge H the same as U or not. Is there an excessive tendency for positive confirmers to judge H the same as U?

This question is made difficult by the fact that once positive confirmers have selected the evidence and have become non- Bayesian Positive Confirmers then it must be

utility) . This solves the problem of risk aversion because the isomorphism between combinations means that they must be treated the same under this broad class of theories.

doubted whether one can assume that the subject does behave like a Bayesian optimiser when using information. It is assumed therefore, as a null hypothesis, that the subject's choice of whether they believe that $H=U$ or $H\neq U$ is independent of whether they have been told whether the combination they have selected is allowed by U or not. Positive Confirmation would then require that there should be a lack of independence and that this should be biased towards the subject saying $H=U$ when the combination selected is allowed by U .

It should be noted that this null hypothesis is in fact far stronger than that required by Bayesianism. Bayesianism would simply require that the proportion of positive confirming subjects stating that $H=U$ when the combination is allowed by U should be equal to the posterior probability of $H=U$. If H is an inner rule and $N>2$ then this probability will be less than 0.5^4 .

6. 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 they came from a wide range of course programmes across all the years. The 111 subjects took part in groups of up to 8 at a time.

In the main part of the experiment, the subject faced a series of six tasks, each of which had the general structure laid out in section three. Before starting these tasks, subjects were given full instructions about the nature of these tasks, about how points were scored and how points were converted into money prizes. These instructions were given orally, with some visual aids to illustrate various points. This was followed by an example of the task. Subjects worked through this task 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 the scoring system. In the first test question they were tested on their understanding of which combinations were allowed by a rule. In the

⁴ An attempt to test for this using symmetric responses across questions is not possible since the numbers of non- positive confirming choices of combinations which are allowed by U is too low for statistical significance.

second test question the subject was asked to give the number of points a person would get for an imaginary previous playing out of the experiment. If a subject got the questions wrong then they were given help. In general, the questions indicated a high level of understanding with more than 95% answering the questions correctly first time.

Each task involved a list of $N = 5$ rules. These rules were identical for each task in the experiment. A list of the rules and combinations is given in table 1 below.

Table 1: Rules and Combinations used in the experiment

Rules used in Experiment	Combinations	Combinations Var.
Rule (i) : There must be at least 2 Red balls	1) 10 Blue; 0 Red	1) 9 Blue; 1 Red
Rule (ii) : There must be at least 4 Red balls	2) 8 Blue; 2 Red	2) 7 Blue; 3 Red
Rule (iii): There must be at least 6 Red balls	3) 6 Blue; 4 Red	3) 5 Blue; 5 Red
Rule (iv): There must be at least 8 Red balls	4) 4 Blue; 6 Red	4) 3 Blue; 7 Red
Rule (v) : They must all be Red balls	5) 2 Blue; 8 Red	5) 1 Blue; 9 Red
	6) 0 Blue; 10 Red	6) 0 Blue; 10 Red

As can be seen, the rules are listed in order going from the most general (Rule (i)) to the least general rule (Rule (v)). In the column next to them are six combinations, one each from each "band" in between the rules. It should be noted that these combinations are evenly spaced out in numerical terms. These combinations are the combinations used in four of the questions in the experiment. The column next to this gives variant combinations used in two of the questions in the experiment. Note that these also one from each band between the rules. An example question in Appendix 2 gives the general layout of how these rules are presented and how the test rule is introduced.

For each task and each participant there was an envelope which contained a slip of paper with one of the five rules on them. For each task, the envelope was dealt from a pack of 10 envelopes containing two copies of each rule. The subject was told that they had to find out whether the rule to be tested (H in the notation used previously) was the same as the rule in the envelope (U in the same notation).

The subjects were then told to choose one of the six combinations and to circle it, also writing down which one they had chosen. Once everyone had written down the combination, then the experimenter went around with the envelope for the subject and task and told the subject whether the combination they had chosen was allowed by the rule in the envelope. No other information was given about the rule in the envelope. The subjects were advised to record this information on their question sheet (see the sample questionnaire in appendix 2)

Once they had been told whether the combination they had chosen was allowed by the rule in the envelope, the subjects were asked a yes/ no question as whether the rule in the envelope was the same as the test rule. When they had answered this question then they went onto the next task.

Notice that the subject did not get any feedback on her actions until she had completed all the tasks. This ensured that her answers to later questions were not influenced by the answers to earlier questions. Some cross- task learning is inevitable in such multiple question experiments but since the order of the tasks was randomised there were no systematic effects.

The six tasks in the experiment were divided into three pairs of symmetrical problems. These pairs of problems were identical in combinations and rules but the rule to be tested in one question in the pair was symmetrical to the rule being tested in the other. The first pair of tasks represented the "base" tasks. In these two tasks the rules being tested were rules (iv) and (ii) (using the notation in table 1) which are symmetric rules. These tasks will be referred to as "Question 1" and "Question 2" respectively (although the actual order of the questions was randomised). This allows a basic test of positive confirmation where choosing positively confirming combinations does have some informative value but is irrational.

The second pair of tasks tests the matching hypothesis. It is possible that the subjects may be choosing a combination simply because it replicates the number of red balls in the test rule and this may explain any positive confirming behaviour in questions 1 and 2.

However, with similar behaviour in this pair of tasks this explanation is not plausible.

While the rules, test rules and bands are the same, the combinations selected out of the bands come from the variant combinations in table 1 and do not have the same numbers

of red balls as in the rules. Subjects therefore will not be able to choose through matching but will have to choose deliberately. In this case the test rules are (iv) and (ii) (i.e. the same as before) and will be referred to as questions 3 and 4 respectively.

In the final pair of tasks positive confirmation is tested at the extremes i.e. with the outermost and innermost rules (rules (i) and (v) respectively). A test of the innermost rule using positive confirmation gives no information at all. The confirmation simply implies that any of the rules in the experiment could be the unknown rule. By contrast testing immediately outside the innermost rule would be a decisive rational test- it results in certain knowledge of whether the test rule is the unknown rule or not. A test of the outermost rule allows for the comparison necessary for the null hypothesis. Apart from the choice of test rule, the rules, combinations and bands are the same as in the "base" question pair. The test of the innermost rule will be referred to as question 5, while the test of the outermost rule is question 6.

The structure of the six questions can be seen in table 2:

Table 2: Questions in experiment

	Test Rule	Matching?	Symmetric question
Question 1	(iv)	Yes	Question 2
Question 2	(ii)	Yes	Question 1
Question 3	(iv)	No	Question 4
Question 4	(ii)	No	Question 3
Question 5	(v)	Yes	Question 6
Question 6	(i)	Yes	Question 5

The second column in the table gives the number of the test rule using the notation in table 1. The third column specifies whether the combination chosen could be explained by the matching hypothesis, while the fourth column specifies the question which is symmetric to that in the first column.

Once the subjects had completed all six tasks then the experiment was finished. The experimenter went round each subject in turn and the subject threw a die to determine

which of the six questions was played out. Once a question was selected then the envelope for that person and question was opened and the rule inside the envelope was compared to the test rule. Each person was given a base of two points to start. If the test and unknown rule were the same and the subject correctly said that they were the same then she gained four points. If the two rules were not the same then she gained one point. Otherwise she gained zero points. The total number of points was added together and the subject rolled the die again. If the number on the die was less than or equal to the total number of points then the subject won a cash prize of £12.

7. Results

Table 3 presents a summary of the choices made by the subjects in the experiment.

Table 3: Combinations chosen by subjects for each question

Comb no.	Corr Rule	Quest 1	Quest 2	Quest 3	Quest 4	Quest 5	Quest 6
1		1	5	4	2	8	12
2	i	13	17	5	14	26	26* [#]
3	ii	17	44* [#]	1	24* [#]	7	11
4	iii	20 [#]	22	15 [#]	28	5	20
5	iv	45*	14	33*	29	18 [#]	28
6	v	15	9	53	14	47*	14

* Immediately Positive Confirming combinations. [#] Bayesian optimal Combinations

In the first column, "Comb no." (=combination number) refers to the number of the combination as given in table 1, while the roman numerals in the "corr rule" (=corresponding rule) refer to the rule number. In each question in the questionnaire, each combination came from a separate band from the others, so effectively "Comb no." labels the bands from the outermost to the innermost. The numbering in "Corr Rule" is arranged so that each rule is aligned in the same row in the table as a combination which is in the band immediately inside that rule. The labels along the top of the table refer to the questions answered, with quest 1, quest 2 etc. referring to question 1, question 2 etc..

The table itself represents the number of subjects who chose each combination within each question. A casual look at the evidence shows that there is substantial evidence for the existence of positive confirmation. In questions 1, 3, and 5 there are substantially more people choosing the immediately positive confirming combination than there are choosing the Bayesian optimal combination. This is particularly clear-cut in question 1 and also in question 5 where 47 of the 111 subjects chose the positive confirmation option- even though this option gives no additional information at all. Noticeably the highest number of people in question 3 actually chose combination 6 rather than the immediate positive confirming combination 5.

Looking at questions 2,4 and 6 we see that the Bayesian optimal combination was chosen by substantial numbers of subjects in all cases, but only in question 2 were the numbers choosing that combination the largest. Instead, there seems to be a spread of combinations chosen, although all are inside the test rule and so can be seen as "Positively Confirming". It is noticeable that question 4 seems to have a different spread of combinations chosen from question 2.

The results given in table 3 suggest that a useful approach to the analysis of this data would be to look at it in general and specific terms. In the latter case one would compare the bands immediately inside and outside test rules. In the former case one would compare all the bands inside and outside the test rule (if they have the same number of rules). This allows both a test of the Bayesian- optimal band as well as of positive confirmation which is not confined to one particular band inside the test rule.

Table 4: Tests of Positive Confirmation- symmetric bands inside/outside test rules

Inside	Outside	Immediate inside/ outside*	All inside/ outside*
Quest 1	Quest 2	14.52 [#]	23.29 [#]
Quest 3	Quest 4	9.756 [#]	23.29 [#]
Quest 5	Quest 6	43.09 [#]	(Same)
Quest 2	Quest 1	10.67 [#]	55.68 [#]
Quest 4	Quest 3	0.29	55.68 [#]
Quest 6	Quest 5	0.2667	24.019 [#]

*Use of McNemar test using χ^2 distribution at 5% level of significance

Significant at 5% level of significance

Table 4 looks at the symmetric bands inside and outside the test rules for each question. In the table symmetric questions are compared with each other and the McNemar test is used to test for differences between the symmetric bands in these questions. The first column in the table specifies the inside band of one question and this is paired with the symmetric outside band of the other question in the second column. The third column give the results for the tests of the choices in the bands immediately inside or outside the test rules while the fourth column gives the results for the tests the choices in all the bands inside or outside the test rules. So, for example, in the first row of the table the figure in the third column tests the choices in the immediate inside band of question 1 with the choices in the immediate outside band of question 2. By contrast, the fourth column tests all the choices in the inside bands of question 1 against all the choices in the outside bands of question 2. Note that the result for question 5 (inside) and question 6 (outside) is the same in both cases.

Using the null hypothesis of equality derived from symmetric neutrality, for the first three rows of the table, there is a significant difference between the numbers of people choosing combinations from the positive confirmation band(s) in questions 1,3 and 5 and from the symmetric bands in questions 2,4 and 6. This holds both for immediately positively confirming bands and for all positively confirming bands. This suggests that there is a general bias towards positive confirmation in all situations, even when this is irrational as in this case.

For the next three rows the story is mixed. Here the choosing of immediately positively confirming combinations is rational and one would expect the differences between these combinations and combinations from symmetric bands to be greater. However, this is not the case. For the immediate inside band of question 2 and immediate outside band of question 1 the difference is significant but this is not the case for the immediate inside bands of questions 4 and 6 and their symmetric bands. By contrast, using all inside and outside bands, the differences are highly significant in all cases. This difference in results can be explained by the fact that, for questions 4 and 6 there is a wide spread of combinations chosen from all bands inside the test rules.

Table 5: Tests of similarity of questions

Comparison between qs	Immediate bands inside*	Distribution of Combinations**
Quest 1 vs Quest 3	3.6	4.115 [#]
Quest 2 vs Quest 4	8.33 [#]	3.283 [#]

* Use of McNemar Test using χ^2 distribution at 5% level of significance

** Use of standardised marginal homogeneity statistic

[#] Significant at 5% level

Table 5 shows the tests for differences between the two base tasks, questions 1 and 2 on one hand and questions 3 and 4 on the other). This uses two tests. The first test simply compares the differences in numbers of subjects choosing immediately inside the test rules of the questions using the McNemar statistic. The second test compares the distributions of choices of combinations across all bands using the standardised marginal homogeneity statistic. In both cases, as stated above the null hypothesis (for the first two rows) is that of equality. All of the statistics, apart from the first row for immediate inside bands, are significant, while the latter is only marginally insignificant. This suggests that there are major differences between behaviour between questions 1 and 3 on one hand and between questions 2 and 4 on the other.

No statistical comparison can be done between questions 5 and 6 and the other questions since they are not isomorphic and so no meaningful null hypothesis can be formed. However, a glance at table 3 suggests that there does seem to be a considerable difference in distributions, particularly between the choices of combinations in questions 2 and 6. In the former the choices are "bunched" in the band immediately inside the test rule, while in the latter they are more spread out.

Another way of looking at the data can be seen in graph 1 where the subjects in questions 1, 3 and 5 are examined. In this case the graph shows positive confirmation defined over all bands inside the test rule and split up into "types" depending on which questions they answered according to positive confirmation (i.e. irrationally). It is easy to see that question 3 is the question with the largest amount of positive confirmation overall while question 5 has the smallest overall amount. However the numbers who

positively confirm in all three cases is quite large. It is also noticeable how few subjects don't have some form of positive confirmation bias.

Having tested for positive confirmation in the search for information, it is necessary to look for positive confirmation in the use of information. Table 6 shows the results for the reactions of subjects to the information selected.

Table 6: Reactions of Positive Confirming subjects to evidence

Question/ Allowed	Subject: H ≠ U	Subject: H = U	χ^2 Statistic (1 df)	Contingency Coefficient
Q1- Not allowed	6	1	8.209 [#]	0.347 [#]
Q1- Allowed	16	37		
Q3- Not allowed	8	1	10.243 [#]	0.432 [#]
Q3- Allowed	26	51		
Q5- Allowed*	13	34	N/A	N/A
Q2- Not Allowed	34	5	20.439 [#]	0.326 [#]
Q2- Allowed	20	30		
Q4- Not Allowed	21	7	6.580 [#]	0.255 [#]
Q4- Allowed	31	36		
Q6- Not Allowed	36	6	8.861 [#]	0.287 [#]
Q6- Allowed	33	24		

* All combinations inside the test rule for question 5 are also inside all the other rules and so cannot distinguish between them.

[#] Significant at the 5% level of significance

In this table, for each question, the analysis is confined to those subjects who picked combinations from positive confirming bands. The first column in the table gives the question and whether the combination picked is allowed by the unknown rule or not. This divides the table into groups of two rows for each question (with the exception of question 5). The second column gives those subjects who declared that the test rule was

not equal to the unknown rule (i.e. $H \neq U$, while the third column gives those subjects who declared that the test rule was equal to the unknown rule (i.e. $H = U$). The fourth column gives the χ^2 test of independence for the two rows for each question. The fifth column gives the χ^2 based Contingency Coefficient measure of association between whether the rule is allowed and whether $H=U$ or not. Question 5 by its very nature cannot have a Positive Confirmation combination which does not allow the test rule so there is only one row and no test statistic.

The table shows significant statistics in all cases, demonstrating that the subjects' judgement is definitely affected by their choice of evidence. This is particularly interesting in the case of questions 1 and 3 where, if the combination is allowed then it increases the chances of the test rule being declared the same as the unknown rule by far more than is justified by the evidence. It should be noted that this is far more than would be justified under a Bayesian updating framework. Even in question 5, where the evidence from positive confirmation is irrelevant, the fact that the chosen combination is allowed by the unknown rule seems to increase the likelihood of the subject stating that the test rule is the same as the unknown rule.

It would be expected that the same thing would be the case with questions 2, 4 and 6, particularly since they include Bayesian choices. Indeed, they are all significant, although the number of those who declare that the test rule and the unknown rule are not equal is comparatively high in all cases. This can be seen in the Contingency Coefficients. All the coefficients are significant, suggesting association between the two variables, but there is more association for questions 1 and 3 than for questions 2, 4 and 6. This result is consistent with Bayesian interpretations⁵, although some of this could be caused by the wide variation in choices of combinations. In Question 6, for example, positive confirmation covers a wide range of bands, not just that immediately inside the test rule, so declaring that the test rule is not the same as the unknown rule is not irrational.

One other point to notice about table 6 is the relatively low number of subjects who, when they found out that the combination they had picked was not allowed by the unknown rule, went on to say that the test rule was the same as the unknown rule. This

demonstrates a high level of consistent thinking in rejecting rules which were proved to be false.

8. Discussion and Conclusions

The results outlined above demonstrate that there is a substantial amount of evidence for the existence of positive confirmation in the choosing of information. In general, when searching for information, it can be said that there is significant evidence for positive confirmation whether it is rational or irrational. In particular, when positive confirmation is not the optimal strategy, then the results are particularly clear-cut.

While positive confirmation seems to dominate in all cases, there do seem to be differences between the questions being asked. In question 3, for example, there is more variation in the combinations chosen than in question 1. Likewise, the variation in combinations chosen in question 6 prevents the test for positive confirmation in the band immediately inside the test rule from being significant. This suggests, in question 1, that matching does have some effect on choice, although the effect seems to be the reverse of that expected. Matching effectively *reinforces* the choice of the band immediately inside the test rule. When matching is not possible then the subjects, while still positively confirming, vary as to which combination they choose. This, however, does not explain the wide variation in question 6 and without more evidence it is hard to see how this can be explained. However it is worth pointing out that the data over all questions does show a high degree of variance and the variation in question 6 may simply be a manifestation of this.

Having noted the effects of variation in the selecting of combinations, it should be pointed out that this variation does not affect the theory that the subjects are choosing evidence according to positive confirmation. Positive Confirmation simply states that subjects would in general select combinations inside the test rule without specifying particular bands⁶ and this has been shown to be the case. The differences between

⁵ Although- see table 4- comparatively few subjects actually choose in the Bayesian optimal band.

⁶ Unlike the Bayesian- optimal combination which *does* specify one particular band.

questions does not affect this hypothesis although it does raise some interesting questions about why there should be so much variation. An investigation of the links between matching, with its ability to reduce this variation, and positive confirmation would be a particularly interesting avenue for research, although unfortunately it is one for which there is insufficient data to investigate in this experiment.

Noticeable among the results is the high level of positive confirmation in question 5. This is an important result because a positive confirming choice in question 5 is a choice of irrelevant evidence. This evidence does not discriminate at all between the test rule and all the other rules but it was still chosen by more than a third of the subjects in the experiment.

As well as the main type of positive confirmation when one is searching for information, it has been possible to test the use which has been made of the information when it has been acquired. Positive confirmation does lead to a consistent response to the decision as to whether the unknown rule is the same as the test rule. In general, there is a tendency to say that the two are the same when the combination is allowed and are not the same when it is not allowed. This suggests that positive confirmation is a consistent, if not rational, means of acquiring information. As well as the general association between the subjects reactions and the information, there is also little evidence of inconsistent behaviour. In particular there is no evidence of subjects accepting a test rule as being the same as an unknown rule when the test rule has been explicitly rejected.

All of these conclusions from the data are consistent with evidence which has been found, in a different experiment, in Jones and Sugden (2001). In that paper, it was also found that the positive confirmation bias existed, even within an incentive-compatible design. In particular, it was found that a subject would choose irrelevant information if it confirmed the rule that the subject was testing even if it was costly. It was also found that information which is interpreted as confirming a rule increases subjects' confidence in the truth of that rule, even if that information had no value.

This experiment has replicated these findings and extended them to situations where selecting positive confirming evidence may no longer be irrelevant, but is still suboptimal. It has been found that, in general, subjects do pick positively confirming evidence, while confirming that large numbers choose it even when it is irrelevant. It has

also been found that people do make consistent use of the information when they find it. However, matching, which was found to have no role in the Jones and Sugden experiment does seem to have some role here in focussing the selection of combinations.

As in the Jones and Sugden experiment, there is no reason to think that the subject would immediately find out that they are making a mistake in using positive confirmation, even if there was feedback between questions. A person who chooses a combination in a positively confirming band would not find out if they were definitely wrong unless their judgement was proved wrong. However this would only happen definitively where there was a situation, such as in question 5, if it turned out that the test rule was not the same as the unknown rule after the subject claimed (as tended to happen in this experiment) that it was. For other inner test rules (as in questions 1 and 3) it is not certain that the subject would learn that easily that positive confirmation was the wrong method. Making wrong judgements in itself would not be persuasive, since in these questions it is perfectly possible (from the subject's point of view) to get the judgement wrong and for positive confirmation to be right. The fact that only a small number of rules is eliminated by acquiring evidence means that there is still an element of risk in getting the right one. In fact, continued "success" for positive confirmation in judgements may even reinforce it as a method of acquiring information⁷

The experiment presented here illustrates that positive confirmation has significant effects on how people search for and use information and that this does not necessarily result in the optimal use of information. Future research in this field would suggest an investigation of whether this phenomenon still exists when subjects are allowed to choose their own rules to test and whether this can be applied to interactive choice settings.

⁷ This would especially be the case if the subjects' method of learning about learning methods was itself positively confirming.

Bibliography

- Becker G.M. , DeGroot M.H. & Marschak J. (1963) "Stochastic Models of behavior"
Behavioral Science vol. 8: 41 - 55
- Börgers T. (1996) "On the relevance of learning and evolution to economic theory"
Economic Journal vol. 106 p. 1374-1385
- Cubitt R.P. & Sugden R. (1998) "The selection of preferences through imitation" Review
of Economic Studies vol. 65 p. 761 -771
- Cheng P.W. & Holyoak K.J. (1989) "On the natural selection of reasoning theories"
Cognition vol. 31 p. 187- 276
- Erev I. & Roth A. (1998) "Predicting How People Play Games: Reinforcement Learning
in Experimental Games with Unique, Mixed Strategy Equilibria" American Economic
Review vol. 88 p. 848- 881
- Evans J. St. B.T. (1972) "Interpretation and Matching Bias in a reasoning task" Quarterly
Journal of Experimental Psychology vol. 24 p. 193- 199
- Fudenberg D. & Levine D.K. (1998) "The Theory of Learning in Games" Cambridge
Mass. London MIT Press
- Grether D.M. (1992) "Testing Bayes' Rule and the Representativeness heuristic: Some
Experimental Evidence" Journal of Economic Behavior and Organisation vol. 17 p. 31-
57
- Grether D.M. (1980) "Bayes' Rule as a descriptive model: The Representativeness
Heuristic" Quarterly Journal of Economics vol. 95 p. 537- 557
- Gorman M.E. & Gorman M.E. (1984) "A Comparison of Disconfirmatory, Confirmatory
and Control Strategies on Wason's 2-4-6 task" Quarterly Journal of Experimental
Psychology Vol. 36A
- Jones M.K. & Sugden R. (2001) "Positive Confirmation bias in the Acquisition of
Information" Theory and Decision vol. 50 p. 59- 99
- Klayman J. & Ha Y-W (1987) " Confirmation Disconfirmation and Information in
Hypothesis testing" Psychological Review vol.94 p. 211-228
- Loomes G. & Sugden R. (1995) "Incorporating a Stochastic element into decision
theories" European Economic Review vol. 39 p. 641-648

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

Mynatt C.R., Doherty M.E. & Tweney R.D. (1978) "Consequences of Confirmation and Disconfirmation in a Simulated Research Environment" Quarterly Journal of Experimental Psychology vol. 30

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

Roth A. & Erev I. (1995) "Learning in Extensive Form games : experimental data and simple dynamic models in the intermediate term" Games and Economic Behavior vol. 8 p.164- 212

Selten R. Sadrieh A. & Abbink K. (1999) "Money does not induce risk neutral behavior but binary lotteries do even worse" Theory and Decision vol. 46 p. 211- 249

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

Appendix 1: Optimal payoff model

Suppose there exists a set \mathbf{V} of N “nested” hypotheses where $H \in \mathbf{V}$ is an arbitrary hypothesis chosen for testing. Suppose j indexes a band $B_j \in \mathbf{B}$ (where \mathbf{B} is the set of all bands and $j \in \mathbf{J}$ where $\mathbf{J} = \{1, 2, \dots, N\}$). Rules are indexed in the same way as the bands they define with j increasing from the innermost band outwards.

Agents are supposed to be testing for the existence of the unknown hypothesis $U \in \mathbf{V}$, where U is selected at random by the experimenter. A *message* $\phi_{k,j}$ ($k \in \mathbf{K}$ where $\mathbf{K} = \{\text{allowed by } U, \text{ not allowed by } U\}$) is the result of picking a combination θ_j and either being told it is allowed or is not allowed by U . The agent is judging whether $U=H$. Denote this occurrence by H_T .

r is the utility acquired as a result of correctly claiming that H_T holds while (to balance Expected Utilities) $r/(N-1)$ is the utility acquired as a result of correctly claiming that H_T does not hold. In order to rationally select information, the probability of H_T is updated using Bayes' Rule.

Given the prior probabilities:

$$P(H_T) = 1/N;$$

$$\text{If } k = \text{"}\theta_k \text{ is allowed by } U\text{"} : P(\phi_{k,j}) = (N-j+1)/N$$

$$\text{If } k = \text{"}\theta_k \text{ is not allowed by } U\text{"} : P(\phi_{k,j}) = (j-1) / N$$

This gives updated probabilities $P(H_T/\phi_{k,j})$. There are eight possible situations:

- i) Testing on Outside; H_T holds; Combination allowed by U .
 $P(H_T/\phi_{k,j}) = 0$
- ii) Testing on Outside; H_T holds; Combination not allowed by U .
 $P(H_T/\phi_{k,j}) = 1/(j-1)$
- iii) Testing on Outside; H_T does not hold; Combination allowed by U .
 $P(H_T/\phi_{k,j}) = 1$
- iv) Testing on Outside; H_T does not hold; Combination not allowed by U .
 $P(H_T/\phi_{k,j}) = (j-2) / (j-1)$
- v) Testing on Inside; H_T holds; Combination allowed by U .
 $P(H_T/\phi_{k,j}) = 1 / (N - j + 1)$
- vi) Testing on Inside; H_T holds; Combination not allowed by U .
 $P(H_T/\phi_{k,j}) = 0$

vii) Testing on Inside; H_T does not hold; Combination allowed by U.

$$P(H_T/\phi_{k,j}) = (N - j) / (N - j + 1)$$

viii) Testing on Inside; H_T does not hold; Combination not allowed by U.

$$P(H_T/\phi_{k,j}) = 1$$

The agent must then decide whether it would be worth more to claim that H_T holds or not. The eight situations are grouped as follows:

i) Testing on Outside, Combination allowed by U.

ii) Testing on Outside, Combination not allowed by U

iii) Testing on Inside, Combination allowed by U

iv) Testing on Inside, Combination not allowed by U

In order to do this, for each of the four groupings the agent must find $\psi(\phi_{k,j})$ as follows for all k:

$$\psi(\phi_{k,j}) = \max (P(H_T / \phi_{k,j}) \times r , 1 - P(H_T / \phi_{k,j}) \times r / ((N-1))) \dots(1)$$

For the next stage the expected utilities for testing on the inside and the outside of H must be calculated. This gives:

$$\alpha(q) = \sum_{k \in K} P(\phi_{k,j}) \psi(\phi_{k,j}) \dots(2)$$

where $\alpha(q)$ is the expected payoff of testing in q bands where $q \in \{\text{Inside, Outside}\}$

Then, going through the algebra the results are:

$$\alpha(\text{Inside Bands}) = r/N [1 + (j-1) / (N-1)] \dots(3)$$

$$\alpha(\text{Outside Bands}) = r/N [1 + (N+1-j) / (N-1)] \dots(4)$$

Appendix 2: Question 1 in experimental questionnaire

Questionnaire- Red/Blue

Envelope 1

Below, on the left, are a set of possible rules one of which is identical to the rule in envelope 1. On the right is a list of combinations of red and blue balls.

Rule (i) : There must be at least 2 Red balls	1) 10 Blue; 0 Red
Rule (ii) : There must be at least 4 Red balls	2) 8 Blue; 2 Red
Rule (iii): There must be at least 6 Red balls	3) 6 Blue; 4 Red
Rule (iv): There must be at least 8 Red balls	4) 4 Blue; 6 Red
Rule (v) : They must all be Red balls	5) 2 Blue; 8 Red
	6) 0 Blue; 10 Red

The rule to be tested is Rule (iv) :

There must be at least 8 Red balls

In order to test this rule you will need to use the combinations as evidence. Ring **one** combination to see if it is allowed by the rule in envelope 1.

Once you have ringed the combination, write down the combination you ringed in the space below:

Please stop here until the experimenter comes around to give you information about your choice.

The combination you ringed [was/ was not] allowed by the rule in envelope 1.

Is Rule (iv) the same rule as that in envelope 1?

YES

NO

Circle your answer

Stop here and only go on to the next question when you are told to do so.

Graph 1

Subjects who Positively Confirm-

Types of subject

