The Suppression of Card Selections in Wason's Selection Task: Evidence that Inference Plays a Role

Aidan Feeney & Simon J. Handley
Centre for Thinking and Language
Department of Psychology
University of Plymouth
Drake's Circus
Plymouth PL4 8AA
Devon
United Kingdom
{afeeney, shandley}@plymouth.ac.uk

Abstract

We report the results of two experiments designed to investigate the role of inference in Wason's selection task. In Experiment 1 participants received either a standard one rule problem or a task which contained an additional rule. This additional rule specified an alternative antecedent. Both groups of participants were asked to select those cards they considered necessary to test whether the rule common to both problems was true or false. The results showed a significant suppression of "q" card selections. In addition there was weak evidence for increased "not-q" selection. In Experiment 2 we manipulated number of rules, as before, and the presence or absence of explicit negation on the cards. Once again "q" card selections were suppressed, but there was no evidence of an increase in "not-q" selection. There was also no effect of type of negation. Our results suggest that inferences about the unseen side of the cards underlie participants' selections. We argue firstly that these findings are inconsistent with current views of selection task performance (Oaksford and Chater, 1994, Evans and Over, 1996) and secondly, that they support accounts which emphasise the role of inference in the task.

Introduction

There is an anomaly in the reasoning literature: in relation to their performance on simple conditional inference tasks people perform poorly on indicative versions of Wason's selection task. This anomaly has led several workers in the field (e.g. Oaksford and Chater, 1994; Evans and Over, 1996; Kirby, 1994) to claim that people's performance on the selection task is best explained in terms of decision making rather than inference. Our goal in this paper is to demonstrate that people do engage in inference when confronted with indicative versions of the selection task.

In its original form Wason's task consisted of four cards, two of which had a letter on the upwards facing side (e.g. A and L), and two of which had numbers (e.g. 3 and 7). Accompanying these cards is an indicative conditional rule such as:

> If there is an A on one side of the card then there is a 3 on the other side

which participants are told may or may not govern what is on each side of the cards. Participants are asked which cards they would need to turn over in order to test whether the rule is true or false. The criterion for solving the task is selection of the A and the 7 cards.

Correct selection rates for indicative selection tasks belie the apparent simplicity of the problem and can be as low as 10% (Wason and Johnson-Laird, 1972). Much more common are selections of the A card on its own, the A and 3 cards, or the A, 3 and not-3 cards. The preponderance of 3 card selections is intriguing. In order to recognise that the 3 card has no bearing on whether the rule is true or false one only has to consider what might be on its unseen side. On the back of the 3 card there may be an A or a not-A. An A on the reverse side of the 3 card offers confirmation of the rule. A not-A, on the other hand, turns out to be irrelevant as the rule concerns those cards with an A only.

People's poor performance on the task in the light of its apparent simplicity and the mismatch between selection task and simple conditional inference task performance has led to a revaluation of the role of inference on the task. We will begin by outlining several alternative accounts of the processes underlying performance on the task. We will characterise these approaches in terms of the role they ascribe to inference. We will then go on to present the results of two experiments which suggest that consideration of what is on the unseen of the cards does affect people's choices on the task.

Non-Inferential Accounts of the Task

In recent years much emphasis has been placed on attempts to integrate the literatures on reasoning and decision making (see Johnson-Laird and Shafir, 1993; Evans and Over, 1996). This has led to accounts of indicative selection task performance which characterise the participants' task as one of decision making. That is, participants have to decide which cards to select when faced with the task. We will briefly present two accounts of the selection task as a task of decision making. Whilst we would emphasise that in many important respects these account are very different, neither suggests that people's card selections on Wason's task are determined by a consideration of what might be on the unseen side of the cards.

The first account is that offered by the heuristic/analytic (Evans, 1989) or dual process (Evans and Over, 1996)

model. This model explains behaviour on indicative selection tasks by positing the existence of an initial stage of processing where the aspects of the task to be attended to (those aspects which contain a high degree of relevance) are unconsciously determined. Evans suggests that attention is mediated by the operation of heuristics. On standard abstract selection tasks which employ affirmative rules the relevant heuristic is the IF heuristic. Evans (1989) argues that the function of if in natural language is to direct attention to the possible state of affairs in which the antecedent is true (causing people to think about the p and q cards in the selection task which correspond to the A and 3 cards in the example above).

Evans (1996) has claimed that participants in selection task experiments select cards based purely on the results of unconscious processing. He further suggests that any explicit inference which participants engage in whilst attempting the selection task takes the form of rationalisations for choosing cards which are attended to. Recently, however, Evans and Over (1996) have suggested that people's selections may sometimes be determined by the explicit inferential system. This suggestion was prompted by the success of several researchers in facilitating performance on the selection task (e.g. Green and Larking, 1995). The significance of the role played by the explicit inferential system remains unclear however.

The second non-inferential account of the selection task is that of Oaksford and Chater (1994) which is, in certain respects, very similar to that of Evans. They share with Evans the view that the selection task is best viewed as a decision making task where participants must make a decision about which cards to turn over. Accordingly, Oaksford and Chater reject logic as a normative standard for subjects' behaviour on the task and instead propose an alternative Bayesian account based upon the principles of optimal data selection.

The basic principle behind optimal data selection is very simple and is illustrated by Oaksford and Chater using the following example. Suppose you are asked to test the rule "If you eat tripe (p) then you feel sick (q)". There are four classes of individuals one could examine to test the rule: individuals who have eaten tripe (*p*); individuals who have not eaten tripe (*not-p*); individuals known to feel sick (*q*); individuals known not to feel sick (*not-q*). The question, of course, is one of deciding which classes of individuals will supply the most useful information. Intuitively, *p*, *q*, and *not-q* individuals are likely to be informative whereas not-p individuals will supply no information whatsoever. Oaksford and Chater have formalised these intuitions using Shannon-Wiener information theory and elements from the probability calculus. Table 1 contains their general predictions, based on their formalisations, about the relationship between the informativeness of any card and the probability of items mentioned in the rule.

Oaksford and Chater make an assumption of rarity - that is the probability of p (eating tripe) and the probability of q (feeling sick) are assumed to be low. Given this assumption they predict the following order of card selection frequencies on the affirmative abstract selection task: p>q>not-q>not-p.

Table 1: The conditions under which Oaksford and Chater predict cards will be informative

Cards	Conditions for Informativeness
p	P(q) is low
q	P(p) & P(q) are low
not-q	P(p) is high
not-p	None

Clearly, Oaksford and Chater's account of the task is non-inferential. Their claim is that people select cards based on the probability of the items mentioned in the rule. Indeed there are several demonstrations in the literature of the importance of the probability of the antecedent and consequent in determining the selections which participants make on the task (e.g. Kirby, 1994)

Inferential Accounts of the Task

In contrast to the non-inferential accounts presented above, the literature also contains what we will call inferential accounts of the indicative selection task. As we did for non-inferential accounts of the indicative selection task, we will also present two inferential accounts: the first offered by Sperber, Cara and Girotto (1995) and the second by Johnson-Laird and Byrne (1991; Johnson-Laird, 1995). Again we hasten to point out that the accounts which we will present are, in many respects, very different. In this context, however, it is the emphasis which both place on inference which is of interest.

Sperber et al's account purports to account for indicative selection task performance in a very simple manner, by assuming firstly, that people understand the task as one of selecting relevant evidence for testing a rule and that this rule may only be tested indirectly (i.e. through its observationally testable consequences). Secondly, inferring some of the consequences of the rule is done spontaneously, as part of the process of comprehension. These spontaneous inferences lead to an interpretation of the rule which contains certain expectations of relevance (i.e. expectations, on the part of the subjects, about what is relevant to the experimenter). For example, on affirmative abstract selection tasks participants, given a minimal presumption of relevance, will infer that there are cards which have both P and Q features. Their selection of cards will reflect a test of this assumption. This will lead to the selection of the p and q cards.

As may be seen from this account, the crucial factor is the spontaneous inferences which subjects make whilst comprehending the rule. This determines both the way in which any conditional rule is found to be relevant and the cards which participants select on an indicative selection task. Interestingly, Oaksford and Chater (1995) claim that their account is compatible with Sperber et al's relevance theoretic account. We would argue that these accounts differ in one crucial respect: Sperber et al claim that people, when attempting the indicative selection task, think about the consequences of turning over certain cards. There

is no such level of explanation in Oaksford and Chater's account.

The final approach to the indicative selection task which will be considered here is one based on mental models (Johnson-Laird and Byrne, 1991; Johnson-Laird, 1995). The model theory assumes that people select only those cards that are explicitly represented in their initial models of the rule. These initial models encode only those items that are mentioned in the rule together with a mental footnote which specifies the manner in which the initial models may be fleshed-out. Furthermore, mental model theory assumes that people select those cards for which the unseen side could have a bearing on the truth or falsity of the rule. Whether people select the p card or the p and q cards depends on their interpretation (and initial encoding) of the rule. A simple conditional interpretation will lead to selection of the p card only as the hidden side of the q card has no bearing on the truth or falsity of the rule. A biconditional interpretation (that a conditional such as If p then q implies its obverse: If q then p) will lead to the selection of the p and q cards as the unseen side of both cards has a bearing on whether the rule is true or false. Although lack of space precludes a more thorough consideration of the model theory, suffice it to say that the model theory, like the relevance theoretic account, claims that people do engage in inference on the selection task. The cards selected, however, depend both on how people interpret and represent the rule as well as the inference they engage in.

We will now consider how it might be possible to discriminate between these two types of account. We will argue that it may be possible to suppress the rate at which certain cards are chosen on the indicative selection task. We will further argue that, if such suppression does occur, it may be explained only in terms of card choice being determined by consideration of the unseen side of the cards.

How to Suppress Card Choices on the Task

It has been known for some time that the rate at which we draw both valid and invalid inferences on a conditional reasoning task may be suppressed. The phenomena of inference suppression was first demonstrated by Rumain, Connell and Braine (1983). Their experiment used a simple conditional inference task where participants were given a series of major and minor premises and were then asked what followed. For example, participants might be told:

If there is a tiger in the box then there is an orange in the box
There is a tiger in the box

It was found that the inclusion of a second major premise such as:

If there is a dog in the box then there is an orange in the box

before the presentation of the minor premise greatly reduced people's tendency to make the invalid conditional inferences. Rumain et al claimed that their manipulation blocked a biconditional interpretation of the first premise ensuring that, for example, when participants were given the premises:

If there is a tiger in the box then there is an orange in the box
If there is a dog in the box then there is an orange in the box
There is an orange in the box

they did not make the affirmation of the consequent (AC) inference that there was a tiger in the box. Various suppression effects have been demonstrated since the publication of Rumain et al's paper (e.g. Byrne, 1989). Our purpose in this paper, however, is not to examine these effects in detail but to the consider the effects of Rumain et al's original manipulation for people's response patterns on the indicative selection task and to discuss these effects in the light of the debate about the role played by inference in contemporary accounts of the task.

We start by assuming that, in the case of the AC inference, the presence of the additional premise in Rumain et al's study caused their participants to realise that there may have been a tiger or a dog in the box Accordingly, to make any inference about the presence of a tiger, given the presence of an orange, would be unsafe.

Now consider an indicative version of the selection task where your are presented with a simple conditional rule such as:

If P then Q

together with four cards labelled P, Q, not-P and not-Q and an additional rule specifying an alternative antecedent such as:

If R then S

How is the presence of the additional rule likely to affect people's card selections? If, as we claim, people's response patterns are determined by a consideration of the possible consequences of turning over, at least some of, the cards, we would expect the presence of the alternative rule to greatly suppress q card selection. This is because on the back of the q card there may be a p or an r. Thus it will have no bearing in determining the truth or falsity of the first rule. Importantly, this prediction rests on the claim that people consider or represent the consequences of turning over the cards and hence, may be characterised as thinking inferentially on Wason's task. We will now present two experiments which test this prediction.

Experiment 1

Method

Participants: 98 first-year students in the psychology of the authors' home university served as participants in this experiment.

Materials: each participant attempted one indicative selection task. The selection tasks used were either abstract or thematic and contained one rule or two. The abstract materials were as follows:

Imagine that your friend has just shown you a strange deck of cards. Each of the cards in the deck has a number on one side and

a letter on the other side. She tells you that the cards in the deck obey the following rule:
If a card has the letter A on one side then it has the number 3 on the other side
You decide to investigate whether this rule is true or false. You get your friend to draw four cards from the deck. These cards are shown below. Please indicate by circling the appropriate card or cards, which ones you need to turn over to decide whether the rule is true or false.
Remember the rule is:
If a card has the letter A on one side then it has the number 3 on the other side

Below this were printed four cards containing a 7, an A, a 3 and an L. These cards were presented in one of two random orders.
The materials for the two rule conditions were identical, except that after presentation of the first rule, participants read the following:

She also tells you that they obey the following additional rule:
If a card has the letter L on one side then it has the number 3 on the other side

The thematic materials were analogous to the abstract materials and concerned an imaginary conversation between the participant and their doctor during which the participant is told that she has been diagnosed with a serious disorder of the immune system. The rule to be tested in this case was:

If a person with this disorder gets pneumonia then they will die

The four cards (again presented in one of two random orders) represented a patient with cancer, a patient with pneumonia, a patient who had died and a patient who had lived. The additional rule in the thematic/two rule condition was:

Of course, if a person gets cancer then they will also die

Procedure: the experiment was carried out at the end of a lecture. Participants received neither payment nor course credit for taking part.

Design: the experiment had a 2 (no. of rules) x2 (materials) between subjects design. These combine to give four conditions: abstract with one or two rules (A1 vs. A2) and thematic with one or two rules (T1 vs. T2)

Results and Discussion

The first analysis performed was on the mean number of cards selected by participants. A 2x2 between subjects Anova was carried out on the total number of card selections. Neither the main effects for number of rules and content nor the interaction was significant. However, the main effect of number of rules approached significance ($F(1, 94) = 3.184$, $p < .08$). The mean number of cards selected by participants in the one rule conditions was 1.94 (S.D. = .74) and 1.69 (S.D. = .66) in the two rules conditions. This result suggests that our number of rules manipulation was successful in suppressing card selections.

An examination of Table 2 suggests that our materials variable had no significant effect on observed patterns of card selection. Accordingly, for the purposes of determining the individual cards responsible for the trend towards

Table 2: Individual card selection frequencies for Experiment 1 as a function of condition

	p	not-p	q	not-q
A1	23 (92%)	5 (20%)	15 (60%)	8 (32%)
T1	20 (80%)	2 (8%)	15 (60%)	9 (36%)
	43 (86%)	7 (14%)	30 (60%)	17 (34%)
A2	14 (61%)	5 (22%)	7 (30%)	12 (52%)
T2	17 (68%)	4 (16%)	8 (32%)	14 (56%)
	31 (65%)	9 (19%)	15 (31%)	26 (54%)

suppression in the total number of cards selected we collapsed across contents. As predicted a significant reduction due to number of rules in q-card selection frequency ($\chi2\ (1) = 8.15$, $p < .005$). There was also a smaller, non-predicted decrease in p card selection frequency ($\chi2\ (1) = 6.07$, $p < .02$) and an increase in not-q card selections ($\chi2\ (1) = 4.05$, $p < .05$).

The significant difference in participant's tendency to select the q card due to number of rules was both highly significant and in the direction predicted. Surprisingly the reduction in q card selections is accompanied by an increase in not q selections. This suggests that our manipulation may serve, not only to suppress true antecedent choices, but also to facilitate the selection of the potentially falsifying not q card. However this increase in selection was not associated with a greater incidence of logically correct selection combinations in the two rule condition. The rate of p and not q selection is no greater with two rules than with one rule (7/50 in the one rule conditions and 9/48 in the two rule conditions). That these results are not due to an increase in logically correct responding is further borne out by the significant and surprising decrease in p card choices.

Although the finding of q card suppression confirms our prediction, it could be argued that the effect of Q being common to both rules may have served to increase its perceived probability. If this is the case, Oaksford and Chater (1994) would predict the observed decrease in p card selections. This finding might also be predicted by the dual process account (Evans and Over, 1996) as the presence of an alternative antecedent may have served to decrease the relevance of the p card.

In summary, this experiment has strongly confirmed our prediction about q-card suppression. However, the finding of both increased not-q card selections and decreased p-card selections are consistent with alternative, non-inferential accounts of performance on the task. In Experiment 2 we test the reliability of the effects found in Experiment 1. We also include a further manipulation

designed to explicitly test a prediction of the dual process model.

Experiment 2

The first aim of this experiment is to test the reliability of the effects on individual card selection frequencies of our number of rules manipulation from Experiment 1. As the content manipulation in Experiment 1 did not lead to any differences in either individual or total card selections it was dropped from this experiment. Instead, all participants received abstract versions of the indicative task. In place of a content manipulation we varied the type of negation used on the cards presented. In Experiment 1 negation was expressed implicitly on the cards. Thus, the not-p and not q-cards were labelled *L* and *7* respectively. In this experiment half of the participants received cards where negation was explicitly expressed i.e. *a card that is not an A* and *a card that is not a 3*.

This manipulation serves to test a prediction of the dual process account. Evans, Clibbens and Rood (1996) have demonstrated that the use of explicit negation on the cards presented in abstract selection tasks leads to an increase in the selection of all cards and a decrease in the selection of those cards which match the items mentioned in the rule. Evans (in press) has claimed that the presence of explicit negation increases the relevance of all of the cards. Hence, the use of explicit negation in this experiment should result in an increase in the total number of cards selected.

Method

Participants: 80 first-year students from the authors' home university served as participants in this study

Materials: again each participant attempted one indicative selection task. All of the tasks used in this experiment were abstract. The materials received by participants in the implicit negation conditions were identical to those used in the abstract conditions of Experiment 1. The materials used in the explicit negation conditions differed only in labelling of the cards presented. The cards presented were labelled *a number which is not 3*, *the letter A*, *the number 3* and *a letter which is not A*. Again the cards were presented in one of two random orders.

Procedure: as for Experiment 1.

Design: the experiment had a 2 (no. of rules) x2 (negation type) between subjects design. These combine to give four conditions: implicit negation with one or two rules (I1 vs. I2) and explicit negation with one or two rules (E1 vs. E2)

Results and Discussion

The results of this experiment were first analysed in terms of the total number of cards selected by participants. A 2x2 between subjects Anova was carried out on the number of card selections. Contrary to the predictions of the dual process account, neither the main effect of Negation Type nor the interaction between Negation Type and Number of Rules was significant. However, the main effect of Number of Rules was highly significant ($F(1, 94) = 12.04$, $p < .001$). The mean number of cards selected was 2.15 (S.D. = .83) in the one rule conditions and 1.6 (S.D. = .55) in the two rules conditions. This finding suggests that our number of rules manipulation was successful in suppressing card selections.

An examination of individual card selection frequencies (displayed in Table 3) suggests that, as was the case in Experiment 1, our number of rules manipulation has had by

Table 3: Individual card selection frequencies for Experiment 2 as a function of condition

	p	not-p	q	not-q
I1	20 (100%)	3 (15%)	15 (75%)	6 (30%)
E1	16 (80%)	7 (35%)	12 (60%)	7 (35%)
	36 (90%)	10 (25%)	27 (68%)	13 (33%)
I2	18 (90%)	3 (15%)	4 (20%)	5 (25%)
E2	17 (85%)	2 (10%)	8 (40%)	7 (35%)
	35 (88%)	5 (13%)	12 (30%)	12 (30%)

far the biggest influence on the results of this experiment. Therefore, for the purposes of our statistical analysis of individual card selection frequencies we collapsed across Negation Type. A significant difference in selection rates for individual cards due to number of rules were found for the q-card ($\chi 2\ (1) = 11.26$, $p < .001$) only. The equivalent statistics for the p and not q cards were $\chi 2\ (1) = .125$ and $\chi 2\ (1) = .058$ respectively. Finally, card combination frequencies revealed that, once again, there was no evidence of our number of rules manipulation leading to an increase in logically correct responding.

The results of this experiment confirm the highly significant suppression of q card selections found in Experiment 1. In addition, the findings suggest that the weaker effects present in Experiment 1 (decreased p card selection and increased not q selection) are unreliable.

General Discussion

The experiments reported in this paper offer strong support for the claim that people engage in inference when attempting Wason's selection task. By providing participants with an additional rule which specified an alternative antecedent we have significantly decreased their tendency to select the q card when testing the experimental rule. We feel that the most likely explanation for this effect is that participants recognise that on the back of the q card there may be a p or a letter which is not a p. Hence, the selection of the q card will have no bearing on determining the truth or falsity of the rule. If this is the case our effect has been caused by participants explicitly considering the unseen side of the cards with which they were presented.

Our findings pose some difficulty for non-inferential accounts of the indicative selection task. For example, the dual process model of Evans and Over (1996; Evans, in press) makes several predictions that are inconsistent with our findings. The presence of an additional rule with a common consequent should serve to increase the relevance of the q card leading to an increased selection rate. Both of

our experiments demonstrate significant suppression of q card selections. In addition, the dual process model predicts increased overall selection rates with the use of explicit negation on the cards. Experiment 2 demonstrates an overall reduction in card selections due to the presence of an additional rule and no evidence for an effect of explicit negation. Evans and Over do claim, however, that people engage in inference on conditional inference tasks. We believe that we have demonstrated people's ability to engage in the same type of thinking on the selection task.

Oaksford and Chater's (1994) information theoretic account fares somewhat better in the light of our findings. It is likely that one effect of our two rule manipulation is to increase the perceived probability of Q. As an examination of Table 1 reveals, the informativeness of the q card is an inverse function of P(p) and P(q). Hence, Oaksford and Chater would predict that as the perceived probability of Q increases, the rate of q-card selection will decrease. This is entirely consistent with our findings. However, as Table 1 also shows, the informativeness of the p card is an inverse function of P(q) on its own. Therefore, Oaksford and Chater's model would also predict a reduced p card selections in our two rules conditions. Although our findings demonstrate a small reduction in Experiment 1, there is no such reduction in p card selections in Experiment 2.

Accounts of the indicative selection task which we have characterised as being inferential in nature can more readily cope with our data. For instance, Johnson-Laird (1995) has claimed that participants who choose the q card on indicative selection tasks do so because their initial representation of the rule does not discount a biconditional interpretation. The additional rule may cue participants to construct a model containing both antecedents as alternative conditions for the consequent (Byrne, 1989). This would ensure that the rule to be tested cannot be interpreted as its obverse and reasoners recognise that turning over the q card will have no implications for the falsity of the rule.

Similarly Sperber, Cara and Girotto (1995) can also offer an explanation of our findings. Recall that their account of the standard abstract task claims that the rule achieves sufficient linguistic relevance by cueing reasoners to assume the existence of cards with P and Q features. People's selections reflect a test of this existential assumption. It is likely, therefore, that our manipulation will cue reasoners to infer the existence of cards with the features mentioned in the experimental rule (P and Q features), and to infer the existence of cards with the features mentioned in the additional rule (R and Q features). Selecting the P card to test for the existence of cases with P and Q features has greater inferential consequences in this context than does selecting the Q card (which may confirm either assumption). Again, it is important to highlight that this account is dependent upon the assumption that reasoners consider the unseen sides of the cards.

To conclude, we feel that although the experimental work demonstrating people's sensitivity to the probabilities of the items they are thinking about is very important, we do not feel that it merits a wholly probabilistic approach to thinking. The experiments presented in this paper demonstrate that, at least under certain conditions, people do think inferentially on the indicative selection task.

References

Byrne, R.M.J. (1989). Suppressing valid inferences with conditionals. *Cognition*, 31, 61-83.

Evans, J.St.B.T. (1983). Linguistic determinants of bias in conditional reasoning. *Quarterly Journal of Experimental Psychology*, 35A, 635-644.

Evans, J.St.B.T. (1989). *Bias in human reasoning: Causes and consequences*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Evans, J.St.B.T (1996). Deciding before you think: Relevance and reasoning in the selection task. *British Journal of Psychology*, 87, 223-240.

Evans, J.St.B.T. (in press). Matching bias in conditional reasoning: Do we understand it after 25 years? *Thinking and Reasoning*.

Evans, J.St.B.T., Clibbens, J. & Rood, B. (1996). The role of implicit and explicit negation in conditional reasoning bias. *Journal of Memory and Language*, 35, 392-409.

Evans, J.St.B.T. & Over, D. (1996). *Rationality and Reasoning*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Green, D.W. & Larking, R. (1995). The locus of facilitation in the abstract selection task. *Thinking and Reasoning*, 1, 183-199.

Johnson-Laird, P.N. (1995). Inference and mental models. In S.E. Newstead & J.St.B.T. Evans (Eds.), *Perspectives on Thinking and Reasoning*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Johnson-Laird, P.N. and Byrne, R.M.J. (1991). *Deduction*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Johnson-Laird, P.N. & Shafir, E. (1993). The interaction between reasoning and decision making: An introduction. *Cognition*, 49, 1-9.

Kirby, K.N. (1994). Probabilities and utilities of fictional outcomes in Wason's four-card selection task. *Cognition*, 51, 1-28.

Oaksford, M. & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608-631.

Oaksford, M. & Chater, N. (1995). Information gain explains relevance which explains the selection task. *Cognition*, 57, 97-108.

Rumain, B., Connell, J. & Braine, M.D.S. (1983). Conversational comprehension processes are responsible for reasoning fallacies in children as well as adults: IF is not the biconditional. *Developmental Psychology*, 19, 471-481.

Sperber, D., Cara, F., and Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, 57, 31-95.

Wason, P.C. & Johnson-Laird, P.N. (1972). *Psychology of reasoning: Structure and content*. London: Batsford.