ELSEVIER

Cognition 57 (1995) 97–108

COGNITION

Discussion

Information gain explains relevance which explains the selection task

Mike Oaksford[a,*], Nick Chater[b,*]

[a]*Department of Psychology, University of Warwick, Coventry, CV4 7AL, UK*
[b]*Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, UK*

1. Introduction

Sperber, Cara and Girotto (1995) argue that relevance theory (Sperber & Wilson, 1986) explains the selection task. The main tenet of relevance theory is that relevant information has the greatest cognitive effects for the least processing effort. They construct experimental materials that they take to vary the cognitive effect and the processing effort required to solve the selection task. They argue that the results of their experiments conclusively support the relevance account, and discount other explanations of selection task performance. In particular, they suggest that their data and their approach are not compatible with Oaksford and Chater's (1994) rational analysis (Anderson, 1990) of the selection task that uses "information gain" to determine card selection. By contrast, in this paper, we argue that the information gain and relevance accounts are compatible, rather than in competition. Oaksford and Chater's notion of expected information gain provides a quantitative measure of relevance appropriate to the selection task. We demonstrate the validity of this interpretation by showing that the information gain account can explain Sperber et al.'s (1995) experimental results.

Why do Sperber et al. (1995) conclude that information gain and relevance approaches are incompatible? First, they contend that the information gain approach does not explain important aspects of the data in the empirical literature, which, they argue, the relevance account can handle. In particular, they argue that the information gain account does not address the facilitation of the "logical" *p*, *not-q* response when the

* Correspondence to either author. M. Oaksford, e-mail: pysad@csv.warwick.ac.uk. N. Chater, e-mail: nick@dragon.psych.ox.ac.uk

0010-0277/95/$09.50 © 1995 Elsevier Science B.V. All rights reserved
SSDI 0010-0277(95)00671-0

consequent of the task rule contains a negation (Evans & Lynch, 1973; Oaksford & Stenning, 1992). However, Oaksford and Chater (1994) provide a detailed quantitative analysis of these experiments, including data from Evans and Lynch (1973), Griggs and Cox (1983), Manktelow and Evans (1979), Oaksford and Stenning (1992), Pollard (1985) and Reich and Ruth (1982). Indeed, Oaksford and Chater show that their theoretically derived expected information gains correlate highly, and significantly, with the observed data.

Second, Sperber et al. (1995) argue that their own data (Experiment 2) are incompatible with Oaksford and Chater's (1994) account. Given the large range of experimental data for which Oaksford and Chater's theory provides a quantitative explanation, it is not clear how to interpret a single anomaly, even if it was completely inexplicable in terms of the theory. Moreover, it is not clear that the relevance account is compatible with the range of data covered by the information gain approach (which provides a comprehensive, quantitative analysis of the majority of the past literature). The Bayesian approach that we adopt in our model of the selection task is in explicit opposition to falsificationism – you can always explain away a single inconsistent result (Duhem, 1914–1954; Quine, 1953). What is important is the ability of a theory to account for the broad pattern of replicable results. In any case, we shall argue that there is a plausible interpretation of Experiment 2, which is compatible with the information gain account.

Third, along with almost all existing accounts of the selection task, they accuse the information gain account of falling "short of either predicting or ruling out good performance (more than 50% correct) on yet untested varieties of the task". Sperber et al. argue that their relevance account does provide predictions. We are at a loss to know what differentiates all these other views from the relevance account in relation to predictive power. In particular, we have made predictions from the information gain account, which we mention in Oaksford and Chater (1994), and which we are currently testing experimentally. Further, we note that Oaksford and Chater (1994) formulated the information gain theory and submitted it for publication before Kirby's (1994) results were available. Oaksford and Chater's (1994) subsequent analysis showed that the information gain theory predicted Kirby's results.

We have suggested that the information gain account may be a way of making a relevance account of the selection task formally precise. We now outline the information gain account, and then show how to apply it to model Sperber et al.'s experiments. Finally, we discuss relevance, information gain and other accounts of the selection task more generally.

2. The information gain approach to the selection task

In Wason's selective task (Wason, 1966, 1968), subjects are instructed to assess whether some evidence is relevant to the truth of falsity of a conditional rule, *if p then q*, where "p" stands for the antecedent clause of the conditional and "q" for the consequent clause. In the standard abstract task, the rule concerns cards, which have a number on one side and a letter on the other. The rule is *if there is a vowel on one side* (p), *then there is an even number on the other side* (q). Four cards are placed before the subject, so that just one side is visible. The visible faces show an "A" (p card), a "K" (*not-p* card), a "2" (q card) and a "7" (*not-q* card). Subjects then select those cards they must turn over to determine whether the rule is true of false. Typical results were: p and q cards (46%); p card only (33%), p, q and *not-q* cards (7%), p and *not-q* cards (4%) (Johnson-Laird & Wason, 1970a). Subjects confront a task that is analogous the scientist's problem of which experiment to perform. Scientists have a hypothesis (the conditional rule) to assess, and they aim to perform experiments (turn cards) likely to provide data (i.e., what is on the reverse of the card) bearing on its truth or falsity.

The "correct" response of turning the p and the *not-q* card derives from a tacit acceptance of Popper's (1959–1935) falsificationist philosophy science which recommends only conducting experiments that can potentially falsify a hypothesis. Subjects should therefore only turn cards that could be logically incompatible with the conditional rule *if p then q*. Subjects should therefore turn the p card, because it might have a *not-q* on the back, and the *not-q* card, because it might have a p on the back. However, subjects typically select the p and q cards. This mismatch has been viewed as throwing human rationality into doubt (see Cohen, 1981; Manktelow & Over, 1993; Stich, 1985, 1990).

Oaksford and Chater (1994) note that contemporary philosophers of science have rejected falsification as unfaithful to the history of science (Koyré, 1957; Kuhn, 1962, Toulmin, 1961) and as in any case unworkable (Duhem, 1914–1954; Putnam, 1974; Quine, 1953). More recent accounts of scientific inference take a Bayesian, probabilistic approach to confirmation (Earman, 1992; Horwich, 1982; Howson & Urbach, 1989). Oaksford and Chater adopt this point of view in providing a rational analysis (Anderson, 1990, 1991) of the selection task that uses the Bayesian theory of optimal data selection (Lindley, 1956; Good, 1966; MacKay, 1992).

Oaksford and Chater (1994) suggest that hypothesis testers should choose experiments (select cards) to provide the greatest possible "expected information gain" in deciding between two hypotheses (i) that the task rule, *if p then q*, is true, i.e., ps are invariably associated with qs; and (ii) that the occurrence of ps and qs is independent. For each hypothesis, Oaksford and Chater (1994) define a probability model that derives from the prior probability of each hypothesis (which for most purposes they assume to be equally likely, i.e., both $= .5$), and the probabilities of p and of q in the task rule. They define information gain as the difference between the uncertainty *before* receiving some data and the uncertainty *after* receiving that data where they measure uncertainty using Shannon–Wiener information. Thus Oaksford and Chater (1994) define the information gain of data D as:

$$\text{Information before receiving } D\colon I(H_i) = -\sum_{i=1}^{n} P(H_i)\log_2 P(H_i)$$

$$\text{Information after receiving } D\colon I(H_i|D) = -\sum_{i=1}^{n} P(H_i|D)\log_2 P(H_i|D)$$

$$\text{Information gain: } I_g = I(H_i) - I(H_i|D)$$

They calculate the $P(H_i|D)$ terms using Bayes' theorem. Thus information gain is the difference between the information contained in the *prior* probability of a hypothesis (H_i) and the information contained in the *posterior* probability of that hypothesis given some data D.

When choosing which experiment to conduct (i.e., which card to turn), the subject does not know what that data will be (i.e., what will be on the back of the card). So they cannot calculate actual information gain. However, subjects can compute *expected* information gain. Expected information gain is calculated with respect to all possible data outcomes, e.g., for the p card, q and *not-q*, and both hypotheses.

Assuming subjects can compute expected information gains, they then have to make a decision about which cards to select. Oaksford and Chater (1994) incorporated two aspects of the decision process in their relevance measure. First, they introduced a noise factor by adding .1 to the information gain for each card. This allows that people may occasionally see the *not-p* card as informational. Second, card selection is a competitive matter – a card should have a greater chance of being chosen the less distinguishable it is from alternatives. Oaksford and Chater (1994) therefore scaled their information gain measure by the mean value for all four cards. They refer to the resulting measure as "scaled expected information gain" ($SE(I_g)$).

Oaksford and Chater (1994) calculated $SE(I_g)$s for each card assuming that the properties described in p and q are rare. They motivate the "rarity assumption" from the observation that it seems to apply to the vast majority of everyday conditional sentences. They also cite support for this view from the literature on other reasoning tasks (Klayman & Ha, 1987; Anderson, 1990). Hence, Oaksford and Chater (1994) argue that people's strategies for dealing with conditional rules will tend, by default, to be adapted to the case where rarity holds.

Adopting the rarity assumption, the order in $SE(I_g)$ is:

$$SE(I_g(p)) > SE(I_g(q)) > SE(I_g(not\text{-}q)) > SE(I_g(not\text{-}p))$$

This corresponds to the observed frequency of card selections in Wason's task: $n(p) > n(q) > n(not\text{-}q) > n(not\text{-}p)$, *where* $n(x)$ denotes the number of cards of type x selected. This account thus explains the predominance of p and q card selections as a rational inductive strategy. This ordering holds only when $P(p)$ and $P(q)$ are both low. Oaksford and Chater note that task manipulations that suggest that this condition does not hold (at least one of $P(p)$ or $P(q)$ is high) leads to alternative orderings, predominantly that:

$$SE(I_g(p)) > SE(I_g(not\text{-}q)) > SE(I_g(q)) > SE(I_g(not\text{-}p))$$

This ordering is more consistent with Popperian falsficationism, where the *p* and and *not-q* instances are favoured. The effect of rarity and its violation will enable us to account for much of Sperber et al.'s results.

Oaksford and Chater (1994) also show how their model generalises to all the main patterns of results in the selection task. Specifically, it accounts for the non-independence of card selections (Pollard, 1985), the negations paradigm (e.g., Evans & Lynch, 1973), the therapy experiments (e.g., Wason, 1969), the reduced array selection task (Johnson – Laird & Wason, 1970b), work on so-called fictional outcomes (Kirby, 1994) and deontic versions of the selection task (e.g., Cheng & Holyoak, 1985) including perspective and rule-type manipulations (e.g., Cosmides, 1989; Gigerenzer & Hug, 1992), and the manipulation of probabilities and utilities in deontic tasks (Kirby, 1994).

3. Modelling Sperber et al.'s results

We now apply the information gain account to Sperber et al.'s four experimental studies in turn, and argue that these studies confirm this account. The basic strategy of these experiments is to show that in a "relevance" condition subjects consistently select the *p* and *not-q* cards, whereas these selections are much less frequently observed in an "irrelevance" condition, where the *p*, *q* card selection dominates. Our approach to modelling Experiments 1–3 will be to show that in the relevance condition the materials violate rarity, whereas they adhere to rarity in the irrelevance cases. We provide a more quantitative analysis of the richer data obtained in Sperber et al.'s Experiment 4.

Experiment 1

Sperber et al.'s Experiment 1 contrasts a relevance condition concerning what they call the "virgin mothers" problem, with an irrelevance condition consisting of a standard abstract selection task. The irrelevance condition uses standard materials, and hence we assume that the default rarity assumption applies, giving the normal ordering: $n(p) > n(q) > n(not\text{-}q) > n(not\text{-}p)$. This is exactly the ordering found in Sperber et al.'s data: $n(p) = 25 > n(q) = 11 > n(not\text{-}q) = 8 > n(not\text{-}p) = 1 (N = 27)$.

The "virgin mothers" problem employs the rule "if a woman has a child, she has had sex". In this rule, both the antecedent and the consequent violate the rarity assumption, because the majority of women have children, and the majority of women have had sex. Therefore, we would predict that *not-q* card selections will exceed *q* card selections leading to the overall pattern: $n(p) > n(not\text{-}q) > n(q) > n(not\text{-}p)$. As before, this is exactly the ordering found in Sperber et al.'s data: $n(p) = 26 > n(not\text{-}q) = 23 > n(q) = 2 > n(not\text{-}p) = 1 (N = 27)$.

Experiment 2

In Experiment 2, both relevance and irrelevance conditions involve contentful materials, concerning the visit to Padua of a group of English schoolchildren. Volunteers are required to look after these children, and there is speculation over the sex and marital status of people who put themselves forward as volunteers. The relevance condition uses the rule: "if a volunteer is male, then he is married". The irrelevance condition uses the rule: "if a volunteer is male, then he is dark haired". Unlike Experiments 1 and 3, it is much less clear how to assign the probabilities in this experiment, because it depends on subjects' assumptions about the people who are likely to put themselves forward in this type of situation. The uncertainty here is paralleled by the uncertainty in Sperber et al.'s account of the task. They assert that "if a volunteer is male, then he is married" is relevant on the grounds that its counterexample is lexicalized (i.e., bachelor); and that "the most salient cognitive effect of the conditional statement is on the presence of bachelors among the volunteers". Although these are perhaps reasonable speculations concerning how subjects represent the problem, these assertions do not follow from any well-specified theory of relevance. Therefore, if the information gain account can also provide a plausible interpretation, then it should be favoured as an account of the computation of relevance in this context.

We suggest that in the volunteering context, subjects assume that male volunteers will be rare (the instructions for the relevance condition explicitly reflect this). So, we argue that rarity holds for the antecedent in both the relevance and the irrelevance conditions. In the irrelevance condition, the consequent is "dark-haired" which is presumably rare.[1] Therefore, we would

[1] One might object that the assumption of rarity for dark hair is not appropriate for the Italian subjects who participated in this study. However, we suspect that the task instructions force an interpretation in which dark hair is relatively rare (i.e., a particularly strict standard of what counts as dark must be in play). This is because the task instructions state: "Mrs. Bianchi, who has strong views on many things, says: 'Men with dark hair love children! I bet you, if a volunteer is male, then he is dark haired." Conversational maxims suggest that utterances such as "Men with dark hair love children!" must be informative. For this utterance to be informative requires that most men have *not* got dark hair, otherwise, very little information will be conveyed because most men will be assumed to love children, irrespective of the statement. This line of thought suggests an interesting possible relationship between the pragmatic principles that relevance theory was designed to explain, and probabilistic measures of information. It may be that pragmatics affects reasoning via its impact on people's subjective probabilities.

predict that q card selections will exceed *not-q* card selections leading to the overall pattern: $n(p) > n(q) > n(\textit{not-q}) > n(\textit{not-p})$. This is exactly the ordering found in Sperber et al.'s data: $n(p) = 16 > n(q) = 12 > n(\textit{not-q}) = 7 > n(\textit{not-p}) = 5 (N = 19)$.

In the relevance condition, the consequent is "married". Because most people are married this violates the rarity assumption. Importantly on the information gain account if either $P(p)$ *or* $P(q)$ is high (or they are both high) then the expected information gain associated with the *not-q* card exceeds that associated with the q card. Therefore because $P(q)$ is high, i.e., the materials violate rarity for the consequent alone, the theory still predicts the ordering: $n(p) > n(\textit{not-q}) > n(q) > n(\textit{not-p})$. As before, this is exactly the ordering found in Sperber et al.'s data: $n(p) = 15 > n(\textit{not-q}) = 13 > n(q) = 5 > n(\textit{not-p}) = 1 (N = 17)$.

Experiment 3

Sperber et al.'s Experiment 3 contrasts two problems about employment. In the irrelevance condition, the rule is: "if a person is older than 65, then this person is without a job". Because most people are younger than 65, and most people are in work, both antecedent and consequence adhere to the rarity assumption, and hence the theory predicts the standard ordering: $n(p) > n(q) > n(\textit{not-q}) > n(\textit{not-p})$. This is exactly the ordering found in Sperber et al.'s data: $n(p) = 15 > n(q) = 10 > n(\textit{not-q}) = 9 > n(\textit{not-p}) = 5 (N = 20)$.

In the relevance condition, the rule is: "if a person is of working age, then this person has a job". Because most people are of working age, and most people have a job, both antecedent and consequent violate the rarity assumption. Therefore the theory predicts the ordering: $n(p) > n(\textit{not-q}) > n(q) > n(\textit{not-p})$. As before, this is exactly the ordering found in Sperber et al.'s data: $n(p) = 19 > n(\textit{not-q}) = 17 > n(q) = 6 > n(\textit{not-p}) = 2 (N = 20)$.

Experiment 4

Sperber et al. used four conditions in Experiment 4 corresponding to all possible combinations of high and low cognitive effects (Ec + |Ec −) and high and low effort (Et + |Et −). The materials used were very similar to those used by Kirby (1994) and involved a machine that is printing double-sided cards with letters on one side and numbers on the other side. The rule used was "if a card has a 6 in the front, it has an E on the back". We interpret all the conditions in this experiment as directly setting the parameters of the information gain account. In showing how we assume that subjects interpret "numbers" as referring to the numerals $(1, 2, \ldots, 8, 9)$.

In the high cognitive effects and low effort (Ec + |Et −) condition subjects are told that the machine prints a 4 or a 6 on the front of a card at random, it then prints an E on the back if there is a 6 on the front, and an E or an A at random if there is a 4 on the front. p (6) and *not-p* (4) are therefore equiprobable and so $P(p) = .5$. When there is a 6 on the front there is always an E printed on the back, so the probability of p, q is .5. When there is a 4 on the front then whether an A or an E gets printed on the back is equiprobable, so the probability of *not-p*, q is .25. Therefore the probability of q, $P(q) = P(p, q) + P(not\text{-}p, q) = .75$.

In the high cognitive effects and high effort condition (Ec + |Et +) subjects are told that the machine prints a number on the front of a card at random, it then prints an E on the back if there is a 6 on the front, and a letter at random if there is not a 6 on the front. The probability of p (6) is therefore 1/9, and *not-p* (*not* 6) is 8/9 and so $P(p) = 1/9$. When there is a 6 on the front there is always an E printed on the back, so the probability of p, q is 1/9. When there is another number on the front then a letter gets printed at random on the back so the probability of *not-p*, q is $8/9 \times 1/26 = 8/234$. Therefore the probability of q, $P(q) = P(p, q) + P(not\text{-}p, q) = 17/117$.

In the low cognitive effects and low effort condition (Ec − |Et −) subjects are told that the machine prints a 4 or a 6 on the front of card at random, it then prints an E or an A at random on the back. Therefore the probability of p (6) is .5, *not-p* (*not* 6) is .5, q(E) is .5, and *not-q* (*not* E) is .5. So $P(p) = P(q) = .5$.

In the low cognitive effects and high effort condition (Ec − |Et +) subjects are told that the machine prints a number on the front of a card at random; it then prints a letter at random on the back. Therefore the probability of p (6) is 1/9, *not-p* (*not* 6) is 8/9, q(E) is 1/26, and *not-q* (*not* E) is 25/26. So $P(p) = 1/9$ and $P(q) = 1/26$.

In the high cognitive effects conditions (Ec + |Et − , Ec + |Et +) subjects are told that the machine has broken down but that Mr. Bianchi has now fixed it. In the low cognitive effects conditions (Ec − |Et − , Ec − |Et +) subjects are told that the machine has broken down and that Mr. Bianchi thinks that the task rule is now in force (rather than the card faces being printed at random as they should be). In both cases an expert informs subjects that the rule is in force. Subjects should therefore assign a low value to the probability that the independence model holds, i.e., $P(M_I)$ should be low. We therefore set $P(M_I)$ to .1 and then used the parameter values derived above to compute scaled expected information gains for each card in each condition of Sperber et al.'s (1995) Experiment 4. However, in our model, the values of $P(p) = 1/9$ and $P(q) = 1/26$ in the Ec − |Et + condition are inconsistent. It is a constraint on our model that $P(q) > P(p)$, otherwise the dependence model cannot hold. A similar problem arises for rules with negated antecedents (see Oaksford & Chater, 1994, pp. 617–618) and was resolved by arguing that subjects must revise $P(p)$ down so that it is less than $P(q)$. Confronting the same situation, this is what we assume subjects do here and so we reset $P(p)$ in the Ec − |Et + condition to 1/28.

Table 1 shows the $SE(I_g)$s for each card in each condition of Sperber et

Table 1
$SE(I_g)$s for each card in each condition of Sperber et al.'s Experiment 4 showing the individual card selection frequencies in parentheses (in each condition, $N = 22$)

	Card			
	p	*not-p*	*q*	*not-q*
Ec + |Et −	1.408(18)	.527(3)	.557(7)	1.508(15)
Ec + |Et +	2.173(17)	.415(2)	.841(7)	.571(10)
Ec − |Et −	1.882(15)	.380(7)	.650(8)	1.088(11)
Ec − |Et +	2.073(17)	.370(5)	1.144(15)	.414(4)

Note: $r(14) = .89$ ($p < 0.0001$). Ec + |Et − = high effects and low effort condition; Ec + |Et + = high effects and high effort condition; Ec − |Et − = low effects and low effort condition; Ec − |Et + = low effects and high effort condition.

al.'s Experiment 4, with the individual card selection frequencies they observed in parentheses. The fit between data and model is very good ($r(14) = .89$, $p < 0.0001$). This result indicates that information gain may well provide an excellent measure of relevance in this task.

Sperber et al. (1995) go on to apply their relevance approach to other versions of the selection task, in particular the recently much studied deontic versions (e.g., Rumelhart, 1980; Griggs & Cox, 1982; Cheng & Holyoak, 1985, 1989; Manktelow & Over, 1987, 1990, 1991; Cosmides, 1989; Gigerenzer & Hug, 1992; Jackson & Griggs, 1990; Johnson-Laird & Byrne, 1991, 1992; Girotto, Mazzocco, & Cherubini, 1992). They argue that their approach is to be preferred because it generalises to these data. However, Oaksford and Chater (1994) also provide a further quantitative measure of relevance based on expected utilities that provides excellent fits to the data on the deontic selection task. So again Oaksford and Chater (1994) provide a more compelling, formal account of relevance in this domain.

4. Conclusions

This commentary has shown that information gain can provide a quantitative account of relevance in the selection task and that consequently Sperber et al.'s relevance approach and our information gain (and expected utility) approach are compatible rather than in competition. Evans (1989, 1993, 1994) has also advocated the view that subjects select those cards in the selection task that they view as relevant or salient. Sperber et al. suggest that Evans fails to "develop an explicit notion [of relevance] of his own". However, recently Over and Evans (1994) have suggested that "epistemic utility" may provide a quantitative measure of relevance in the same way as information gain. It remains to be seen whether epistemic utility can be appropriately formalised and applied to the range of selection task results in the same way as Oaksford and Chater's information gain and expected utility measures. Nevertheless the goal of uncovering suitable relevance measures now seems firmly established.

Why are relevance measures needed? The principal reason concerns the computational intractability of current theories of reasoning (Chater & Oaksford, 1990, 1993; Oaksford & Chater, 1991, 1992, 1993, 1995). All current theories tacitly assume that subjects only represent the most relevant or plausible information from which to draw inferences. In artificial intelligence (AI) the problem of retrieving relevant information from memory in order to draw inferences is known as the frame problem (Glymour, 1987). This problem has bedevilled work in AI knowledge representation since the 1960s (McArthy & Hayes, 1969). However, people do not seem to be prone to these problems – from the vast store of world knowledge people seem to unerringly access the most relevant and plausible information to solve a problem or to interpret a situation. As Sperber and Wilson (1986) identified, what linguistics and psychology requires is a well-defined theory of relevance. As Oaksford and Chater's (1994) model reveals, developing formal relevance measures may also resolve many outstanding problems in the psychology of reasoning.

Acknowledgements

We thank Dan Sperber, Francesco Cara and Vittorio Girotto for providing us with the data from their experiments without which we could not have written this reply.

References

Anderson, J.R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.

Anderson, J.R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, *14*, 471–517.

Chater, N., & Oaksford, M. (1990). Autonomy, implementation and cognitive architecture: a reply to Fodor and Pylyshyn. *Cognition*, *34*, 93–107.

Chater, N., & Oaksford, M. (1993). Logicism, mental models and everyday reasoning. *Mind and Language*, *8*, 72–89.

Cheng, P.W., & Holyoak, K.J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391–416.

Cheng, P.W., & Holyoak, K.J. (1989). On the natural selection of reasoning theories. *Cognition*, *33*, 285–313.

Cohen, L.J. (1981). Can human irrationality be experimentally demonstrated? *Behavioural and Brain Sciences*, *4*, 317–370.

Cosmides, L. (1989). The logic of social exchange: has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, *31*, 187–276.

Duhem, P. (1954). *The aim and structure of physical theory*. Princeton, NJ: Princeton University Press. (Originally published 1914.)

Earman, J. (1992). *Bayes or bust?* Cambridge, MA: MIT Press.

Evans, J.St.B.T. (1989). *Bias in human reasoning: Causes and consequences.* Hillsdale, NJ: Erlbaum.

Evans, J.St.B.T. (1993). Bias and rationality. In K.I. Manktelow, & D.E. Over (Eds.), *Rationality* (pp. 6–30). London: Routledge.

Evans, J.St.B.T. (1994) Relevance and reasoning. In S.E. Newstead & J.St.B.T. Evans (Eds.), *Perspectives on thinking and reasoning.* Brighton: Erlbaum.

Evans, J.St.B.T., & Lynch, J.S. (1973). Matching bias in the selection task. *British Journal of Psychology, 64,* 391–397.

Gigerenzer, G., & Hug, K. (1992). Domain-specific reasoning: social contracts, cheating, and perspective change. *Cognition, 43,* 127–171.

Girotto, V., Mazzocco, A., & Cherubini, P. (1992). Judgements of deontic relevance in reasoning: a reply to Jackson and Griggs. *Quarterly Journal of Experimental Psychology, 45A,* 547–574.

Glymour, C. (1987). Android epistemology and the frame problem: comments on Dennett's "Cognitive Wheels". In Z.W. Pylyshyn (Ed.), *The robot's dilemma: The frame problem in artificial intelligence* (pp. 65–76). Norwood, NJ: Ablex.

Good, I.J. (1966). A derivation of the probabilistic explication of information. *Journal of the Royal Statistical Society, Series B, 28,* 578–581.

Griggs, R.A., & Cox, J.R. (1982). The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology, 73,* 407–420.

Griggs, R.A., & Cox, J.R. (1983). The effects of problem content and negation on Wason's selection task. *Quarterly Journal of Experimental Psychology, 35A,* 519–533.

Horwich, P. (1982). *Probability and evidence.* Cambridge, UK: Cambridge University Press.

Howson, C., & Urbach, P. (1989). *Scientific reasoning: The Bayesian approach.* La Salle, IL: Open Court.

Jackson, S.L., & Griggs, R.A. (1990). The elusive pragmatic reasoning schemas effect. *Quarterly Journal of Experimental Psychology, 42A,* 353–373.

Johnson-Laird, P.N., & Byrne, R.M.J. (1991). *Deduction.* Hillsdale, NJ: Erlbaum.

Johnson-Laird, P.N., & Byrne, R.M.J. (1992). Modal reasoning, models, and Manktelow and Over. *Cognition, 43,* 173–182.

Johnson-Laird, P.N., & Wason, P.C. (1970a). A theoretical analysis of insight into a reasoning task. *Cognitive Psychology, 1,* 134–148.

Johnson-Laird, P.N., & Wason, P.C. (1970b). Insight into a logical relation. *Quarterly Journal of Experimental Psychology, 22,* 49–61.

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

Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation and information in hypothesis testing. *Psychological Review, 94,* 211–228.

Koyré, A. (1957). *From the closed world to the infinite universe.* Baltimore, ML: Johns Hopkins University Press.

Kuhn, T.S. (1962). *The structure of scientific revolutions.* Chicago: University of Chicago Press.

Lindley, D.V. (1956). On a measure of the information provided by an experiment. *Annals of Mathematical Statistics, 27,* 986–1005.

MacKay, D.J.C. (1992). *Bayesian methods for adaptive models.* Unpublished doctoral dissertation, California Institute of Technology, Passedena, CA.

Manktelow, K.I., & Evans, J.StB.T. (1979). Facilitation of reasoning by realism: effect or non-effect. *British Journal of Psychology, 70,* 477–488.

Manktelow, K.I., & Over, D.E. (1987). Reasoning and rationality. *Mind and Languagе, 2,* 199–219.

Manktelow, K.I., & Over, D.E. (1990). Deontic thought and the selection task. In K.J. Gilhooly, M. Keane, R. Logie, & G. Erdos (Eds.), *Lines of thought: Reflections on the psychology of thinking.* Chichester: Wiley.

Manktelow, K.I., & Over, D.E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition, 39,* 85–105.

Manktelow, K.I., & Over, D.E. (1993). *Rationality: Psychological and philosophical perspectives*. London: Routledge.
McCarthy, J.M., & Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.), *Machine intelligence* (Vol. 4), New York: Elsevier.
Oaksford, M., & Chater, N. (1991). Against logicist cognitive science. *Mind and Language*, *6*, 1-38.
Oaksford, M., & Chater, N. (1992). Bounded rationality in taking risks and drawing inferences. *Theory and Psychology*, *2*, 225-230.
Oaksford, M., & Chater, N. (1993). Reasoning theories and bounded rationality. In K.I. Manktelow, & D.E. Over (Eds.), *Rationality* (pp. 31-60). London: Routledge.
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). Theories of reasoning and the computational explanation of everyday inference. *Thinking and Reasoning*, *1*, 121-152.
Oaksford, M., & Stenning, K. (1992). Reasoning with conditionals containing negated constituents. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *18*, 836-854.
Over, D.E., & Evans, J.St.B.T. (1994). Hits and misses: Kirby on the selection task. *Cognition*, *52*, 235-243.
Pollard, P. (1985), Nonindependence of selections on the Wason selection task. *Bulletin of the Psychonomic Society*, *23*, 317-320.
Popper, K.R. (1959). *The logic of scientific discovery*. London: Hutchinson. (Originally published 1935.)
Putnam, H. (1974). The "corroboration" of theories. In P.A. Schilpp (Ed.), *The Philosophy of Karl Popper* (Vol. 1, pp. 221-240). La Salle, IL: Open Court.
Quine, W.V.O. (1953). Two dogmas of empiricism. In W.V.O. Quine (Ed.), *From a logical point of view* (pp. 20-46). Cambridge, MA: Harvard University Press.
Reich, S.S., & Ruth, P. (1982). Wason's selection task: verification, falsification and matching. *British Journal of Psychology*, *73*, 395-405.
Rumelhart, D.E. (1980). Schemata: the building blocks of cognition. In R.J. Spiro, B.C. Bruce, & W.F. Brewer (Eds.), *Theoretical issues in reading comprehension*. Hillsdale, NJ: Erlbaum.
Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, *57*, 31-95.
Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford: Basil Blackwell.
Stich, S. (1985). Could man be an irrational animal? *Synthese*, *64*, 115-135.
Stich, S. (1990). *The fragmentation of reason*. Cambridge, MA: MIT Press.
Toulmin, S. (1961). *Foresight and understanding*. London: Hutchinson.
Wason, P.C. (1966). Reasoning. In B. Foss (Ed.), *New horizons in psychology*. Harmondsworth, UK: Penguin.
Wason, P.C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*, 273-281.
Wason, P.C. (1969). Regression in reasoning. *British Journal of Psychology*, *60*, 471-480.