# *Rational Models of Cognition*


Edited by

**Mike Oaksford**
*School of Psychology,*
*University of Wales, Cardiff*

and

**Nick Chater**
*Department of Psychology,*
*University of Warwick*



Oxford · New York · Tokyo
OXFORD UNIVERSITY PRESS
1998

# 17 *A revised rational analysis of the selection task: exceptions and sequential sampling*

Mike Oaksford and Nick Chater

Results in the psychology of reasoning appear to show that people make many errors when confronted with tasks having superficially obvious logical solutions. The task most often used to illustrate this point in both the philosophical and the psychological literature is Wason's (1966, 1968) selection task. In the selection task an experimenter presents participants with four cards, each with a number on one side and a letter on the other, and a rule of the form *if p then q*, e.g. *if there is a vowel on one side* (*p*), *then there is an even number on the other side* (*q*). The four cards show an 'A' (*p* card), a 'K' (*not-p* card), a '2' (*q* card) and a '7' (*not-q* card). Participants have to select those cards that they must turn over to determine whether the rule is true or false. Logically, participants should select only the *p* and the *not-q* cards. However, as few as 4% of participants make this response, other responses being far more common (*p* and *q* cards (46%); *p* card only (33%), *p*, *q* and *not-q* cards (7%), *p* and *not-q* cards (4%) (Johnson-Laird and Wason, 1970*a*).

Oaksford and Chater (1994) provided a *rational analysis* (Anderson, 1990) of the selection task based on Bayesian optimal data selection (henceforth, the ODS model; Lindley, 1956; Fedorov, 1972; Mackay, 1992). They argued that participants' behaviour reflects a rational strategy of optimizing the expected amount of information gained by turning each card. On this view, the selection task is not a logical reasoning task but a task of probabilistic optimal data selection in inductive hypothesis testing. Oaksford and Chater (1994) also generalize their optimal data selection model to all the main experimental results on the selection task. Specifically, it accounts for the non-independence of card selections (Pollard, 1985), the negations paradigm (e.g. Evans and Lynch, 1973), the therapy experiments (e.g. Wason, 1969), the reduced array selection task (RAST) (Johnson-Laird and Wason, 1970*b*), and work on so-called fictional outcomes (Kirby, 1994). Oaksford and Chater (1994) also showed how a related maximum expected utility model accounts for deontic versions of the selection task (e.g. Cheng and Holyoak, 1985), where participants must reason about how one ought to behave, including perspective and rule-type manipulations (e.g. Cosmides, 1989;



Gigerenzer and Hug, 1992), and the manipulation of probabilities and utilities (Kirby, 1994).

The status of Oaksford and Chater's model is contentious (see, e.g. Almor and Sloman, 1996; Evans and Over, 1996*a*; Laming, 1996; and for a reply, Oaksford and Chater, 1996). On the one hand, support for this model derives from Oaksford and Chater's (1995*b*) re-interpretation of Sperber *et al.*'s (1995) results, and the recent results of Manktelow *et al.*, (1995, see, Oaksford *et al.*, 1997). In particular, Oaksford *et al.* (1997*a*) have shown that probabilistic manipulations in the 'RAST', (Girotto, 1988; Girotto *et al.* 1988, 1989; Johnson-Laird and Wason, 1970*b*; Light *et al.*, 1989; Wason and Green, 1984)—where participants only choose between the *q* and the *not-q* cards—follow the predictions of ODS. On the other hand, Evans and Over (1996*a*) argue that the ODS model fails to capture the data from Kirby (1994) and that data from Pollard and Evans (1983) falsify the ODS model. Moreover, the recent work of Oaksford *et al.* (1997), although generally supportive of ODS, revealed some effects of sequential sampling in the RAST that may require some revisions to the model.

The goal of this chapter is to show that some straightforward revisions of the ODS model to deal with exceptions and sequential sampling can: (i) meet the objections raised by Evans and Over (1996*a*), and (ii) explain the effects of sequential sampling in the RAST, observed by Oaksford *et al.* (1997). We first introduce the ODS framework.

## Optimal data selection

In this section we outline Oaksford and Chater's ODS model (we refer the reader to Oaksford and Chater, 1994, for the complete description of the model).

In Wason's selection task (Wason, 1966, 1968), participants confront a problem that is analogous to 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. Oaksford and Chater's (1994) model is based on contemporary Bayesian accounts of scientific inference that reject Popper's (1959) falsificationist view that only potentially falsifying evidence should be sought (Horwich, 1982; Howson and Urbach, 1989; Earman, 1992). Bayesian accounts adopt an explicitly *subjective* as opposed to a *frequentist* approach to probability. On the subjective interpretation, probabilities are degrees of belief (Keynes, 1921; Ramsey, 1931) rather than limiting frequencies (e.g. von Mises, 1939). Oaksford and Chater's model is about how prior beliefs affect judgements about the most informative data to select.[1] In particular, peoples' prior beliefs about the probabilities of the antecedents, *p*, and consequents, *q*, of rules, *if p then q*, play a central role.

Oaksford and Chater (1994) suggest that hypothesis testers should choose experiments (select cards) that 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* (although *qs* are not invariably associated with *ps*), and (ii) that the occurrence of *ps* and *qs* are independent. Participants' prior degree of belief in (ii) is $P(M_I)$ and their prior degree of belief in (i) is $P(M_D)$, i.e. $1 - P(M_I)$. Where '$M_I$' refers to the contingency table representing independence ($I$) between $p$ and $q$, and '$M_D$' refers to the contingency table representing a dependency ($D$) between $p$ and $q$ (see Table 17.1). For most purposes Oaksford and Chater assume that these are equally likely, i.e. $P(M_I) = 0.5$. For each hypothesis,

**Table 17.1 (a) Shows the contingency table appropriate for the dependence model $M_D$, where there is an exceptionless dependency between the *p* and *q* (*b*) Shows the equivalent table for the independence model $M_I$. *a* corresponds to the probability of *p*, *P*(*p*), and *b* corresponds to the probability of *q* in the absence of *p*, *P*(*q*|*not-p*)**

| $M_D$ | q | not-q | $M_I$ | q | not-q |
|---|---|---|---|---|---|
| *p* | *a* | 0 | *p* | *ab* | *a*(1 − *b*) |
| *not-p* | (1−*a*)*b* | (1 − *a*) (1 − *b*) | *not-p* | (1 − *a*)*b* | (1 − *a*) (1 − *b*) |

Oaksford and Chater (1994) define probability models ($M_I$ and $M_D$) that derive from participants' prior beliefs about 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 (Wiener, 1948; Shannon and Weaver, 1949). This is the same approach to optimal data selection proposed by Lindley (1956). Thus, Oaksford and Chater (1994) define the information gain of data $D$ as:

$$\text{information before receiving } D\text{:} \quad I(H_i) = -\sum_{i=1}^{n} P(H_i)\log_2 P(H_i) \qquad 17.1$$

$$\text{information after receiving } D\text{:} \quad I(H_i|D) = -\sum_{i=1}^{n} P(H_i|D)\log_2 P(H_i|D) \qquad 17.2$$

$$\text{information gain:} \quad I_g = I(H_i) - I(H_i|D) \qquad 17.3$$

The posterior probabilities ($P(H_i|D)$) are calculated 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 (that is, which card to turn), participants do not know what that data will be (that is, what will be on the back of the card). So they cannot calculate actual information gain. However, they 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.



Given the expected information gains associated with each card, a decision has to be made about which cards to select. Oaksford and Chater (1994) incorporated two aspects of the decision process in their measure. First, they introduced a noise factor by adding 0.1 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. To reflect this Oaksford and Chater (1994) scaled their information gain measure by the mean value for all four cards. Consequently, card choice is relative to the total expected information gain available and is not determined by the absolute $E(I_g)$ value alone. The higher the proportion of the total available $E(I_g)$ a card possesses the more likely it is to be selected. Oaksford and Chater refer to this derived 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, i.e. they have a low probability of occurrence. Take the rule *all ravens are black*, for example, the probability that any given bird is a raven is low, as is the probability that it is black. The 'rarity assumption' seems to apply to the vast majority of everyday categories that are used to construct hypotheses about the world. Moreover, there is evidence that people adopt this assumption from the literature on other reasoning tasks (Klayman and Ha, 1987; Anderson, 1990, see, Oaksford and Chater, 1994). Oaksford and Chater (1996) point out that further evidence for the rarity assumption comes from the normative literature on Bayesian epistemology (e.g. Horwich, 1982; Howson and Urbach, 1989). Making a rarity assumption resolves the ravens paradox of non-Bayesian confirmation theory (Goodman, 1954), whereby non-black, non-ravens, e.g. a pair of white socks, must confirm the hypothesis that all ravens are black. Consequently, there are strong normative and empirical grounds for Oaksford and Chater's (1994) assumption that people's strategies for dealing with conditional rules are adapted to the case where rarity holds.

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

$$SE(I_g(p)) > SE(I_g(q)) > SE(I_g(not{-}q)) > SE(I_g(not{-}p)) \qquad 17.4$$

This corresponds to the observed frequency of card selections in Wason's task: *n*(*p*) > *n*(*q*) > *n*(*not-q*) > *n*(*not-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:

$$SE(I_g(p)) > SE(I_g(not{-}q)) > SE(I_g(q)) > SE(I_g(not{-}p)) \qquad 17.5$$

This ordering is more consistent with Popperian falsificationism, which favours the $p$ and *not-q* instances. The effects of rarity and its violation permit us to explain the range of results we outlined above and make definite predictions in the RAST.

## Implementation and sensitivity

The ODS model provides a rational analysis (Anderson, 1990, 1994) of the selection task that suggests that manipulating $P(p)$ or $P(q)$ should lead to predictable variations in the proportions of cards selected. However, as Oaksford and Chater (1994) and Oaksford *et al.* (1997) discuss, the level of detail at which the model can make predictions also depends on how the cognitive system *implements* this model (see also, Anderson, 1990). As Oaksford *et al.* (1997) argue, the critical question is how sensitive can we expect people to be to changes in $P(p)$ or $P(q)$ and consequently to changes in the $SE(I_g)$ values?

At one extreme the cognitive system may implement the rational analysis directly, i.e. it may perform all the computations specified by the model. If this is the case then varying the parameters of the model should lead to card selections that directly mirror the resulting $SE(I_g)$ values. At the other extreme, the cognitive system may implement this analysis via a hard-wired and cognitively impenetrable (Pylyshyn, 1984) heuristic that has evolved to deal with an environment where rarity is the norm. If this is the case then, although our model would explain why selecting the $p$ and $q$ cards is an adaptive rational strategy, we could not predict any performance variation in response to varying the model's parameters.

To explain the data, Oaksford and Chater (1994) already assume that people are sensitive to manipulations of $P(p)$ and $P(q)$, and consequently we regard the second possibility as implausible. We also regard the first possibility as implausible because, as Oaksford and Chater (1994) argue, the full Bayesian analysis it assumes is likely to prove computationally intractable when scaled up to real human reasoning (see also, Chater and Oaksford, 1990; Oaksford and Chater, 1991, 1992, 1993, 1995*a*). In sum, the truth must lie somewhere between these two extremes of perfect sensitivity to changes in $P(p)$ or $P(q)$ and no sensitivity to such changes.

Oaksford *et al.* (1997) argue that sensitivity may depend on a variety of factors. If people compute and mentally represent something analogous to $SE(I_g)$ values, then noise and/or imperfect transduction may lead to reduced sensitivity. Moreover, there is quite a broad region where $SE(Ig(not-q)) \approx SE(I_g(q))$ as $P(p)$ or $P(q)$ vary. Depending on the discriminability between the mental analogues of $SE(I_g(not-q))$ and $SE(I_g(q))$, this could lead to quite a broad region of uncertainty about which card to choose. Oaksford *et al.* (1997) explore the possible consequences of this uncertainty for performance on the reduced array version of the selection task.

One immediate consequence, which Oaksford and Chater (1994) pointed out, is that it may be difficult to override habitual strategies by explicit instruction in psychological experiments. In particular, in the selection task we would expect it to be difficult for participants to violate rarity to which Oaksford and Chater (1994) assume the cognitive system is adapted. Evans and Over (1996*b*, p. 20) reiterate this point, arguing that many 'habitual methods of reasoning . . . will not easily be modified by presentation of verbal instructions for the sake of experiment'. As we will see later on, despite the apparent unanimity on this point, it will provide a point of disagreement between us and Evans and Over.



## Exceptions and Evans and Over

In this section we observe that some recent criticisms of the ODS model by Evans and Over (1996*a*) centre on the fact that, in the version presented in Oaksford and Chater (1994), exceptions were not permitted in the dependence model, i.e. $P(not\text{-}q|p)$ was 0. Evans and Over (1996*a*) argue that a consequence of not allowing exceptions is that the ODS model cannot really explain Kirby's (1994) data which, Oaksford and Chater (1994) argued, provided strong support for ODS, nor can it account for data reported by Pollard and Evans (1983).

Most everyday generalizations about the world, such as *birds fly*, admit exceptions. Indeed, elsewhere we have discussed at length the possible consequences of exceptions for the adequacy of logicist approaches to cognition in general (Chater and Oaksford, 1990, 1993; Oaksford and Chater, 1991, 1992) and to human reasoning in particular (Oaksford and Chater, 1991, 1992, 1993, 1995*a*). Our own previous research therefore clearly indicates the need to consider the consequences of allowing exceptions in the ODS dependence model. In this section we make this modification and re-model the data reported in Kirby (1994*a*) and in Pollard and Evans (1983).

### Kirby (1994)

Kirby's experimental instructions described a device for printing cards that had just made an error after printing out 100 cards (in experiment 1) or 10 cards (in experiment 2). Evans and Over note that this means that participants know that any exceptionless and universal rule describing the device's behaviour must be false. If so, they argue, the ODS model cannot apply because there is no information to be gained—the rule is known to be false from the outset. Consequently, they argue that for our model to apply to Kirby's experiments at all, participants must interpret the rule not as universal (i.e. applying to all cards that the device may print), but as applying only to the four cards presented in the task. The subject's task is to decide whether these particular cards contain errors. On this assumption, they use the error rate specified in the instructions to compute $P(M_I)$ values that differ from those used by Oaksford and Chater (1994).

In the original ODS model, Oaksford and Chater (1994) always assume that participants regard the four cards in the selection task as a sample from a larger population of cards, over which the conditional rule is defined. Evans and Over are correct to point that this interpretation apparently cannot apply to Kirby's experiments, where participants know that there are exceptions to the rule. We have argued elsewhere that people do not interpret everyday conditional rules as exceptionless (Chater and Oaksford, 1990, 1993; Oaksford and Chater, 1991, 1992, 1993, 1995*a*; Oaksford, 1993). Any everyday rule, such as *birds fly*, *if you put money in the coke machine, you get a coke*, and so on, succumb to indefinitely many exceptions, such as ostriches, penguins and broken or empty coke machines. In the original ODS model, Oaksford and Chater (1994) did not allow for exceptions for simplicity, and because they did not appear necessary to model the data.

Exceptions can be straightforwardly incorporated in our model by the introduction of an exception parameter, $e$, in the dependence model. $e$ is the probability of *not-q* given $p$, in the dependence model $M_D$ ($P(not-q|p, M_D)$) This involves the following changes to the $p$, $q$ and $p$, *not-q* cells of Oaksford and Chater's original dependence model (see Table 17.1 above): $P(p, q|M_D) = a(1-e)$ and $P(p, not-q|M_D) = ae$. We assume that $e$ has a small fixed value—rules will tolerate some exceptions, but must be rejected in the face of large numbers of exceptions. Once exceptions are allowed, then there is a genuine question concerning whether the rule holds or not, and Oaksford and Chater's information gain analysis can therefore be applied. This provides an intuitively obvious alternative to Evans and Over's analysis of the task, from which they derive their $P(M_I)$ values.

We now show that adding exceptions to Oaksford and Chater's (1994) account does not substantially alter its predictions. Following Evans and Over, we assume that subjects interpret the error rate to refer to the proportion of *not-q*s associated with *p*s. $e$ corresponds to the proportion of such cases that would be predicted from $M_D$. Apart from the fact that $e$ must be small (because if there are too many exceptions then the rule will presumably be rejected), we have no grounds to set $e$ at any particular value. We do assume, however, that $e$ is large enough to tolerate the error rates in Kirby's experiments, so that subjects are likely to believe the rules to be true (that is, $P(M_D)$ is high—specifically, we set $P(M_D) = 0.99$ throughout). Fortunately, our predictions do not appear sensitive to the precise value of $e$. We set $e = 0.1$ and $e = 0.01$, which happen to correspond to the error rates used by Kirby. However, note that $e$ is the tolerance to exceptions of the participant, and is not directly determined by the experimental set-up. As in Oaksford and Chater's original analysis, we set $P(p)$ directly from Kirby's instructions. Oaksford and Chater somewhat unrealistically set $P(q) = P(p)$. Adopting this assumption here allows good fits with the data, but we take the opportunity here to show that this assumption is not necessary. The basic restriction on our model is that $P(p) \leq P(q)$. We somewhat arbitrarily set $P(q) = 2/1001$ when $P(p) = 1/1001$ and $P(q) = 1000.5/1001$ when $P(p) = 1000/1001$ in Kirby's experiment 1. We also set $P(q) = 2/100$ when $P(p) = 1/100$, $P(q) = 75/100$ when $P(p) = 50/100$, and $P(q) = 95/100$ when $P(p) = 90/100$ in Kirby's experiments 2 and 3. These values were not optimized to produce the best fits with the data.

As in Oaksford and Chater (1994), for each experiment we generated $SE(I_g)$ values for each card, and computed the Spearman rank order correlation coefficient with the selection frequencies in Kirby. In experiment 1, correlations were positive, but weaker than found by Oaksford and Chater (1994), and not statistically significant (there are only eight data points in the analysis). We found that $\rho(N = 8) = 0.69$, $p = 0.07$, with $e = 0.01$, and that $\rho(N = 8) = 0.60$, $p = 0.14$, with $e = 0.1$. In experiment 2, the correlations were comparable with those obtained by Oaksford and Chater: $\rho(N = 12) = 0.92$, $p < 0.01$, with $e = 0.01$ and $\rho(N = 12) = 0.84$, $p < 0.01$, with $e = 0.1$. In experiment 3, introducing exceptions produces much better fits than Oaksford and Chater (1994) obtained: $\rho(N = 12) = 0.80$, $p < 0.01$, with $e = 0.01$ and $\rho(N = 12) = 0.85$, $p < 0.01$ with $e = 0.1$. The fit across the three experiments is very similar to Oaksford and Chater (1994). The principal difference is that Oaksford and Chater found the worst fit to experiment 3, while the current analysis provides the worst fit for experiment 1. Interestingly, Over and Evans (1994) argue that experiment 1 is the least interpretable of Kirby's experiments, and hence attention should focus on experiments 2 and 3.[1]



Evans and Over (1996*a*) raise an important issue about the compatibility of our model with Kirby's data concerning the presence of exceptions. The revised ODS model includes the possibility of exceptions to take account of Evans and Over's observation. Importantly, the model fits with the empirical data are comparable with Oaksford and Chater's original analysis. Moreover, incorporating exceptions is consistent with our previous work, which emphasizes that everyday conditionals always admit exceptions (e.g. Chater and Oaksford, 1990, 1993; Oaksford and Chater, 1991, 1992, 1993, 1995*a*).

## Pollard and Evans (1983)

Evans and Over (1996*a*) argue that Pollard and Evans' (1983) data falsifies our model. Pollard and Evans gave participants an initial learning phase with seven *p*, *q* cards, one *p*, *not-q* card, seven *not-p*, *q* cards, and seven *not-p*, *not-q* cards. Participants predict what is on the reverse of each card, before turning it over. The experimenter then draws four cards from this set and participants perform the selection task. In a 'usually true' condition participants consider a rule of the form *if p then q*, and in a 'usually false' condition they consider a rule of the form *if p then not-q*.

Evans and Over assume that participants can directly estimate the probabilities used in our model from the card frequencies, so that $P(p) = 0.36$ and $P(q) = 0.64$ and they argue for a way of calculating priors that would lead participants to set $P(M_I) = 0.242$ in the usually true condition, and $P(M_I) = 0.940$, in the usually false conditions. Putting these values into our model, we find the scaled $E(I_g)$ for the *not-q* card is 0.198 in the usually true condition and 0.039 in the usually false condition.[2] This predicts more *not-q* card selections in the usually true condition than in the usually false condition, whereas Pollard and Evans obtained the opposite result.

Evans and Over (1996*a*) argue that this result falsifies the ODS model. However, there are several reasons to doubt the significance of this isolated and unreplicated result. First, the conclusion that Evans and Over wish to draw from this result is only clear from a falsificationist perspective (Popper, 1959) where a single predictive failure is taken to falsify a hypothesis. However, the ODS model (see Oaksford and Chater, 1994, 1996) and Evans and Over (1996*b*) both explicitly reject Popper's falsificationist philosophy. It is inconsistent of Evans and Over (1996*a*) to employ a falsificationist argument against ODS, that they reject in their own theoretical work (Evans and Over, 1996*b*). Second, Evans and Over (1996*b*) have suggested (see above) that altering people's habitual strategies by explicit instruction may be very difficult. However, in the case of Pollard and Evans' experiments, Evans and Over (1996*a*) are happy to assume that the learning phase has overridden participants' habitual rarity values for $P(p)$ and $P(q)$. This would be legitimate if Pollard and Evans (1983) provided evidence that their procedure leads their participants to abandon these habitual values. However, they did not report any measures that could bear on this question.[3]

We now argue that it is quite consistent for participants in Pollard and Evans' experiment to succeed on the learning task while retaining their habitual rarity values for $P(p)$ and $P(q)$. In Pollard and Evans' (1983) learning task, participants are required to predict what is on the other side of a card. Participants are not asked to estimate the proportion of *p* or *q* cards they have seen. Therefore, the prediction task relies on knowledge of the conditional probabilities $P(q|p)$ only, not on knowledge of $P(p)$ or $P(q)$. Consequently, it is reasonable to argue that although participants may accurately estimate $P(q|p)$, because this is the quantity on which success at the task relies, they will be *un*able to estimate accurately $P(p)$ or $P(q)$, because the task does not draw their attention to these values. Indeed, it seems reasonable to suggest that although participants estimate $P(q|p)$ from the data, they continue to adhere to their habitual strategy of assuming rarity values for $P(p)$ and $P(q)$. This is consistent with our revised model with exceptions. Note that on this model $P(q|p, M_D)$ is not determined by $P(p)$ and $P(q)$, this is because $P(q|p, M_D) = 1 - e$, i.e. it depends on the value of the exceptions parameter alone. Consequently, $P(q|p)$ can be regarded as simply providing participants with an estimate of *e*.[4]

None the less, it could be argued that participants estimate $P(p)$ and $P(q)$ implicitly. However, the standard Bayesian assumption would be that their estimates should depend on prior beliefs (in this case, our rarity assumption) as well as current observations. One psychological account of how this occurs is that updating beliefs depends on the sample size on which the priors are based (Gigerenzer, 1994). For example, if your prior $P(p)$ is 0.2 based on $N = 1000$, then observing one *p* card won't alter your beliefs very much, i.e. $P(p)$ now equals 0.2008. Alternatively, if your prior $P(p)$ is 0.2 based on $N = 5$ then observing one *p* card will have a far greater effect, i.e. $P(p)$ now equals 0.33. Oaksford and Chater argue that the rarity assumption is ubiquitous. Consequently, the sample sizes on which participants base their prior beliefs about $P(p)$ and $P(q)$ are presumably very large. This suggests that the experience participants have of the materials in Pollard and Evans is unlikely to move their beliefs about $P(p)$ and $P(q)$ very far from their default rarity values.

We have argued that all of Evans and Over's predictions using the ODS model rely on assumptions about what participants have implicitly learned about $P(p)$ and $P(q)$, about which there is no empirical evidence. Hence their claim that these tasks unambiguously falsify Oaksford and Chater's model seems premature. Evans and Over's arguments would make a strong case, if there were no reasonable set of parameters in Oaksford and Chater's model that could capture Pollard and Evans (1983) data. We now show that there is such a reasonable set of parameters. But first we consider a consequence of introducing the exceptions parameter.

Evans and Over manipulate $P(M_I)$ in order to model whether people believe the task rule to be true or false. They assume that a low probability rule (i.e. where $P(q|p)$ is low) will not be believed. However, once exceptions are admitted the possibility must be allowed that a low probability rule may be believed to be in force. A classic example from the philosophy of science illustrates that this possibility is intuitively reasonable (Bromberger, 1965). Paresis is an unpleasant disease that someone can only contract if they have latent untreated syphilis. However, the probability of catching paresis given that you have syphilis is only about 0.01, i.e. the probability of an exception is 0.99. Nonetheless physicians have a strong belief that this weak relationship holds! In the context of modelling Pollard and Evans experiments this means that simply because $P(q|p)$ was low or high cannot be unambiguously



interpreted as resulting in a correspondingly low or high belief in the rule, especially where that rule might be interpreted as a cause–effect relationship.

Although, the presence of a large number of exceptions does not always lead to disbelief in a rule, along with Evans and Over, we will assume that they go together.[5] Consequently, we model the believed true condition as having a small number of exceptions *and* as having a low value for $P(M_I)$ and we model the believed false condition as having a large number of exceptions *and* as having a high value for $P(M_I)$. The precise values of $P(M_I)$ are not critical, but for the sake of argument we use Evans and Over's estimates for $P(M_I)$: in the believed true condition, $P(M_I) = 0.242$, and in the believed false condition, $P(M_I) = 0.940$. For *e* we will use the empirical values: in the believed true condition, $e = 0.125$, and in the believed false condition, $e = 0.875$. For the reasons we outlined above, we retain the default rarity values for $P(p)$ and $P(q)$, ignoring the possibility of implicit learning ($P(p) = 0.1$ and $P(q) = 0.11$). Table 17.2 shows the $SE(I_g)$s for each card in the usually true and the usually false conditions. Fits to the standard *if p then q* rule form were good, $r(6) = 0.87$, $P < 0.005$, as were the fits to the overall data, including all rules types, $r(6) = 0.86$, $P < 0.01$. These close fits suggest that Pollard and Evans experiments can be viewed as consistent with the ODS model.

**Table 17.2 $SE(Ig)$s for each card in the usually true (True) and the usually false (False) condition of Pollard and Evans' (1983) experiment 2 showing the frequency of card selections for the standard *if p then q rule* (If *p*, *q*) and for all four rules (All) collapsed**

| Condition | | Cards | | | |
|---|---|---|---|---|---|
| | | *p* | *not-p* | *q* | *not-q* |
| True* | $SE(I_g)(e = 0.125)$ | 1.990 | 0.365 | 1.190 | 0.451 |
| | If *p*, *q* | 96 | 13 | 83 | 29 |
| | All | 95 | 17 | 77 | 35 |
| False† | $SE(I_g)(e = 0.875)$ | 1.001 | 0.999 | 1.001 | 0.999 |
| | If *p*, *q* | 75 | 33 | 58 | 50 |
| | All | 78 | 33 | 56 | 58 |

*$P(M_I) = 0.242$; †$P(M_I) = 0.940$; $P(p) = 0.1$; $P(q) = 0.11$. For the *if p, then q* rule (if *p*, *q*), $e = 0.125$, $r(6) = 0.87$, $P < 0.005$. For all rules collapsed (All), $e = 0.875$, $r(6) = 0.86$, $P < 0.01$.

## Conclusions

In this section we have shown that a revised ODS model is compatible with the empirical data cited by Evans and Over (1996*a*), under reasonable parameter values. Importantly, the only real change required was the addition of exceptions, a change that is consistent with our continued insistence that the possibility of exceptions has profound consequences for theories of reasoning. Our proposal for revising ODS and the parameter values we recommend, like those suggested by Evans and Over, must ultimately be assessed by further empirical research to provide a more rigorous test of Oaksford and Chater's and other related models.

## Sequential sampling in the reduced array selection task

In the last section we showed how modifying the ODS model to take account of exceptions successfully captured some aspects of the empirical data that the original model seemed unable to explain. In this section we explore whether the recent results of Oaksford *et al.* (1997) in the RAST similarly require modifications to the ODS model.

In the RAST participants choose between the *q* and *not-q* options only (hence 'reduced array'; Johnson-Laird and Wason, 1970*b*; Wason and Green, 1984). The stimuli in the original RAST consisted of 30 coloured shapes. The experimenter informs the participants that there are 15 black shapes and 15 white shapes, each of which is a triangle or a circle. The shapes are in two boxes one containing the white shapes and the other containing the black shapes. On being presented with a test sentence, e.g. *all the triangles are black*, participants have to assess the truth or falsity of the sentence by asking to see the *least* number of black or white shapes. In Johnson-Laird and Wason (1970*b*), although all participants chose some confirmatory black shapes (no participant chose more than nine), they all chose all 15 potentially falsificatory white shapes. Thus, where participants in effect perform multiple selection tasks, they tend to show falsificatory behaviour. Wason and Green (1984) report a variant on the RAST (see, Oaksford and Chater, 1994) and Girotto and Light and their colleagues (Girotto, 1988; Girotto *et al.*, 1988, 1989; Light *et al.*, 1989) have used it in developmental studies using thematic content.

Oaksford and Chater (1994) suggest the following explanation for the basic findings on the RAST. The RAST makes explicit that the rule applies to a limited domain of cards or shapes that the experimenter describes as being in a box or in a bag (or in Wason and Green (1984) 'under the bar'). The experimenter also informs participants that in this limited domain there are equal numbers of *q* and *not-q* instances. It follows that $P(q)=P(not\text{-}q)=0.5$, violating the rarity assumption. If participants are sensitive to these experimentally given frequencies, then this leads to a value of $SE(I_g(not\text{-}q))$ which is higher than $SE(I_g(q))$. Consequently, ODS predicts more *not-q* card selections than *q* card selections as is typically observed in the RAST.

Oaksford *et al.* (1997) tested this explanation of performance on the RAST by systematically varying $P(q)$. They did this by using stacks of cards rather than boxes of coloured shapes. The number of cards in each stack was varied to achieve the probability manipulation. By varying these probabilities they were able to show that the proportions of *q* and *not-q* cards selected varied according to the ODS model, i.e. as $P(q)$ falls, *q* card selections rise and *not-q* card selections fall.

A crucial feature of the RAST, that distinguishes it from the standard selection task, is that it involves sequential sampling, i.e. participants can look at many cards and at each trial they can turn the card to see what is on the other side. This procedure raises many questions about the resulting sequence of selections that subjects make that are not addressed by the current ODS model. This is pressing because Oaksford *et al.*'s (1997) results revealed some minor discrepancies with the ODS account. Specifically, in a medium $P(q)$ condition, where $P(q)$ was set to 0.5, the proportion of *q* cards selected was the same as (experiment 1) or higher (experiment 3) than the proportion of *not-q* cards selected. However, according to ODS, at this value of $P(q)$



participants should have preferred the *not-q* card to the *q* card. This result was robust and therefore in need of explanation. Oaksford *et al.* suggested that the process of sequential sampling may explain this discrepancy. Here we explore some possible models for how sequential sampling should proceed in the RAST.

Two critical questions are raised by the introduction of sequential sampling. First, what is the stopping criterion? That is, when do people stop sampling, satisfied that the rule is true or false? Second, are there any trial-by-trial effects on the sequence of cards selected, that might discriminate between theories? Oaksford *et al.* (1997) were not primarily concerned with these questions and did not record the precise sequence of cards selected (they were concerned with the total number of each card selected and the initial card selected). We will look at three possible models of sequential sampling behaviour in the RAST: probability matching, Bayesian revision and epistemic utility models.

### Probability matching

The first possibility we look at is probability matching. Our initial assumption in Oaksford *et al.* (1997) was that the proportion of trials on which participants choose the *q* card or the *not-q* card in the RAST would reflect the underlying expected information gains of the cards computed at the beginning of the experiment. A rationale for this assumption follows from a straightforward adaptation of Myerson and Miezen's (1980) model of matching behaviour (Herrnstein, 1961) during foraging. Myerson and Miezen were concerned to model the finding that animals tend to distribute their time spent foraging at particular patches in proportion to the prevalence of food at those patches. We assumed that people tend to distribute the trials at which they select particular card types (*q* or *not-q*) in proportion to the expected information gains of those card types.

In order to show how this account could apply to selection behaviour in the RAST, we translate Myerson and Miezen's (1980) model into the language of the RAST and ODS. The first assumption is that for each available card, in the RAST the *q* and *not-q* card, people can compute a relative preference variable based on the expected information gains. Gallistel (1989) calls this 'relative patch affinity' and computing this quantity corresponds directly to Oaksford and Chater's (1994) computation of *scaled* expected information gains.[6] Given the two cards in the RAST we calculate:

$$SE[I_g(q)] = \frac{E[I_g(q)]}{E[I_g(q)] + E[I_g(\neg q)]}, \quad SE[I_g(\neg q)] = \frac{E[I_g(\neg q)]}{E[I_g(q)] + E[I_g(\neg q)]} \qquad 17.6$$

The second assumption is that probability matching is a consequence of modelling the distribution of time spent at different food patches as a random Poisson process where the rate parameter is proportional to patch affinity. Because the RAST involves a small number of discrete trials, and hence the Poisson process is not appropriate, we show how probability matching also arises in the RAST from modelling which card is chosen at each trial using the binomial distribution where the probability of selecting a card on a given trial is given by the $SE[I_g()]$ value. Given *n* trials (the actual number will vary from participant to participant) the mean number of trials on which a participant will select either card in the RAST is therefore:

$$\begin{aligned}\mu_q &= n.SE[I_g(q)]\\ \mu_{\neg q} &= n.SE[I_g(\neg q)]\\ &= n(1 - SE[I_g(q)])\\ &= n - \mu_q\end{aligned} \qquad 7.7$$

It follows from these equations that the ratio of the expected number of $q$ card and *not-q* card selections is equivalent to the ratio of the expected information gains, that is:

$$\frac{\mu_q}{\mu_{\neg q}} = \frac{SE[I_g(q)]}{SE[I_g(\neg q)]} = \frac{E[I_g(q)]}{E[I_g(\neg q)]} \qquad 17.8$$

Such an account makes a variety of predictions about the structure of sequential samples in the RAST that can be empirically tested. For example, predictions can be made about the probability of selecting a card on any trial, given any initial set of $SE[I_g()]$ values, and about the likelihood of particular sequences containing different numbers of $q$ and *not-q* card selections. One prediction, well confirmed in the data, is that people often sample from both stacks rather than sticking with the stack ODS initially recommends. This is understandable if the trial sequence is generated by a random process.

There are some problems with this account, however. First, this model provides no stopping criterion. In the RAST participants are only allowed to stop sampling when they have selected all the cards from one stack. However, they are not told that this is the case. Nonetheless, exhausting a stack is a natural landmark in the course of the experiment that might be expected to prompt participants to suggest that they terminate the experiment. This is not a rational response to the sequence of evidence that they have just experienced. So one problem with this proposal is that it removes any rational content from the decision to stop sampling, i.e. there is no *rational* stopping rule. Of course, this conceptual point does not invalidate the ODS model, especially if the predictions for sequential samples based on initial $SE[I_g()]$ values proved to be correct. However, the lack of a rational stopping rule does not sit very well with the claim that people's data selection behaviour can be regarded as rational (Oaksford and Chater, 1994).

Second, in Oaksford *et al.* (1997) the values of $P(q)$ chosen for the low and medium $P(q)$ conditions produced roughly symmetrical $SE[I_g()]$ values for these conditions. In the low condition, calculated as above, $SE[Ig(q)]=0.799$ and $SE[I_g(not\text{-}q)]=0.201$; and in the medium condition, $SE[I_g(q)]=0.282$ and $SE[I_g(not\text{-}q)]=0.718$ (calculated with $P(p)=0.1$, $e=0.1$ and $P(M_I)=0.5$). In the low condition this means that $q$ cards should be preferred to *not-q* cards in the ratio of 4:1 (approximately) whereas in the medium condition the *not-q* card should be preferred to the $q$ card in the ratio of 2.5:1 (approximately; for completeness in the high condition, the *not-q* card should be preferred to the $q$ card in the ratio of 31:1, approximately). However, a robust finding in Oaksford *et al.*'s experiments was that in the medium $P(q)$ condition, participants chose the same (experiment 1) or more (experiment 3) $q$ than *not-q* cards. The current probability matching model seems unable to explain this result.



## Bayesian revision

As Oaksford and Chater (1996; see also Oaksford *et al.*, 1997) commented, the ODS model only becomes a model of hypothesis testing in the context of sequential sampling (Fedorov, 1972; see also, Jessop, 1996; Laming, 1996). During sequential sampling participants can revise their priors, re-compute $SE[I_g()]$s to determine the next card to select and so on until the odds on either the dependence model or the independence model being true are so high that no more evidence need be collected, that is, either $P(M_I) \approx 0$, or $P(M_I) \approx 1$ (see, Fedorov, 1972). As Jessop (1996) outlines, the decision process can be described more formally using Bayesian decision trees. He also shows that when applied to a sequential version of the standard four card selection task, where each card is turned as it is selected, ODS predicts the empirically observed sequence of card selections (this relies on what Jessop calls a 'sequential rarity' assumption).

The immediate advantage of this approach is that there is a rational stopping criterion, i.e. stop when either $P(M_I) \approx 0$, $P(M_I) \approx 1$. However, there are some immediate problems in applying this natural extension of the ODS model to the RAST. First, in the standard RAST procedure, participants are not allowed to stop sampling until one stack is exhausted. Therefore, before this model can be tested properly experiments need to be conducted where people are allowed to stop sampling naturally, i.e. when they want to. Second, the earliest research on Bayesian inference shows that people are 'conservative', insofar as they revise their beliefs more slowly than Bayes's theorem dictates that they should (Edwards, 1968). Moreover, there is likely to be individual variation in the degree of conservativeness. We would therefore need to include a learning rate parameter to model actual human performance. Third, as Oaksford and Chater (1994) observed, the order over $SE[I_g()]$ values is insensitive to variation in $P(M_I)$. On the assumption that the experimental set up in the RAST fixes $P(p)$ and $P(q)$ at the outset, all that varies on a trial by trial basis is $P(M_I)$. Consequently, the order of $SE[I_g()]$ values and hence the card that ODS recommends selecting will not change over trials. For example, in the low condition ODS will recommend selecting the $q$ card at every trial until $P(M_I) \approx 0$ (in the RAST none of the cards are falsifying instances which is another factor that needs to be varied in these experiments). As we noted above, one advantage of the probability matching model was that occasional response alternations, which are observed in the data, are predicted. However, a straightforward application of Bayesian revision with ODS cannot predict these response alternations. Finally, a second consequence of the ODS model's insensitivity to variation in $P(M_I)$ is that it will not be able to capture Oaksford *et al.*'s results for the medium condition. In this condition, this model recommends selecting the *not-q* card at every trial until $P(M_I) \approx 0$. Consequently, it could not predict *any* $q$ card selections, let alone that they can be in the majority.

Despite these apparent limitations there is an assumption in this application of the ODS model that may well be violated in the RAST. We have assumed that only prior beliefs, $P(M_I)$, in the two hypothesis (models) are revised trial by trial—we assume that $P(p)$ and $P(q)$ are fixed at the beginning of the experiment and that they remain fixed throughout. In Oaksford *et al.* (1997) participants are told the values of $P(p)$ and $P(q)$ in the form of frequency statements and they also have stacks of cards in front of them that concretely reflect these frequencies. Although this may successfully encourage participants to utilize these probabilities, in their initial assessment of which card to select, it may not, and perhaps should not, prevent participants from actively updating these probabilities when they begin to sample the actual cards. Oaksford *et al.* (1997) speculated that subjects may revise $P(p)$ and $P(q)$ trial by trial in the RAST and that this may explain the apparent failure of the ODS model to predict the results for the medium $P(q)$ condition. Here we show that this is indeed the case.

How should participants update their beliefs about $P(M_I)$, $P(p)$ and $P(q)$? We make two assumptions. First, as we mentioned above, people are conservative in revising their beliefs. Consequently, we assume that participants revise their degree of belief in a hypothesis by only a half of what Bayes's theorem would recommend. So if Bayes recommended that your degree of belief in a hypothesis should be revised from 0.5 to 0.3, i.e. it should be decreased by 0.2, we assume that you only revise your degree of belief by half this amount, i.e. from 0.5 to 0.4. More formally, on trial $n$ your conservative degree of belief in the independence model, $\text{Cons}P(M_I)_n$, is:[7]

$$\text{Cons}P(M_I)_n = P(M_I)_n + \frac{1}{2}(\text{Cons}P(M_I)_{n-1} - P(M_I)_n) \qquad 17.9$$

Second, we assume that if people revise their degrees of belief about $P(p)$ and $P(q)$, this is because they lack confidence in the values they have been given. We embody this lack of confidence by assuming that participants regard the values they are given for $P(p)$ and $P(q)$ as being based on a small sample size. As we argued above in our discussion of Pollard and Evans (1983), this will influence the magnitude of the effects of sequential sampling on participants estimates of $P(p)$ and $P(q)$. In the simulations we report we assume that participants base the initial values of $P(p)$ and $P(q)$ on a sample of six cards.

Table 17.3 shows the predicted sequence of card selections for the low $P(q)$ condition in Oaksford *et al.* (1997). The first row of the table shows the initial parameter values and the consequent $SE[I_g()]$s upon which the decision about which card to select on the first trial is based. The parameters are updated on each trial as we described in the last paragraph. We show $P(p)$ and $P(q)$ as fractions where the denominator always reflects the sample size on which that value is assumed to be based. Whenever a $q$ card is selected it is also a $p$ card, so the numerators of both $P(p)$ and $P(q)$ are incremented by 1 as well as both their denominators. In contrast, whenever a *not-q* card is selected it is also a *not-p* card, therefore only the denominators of both $P(p)$ and $P(q)$ are incremented by 1. For these simulations $P(p) = P(q)$. This is not a critical assumption and seems reasonable given the low



sample size on which participants assume the estimates of $P(p)$ and $P(q)$ are based. Which card is selected on trial $n$ is determined by which card has the highest $SE[I_g()]$ value on trial $n-1$. As a stopping rule we have used a 99% confidence level in the dependence model. Consequently, we have stopped sampling when $\text{Cons}P(M_I) < 0.01$. Observe that when $P(p)$ and $P(q)$ are also updated, response alternations away from the card initially recommended are predicted by the ODS model (trials 3, 4, 6, 8, 11). Although for the low $P(q)$ condition $q$ card selections predominate (10 of 15 trials) the model also predicts that a *not-q* card will be selected on five trials. These predictions are consistent with Oaksford *et al.*'s finding that in the low $P(q)$ condition participants select more $q$ cards than *not-q* cards.

**Table 17.3 The sequence of card selections predicted by the ODS model for the low *P(q)* condition in Oaksford *et al.* (1997) when *P(p)* and *P(q)* are updated trial-by-trial in the RAST**

| Trial no. | Card selected | $P(p)$ | $P(q)$ | $P(M_I)$ | $\text{Cons}P(M_I)$ | $SE[I_g(q)]$ | $SE[I_g(\neg q)]$ |
|---|---|---|---|---|---|---|---|
| | | 1/6 | 1/6 | 0.5 | 0.5 | 0.720 | 0.280 |
| 1 | *q* | 2/7 | 2/7 | 0.196 | 0.348 | 0.515 | 0.485 |
| 2 | *q* | 3/8 | 3/8 | 0.166 | 0.257 | 0.397 | 0.603 |
| 3 | *not-q* | 3/9 | 3/9 | 0.179 | 0.218 | 0.481 | 0.519 |
| 4 | *not-q* | 3/10 | 3/10 | 0.157 | 0.188 | 0.550 | 0.450 |
| 5 | *q* | 4/11 | 4/11 | 0.075 | 0.131 | 0.483 | 0.517 |
| 6 | *not-q* | 4/12 | 4/12 | 0.088 | 0.110 | 0.543 | 0.457 |
| 7 | *q* | 5/13 | 5/13 | 0.044 | 0.077 | 0.498 | 0.502 |
| 8 | *not-q* | 5/14 | 5/14 | 0.049 | 0.063 | 0.551 | 0.449 |
| 9 | *q* | 6/15 | 6/15 | 0.025 | 0.044 | 0.520 | 0.480 |
| 10 | *q* | 7/16 | 7/16 | 0.019 | 0.031 | 0.497 | 0.503 |
| 11 | *not-q* | 7/17 | 7/17 | 0.018 | 0.025 | 0.543 | 0.457 |
| 12 | *q* | 8/18 | 8/18 | 0.011 | 0.018 | 0.526 | 0.474 |
| 13 | *q* | 9/19 | 9/19 | 0.008 | 0.013 | 0.513 | 0.487 |
| 14 | *q* | 10/20 | 10/20 | 0.006 | 0.010 | 0.504 | 0.496 |
| 15 | *q* | 11/21 | 11/21 | 0.005 | 0.008 | | |

On trial $n$, $\text{Cons}P(M_I)_n = P(M_I)_n + \frac{1}{2}(\text{Cons}P(M_I)_{n-1} - P(M_I)_n)$; $e = 0.01$

Table 17.4 shows the predicted sequence of card selections for the medium $P(q)$ condition in Oaksford *et al.* (1997). Oaksford *et al.* found that participants selected the same proportion (experiment 1) or more (experiment 3) $q$ cards than *not-q* cards in the medium $P(q)$ condition. This pattern of results appeared contrary to the predictions of the ODS model because if $P(p)$ and $P(q)$ remain fixed then ODS can only recommend selecting *not-q* cards in this condition. However, as Table 17.4 reveals, when $P(p)$ and $P(q)$ are updated trial-by-trial, before the stopping criterion is reached (trial 16) the ODS model will recommend selecting equal numbers of $q$ cards and *not-q* cards. If one more trial is conducted then in this condition, participants will indeed be predicted to select more $q$ cards than *not-q* cards. The reason for this trial-by-trial behaviour is as follows. Taking the low $P(q)$ condition, as participants select $q$ cards their estimates of $P(p)$ and $P(q)$ will go up because these are all $p$, $q$ instances. If the stopping criterion is not reached before rarity is violated the model is bound to predict that participants should select some *not-q* cards. Taking the medium and high $P(q)$ conditions, as participants select *not-q* cards their estimates of $P(p)$ and $P(q)$ will go down because these are all *not-p*, *not-q* instances. A further factor is that although in the normal range $P(M_I)$ has little effect on the ordering of $SE[I_g()]$ s, as $P(M_I) \rightarrow 0$, it would appear that rarity is relaxed, i.e. higher values of $P(q)$ can still lead to $SE[I_g(q)] > SE[I_g(not-q)]$. This factor is responsible for the prediction of long sequences of $q$ card selections as $P(M_I)$ approaches the stopping criterion. An interesting example is when $P(p) = P(q) = 8/18$. In the medium $P(q)$ condition (see Table 17.4) this occurs when $\text{Cons}P(M_I) = 0.028$ and consequently $SE[I_g(q)] = 0.496$ and $SE[I_g(not-q)] = 0.504$, so the ODS model recommends selecting a *not-q* card. In contrast, in the high $P(q)$ condition (see Table 17.5) $P(p) = P(q) = 8/18$ when $\text{Cons}P(M_I) = 0.014$ and consequently $SE[I_g(q)] = 0.542$ and $SE[I_g(not-q)] = 0.458$, so the ODS model recommends selecting a $q$ card.

**Table 17.4 The sequence of card selections predicted by the ODS model for the medium *P(q)* condition in Oaksford *et al.* (1997) when *P(p)* and *P(q)* are updated trial-by-trial in the RAST**

| Trial no. | Card selected | $P(p)$ | $P(q)$ | $P(M_I)$ | $\text{Cons}P(M_I)$ | $SE[I_g(q)]$ | $SE[I_g(\neg q)]$ |
|---|---|---|---|---|---|---|---|
| | | 3/6 | 3/6 | 0.5 | 0.5 | 0.152 | 0.848 |
| 1 | *not-q* | 3/7 | 3/7 | 0.337 | 0.419 | 0.257 | 0.743 |
| 2 | *not-q* | 3/8 | 3/8 | 0.294 | 0.356 | 0.358 | 0.642 |
| 3 | *not-q* | 3/9 | 3/9 | 0.258 | 0.307 | 0.444 | 0.556 |
| 4 | *not-q* | 3/10 | 3/10 | 0.229 | 0.268 | 0.517 | 0.483 |
| 5 | *q* | 4/11 | 4/11 | 0.120 | 0.194 | 0.445 | 0.555 |
| 6 | *not-q* | 4/12 | 4/12 | 0.134 | 0.164 | 0.509 | 0.491 |
| 7 | *q* | 5/13 | 5/13 | 0.070 | 0.117 | 0.462 | 0.538 |
| 8 | *not-q* | 5/14 | 5/14 | 0.076 | 0.097 | 0.518 | 0.482 |
| 9 | *q* | 6/15 | 6/15 | 0.040 | 0.068 | 0.487 | 0.513 |
| 10 | *not-q* | 6/16 | 6/16 | 0.042 | 0.055 | 0.536 | 0.464 |
| 11 | *q* | 7/17 | 7/17 | 0.022 | 0.039 | 0.513 | 0.487 |
| 12 | *q* | 8/18 | 8/18 | 0.017 | 0.028 | 0.496 | 0.504 |
| 13 | *not-q* | 8/19 | 8/19 | 0.016 | 0.022 | 0.541 | 0.459 |
| 14 | *q* | 9/20 | 9/20 | 0.010 | 0.016 | 0.527 | 0.473 |
| 15 | *q* | 10/21 | 10/21 | 0.007 | 0.012 | 0.516 | 0.484 |
| 16 | *q* | 11/22 | 11/22 | 0.006 | 0.009 | 0.507 | 0.493 |
| 17 | *q* | 12/23 | 12/23 | 0.005 | 0.007 | | |

On trial $n$, $\text{Cons}P(M_I)_n = P(M_I)_n + \frac{1}{2}(\text{Cons}P(M_I)_{n-1} - P(M_I)_n)$; $e = 0.01$

Table 17.5 shows the predicted sequence of card selections for the high $P(q)$ condition in Oaksford *et al.* (1997). As Oaksford *et al.* (1997) found, there are more *not-q* than $q$ cards predicted in this condition.



**Table 17.5 The sequence of card selections predicted by the ODS model for the high *P(q)* condition in Oaksford *et al.* (1997) when *P(p)* and *P(q)* are updated trial-by-trial in the RAST**

| Trial no. | Card selected | $P(p)$ | $P(q)$ | $P(M_I)$ | $\text{Cons}P(M_I)$ | $SE[I_g(q)]$ | $SE[I_g(\neg q)]$ |
|---|---|---|---|---|---|---|---|
| | | 5/6 | 5/6 | 0.5 | 0.5 | 0.006 | 0.994 |
| 1 | *not-q* | 5/7 | 5/7 | 0.164 | 0.332 | 0.042 | 0.956 |
| 2 | *not-q* | 5/8 | 5/8 | 0.129 | 0.231 | 0.115 | 0.885 |
| 3 | *not-q* | 5/9 | 5/9 | 0.103 | 0.167 | 0.208 | 0.792 |
| 4 | *not-q* | 5/10 | 5/10 | 0.083 | 0.125 | 0.301 | 0.699 |
| 5 | *not-q* | 5/11 | 5/11 | 0.067 | 0.096 | 0.384 | 0.616 |
| 6 | *not-q* | 5/12 | 5/12 | 0.055 | 0.075 | 0.457 | 0.543 |
| 7 | *not-q* | 5/13 | 5/13 | 0.045 | 0.060 | 0.517 | 0.483 |
| 8 | *q* | 6/14 | 6/14 | 0.025 | 0.043 | 0.485 | 0.515 |
| 9 | *not-q* | 6/15 | 6/15 | 0.025 | 0.034 | 0.538 | 0.462 |
| 10 | *q* | 7/16 | 7/16 | 0.014 | 0.024 | 0.515 | 0.485 |
| 11 | *q* | 8/17 | 8/17 | 0.011 | 0.018 | 0.496 | 0.504 |
| 12 | *not-q* | 8/18 | 8/18 | 0.010 | 0.014 | 0.542 | 0.458 |
| 13 | *q* | 9/19 | 9/19 | 0.006 | 0.010 | 0.530 | 0.470 |
| 14 | *q* | 10/20 | 10/20 | 0.005 | 0.008 | | |

On trial $n$, $\text{Cons}P(M_I)_n = P(M_I)_n + \frac{1}{2}(\text{Cons}P(M_I)_{n-1} - P(M_I)_n)$; $e = 0.01$

In summary, a revised ODS model that allows that during sequential sampling participants are learning both about which hypothesis is true and about the distribution of $p$ and $q$ cards seems able to model the medium $P(q)$ condition in the RAST. ODS also captures the basic findings for the low and medium $P(q)$ conditions as well as predicting the response alternations Oaksford *et al.* (1997) noticed in participants card selections. Importantly, ODS also embodies a rational stopping rule which means that unlike the probability matching model, participants can be interpreted as responding rationally to the sequence of evidence that they experience.

Finally, if we allow revision of $P(p)$ and $P(q)$ and a low sample size in the RAST, then how can we justify our argument that people neither assume a low sample size nor revise their beliefs about $P(p)$ and $P(q)$ in the learning phase of Pollard and Evans' (1983) experiments? There is no inconsistency because, as we have argued, in Pollard and Evans' learning phase participants are not trying to *infer* whether the rule is true or false, they are trying to *predict* what is on the other side of the cards. It is only after they have performed this task, which concentrates attention on $P(q|p)$ and not on $P(p)$ and $P(q)$, that they are unexpectedly confronted with an inference task. In contrast, in the RAST learning and inference are integrated. Consequently, participants are aware as they sample the cards that they are looking for information to determine the truth or falsity of the rule. Although $P(p)$ and $P(q)$ are relevant to this task, these probabilities are not relevant to the prediction task in Pollard and Evans learning phase. As we noted above, Evans and Over could argue that people may none the less learn $P(p)$ and $P(q)$ implicitly. However, this seems inconsistent with Evans and Over (1996*b*) own view that relevance judgements are implicit cognitive processes. If this is so then presumably if $P(p)$ and $P(q)$ are irrelevant to the explicit task participants confront, then information about these probabilities will also be judged irrelevant at the implicit level. Therefore, it would seem inconsistent of Evans and Over to argue that participants learn $P(p)$ and $P(q)$ implicitly.

## Epistemic utilities

We have shown that on a disinterested model of inquiry (see Chater *et al.*, chapter 20) revising the base probabilities in line with experience leads to predictions that are in line with actual performance in the RAST. In Oaksford *et al.* (1997) we suggested that these results may require the introduction of explicit utilities with respect to different evidence types and that it was these utilities that were changing as a result of sequential sampling. Our rationale for such an account, which we suggested may be consistent with Evan and Over's notion of 'epistemic utility', was as follows.

Oaksford *et al.* (1997) argued that according to ODS, in both the medium and high $P(q)$ conditions participants attention is initially directed towards the *not-q* card stack. All the cards in this stack are *not-p*, *not-q* cards, i.e. although examining these cards will confirm the rule (i.e. $P(M_D)$ rises), they provide participants with no positive instances of the rule. This is like arguing that although you have never seen a raven you are confident that all ravens are black because all the non-black things you have examined were also not ravens. This is the raven's paradox of classical confirmation theory (see Goodman, 1954). As discussed by Oaksford and Chater (1996), the raven's paradox is eliminated by Oaksford and Chater's (1994) default rarity assumption (see also Horwich, 1982; Howsen and Urbach, 1989). In the low $P(q)$ condition participants attention is focused on the $q$ cards which are all $p$, $q$ instances, i.e. participants are looking for blue triangles (black ravens). However, when rarity is not in force, i.e. in the medium and high $P(q)$ conditions, participants may want to see some blue triangles as well as red circles to confirm that the rule is not just vacuously true. Consequently, although participants will initially be guided to the *not-q* stack in the medium and high $P(q)$ conditions by ODS they may at some time decide to search for positive instances of the rule.

However, participants are more likely to look for positive instances in the medium condition than in the high condition. $SE[I_g(not-q)]$ is much higher in the high $P(q)$ condition than in the medium $P(q)$ condition. Consequently, it will take more data to overcome the attentional focus on the *not-q* card in the high $P(q)$ condition than in the medium $P(q)$ condition. However, in the RAST the number of data points are limited to a fixed number (in Oaksford *et al.*'s experiments 3 and 4 it is 10). Therefore the point where attention is refocused is more likely to be met in the medium $P(q)$ condition than in the high $P(q)$ condition.

Conversely, in the low $P(q)$ condition after initially focusing on the $q$ cards participants attention may be diverted towards the *not-q* cards. From our Bayesian perspective both $p$, $q$ and $p$, *not-q* cases may provide important classes of evidence. After examining a certain number of $p$, $q$ cases attention may shift to looking for $p$, *not-q* cases. This could explain the differences between Oaksford *et al.*'s experiments 1 and 3 in the low $P(q)$ condition. In experiment 1 in the low $P(q)$ condition there were five $q$ cards, whereas in experiment 3 there were 10. It is therefore more likely that participants will exhaust this stack while their attention is still focused on $p$, $q$ instances in experiment 1 than in experiment 3. Because these instances are now



exhausted participants declare the rule true and terminate the experiment. Consequently, a far lower proportion of *not-q* cards should be selected in the low $P(q)$ condition in experiment 1 than in experiment 3 which is consistent with the observed differences between these experiments.

Oaksford *et al.* (1997) suggested that current epistemic utility theory does not address how these changes of attention come about. However, one proposal would be to allow different utilities of evidence in a similar manner to Oaksford and Chater's (1994, pp. 621–5) maximum expected utility model of the deontic selection task. Modelling changes of attention due to sequential sampling would involve making the utilities for each evidence type some decreasing function of the number of instances of that evidence type observed. If all evidence were equally weighted initially, such a model would make the same predictions as ODS for initial card selections. However, after accumulating data, other evidence types will come to have higher utility, as the utility of the evidence type initially recommended by ODS falls off.

In retrospect and given the obvious success of the ODS model in explaining sequential sampling in the RAST, we are no longer convinced that such an account is either necessary or indeed conceptually coherent. Because Evans and Over (1996*a*, *b*) used the term epistemic *utility* we have assumed that their intention is to introduce an explicit utility function to explain people's data selections in tasks like the selection task. Taking this at face value led to the suggestion that people may be revising their utilities for different evidence classes instead of the probabilities $P(q)$ and $P(p)$. However, we no longer believe that this view is conceptually coherent as we now argue.

The proposal is to introduce utilities with respect to evidence classes, i.e. when $p$, $q$; p, *not-q*; *not-p*, $q$; and *not-p*, *not-q* instances are taken to bear on the truth or falsity of the rule, and that the utilities of different instances can vary as a function of encountering these instances. There seem to be two problems with this proposal. First, in a standard utility function that relates the outcomes of actions to utilities, the utility function defines a person's goals. So, for example, if your goal is to detect violators of the rule *if you are drinking beer you must be over 18 years of age*, then you will attach a high utility to $p$, *not-q* instances. In contrast, in disinterested inquiry the goal is to determine whether a rule is true or false. As we have seen, which instances are most relevant to this goal may vary depending on the probabilities $P(M_I)$, $P(q)$ and $P(p)$. However, fulfilling this goal seems totally independent of whatever utilities you assign to these evidence classes. So for example, if you are a Bayesian you will be interested in whatever evidence class affords the most discrimination between the hypothesis under test (in ODS, the dependence model) and the foil (in ODS, the independence model). And in a selection task you will be interested in looking at the cards that are most likely to reveal these evidence classes. These decisions rely only on the probabilities not on any assignment of utilities.

The second problem relates to allowing these utilities to change as a function of the number of trials carried out. We can illustrate the problem more clearly by showing the consequences of allowing this to happen in a context where assigning utilities is appropriate. Let us suppose that you are enforcing the under age drinking rule mentioned in the last paragraph. You therefore want to catch violators of the rule, i.e. people who are drinking beer but who are under 18 ($p$, *not-q* cases), and so you attach a high utility to these cases. However, on the proposal made above, each violator you detect will lead to a decrement in the utility associated with violators. Consequently, the more successful you are at achieving your goal of detecting violators, the more likely you are to spontaneously change your goals so that you stop looking for violators. Although perhaps reflecting something of the perversity of human nature, this behaviour is unlikely to impress your supervisor.

We have now argued that introducing utilities with respect to evidence classes and allowing these to change over trials seems to lead to conceptual problems. How then can we make sense of Evans and Over's (1996*a*, p. 363) assertion that an overarching concept of 'epistemic utility . . . underlie[s] all human hypothesis testing and reasoning'. One way is to look at the actual measures that they have proposed as measures of epistemic utility. In, for example, Evans and Over (1996*a*, *b*) and Over and Jessop (see Chapter 18) our own information gain measure and the log-likelihood ratio are labelled as attempts to define a notion of epistemic utility. There are two points to make. First, these measures are not in competition as Evans and Over (1996*a*, *b*) claim they are. As Over and Jessop (see Chapter 18) point out, there may be more than one foil hypothesis, i.e. a hypothesis tester may be trying to discriminate between more than just two hypotheses. However, the log-likelihood ratio is restricted to just two hypotheses whereas the information gain measure used in the ODS model generalizes to $n$ hypotheses (see also Jessop, 1996). Consequently, the log-likelihood ratio is not a serious theoretical competitor. Second, neither of these measures makes any reference to an explicit utility function, they are both based purely on probabilities. Therefore, although Evans and Over (1996*a*, *b*) may choose to refer to these measures as 'epistemic *utilities*', this label has no descriptive or theoretical content.

Finally, the reason for introducing utilities was not only to account for the medium $P(q)$ condition in Oaksford *et al.* (1997) but also to fulfil the reasonable requirement that one should not believe rules just because there are no falsifying cases, some positive instances are also required. Of course in the ODS model this problem is resolved by the rarity assumption ($P(p)$ and $P(q)$ are low) which means that for hypotheses made up of predicates describing real world properties people will always begin sampling positive instances. Notice also that in sequential sampling, even if $P(p)$ or $P(q)$ are high, a consequence of conservative revision is that people have enough time before $P(M_I) \approx 0$ or 1, eventually to sample some positive instances. Of course these results for the high and medium $P(q)$ conditions rely on the structure of the particular samples used in the RAST. The obvious next step is to use differently structured samples to see what different models predict and how the empirical results come out.

In summary, proposing a changing utility function to model behaviour in the RAST is (i) unnecessary to model the data, as we showed in the previous section, and is (ii) conceptually problematic along with Evans and Over's (1996*a*, *b*) proposal that people based their judgements on a measure they refer to as 'epistemic utility'.



## Conclusions

In this chapter we have argued for a revised ODS model which allows exceptions and when applied to sequential sampling employs conservative Bayesian revision. The revised model straightforwardly meets the objections raised by Evans and Over (1996) that the original model could not explain data from Kirby (1994) or from Pollard and Evans (1983). Importantly, the revised model generalizes naturally to the context of sequential sampling, explaining the apparently aberrant results of Oaksford *et al.* (1997). These latter results bear on the important issue of the rationality of human learning (Shanks, 1995, also see Chapter 15). They appear to show that various trial-by-trial effects in learning about a rule may be susceptible to the same rational analysis as Oaksford and Chater (1994) applied to the selection task. As Oaksford and Chater (1994) observed, this leaves the exciting possibility of unified rational explanations of selection task results, causal reasoning and instrumental learning (see also, Glymour and Cheng, Chapter 14; Over and Jessop, Chapter 18).

# Notes

1. Finally, we clear up two remaining points. First, Evans and Over argue that we do not predict changes for the *not-p* card that Kirby (1994) observed. This is not true. We do predict some of the movements observed for this card. Evans and Over reach this conclusion because they do not *scale* their $E(I_g)$ values, as Oaksford and Chater's model specifies. Second, Evans and Over argue that we should distinguish benefits from costs in modelling Kirby's deontic experiment 4. However, such additional factors need only be introduced if there are empirical reasons to do so. Modelling Kirby's data did not require a distinction between costs avoided and benefits gained. This distinction may nonetheless be needed to model other data sets.

2. These are not the same as Evans and Over's calculated values of 0.298 (usually true) and 0.062 (usually false). The discrepancy seems due to unimportant differences in the fine detail of Evans and Over's calculations. The important point is that we agree with Evans and Over on the ball-park figures and hence their arguments stand. Our model is implemented in *Mathematica* (Wolfram, 1991) which makes exploring its properties particularly easy. Anyone wishing to obtain a copy of the model can do so by contacting Mike Oaksford (oaksford@cardiff.ac.uk).

3. A further issue, which we will not make to much of, is that it seems impossible for Evans and Over (1996*a*) to have used our model to calculate expected information gains. Apparently they did not use values of $P(p)$ and $P(q)$ directly in their calculation but rather used the likelihoods, i.e. the probabilities of the data given the hypothesis (Evans and Over, personal communication). In Pollard and Evans (1983) materials, the probability of $q$ given $p$ ($P(q|p)$, is 0.875, which Evans and Over apparently used to compute expected information gains. But if people know that $I(q|p) = 0.875$, then they already know that the dependence model cannot hold because in the original ODS dependence model $P(q|p)$, i.e. $P(q|p, M_D)$, is always 1. Consequently, we do not know how Evans and Over (1996*a*) could have calculated expected information gains using our original model. Nevertheless, using $P(p)$ and $P(q)$ in the calculations, as we did above, leads to a similar conclusion which we must still address.

4. Knowledge of conditional probabilities without knowledge of absolute probabilities is very common. For example, although you may know that the conditional probability of thunder given lightening is very near 1, you may have no idea how frequent either lightening or thunder are, other than, of course, that they are pretty rare.

5. Another approach would be to allow $P(M_I)$ to be a free parameter, and see what value of this parameter gives the best fit to the data. It would appear that much better fits than we report here are obtainable when $P(M_I)$ is low, i.e. when both the low exception ('believed true') and high exception ('believed false') rules are believed true. This clearly makes the strong prediction that people do indeed divorce their judgements of belief from the number of exceptions *per se*.

6. In Oaksford and Chater (1994, 1995) the scaling factor was the average expected information gain over the four cards. Moreover we added a fixed amount to all $E[I_g()]$s to reflect the presence of the *not-p* card. Although $E[I_g(not{-}p)] = 0$, its mere presence means that it might get selected and this should vary with the likelihood of the other cards being selected. This procedure is unnecessary in the standard RAST where only the $q$ and *not-q* cards are available.

7. More generally we can include a learning rate parameter, $\eta$, that may vary from participant to participant:

$$\mathrm{Cons}P(M_I)_n = P(M_I)_n + \eta\,(\mathrm{Cons}P(M_I)_{n-1} - P(M_I)_n)$$

## References


Almor, A. and Sloman, S. A. (1996). Is deontic reasoning special? *Psychological Review*, **103**, 374–80.

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

Anderson, J. R. (1994). *Rules of the mind*. Lawrence Erlbaum Associates. Hillsdale, NJ.

Bromberger, S. (1965). An approach to explanation. In *Studies in analytical philosophy*, (ed. R. J. Butler) pp. 72–105. Basil Blackwell, Oxford.

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

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

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

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.

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

Edwards, W. (1968). Conservatism in human information processing. In *Formal representation of human judgment* (ed. B. Kleinmuntz) pp. 17–52. Wiley, New York.

Evans, J. St.B. T. and Lynch, J. S. (1973). Matching bias in the selection task. *British Journal of Psychology*, **64**, 391–7.

Evans, J. St.B. T. and Over, D. E. (1996*a*). Rationality in the selection task: epistemic utility versus uncertainty reduction. *Psychological Review*, **103**, 356–63.

Evans, J. St.B. T. and Over, D. E. (1996*b*). *Rationality and reasoning*. Psychology Press, Hove, East Sussex.

Fedorov, V. V. (1972). *Theory of optimal experiments*. Academic Press, London.

Gallistel, C. R. (1989). *The organisation of learning*. MIT Press, Cambridge, MA.

Gigerenzer, G. (1994). Why the distinction between single-event probabilities and frequencies is relevant for psychology (and vice versa). In *Subjective probability* (ed. G. Wright and P. Ayton) pp. 129–61. Wiley, New York.

Gigerenzer, G. and Hug, K. (1992). Domain-specific reasoning: social contracts, cheating, and perspective change. *Cognition*, **43**, 127–71.

Girotto, V. (1988). Pragmatic knowledge and deductive reasoning in children. *Giornale Italiano di Psicologia*, **15**, 287–314.

Girotto, V., Light, P., and Colbourn, C. J. (1988). Pragmatic schemas and conditional reasoning in children. *Quarterly Journal of Experimental Psychology*, **40A**, 469–82.

Girotto, V., Blaye, A., and Farioli, F. (1989). A reason to reason: pragmatic basis of children's search for counterexamples. *Cahiers de Psychologie Cognitive*, **9**, 297–321.

Goodman, N. (1954). *Fact, fiction and forecast*. Harvard University Press, Cambridge, MA.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour*, **24**, 107–16.

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





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

Jessop, A. (1996). A sequential analysis of the selection task. Unpublished manuscript, Business School, University of Durham.

Johnson-Laird, P. N. and Wason, P. C. (1970*a*). A theoretical analysis of insight into a reasoning task. *Cognitive Psychology*, **1**, 134–48.

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

Keynes, J. M. (1921). *A treatise on probability*. Macmillan, London.

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

Klayman, J. and Ha, Y. (1987). Confirmation, disconfirmation and information in hypothesis testing. *Psychological Review*, **94**, 211–28.

Laming, D. (1996). On the analysis of irrational data selection: a critique of Oaksford and Chater (1994). *Psychological Review*, **103**, 364–73.

Light, P., Blaye, A., Gilly, M., and Girotto, V. (1989). Pragmatic schemas and logical reasoning in 6- to 8-year old children. *Cognitive Development*, **4**, 49–64.

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 Ph.D. Thesis, California Institute of Technology, Pasadena, CA.

Manktelow, K. I., Sutherland, E. J., and Over, D. E. (1995). Probabilistic factors in deontic reasoning. *Thinking and Reasoning*, **1**, 201–20.

Mises, R. von (1939). *Probability, statistics and truth*. Allen and Unwin, London.

Myerson, J. and Miezen, F. M. (1980). The kinetics of choice: an operant systems analysis. *Psychological Review*, **87**, 160–74.

Oaksford, M. and Chater, N. (1991). Against logicist cognitive science. *Mind & Language*, **6**, 1–38.

Oaksford, M. and Chater, N. (1992). Bounded rationality in taking risks and drawing inferences. *Theory & Psychology*, **2**, 225–30.

Oaksford, M. and Chater, N. (1993). Reasoning theories and bounded rationality. In *Rationality*, (ed. K. I. Manktelow and D. E. Over), pp. 31–60. Routledge, London.

Oaksford, M. and Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, **101**, 608–31.

Oaksford, M. and Chater, N. (1995*a*). Theories of reasoning and the computational explanation of everyday inference. *Thinking and Reasoning*, **1**, 121–52.

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

Oaksford, M. and Chater, N. (1996). Rational explanation of the selection task. *Psychological Review*, **103**, 381–91.

Oaksford, M., Chater, N., Grainger, B., and Larkin, J. (1997). Optimal data selection in the reduced array selection task (RAST). *Journal of Experimental Psychology: Learning, Memory & Cognition*, **23**, 441–58.

Pollard, P. (1985). Nonindependence of selections on the Wason selection task. *Bulletin of the Psychonomic Society*, **23**, 317–20.

Pollard, P. and Evans, J. St. B. T. (1983). The effect of experimentally contrived experience on reasoning performance. *Psychological Research*, **45**, 287–301.

Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson, London.
Pylyshyn, Z. W. (1984). *Computation and cognition*. MIT Press, Cambridge, MA.
Ramsey, F. P. (1931). *The foundations of mathematics and other logical essays*. Routledge and Kegan Paul, London.
Shannon, C. E. and Weaver, W. (1949). *The mathematical theory of communication*. University of Illinois Press, Urbana, IL.
Sperber, D., Cara, F., and Girotto, V. (1995) Relevance theory explains the selection task. *Cognition*, **57**, 31–95.
Wason, P. C. (1966). Reasoning. In *New horizons in psychology*, (ed. B. Foss), pp. 135–51. Penguin, Harmondsworth, Middlesex.
Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, **20**, 273–81.
Wason, P. C. (1969). Regression in reasoning. *British Journal of Psychology*, **60**, 471–80.
Wason, P. C. and Green, D. W. (1984). Reasoning and mental representation. *Quarterly Journal of Experimental Psychology*, **36A**, 597–610.
Wiener, N. (1948). *Cybernetics*. Wiley, New York.
Wolfram, S. (1991). *Mathematica* [Computer program]. Addison-Wesley, Reading, MA.