

## Rational Explanation of the Selection Task

Mike Oaksford  
University of Warwick

Nick Chater  
University of Oxford

M. Oaksford and N. Chater (O&C; 1994) presented the first quantitative model of P. C. Wason's (1966, 1968) selection task in which performance is rational. J. St B T Evans and D. E. Over (1996) reply that O&C's account is normatively incorrect and cannot model K. N. Kirby's (1994b) or P. Pollard and J. St B T Evans's (1983) data. It is argued that an equivalent measure satisfies their normative concerns and that a modification of O&C's model accounts for their empirical concerns. D. Laming (1996) argues that O&C made unjustifiable psychological assumptions and that a "correct" Bayesian analysis agrees with logic. It is argued that O&C's model makes normative and psychological sense and that Laming's analysis is not Bayesian. A. Almor and S. A. Sloman (1996) argue that O&C cannot explain their data. It is argued that Almor and Sloman's data do not bear on O&C's model because they alter the nature of the task. It is concluded that O&C's model remains the most compelling and comprehensive account of the selection task.

Research on Wason's (1966, 1968) selection task questions human rationality because performance is not "logically correct." Recently, Oaksford and Chater (O&C; 1994) provided a *rational analysis* (Anderson, 1990, 1991) of the selection task that appeared to vindicate human rationality. O&C argued that the selection task is an inductive, rather than a deductive, reasoning task: Participants must assess the truth or falsity of a general rule from specific instances. In particular, participants face a problem of *optimal data selection* (Lindley, 1956): They must decide which of four cards ( $p$ ,  $not-p$ ,  $q$ , or  $not-q$ ) is likely to provide the most useful data to test a conditional rule, *if  $p$  then  $q$* . The "logical" solution is to select the  $p$  and the  $not-q$  cards. O&C argued that this solution presupposes falsificationism (Popper, 1959), which argues that only data that can disconfirm, not confirm, hypotheses are of interest. In contrast, O&C's rational analysis uses a *Bayesian* approach to inductive confirmation (Earman, 1992; Horwich, 1982; Howson & Urbach, 1993) and, specifically, to optimal data selection (Lindley, 1956; MacKay, 1992). According to this approach, people base card selections on *expected information gain*, or  $E(I_g)$ .

O&C's account contrasts with previous accounts of the selection task in three ways (a partial exception being Kirby, 1994b). First, it provides an explicit alternative to the logical view of rational behavior in the task. Second, O&C provide a formal model so that predictions are derived mathematically rather than by appeal to intuition. Third, O&C provide quantitative fits to a large range of empirical data. In this article, we reply to commentaries on our model by Evans and Over (E&O), Laming, and Almor and Sloman (A&S).

### Evans and Over

E&O argue that our model is inadequate for two reasons. First, it is not normatively justified because the  $E(I_g)$  measure has some counterintuitive properties. Second, it is descriptively inadequate to Kirby's (1994b) and Pollard and Evans's (1983) results. We respond to these points in turn and then consider E&O's residual arguments.

### *Is Oaksford and Chater's Theory Normatively Justified?*

E&O suggest that "even as a normative proposal, [O&C's] approach has serious problems" (1996, p. 358). For example, if you begin with  $P(H) = .25$  and, after turning the card, change to  $P(H) = .75$ , your amount of uncertainty is the same as before, and hence information gain ( $I_g$ ) is 0. Thus,  $I_g$  seems an inappropriate information measure because you have learned something from turning the card. E&O also note that turning a card can lead to less certainty about whether the rule is true; again, turning the card intuitively provides information, although  $I_g$  is negative. This is a minor matter because expected information gain,  $E(I_g)$ , is always nonnegative (see the Appendix), and O&C's calculations all concerned expected values. Nonetheless, E&O do point out an unattractive feature of O&C's  $E(I_g)$  measure.

We can take E&O's insight into account by using an alternative to the  $E(I_g)$  measure. Intuitively, E&O's point is that a card's informativeness depends on the magnitude of the difference between one's degree of belief in the rule before and after turning the card. If one must revise one's belief from certainty that the rule is true or false to less certainty, then the card has still been informative.

We can formalize this suggestion by comparing the probability distributions representing the new and old degrees of belief. These probability distributions contain just two values:  $P(M_D)$ , the probability that the conditional is true, and  $P(M_I)$ , the probability that the antecedent and consequent are independent. To measure the difference between the new and the old distributions, we use the standard information-theoretic mea-

---

Correspondence concerning this article should be addressed to Mike Oaksford, Department of Psychology, University of Warwick, Coventry CV4 7AL England, or to Nick Chater, Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD England. Electronic mail may be sent via Internet to [psyad@csv.warwick.ac.uk](mailto:psyad@csv.warwick.ac.uk) or [nick@psy.ox.ac.uk](mailto:nick@psy.ox.ac.uk).

sure: the Kullback–Liebler distance,  $D$ , between new and old probability distributions (Kullback & Liebler, 1951; see the Appendix).<sup>1</sup>  $D$  is always nonnegative and is zero only when the two distributions are identical (i.e., turning the card has not led to any revision of previous beliefs). Specifically,  $D$  is positive in the cases E&O mention, where  $I_g$  is negative or zero.

Taking on board E&O's point, then, we can switch from  $I_g$  to  $D$  to assess the informativeness of a card. Remarkably, this requires no change whatever to our original analysis. Although the new and old measures are different, their expected value is always the same (we prove this in the Appendix). Because we base all of our predictions on expected information gain, this means that we can switch to expected Kullback–Liebler distance with no theoretical revision whatsoever (apart from expository differences). In somewhat different forms, this result is well-known in the information-theoretic literature (e.g., see Cover & Thomas, 1991; MacKay, 1992).

### *Is Oaksford and Chater's Theory Descriptively Adequate?*

O&C's model always assumes that participants interpret the four cards in the selection task as a sample from a larger population of cards, over which the conditional rule is defined. E&O observe that this interpretation does not seem to apply to Kirby's or Pollard and Evans's experiments, in which participants know there are exceptions to the rule. Consequently, they argue that the rule can apply only to the four cards. On this assumption, they then generate predictions from O&C's model that seem to conflict with the data. We make three points here. First, the occurrence of exceptions does not entail that an exceptionless rule must apply just to the four cards. For example, in Kirby's experiments the rule can apply to the cards the machine subsequently produces.<sup>2</sup> Second, we have argued elsewhere that everyday conditional rules are not interpreted as exceptionless (Chater, 1993; Chater & Oaksford, 1990, 1993; Oaksford, 1993; Oaksford & Chater, 1991, 1992, 1993, 1995b). It is straightforward to produce a more realistic model by incorporating an exception parameter. This is not the place to modify our existing model; when this is done, however, the model's predictions seem to be unchanged, and the fits appear comparable to that of the original model.<sup>3</sup> Third, to derive their predictions for Pollard and Evans's experiment, E&O assume that participants estimate  $P(p)$  and  $P(q)$  from the data. However, Pollard and Evans's learning phase involved a prediction task that focuses attention on  $P(q|p)$ .  $P(q|p)$  does not determine  $P(p)$  and  $P(q)$ . Consequently, it is reasonable to argue that participants adopt default rarity values for  $P(p)$  and  $P(q)$  in computing information gain. This would bring O&C's predictions into line with Pollard and Evans's data. Therefore, Pollard and Evans's results may not be inconsistent with O&C's model.

### *Residual Arguments*

E&O make three residual points, which we briefly address. First, they distinguish between "*Rationality*<sub>1</sub>: reasoning or acting in such a way as to achieve one's goals [and] *Rationality*<sub>2</sub>: reasoning or acting in conformity with a relevant normative system such as formal logic or probability theory" (1996, p. 357).

They then note that "it may . . . appear that [O&C] have provided a rational<sub>1</sub> . . . account of the problem. On reflection we fear that this is not so. They have in fact substituted one rational<sub>2</sub> analysis for another" (1996, p. 357). This suggests that rational<sub>2</sub> analysis is not a good thing. However, because E&O give no argument supporting this claim, there is nothing to which we can reply. Also, E&O give no argument why our model is not a rational<sub>1</sub> account. In O&C's model, a person's goal is to reduce his or her uncertainty in indicative selection tasks and to maximize expected utility in deontic tasks. Consequently, O&C's account appears to be a rational<sub>1</sub> theory by E&O's own definition.<sup>4</sup>

Second, the title of E&O's article suggests that there is an alternative account, "epistemic utility." However, this concept is not defined in the article.<sup>5</sup> E&O state that "intuitively, people's subjective epistemic utility is measured by the relevance of some data for them given their goals" (1996, p. 358). Without an account of relevance or goals, this is uncontroversial: O&C and E&O agree that explaining the selection task involves specifying people's goals and specifying what relevance means. For O&C, a person's goal is to reduce uncertainty, and relevance means expected reduction in uncertainty. E&O suggest that the absolute value of log-likelihood ratios could be used as a measure of epistemic utility. However, this measure is insensitive to goals and, hence, is not a measure of epistemic utility by their criterion.

Third, E&O state that O&C "provide no psychological theory to explain subjects' selection whatever" (1996, p. 356). However, E&O do not explain what they mean by "psychological theory," why a psychological theory (in their sense) is a good thing, or why our theory is not psychological. It is therefore difficult to respond to this claim. Our model quantitatively fits data from a wide range of experiments and hence appears to be a psychological model of sorts. Perhaps E&O are using "psychological theory" to mean an algorithmic-level account rather

<sup>1</sup> Note that it is not a true distance (e.g., it is not symmetrical).

<sup>2</sup> However, participants could interpret the rule as exceptionless but as applying to some set of cards not including those that have been shown to include errors, such as cards that the computer will print in the future. Intuitively, this is analogous to a person checking whether a machine is now working after observing a breakdown. Consequently, O&C's original model of Kirby's data could apply to participants' interpretation of the experimental setup.

<sup>3</sup> This may be unsurprising given that we have allowed ourselves the luxury of an extra parameter. However, the model's predictions turn out to be insensitive to large variations of this parameter. Consequently, its function is to achieve a better mapping between task and model, not to achieve better data fits.

<sup>4</sup> Discussions of the rational<sub>1</sub>–rational<sub>2</sub> distinction (Evans, 1993; Evans, Over, & Manktelow, 1993) do not appear to be consistent. Evans (1993) identified rationality<sub>1</sub> as rationality of purpose and rationality<sub>2</sub> as rationality of process and claimed that "the notion of maximizing utility is clearly a case of rationality<sub>1</sub>" (p. 8). O&C used maximizing utility to explain deontic selection tasks, but E&O argue that this account is a rational<sub>2</sub> theory. It is also unclear why E&O imply that rational<sub>2</sub> explanation is a bad thing, given that according to Evans (1993), rationality<sub>2</sub> explanation is often successful in psychology (he cited the example of learning theory).

<sup>5</sup> Furthermore, we could not find a formal account of epistemic utility to compare with  $E(I_g)$  in any of the references E&O cite.

than a rational analysis. Following Anderson (1990, 1991, 1994), O&C assume that a complete psychological theory requires both levels of explanation but that rational analysis is prior to the algorithmic level (Oaksford & Chater, 1995b). Once a rational analysis has been specified, two questions arise: Are there algorithmic-level accounts that implement the rational analysis? and If there are many such accounts, how can they be distinguished empirically? Regarding the first question, because our rational analysis involves simple mathematical relationships, we can provide many different algorithmic-level accounts. Regarding the second question, because our rational analysis already captures the bulk of the empirical data, any implementation would capture these data. Therefore, without additional data, speculation at the algorithmic level seems premature.

### Summary

E&O provide a valuable critique that has required us to modify our account. First, we accommodated their objections to information gain without theoretical revision by using expected Kullback–Liebler distance. Second, we suggested that modifying our model may provide a better account of tasks in which participants know that there are exceptions (e.g., the tasks of Kirby, 1994b, and Pollard & Evans, 1983).

### Laming

Laming (1996) argues that O&C make implausible psychological assumptions and that a “correct” Bayesian analysis of the selection task makes the same predictions as the logical solution. We discuss Laming’s arguments in the order they arise.

### *How to Construct a Psychological Theory*

Here we address each point in Laming’s *Constructing Psychological Theories* section, which provides an overview of Laming’s arguments. First, Laming argues that O&C rely on arbitrary and psychologically implausible assumptions. O&C’s assumptions are not arbitrary but were derived from the theory of optimal data selection (Chaloner & Verdinelli, 1994; Good, 1960; Lindley, 1956; Luttrell, 1985; MacKay, 1992) and Bayesian epistemology (Earman, 1992; Horwich, 1982; Howson & Urbach, 1993; Mackie, 1963).<sup>6</sup> What is remarkable is that these assumptions, derived to solve normative problems in statistics and in epistemology, also make accurate predictions in the selection task. Furthermore, as we show later, each of our assumptions has a psychological justification.

Second, Laming objects to O&C’s information measure (Shannon–Wiener information) because “to be psychologically meaningful, the measure of information has to relate to *the question put to the subjects* [italics added]” (1996, p. 365). We argue that to be psychologically meaningful, an information measure has to relate to the question that participants *think* they have been asked. Experimenters cannot legislate for how people understand psychological tasks. The interpretation people adopt is an empirical matter that must be determined by fitting theoretical models to data. Because O&C’s model fits the data and Laming’s does not, it would seem that participants

may indeed interpret the problem as one of optimal data selection.

Third, Laming notes correctly that O&C use the subjective, rather than the “objective,” interpretation of probability. He argues that the subjective–objective distinction is “irrelevant to the validity of [O&C’s] theory” (Laming, 1996, p. 365). We argue, on the contrary, that the subjective interpretation is crucial and that failure to realize this leads Laming to misunderstand our model and to propose an inappropriate alternative. On the frequentist interpretation (Laming’s objective interpretation), probabilities are limiting frequencies in a repeated experiment (e.g., von Mises, 1939). Accordingly, probabilities can be assigned only to events that are repeatable so that limiting frequencies are defined. The frequentist view underlies classical approaches to hypothesis testing (e.g., Fisher, 1922; Neyman & Pearson, 1928). On the subjective interpretation, probabilities are degrees of belief (Keynes, 1921; Ramsey, 1931). Accordingly, probabilities can be assigned to all statements, including those describing unrepeatable events. Consequently, the probability that, for example, Oswald shot Kennedy is well defined, whereas on the frequentist interpretation it is not. The subjective interpretation underlies the Bayesian approach (Cox, 1946; de Finetti, 1937; Good, 1960; Lindley, 1971; Ramsey, 1931; Rosenkrantz, 1981).

Finally, Laming argues that our data fits are not impressive, assuming we set parameters arbitrarily. We show later that we set parameter values not arbitrarily but by reference to the literature on Bayesian epistemology. Moreover, we show that these parameter values, as with our other assumptions, are psychologically plausible. We now turn to Laming’s specific points.

### *Optimal Data Selection and Testing of Statistical Hypotheses*

The role of Laming’s *Testing Statistical Hypotheses* tutorial section seems to be twofold. First, it gives the impression that O&C’s optimal data selection account is suspect. Second, it provides the background for Laming’s “correct” Bayesian analysis. We address these issues in turn.

First, O&C’s analysis is not suspect but is a straightforward application of a Bayesian measure of the information provided by an experiment, introduced by Lindley (1956), one of the world’s leading Bayesian statisticians. Lindley (1956, p. 987) argued that

the measure of information [provided by an experiment] is given by Shannon’s function [i.e., Shannon–Wiener information] . . . [and that] prior probability distributions are . . . basic to the study. It seems obvious to the author that prior distributions, though usually anathema to the statistician, are essential to the notion of experimental information. To take an extreme case, if the prior distribution is concentrated on a single parameter value [or a

<sup>6</sup> In our original submission to this journal, we clearly outlined the origins of our assumptions. However, for reasons of journal space, the reviewers suggested, and we agreed, that this material should be left out. Consequently, we welcome Laming’s critique for the opportunity it affords us to make the origins of our assumptions explicit.

single hypothesis], that is if the state of nature is known, then no experiment can be informative.<sup>7</sup>

Thus, O&C simply have applied long-standing ideas from Bayesian statistics,<sup>8</sup> and Laming's suggestion that O&C's information measure is suspect misses the mark.

Second, Laming appears to misrepresent the Bayesian approach. He states that the use of priors is the essence of Bayesian statistics. However, as we discuss later, this misrepresents the Bayesian approach, which actually depends on the subjective interpretation of probability (Howson & Urbach, 1993; Lindley, 1971). Furthermore, Laming recommends estimating parameters using maximum likelihood, which has no Bayesian justification (Lindley, 1971). This is particularly inappropriate in the selection task, in which no data are available on which to base such estimates. These problems lead Laming to his "correct" Bayesian analysis, which, as we show later, is not really Bayesian.

### Psychological Assumptions

We now come to the core of Laming's argument: that our model does not apply to the selection task and that it fits the data only by using arbitrary and psychologically implausible assumptions. He makes six specific points to which we reply individually. First, however, we outline two important issues bearing on Laming's arguments.

### Rational Analysis and Task Interpretation

Laming states that the task set cannot be captured by O&C's model. This presupposes that the purpose of rational analysis is to specify what people should do given the task description; rational analysis has only a normative function. However, as we have argued (Oaksford & Chater, 1995b), the purpose of rational analysis is to characterize the task participants *think* they have been assigned. A rational analysis must be both normatively justified and descriptively adequate. In practice, this means that the focus is on modeling people's actual behavior rather than deriving models of the experimenter's preconceived ideas about what the task investigates.

The distinction between the task set as viewed by experimenter and participant is familiar in the reasoning literature. For example, critiques of Piaget's reasoning studies (Bower, 1974; Bryant & Trabasso, 1978; Donaldson, 1978; Harris, 1975) argued that many tasks were not understood by children. If the tasks were presented in a more child-centered way, reasoning previously absent would emerge. Another example is Smedslund's (1970; see Evans, 1993) observation that one cannot assess whether people reason logically independent of their task interpretation. A final example is the observation that "errors" in probabilistic reasoning may occur because the materials violate people's natural ways of representing probabilistic information (Birnbau, 1983; Gigerenzer, Hell, & Blank, 1988; Gigerenzer & Murray, 1987) so that the task participants tackle is not the task that the experimenter intended.

In summary, rational analysis characterizes how participants both interpret and solve a problem. Consequently, Laming's claim that O&C do not model the task people have been assigned is irrelevant because this was not their goal.

### Bayesian Epistemology and Rarity

Laming argues that our data fits rely on setting parameters arbitrarily. We now show that our rarity assumption, which determines our parameter values, is not arbitrary but derives directly from the literature on Bayesian epistemology.

In the quotation to follow, an influential Bayesian epistemologist (Horwich, 1982) discusses Mackie's (1963) solution to one of the paradoxes of confirmation theory (Goodman, 1983). The "ravens paradox" is that non-Bayesian confirmation theory entails that a nonblack nonraven (e.g., a pink flamingo) confirms the hypothesis that all ravens are black.

The central idea of Bayesian accounts is that our background assumptions concerning the proportion of ravens and black objects in the universe affect the extent to which hypotheses are confirmed by various kinds of evidence. Suppose we believe that the proportion of things which are ravens is very small: call it  $x$ ; and the proportion of black things  $y$ . Then our relevant background assumptions may be represented by the following table:

	R	not-R
B	$xy$	$(1-x)y$
not-B	$x(1-y)$	$(1-x)(1-y)$

Thus we suppose that the subjective probability of observing a black raven  $P(BR)$ , is  $xy$ ; and similarly,  $P(BnotR) = (1-x)y$ ,  $P(notBR) = x(1-y)$ , and  $P(notBnotR) = (1-x)(1-y)$ .

Now consider the table which according to Mackie, would represent the further supposition—All ravens are black:

	R	not-R
B	$x$	$y-x$
not-B	0	$1-y$

If  $H$  is true, there are no non black ravens. (Horwich, 1982, p. 56)

Mackie's argument implies that although a nonblack nonraven "will tend to confirm 'All ravens are black,' it will do so only to a negligible degree and will not carry as much weight as the observation of a black raven" (Horwich, 1982, p. 57), as long as  $x \approx 0$  (i.e., if rarity holds). Thus, contrary to Laming, O&C's parameters were not set simply to fit the data. O&C's goal was to determine whether Bayesian models that resolve conceptual problems in epistemology could also model human behavior.

### Response to Laming's Arguments

We now take the six specific assumptions that Laming identifies and show that they are normatively justified and that they make psychological sense.

<sup>7</sup> Note that Lindley, somewhat confusingly but for sound reasons (see Lindley, 1956, p. 989), introduced a sign reversal. O&C followed this convention, which caused some confusion, as pointed out by both E&O and Laming. In the Appendix and elsewhere (Oaksford & Chater, 1995a, 1995b), we adopt the standard convention of not reversing the sign.

<sup>8</sup> Laming, by contrast, recommends against using Shannon's measure and the use of prior distributions.

*Shannon–Wiener information.* Laming's criticism of O&C's use of Shannon–Wiener information has several problems. First, as we have shown, it is standard in Bayesian optimal data selection (Good, 1960; Lindley, 1956, 1971; MacKay, 1992). We suspect that Laming's objection derives from his view that optimal data selection does not apply to the task participants are assigned. But as we have already noted, O&C's goal was to model the task participants think they have been assigned.

Second, Laming claims that using O&C's measure is statistically inappropriate. He algebraically transforms our  $E(I_g)$  measure into his Equation 10, which measures “the expected information from a single event in favor of the communication channel being functional ( $H_i$  and  $D_k$  related) and against the alternative that they are independent” (Laming, 1996, p. 370–371). Note that Laming's Equation 10 is expected Kullback–Liebler distance (see our Equation A4 in the Appendix). Hence, Laming's analysis confirms our own. However, Laming argues that his interpretation using communication channels invalidates our model. This argument rests on the false assumption that if a formula has one interpretation, it cannot have another. Rather than invalidating our original interpretation, Laming has simply shown that our measure has yet another interpretation.

Third, Laming objects that O&C's measure does not discriminate between hypotheses. However, it does discriminate between hypotheses (Fedorov, 1972, chaps. 6 and 7). Discrimination depends on sequential sampling and recomputation of information gain to determine the optimal data to select next. This involves iteratively recomputing the priors at each stage in the standard Bayesian way. By selecting data using  $E(I_g)$ , the posteriors converge on the true hypothesis using the minimum number of observations. Thus, O&C's measure does discriminate between hypotheses. Of course, the selection task is not a sequential sampling task; participants never see the data. Nevertheless, Bayesian hypothesis testers should use their prior beliefs to select data that will optimize discrimination between models in the long run.

Laming also argues that O&C are faced with a paradox: that participants must already possess the information they should derive from the data. The “paradox” arises because Laming uses a frequentist interpretation of probabilities, whereas O&C use a subjectivist interpretation. According to the frequentist interpretation, the probability of uncovering a particular number or letter after having turned the card— $P(D_k|H_i)$ —must be either 0 or 1 (as Laming notes later). This is because however many times one turns the card, it will give the same result, and hence the limiting frequencies can take only the value 0 (one never reveals the number or letter) or 1 (one always reveals the number or letter). However, in the selection task, participants do not know what is on the other side of the card and, hence, cannot assign these probabilities (in the frequentist sense). But O&C used these probabilities in their calculations. Laming concludes that O&C's account assumes that participants must know what is on the back of the card, even before they have turned it.

Laming's difficulty is inevitable on his frequentist interpretation. On the subjectivist interpretation, however, there is no difficulty. The  $P(D_k|H_i)$  capture degrees of belief about what is on the back of the card, before it is turned. Because participants

are not certain what is on the back of the card, these probabilities will take intermediate values rather than being 0 or 1, depending on prior knowledge. This approach is standard in Bayesian statistics (e.g., Lindley, 1971). It also makes psychological sense, reflecting the psychologically reasonable assumption that prior knowledge will affect where one looks for evidence. In summary, O&C's information measure makes both normative and psychological sense.

*Rationality.* Laming objects that our rational recommendations do not provide a perfect fit with the experimental data (e.g., only 89% of participants choose the  $p$ -card, which is the most informative card). But requiring a perfect fit between theory and data seems entirely unreasonable and is not demanded of any other psychological theory.

Laming also states that

a rational Bayesian theory ought to look like this: Calculation shows that some particular card offers the greatest expected gain of information, and that card is the universal first choice. Depending on what is discovered on the underside of the card, one or other of the remaining cards is chosen next because it offers the greatest expected gain of information of those remaining. (1996, p. 370)

Laming is correct that a Bayesian analysis of the task assumes sequential sampling. But he is wrong to conclude that such an analysis is inappropriate to the selection task, in which participants choose cards without turning them over. As mentioned in the previous section, it is perfectly rational to select data to minimize the length of a sequential sample required to discriminate hypotheses before that sample becomes available.

*Bayesian analysis.* We are unclear about Laming's argument here. He appears to believe that, for the Bayesian, priors must be set from previous data, if they are not to reflect mere bias. Because, in the selection task, the participant sees no data, he or she assumes that priors cannot meaningfully be set. But Bayesian analysis must always begin from some priors before data are observed, on pain of infinite regress. The question of how priors should be set to take account of general knowledge is a major issue in Bayesian statistics (Berger, 1985; Box & Tiao, 1973; Lindley, 1971). Moreover, O&C have argued that people have a great deal of prior knowledge about conditionals (e.g., that rarity almost always holds), which is taken to be relevant to the task.

*Characterization of the task.* Here, Laming's objection seems to arise from his frequentist interpretation of O&C's probabilities. He imagines the situation in which there are many vowels, some with odd and some with even numbers on their undersides; in this context, the probability that a randomly chosen vowel has an even number on the back may lie between 0 and 1 (if one repeatedly chooses a random card with a vowel uppermost many times, the limiting frequency will be in proportion to the number of vowels with odd and even numbers on their undersides). Laming argues that O&C are implicitly committed to this setup if their intermediate probability values are to make sense.

Because O&C's account is Bayesian, however, all probabilities are degrees of belief, and hence no fictitious repeated experiments need be imagined to make sense of O&C's probabilities. Furthermore, there is evidence that participants do interpret the cards as being drawn from a larger population when only con-

fronted with four cards (Beattie & Baron, 1988). Moreover, when the experimenter draws the four cards from a larger pack in front of each participant before he or she performs the task (Evans & Lynch, 1973; Oaksford & Stenning, 1992), the results are the same as in the standard task.

Laming also argues that O&C do not consider the full range of possible hypotheses. O&C compare the rule with a particular independence model rather than a fully general "foil" model. This assumption was not introduced arbitrarily to fit the data. As the earlier quotation from Horwich reveals, Mackie used the same characterization of people's background knowledge to resolve the ravens paradox. Moreover, although Laming downplays explaining the data, that our simple model accurately captures the empirical results must be a virtue. Other researchers may propose alternative rational analyses, should these be necessary to capture further empirical data.

Laming also objects to O&C's assumption that participants discriminate between two particular instances of  $M_D$  and  $M_I$  rather than comparing these models in the abstract. This is reasonable because the values of  $a$  and  $b$  reflect particular degrees of belief in the antecedent being true and in the consequent being true when the antecedent is false. This assumption is psychologically innocuous. These values relate to people's degrees of belief about the proportions of various properties in their environments. It is psychologically reasonable to assume that people have access to this information. This assumption also makes normative sense; it resolves an important paradox in the logic of confirmation. In summary, our choice of models makes both normative and psychological sense.

*Identification of model parameters.* Laming objects that O&C equate parameters  $a$  and  $b$  between models. As shown in the quotation from Horwich, Mackie (1963) made a similar assumption, except that Mackie equated  $P(p)$  and  $P(q)$  between models, whereas O&C equate  $P(p)$  ( $a$ ) and  $P(q|not-p)$  ( $b$ ). Laming's objection is unclear because he endorses our rationale for equating these parameters, as we now show.

Equating  $a$  between models, as Laming notes, is equivalent to asserting that the antecedent ( $p$ ) has the same probability in each. As O&C argue, if, by contrast, the probability of  $p$  were, say, higher in  $M_D$  than in  $M_I$ , this would mean that observing  $p$  and  $not-p$  instances alone (without being able to see both sides of the cards) would discriminate between models (by the application of Bayes's theorem). Laming's response is puzzling: "Not true. . . . [The conditional rule] says nothing about the relative frequencies of vowels [ $p$  cards] and consonants [ $not-p$  cards]" (1996, p. 369). This is puzzling because Laming agrees that the conditional rule says nothing about the frequencies of  $p$  and  $not-p$  cards, which implies that it should not be possible to discriminate between models by observing one side of the cards. It is this intuition that requires equating the parameter  $a$  between models.

Similarly, we equated  $b$ , the probability of  $q$  in the absence of  $p$ , between models. Laming (1996, p. 369) argues that our "models are formulated the way they are in order to accommodate the relatively uncommon selection of the 'K.'" He then proposes alternative models in which the 2 card receives zero information gain rather than the K card, as in O&C. The suggestion is that O&C's decision to keep  $b$  constant between models was made solely to fit the data. However, O&C's assumption was constrained both psychologically and norma-

tively. Psychologically, it reflects the finding that participants regard false antecedent instances (i.e., the  $not-p$  cases) as irrelevant to the truth or falsity of a conditional rule. This has been established with an independent experimental paradigm: the truth-table task (Evans, 1972; Evans & Newstead, 1977; Johnson-Laird & Tagart, 1969). Furthermore, normatively, Quine (1959) suggested that conditional sentences do not assert a conditional but, rather, assert the consequent,  $q$ , conditional on the antecedent,  $p$ . From this logical point of view, cases in which the antecedent is false,  $not-p$  cases, are irrelevant to the truth or falsity of a conditional rule. No such evidence or normative proposals exist in support of the models Laming proposes in which the 2 card has zero information gain. Consequently, Laming's alternative model is irrelevant.

*The rarity assumption.* Laming objects to the rarity assumption, that  $P(p)$  and  $P(q)$  are low. He suggests that the rarity assumption has a bizarre consequence in the standard selection task. If the antecedent and consequent of the rule *If there is a vowel on one side of the card there is an even number on the other* are rare, then most cards must have consonants on one side and odd numbers on the other. However, the rarity assumption again makes perfect normative and psychological sense.

Normatively, the quotation from Horwich reveals that the rarity assumption is critical to Mackie's resolution of the ravens paradox. Moreover, Horwich's (1982) own analysis of this paradox assumes that  $P(not-p \ \& \ not-q) \approx 1$ . Consequently, the assumption that Laming appears to find bizarre is precisely the one that allows Bayesian confirmation theory to avoid paradox. Again, we based our assumptions on Bayesian epistemology and did not introduce them simply to fit the data. What is remarkable is that an assumption derived for this normative purpose should prove so valuable in modeling empirical data.

Psychologically, O&C argued that people's everyday encounters with conditionals influence their behavior in the selection task and that, in everyday contexts, rarity almost invariably holds. Thus, everyday strategies for hypothesis testing may be adapted to an environment where rarity is the norm. Moreover, O&C assumed that these default strategies are a major influence on behavior, even when participants do not know whether rarity holds. O&C (1994, pp. 627-628) provided two lines of experimental support for this claim: that explaining results on Wason's (1960) 2-4-6 task (Klayman & Ha, 1987) and causal reasoning (Anderson, 1990) both require rarity. More generally, we have argued extensively that people transfer their reasoning strategies from the everyday world to the laboratory (Chater & Oaksford, 1990, 1993; Oaksford & Chater, 1991, 1992, 1995a, 1995b). In summary, contrary to Laming, the rarity assumption is normatively and psychologically reasonable.<sup>9</sup>

### Data Coverage

Laming then argues that

all these assumptions [in O&C's rational analysis] are invoked to match merely the rank order of the frequencies with which the

<sup>9</sup> Klayman and Ha's minority phenomena assumption is somewhat less restrictive than our rarity assumption, specifying only that probabilities are less than one half.

different cards are selected for inspection [in the standard selection task]. . . . Moreover, if that rank order had been other than it is, it would simply have dictated different parameter values and assumptions. For that reason there is no need to examine O&C's treatment beyond the basic experimental paradigm. (1996, p. 371).

Laming's argument is wrong in two respects. First, O&C's parameters are not set arbitrarily. Consequently, according to his own reasoning, he must consider the other data that O&C explain. Second, even if O&C's parameters were set to explain the rank order in the standard task, then the other data would provide a test of O&C's model.

O&C have shown good fits with data from most of the studies reported on the selection task since Wason's (1966, 1968) original papers. For example, O&C's model captures the associations between card selections observed in abstract selection tasks (Pollard, 1985), data from the reduced array selection task (Johnson-Laird & Wason, 1970), the negations paradigm (e.g., Evans & Lynch, 1973; Manktelow & Evans, 1979), tasks with "fictional" outcomes (Kirby, 1994b), therapy experiments (Wason, 1969; Wason & Johnson-Laird, 1970), and a range of thematic selection tasks results (e.g., Cheng & Holyoak, 1985; Cosmides, 1989; Gigerenzer & Hug, 1992; Manktelow & Over, 1991). These studies do not represent a homogeneous set consisting of many near replications. On the contrary, they show that varying the nature of the task produces radically different results. O&C's rational analysis explains this variation. Furthermore, no other account of the selection task attempts this breadth of data coverage. We now turn to the second part of Laming's argument, that a "correct" Bayesian analysis confirms the standard logical solution.

### Laming's "Correct" Bayesian Analysis

Laming provides a "correct" Bayesian analysis of the selection task. He assumes that the conditional probability of an odd or even number on the back of, say, the A card must be one or zero, depending on whether the underside actually is odd or even. For Laming, that participants do not know whether the underside of the card is odd or even is not grounds for some intermediate probability because he does not interpret probabilities as degrees of belief. Laming argues that participants should turn only the A and 7 cards in the standard task, in line with the standard logical account.

Laming's "correct" Bayesian analysis is mathematically correct, but it is not Bayesian because it begins by rejecting the fundamental principle of Bayesian statistics, that probabilities are degrees of belief. Laming states that "the essence of Bayesian analysis is the inclusion of the priors, not that they be subjective" (1996, p. 366). This is a common misunderstanding against which Bayesians often warn (e.g., Howson & Urbach, 1993; Lindley, 1971). Comparing the Bayesian approach with the frequentist view, Howson and Urbach (1993, p. 11) stated that,

the other strand of inductive probability treats the probabilities as a property of our attitude toward them; such probabilities are then interpreted, roughly speaking, as measuring degrees of belief. This is called the *subjectivist or personalist interpretation*. The scientific

methodology based on this idea is usually referred to as the methodology of *Bayesianism*.

Laming confuses the use of Bayes's theorem (an uncontroversial theorem of probability theory) and Bayesian statistics (a vigorous, although controversial, approach to statistical inference).

In summary, Laming's "correct" Bayesian account poses no problems for O&C's rational analysis. Laming grants that his account does not fit the empirical data; for Laming, participants' behavior is simply not rational. However, O&C's rational analysis shows that behavior can be viewed as rational. It does not, and could not, show that behavior is rational on any defensible view of rationality. Therefore, the fact that Laming's non-Bayesian account gives different prescriptions is irrelevant.

### Summary

Laming misrepresents O&C's rational analysis because he is concerned with the task set rather than the task that participants think they have been assigned and because he misinterprets the statistical basis of our theory. Rational analysis must be normatively justified and descriptively adequate. O&C's account is normatively justified because it is based on Bayesian optimal data selection. It is descriptively adequate because it provides fits to a wide range of data without setting parameters arbitrarily.

### Almor and Sloman

A&S argue that O&C's model cannot account for data in cases in which  $p$  and  $not-q$  card responses are elicited without using deontic materials. A&S use four rules that they claim are not deontic and for which it is not clear whether O&C's rarity assumption holds. However, the logical  $p$  and  $not-q$  card response predominates for these rules. A&S conclude that O&C cannot explain these data. Moreover, they argue that their results are not compatible with any theory that uses the distinction between deontic and indicative tasks to explain so-called "facilitation" effects (i.e., choosing the  $p$  and  $not-q$  cards).<sup>10</sup>

A&S raise the important issue of how to explain  $p$  and  $not-q$  responses in nondeontic selection tasks. Such results threaten any theory that rules out  $p$  and  $not-q$  responding for nondeontic tasks. They are less threatening to O&C's model because it also allows  $p$  and  $not-q$  responses when the materials violate rarity. However, as A&S argue, it is unclear whether their materials do violate rarity.

There have been other demonstrations of  $p$  and  $not-q$  responses in abstract tasks without violating rarity (Green, 1995; Green & Larking, 1995; Platt & Griggs, 1993, 1995). However, they are problematic only on a strong interpretation of O&C's claim that violating rarity is not only sufficient but also necessary for the  $p$  and  $not-q$  response. Although A&S and these other experiments suggest that rarity violation may not be a necessary

<sup>10</sup> Note that if O&C are right, the view underlying this terminology—that participants' performance is facilitated from an initially irrational baseline—is wrong (see also Manktelow & Over, 1987). A&S are careful to avoid this misleading terminology.

condition for the *p* and *not-q* response, they do not question that rarity violation is a sufficient condition. Moreover, there is evidence that rarity violation is indeed sufficient for the *p* and *not-q* response (Kirby, 1994a, 1994b; Oaksford & Chater, 1995a; Sperber, Cara, & Girotto, 1995). Consequently, inducing high *p*, *not-q* selections without violating rarity is consistent with O&C's account.

However, we argue that A&S obtain high *p*, *not-q* selections only by altering the task. A&S's materials are either analytic (i.e., true by definition), and so O&C's model does not apply, or deontic, and so O&C's maximum expected utility model applies. We first compare A&S's experiments with other studies revealing the *p* and *not-q* response in the abstract task.

### *The p and not-q Response in the Abstract Task*

Other experiments revealing high *p*, *not-q* selections have used manipulations to force a logical interpretation of the rule (Green, 1995; Green & Larking, 1995; Platt & Griggs, 1993, 1995). For example, Platt and Griggs (1995, p. 60) explicitly provided the logical interpretation, telling participants that "a card with an A on its letter side can only have a 4 on its number side, but a card with a B on its letter side can have either a 4 or 5 on its number side." They also told participants to look for cards that violated this rule.<sup>11</sup> Green (1995) first told participants to imagine and write down all of the different possible combinations of letters and numbers for each card. Participants then had to imagine which combinations could violate the rule. Finally, they were asked to indicate which cards had such a combination. With this amount of coercion, observing high *p*, *not-q* selections is not surprising. What is more surprising is how few participants gave the *p* and *not-q* response. In most of Green's (1995) experiments, in the full externalization condition (outlined earlier) more than 50% of participants still did not make the *p* and *not-q* response. Platt and Griggs (1993, 1995) and Green and Larking (1995) obtained similar results. It seems that participants' natural reasoning strategies are very resistant even to these quite extreme attempts to force a logical interpretation in the abstract selection task.<sup>12</sup>

A&S's experiments contrast with these studies because A&S do not use any additional instructions to force a logical interpretation, but they achieve similar-sized effects. Therefore, A&S's manipulations are of more theoretical interest.

### *Deontic and Analytic Rules*

We argue that A&S's rules are either deontic or analytic.<sup>13</sup> However, even in their "abstract" experiment, A&S cue participants into realistic settings that have plausible deontic interpretations. Consequently, we could argue that our maximum expected utility account of the deontic selection task explains all of A&S's results. Our model predicts the *p* and *not-q* response for obligation rules and an enforcer's perspective that could reasonably characterize A&S's materials. However, as we noted earlier, we believe that the analytic nature of two of A&S's rules affects their results.

One can compare standard rules with each of A&S's rules by

asking what the reaction would be to a counterexample. Consider two standard rules:

1. If A on one side then 2 on the other.  
A3 implies that the rule is false.
2. If it's a raven then it's black.  
White raven implies that the rule is false.

Rule 1 is a standard selection task rule. As does Rule 2, it represents a claim about the way the world is. The reaction to the counterexample, A3, is that the rule is false.<sup>14</sup> Compare this with the reaction to a *p*, *not-q* instance of A&S's rules:

3. If a large object is stored then a large container must be used.  
Large object in a small container implies a contradiction.
4. If the weak force wins the strong force must have been weakened first.  
Weak force wins, strong force not weakened implies a contradiction.

In both Rule 3 and Rule 4, the counterexamples seem to violate the meaning of "large object" and "strong force." For example, large objects require large containers; otherwise, they would not be large. Similarly, strong forces overcome weak forces; otherwise, they would not be strong. This contrasts with the indicative rules in Rules 1 and 2.

Another test is to append "It must be the case that" to Rules 1–4. Although this results in true sentences for Rules 3 and 4, it is nonsense for Rule 2 (i.e., it is simply not true that "It must be

<sup>11</sup> As Platt and Griggs (1995) observed, the use of much modal terminology (i.e., "can" and "can only") and the violation instruction may well have induced a deontic context that produced the facilitation.

<sup>12</sup> Two of the studies just discussed, Green (1995) and Platt and Griggs (1995), claim to show that probabilistic manipulations fail to have the effects predicted either by Kirby (1994b) or by O&C (1994). However, in both cases the experimenters embedded the probabilistic manipulation in other manipulations, which we have outlined, designed to force a logical interpretation of the rule. Consequently, how these data bear on our model is obscure; other factors so confound the data as to make them uninterpretable. Moreover, as Platt and Griggs (1995) conceded, they cannot be sure that participants' subjective probabilities were appropriately calibrated to the letter and number frequencies used in these experiments. This is especially true because they made no distinction between type and token frequencies, which O&C (1994) argued may be an important factor. Green (1995) and Platt and Griggs (1995) assumed that what matters in terms of people's everyday hypothesis testing is that there are 5 vowels and 21 consonants (i.e., the frequencies of letter types). However, as O&C (1994) argued, it is more likely that people's prior experience with particular letter and number tokens provides the priors they use in optimal data selection.

<sup>13</sup> All of A&S's rules also use the modal "must" in the consequent. In contrast, the rules used in other studies eliciting *p* and *not-q* responses were explications of standard abstract rules. However, A&S used this modal in all of the rules in their experiments, so this is unlikely to be a factor.

<sup>14</sup> However, participants may not interpret the occurrence of a falsifying instance, A3, immediately as meaning that the rule is false, as A&S's Rule 9 reveals. It makes sense to seek evidence for the truth or falsity of this generalization; however, observing a white raven would not necessarily lead one to reject Rule 9 as a very useful rule. As we have argued (Oaksford & Chater, 1991, 1992, 1993, 1995b), most of the rules that make up people's world knowledge admit some exceptions.



the case that if it's a raven then it's black" and similarly for Rule 1). Rules 1 and 2 make contingent claims about how the world might be. Rules 3 and 4, in contrast, make definitional or analytic claims about how the world must be for these terms to apply. Analyticity matters for O&C's account because it is about how people optimally select data to determine the truth of a rule. But when a rule is analytically true— $P(M_D) = 1$ —there is no uncertainty, and so no data (no card selections) can reduce it. Consequently, optimal data selection does not apply to analytic materials in our model. It is therefore not surprising that A&S's results differed from results in standard selection tasks. No current theory of the selection task makes predictions when the conditional rule is analytic. Consequently, A&S's experiments require a novel theoretical analysis from any point of view. A&S's remaining rules have a different interpretation:

5. If a product gets a prestigious prize then it must have a distinctive quality.  
Winning prize has no distinctive quality implies that the rule is still in force.
6. If the product breaks then it must have been used under abnormal conditions.  
Product breaks under normal conditions implies that the rule is still in force.

We argue that A&S's contexts encourage participants to understand Rules 5 and 6 deontically. For example, in Rule 5, participants adopt the perspective of a journalist investigating prize-winning products. The criterion for winning the prize (having a distinctive quality) defines a norm (i.e., which products ought to win prizes). The rule is deontic. Specifically, the journalist is interested in whether these norms really determine which products win prizes (rather than, for example, prizes being awarded by corrupt means).

A final test is to append "It should be the case that" to Rules 2–6. Although this makes sense for Rules 5 and 6, it is nonsense for Rule 2 (i.e., "It should be the case that if it's a raven then it's black"). It is equally nonsensical for Rules 3 and 4 (e.g., it is not that one *should* store large objects in large containers; rather, one *must* store large objects in large containers). Given the deontic interpretation of Rules 5 and 6, we can explain these data using our maximum expected utility model.

All of the rules that A&S use differ from the rules normally used in the abstract selection task. For A&S's rules, it does not make sense to collect information to determine whether they are true or false. Therefore, O&C's optimal data selection model could not apply to any of them and, hence, is not challenged by A&S's results.

### Conclusion

E&O and Laming have given us the chance to elaborate the theoretical foundations and empirical consequences of our account. Our measure, to which E&O and Laming object for different reasons, is standard in Bayesian optimal data selection and can be reinterpreted to meet E&O's concerns. Furthermore, the assumptions to which Laming objects make sound normative sense, being derived from Bayesian epistemology. Our model also makes sound psychological sense, both because its assumptions are psychologically reasonable and because it is

consistent with further data that E&O believe to be problematic. Furthermore, A&S's data showing that analytic rules also elicit high  $p$ , *not-q* selections do not question that our model provides a sufficient condition for the  $p$  and *not-q* response in the abstract task. Nor do they question the theoretical distinction many researchers in this area have drawn between abstract and deontic tasks. In summary, E&O, Laming, and A&S do not provide grounds to question the view that our model provides the most compelling and comprehensive explanation of the selection task currently available.

### References

- Almor, A., & Sloman, S. A. (1996). Is deontic reasoning special? *Psychological Review*, *103*, 374–380.
- 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.
- Anderson, J. R. (1994). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Beattie, J., & Baron, J. (1988). Confirmation and matching biases in hypothesis testing. *Quarterly Journal of Experimental Psychology*, *40A*, 269–297.
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. New York: Springer-Verlag.
- Birnbaum, M. H. (1983). Base rates in Bayesian inference: Signal detection analysis of the cab problem. *American Journal of Psychology*, *96*, 85–94.
- Bower, T. G. R. (1974). *Development in infancy*. San Francisco: Freeman.
- Box, G. E. P., & Tiao, G. C. (1973). *Bayesian inference in statistical analysis*. Reading, MA: Addison-Wesley.
- Bryant, P. E., & Trabasso, T. (1978). Transitive inferences and memory in young children. *Nature*, *232*, 456–458.
- Chaloner, K., & Verdine, I. (1994). *Bayesian experimental designs: A review* (Tech. Rep. No. 599). Pittsburgh, PA: Department of Statistics, Carnegie Mellon University.
- Chater, N. (1993). Mental models and non-monotonic reasoning. *Behavioral and Brain Sciences*, *16*, 340–341.
- 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: Reply to Garnham. *Mind and Language*, *8*, 72–89.
- Cheng, P. W., & 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.
- Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory*. New York: Wiley.
- Cox, R. T. (1946). Probability, frequency and reasonable expectation. *American Journal of Physics*, *14*, 1–13.
- de Finetti, B. (1937). Foresight: Its logical laws, its subjective sources. In H. E. Kyburg & H. E. Smokler (Eds. and Trans.), *Studies in subjective probability* (pp. 55–118). Chichester, England: Wiley.
- Donaldson, M. (1978). *Children's minds*. Glasgow, Scotland: Collins.
- Earman, J. (1992). *Bayes or bust?* Cambridge, MA: MIT Press.
- Evans, J. St B T (1972). Interpretation and 'matching bias' in a reasoning task. *Quarterly Journal of Experimental Psychology*, *24*, 193–199.
- 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, & Lynch, J. S. (1973). Matching bias in the selection task. *British Journal of Psychology*, 64, 391-397.
- Evans, J. St B T, & Newstead, J. S. (1977). Language and reasoning: A study of temporal factors. *Cognition*, 8, 265-283.
- Evans, J. St B T, & Over, D. E. (1996). Rationality in the selection task: Epistemic utility versus uncertainty reduction. *Psychological Review*, 103, 356-363.
- Evans, J. St B T, Over, D. E., & Manktelow, K. I. (1993). Reasoning, decision making and rationality. *Cognition*, 49, 165-187.
- Fedorov, V. (1972). *Theory of optimal experiments*. London: Academic Press.
- Fisher, R. A. (1922). On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London, Series A*, 222, 309-368.
- Gigerenzer, G., Hell, W., & Blank, H. (1988). Presentation and content: The use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 513-525.
- Gigerenzer, G., & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition*, 43, 127-171.
- Gigerenzer, G., & Murray, D. J. (1987). *Cognition as intuitive statistics*. Hillsdale, NJ: Erlbaum.
- Good, I. J. (1960). Weight of evidence, corroboration, explanatory power, information and the utility of experiments. *Journal of the Royal Statistical Society, Series B*, 22, 319-331.
- Goodman, N. (1983). *Fact, fiction and forecast* (3rd ed.). Cambridge, MA: Harvard University Press.
- Green, D. W. (1995). Externalisation, counter-examples and the abstract selection task. *Quarterly Journal of Experimental Psychology*, 48A, 424-446.
- Green, D. W., & Larking, R. (1995). The locus of facilitation in the abstract selection task. *Thinking and Reasoning*, 1, 183-199.
- Harris, P. (1975). Development of search and object permanence during infancy. *Psychological Bulletin*, 82, 332-344.
- Horwich, P. (1982). *Probability and evidence*. Cambridge, England: Cambridge University Press.
- Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach* (2nd ed.). La Salle, IL: Open Court.
- Johnson-Laird, P. N., & Tagart, J. (1969). How implication is understood. *American Journal of Psychology*, 82, 367-373.
- Johnson-Laird, P. N., & Wason, P. C. (1970). Insight into a logical relation. *Quarterly Journal of Experimental Psychology*, 22, 49-61.
- Keynes, J. M. (1921). *A treatise on probability*. London: Macmillan.
- Kirby, K. N. (1994a). False alarm: A reply to Over and Evans. *Cognition*, 52, 245-250.
- Kirby, K. N. (1994b). 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.
- Kullback, S., & Liebler, R. A. (1951). Information and sufficiency. *Annals of Mathematical Statistics*, 22, 79-86.
- Laming, D. (1996). On the analysis of irrational data selection: A critique of Oaksford and Chater (1994). *Psychological Review*, 103, 364-373.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27, 986-1005.
- Lindley, D. V. (1971). *Bayesian statistics: A review*. Philadelphia: Society for Industrial and Applied Mathematics.
- Luttrell, S. P. (1985). The use of transinformation in the design of data sampling schemes for inverse problems. *Inverse Problems*, 1, 199-218.
- MacKay, D. J. C. (1992). Information-based objective functions for active data selection. *Neural Computation*, 4, 590-604.
- Mackie, J. L. (1963). The paradox of confirmation. *British Journal for the Philosophy of Science*, 38, 265-277.
- Manktelow, K. I., & Evans, J. St B 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 Language*, 2, 199-219.
- Manktelow, K. I., & Over, D. E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition*, 39, 85-105.
- Neyman, J., & Pearson, E. S. (1928). On the use of interpretation of certain test criteria for purposes of statistical inference. *Biometrika*, 20, 175-240, 263-294.
- Oaksford, M. (1993). Mental models and the tractability of everyday reasoning. *Behavioral and Brain Sciences*, 16, 360-361.
- 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. (1995a). Information gain explains relevance which explains the selection task. *Cognition*, 57, 97-108.
- Oaksford, M., & Chater, N. (1995b). 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, 835-854.
- Platt, R. D., & Griggs, R. A. (1993). Facilitation in the abstract selection task: The effects of attentional and instructional factors. *Quarterly Journal of Experimental Psychology*, 46A, 591-613.
- Platt, R. D., & Griggs, R. A. (1995). Facilitation and matching bias in the abstract selection task. *Thinking and Reasoning*, 1, 55-70.
- Pollard, P. (1985). Nonindependence of selections on the Wason selection task. *Bulletin of the Psychonomic Society*, 23, 317-320.
- Pollard, P., & 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*. London: Hutchinson.
- Quine, W. V. O. (1959). *Methods of logic*. New York: Holt.
- Ramsey, F. P. (1931). *The foundations of mathematics and other logical essays*. London: Routledge & Kegan Paul.
- Rosenkrantz, R. D. (1981). *Foundations and applications of inductive probability*. Atascadero, CA: Ridgeview.
- Smedslund, J. (1970). On the circular relationship between logic and understanding. *Scandinavian Journal of Psychology*, 11, 217-219.
- Sperber, D., Cara, F., & Girotto, V. (1995). Relevance explains the selection task. *Cognition*, 57, 31-95.
- von Mises, R. (1939). *Probability, statistics and truth*. London: Allen & Unwin.
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12, 129-140.
- Wason, P. C. (1966). Reasoning. In B. Foss (Ed.), *New horizons in psychology* (pp. 135-151). Harmondsworth, Middlesex, England: 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.
- Wason, P. C., & Johnson-Laird, P. N. (1970). A conflict between selecting and evaluating information in an inferential task. *British Journal of Psychology*, 61, 509-515.

## Appendix

## Proof of the Equivalence of Expected Information Gain and Expected Kullback–Liebler Distance

Consider hypotheses,  $h$ , and data,  $d$ . The uncertainty associated with  $h$  before the data are collected is

$$-\log_2[P(h)]. \quad (\text{A1})$$

This is sometimes known as the surprisal of  $h$ . The uncertainty associated with  $h$  after the data are collected is

$$-\log_2[P(h|d)] \quad (\text{A2})$$

(i.e., the same as Equation A1 but with the appropriate revision of the probability). Information gain,  $I_g$ , is therefore

$$\begin{aligned} -\log_2[P(h)] - \{-\log_2[P(h|d)]\} \\ = \log_2[P(h|d)] - \log_2[P(h)]. \quad (\text{A3}) \end{aligned}$$

We are interested in the expectation of this quantity with respect to the joint distribution of  $h$  and  $d$ . Expected information gain,  $E(I_g)$ , is

$$\begin{aligned} E(I_g) = \sum_{h,d} \{ \log_2[P(h|d)] - \log_2[P(h)] \} \\ = \left\{ \log_2 \left[ \frac{P(h|d)}{P(h)} \right] \right\} \\ = \sum_{h,d} P(h,d) \log_2 \left[ \frac{P(h|d)}{P(h)} \right]. \quad (\text{A4}) \end{aligned}$$

The calculations in Equation A4 simply write out the expectation explicitly.

Now we turn to our new approach, which is based on the difference between new and old distributions. The Kullback–Liebler distance from a distribution  $P'(x)$  and a distribution  $P(x)$  is

$$D(P',P) = \sum_k P'(x_k) \log_2 \left[ \frac{P'(x_k)}{P(x_k)} \right]. \quad (\text{A5})$$

The distribution of interest here is the distribution of belief in the available hypotheses,  $h$ . The new distribution is given by the  $P(h|d)$  values that take the data into account; the old distribution is given by the  $P(h)$  values. Applying Equation A5, the Kullback–Liebler distance from the new to the old distribution is

$$D(P^{new}, P^{old}) = \sum_h P(h|d) \log_2 \left[ \frac{P(h|d)}{P(h)} \right]. \quad (\text{A6})$$

We are already summing over hypotheses, so we need take expectations only over data (taking expectations over the joint distribution of  $h$  and  $d$  produces the same result). The expected value of  $D$ ,  $E(D)$ , is given by

$$\begin{aligned} E(D) &= \sum_d \left\{ \sum_h P(h|d) \log_2 \left[ \frac{P(h|d)}{P(h)} \right] \right\} \\ &= \sum_d P(d) \sum_h P(h|d) \log_2 \left[ \frac{P(h|d)}{P(h)} \right] = \sum_{h,d} P(h,d) \log_2 \left[ \frac{P(h|d)}{P(h)} \right]. \end{aligned}$$

Because all of the calculations in O&C (1994) involve  $E(I_g)$ , we can adopt  $E(D)$  without altering any substantive aspect of the original analysis. Furthermore,  $D$  has none of the counterintuitive properties that E&O point out for  $I_g$ . Specifically, Kullback–Liebler distance is nonnegative, unlike information gain, which may be positive or negative. Note also that expected Kullback–Liebler distance is nonnegative (because the expectation of a nonnegative quantity is nonnegative), and therefore expected information gain must also be nonnegative.

Received August 9, 1995

Revision received November 17, 1995

Accepted November 28, 1995 ■