Psychological Review
1999, Vol. 106, No. 1, 223–227

0033-295X/99/$3.00

Information Gain and Decision-Theoretic Approaches to Data Selection: Response to Klauer (1999)

Nick Chater
University of Warwick

Mike Oaksford
Cardiff University

K. C. Klauer (1999) argued that a Bayesian decision-theoretic rational analysis of Wason's selection task is preferable to an information gain account (M. Oaksford & N. Chater, 1994) because it has a better normative justification and may provide a better fit with the empirical data. The authors argue that Klauer's proposal and their proposal are equally well justified from a normative perspective and that, where the predictions of the 2 approaches diverge, the existing empirical evidence is consistent with the information gain approach. However, more empirical research is required to decide between these 2 accounts.

Klauer's (1999) elegant and tightly argued article proposes two novel rational analyses of Wason's selection task (Wason, 1966, 1968) based on a Bayesian decision-theoretic approach to optimal data selection. One account applies to sequential decision problems, and the other applies to nonsequential decision problems. Klauer argues that these accounts are better justified and provide better fits with the empirical data than our rational analysis (Oaksford & Chater, 1994). Klauer's proposals are important and will stimulate further theoretical and empirical research. However, we argue that there are no normative grounds to prefer a decision-theoretic approach over our information gain model and that, where the predictions of Klauer's approach and our approach diverge, the data are not inconsistent with information gain. Before outlining where the decision-theoretic and information gain approaches diverge, it is important to stress that they have many common features.

Common Features

First, Klauer's approach and our approach both adopt the framework of rational analysis (Anderson, 1990, 1991; Oaksford & Chater, 1998b), which explains behavior as (perhaps approximately) adapted to achieving some goal, given some environment. Rational analysis must be both normatively justified (i.e., it must provide a justified way of achieving the goal in the given environment) and descriptively adequate (i.e., its predictions should fit with the empirical data). In contrast, most accounts of the selection task assume that performance is not normatively justified (e.g., Evans, 1983, 1984, 1989; Johnson-Laird & Byrne, 1991; Rips, 1994; Wason, 1968).

Second, Klauer's model and ours both view the selection task as a problem of optimal data selection. That is, they assume that the goal of card selection is to discriminate as well as possible between alternative hypotheses, and they also agree on the particular probabilistic formulation of these hypotheses.

Third, Klauer's model and our model adopt a Bayesian framework that associates numbers with degrees of belief and applies the probability calculus to these numbers. Thus, both the information gain and decision-theoretic approaches would be inadmissible according to Laming (1996), who denied the appropriateness of the Bayesian framework.

Finally, the models agree on most of their empirical predictions (see also Evans & Over, 1996; Nickerson, 1996), both predicting a range of effects when probabilities are varied experimentally. In contrast, most accounts of the selection task assume that these factors should have no effect on card selections (e.g., Evans, 1984, 1989; Johnson-Laird & Byrne, 1991; Rips, 1994; Sperber, Cara, & Girotto, 1995; Wason, 1968). Empirical studies have confirmed that probabilistic factors do affect selection task performance (Green, Over, & Pyne, 1997; Kirby, 1994; Manktelow, Sutherland, & Over, 1995; Oaksford & Chater, 1995; Oaksford, Chater, Grainger, & Larkin, 1997).

Thus, in the space of possible accounts of the selection task. Klauer's Bayesian decision-theoretic accounts and our information gain model are close neighbors. Nonetheless, we agree with Klauer that the differences between these accounts are important in terms of both normative justification and descriptive adequacy.

Normative Justification

There are two starting points for analyzing any empirical inquiry, with fundamentally different goals (see Chater, Pickering, & Crocker, 1998, for a detailed discussion). First, decision-theoretic inquiry begins with the goal of rational choice, the goal is to make the best possible decisions (Berger, 1985; Lindley, 1971, chap. 4). According to this viewpoint, inquiry is worthwhile only if it leads to better decisions. Second, disinterested inquiry begins with the goal of searching for a framework for rational thought; the goal is to know as much as possible about the world (Earman, 1992; Howson & Urbach, 1989; Lindley, 1971, chap. 12). According to this viewpoint, inquiry is valuable for its own sake, because it

Correspondence concerning this article should be addressed to Nick Chater, Department of Psychology, University of Warwick, Coventry, CV2 7AL, England, or to Mike Oaksford, School of Psychology, Cardiff University, P.O. Box 901, Cardiff, CF1 3YG, Wales, United Kingdom. Electronic mail may be sent to oaksford@cardiff.ac.uk.

leads to knowledge: Inquiry is "disinterested." Both of these viewpoints can be formulated precisely.

The decision-theoretic approach to inquiry is formalized as part of Bayesian decision theory (Berger, 1985) and is the basis of Klauer's account. Here the measure of the value of inquiry is the expected utility, over and above the costs of the inquiry. The disinterested approach was formalized in Lindley's (1956) account of optimal data selection, which was closely related to Good's (1950) measure of "weight of evidence" and was the motivation for our (Oaksford & Chater, 1994) model (see Oaksford & Chater, 1996). Here the measure of the value of inquiry is expected information gain, where information is defined according to standard information theory (Cover & Thomas, 1991). Thus, both theories are equally normatively justified, both corresponding to formally well-defined normative theories in Bayesian statistics.[1]

These approaches define different notions of the value of turning a card in the selection task, one in terms of utilities and the other in terms of information gain. But the pragmatics of the task instructions clearly indicate that there must, in some sense, be constraints on how many cards are turned; otherwise, the task can be solved trivially by turning every card. One attractive feature of Klauer's decision-theoretic approach is that it provides a simple way of characterizing this constraint in terms of costs (negative utilities). By using these costs as a free parameter and optimizing the fit to the data, Klauer can model the number of cards that participants turn. This is an advance over our model (Oaksford & Chater, 1994). We focused on deriving an informational ordering over the four cards and simply assumed that participants would choose the most informative cards in this ordering. We left the total number of cards turned to be determined by the pragmatics of the instructions rather than including this in our model. Consequently, we predicted that if participants decided to turn just one card, this would be the most informative single card, *p;* if they decided to turn two cards, they would turn the two most informative cards, *p* and *q;* and, if they decided to turn three cards, they would turn the three most informative cards, *p, q,* and *not q* (assuming rarity). It is in this sense that we discussed "optimal" data selection: Participants choose the maximally informative cards, given their choice of the number of cards to select. Klauer goes further to the extent that he also models how the participants choose how many cards they should select, given the pragmatic constraints of the task. He models this choice using a decision-theoretic approach that associates turning each card with a cost.[2]

We have argued that decision-theoretic and information gain rational analyses have equal normative status. Klauer's decision-theoretic approach has the advantage over our approach that it models how many cards are chosen rather than merely the frequency ordering over them. But it has the disadvantage not shared by our model that it requires the specification of utilities, whereas in the abstract selection task no utilities are mentioned, and no actions are performed.[3] From a normative perspective, each side seems to have its advantages and disadvantages. But rational analysis is, of course, intended to provide not only a normatively justified, but also a descriptively adequate, account of cognitive processes. Therefore, we now turn to the empirical data.

Descriptive Adequacy

Although Klauer's two decision-theoretic rational analyses and our information gain model make many common predictions, there are points of contrast.

Klauer's Sequential Model Versus Information Gain

Klauer identifies a qualitative difference between the predictions of the sequential decision-theoretic model and information gain: When the "rarity" assumption holds—that is, $P(p)$ and $P(q)$ are small—and participants believe that rule is most likely to be false—that is, $P(M_D) < .5$—the sequential decision-theoretic model predicts more *not q* than *q* card selections. Information gain makes the opposite prediction.

Klauer cites studies that appear to favor the decision-theoretic model (Fiedler & Hertel, 1994; Love & Kessler, 1995; Pollard & Evans, 1983). However, he also cites studies that appear to show no effects of manipulating believability (Green & Over, 1997; Oaksford, Chater, & Grainger, 1997). Klauer therefore concludes that although "a number of studies suggest that the predictions made by the new procedure are descriptively more adequate . . . the issue is still under debate" (1999, p. 221). Thus, Klauer views the evidence as inconclusive.

In contrast, we argue that the studies appearing to support the decision-theoretic model all suffer from the same problem: They vary the exceptions that a rule has but not its believability. Fiedler and Hertel (1994) used a falsification instruction implicitly em-

[1] It is important to realize that Bayesian approaches are characterized by their adoption of a subjective view of probability (see Oaksford & Chater, 1996, for discussion), not by the adoption of a decision-theoretic approach. Many Bayesians explicitly do not adopt a decision-theoretic stance. For example, Earman (1992) adopted a non-decision-theoretic Bayesian perspective on the philosophy of science.

[2] This approach was also used in our (Oaksford & Chater, 1994) account of the deontic selection task. An interesting research question is how Klauer's approach might be extended to the deontic selection task and whether such an extension would differ from our decision-theoretic account of this task.

Another possible reason for why Klauer assumes information gain is the wrong measure is that it does not minimize the length of a sequential sample. Klauer points out that our (Oaksford & Chater, 1996) passing comment that this was the case is incorrect, and he gives a valuable analysis of the problem of converging on the true hypothesis. However, it should be clear that this was not the motivation for using the information gain measure; indeed, in Oaksford and Chater (1994), the task was not modeled in sequential terms, and the question of convergence did not arise. Thus, Klauer (1999, p. 217) is wrong to assume that we accept minimizing the number of observations as the performance criterion appropriate to sequential sampling; we simply erroneously assumed that this followed as a side effect of the information gain approach.

[3] However, as Klauer points out, tasks with explicit costs can be imagined and might be treated from a decision-theoretic viewpoint. We agree with Klauer that exploring such tasks is an important empirical program. Note also that turning a card does not constitute an action in a decision-theoretic sense but is part of the data collection process. For example, with respect to the rule *If you eat tripe, you become sick,* the actions of interest might include whether people should be allowed to eat tripe, whether one personally will eat tripe, and so on, not whether a particular piece of evidence concerning the rule is selected.

phasizing the possibility of a *p, not q* case (i.e., emphasizing exceptions). Love and Kessler's (1995, p. 153) focusing manipulation was "implemented in two ways, either by emphasizing the significance of the rule violations or by discussing the probability of *not-q* occurring with *p*" (i.e., emphasizing exceptions). Similarly, in Pollard and Evans (1983), the learning phase introduced two rules that had either few or many exceptions.

However, there are good reasons to believe that whether a rule has exceptions and whether one believes the rule can be orthogonal judgments (Green & Over, 1997; Oaksford & Chater, 1998b). In our model and the Klauer model, the rule (i.e., M_D) allows no exceptions, so any exception leads to definitive falsification, that is, $P(M_D) = 0$. But this rigid interpretation may not be psychologically appropriate. For example, people believe that birds fly, although they are aware that ostriches, penguins, and injured birds cannot fly (i.e., *birds fly* is a default rule). From a variety of perspectives in cognitive psychology, philosophy of language, philosophy of science, and artificial intelligence, there is general agreement that almost all commonsense rules are default rules (Barwise & Perry, 1983; Goodman, 1954/1983; Holland, Holyoak, Nisbett, & Thagard, 1986; Holyoak & Spellman, 1993; Minsky, 1977; Oaksford & Chater, 1991, 1993, 1998a; Winograd & Flores, 1986). That is, they allow exceptions. Thus, a more realistic version of the decision-theoretic and information gain accounts should allow rules to be probabilistic rather than rigid. We have made this modification (Oaksford & Chater, 1998b), as have Over and Jessop (1998). Once exceptions are allowed, the connection between exceptions and believability can be broken, so it cannot always be assumed that manipulating exceptions also manipulates believability.

An extreme example from the philosophy of science illustrates this point: Physicians believe that the best predictor of whether someone has paresis is whether he or she has syphilis, even though the probability of someone having paresis given that he or she has syphilis is only about .01 (i.e., 99% of cases are counterexamples; Bromberger, 1965). Crucially, a cause (*c*) is present when it makes its effects (*e*) more probable, that is, when $P(e|c) > P(e)$ (Cheng, 1997; Salmon, 1970), not simply when $P(e|c)$ is high. In the philosophy of science, this point was raised as a criticism of Hempel's (1965) account of inductive–statistical explanation (Bromberger, 1965; Salmon, 1970). Green and Over (1997) made the same point in introducing their Experiment 2. Moreover, Cheng (1997) based her probabilistic contrast model of causal induction on this idea. Highly believed rules in which the probability of the consequent given the antecedent is low also occur in everyday life. For example, many people believe quite strongly that allowing children to walk home from school increases their chance of being abducted, although the probability of a child being abducted while walking home from school is tiny. In sum, the relationship between number of exceptions and degree of belief in a rule can be broken.

Thus, testing whether low belief in the rule can lead to more *not q* card selections requires manipulating believability directly rather than manipulating exceptions. This was done in two of the studies Klauer cites. In Green and Over's (1997) Experiment 2, participants selected evidence about a causal conditional rule concerning the symptoms revealed by an imaginary disease. Participants were also told to imagine that they had "already examined a large number of patients and have found a lot of evidence that the claim is true (false)." In this procedure, believability was varied directly rather than by emphasizing exceptions. Green and Over found that an almost identical proportion of participants turned the *not q* card in the believed true (55.3%) and believed false (54.9%) conditions. A similar result was obtained by Oaksford, Chater, and Grainger (1997; see also Oaksford & Chater, 1998a). In their Experiments 1 and 2, they had participants rate the believability of rules that they then tested in a selection task. In neither experiment was there any indication that disbelieving the rule led to more *not q* card selections. In Experiment 1, 29.7% of participants selected the *not q* card in the high belief condition, and 25% did so in the low belief condition. In Experiment 2, 33.1% of participants selected the *not q* card in the high belief condition, as compared with 33.4% in the low belief condition. The absence of any effect of believability is not consistent with the sequential decision-theoretic model, but it is consistent with the information gain model. However, there is clearly much scope for more empirical research on the effects of believability.[4]

Klauer's Nonsequential Model Versus Information Gain

Klauer suggests that his nonsequential model is most appropriate to the selection task. Therefore, it seems reasonable to expect that this model should be most consistent with the empirical data. Klauer shows that the nonsequential model can account not only for the standard selection task rule *if p, then q,* where the *p* and *q* card selections predominate, but also for the negations paradigm (Evans & Lynch, 1973), where negations are included in the rules. We modeled rules such as *if p, then not q,* where the falsificatory response dominates, by assuming that negated categories are high probability categories. Klauer shows that, with the nonsequential model, the best fit parameter values for $P(p)$ and $P(q|not\ p)$ bear out this assumption. However, another parameter, the ratio $P(H_0)l_0/P(H_1)l_1$, also changes between these two rules.

In modeling the standard *if p, then q* rule, the best fit value for this parameter was 13.3, whereas, for the *if p, then not q* rule, the best fit value was 1.44. These values mean that, assuming that the losses, l_0 and l_1, do not change between rules, $P(H_0)$ is higher for the standard rule than for the negated consequent rule. Consequently, the nonsequential model appears to predict that people make more apparently falsificatory selections when they have

[4] Why do studies manipulating exceptions show increased *not q* card selections? We suggest that they mirror the earlier "therapy" experiments (Wason, 1969; Wason & Golding, 1974; Wason & Johnson-Laird, 1970) showing increased *not q* card selections as various therapies focused attention on falsifying cases, just as in two of the studies cited by Klauer (Fiedler & Hertel, 1994; Love & Kessler, 1995). Oddly, Love and Kessler (1995, p. 164) suggested that these early attempts at "directing subjects to consider the falsifying selection of *not-q* were not successful." However, in Wason (1969), 42% of participants finally made the *p, not q* response, which compares favorably with Love and Kessler (1995), who, using a similar procedure (Experiment 2, Task 3), found just 35% *p, not q* responses. We explained the therapy experiments as participants lowering their criterion for card selection down the informativeness scale (see Oaksford & Chater, 1994, pp. 620–621), which can also explain Fiedler and Hertel's (1994) and Love and Kessler's (1995) results. We suggested (Oaksford & Chater, 1996) and have shown (Oaksford & Chater, 1998b) that Pollard and Evans's (1983) study can be captured by the information gain account if exceptions are allowed.

relatively high belief in the rule, exactly the opposite prediction to that of the sequential model. This feature of the nonsequential model is, of course, consistent neither with the results Klauer reports in favor of his sequential model (Fiedler & Hertel, 1994; Love & Kessler, 1995; Pollard & Evans, 1983) nor with the more recent results we have discussed in which believability was manipulated directly (Green & Over, 1997; Oaksford, Chater, & Grainger, 1997). That is, the existing data offer no support for the nonsequential decision-theoretic model where its predictions diverge from information gain.[5]

Conclusion

Klauer proposes two important rational analyses of the selection task set in a decision-theoretic framework, arguing that these are normatively preferable and may be descriptively preferable to our information gain account. We have argued that the normative status of the two approaches is equal; they are simply derived from different, equally important models of inquiry. Information gain is derived from the general goal of "disinterested inquiry," whereas the decision-theoretic approach is derived from the goal of making better decisions. We have suggested that each has advantages and disadvantages with respect to modeling the standard selection task. On the one hand, information gain does not model costs so as to choose the number of cards to select, in contrast to the decision theoretic approach. On the other hand, in contrast to information gain, the decision-theoretic approach requires explicit utilities or actions that are not introduced in the standard task. We have also argued that, contrary to Klauer's suggestion, the existing data are consistent with the information gain approach, although they apparently contradict his own nonsequential model. Clearly, more empirical research is required to determine which kind of optimal data selection account best captures Wason's selection task. However, that some version of this approach is appropriate to the standard selection task is now firmly established. Moreover, Klauer's stimulating article opens up new theoretical and experimental directions, not only in relation to the selection task but in relation to the general problem of the psychology of human inquiry.

[5] We thank Karl Christoph Klauer (personal communication, March 11, 1998) for confirming this general pattern. In the model, the effect of prior beliefs appears complex, but note that the empirical data show no effects of this parameter.

References

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

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

Barwise, J., & Perry, J. (1983). *Situations and attitudes.* Cambridge, MA: MIT Press.

Berger, J. O. (1985). *Statistical decision theory and Bayesian analyses.* New York: Springer-Verlag.

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

Chater, N., Pickering, M., & Crocker, M. (1998). The rational analysis of inquiry: The case of parsing. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 441–468). Oxford, England: Oxford University Press.

Cheng, P. W. (1997). From co-variation to causation: A causal power theory. *Psychological Review, 104,* 367–405.

Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory.* New York: Wiley.

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

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

Evans, J. S.B. T. (1984). Heuristic and analytic processes in reasoning. *British Journal of Psychology, 75,* 451–468.

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

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

Evans, J. S.B. T., & Over, D. E. (1996). Rationality in the selection task: Epistemic utility versus uncertainty reduction. *Psychological Review, 102,* 356–363.

Fiedler, K., & Hertel, G. (1994). Content-related schemata versus verbal-framing effects in deductive reasoning. *Social Cognition, 12,* 129–147.

Good, I. J. (1950). *Probability and the weighting of evidence.* London: Griffin.

Goodman, N. (1983). *Fact, fiction and forecast.* Cambridge, MA: Harvard University Press. (Original work published 1954)

Green, D. W., & Over, D. E. (1997). Causal inference, contingency tables and the selection task. *Cahiers de Psychologie Cognitives, 16,* 459–487.

Green, D. W., Over, D. E., & Pyne, R. A. (1997). Probability and choice in the selection task. *Thinking and Reasoning, 3,* 209–236.

Hempel, C. G. (1965). *Aspects of scientific explanation and other essays in the philosophy of science.* New York: Free Press.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. (1986). *Induction: Processes of inference, learning and discovery.* Cambridge, MA: MIT Press.

Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology, 44,* 265–315.

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

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

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

Klauer, K. C. (1999). On the normative justification for information gain in Wason's selection task. *Psychological Review, 106,* 216–223.

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. *Annuals of Mathematical Statistics, 27,* 986–1005.

Lindley, D. V. (1971). *Bayesian statistics: A review.* Philadelphia: Society for Industrial and Applied Mathematics.

Love, R. E., & Kessler, C. L. (1995). Focusing in Wason's selection task: Content and instruction effects. *Thinking and Reasoning, 1,* 153–182.

Manktelow, K. I., Sutherland, E. J., & Over, D. E. (1995). Probabilistic factors in deontic reasoning. *Thinking and Reasoning, 1,* 201–220.

Minsky, M. (1977). Frame system theory. In P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking: Readings in cognitive science* (pp. 355–376). Cambridge, England: Cambridge University Press.

Nickerson, R. S. (1996). Hempel's paradox and Wason's selection task: Logical and psychological puzzles of confirmation. *Thinking and Reasoning, 2,* 1–32.

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

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

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

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

Oaksford, M., & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review, 103,* 381–391.

Oaksford, M., & Chater, N. (1998a). *Rationality in an uncertain world: Essays on the cognitive science of human reasoning.* Hove, Sussex, England: Psychology Press.

Oaksford, M., & Chater, N. (1998b). A revised rational analysis of the selection task: Exceptions and sequential sampling. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 372–398). Oxford, England: Oxford University Press.

Oaksford, M., Chater, N., & Grainger, B. (1997). *Probability effects in data selection.* Unpublished manuscript, School of Psychology, Cardiff University, Cardiff, Wales.

Oaksford, M., Chater, N., Grainger, B., & Larkin, J. (1997). Optimal data selection in the reduced array selection task (RAST). *Journal of Experimental Psychology: Learning, Memory and Cognition, 23,* 441–458.

Over, D. E., & Jessop, A. (1998). Rational analysis of causal conditionals and the selection task. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 399–414). Oxford, England: Oxford University Press.

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

Rips, L. J. (1994). *The psychology of proof.* Cambridge, MA: MIT Press.

Salmon, W. (1970). *Statistical explanation and statistical relevance.* Pittsburgh, PA: University of Pittsburgh Press.

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

Wason, P. C. (1966). Reasoning. In B. M. Foss (Ed.), *New horizons in psychology* (pp. 135–151). Harmondsworth, Middlesex, England: Penguin Books.

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., & Golding, E. (1974). The language of inconsistency. *British Journal of Psychology, 65,* 537–546.

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.

Winograd, T., & Flores, F. (1986). *Understanding computers and cognition.* New York: Academic Press.

Received March 16, 1998
Revision received August 5, 1998
Accepted August 5, 1998 ■