Johnson-Laird, P.N., Byrne, R.M.J., & Schaeken, W.S. (1992). Propositional reasoning by model. *Psychological Review, 99*, 418–439.
Johnson-Laird, P.N., & Savary, F. (1995). How to make the probable seem impossible. In J. Moore & J. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates Inc.
Liberman, N., & Klar, Y. (in press). Hypothesis testing in Wason's selection task: Social exchange, cheating or task understanding. *Cognition*.
Love R., & Kessler, C. (1995). Focusing in Wason's selection task: content and instruction effects. *Thinking and Reasoning, 1*(2), 153–182.
Manktelow, K.I., & Over, D.E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition, 39*, 85–105.
Oakhill, J.V., Johnson-Laird, P.N., & Garnham, A. (1989). Believability and syllogistic reasoning. *Cognition, 31*, 117–140.
Oaksford, M., & Stenning, K. (1992). Reasoning with conditionals containing negated constituents. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 835–854.
Politzer, G., & Nguyen-Xuan, A. (1992). Reasoning about conditional promises and warnings: Darwinian algorithms, mental models, and relevance judgements or pragmatic schemas. *Quarterly Journal of Experimental Psychology, 44*, 401–412.
Sperber, D., Cara, F., & Girotto, V. (in press). Relevance theory explains the selection task. *Cognition*.
Wason, P.C. (1966). Reasoning. In B.M. Foss (Ed.), *New horizons in psychology*. Harmondsworth, UK: Penguin.

Two and Three Stage Models of Deontic Reasoning

Mike Oaksford, University of Warwick, UK, and
Nick Chater, University of Oxford, UK

Commentary on "Pragmatic Reasoning With a Point of View" by Keith J. Holyoak and Patricia W. Cheng

Holyoak and Cheng (this issue; henceforth "H & C") provide a computational-level analysis (Marr, 1982) of deontic reasoning that corrects a probably too simplistic view of deontic rules (Cosmides, 1989) by introducing important ideas from jurisprudence. This analysis addresses the frequently cited criticism of pragmatic reasoning schema (PRS) theory that it does not account for the selection of the *not-p* and *q* cards in some versions of the thematic selection task. H & C suggest that people possess two PRSs, one the original permission schema from Cheng and Holyoak (1985) and an obligation schema derived from Politzer and Nguyen-Xuan (1992). They observe that these schemas are interdefinable because the antecedents and consequents of the rules that make them up involve rights and duties that are complementary. Rights and duties implicitly introduce two individuals who have different roles—e.g. employer and employee. By focusing on their rights, one of the individuals may interpret a rule as a permission whereas the other may interpret it as an obligation.

Different rules mediate schema access—P3 for the permission schema and O1 for the obligation schema. For a permission schema access via P3 leads to clause order inversion of the P1 rule, turning an *if p then q* rule into an *if q then p* rule. Using the rules in the permission schema with this clausal inversion yields the *not-p*, *q* selection pattern.

H & C also note that their approach is consistent with subjective utility approaches to deontic reasoning (Kirby, 1994a; Manktelow & Over, 1991; Oaksford & Chater, 1994). They point to an important distinction, which we focus on in this commentary, between the factors that lead people to take a particular perspective and the effects of taking that perspective. Consistent with subjective utility approaches H & C argue that taking a particular perspective affects the utilities people assign to the various possible outcomes determined by a deontic rule. We find H & C's computational-level proposals for what gets computed in taking perspectives both compelling and novel. However, H & C also propose an algorithmic-level account of how people make their actual card selections using pragmatic reasoning schemas (PRS) with which we take issue in this commentary.

THREE-STAGE MODEL

H & C's account suggests a three-stage model of the processes involved in deontic reasoning:

(i) *Perspective Assignment*: Determining which perspective to adopt.
(ii) *Deontic Inference*: From the perspective determined in (i) decide which deterministic inferences to make.
(iii) *Inference Withdrawal*: Dependent on the probabilities and utilities assigned in (i), withdraw some inferences made in (ii).

So for example, take H & C's day-off rule:

> If an employee works on the weekend, then that person gets a day off during the week (1)

Employers interpret (1) as:

> If an employee works on the weekend, then that person *may* take a day off during the week (2)

because they focus on their rights in this contractual arrangement. (1) also leads to an assignment of high utility to cases of abuse where an employee takes a day off without having worked at the weekend, i.e. the *not-p*, *q* case. (2) leads people to access the permission schema via P3 which creates deterministic permission rules with appropriately inverted antecedent and consequent clauses. These deterministic rules yield the now familiar *not-p, q* card selections. H & C suggest that the degree to which subjects draw these deterministic inferences will depend on their expected utility. So, if the expected utility of using the rule to guide card-turning behaviour is low then subjects will tend not to turn the card. The expected utility of turning, for example, the *not-p* may be low for two reasons. First, it may be very unlikely that someone who doesn't work at the weekend takes a day off. Second, it may be that although the benefits of detecting violators are high, the costs of false alarms are not small enough for subjects to discount them. For example, it may be important in a company where there is a strong union not to offend innocent workers and hence to avoid cases other than *not p, q*. According to H & C's PRS account, low expected utility of turning the *not-p* would have to lead to the inference based on P4—that this card must be turned over—being withdrawn.

TWO-STAGE MODEL

In this commentary we argue that the three-stage model implied by H & C (this issue, see also Stevenson & Over, 1995) is unparsimonious and that a more parsimonious two-stage model is consistent with Oaksford and Chater's (1994) formal computational-level analysis of the deontic selection task:

(i) *Perspective Assignment*: Determining which perspective to adopt.
(ii) *Maximise Expected Utility*: Calculate expected utilities for choices based on background knowledge and (i).

Reasoners determine the case to assign positive utility according to H & C's account (i), but then base card selection purely on the expected utility calculation. Using H & C's example (1) we now show how Oaksford and Chater's (1994) computational-level analysis can account for card selection after perspective assignment.[1] We label the antecedent of (1) the condition (*con*) and the consequent the action (*act*).

Oaksford and Chater (1994) use two contingency tables to describe peoples' expectations about the probabilities of the various instances. The first represents the contingencies on the assumption that people obey the rule (Table 1a) and the second represents the contingencies on the assumption that people ignore the rule (Table 1b), i.e. they behave as if the condition and action are independent events. People may take more than these two attitudes towards a regulation, e.g. they may assume that there are "negative" people who always do the exact opposite of a regulation. However, Oaksford and Chater (1994) assume that 1(b) is a sufficiently general "catch-all". The parameters of the model are $P(M_I)$, the probability that people are ignoring the rule, hence $P(M_D)$, the probability that people are obeying the rule, $= 1 - P(M_I)$; $P(con)$, the probability of the condition; and $P(act)$, the probability of the action.[2]

TABLE 1
Contingency Tables for M_D and M_I

1(a): M_D	*act*	$\overline{act}$
con	a	0
$\overline{con}$	$(1-a)b$	$(1-a)(1-b)$

(b): M_I	*act*	$\overline{act}$
con	ab	$a(1-b)$
$\overline{con}$	$(1-a)b$	$(1-a)(1-b)$

1(a) shows the table of probabilities appropriate for when the rule is obeyed M_D, 1(b) shows the equivalent table for when the rule is ignored M_I.
a corresponds to the probability of *con*, $P(con)$, and b corresponds to the probability of *act* in the absence of *con*, $P(act|\overline{con})$.

TABLE 2
Utilities of Card Combinations for the Employer and Employee Perspectives

Employer	*act*	$\overline{act}$
con	–0.1	–0.1
$\overline{con}$	5–0.1	–0.1

Employee	*act*	$\overline{act}$
con	–0.1	5–0.1
$\overline{con}$	–0.1	–0.1

Oaksford and Chater (1994) assign utilities to the various instances in the contingency tables in Table 1. Table 2 shows these utilities for H & C's example (1). We assigned a small negative utility (–0.1) to every combination of cards, because of the assumption of a fixed cost for turning any card. For the employer's perspective, we assigned a large positive utility (+5) to finding cases where the *action* occurs but the *condition* is not satisfied. For the employee's perspective we assigned a large positive utility (+5) to finding cases where the *condition* is satisfied but the *action* is not performed. The numerical values are arbitrary—all that is important is that the positive utility is large in comparison to the cost for turning over a card.

Expected utilities (*EU*) can then be calculated for each card using equations (3) to (7)

$$EU(con) = P(act|con)U(con,act) + P(\overline{act}|con)U(con,\overline{act}) \quad (3)$$

$$EU(\overline{con}) = P(act|\overline{con})U(\overline{con},act) + P(\overline{act}|\overline{con})U(\overline{con},\overline{act}) \quad (4)$$

$$EU(act) = P(con|act)U(con,act) + P(\overline{con}|act)U(\overline{con},act) \quad (5)$$

$$EU(\overline{act}) = P(con|\overline{act})U(con,\overline{act}) + P(\overline{con}|\overline{act})U(\overline{con},\overline{act}) \quad (6)$$

Where the conditional probabilities $P(x|y)$ are the expected values calculated with respect to the two contingency tables:

$$P(x|y) = P(x|y, M_I)P(M_I) + P(x|y, M_D)P(M_D) \tag{7}$$

In equations (3)–(6) the expected utility of each card is calculated as the weighted sum of the utilities of each possible outcome given the visible face of the card. The weights are the probabilities of each outcome.

Oaksford and Chater (1994) fixed $P(M_I)$ at 0.5 on the reasonable assumption that subjects are uncertain whether people are obeying or disobeying the rule. They then calculated expected utilities for each card by averaging over all pairs of values for $P(con)$ and $P(act)$ in the range 0.1 to 0.9 at 0.1 intervals. We show these values in Table 3 for the employer's and the employee's perspective, assuming a permission rule of the form *if con, then may act*. Table 3 reveals the same behaviour as that observed by H & C.

H & C argue that for their employee-O1 and employer-P3 conditions subjects interpret the ambiguous task rule as a permission, i.e. *if con, then may act.* Consistently assigning *p* to *con* and *q* to *act*, Table 3 reveals that the following card selections maximise expected utility. In the employer-P3 condition, subjects should select the *not-p* and *q* cards, and in the employee-O1 condition, subjects should select the *p* and *not-q* cards. In the employer-O1 condition subjects are presented with the rule: *An employee must have worked on the weekend if the person takes a day off during the week.* This rule is equivalent to the obligation rule, *if the person takes a day off during the week, they must have worked at the weekend.* Notice that this reverses the clausal order, i.e. this is an *if act* (*q*), *then con*(*p*) rule. But maximising expected utility still involves turning the *not-p* and *q* cards because of the clausal inversion from the permission to the obligation rule. In sum, if subjects are maximising expected utility then Oaksford and Chater's (1994) model predicts just the results H & C

TABLE 3
Average Expected Utilities

Card Face	*(i) Employer*	*(ii) Employee*
act	+1.20	−0.10
$\overline{act}$	−0.10	+2.31
con	−0.10	+2.23
$\overline{con}$	+1.03	−0.10

Average expected utilities for each card face (*action, not-action, condition, not-condition*) (i) the employer's perspective, and (ii) the employee's perspective.

found in their experiment. Thus it appears that both the two- and the three-stage models can explain H & C's results. General criteria of parsimony therefore suggest that we should prefer the two-stage model.

There is also data where expected utility is clearly doing all the explanatory work. Kirby (1994a; see also Over & Evans, 1994; Kirby, 1994b) explicitly varied the probabilities and utilities in the deontic selection task. He used an obligation rule and the equivalent of an employer's perspective. Kirby found that when the cost of a false alarm increased fewer subjects selected the *not-q* card and as the cost decreased more subjects selected the *not-q* card. Moreover, selection of this card decreased as the probability of disobeying the rule decreased. Oaksford and Chater (1994) show that this behaviour is consistent with their optimal data selection model. Importantly the expected utility approach can capture much more of the variation in the data on the assumption that the proportion of subjects choosing a card directly reflects its expected utility. In sum, we can explain all the relevant data without invoking explicit rule-based accounts of deontic inference, i.e. without invoking the second stage (ii) of the three-stage model. Thus although we agree that H & C's computational-level account of perspective assignment provides a better and more comprehensive account than alternatives, we believe that the algorithmic theory based on PRSs about how this affects deontic inference is not needed to explain the data.

Several issues arise from this discussion. First, an account of perspective shifts is largely independent of the PRS account of deontic inference. Although both the two- and three-stage models of deontic reasoning require an account of perspective assignment, they use different processes from then on. Therefore we need not tie theories of perspective assignment to any particular theory of reasoning, be it mental logics, mental models, heuristics, or pragmatic reasoning schemas.

Second, although there have been informal (Manktelow & Over, 1991) approaches suggesting that subjective utility may play an important role in deontic reasoning, recent formal accounts (Kirby, 1994a; Oaksford & Chater, 1994) are crucial to demonstrating the viability of this approach. Oaksford and Chater's formal model is the most comprehensive account, explaining a wide range of data concerning the non-independence of card selections (Pollard, 1985), the negations paradigm (e.g. Evans & Lynch, 1973), the therapy experiments (e.g. Wason, 1969), the reduced array selection task (Johnson-Laird & Wason, 1970), work on so-called fictional outcomes (Kirby, 1994a) and deontic versions of the selection task, including perspective and rule-type manipulations (e.g. Cheng & Holyoak, 1985), and the manipulation of probabilities and utilities in deontic tasks (Kirby, 1994a). It remains to be seen if other proposals such as epistemic utility (Manktelow & Over, 1991; Over & Evans, 1994) can be formalised and shown to provide a better account of the data on deontic reasoning tasks and other versions of Wason's (1966, 1968) selection task.

Third, the three-stage model shares some common characteristics with some proposals of Stevenson and Over (1995). In the context of the conditional inference paradigm, they argue that deductive reasoning and probabilistic reasoning are complementary. They consider the following example:

If John goes fishing, he will have a fish supper.	(8)
If John catches a fish he will have a fish supper.	(9)
John will not have a fish supper.	(10)

Presenting (8) and (10) together (without 9) yields high levels of *modus tollens* inferences to "John did not go fishing" (Byrne, 1989). Adding (9) however, suppresses the number of these inferences that people draw. Stevenson and Over (1995) argue that this is because having "extended our beliefs by performing an instance of ... *modus tollens* ... we [may] get more information ... leading us to doubt the major premise ... [and so] ... we could be led to express doubt about the conclusion." In other words, probability judgements may lead to inferences being withdrawn.

However, the data (Byrne, 1989; Cummins, Lubart, Alksnis, & Rist, 1991; Stevenson & Over, 1995) seems perfectly compatible with a purely probabilistic model. Subjects simply select the conclusion that has the highest conditional probability given the minor premise. We can represent the rule as a probability model as in Table 1(a) and 1(b). So for example, from 1(a) the conditional probability of p given *not-q* is zero, whereas the conditional probability of *not-p* given *not-q* is 1.[3] However, if we allow that the rule is uncertain, i.e. the conditional probability of q given p is less than 1 (i.e. the p, *not-q* cell is non-zero), then these probabilities will vary. In sum, it seems highly likely that we can construct a purely probabilistic model of these data. In consequence, models that include a deductive component and a probability component are unparsimonious—it could be probabilities all the way down!

To summarise, we have argued that two-stage models that do not include a rule-governed or deductive stage of deontic inference provide more parsimonious accounts of deontic reasoning than three-stage models. We agree that an account of perspective assignment is a necessary component of deontic reasoning and that H & C have by far the best account to date of how to achieve this. However, we doubt that H & C's PRS account of deontic inference is necessary to account for the data on deontic reasoning performance.

NOTES

[1]For the general case of deontic reasoning, see Oaksford and Chater (1994), pp.621–625).
[2]$P(act)$ can be calculated from $P(M_l)$, $P(con)$, and $P(act|\overline{con})$ (Oaksford & Chater, 1994, Equation 8).
[3]From Table 1(a) the conditional probability of p given *not-q* $= \frac{0}{(1-a)(1-b)}$ and the conditional probability of *not-p* given *not-q* $= \frac{(1-a)(1-b)}{(1-a)(1-b)}$.

REFERENCES

Byrne, R.M.J. (1989). Suppressing valid inferences with conditionals. *Cognition*, *31*, 1–21.

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.

Cummins, D.D., Lubart, T., Alksnis, O., & Rist, R. (1991). Conditional reasoning and causation. *Memory & Cognition*, *19*, 274–282.

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

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

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

Kirby, K.N. (1994b). False alarm: A reply to Over and Evans. *Cognition*, *52*, 245–250.

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

Marr, D. (1982). *Vision.* San Francisco, CA: W.H. Freeman & Co.

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

Over, D.E., & Evans, J.St.B.T. (1994). Hits and misses: Kirby on the selection task. *Cognition*, *52*, 235–243.

Politzer, G., & Nguyen-Xuan, A. (1989). Reasoning about promises and warnings: Darwinian algorithms, mental models, relevance judgements or pragmatic schemas. *Quarterly Journal of Experimental Psychology*, *44*, 402–421.

Pollard, P. (1985). Nonindependence of selections on the Wason selection task. *Bulletin of the Psychonomic Society*, *23*, 317–320.

Stevenson, R.J., & Over, D.E. (1995). Deduction from uncertain premises. *Quarterly Journal of Experimental Psychology*, *48A*, 613–643.

Wason, P.C. (1966). Reasoning. In B. Foss (Ed.), *New horizons in psychology.* Harmondsworth, UK: Penguin.

Wason, P.C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*, 273–281.

Wason, P.C. (1969). Regression in reasoning. *British Journal of Psychology*, *60*, 471–480.

On the Relationship between Pragmatic Schemas and Mental Logic

David P. O'Brien, City University of New York, USA

Commentary on "Pragmatic Reasoning With a Point of View" by Keith J. Holyoak and Patricia W. Cheng

Holyoak and Cheng (this issue, p.304) noted that "content effects of the sort observed in research on reasoning about regulations clearly lie beyond the scope of existing psychological models of reasoning based on variants of formal logic (e.g. Braine & O'Brien, 1991). In general, proponents of the logic-based