Social roles and utilities in reasoning with deontic conditionals*

K.I. MANKTELOW
D.E. OVER
Sunderland Polytechnic

Received August 23, 1990, final revision accepted January 1, 1991

Abstract

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

A set of experiments is reported in which a new formulation of deontic thinking is tested. This is that people represent subjective utilities inherent in conforming to or violating deontic statements, along with the social dynamics of these statements. The experiments used Wason's selection task and tested people's understanding of conditional permission. In the first two experiments, familiar scenarios referring to family interactions were used. In the third, an imaginary business content was used. In both cases it was apparent that people's thinking depended on their representation of the utilities associated with the agent of a permission statement (the party who lays down the rule) and the actor (the party whose behaviour is its target). The results are discussed as favouring an explanation in terms of mental models, rather than the schema theories which have dominated this field hitherto.

Introduction

Empirical research on conditional reasoning has been primarily concerned, over the last decade, with a single issue: the effects of content. Although content effects on reasoning have been explored for much longer than that, at least as far back as Wilkins (1928), interest in this topic has greatly increased because research on it has suggested an alternative to the traditional view of reasoning. That view, exemplified by Piagetian theory (e.g., Inhelder & Piaget, 1958) and others (e.g., Braine & Rumain, 1983; MacNamara,

*Thanks to Vittorio Girotto and two referees for helpful comments on an earlier draft of this paper. Requests for reprints should be sent to K.I. Manktelow, Faculty of Arts and Social Studies, Sunderland Polytechnic, Douro House, Douro Terrace, Sunderland SR2 7DX, U.K.

0010-0277/91/$6.80 © 1991 – Elsevier Science Publishers B.V.

1986), holds that the human mind contains content-independent inference rules which people apply when making judgements. Content effects obviously pose problems for this view: reasoning cannot be content-independent if it is affected by content. Attempts by inference rule theorists to account for this by proposing content-specific constraints fatally weaken the theory's parsimony and testability (see Evans, 1989; Manktelow & Over, 1987, 1990a).

The alternative is to abandon tradition and to posit the existence of content-dependent mental processes. Two general proposals of this kind have arisen: schema theories, as in Cheng and Holyoak (1985), and the theory of mental models, as in Johnson-Laird (1983) and Johnson-Laird and Byrne (1991). The research reported in this paper has been influenced by both of these particular proposals, as well as by Cosmides (1989). The time has come, we feel, to start to put the pieces together to explain *deontic reasoning*, which is essential to much human thought, but which has been largely ignored by the traditional psychological view. Indeed, it is not too much to say that recent work has taken a great step forward just by bringing out the importance of deontic concepts for cognitive psychology.

Most of this work has depended on Wason's (1968) selection task as its experimental technique. We shall follow this trend here, though it should be stated at the outset that our purpose is not to attempt to explain the task, but to use it, rather in the manner of other paradigms such as the Stroop effect or the memory span, as an investigative instrument. We shall therefore first outline the task, then show how its use in the exploration of content effects in conditional reasoning in general has opened up the study of deontic reasoning.

Wason's selection task, as originally given, concerned four cards and an indicative conditional of the general form, *if p then q*, about those cards. Each card revealed whether the antecedent of this conditional was true or false on one side, and whether its consequent was true or false on the other side. Only one side of each of the four cards was visible to the subjects. In symbols, one card made p visible (showing that the antecedent was true), one made $\bar{p}$ visible (showing that the antecedent was false), one made q visible (showing that the consequent was true), and one made $\bar{q}$ visible (showing that the consequent was false). The subjects were asked to indicate just those cards they would have to turn over to discover whether the conditional was true or false. Given the assumption that the conditional was to be interpreted as a material one, equivalent to *not-p or q*, the logically correct response was to choose the p card and the $\bar{q}$ card. Only these cards might have the falsifying combination of p on one side and $\bar{q}$ on the other.

The early experiments used "abstract" conditionals, such as *If there is a vowel on one side of the card, then there is an even number on the other*, in

which case the logically correct response was, say, a card showing an A along with a card showing a 3. The result was that very few people got the answer right in this kind of case (Wason & Johnson-Laird, 1972; though see Griggs, 1989). However, it was later found that subjects were more successful when certain more realistic contents for the conditionals were used. To cut a long story short, contents which succeeded in facilitating what was supposed to be the logically correct response appeared to be those the subjects had some familiarity with, in some sense. It is the nature of this familiarity which researchers have debated. The issue can be illustrated by using a content which reliably produces high levels of apparently good performance: the "Sears" version of the task devised by D'Andrade (see Rumelhart, 1980).

In this version of the task, subjects imagine they are department store workers supplied with the following rule: *If any purchase exceeds $30, then the receipt must have the signature of the departmental manager on the back.* Four receipts are placed before them, showing, say, $75 ($p$), $25 ($\bar{p}$), a signature ($q$), and no signature ($\bar{q}$). Most people rightly select the $75 and the unsigned receipts as being those which could potentially break the rule. Now, is this because they know about department stores, or generally about *regulations*, a certain type of deontic construction? Evidence from the work of Cheng and Holyoak (1985, 1989; Cheng, Holyoak, Nisbett, & Oliver, 1986) strongly favours the latter.

Before more is said about this topic, there are some very important points which have been obscured in the literature. Sometimes the conditionals in the original, abstract selection tasks were misleadingly called "rules". Those, however, were not guides for correct behaviour; those abstract indicative conditionals purported to describe matters of fact, such as a supposed relationship between vowels and even numbers on the cards. Thus a card with an A on one side and a 3 on the other *falsified* the abstract indicative conditional we have used as an example. In complete contrast, finding an unsigned receipt for $75 does *not* falsify the regulation of the department store: what it does show is that someone has violated, broken, or failed to conform to that regulation. This is a rule in the proper, deontic sense: it is normative in nature and seeks to guide or govern behaviour. In technical terms, this particular rule is a conditional obligation: it states what store workers must do – how they ought to behave – if a purchase exceeds $30. Moreover, we must clearly distinguish, in this kind of case, between the person or institution which lays down the rule, and whoever is supposed to conform to it. Ordinary people must at least have some knowledge of this distinction in order to engage in correct deontic reasoning, and as we shall see, cognitive psychology has to have some account of this knowledge.

Deontic conditional reasoning

Deontic thinking takes place when we consider what we may (or are allowed to) do, or what we ought to (must or should) do, rather than what was, is, or will actually be the case. So basic forms of deontic thought are those concerned with permissions and obligations. Cheng and Holyoak (1985) are responsible for introducing deontic considerations into research on conditional reasoning, a development which has invigorated the field. They propose that there are schemas acquired by induction from experience, and stored in long-term memory, for deontic reasoning about what they take to be permission rules.

They found that subjects produced a high level of correct performance on the selection task when it was designed to evoke what they termed a permission schema; this in turn they proposed to be a subset of a general class of these thought processes, called pragmatic reasoning schemas. The nub of their explanation is that any problem whose content clearly indicates that it is concerned with permission will evoke the permission schema. They hold that the schema's constituent production rules will result in the facilitation of correct responses. For example, in Cheng and Holyoak (1985, expt. 1) there is a rule about which inoculations people must have recorded on cards in order to enter a country:

> If the form says ENTERING on one side, then the other side includes cholera among the list of diseases.

Four cards showed ENTERING (p), TRANSIT ($\bar{p}$), "Cholera, typhoid, hepatitis" (q), and "Typhoid, hepatitis" ($\bar{q}$). The correct response, the p and $\bar{q}$ cards, was produced by almost everyone when a suitable rationale for this rule was given. Crucially, even when much specific content was removed, and the rule referred to "precondition P" and "action A" in a deontic problem, high rates of correct responding were observed. (Even this general content suggested, by bringing in an authority who laid it down, differences in utility between conforming to the rule and violating it, which may be important; see below.)

Cheng and Holyoak's work introduced the first schema theory of deontic reasoning, and generated successfully tested hypotheses, yet it has not gone unchallenged. Cosmides (1989) propounds a different sort of theory of this kind of reasoning. She argues that the understanding exhibited in deontic selection tasks is a product not of inductive learning, but of innate thought processes which she terms Darwinian algorithms.

The kind of innate algorithms concerned in deontic thinking depend, she contends, on the dynamics of social exchange. Understanding the rules of social exchange has been, according to Cosmides, essential to human survival,

and this in turn necessitates a ready appreciation of the benefit–cost relationships in social exchange. The fundamental relationship can be expressed in the following generic rule:

If you take a benefit, then you pay a cost.

Cosmides' case is that everyone understands conditionals of this form, and also, in effect, understands that to obtain a benefit without paying a cost is cheating. She argues that this understanding was adaptive under primitive conditions, and consequently people are innately sensitive to cheaters. Cosmides translates her evolutionary view straightforwardly into the selection task: given a rule of the above form, and a context suggesting the underlying cost–benefit structure, she predicts, and observes, that people readily detect the possibility of cheaters: those who take their benefit (p) without paying a cost ($\bar{q}$). She infers that subjects produce the correct response not by content-independent reasoning, but by looking for cheaters; and that in Cheng and Holyoak's experiments, facilitation of the correct response came about because their contents happened to fit her proposed benefit–cost structure.

Cosmides' thesis has been vigorously contested by Cheng and Holyoak (1989). They concentrate mostly on the last point: Cosmides' claim for the primacy of her benefit–cost relationship in promoting correct solutions. It appears, contrary to this claim, that reliable facilitation of correct responses can be produced by contents which do not have this property. Consider a form of the selection task devised by Manktelow and Over (1990b) and discussed by Cheng and Holyoak (1989). The conditional reads:

If you clear up spilt blood, you must wear rubber gloves.

Using this material in a selection task resulted in a high degree of correct responses, and yet this rule cannot plausibly be seen as relating a benefit (which would have to be clearing up blood in Cosmides' scheme) to a cost (wearing rubber gloves). What seems to be going on here is something more general. Assuming that they are to clear up spilt blood, people would greatly prefer wearing rubber gloves and avoiding a serious disease to not wearing the gloves and getting the disease. It is the significant difference in *utility* between the outcomes of these possible actions which, we hold, leads people to look for cases in which someone clears up spilt blood but does not wear rubber gloves (see Manktelow & Over, 1990a, 1990b).

In our view, Cosmides' work is not sufficiently general, but the work of Cheng and Holyoak, as important as it has been, also has shortcomings, which we shall bring out below. Our main object, however, is not to arbitrate between these theorists; it is rather to avoid some limitations of both approaches. We do this by looking seriously at the semantics of deontic state-

ments, and by taking properly into account the full pragmatic and social background of most deontic discourse. We shall then go on to describe empirical tests of this approach, concentrating in this paper on conditional permissions.

Utilities and social structures in deontic thinking

According to Cheng and Holyoak, certain deontic problems or questions will evoke a particular reasoning schema, consisting of a set of production rules which will determine reasoning behaviour. There are difficulties with interpreting their production rules, especially as these contain deontic terms, like "may" and "must", themselves. It is unclear what role these terms are supposed to play in the postulated mental system based on the schema, unless they are assumed to be logical constants in a kind of mental formulation of some deontic logic. But the first critical point we really want to pursue here concerns their classification or taxonomy of deontic rules.

Cheng and Holyoak (1985) define a "permission rule" as one in which the consequent specifies a precondition which *ought* to be met if an action specified in the antecedent is to be taken (see also Cheng & Holyoak, 1989). The immigration rule above would be an example. But as our emphasis on "ought" in this definition suggests, such a rule is properly called a *conditional obligation*. A proper *conditional permission* has a different logical form; an example of this form is given by the statement that one *may* enter a country if one has been inoculated against cholera. As a matter of fact, no country allows people to enter it merely because they have been inoculated against cholera. A *true* conditional permission would state that people may enter the country if they are citizens of it. With the right kind of pragmatic background, the use of a conditional obligation about a "precondition" can suggest the truth of the conditional permission stating that the action may be taken if this condition is fulfilled. As will become clear below, it is important to stress the significance of pragmatic factors in experiments on deontic reasoning, and to acknowledge that people must have complex representations of semantic, pragmatic, and social facts, in order to reason well deontically. To begin with, however, pragmatic points must be kept separate from semantic ones, and obligation and permission statements, whether conditional or not, do differ semantically.

Although Cosmides' theory has been criticized in its particulars, it does contain a crucial insight which brings it close to the semantics of deontic statements, and which forms no part of the non-semantic theory of Cheng and Holyoak. Cosmides sees the importance of subjective utilities (her benefits and costs) in deontic reasoning. This point seems to be based on her

interesting evolutionary perspective, rather than on a semantic one, but nevertheless it can be given support from the semantic point of view. We cannot go into the technicalities of the semantics of deontic statements here (see Jackson, 1985; Manktelow & Over, 1990a, ch. 4), but consider, informally and in simplified cases, how we might try to determine the truth or falsity of some deontic statements.

Let us suppose we are working in a hospital and want to know whether we *ought to* wear rubber gloves at some point. We imagine a course of action in which we wear the gloves and compare it to one in which we do not wear the gloves. It may be clear that the former action would protect us from a serious disease which the latter action would not protect us from. This would mean that the former action had a much higher utility than the latter, and we would conclude that it was true that we ought to wear gloves in this case. Later as we are about to go home, we may wonder whether we *may* take off our gloves. Then we are likely to infer that taking off the gloves will not have a lower utility than keeping them on, and so it is true that we may take them off. It should be stressed that we are talking about strictly *subjective* utilities here: a person's own beliefs and values may lead to their preferring bare hands in any circumstances, no matter what informed opinion might say. But as in decision theory, these judgements are not to reflect only selfish desires. If the hospital has a regulation requiring the wearing of gloves for everyone's benefit, this will affect our utility judgements. We generally prefer to do what legitimate authorities ask us to.

When we come to consider whether *conditional* obligations or permissions are true, we simply imagine that we are confronted with a possible case in which the antecedent of the conditional holds, for example one in which we are about to clean up spilt blood, and then ask about the utilities of various actions, such as wearing or not wearing gloves, in that case.

The schemas of Cheng and Holyoak do not provide any way of assessing the utilities of possible actions, or representing which outcomes of these actions are preferred to others. In fact, their schemas seem to be more concerned with what it is to conform to a conditional obligation which is already thought of as true, and for that reason alone cannot fully explain what it is to understand and reason with deontic statements. The semantics we have sketched for these statements suggests that an approach based on mental models would be better, but one which radically extends the influential work of Johnson-Laird (1983) and even its recent development in Johnson-Laird and Byrne (1991). Utilities must at least be represented or assessed in these models, if they are to be used in a theory of people's basic semantic understanding of deontic statements. Once we account for this underlying understanding, we can hope to explain why people accept some deontic statements

but reject others, and how they understand the deontic statements other people accept in some scenario or rationale. Only at that point can an account of the difference between conforming to a conditional obligation and violating it have an adequate foundation.

A theory of mental models, or any other approach, has to be even more complex to account for the social and pragmatic factors in deontic reasoning. The importance of these factors for this reasoning is a point which needs to be emphasized. In some contexts, as a result of certain perhaps implicit social agreements, one course of action will be better than another because we would be cheated, in Cosmides' sense, in the latter but not the former. More generally, explicit or implicit laws, rules and regulations provide an essential background, by affecting our subjective utilities, for deontic reasoning in most social settings. As hospital workers, we may place a low utility on certain courses of action because taking these would mean breaking a hospital regulation. We may not want to do that just because we are good professionals, or because we shall get the sack if we are found out.

Further distinctions are necessary, as we shall demonstrate in the experiments we report below, in mental models which are to facilitate deontic reasoning in social settings. A law, rule or regulation is often laid down by an authority we shall term the *agent* (cf. Cheng & Holyoak, 1985). These deontic statements are supposed to guide the behaviour of an individual or group we shall simply term the *actor*. These social roles, the agent and the actor, must be represented separately, since their subjective utilities, the benefits and costs of actions as they see them, can be very different. When an agent lays down a rule with the benefit–cost structure Cosmides describes, the actor would cheat by taking the benefit without paying the cost, but the agent would cheat by extracting the cost without supplying the benefit. More generally, agents, as well as actors, can violate or fail to conform to rules which they themselves lay down.

Suppose an agent utters a conditional permission of the form *You may do x if you do y first*. It might be clear in the context of this utterance that the only way the actor gets permission to do x is by doing y; this would have to be inferred pragmatically. The actor could get some utility from x but not from y, and he might then try to do x without doing y. On the other hand, the agent might get some benefit when the actor does y but not x, and so the agent may try to prevent the actor from doing x after doing y. Thus to understand possible contexts of this deontic form, various elements must be represented: the social roles of agent and actor, their respective subjective utilities, and pragmatic factors. To reason successfully about this sort of deontic context, people would have to be able to keep the agent, the actor, and their respective subjective utilities separate, and have the means to answer

the different questions which could be asked about whether the agent or the actor was failing to conform to the rule in the context.

The purpose of the experiments reported in this paper is to demonstrate that people do display this sort of detailed understanding in their deontic reasoning. We shall do this by concentrating on proper conditional permissions.

Experiment 1

The rule in this experiment was a conditional permission statement said to have been made by a mother to her young son:

If you tidy your room then you may go out to play.

In this kind of case, it is natural to assume, without special background information, that the mother prefers tidy rooms to untidy rooms, and hence would place a higher utility on the boy's possible action of tidying his room than on his leaving it untidy; and that the boy prefers going out to play rather than being kept in, and so would place a greater utility on going out to play than staying in. There are many deontic cases generally like this in which mother and son would accept that the above conditional permission was true, but it also seems natural, without further information, to infer pragmatically, from the utterance of the rule, that the only way the son can get permission to go out is to tidy his room.

In our terminology, the mother is the *agent*, and the son is the *actor*. The four cards in the task show whether the son tidied up his room (p) or not ($\bar{p}$), and whether he went out to play (q) or not ($\bar{q}$). With the assumptions we have made, there are really *four* possible outcomes in which there can be a failure to conform to the rule in some sense:

Case 1. The agent (mother) sees p is true but does not allow q (we might say that she is unfair).
Case 2. The agent does not see p is true but does allow q (she is weak).
Case 3. The actor (son) makes p true but does not make q true (he practises self-denial or unknowingly cheats himself).
Case 4. The actor does not make p true but makes q true (he cheats his mother).

In this experiment, we looked at cases 1 and 4. Case 4 is similar to a Cosmides benefit–cost scenario (the son takes the benefit of going out without paying the cost of tidying his room), but case 1, while undoubtedly a social contract type of scenario, does not strictly fit her scheme, because that does not

explicitly admit the presence of both agents and actors. Our hypothesis was that subjects would be able to model or represent the above state of affairs in what we have described as the natural way, and specifically be able to distinguish in their reasoning cases 1 and 4. Case 1 may arise in a selection task where subjects are asked about cards that might reveal whether *the mother* fails to conform to the rule in the way described. Case 4 may arise in a selection task where subjects are asked about cards that might reveal whether *the son* fails to conform to the rule in the way described there. Therefore, our predictions were that, in a case 1 selection task, subjects would tend to choose just the p and $\bar{q}$ cards, while in a case 4 selection task they would tend to choose just the $\bar{p}$ and q cards.

Method

Two versions of the selection task were constructed using the permission rule given above. One was designed so that subjects would think about the behaviour of the agent (case 1 above), and the other was designed so that they would think about the actor (case 4). The appropriate task instructions were as follows.

Case 1

"Imagine that you are a boy who has been given a rule by your mother about keeping the house tidy while she is at work. The rule is:

If you tidy your room, then you may go out to play.

You have not been seeing eye to eye with your mother lately because you think that sometimes she makes rules like this but doesn't keep to them. So you decide to record on cards what she does on the first few days after she introduces this rule. You record on one side of each card whether you tidied up your room or not, and on the other side whether your mother let you go out to play or not."

The subject was then asked to select which cards would show whether the mother had broken the rule. The four cards showed "I tidied my room" (p), "I did not tidy my room" ($\bar{p}$), "Mother let me go out to play" (q), "Mother did not let me go out to play" ($\bar{q}$).

Case 4

"Imagine that you are a working mother who has given her son a rule about keeping the house tidy while you are out at work. The rule is:

If you tidy your room, then you may go out to play.

Because you're not sure whether your son will obey your rule or not, you ask your neighbour to keep an eye on him for the first few days. You have given her cards on which for each day she records on one side whether he tidies his room or not, and on the other side whether he goes out to play or not."

This time the subject had to select those cards which would show whether the boy had broken the rule. The cards showed "Your son tidied his room" (p), "Your son did not tidy his room" ($\bar{p}$), "Your son went out to play" (q), "Your son did not go out to play" ($\bar{q}$).

Procedure

Subjects were 45 third-year arts and social science students studying psychology at Sunderland Polytechnic. None had prior experience of the selection task. They were given the task during one of their classes. Each subject received one of the two versions, at random. In both versions, four different random orders of the four cards were used on the test sheets. Subjects indicated their choice of cards by placing a tick in an answer box for each card chosen, and a cross for each one they decided not to choose.

Results and discussion

The individual card choices and selection combinations are shown in Table 1. It can be seen that the prediction for case 4 was strongly confirmed (χ^2 for the individual cards = 18.21, $p < .01$); the predicted combination was overwhelmingly the most frequently selected. The results for case 1 were not as clear: there was no difference between individual cards ($\chi^2 = 1.43$), and the predicted combination, while frequent, was not the modal one.

Table 1. *Frequency of individual card selections and selection combinations under two scenarios in Experiment 1 using the rule*: If you tidy your room then you may go out to play

Cards				Combinations								
1. *Son breaks the rule* ($N = 19$)												
p	$\bar{p}$*	q*	$\bar{q}$	$\bar{p}q$*	$q\bar{q}$	$p\bar{q}$	$\bar{p}$	q	All			
2	15	17	4	13	2	1	1	1	1			
2. *Mother breaks the rule* ($N = 26$)												
p*	$\bar{p}$	q	$\bar{q}$*	$p\bar{q}$*	All	$p\bar{p}\bar{q}$	pq	$q\bar{q}$	$\bar{p}q\bar{q}$	$\bar{p}q$	$\bar{p}\bar{q}$	$p\bar{p}$
18	15	15	21	6	8	2	2	2	2	1	1	1

*Predicted patterns.

The case 4 result is unusual in that only once before has this particular response ($\bar{p}q$) been predicted and observed, by Cosmides (1989), using so-called "switched" rules. A switched rule was one which took the form "If you pay a cost then you take a benefit": a cheater is still someone who takes a benefit without paying a cost, but who would in this case be revealed in a $\bar{p}q$ instance. There is some similarity to case 4 here. However, Cosmides' rules did not have the form of a conditional permission, and yet her scenarios are so elaborate and suggestive that subjects might have normalized or interpreted them in that way (Manktelow & Over, 1987, 1990a; see Johnson-Laird & Byrne, 1991, for a similar point). The materials in the present experiments were deliberately designed to be as simple as possible while still displaying the actual form of a conditional permission, to avoid this criticism.

The results for case 1 do not clearly follow the predicted pattern. One possible reason for this is that the procedures for cases 1 and 4 were not equivalent. In case 4, a role for three parties was provided: the mother, the son, and a neighbour recording the boy's conduct. The subject adopted the role of the mother. In case 1, there was no character corresponding to the neighbour, and the subject took the part of the boy examining his own records. It seems more plausible to examine records which someone else has compiled (as with mother and neighbour in case 4), and so a "third party" was introduced in a second version of case 1.

Experiment 2

The scenario was modified to include a third party, as in case 4. The wording of the rule itself was also changed to make it clearer that it was the mother's actions which were under examination. The changed parts of the instructions were as follows:

"Imagine that you have a little brother who has been given a rule by your mother about keeping the house tidy. The rule is:

If you tidy your room, then I will let you go out to play.

Your brother has been getting angry: he thinks that sometimes Mother has been unfair to him and not kept her side of the rule while he has kept his. So he has recorded on cards what she has done over the last few days. On one side of each card he has written down whether he tidied his room or not, and on the other side whether Mother let him go out to play or not."

Procedure

Apart from the modified content, the procedure was the same as in Experiment 1. Subjects were 27 second and third-year arts and social science stu-

dents taking psychology courses at Sunderland Polytechnic. They performed the task in groups of 5–15.

Results and discussion

The individual choices and selection combinations are presented in Table 2. It can be seen that this time the initial prediction that subjects would select the p and $\bar{q}$ cards was strongly confirmed (χ^2 for individual cards = 24.05, $p < .01$), with the predicted combination by far the most frequently selected.

Table 2. *Frequency of individual card selections and selection combinations in Experiment 2 using the rule*: If you tidy your room then I will let you go out to play

Cards				Combinations					
Mother breaks the rule									
p^*	$\bar{p}$	q	$\bar{q}^*$	$p\bar{q}^*$	$p\bar{p}$	pq	p	$q\bar{q}$	All
26	5	5	19	17	4	3	1	1	1

*Predicted patterns. N = 27.

Experiment 2 shows the importance of structuring the task properly so as to give subjects a clear and plausible role in its solution. Overall, apart from this procedural point, the main predictions of these experiments have been strongly confirmed. However, these tests are incomplete in at least two respects. Firstly, we have not looked at cases 2 and 3 in these first two experiments. Secondly, we have used highly familiar content, and this lays these studies open to the charge that the experiments may have demonstrated a simple "memory-cueing" effect, with people just retrieving well-remembered experiences from their own lives. To test the proposition that it is understanding of the deontic context, with its utilities and social roles, which is responsible for these results, rather than the family content, a scenario from outside subjects' actual life experience is needed. Both these points are addressed in the next experiment.

Experiment 3

In this experiment, the same scenario was used for all four outcomes corresponding to cases 1–4 above. Because fewer assumptions could be made about what the subjective utilities would be taken to be, these were stated

explicitly in the task content. One advantage of using less familiar materials was that exactly the same rule could be used in all four cases, thus enabling a purer test of the effects of deontic context. The scenario concerned a shop which was offering a free gift to customers, this offer being expressed in the conditional permission statement:

If you spend more than £100, then you may take a free gift.

The shop is therefore the agent, and the customer is the actor. The four possible outcomes, in which there is some sort of failure to conform to this rule, or what is strongly suggested by it pragmatically, are as follows:

Case 1: The shop takes enough money but does not issue the gifts.
Case 2: The shop takes insufficient money but issues the gifts.
Case 3: The customers pay enough but do not take the gifts.
Case 4: The customers do not pay enough but take the gifts.

The specific predictions for the selection task were therefore:

Case 1: the p and $\overline{q}$ cards; case 2: $\overline{p}$ and q; case 3: p and $\overline{q}$; case 4: $\overline{p}$ and q.

Procedure

Subjects were 52 second-year communication studies and English studies students at Sunderland Polytechnic. None had experience of the selection task. They were randomly divided into four separate groups. A complete set of instructions for case 1 follows; the italicized sections indicate those parts which were changed in the other conditions. These were the only parts of each task which differed from the others. They were not emphasized on the test sheets.

"You are working as a company manager in an industrial town during an economic depression. Times are hard, people don't have much money, and your firm is struggling to survive.

Your firm has devised a way to try to keep its business, by getting people to spend more in its shops in exchange for valuable free gifts. Of course, the firm does not want to give things away, but is willing to try this to increase its sales. Similarly, people don't want to spend more, but the gifts are very attractive.

The firm's offer is expressed in this rule, which is given to customers:

If you spend more than £100, then you may take a free gift.

You have been brought in because you have been told that in one of the firm's shops the offer has run into a problem: *you suspect that the store has not given some customers what they are entitled to.*

You have four customers' receipts in front of you. Each receipt shows on one side how much the customer spent, and on the other side whether they took their free gift. Which of the receipts must you turn over to see *whether this store has not given customers what they are entitled to?*"

For the other conditions, the changed parts of the instructions were:

Case 2: ... you suspect that the store has been giving some customers more than they are entitled to.
... whether this store has given some customers more than they are entitled to?
Case 3: ... you suspect that some customers in this store might not have taken what they are entitled to.
... whether some customers might not have taken what they are entitled to?
Case 4: ... you suspect that some customers in this store might have taken more than they are entitled to.
... whether some customers might have taken more than they are entitled to?

As in the previous experiment, subjects ticked answer boxes on the test sheets for those receipts which they chose to examine, and placed crosses in boxes for those they chose not to.

Results and discussion

The individual selections and card combinations are shown in Table 3. Here it can be seen that the predictions in each of the four cases were strongly confirmed: for individual cards, case 1, $\chi^2 = 14.84$, $p < .01$; case 2, $\chi^2 = 14.63$, $p < .01$; case 3, $\chi^2 = 10.67$, $p < .02$; case 4, $\chi^2 = 19.96$, $p < .001$. In all four cases the predicted combination was the modal choice, with no other receiving more than two choices across all four conditions.

These results strikingly confirm the analysis of deontic reasoning presented here, and demonstrate the importance of systematically accounting for judgements about subjective utility, social roles, and pragmatic factors in explaining this centrally important mode of human thought. People are clearly able to distinguish between social roles and associated subjective utilities in deontic contexts, and infer the answers to appropriate questions about them. A deontic reasoning task, of the type we have described, evidently evokes a natural form of thinking that people can bring to bear even in quite artificial situations, such as a paper-and-pencil task about a scenario with which they have no actual experience. It should be borne in mind that this experiment used exactly the same lexical content through all four conditions in both the central conditional rule and its surrounding story, with the exception of two

Table 3. *Frequency of individual card selections and selection combinations under the four conditions in Experiment 3 using the rule*: If you spend more than £100, then you may take a free gift

Cards				Combinations						
Case 1: The firm cheats the customers										
p^*	$\bar{p}$	q	$\bar{q}^*$	$p\bar{q}^*$	p	$\bar{p}\bar{q}$	p	All		
12	2	1	10	8	1	1	2	1		
Case 2: The firm cheats itself										
p	$\bar{p}^*$	q^*	$\bar{q}$	$\bar{p}q^*$	$p\bar{p}q$	q	$p\bar{p}$	pq		
4	11	12	0	8	2	1	1	1		
Case 3: The customers cheat themselves										
p^*	$\bar{p}$	q	$\bar{q}^*$	$p\bar{q}^*$	$pq\bar{q}$	$p\bar{p}\bar{q}$	$\bar{p}$	p	$\bar{q}$	$\bar{p}q$
10	2	2	10	7	1	1	1	1	1	1
Case 4: The customers cheat the firm										
p	$\bar{p}^*$	q^*	$\bar{q}$	$\bar{p}q^*$	$p\bar{p}q$	pq				
2	12	13	0	11	1	1				

*Predicted patterns.

half-sentences. This apparently slight manipulation was sufficient to trigger the different treatments of the task which were predicted.

General discussion

These experiments have demonstrated that ordinary deontic thought is a richer and more complex aspect of cognition than previous empirical research has shown. In light of these findings, it is necessary to discuss two topics. The first is the status of the facilitation effect in experimental work on deontic reasoning. The second is the question of what kind of psychological theory we should have of deontic reasoning.

The facilitation effect

Most investigators of content effects in conditional reasoning have clung to a simplistic notion of the facilitation effect. This effect was originally said to occur when subjects tended to pick just the p card (showing the antecedent of the conditional true) and the $\bar{q}$ card (showing the consequent false). Such

a response was taken to be the correct one for determining whether the indicative conditionals in the early selection tasks were false. It was assumed that these conditionals should be seen as material ones, and therefore that extensional propositional logic and its semantics should be applied to them. Later, as experiments on realistic forms of the selection task were run, it was not noticed for a long time that the subjects would naturally interpret many of the realistic conditionals, not as indicative, but rather as deontic. This oversight was apparently "facilitated" in the experimenters by their unfortunate tendency to call both indicative and deontic conditionals "rules". Recently, the deontic nature of many selection tasks has started to be recognized, but the old assumptions or presuppositions about the facilitation effect have not been re-examined.

First of all, as we have already pointed out, subjects are not even doing the same thing in a deontic selection task as they are in an indicative one. In the former, they are not trying to point out which cards might reveal whether the conditional is false. A deontic selection task is based on a conditional obligation or permission, and what we say about the truth or falsity of these depends on our preferences among various possible outcomes or states of affairs: these preferences fix the utilities we attach to these outcomes or states of affairs. As far as we are aware, nothing like the selection task has ever contained information *on cards* about preferences or utilities, which would indicate whether a conditional obligation or permission was true or false. It would indeed be very interesting to see what experiments of this general type would reveal. However, the deontic selection tasks which have been constructed so far have all presupposed the truth of the relevant deontic conditionals, or explained why these conditionals are taken to be true in some background story or scenario. What the subjects are, in effect, asked to do in these experiments is to look for possible violations or failures to conform to the deontic conditionals.

The notion of violation, or the weaker one of failure to conform to an obligation or permission, is not a strictly logical or semantic one at all. Propositions are true or false, and arguments are valid or invalid; but only *people*, in particular contexts, can violate or fail to conform to obligations and permissions. The reference to contexts here is important, as there can be good reasons, excuses, or mitigating circumstances at times when people do not follow a rule. These further concepts would also be interesting to investigate in experiments on reasoning. But even as far as present deontic selection tasks are concerned, we should drop the narrow idea of a unique logically correct response, independent of content and context, and speak more broadly instead of various *rational* responses for differing contents and contexts. Return to case 4 of our Experiment 1 as an example. The rational

selections for the content expressed there, in the kind of context described, are surely just the $\bar{p}$ card and the q one: these are the two cards the mother would need to turn over to discover whether her son went out to play without tidying his room. Since that is the response we actually tended to get, we could say that the correct (rational) response was "facilitated" by the wording of the content there, and the specification of the context, for example by bringing in the neighbour as a "third party".[1]

Reasoning by schema or mental model?

Reflecting on the richness of ordinary deontic reasoning leads us to ask whether it can be easily and naturally accounted for in a schema-type theory. Particular questions can be asked, as we have already pointed out, about Cheng and Holyoak's schemas. These consist of sets of production rules which are hard to interpret, primarily because they themselves contain deontic terms. Of course, we cannot have a psychological theory which purports to explain deontic reasoning in terms of unanalysed deontic concepts. One could hold that the conditionals in these schemas are explicitly formulated in some mental syntactic system. This would mean that the schemas were a kind of mental deontic logic, with the deontic terms in them – the mental equivalents of "may" and "must" – appearing as logical constants. But as we have also suggested, this view, whatever its other advantages or disadvantages, would not be general enough. At best, such a mental system might yield inferences about how to conform to deontic conditionals which are already accepted as true. It would not supply a means of finding out which of these conditionals were true, or in some sense acceptable, in the first place, either because of some individual prudential or moral reflections, or because they have been laid down by some appropriate authority. Cheng and Holyoak's schemas contain nothing about preferences, benefits and costs, and utilities, and consequently they do not tell us why an agent would lay down a rule and an actor accept it.

Cosmides' schema theory refers to benefits and costs, and so to utilities, but in a way which is really too simple. Her proposed schema essentially consists of an algorithm which looks for cheaters; that is, people who take

[1] A point worth noting here is that even young children seem to be particularly adept at deontic reasoning, even when it is tested by the selection task, as Girotto and his colleagues have shown (see Girotto, Blaye, & Farioli, 1989; Girotto, Light, & Colbourn, 1988). Indeed, in recent research independent of ours, they have shown that children can also appreciate different forms of failure to conform to deontic rules. They do this by using a distinction between those who make and those who receive promises, a distinction similar to the one presented here between agents and actors (see Light, Girotto, & Legrenzi, 1990).

benefits without paying costs. As such, it equally fails to be a way of inferring which rules or regulations are to be accepted as true. An immigration regulation may state that we must pay the cost of being inoculated against cholera if we are to have the benefit of entering a country. But why do we accept this as true, as a good rule, in the first place? That is the more fundamental question, and its answer should help to explain why we would be so keen to spot cheaters who broke the rule. Again, her theory does not explicitly distinguish between authorities or agents, who have some sort of right to lay down such a rule, and the actors who must follow it if they are not to be called cheaters. Finally, we have explained above our experimental evidence that Cosmides' view is not general enough.

The obvious alternative to a schema theory of reasoning is the theory of mental models, originated by Johnson-Laird (1983), and most recently expounded by Johnson-Laird and Byrne (1991). This theory asserts the existence of general model-constructing and inference-making procedures in the mind, rather than the existence of domain-specific schemas. As far as the selection task is concerned, Johnson-Laird and Byrne propose that without a realistic content people will tend to construct a model which leads them to consider only the cards actually mentioned in the rule (the p and q cards). The standard facilitation effect comes about, they hold, when a content causes the negation of the consequent ($\overline{q}$) to be explicitly represented in the mental model. (It is unclear why this should be accompanied by subjects' ceasing to pick the q card, but that is not a crucial point for the present discussion.)

In the present study, we have seen that, at least as far as deontic contexts are concerned, it can be misleading to talk in terms of the facilitation effect, when that means the choice of only the p and $\overline{q}$ cards. We have constructed deontic contexts in which the choice of only the $\overline{p}$ and q cards is the rational one, and the one we tend to get. Cosmides (1989) reports a curious choice of just $\overline{p}$ and q and cards as a response to so-called "switched" rules, reading "If you pay a cost then you take a benefit", such as "If a man has a tattoo on his face then he eats cassava root". The story in which this rule was embedded made it clear that cassava was a highly desired commodity – people were said to "crave" it. In a selection task, people picked cards representing the "cheater" who ate cassava but had no tattoo. Johnson-Laird and Byrne (1990, ch. 4) argue that subjects do this because they treat this rule as meaning "A man may eat cassava root only if he has a tattoo on his face"; in other words, the subjects are supposed to "switch" the antecedent and consequent back again, and introduce a "may" into the antecedent. (For a somewhat similar but more pragmatic account of these "switched" cases, see Manktelow & Over, 1987.)

Johnson-Laird and Byrne's comment on these "switched" cases does not apply to the present results. There is not even a hint of anything odd, curious, illogical, or irrational about the choice of only the $\bar{p}$ and q cards in the deontic contexts where we have found this response. However, Johnson-Laird and Byrne do make it clear that certain contexts will result in what they term the "fleshing out" of a person's mental models. We would argue that providing an account of this "flesh" is of central importance in explaining deontic thinking: Johnson-Laird and Byrne's theory does not bring out the importance of utility judgements for deontic reasoning, nor the wide range of pragmatic factors which influence this reasoning. What has to be represented in mental models in the present context are not just possible states of affairs, and possible outcomes for various actions, but the preferences people have among these possibilities. These preferences would be the basis for holding that the people attached different utilities to the outcomes of relevant actions they or others might perform. In this way, mental models could be used to explain why people accept certain deontic statements as true. Not only that, but mental models would be a suitable way of distinguishing between the social roles of actors and agents, and of representing what it is to violate or to fail to conform to these statements.

It would be very natural to use mental models to represent the essentially semantic understanding of what makes a deontic statement true. Such models were originally introduced by Johnson-Laird (1983) to account for people's semantic understanding of truth in a possible state of affairs and of the validity of arguments. It could be doubted that mental models, as subjective constructions, could ever fully represent an objective notion of truth in real states of affairs. But few of us think that deontic statements are true in an objective sense. For that to be so, there would have to be objective deontic facts; and however much we may prefer, say as parents, tidy to untidy rooms, we do not think it is a fact of nature, independent of what we think and feel, that the former is better than the latter. It is rather that we need a subjective concept of truth for deontic statements, and that is just what we shall get if we explain this notion in terms of subjective utility and mental models.

Another subjective notion which is important here, and which would benefit from further research, is that of subjective probability. We can be more or less confident that our actions will lead to certain outcomes, which may have more or less utility for us. For instance, we can be more or less confident that we shall contract a serious disease if we clean up spilt blood without wearing rubber gloves. This kind of fact about ourselves can have an immense effect on our deontic reasoning, on what we think we ought to do or may do, and on our consequent decision making. Thus subjective probabilities must also be represented in an adequate theory of deontic reasoning. This

reinforces the view that such a theory cannot be a relatively simple, specialized schematic one, and our preference for a theory based on mental models: it may be possible to give these models weights representing subjective probabilities. It is often assumed in the literature that the study of conditional reasoning fits in neatly somewhere under the broad heading of research on deductive reasoning. But if our view is correct, then the study of deontic conditional reasoning must also reach out and make contact with the vast literature on subjective utility, subjective probability, and decision making.

References

Braine, M.D.S., & Rumain, B. (1983). Logical reasoning. In J.H. Flavell & E.M. Markman (Eds.), *Handbook of child psychology*, 4th Edn. Chichester: Wiley.

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

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

Cheng, P.W., Holyoak, K.J., Nisbett, R.E., & Oliver, L.M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology, 18*, 293–328.

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.

Evans, J.StB.T. (1989). *Bias in reasoning: Causes and consequences*. Hove, Sussex: Erlbaum.

Girotto, V., Blaye, A., & Farioli, F. (1989). A reason to reason: Pragmatic basis of children's search for counterexamples. *European Bulletin of Cognitive Psychology, 9*, 297–321.

Girotto, V., Light, P.H., & Colbourn, C. (1988). Pragmatic schemas and conditional reasoning in children. *Quarterly Journal of Experimental Psychology, 40A*, 469–482.

Griggs, R.A. (1989). To "see" or not to "see": That is the selection task. *Quarterly Journal of Experimental Psychology, 41A*, 517–529.

Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking*. New York: Basic Books.

Jackson, F. (1985). On the semantics and logic of obligation. *Mind, 94*, 177-196.

Johnson-Laird, P.N. (1983). *Mental models*. Cambridge: Cambridge University Press.

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

Light, P.H., Girotto, V., & Legrenzi, P. (1990). Children's reasoning on conditional promises and permissions. *Cognitive Development, 5*, 369–383.

MacNamara, J. (1986). *A border dispute: The place of logic in psychology*. Cambridge, MA: MIT Press.

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

Manktelow, K.I., & Over, D.E. (1990a). *Inference and understanding*. London: Routledge.

Manktelow, K.I., & Over, D.E. (1990b). Deontic thought and the selection task. In K. Gilhooly, M. Keane, R. Logie, & G. Erdos (Eds.), *Lines of thinking: Reflections on the psychology of thought* (Vol. 1). Chichester: Wiley.

Rumelhart, D.E. (1980). Schemata: The building blocks of cognition. In R.J. Spiro, B.C. Bruce, & W.F. Brewer (Eds.), *Theoretical issues in reading comprehension*. Hillsdale, NJ: Erlbaum.

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

Wason, P.C., & Johnson-Laird, P.N. (1972). *Psychology of Reasoning: Structure and Content*. London: Batsford.

Wilkins, M.C. (1928). The effect of changed material on the ability to do formal syllogistic reasoning. *Archives of Psychology, 16*, No. 102.