

## Is there a faculty of deontic reasoning? A critical re-evaluation of abstract deontic versions of the Wason selection task

**Laurence Fiddick**

*Max Planck Institute for Human Development, Berlin, Germany*

Undoubtedly one of the most important studies in evolutionary cognitive psychology is Cosmides' (1989) analysis of content effects on the Wason selection task in terms of adaptive cheater detection. However, in her landmark paper, Cosmides ventured beyond the confines of the selection task to argue for a bold new synthesis of evolutionary biology and cognitive psychology using Marr's (1982) concept of a computational theory as the bridge between these two disciplines. Following Marr's lead, Cosmides argued that cognitive psychology could not make progress unless it was informed by task analyses of the problems the mind was designed to solve. These task analyses would ultimately be supplied by evolutionary theory, which specializes in the study of natural design. Social Contract Theory, Cosmides' specific account of content effects on the Wason selection task, was merely one illustrative part of a larger programme, evolutionary psychology. Unfortunately, the double message of Cosmides' (1989) paper has generated a considerable amount of confusion as some appear to have confused Social Contract Theory as a general theory of reasoning whereas the true scope of the theory is much narrower. In this chapter I hope to dispel some of the confusion surrounding Social Contract Theory and, in the process, clarify where some of the difficulties facing the theory lie and where they do not.

Cosmides (1989) illustrated the benefits offered by an evolutionary perspective by making an example of the then puzzling literature on the Wason selection task. The Wason selection task was originally developed by Peter Wason (1968) as a test of people's ability to logically falsify a hypothesis. By then, the selection task had already generated a large literature, for while it is

disarmingly simple, the selection task has proven to be notoriously difficult for subjects to solve. In the task, subjects are given a conditional rule of the form: "If P then Q", which applies to four cards with information on both sides. One side of each card states whether or not "P" is true, and on the other side of the card it states whether or not "Q" is true. Only one side of each card is showing (depicting the information "P", "not-P", "Q" and "not-Q") and the subjects' task is to determine which, if any, of the four cards they would need to turn over to determine if the rule has been violated. Since the rule is only logically violated when "P" is true and "Q" is false ("P & not-Q"), the solution is straightforward: Subjects need to turn over the "P" and "not-Q" cards since these cards, and only these cards, could potentially be instances of "P & not-Q". Yet, on standard versions of the task, typically fewer than 10 per cent of subjects solve it correctly.

This discrepancy between the apparent simplicity of the task and the overwhelming failure of subjects to solve it correctly has generated a large literature dedicated to figuring out what the source of the difficulty is and what, if anything, can be done to improve performance. Initially, some studies seemed to suggest that when more realistic rules are employed (Wason (1968) had employed the abstract rules such as: "If there is a D on one side of any card, then there is a 3 on its other side") performance on the task improves (e.g. Johnson-Laird, Legrenzi, & Legrenzi, 1972; Wason & Shapiro, 1971); however, later studies employing realistic rules failed to replicate the improvements (e.g. Manktelow & Evans, 1979). Later still there was some indication that prior experience with the rules was crucial to improved performance on the task (Griggs & Cox, 1982), although this interpretation has since been ruled out (Cosmides, 1989). Hence, by the early 1980s there was a confusing pattern of "content effects" whereby versions of the selection task employing some thematic contents but not others inexplicably enhanced performance on the selection task.

### SOCIAL CONTRACT THEORY

It is in this context that Cosmides (1985, 1989) first proposed Social Contract Theory (SCT), formulated on the basis of an evolutionary task analysis of social cooperation—the computational theory of social exchange (Cosmides & Tooby, 1989). Drawing upon prior evolutionary task analyses of reciprocity suggesting that social cooperation among nonkin could evolve only if altruists could detect and punish/exclude cheaters (Axelrod, 1984; Axelrod & Hamilton, 1981; Trivers, 1971), the computational theory of social exchange (Cosmides & Tooby, 1989) hypothesizes that humans possess an evolved "look for cheaters" algorithm that is activated in situations involving the reciprocal exchange of benefits, that is, social exchange. In the Wason selection task, an offer to engage in social exchange can be expressed by a

social contract rule of the form: “If you take the benefit then you must pay the cost”, where a cheater is someone who “takes the benefit” but “doesn’t pay the cost”. Previous studies demonstrating improved logical performance had, according to Cosmides (1989), unwittingly employed social contract rules and triggered the “look for cheaters” algorithm. This mental algorithm would then lead subjects to select the “benefit” (“P”) and “no-cost” (“not-Q”) cards and, thereby, fortuitously make the logically correct selection.

In support of this analysis, Cosmides (1989) constructed parallel versions of the selection task, one in which the context of the rule gave it the cost/benefit structure of a social contract and another lacking the cost/benefit interpretation. As predicted by SCT, subjects selected the logically correct “benefit” (“P”) and “no-cost” (“not-Q”) cards when the rule had the cost/benefit structure of a social contract, but failed to do so when it did not—a result that has subsequently been replicated in numerous experiments (Cosmides & Tooby, 1992; Gigerenzer & Hug, 1992; Platt & Griggs, 1993).

Cosmides (1989) presented SCT as an illustrative example of the results that can be achieved by adopting the evolutionary psychological framework, a metatheoretical framework for conducting all psychological research. Unfortunately, confusion quickly arose as to the precise scope of Cosmides’ claims, with many seeming to interpret SCT—and not the evolutionary psychological framework—as a general theory of the mind, particularly reasoning, and not as a more restricted theory of social cooperation. If it were to be argued that Cosmides proposed a theory of reasoning, *tout court*, then that theory of reasoning would have to be the larger framework, evolutionary psychology, which argues for an adaptive, multimodular view of not only reasoning, but of all mental processes. Yet, the claim is routinely made that SCT, and not evolutionary psychology *per se*, is faulty or at best incomplete because it cannot account for instances of content effects on nonsocial contract versions of the selection task (Almor & Sloman, 1996; Cheng & Holyoak, 1989; Girotto, Blaye, & Farioli, 1989a; Manktelow & Over, 1990; Politzer & Nguyen-Xuan, 1992)<sup>2</sup> with reference often made to Cosmides’

<sup>1</sup> One constraint upon a well-formulated social contract is that the benefits of social cooperation outweigh the costs,  $B > C$  (Trivers, 1971). This constraint is easily satisfied when  $C = 0$ , such that a social contract can be well-formed even when the “cost” is a costless requirement. See Appendix A.1 in Fiddick, Cosmides, and Tooby (2000) for further clarification.

<sup>2</sup> A peculiar corollary is the widely held view that Gigerenzer and Hug (1992) have proposed an alternative theory of deontic reasoning based upon a cheating option that is in some sense an alternative to SCT, where Gigerenzer and Hug’s proposal remains viable even though SCT is falsified (e.g. Almor & Sloman, 2000; Girotto, 1991; Holyoak & Cheng, 1995). This error, I would suggest, is the result of a tendency to view SCT simply as a theory of content effects in reasoning, instead of as a functionalist theory of cooperation. The function of the “look for cheaters” algorithm is specified at the level of the computational theory, in this case the computational theory of social exchange (Cosmides & Tooby, 1989) from which Gigerenzer and Hug’s (1992) theory is derived.

(1989) claim that “robust and replicable content effects are found only for rules that are standard social contracts”. Invariably, this claim is taken as a prediction of SCT, specifically the prediction that people should be competent only at reasoning about social contracts. However, Cosmides was not making a prediction, as the full context of the claim (which is never given) makes clear:

Previous results on the Wason selection task are consistent with a social contract interpretation (for a detailed review, see Cosmides, 1985). Robust and reliable content effects are found only for rules that relate terms that are recognizable as benefits and costs in the format of a standard social contract.  
(Cosmides 1989, pp. 199–200)

Moreover, social contract theory explains the apparently contradictory literature attempting to stalk the “elusive” content effect on the Wason selection task: robust and reliable content effects are found only for rules that are standard social contracts – the only rules for which the predicted social contract response is also the logically falsifying response.  
(Cosmides, 1989, pp. 262)

As these quotes make perfectly clear, Cosmides was summarizing the past selection task literature, and not making a prediction about selection task performance in nonsocial contract domains. While there may be reasonable grounds for objecting to Cosmides’ assessment of the selection task literature circa the late 1980s, SCT cannot be faulted for the failure of a prediction it does not make. Neither SCT, nor the larger evolutionary psychology framework, predicts that that logically correct performance cannot be elicited on nonsocial contract versions of the selection task. Yet many of the arguments against SCT have tended to focus on the theory’s failure to account for enhanced reasoning on nonsocial contract rules.

### Does reasoning about precautions falsify Social Contract Theory?

People do routinely solve certain nonsocial contract versions of the Wason selection task correctly. In particular, numerous studies have demonstrated enhanced levels of performance on precautionary versions of the Wason selection task (Cheng & Holyoak, 1989; Fiddick, 1998; Fiddick, Cosmides, & Tooby, 2000; Girotto, Gilly, Blaye, & Light, 1989b; Love & Kessler, 1995; Manktelow & Over, 1990). Whereas a social contract is a rule of the form: “If you take the benefit then you must pay the cost”, a precaution is a rule of the form: “If the hazard exists then you must use protection”, and all parties to this intellectual dispute agree that precautions do not fall within the scope of SCT (Cheng & Holyoak, 1989; Fiddick, 1998; Fiddick et al., 2000; Girotto et al., 1989a; Manktelow & Over, 1990; Politzer & Nguyen-Xuan, 1992).

In and of itself, these findings are damaging to SCT only to the extent that it is a general theory of reasoning, which it is not. Indeed, one of the most controversial aspects of SCT is the proposal that the reasoning mechanism involved is a module—a reasoning instinct—that is activated by a narrow range of contexts, in this case social relationships involving the exchange or regulation of benefits. It is the very domain-specificity of the proposed mechanism that places precautions beyond the scope of SCT (see Chapter 1). One should ask instead whether the evolutionary psychology framework, more generally, and not SCT specifically, can account for people’s reasoning in nonsocial contract domains. With respect to precaution rules, Fiddick (1998; Fiddick et al., 2000), following Cosmides and Tooby (1992), has argued that people possess additional reasoning instincts for managing hazards. Hence, the real question with precautionary rules is not whether they fall within the domain of SCT—all sides agree that they do not—but whether people reason about social contracts and precautions using a common mental mechanism or whether these rules invoke distinct reasoning mechanisms (Fiddick, 1998).

The claim that enhanced performance on precaution problems argues against SCT merely begs the question for it assumes, but never demonstrates, that a common mental mechanism underlies reasoning about both social contracts and precautions. While it is true that both social contracts and precautions tend to elicit logically correct performance on the selection task, so do abstract rules with negated consequents<sup>3</sup> but few researchers feel compelled to argue that all three types of rules are handled by the same psychological mechanism (Sperber, Cara, & Girotto, 1995, being a notable exception). Likewise, few are tempted to argue that the logical similarity of responses elicited by abstract letter and number rules and nondeontic thematic rules suggests their psychological equivalence. Hence, the logical analysis of selection task performance provides a weak assay of people’s understanding of rules. Still, social contracts and precautions, but not abstract rules with negated consequents, intuitively seem to be related types of rules and the observation that both types of rules elicit similar performance on the selection task is consistent with the proposal that a common reasoning faculty is operative.

### THE DEONTIC ALTERNATIVE

What social contracts and precautions have in common is that they are both deontic rules, rules specifying actions that one is obligated or entitled to perform, and it is this feature of social contracts and precautions that is

---

<sup>3</sup> In a finding that has often been replicated, Evans and Lynch (1973) demonstrated that people routinely provide the correct logical answer to versions of the selection task employing abstract rules of the form: “If P then NOT Q”.

widely held to account for enhanced performance when they are embedded in the Wason selection task (Cheng & Holyoak, 1985, 1989; Cummins, 1996a, 1996b; Girotto, 1991; Manktelow & Over, 1990, 1991, 1995). The most influential of these accounts is Cheng and Holyoak's (1985) Pragmatic Reasoning Schemas Theory (PRST). This proposes that people reason about practical real-world problems using abstract knowledge structures, pragmatic reasoning schemas, that are compiled from personal experiences with different problem domains. One problem domain that people have considerable experience with is rules regulating behaviour: permissions and obligations. A permission, for example, is a rule with the abstract form: "If the action is to be taken then the precondition must be satisfied". Experience with such rules leads to the construction of a *permission schema* composed of four production rules that guides inferences about them. The four rules of the permission schema are the following (Holyoak & Cheng, 1995, p. 291):

- P1: If the action is to be taken, then the precondition must be satisfied.
- P2: If the action is not to be taken, then the precondition need not be satisfied.
- P3: If the precondition is satisfied, then the action may be taken.
- P4: If the precondition is not satisfied, then the action must not be taken.

When the conditional rule employed in a selection task is a permission matching rule P1 of the permission schema: "If the action is to be taken then the precondition must be satisfied", the permission schema becomes activated. Logically correct performance follows as Rule P1 of the schema causes subjects to select the "action to be taken" ("P") card and Rule P4 causes subjects to select the "precondition is not satisfied" ("not-Q") card. Rules P2 and P3 indicate that it is irrelevant that the "action is not to be taken" ("not-P") and "precondition is satisfied" ("Q") cards, respectively. Therefore, the cards corresponding to these conditions need not to be selected. Like SCT, PRST also predicts that social contracts *qua* permission rules will elicit a high level of logically correct "P & not-Q" selections on the Wason selection task. However, unlike SCT, PRST also predicts that precautions, too, will elicit logically correct performance, since they also match the form of a permission rule.

In one of the key tests of the theory, Cheng and Holyoak (1985, Experiment 2) presented subjects with an abstract permission version of the selection task. The rule employed in this problem stated: "If one is to take action 'A', then one must first satisfy precondition 'P'", which matches Rule P1 of the permission schema. Performance with this rule was contrasted with performance on a version of the selection task employing the abstract, nondeontic conditional: "If a card has an 'A' on one side, then it must have a '4' on the other side", which fails to map on to permission schema. As predicted, 61 per cent of subjects correctly solved the abstract permission problem compared

to only 19 per cent correct on the nondeontic problem. This finding has been challenged in studies claiming that the effect is vulnerable to minor alterations in the presentation of the problem (Jackson & Griggs, 1990) or changes in instructions (Noveck & O'Brien, 1996). However, these objections are either orthogonal to or predicted by the PRST, and the effect of these manipulations disappears when more carefully designed problems are employed (Giroto, Mazzocco, & Cherubini, 1992; Kroger, Cheng & Holyoak, 1993).

### ABSTRACT DEONTIC RULES ARE PROBLEMATIC FOR SOCIAL CONTRACT THEORY

The demonstration that subjects can correctly solve abstract deontic versions of the Wason selection task is widely held to be amongst the most compelling evidence in support of PRST, yet in the debate between advocates of SCT and those of PRST, scant attention is paid to the deep difficulty that these rules pose to SCT. Yet the reason these rules pose a difficulty for SCT is straightforward. Subjects' ability to reason successfully about abstract deontic rules directly challenges the assumption that people rely upon more specialized reasoning mechanisms, be they a "look for cheaters" algorithm or procedures for managing hazards. Whereas social contracts and precautions require one to *assume* a common deontic schema, abstract deontic rules would appear to *directly map* onto the proposed schemas. Furthermore, advocates of SCT cannot simply propose an abstract deontic instinct without calling into question the rationale for distinct social contract and precaution mechanisms, since they would thereby be redundant.

Despite the obvious similarities between the abstract deontic rules employed in the selection task and the production rules in the proposed schemas, the importance of these findings depends upon whether subjects are reasoning with the abstract form of the rules as explicitly stated in the problem materials or whether they are reasoning with some other representation of the rules. PRST is only supported, and SCT flawed, to the extent that subjects do, in fact, map the abstract rules directly onto the hypothesized schemas. Although abstract deontic rules bear an obvious resemblance to the hypothesized production rules, it nevertheless remains an untested assumption that these rules are closer to the internal mental representation than are social contracts and precautions.

Cosmides (1989) has claimed that subjects interpret abstract permission rules as social contract rules, thereby triggering a search for cheaters as predicted by SCT. According to Cosmides (1989, p. 239) an abstract permission is implicitly a social contract because:

... saying that one must fulfill or satisfy a precondition in order to be permitted to do something is just another way of saying that one must pay a cost or meet a

requirement . . . In addition, saying that someone is *permitted* to take action A linguistically marks “action A” as a rationed benefit: it implies that the person *wants* to take action A (your mother permits you to get ice cream, she does not “permit” you to be spanked), and it implies that the person doing the permitting has the power to forbid action A (emphasis in the original).

As Cosmides’ (1989) analysis of abstract permission problems suggests, there are at least two ways to interpret people’s reasoning about abstract permissions. Either people could be mentally representing abstract permissions in a manner very close to their surface form (as Rule P1 of the permission schema) and then invoking the appropriate schema; or they could be mentally representing the abstract permissions in a manner more distant from their surface form (as a social contract) and then invoking a “look for cheaters” algorithm. While the basic finding that people reason correctly with an abstract problem is impressive, it is not clear whether people are interpreting the rules as explicit permissions or implicit social contracts.

### How do people interpret abstract deontic rules?

There are two possibilities to be considered: (1) people interpret the abstract permissions as explicitly stated in the problems and map them directly onto the deontic schemas postulated by PRST; or (2) people interpret the abstract permissions as implicit social contracts and thereby activate the “look for cheaters” algorithm postulated by SCT. One way to test between these possibilities is to use the following principle: the closer the surface form of a rule matches the representations processed by a cognitive mechanism, the more likely that mechanism will be invoked. For example, Jackson and Griggs (1990, Experiment 2) found that performance on abstract permission problems decreased when the cards representing “not-P” and “not-Q” used implicit negatives, for example, rather than stating “Has not taken action A”, the “not-P” card read “Has taken action B”. Kroger et al.’s (1993, p. 622) explanation for this result was that: “explicit negatives will make it easier to match the *not-q* case to Rule 4 of the permission schema. Accordingly, removing [this factor] should diminish facilitation for the abstract permission rule”. By the same logic, given two equally abstract rules—both of which are hypothesized to feed into the same mechanism—that which elicits the greatest amount of facilitation should be closer to the form of representation processed by the underlying mechanism.

PRST and SCT propose that abstract permissions are mentally represented in different ways. PRST proposes that the abstract permission: “If one is to take action ‘A’, then one must first satisfy precondition ‘P’”, is mapped onto Rule P1 of the permission schema: “If the action is to be taken, then the precondition must be satisfied”. SCT proposes that the abstract permission is

mapped onto the representation of a social contract: “If the benefit is accepted, then the cost must be paid”. Hence, SCT predicts that more of a transformation needs to be made in order to map the abstract permission onto the underlying mechanism, whereas PRST predicts the opposite—more of a transformation needs to be made to map the abstract social contract. The end result is that PRST predicts that performance on an abstract permission problem will be better than that observed on an abstract social contract problem. SCT predicts the opposite.

### IS A SOCIAL CONTRACT A PERMISSION OR IS A PERMISSION A SOCIAL CONTRACT?

Past studies of abstract deontic versions of the selection task have featured abstract permissions (Cheng & Holyoak, 1985; Girotto et al., 1992; Jackson & Griggs, 1990; Kroger et al., 1993; Noveck & O’Brien, 1996), abstract obligations (Girotto et al., 1992; Jackson & Griggs, 1990; Noveck & O’Brien, 1996), and an abstract precaution (Cheng & Holyoak, 1989). There have been no published studies in which an abstract social contract rule was employed. Hence, it is difficult to judge the relative ease with which people reason about abstract permissions, on the one hand, and abstract social contracts on the other.

I attempted to fill this empirical gap by testing subjects on abstract social contract, permission, obligation and precaution versions of the Wason selection task. As outlined above, PRST predicts that people will perform better when reasoning about abstract permissions than when reasoning about abstract social contracts. SCT predicts the opposite. Similar predictions can be made with respect to abstract precautions and their corresponding abstract deontic rules. PRST predicts that people will perform better when reasoning about abstract permissions and obligations than when reasoning about a corresponding abstract precaution. The more domain-specific view of precautions proposed by Fiddick (1998; Fiddick et al., 2000) predicts the opposite: performance will be higher on abstract precaution problems than on abstract permission and obligation problems since abstract precautions are easier to input into the hypothesized hazard management mechanism. Hereafter, I shall refer to abstract precautions and social contracts collectively as *adaptive* rules, and abstract permissions and obligations as *deontic* rules. In summary, PRST predicts that performance will be higher on the deontic rules and the evolutionary psychology framework<sup>4</sup> predicts that performance will be higher on adaptive rules. I tested these contrasting

---

<sup>4</sup> Comprised, in this case, of SCT and additional account of precautions (Fiddick, 1998; Fiddick et al., 2000).

predictions by giving three groups of subjects, who were naïve to the selection task, abstract versions of the task.

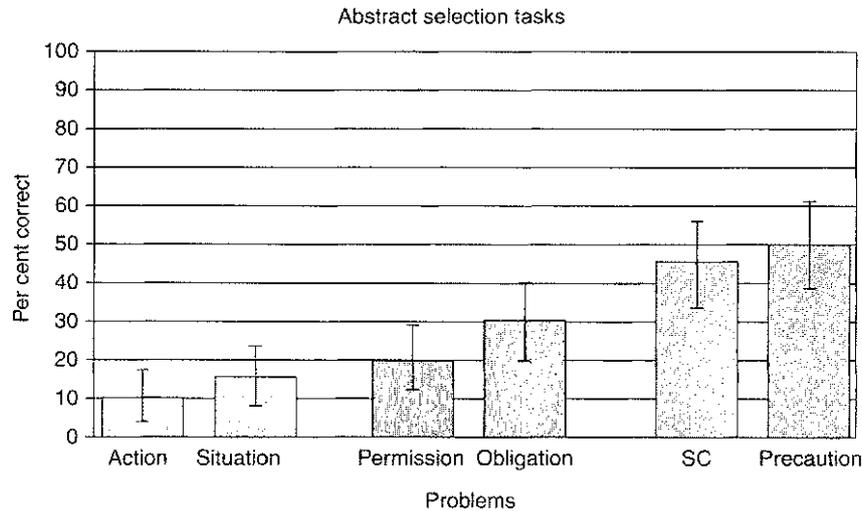
The first group of 20 subjects received two abstract *deontic* versions of the selection task. These were Cheng and Holyoak's (1985, Experiment 2) abstract permission problem and Jackson and Griggs' (1990, Experiment 1) abstract obligation problem. The abstract permission problem featured the rule: "If one is to take action A, then one must first satisfy precondition P", whereas the abstract obligation problem featured the rule: "If situation I arises, then action A must be taken".

A second group of 20 subjects received two abstract *adaptive* versions of the selection task. These were a variant of Cheng and Holyoak's (1989) abstract precaution problem and an abstract social contract problem of my own design. The abstract precaution problem featured the rule: "If one is to take the dangerous action D, then one must have protection P", whereas the abstract social contract problem featured the rule: "If the benefit B is taken, then cost C must be paid". The abstract social contract problem was identical to the Cheng and Holyoak's abstract permission problem other than the change in rule and the necessary changes to the cards and supporting scenario that this required.

Finally, a third group of 20 subjects received two abstract *control* problems to assess baseline levels of performance on the selection task. Where others have contrasted performance on abstract deontic versions of the selection tasks with performance on abstract letter and number versions of the selection task (e.g. Cheng & Holyoak, 1985, Experiment 2; Girotto et al., 1992; Jackson & Griggs, 1990; Kroger et al., 1993; Noveck & O'Brien, 1996), I chose to devise control problems that more closely paralleled the deontic and adaptive versions of the task. Hence, this last group of subjects received selection tasks employing abstract rules that described people's actions. No modal verbs were employed, nor were subjects cued to adopt the perspective of an authority enforcing the rule, but they were still instructed to look for violations. In the abstract *action* problem, subjects read that, "The following rule describes people's behavior: 'If one takes action B, then one takes action A'." The rule in this task was designed to match the abstract permission. In the abstract *situation* problem, subjects read that "The following rule describes people's behavior: 'If situation I arises, then one takes action A'." The rule in this task was designed to match the abstract obligation.

### Do people reason better with abstract deontic or abstract adaptive rules?

As predicted by the evolutionary psychology framework, performance was highest on the adaptive problems (Fig. 2.1) with 50 per cent of subjects correctly solving the precaution problem and 45 per cent correctly solving the



**Figure 2.1** Percentage of correct responses on the abstract selection tasks. The error bars represent one standard deviation above and below the mean. SC, social contract.

social contract (SC) problem. Performance on the deontic problems was midway between that on the adaptive problems and the control problems: 30 per cent correct on the obligation problem and only 20 per cent correct on the permission problem. Performance was lowest on the control problems, 15 per cent on the situation problem and 10 per cent correct on the action problem. On average, subjects did significantly better on the adaptive problems than on the deontic problems. Summing performance across the adaptive problems and the deontic problems, the adaptive problems were correctly solved 48 per cent of the time, compared to 25 per cent for the deontic problems. While subjects performed significantly better on the adaptive problems than their matched controls,<sup>5</sup> there was no significant difference in performance between the deontic problems and their matched controls.<sup>6</sup>

The results of this experiment were in general agreement with the evolutionary psychology framework. Subjects appeared to find it easier to map abstract social contract and precaution rules onto the mechanisms postulated to underlie deontic reasoning. This suggests that the surface form of abstract social contracts and precautions is closer to the input conditions of the underlying mechanisms than is the surface form of abstract permissions and

<sup>5</sup> Abstract social contract versus action ( $Z = 2.48, p < .01, h = .83$ ); abstract precaution versus action ( $Z = 2.76, p < .005, h = .93$ ). All  $Z$  scores reported are the result of a test of proportions (Blalock, 1972).  $h$  is the effect size for a test of proportions:  $h = .20$ , small effect;  $h = .50$ , medium effect;  $h = .80$ , large effect (Cohen, 1977).

<sup>6</sup> Abstract permission versus action ( $Z = 0.89, p > .10, h = .28$ ); abstract obligation versus situation ( $Z = 1.14, p > .10, h = .36$ ).

obligations. Consequently, it is problematic to conclude that subjects in previous studies were mentally representing abstract permissions and obligations in a manner similar to their surface form. The purportedly strong evidence in favour of PRST is, at best, ambiguous in its support for the theory.

### EMOTIONS AS A CUE TO INTERPRETATION

Although the previous experiment raised some doubt about whether people interpret abstract permissions and obligations explicitly as stated in the problem materials, it sheds no positive light on how people do, in fact, interpret these rules. The results provide little direct evidence in support of Cosmides' (1989) claim that people interpret abstract permissions as social contracts. Again, the cards that people select on the Wason selection task is a poor guide in determining how people interpret a rule. Consider, again, social contracts and conditionals with negated consequents. They both routinely elicit "P & not-Q" selections on the Wason selection task, but this provides little reason for believing that people are interpreting both types of rules in the same way. Hence, the fact that both abstract social contracts and abstract permissions elicit "P & not-Q" selection is no guarantee that people reason about both types of rules using the same mental mechanism. Although unlikely, people could be interpreting the abstract permissions as precautions, for example, which would also lead them to select the "P & not-Q" cards. Less ambiguous evidence is required to substantiate Cosmides' proposal.

Emotional reactions to rule violations can potentially provide an alternative means of categorizing deontic rules. Rozin, Lowery, Imada and Haidt (1999) have found that different moral codes are associated with different emotions. One of the methods they used was to give people descriptions of moral violations and have them select which facial expression a person would show if that person witnessed the violation. This same method could potentially be employed to assess how people interpret abstract permissions and obligations.

Indeed, Rozin et al.'s method can be used to dissociate social contracts from precautions. Using a modification of this method for the Wason selection task, I have found that people associate different emotions with violations of social contracts and precautions (Fiddick, unpublished). Subjects were presented with a rule embedded in a story, as with the Wason selection task. However, rather than instructing subjects to select cards representing potential violations, they were informed that the rule has been broken and their task was to indicate who, among an array of four people, saw the violation occur. The faces of the people varied in the emotions that they expressed. The results of this study indicated that people associate different types of rules with different emotions. In principle, the same method could easily be applied to abstract rules providing a converging line of evidence for

the psychological equivalence of social contracts and permissions. Besides eliciting logically identical performance on the selection task, abstract social contracts and permissions might also elicit the same emotional reactions when they are violated.

The emotions associated with the rule violations were anger for social contracts and fear for precautions. This pattern of emotional reactions is readily explained from an evolutionary point of view. Consider first social contracts and anger. The goal in detecting cheaters, according to SCT, is to be able to punish or exclude them. It is this element of punishment that makes anger a suitable emotion to express, for as Lazarus (1991, p. 225) describes anger:

I suggest that anger, in contrast with fright and anxiety, is potentiated by an appraisal that the demeaning offense *is best ameliorated by attack*; in effect, the individual *evaluates her coping potential of mounting an attack favorably*, which is also the innately given action tendency (emphasis added).

In short, anger is the emotion that is elicited when one seeks to punish others for a harmful wrongdoing. Consider, on the other hand, precautions and fear. According to the evolutionary analysis of precautions (Fiddick, 1998; Fiddick et al. 2000), the goal is to manage hazards. It is this element of avoiding or preventing injury that makes fear a suitable emotion, for as Lazarus (1991, p. 238) states: “In both fright and anxiety, the action tendency is avoidance or escape, in contrast with approach or attack”. While violations of both social contracts and precautions may involve some form of loss or harm, the assignment of blame is an important component of cheater detection that is typically absent from hazard management. Likewise, blame is an important component of anger, but is absent from fear (Lazarus, 1991). In short, there are good functional grounds for predicting that anger will be associated with violations of social contracts and that fear will be associated with violations of precautions.

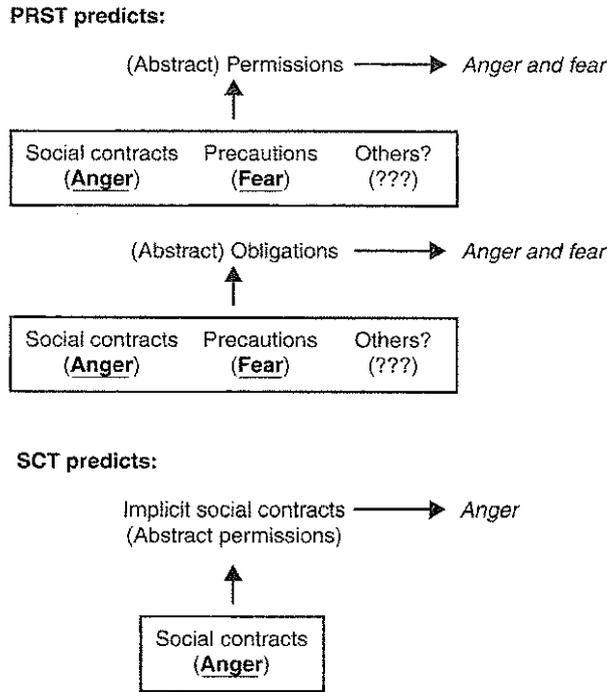
Moreover, this same line of argument suggests that there might be different types of precautions. There are grounds for distinguishing precautions against physical injury, from precautions against infections and social aggression. Whereas the threat of physical injury is typically associated with fear, the threat of infectious contamination is more closely associated with disgust. Similarly, Lazarus (1991) has argued that it might be best to distinguish between “fright” as elicited specifically by the threat of physical injury and “fear”, which may include the threat of social aggression. Intuitively, these three types of hazards are quite distinct and may impose different computational demands, but whether there are characteristic differences in prudential reasoning about these three domains of hazards will have to remain an open question for future research. For now I will simply assume that precautionary rules employed in these studies are interpreted in terms of physical injury.

Instead, I investigated which emotions are associated with abstract permissions and obligations by devising some abstract adaptive and abstract deontic versions of the emotion selection task. I presented 21 subjects with four abstract versions of the emotion selection task. These included an abstract social contract task, an abstract precaution task, an abstract permission task, and an abstract obligation task. The tasks repeated the rules and stories employed in the Wason selection tasks in the previous experiment, but without the cards and their supporting statements. In their place were the pictures of four faces with each face depicting a different emotion. The four emotions were anger, disgust, fear, and happiness. The pictures were scanned, greyscale reproductions from Matsumoto and Ekman's (1988) Japanese and Caucasian Facial Expression of Emotion (JACFEE) slides, however, only Caucasian faces (of both sexes) were used. The accompanying story explained that "Recently someone observed the rule being broken. Indicate the person, who you think, saw the rule being broken".

#### What do PRST and the evolutionary framework predict for the deontic rules?

Although neither PRST nor SCT make explicit predictions about the emotional reactions that people will have in response to deontic rule violations, deducing the predictions that they would have to make given previous findings with adaptive rule violations is fairly straightforward (Fig. 2.2). Recall that violations of social contracts elicit angry reactions whereas violations of precautions elicit fearful reactions (indicated in the boxes in Fig. 2.2). According to PRST, permissions and obligations subsume both social contracts and precautions. This suggests two alternative predictions. The first alternative assumes that people mentally translate permission and obligations into social contracts and precautions before completing the task (this is depicted in Fig. 2.2 under the heading "PRST predicts"). Under this scenario, violations of permissions and obligations should elicit a mixture of anger and fear (the predicted emotions are given in italics). Some people will interpret the rules as social contracts and select the angry face, while others will interpret the rules as precautions and select the fear face. However, subjects may be more confused on the abstract obligation task because the antecedent, "situation I arises" suggests no clear interpretation as either a benefit or a danger. Should such confusion arise, one might predict a random pattern of responding on the abstract obligation task.

Indeed, subjects might fail to mentally translate either deontic rule into a social contract or precaution and instead try to complete the task on the basis of the abstract forms of the rules. Under this second alternative, PRST predicts that people will be confused, or otherwise unable to complete the task in any principled manner, and will select one of the four emotions at random for



**Figure 2.2** Emotions associated with abstract deontic rules according to PRST and SCT. Emotions that have previously been demonstrated to be associated with social contracts (anger) and precautions (fear) are underlined and in bold. Emotions predicted to be associated with abstract permissions and obligations are indicated in italics.

both the permission and the obligation rules, due to the very abstract nature of the rules. Since this basically amounts to random behaviour, I have not modelled this option in Fig. 2.2.

PRST cannot, however, argue that people will have a default tendency to associate abstract permissions with anger. Not only would this concede Cosmides’ (1989) claim that people interpret abstract permissions as though they were social contracts, but it would also weaken the claim that precautions are readily interpreted as permissions since these rules would then evoke different emotions—a distinction that would not appear to be warranted by PRST.

SCT predicts a different pattern of performance (see Fig. 2.2 “SCT predicts”). According to Cosmides’ analysis of abstract permissions, people interpret these rules as social contracts, so there should be a clear preference for selecting the angry face. However, SCT offers no clear predictions for performance with an abstract obligation, nor does a theory of hazard management (Fiddick, 1998; Fiddick et al., 2000). The problem, suggested above,

is that the abstract obligation offers no clear interpretation as either a social contract or a precaution.

### What did subjects do?

In a replication of previous findings, a majority of subjects selected the angry face for the abstract social contract problem (67 per cent of selections) and the fear face for the abstract precaution problem (86 per cent of selections, Fig. 2.3). The interesting question, though, is what emotions subjects selected for the abstract deontic rules. Did they select a mixture of emotions as predicted by PRST or did they select a single emotion—anger—as predicted by SCT, at least with respect to the permission rule? The modal response of the subjects was to select the angry face for both the abstract permission problem (43 per cent of selections) *and* the abstract obligation problem (48 per cent of selections). The next most frequent selections were the disgust face for the abstract permission (29 per cent of selections) and the fear face for the abstract obligation (24 per cent of selections). The remaining emotions were selected by less than 20 per cent of the subjects for both deontic rules.

This experiment provides additional support for SCT's account of abstract permission rules. Not only did the emotions selected on the abstract permission problem match that predicted by SCT, but the percentage of subjects selecting the angry face for the abstract permission problem was slightly lower than that observed on the abstract social contract problem. This is precisely the pattern that would be expected if people interpreted the abstract permission as an implicit social contract.

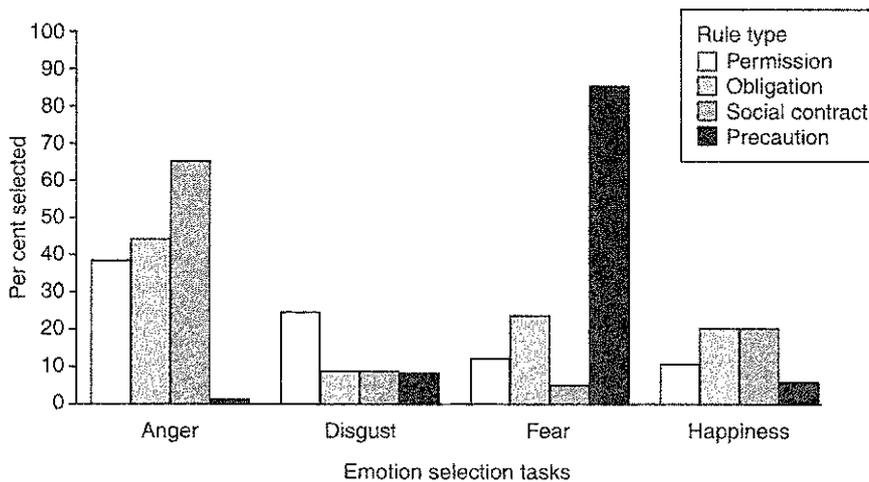


Figure 2.3 Percentage of participants selecting an emotion for each of the emotion selection tasks.

contract is a better exemplar of a social contract than the abstract permission, the former should elicit a clearer pattern of emotion selections than the latter.

Neither of PRST's alternative set of predictions was supported. Subjects were not guessing when they completed the deontic problems (ruling out the second alternative, performance would be random) and they showed a clear preference for selecting the anger face for both deontic rules (ruling out the first alternative, a bimodal response pattern for both rules).

Although the results of this experiment disconfirm PRST and lend support to Cosmides' (1989) claim that people interpret abstract permission rules as social contracts, it is not clear whether Cosmides' analysis is entirely correct. The problem is that a slightly higher percentage of participants selected the angry face on the abstract obligation problem than was observed on the abstract permission problem, but the abstract obligation lacks precisely those features that Cosmides invoked to interpret the permission as a social contract. For example, while it might be plausible to suggest that one is "permitted" to do things one considers a benefit, it is less straightforward to assume that a "situation arising", and presumably beyond one's control, is a benefit that obliges one to take an action. This is not to deny that participants are interpreting the abstract permission rule as a social contract. I merely wish to suggest that participants might be interpreting both the abstract permission and the abstract obligation as social contracts for some reason other than the one provided by Cosmides.

One possible alternative explanation is a cultural bias in the West to interpret rights and duties in contractual terms (Shweder, Mahapatra, & Miller, 1987). Given rules stating vague permissions and obligations, the subjects in this study, who were all highly educated Westerners, might invoke a default cultural bias to interpret deontic rules in contractual terms. Such a bias would lead a sizeable proportion of subjects to view both abstract permissions and abstract obligations as contracts.

### WHAT IS THE SIGNIFICANCE OF THESE RESULTS FOR THE STUDY OF DEONTIC REASONING?

While there are major differences in theoretical explanations of deontic reasoning (Almor & Sloman, 1996; Cheng & Holyoak, 1985; Cosmides, 1989; Cummins, 1996a; Manktelow & Over, 1991; Oaksford & Chater, 1994; Sperber et al., 1995), recent studies on abstract deontic rules have done little to assess the relative merits of these different positions. For the most part, these studies have either focused on theoretically trivial manipulations, such as whether "not-P" and "not-Q" are stated explicitly or implicitly on the cards (Giroto et al., 1992; Jackson & Griggs, 1990; Kroger et al., 1993); or they have focused on manipulations already known to have an effect on reasoning

about deontic rules, such as whether one is instructed to look for violations of an existing rule or instructed to look for violations in order to establish the rule's existence (Noveck & O'Brien, 1996; cf. Gigerenzer & Hug, 1992). In general, these studies appear to have been motivated by an attempt to either undermine or support what is taken to be some of the best evidence in favour of PRST, but without, thereby, advancing any alternative account of deontic reasoning. No studies that I am aware of have attempted to use abstract rules to decide among rival accounts of deontic reasoning. This is an oversight because the demonstration that people reason correctly on abstract deontic versions of the selection task is highly problematic for some rival accounts of deontic reasoning.

The findings presented here do not simply end with a negative verdict against PRST, they also give positive support for a specific class of deontic reasoning theories. The abstract social contract and abstract permission rules that I have employed here are considered to be paraphrases by all who accept the psychological reality of deontic reasoning, so the question inevitably arises: Which version is paraphrastic and which version is primitive? Of course, neither rule may be psychologically primitive—both may be tokens of some other type of rule—but until an alternative candidate is proposed, the current state of the field can be roughly divided in two. On the one side there are those who assume that deontic concepts are psychologically primitive (Cheng & Holyoak, 1985; Cummins, 1996a) and on the other there are those who assume that deontic concepts can be further decomposed into costs and benefits or other more specific terms (Cosmides, 1989; Manktelow & Over, 1991; Oaksford & Chater, 1994). The results presented here would appear to support the latter position, but not everyone accepts the psychological reality of deontic reasoning—at least not as evidenced by the Wason selection task.

### Nondeontic accounts of abstract deontic rules

Some researchers have questioned the need to postulate domain-specific reasoning mechanisms to account for people's reasoning on abstract deontic versions of the Wason selection task (Almor & Sloman, 1996; Sperber et al., 1995). These researchers would not necessarily deny that people possess a deontic reasoning competence, only that it has not been demonstrated by people's performance on the selection task. To bolster these claims, they have presented evidence that it is possible to elicit high levels of "P & not-Q" selections on abstract nondeontic versions of the selection task. Given that the interpretation of the results that I have adopted here assumes that some more specialized forms of reasoning underlie performance on abstract deontic versions of the selection task, it is worth considering these claims in more detail.

Almor and Sloman (1996), for example, have demonstrated enhanced levels of “P & not-Q” selections with abstract nondeontic versions of the Wason selection task employing rules such as: “If a large object is stored, then a large container must be used”. Almor and Sloman (1996, 2000) interpret their results as being problematic for the deontic reasoning theories because it undermines what they take to be a prediction of the theories, that only deontic rules will elicit high levels of “P & not-Q” selections. As foreshadowed in my discussion of SCT and precautions, this is not a general prediction made by SCT, nor is it a prediction made by other theories of deontic reasoning. As is the case with SCT and precautions, the error is in taking the aims and claims of the deontic theorists out of context. The theory’s claims were largely retrodictive in attempting to account for the differences between those rules that *had* and those rules that *had not* elicited high levels of normatively correct performance on the Wason selection task—a distinction that had roughly corresponded to the difference between deontic and nondeontic rules when the theories were proposed.

No deontic theory has proposed that nondeontic rules would always fail to elicit correct performance on the selection task. Any number of theoretically trivial manipulations, such as negating the consequent or explicitly prompting subjects to look for instances of “P & not-Q”, could conceivably improve performance on nondeontic versions of the selection task. The real question is whether or not people reason differently about deontic and nondeontic rules regardless of whether nondeontic rules elicit “P & Q” selections or “P & not-Q” selections. Of course, logically different performance on the selection task constitutes *prima facie* evidence for different reasoning processes, but even so there is no guarantee that different reasoning processes are at play. Indeed, the “matching bias” explanation for the effect of negated consequents proposes that there really is no difference in subjects’ reasoning (or lack thereof) despite the logical difference in performance with standard conditionals. Subjects are simply ignoring the negation and selecting the same cards that they would with non-negated consequents, with performance guided in both cases by the topic of the rule (Evans, 1989). What is required to undermine the deontic theories is not simply to demonstrate logically identical performance on both deontic and nondeontic versions of the selection task, but to provide a convincing argument that performance on both versions of the task is guided by the same principles.

Almor and Sloman’s (1996, p. 379) proposed explanation is that:

According to our intuitions—the same intuitions that have guided us in constructing the problems we used—people’s performance in the selection task is governed by their beliefs about the dependence relations in the problem. This

belief causes them to form probabilistic expectations that they can test as they select cards.

But it borders on a tautology to claim that in a test of *conditional* reasoning beliefs about dependence relations governs subjects' performance. Nor does it provide any further insight to note that when subjects' selections conform to the normative interpretation of the rule as the material conditional that relevant expectations are those that "happen to have a structure like conditional implication" (Almor & Sloman, 1996, p. 375). Other than demonstrating that people can solve nondeontic versions of the selection task at levels comparable to those found with deontic versions, Almor and Sloman provide no independent confirmation or argument that people do have these expectations when solving both deontic and nondeontic versions of the selection task.

A more promising line of explanation is provided by Sperber et al.'s (1995; see also Liberman & Klar, 1996; Love & Kessler, 1995) Relevance Theory account of the selection task. Sperber et al. present a detailed account of the principles guiding performance on both deontic and nondeontic versions of the selection task. The key to enhanced performance on all versions of the selection task is whether or not the conversational pragmatics of the task: (1) make violating instances of "P & not-Q" manifest (although not necessarily explicitly); and (2) highlight the relevance of detecting violations.

The device that they employ to elicit correct performance on nondeontic versions of the task is to place the conditional rule in the context of a denial. When it has been alleged that "P & not-Q" is true, one can deny the claim by stating: "If P then Q", which is logically equivalent to claiming "It is not the case that P & not-Q". Most importantly for present concerns, Sperber et al. (1995, Experiment 4, The Machine Problem) demonstrated this with an abstract letter and number task. In one condition of their experiment it was alleged that a machine has mistakenly produced cards with the features "6-on-the-front & not-E-on-the-back". A repairman then fixed the machine, denying that the machine is still malfunctioning by stating: "If a card has a 6 on the front, it has an E on the back". As predicted, subjects performed well, with 57 per cent correctly solving the task. Further strengthening their claims to have isolated the relevant variables, when they systematically removed elements of the scenario that highlighted violating instances of "P & not-Q" and the relevance of detecting violations, they observed that performance decreased accordingly.

Although providing an elegant and convincing demonstration of their proposal with respect to nondeontic versions of the selection task, this experiment does not in itself provide a convincing demonstration that the same principles are at play in deontic versions of the selection task. While

the deontic theories agree with Sperber et al.'s analysis in suggesting that there is some practical utility to detecting deontic rule violations,<sup>7</sup> they part ways with the Relevance Theory account over the assumption that the context needs to highlight violating instances. The question is really how much structure must be supplied by the scenario and, conversely, how much can be supplied by knowledge structures in the mind of the reasoner. The deontic theories assume that a large amount of the relevant structure is built in to the mind of the reasoner, while the Relevance Theories assume that the supporting scenario supplies almost all of this structure.

Consider, in this light, Cheng and Holyoak's (1989, p. 289) abstract precaution scenario, a close variant of which was employed in the experiments reported here:

Suppose you are responsible for ensuring whether people who are about to engage in certain hazardous activities have taken the precautionary measures necessary for protecting them from harmful effects inherent in those activities. The precautions take the general form: *If one is to take the hazardous activity H, then one must have protection P*, where H is any hazardous activity and P is the appropriate protection for the particular activity.

While it is apparent from this scenario that there would be some practical utility to detecting violations of the rule, neither the rule nor the surrounding context highlights "P & not-Q" as violations of the rule. Contrast this with Sperber et al.'s "machine problem" where the story stated that: "the machine has printed cards it should not have printed. On the back of the cards with a 6, the machine has not always printed an E: sometimes it has printed an A instead of an E" (Sperber et al. 1995, p. 75). High levels of violation detection result in both cases, but it requires more coaxing on the part of the experimenter in the case of nondeontic versions of the selection task.

Expanding upon Sperber et al.'s own metaphor, it would appear that it is possible to cook up nondeontic versions of the selection task that elicit high levels of violation detection by following their proposed recipe, but following such a strict recipe is not required in the case of deontic versions of the task. When the supporting devices employed by the Sperber et al.'s (1995) "relevant" versions of the selection task are removed, the deontic content effect still persists (Fiddick et al., 2000). Moreover, many of these relevance cues are generally absent from abstract deontic versions of the selection task, so it is not surprising that Relevance Theoretic accounts of the selection task (e.g. Liberman & Klar, 1996; Love & Kessler, 1995; Sperber et al., 1995) have

---

<sup>7</sup> Strictly speaking, Sperber et al. (1995) argue that the detection of violations has to have some degree of "cognitive effect". However, this concept of cognitive effect is rather vaguely defined and does not seem to rule out cognitive effects derived from detecting other outcomes like mutual cooperation (see Fiddick et al., 2000, for further discussion).

virtually ignored abstract deontic versions of the task in their account of the “deontic content” effect.<sup>8</sup> Hence it would appear that there is a genuine deontic content effect worth accounting for. Subjects are competent at detecting deontic rule violations over a wide range of pragmatic contexts, whereas with nondeontic versions of the selection task, various *ad hoc* contextual elements need to be incorporated before substantial levels of violation detection are elicited.

The brittleness of the Relevance Theories is further underlined by the emotion study presented here. There are, as I have suggested above, good evolutionary grounds for predicting that anger will be associated with violations of social contracts and fear will be associated with violations of precautions. The evolutionary analysis of deontic reasoning was, therefore, able to provide some guidance beyond the Wason selection task, but it is not clear if or how the Relevance Theoretic perspective can provide any guidance in making predictions for the emotion study. The finding that different emotions are associated with social contracts and precautions not only reinforces the view that social contracts and precautions are psychologically distinct, but it further suggests that the psychology of cooperation and hazards is far richer than the simple dichotomy of “reasoning” versus “interpretation” favoured by the Relevance Theorists.

The modular view of mind espoused by evolutionary psychology does not only entail domain-specificity, for example, the psychological dissociation between social contracts and precautions. It also suggests the “vertical”, as opposed to “horizontal”, organization of the mind in which special-purpose representations, inferences, memories, and even emotions are all integrated to bring about adaptive functioning (Fodor, 1983; Gigerenzer, 1995; Tooby & Cosmides, 1990). In putting the disparate pieces of the functional puzzle together, a considerable degree of constraint is placed upon enumeration of psychological adaptations, potentially solving evolutionary psychology’s “grain problem” (see Chapter 3). Hence, the parallel dissociation between social contracts and precautions in terms of both elicited emotions, as presented here, and reasoning processes, as studied previously (Fiddick, 1998, unpublished; Fiddick et al., 2000), suggests that the proposed characterization of these rules and the adaptive problems they map onto is couched at the right grain of analysis.

In contrast, while the Relevance Theoretic approach is not only plausible but may even have some rigour in the context of one highly constrained task, such as the selection task, these accounts threaten to degenerate into the trivial claim that people act on the basis of their representation of the materials provided as additional methodologies demonstrate systematic

---

<sup>8</sup> For a more extended defence of SCT against the relevance theoretic accounts of the selection task, see Fiddick et al. (2000).

dissociations in the psychology of cooperation and hazards. Hence, nonevolutionary accounts of reasoning are likewise susceptible to the grain problem.

The constraints of vertical integration also come to the fore in the assessment of some other deontic theories that I have not directly addressed so far: Cummins' (1996a, 1999) Dominance Theory and assorted Decision Theoretic accounts of deontic reasoning (Kirby, 1994; Manktelow & Over, 1991, 1995; Oaksford & Chater, 1994).

### Dominance Theory and deontic reasoning

Aligned on the side of PRST is Cummins' (1996a, 1999) Dominance Theory. Where PRST claims that permissions and obligations are rules imposed by an authority for a social purpose, Cummins gives the theory an evolutionary gloss by proposing that social living within dominance hierarchies has shaped the human mind (and that of other species as well) for reasoning about the "social code"—rules imposed and enforced by high-ranking individuals to maintain their priority of access to fitness-enhancing resources. In contrast to PRST, which proposes that deontic reasoning schemas are compiled from extensive ontogenetic experience with social rules and regulations, Dominance Theory hypothesizes that the deontic concepts of permission, obligation and prohibition are innately built-in to the human mind, enabling fast-track learning of the rules and regulations of one's community. Cummins' proposal highlights the diversity of opinion found within the evolutionary psychology community, even within a narrowly restricted topic such as deontic reasoning. Besides tracing deontic reasoning to a different set of selection pressures than does SCT, Cummins also places a greater emphasis on a comparative approach than is found within Cosmides and Tooby's programme.

Given the close affinities between Dominance Theory and PRST, the former suffers the same problems as the latter in accounting for the findings presented here. Namely, the theory has a difficult time accounting for why subjects perform better on the abstract adaptive problems than the abstract deontic problems. However, it is the results of the emotion study that pose the greatest difficulty for Dominance Theory as it is presently formulated. While Dominance Theory gives a reasonable account of people's reasoning about social contracts—rules regulating access to benefits—reasoning about precautions deviates from the theory in two respects. First, the dissociation in emotions evoked by violations of social contracts and precautions is not predicted by Dominance Theory given that it, like PRST, treats social contracts and precautions as psychologically equivalent rules—deontic rules pure and simple. Second and more specifically, the emotion study suggests that people are less prone to punish violations of precautions than violations of social contracts. This is evidenced by their tendency to get angry, a prelude

to punishing the transgressor, when social contracts are violated but not when precautions are violated. Hence, it is unlikely that enforcement of the “social code” motivates people’s reasoning about precautions quite like Dominance Theory predicts. The theory would be more consistent, it would seem, if it abandoned the claim that deontic reasoning is a unified phenomenon encompassing both social contracts and precautions, and concentrated on social contracts alone.

### Decision-theoretic accounts of deontic reasoning

Allied with SCT, at least with respect to abstract deontic rules, are various decision-theoretic accounts of deontic reasoning (Kirby, 1994; Manktelow & Over, 1991, 1995; Oaksford & Chater, 1994). Like SCT, these theories stress the importance of perceived costs and benefits and, hence, are similarly challenged to explain why subjects succeed in their reasoning about abstract deontic rules when it is not readily apparent how detecting violations of these rules would increase one’s subjective utility. While Cosmides’ (1989) interpretation of abstract permissions in terms of social contracts provides a common defence for both these decision-theoretic accounts and SCT and these theories, thereby, derive some support from the abstract selection task study presented here, these utilitarian theories find little support in the emotion study. Like Dominance Theory, the decision-theoretic approach is generally taken to apply uniformly to both social contracts and precautions (Manktelow & Over, 1991, 1995), but as the results of the emotion study demonstrate, social contracts and precautions are clearly dissociable in terms of their associated emotions. This suggests a psychological differentiation of these rules not captured by the decision theoretic approach.

While the results of the emotion study suggest that social contracts and precautions should be distinguished, the opposite conclusion is suggested by Manktelow and Over’s (1990, 1991, 1995) semantic analysis of deontic statements. Borrowing the methods and results of analytic philosophy, Manktelow and Over have asked “what it means to say” that one “may” or “must” perform an act. They conclude that any proper semantic analysis of deontic statements must ultimately refer to people’s preferences/perceived utilities and on these grounds there is no distinction between social contracts and precautions. In either case, we prefer that the rules are not violated and so we say that “If one does X, then one **MUST** do Y”, whereas when we are indifferent to the outcome we say “one **MAY** do Y”. How is vertical integration possible when the emotions pull one way and semantic intuitions pull the other?

One way of resolving this impasse is to consider the design constraints that the various psychological faculties are subject to. Whereas the emotions are functionally integrated with cognition and motivation (Lazarus, 1991)

suggesting a richness of psychological articulation, our linguistic abilities, which I am assuming semantics to be a part of, are under a strong design constraint to be economical, for example, information has to pass through a serial channel bottleneck (Pinker & Bloom, 1990), suggesting that the semantic system might be impoverished in comparison (although this representational poverty may be compensated for by pre- and postcommunicative inference, see Sperber & Wilson, 1995). The upshot is that the semantic system might, simply for communicative purposes, group things like social contracts and precautions (not to mention *necessity* and *possibility*, which also share the modal operators “must” and “may”) that the rest of the mind distinguishes. Although we may have easier cognitive access to the semantic system, it may provide only a rough map of the mind’s design.

## CONCLUSION

Despite frequent pronouncements of its demise, SCT remains a viable theory of people’s reasoning about social exchange. Over the past decade there have been refinements in the predictions of the theory, especially with respect to performance on the Wason selection task (Gigerenzer & Hug, 1992), but these changes have been within the spirit of the core theory—the computational theory of social exchange (Cosmides & Tooby, 1989). The theory has, unfortunately, had its fate too closely tied to findings from the Wason selection task. This narrow focus on a single method has distorted the intellectual debate about social cooperation. It has provided a false sense of psychological cohesion between social cooperation and hazards due to the logical analysis of selection task performance under which both social contracts and precautions are interpreted as eliciting the same pattern of responding. It has also encouraged alternative accounts of cheater detection based solely on methodological artifacts limited to the selection task. Moreover, it has obscured other facets of the psychology of social cooperation and hazards, such as associated emotions. Yet there is also much we have learned from using the selection task. What is called for is not an abandonment of the selection task as some have called for (Sperber et al., 1995), but converging lines of evidence from different methods and the second experiment reported here is, I hope, a step in that direction.

## ACKNOWLEDGEMENTS

I would like to thank Jörn Schultz and Stefan Krauss for their help in translating the problems and conducting the experiments and Denise Cummins, David Over and Melissa Rutherford for helpful comments. I would also like to acknowledge the generous support of the ESRC Centre for Economic Learning and Social Evolution, University College London, where I was a fellow while writing this chapter.

## REFERENCES

- Almor, A., & Sloman, S. (1996). Is deontic reasoning special? *Psychological Review*, *103*, 374–380.
- Almor, A., & Sloman, S. (2000). Reasoning versus text processing in the Wason selection task: A nondeontic perspective on perspective effects. *Memory & Cognition*, *28*, 1060–1070.
- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, *211*, 1390–1396.
- Blalock, H. (1972). *Social statistics* (2nd ed.). New York: McGraw Hill.
- Cheng, P., & Holyoak, K. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391–416.
- Cheng, P., & Holyoak, K. (1989). On the natural selection of reasoning theories. *Cognition*, *33*, 285–313.
- Cohen, J. (1977). *Statistical power analysis for the behavioral sciences* (revised Ed.). New York: Academic.
- Cosmides, L. (1985). *Deduction or Darwinian algorithm?: An explanation of the "elusive" content effect on the Wason selection task*. Doctoral dissertation, Harvard University. University Microfilms 86-02206.
- 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.
- Cosmides, L., & Tooby, J. (1989). Evolutionary psychology and the generation of culture, part II. Case study: A computational theory of social exchange. *Ethology and Sociobiology*, *10*, 51–97.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind* (pp. 163–228). New York: Oxford University Press.
- Cummins, D.D. (1996a). Evidence for the innateness of deontic reasoning. *Mind & Language*, *11*, 160–190.
- Cummins, D.D. (1996b). Evidence of deontic reasoning in 3- and 4-year-olds. *Memory & Cognition*, *24*, 823–829.
- Cummins, D.D. (1999). Cheater detection is modified by social rank. The impact of dominance on the evolution of cognitive functions. *Evolution & Human Behaviour*, *20*, 229–248.
- Evans, J.St.B.T. (1989). *Bias in human reasoning: Causes and consequences*. Hove, UK: Lawrence Erlbaum Associates Ltd.
- Evans, J.St.B.T., & Lynch, J. (1973). Matching bias in the selection task. *British Journal of Psychology*, *64*, 391–397.
- Fiddick, L. (1998). *The deal and the danger: An evolutionary analysis of deontic reasoning*. Unpublished doctoral dissertation. University of California, Santa Barbara.
- Fiddick, L. (under review). Domains of deontic reasoning: Resolving the discrepancy between the cognitive and moral reasoning literatures.
- Fiddick, L., Cosmides, L., & Tooby, J. (2000). No interpretation without representation: The role of domain-specific representations and inferences in the Wason selection task. *Cognition*, *77*, 1–79.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Gigerenzer, G. (1995). The taming of content: Some thoughts about domains and modules. *Thinking and Reasoning*, *1*, 324–333.
- Gigerenzer, G., & Hug, K. (1992). Domain specific reasoning: Social contracts, cheating, and perspective change. *Cognition*, *43*, 127–171.
- Giroto, V. (1991). Reasoning on deontic rules: The pragmatic reasoning schemas approach. *Intellectica*, *11*, 15–52.
- Giroto, V., Blaye, A., & Farioli, F. (1989a). A reason to reason: Pragmatic basis of children's search for counterexamples. *Cahiers de Psychologie Cognitive*, *9*, 297–321.

- Giroto, V., Gilly, M., Blaye, A., & Light, P. (1989b). Children's performance in the selection task: Plausibility and familiarity. *British Journal of Experimental Psychology*, *80*, 79-95.
- Giroto, V., Mazzocco, A., & Cherubini, P. (1992). Judgements of deontic relevance in reasoning: A reply to Jackson and Griggs. *Quarterly Journal of Experimental Psychology*, *45A*, 547-574.
- Griggs, R., & Cox, J. (1982). The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology*, *73*, 407-420.
- Holyoak, K., & Cheng, P. (1995). Pragmatic reasoning with a point of view. *Thinking and Reasoning*, *1*, 289-313.
- Jackson, S., & Griggs, R. (1990). The elusive pragmatic reasoning schemas effect. *Quarterly Journal of Experimental Psychology*, *42A*, 353-373.
- Johnson-Laird, P., Legrenzi, P., & Legrenzi, M. (1972). Reasoning and a sense of reality. *British Journal of Psychology*, *63*, 395-400.
- Kirby, K. (1994). Probabilities and utilities of fictional outcomes in Wason's four-card selection task. *Cognition*, *51*, 1-28.
- Kroger, K., Cheng, P., & Holyoak, K. (1993). Evoking the permission schema: The impact of explicit negation and a violation-checking context. *Quarterly Journal of Experimental Psychology*, *46A*, 615-635.
- Lazarus, R. (1991). *Emotion and adaptation*. Oxford: Oxford University Press.
- Lieberman, N., & Klar, Y. (1996). Hypothesis testing in Wason's selection task: Social exchange cheating detection or task understanding. *Cognition*, *58*, 127-156.
- Love, R., & Kessler, C. (1995). Focusing in Wason's selection task: Content and instruction effects. *Thinking and Reasoning*, *1*, 153-182.
- Manktelow, K., & Evans, J.St.B.T. (1979). Facilitation of reasoning by realism: Effect or non-effect? *British Journal of Psychology*, *70*, 477-488.
- Manktelow, K., & Over, D. (1990). Deontic thought and the selection task. In K. Gilhooly, M. Keane, R. Logie, & G. Erdos (Eds.), *Lines of thought: Reflections of the psychology of thinking* (pp. 153-164). London: Wiley.
- Manktelow, K., & Over, D. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition*, *39*, 85-105.
- Manktelow, K., & Over, D. (1995). Deontic reasoning. In S. Newstead & J.St.B.T. Evans (Eds.), *Perspectives on thinking and reasoning: Essays in honour of Peter Wason* (pp. 91-114). Hove, UK: Lawrence Erlbaum Associates Ltd.
- Matsumoto, D., & Ekman, P. (1988). *Japanese and Caucasian facial expressions of emotion (JACFEE) and neutral faces (JACNeuF)* [slides]. San Francisco: Intercultural and Emotion Research Laboratory, Department of Psychology, San Francisco State University.
- Noveck, I., & O'Brien, D. (1996). To what extent do pragmatic reasoning schemas affect performance on Wason's selection task? *Quarterly Journal of Experimental Psychology*, *49A*, 463-489.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*, 608-631.
- Pinker, S., & Bloom, P. (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, *13*, 707-784.
- Platt, R., & Griggs, R. (1993). Darwinian algorithms and the Wason selection task: A factorial analysis of social contract selection task problems. *Cognition*, *48*, 163-192.
- Politzer, G., & Nguyen-Xuan, A. (1992). Reasoning about conditional promises and warnings: Darwinian algorithms, mental models, relevance judgments or pragmatic schemas? *Quarterly Journal of Experimental Psychology*, *44A*, 401-421.
- Rozin, P., Lowery, L., Imada, S., & Haidt, J. (1999). The CAD triad hypothesis: A mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). *Journal of Personality and Social Psychology*, *76*, 574-586.
- Shweder, R., Mahapatra, M., & Miller, J. (1987). Culture and moral development. In J. Kagan &

- S. Lamb (Eds.), *The emergence of moral concepts in early childhood* (pp. 1–83). Chicago, IL: University of Chicago Press.
- Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, *57*, 31–95.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Cambridge, MA: Blackwell.
- Tooby, J., & Cosmides, L. (1990). The past explains the present: Emotional adaptations and the structure of ancestral environments. *Ethology and Sociobiology*, *11*, 375–424.
- Trivers, R. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, *46*, 35–57.
- Wason, P. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*, 273–281.
- Wason, P., & Shapiro, D. (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology*, *23*, 63–71.