

What Determines the Acceptability of Deontic Health and Safety Rules?

Linden J. Ball (L.Ball@lancaster.ac.uk)
Department of Psychology, Lancaster University, UK

David Alford (D.Alford@lancaster.ac.uk)
Department of Psychology, Lancaster University, UK

Abstract

On what basis do people provide acceptability or ‘goodness’ judgements for *deontic* conditional rules of the form *if p then must q* and *if p then may q*? Using a decision theoretic analysis, Over et al. (2004) propose that such conditionals are judged as acceptable to the extent that the *p&q* possibility is preferred to the *p¬-q* possibility. Their empirical evidence upholds this ‘conditional expected utility’ (CEU) prediction for conditional obligations and permissions relating to everyday activities (e.g., *If you wash the dishes then you must wear rubber gloves*). We report two experiments examining Over et al.’s CEU hypothesis in relation to real-world deontic rules concerning everyday health and safety issues (e.g., *If you are in a moving car then you must have your safety-belt fastened*). We propose that Over et al.’s CEU hypothesis provides a compelling account of our findings.

Keywords: Deontic reasoning; permission and obligation rules; health and safety; conditional expected utility.

Introduction

The nature of deductive reasoning has long been of interest to cognitive scientists, and one logical form—the *conditional*—continues to attract considerable research attention. It has become increasingly apparent that there are, in fact, two distinct types of conditional: ‘indicative’ and ‘deontic’. Indicative conditionals are rules that are used to describe matters of fact. Take the case of medical studies indicating that the hand-arm vibration (HAV) associated with excessive, continuous use of vibrating power tools (VPTs) can lead to the medical condition ‘vibration white finger’. With such medical evidence in mind, a consultant might assert the following conditional to a patient who uses vibrating equipment as part of their work: *If you use VPTs excessively then you will damage your health*. This latter conditional is an indicative conditional since it expresses a factual, cause-effect relation between an antecedent condition and a consequent state of affairs.

In contrast to indicative conditionals, deontic conditionals are used to *direct* people’s behaviour. For example, a company-based health and safety (H&S) inspector who has concerns about the health implications for workers of using VPTs may assert the deontic rule *If you have operated a VPT for 15 minutes then you must have a break*. Deontic conditionals such as this are ‘proper’ rules, since they provide strong guidance on what must be done by specifying social regulations, laws and moral rules. Such conditionals are easy to spot as they tend to contain the

deontic modal operators *should*, *ought to*, and *must* (expressing obligation), and *may* (expressing permission).

Our interest in the present paper closely follows that of Over, Manktelow, and Hadjichristidis (2004), who questioned why someone might consider deontic rules as being ‘good’ rules in the first place. In other words, what determines the acceptability of some deontic conditionals and the unacceptability of others? Our aim in the present paper is to address such issues by focusing on obligations and permissions that relate to H&S situations that pertain to everyday, real-world activities such as using road vehicles.

What Makes a Deontic Conditional a Good Rule?

In examining what makes deontic conditionals acceptable we follow Over et al. (2004) in adopting a *decision theoretic perspective* on this issue (cf. Perham & Oaksford, 2005), which claims that it is vital to ascertain people’s *preferences* for certain possible actions over other possible actions. The importance of preference judgements in determining deontic rule acceptability is, perhaps, best illustrated in relation to the conditional obligation previously specified: *If you have operated a VPT for 15 minutes then you must have a break*. Why would a H&S inspector stipulate this rule about VPT operation to members of the workforce? Under Over et al.’s (2004) account, the answer depends on the inspector’s preferences among the following four logical possibilities:

- | | |
|----------------------|---|
| (<i>p&q</i>) | The VPT is operated for 15 minutes and the worker has a break |
| (<i>p&¬q</i>) | The VPT is operated for 15 minutes and the worker does not have a break |
| (<i>¬p&q</i>) | The VPT is not operated for 15 minutes and the worker has a break |
| (<i>¬p&¬q</i>) | The VPT is not operated for 15 minutes and the worker does not have a break |

A H&S inspector will clearly have preferences among these possibilities, such that they are likely to prefer the *p&q* situation to the *p&¬q* situation. Thus, the H&S inspector sees the expected utility of 15 minutes of VPT operation being reached and the worker having a break as greater than the expected utility of 15 minutes of VPT operation being reached and the worker not having a break. Without such a preference there would be little point in the inspector asserting the rule as an obligation for the workforce to follow. Indeed, it is the very preference for *p&q* over *p&¬q* that renders a conditional obligation acceptable in the first

place. For Over et al. (2004) this is the crux of what it means for a deontic conditional to be seen as a good rule. Clearly, however, the preferences of individuals within the workforce who actually use VPTs will be *subjective*, and may differ from the preferences of the inspector. Some workers may even have reversed $p \& q$ and $p \& \neg q$ preferences compared with the inspector, whereby they view the expected utility of 15 minutes of VPT operation being reached and not having a break as greater than the expected utility of 15 minutes of VPT operation being reached and having a break. Such reversed preferences may arise simply because the operative places a high value on timely completion of contracted jobs. For an individual who views the $p \& \neg q$ situation as having more expected utility than the $p \& q$ situation, a conditional obligation rule will be seen as a 'bad' rule. Such differences in preferences may depend on a person's role or rank in the workplace (e.g., Manktelow & Over, 1991, 1995) or on the perceived likelihood of various consequences (e.g., Manktelow, Sutherland, & Over, 1995).

The Conditional Expected Utility Hypothesis

Over et al. (2004) refer to their decision theoretic analysis of deontic rule acceptability as the *conditional expected utility* (henceforth CEU) hypothesis, since people are considered to accept such rules based on judgements of the likelihood (P) and utility (U) of the rule's consequent (q) given that the antecedent (p) holds, that is, $P.U(q|p)$. It is the ratio between $P.U(p \& q)$ and $P.U(p \& \neg q)$ —reflected in the preference for $p \& q$ over $p \& \neg q$ —that provides a value for $P.U(q|p)$.

One key strength of the CEU hypothesis is that it generalises the 'suppositional theory' of the indicative conditional (e.g., Edgington, 1995; Evans & Over, 2004) to the deontic domain. The suppositional theory proposes that an indicative rule of the form *if p then q* imparts the belief that there is a relation between p and q , and that this belief is captured by a person's judgement of the *conditional probability*, q given p , that is, $P(q|p)$. Take the indicative conditional: *If you use VPTs excessively then you will damage your health*. An individual who is interpreting this rule is claimed to 'suppose' the antecedent by a process dubbed the 'Ramsey test' (Ramsey, 1990), whereby they hypothetically add the antecedent to their stock of knowledge. They then mentally simulate the consequent of the rule in the context of the antecedent. How much the individual believes in the consequent under the antecedent determines the *strength* of the conditional (i.e., its acceptability) as is reflected in the conditional probability judgement, $P(q|p)$, such as the probability of damaging your health given that you use VPTs excessively.

This suppositional account of the indicative conditional has been well supported in recent studies (e.g., Evans, Handley & Over, 2003; Oberauer & Wilhelm, 2003). Over et al.'s (2004) account of how people determine the acceptability of deontic conditionals reflects an ingenious extension of the Ramsey test to deontic reasoning. So far, however, we have only been considering Over et al.'s (2004) CEU hypothesis as it relates to the acceptability of

conditional obligation rules. But how does the hypothesis relate to the acceptability of conditional permission rules?

Consider the following permission rule: *If you have operated a VPT for about 15 minutes then you may have a break*. According to the CEU account, this conditional permission would be acceptable for someone when either: (1) they definitely preferred the $p \& q$ situation to the $p \& \neg q$ situation; or (2) they were indifferent between the two situations (see Manktelow & Over, 1995). Consider an occasion when an H&S inspector simply wanted to assess whether workers might be inclined to take breaks from VPTs if this was highlighted as a permissible action. The inspector might perceive a small health benefit for workers if they took a break after 15 minutes (the $p \& q$ situation) but might also recognise the small benefit in productivity that would arise from workers not taking a break (the $p \& \neg q$ situation). Preference judgements might be finely balanced between these two possibilities, but such a state of near 'indifference' would still be an occasion when this rule would be endorsed by the inspector (cf. Over et al., 2004). It is important to note, however, that in general, it would be pragmatically anomalous for someone like the inspector to assert a relatively weak conditional permission when they actually wanted to assert a relatively strong conditional obligation. Thus a conditional permission would rarely be stated if a conditional obligation was viewed as a better rule.

Over et al. (2004) not only present a compelling account of the acceptability conditions for obligation and permission rules, but they also report two experiments that support predictions of the CEU hypothesis. In particular, they demonstrate that deontic rules of the form *if p then must q* and *if p then may q* are judged as good to the extent that the $p \& q$ possibility is preferred to the $p \& \neg q$ possibility. In addition, they reveal an unexpected difference in the evaluation of *must* and *may* rules, with the former being judged as better rules. Intriguingly, this difference between *must* and *may* rules was seen to generalise from contexts where rules were presented as being self-imposed personal rules (Over et al., 2004, Experiment 1) to contexts where rules were being expressed by an agent in a position of authority in order to regulate the actions of an actor who was the target of the rule (Over et al., 2004, Experiment 2).

Experiment 1

Despite its clear strengths, however, Over et al.'s (2004) research focused almost entirely on fairly mundane situations where rule violations are linked only with relatively mild moral consequences (e.g., the potential for being seen as a 'free rider' when meeting up with friends without enough money). In contrast, our experimental work is concerned with deontic rules pertaining to important, real-world H&S contexts that are associated with the possibility of serious consequences for rule violation, including, in the extreme case, the potential for loss of life.

Experiment 1 utilised the paradigm pioneered by Over et al. (2004), and applied this to examine issues concerning the acceptability of conditional obligations and permissions

pertaining to real-world H&S situations. The experiment manipulated two factors: (1) whether rules were *obligations* or *permissions*; and (2) whether H&S content was *present* or *absent* within rules. All scenarios asked participants to adopt an actor's perspective on rules that were being stipulated by an individual in authority. Where H&S content was present we used a rule-set involving prudential obligations and permissions concerning physical well-being, protection and safety. Examples of such rules are: *If you are riding a motorbike then you must wear protective clothing* (obligation) and *If you are wearing a bicycle helmet then you may ride your bike* (permission). Note that none of these rules was associated with a financial or legal penalty for violation; they were all commonsense rules that prevail in the UK. For the cases in which H&S content was absent we mainly used Over et al.'s (2004) rule-set of prudential obligations and permissions, although we had to replace three of the original obligation rules as these had loose associations to H&S issues. Examples of the rules from this set that we used in Experiment 1 are: *If you meet up with your friends then you must have enough money* (obligation) and *If you extend your overdraft limit then you may buy some extra luxuries* (permission).

Our detailed objectives were: (1) to replicate Over et al.'s observation of a difference in goodness judgements for *must* versus *may* rules; (2) to determine whether any *must/may* difference in perceived rule goodness generalises to rules where H&S content is present; (3) to examine the novel hypothesis that rules where H&S content is present may be judged as better than rules where H&S content is absent (given the more salient personal costs in the former in terms of well-being and safety); (4) to replicate Over et al.'s finding that the preference for the $p \& q$ possibility over the $p \& \neg q$ possibility is greater for *must* rules than for *may* rules; and (5) to validate the general CEU prediction that deontic conditionals will be judged as good to the extent that the $p \& q$ possibility is preferred to the $p \& \neg q$ possibility.

Method

Thirty individuals volunteered via a local participant panel. All were competent English speakers who were familiar with UK culture. None had received prior training in reasoning or logic. Participants were tested individually and received £3 for tackling two tasks. One task involved providing goodness ratings for each rule in a set of 12 obligation rules (6 with H&S content and 6 without H&S content) and 12 permission rules (6 with H&S content and 6 without). Presentation order of these rules was randomised for each participant. Response sheets for these rules contained the following instructions (cf. Over et al., 2004):

"In everyday life we are often given rules by people about what we must do, or what we may do, in certain situations. An example of a 'must' rule might be: *If you go for a run then you must take a shower*. And an example of a 'may' rule might be: *If you go for a run then you may eat some chocolate cake*. One question we can ask about these rules is whether we think they are good rules or not. On the next

few pages you will see a set of the kinds of rules that you might be given. For each rule you will be asked to provide a rating in terms of how good a rule you think it is. For example, consider the following rule: *If you go for a run then you must take a shower*. How good a rule is this?"

Following this question a 11-point scale was presented, with 0 labelled 'a very bad rule', 5 'neutral' and 10 'a very good rule'. The instructions then asked participants to turn over the page and consider each of the given rules.

The other task involved giving participants all four logical possibilities for a particular rule ($p \& q$, $p \& \neg q$, $\neg p \& q$, and $\neg p \& \neg q$) and asking them to arrange these possibilities in order of preference (where 1 was most preferable and 4 least preferable). Presentation of the four possibilities for each rule was randomised for all rules and for all participants. The instructions for this task read:

"On the following pages you will see lists of four possible situations. Please look at each list and give each situation a number to show the order in which *you yourself* would prefer these situations to occur. For example, here is a list of four possible situations:

The possible situations are:	Preference order:
You go for a run and you take a shower
You go for a run and you don't take a shower
You don't go for a run and you take a shower
You don't go for a run and you don't take a shower

If the best situation for you would be that you don't go for a run and don't take a shower, give this a preference score of 1; if the next is that you go for a run and take a shower, give this a preference score of 2; and so on".

The instructions then directed participants to turn over the page and begin. No time limit was set for the experiment. The order of presentation of the two experimental tasks was counterbalanced so that half received the rule-goodness task first followed by the situation-preference task, whilst the other half received these tasks in the reverse order.

Results and Discussion

Rule goodness scores Initial analyses of the rule-goodness scores (see Table 1) related to our first three experimental objectives: (1) to replicate Over et al.'s observation of a difference in goodness judgements for *must* (obligation) versus *may* (permission) rules; (2) to establish whether any difference between *must* and *may* rules generalises across rules where H&S content is present; and (3) to determine whether rules where H&S content is present are judged as better than rules where H&S content is absent.

Table 1 shows mean goodness scores for rules in Experiment 1. A two-way ANOVA revealed a main effect of Deontic Rule Form, $F(1, 29) = 60.19$, $MSE = 0.99$, $p < .001$, replicating Over et al.'s finding that obligation rules are judged to be better rules than permission rules. Our mean goodness scores for *must* and *may* rules without H&S content were nearly identical to those obtained by Over et al. (2004, Experiment 2) for an essentially equivalent rule-

set. Our analysis also revealed a reliable main effect of Content, $F(1, 29) = 173.66$, $MSE = 0.69$, $p < .001$, indicating that H&S rules were judged as better than non-H&S rules. The interaction was not reliable, $F(1, 29) = 0.61$, $MSE = 0.58$, $p = .44$. Thus, it appears that the effects of Deontic Rule Form and Content are additive.

Comparing goodness scores and preference scores Our next analyses aimed to replicate Over et al.'s finding that the preference for the $p \& q$ over the $p \& \neg q$ possibility is greater for *must* than *may* rules. Table 2 shows the mean differences in preference ratings for the $p \& q$ and $p \& \neg q$ possibility for the various rule types used. ANOVA revealed a predicted effect of Deontic Rule Form, $F(1, 29) = 31.01$, $MSE = 0.37$, $p < .001$, with *must* rules showing a stronger preference for the $p \& q$ possibility over the $p \& \neg q$ possibility than *may* rules. ANOVA also revealed a significant effect of Content, $F(1, 29) = 6.86$, $MSE = 0.40$, $p = .014$, indicating that rules where H&S content was present were associated with a stronger preference for the $p \& q$ possibility over the $p \& \neg q$ possibility than rules where H&S content was absent. The interaction was not significant, $F(1, 29) = 0.36$, $MSE = 0.20$, $p = .55$, suggesting that the effects of Deontic Rule Form and Content are, again, additive in nature. Indeed, the pattern of effects for the present preference scores closely parallels that for the rule-goodness scores, further validating Over et al.'s CEU account of what renders deontic rules acceptable: They are acceptable to the extent that the $p \& q$ situation is preferred over the $p \& \neg q$ situation. Moreover, such acceptability judgements and $p \& q$ versus $p \& \neg q$ preferences appear sensitive to the form and content of rules in ways that make sense in terms of associated utilities of the $p \& q$ and $p \& \neg q$ situations.

To validate the CEU hypothesis further we used our dataset to test another of Over et al.'s predictions: that irrespective of rule content (H&S present vs. absent) or rule form (*must* or *may*), deontic rules will be judged as good to the extent that the $p \& q$ possibility is preferred to the $p \& \neg q$ possibility. To test this prediction we computed Spearman rank-order correlation coefficients for each individual's goodness ratings for a rule and their $(p \& \neg q) - (p \& q)$ preference ratings for the situations associated with the rule. This was done collapsing across Deontic Rule Form and Content, to produce 30 correlations (one per participant). Of these, 27 were in the predicted (positive) direction, which was reliable with a binomial test, $p < .001$, two-tailed.

Table 1. Mean goodness scores for rules in Experiment 1.

Content	Deontic Rule Form			
	Obligation		Permission	
	M	SD	M	SD
H&S Present	8.91	0.84	7.61	1.02
H&S Absent	7.05	1.18	5.54	1.25

Note. SD = Standard deviation.

Table 2. Mean difference in the rating of the $p \& q$ situation and the rating of the $p \& \neg q$ situation in Experiment 1.

Content	Deontic Rule Form			
	Obligation		Permission	
	M	SD	M	SD
H&S Present	1.63	0.68	1.07	0.53
H&S Absent	1.38	0.54	0.72	0.70

Notes. Difference scores for preference ratings are calculated in terms $(p \& \neg q) - (p \& q)$. Ratings were subtracted in this way because of the order manipulation, where higher numbers indicate a lower preference (cf. Over et al., 2004). SD = Standard deviation.

Experiment 2

Experiment 1 replicates key findings from Over et al. (2004), such as the greater acceptability of *must* versus *may* rules, and further supports the CEU account of the acceptability of conditional obligations and permissions. Experiment 1 also extends Over et al.'s research by generalising the CEU theory to deontic rules embodying H&S content, where personal costs associated with rule-related situations are potentially high, extending beyond mere vulnerability to the possibility of injury or death. We observed that when H&S content is present in obligation and permission rules then they are judged as better than equivalent rules where H&S content is absent. Likewise, H&S rules are associated with higher mean differences in the rating of their associated $p \& q$ versus $p \& \neg q$ situations when compared to non-H&S rules. These effects all seem readily interpretable in terms of the CEU hypothesis.

We note, however, that no rules in Experiment 1 were associated with financial or legal penalties for violation. The question that then arises is whether financial or legal costs associated with rules might render them more acceptable. Indeed, penalties are typically associated with rules that authorities hold as particularly important for the protection and safety of a society's members. Thus, in the UK a H&S rule such as *If you are in a moving car then you must have your safety-belt fastened* would be expected to be judged as highly acceptable and produce strong preferences for the $p \& q$ over the $p \& \neg q$ situation. Experiment 2, therefore, examined the same CEU predictions as in Experiment 1, but with materials encompassing fiscal or legal penalties for rule violation. We also manipulated the same factors as previously: whether rules were obligations or permissions; and whether H&S content was present or absent.

Method

Participants were 30, English-speaking volunteers who were familiar with UK culture. All were tested alone and received £3. Equivalent instructions, procedures and controls were used to Experiment 1. For the cases where H&S content was present, the rule-set involved prudential obligations (six rules) and permissions (six rules) concerning well-being, protection and safety. Unlike Experiment 1, these rules were

linked to the possibility of monetary or legal penalties for violation. Examples are: *If you are using a mobile phone when driving then it must be a hands-free device* (obligation) and *If you plan to take a taxi home from the pub later then you may go out drinking alcohol for the evening* (permission). For the cases where H&S content was absent we used a rule-set of prudential obligations (six rules) and permissions (six rules) that were also associated with financial or legal repercussions if violated. Examples are: *If your library book has reached its return date then you must take it back to the library* (obligation) and *If your TV licence is valid then you may watch your TV* (permission).

Results and Discussion

Rule goodness scores The rule-goodness scores in Experiment 2 (see Table 3) were marginally higher than in Experiment 1, supporting the view that rules linked to financial and legal penalties are judged as more acceptable than those not linked to such penalties. A two-way ANOVA revealed a main effect of Deontic Rule Form, $F(1, 29) = 32.41$, $MSE = 1.19$, $p < .001$, replicating the Experiment 1 result that *must* rules are judged as better than *may* rules. There was also a reliable main effect of Content, $F(1, 29) = 29.75$, $MSE = 0.46$, $p < .001$, indicating that rules relating to H&S were judged as better than those unrelated to H&S. In addition, and unlike Experiment 1, the interaction was reliable, $F(1, 29) = 7.61$, $MSE = 0.74$, $p = .01$, and appears to arise because the separation in goodness evaluations for *must* versus *may* rules is more marked on the H&S contents than the non-H&S contents. This attests to the apparent potency of conditional obligations embodying H&S content and possible financial or legal costs for violation.

Table 3. Mean goodness scores for rules in Experiment 2.

Content	Deontic Rule Form			
	Obligation		Permission	
	M	SD	M	SD
H&S Present	9.13	0.66	7.57	1.27
H&S Absent	8.02	1.15	7.32	1.18

Note. SD = Standard deviation.

Comparing goodness scores and preference scores Our next analyses examined the hypothesis that the preference for the $p \& q$ possibility over $p \& \neg q$ will be greater for *must* than *may* rules. Table 4 shows the mean differences in preference ratings for the $p \& q$ possibility and the $p \& \neg q$ possibility for the four types of rule used in Experiment 2. A 2×2 ANOVA revealed a reliable main effect of Deontic Rule Form, $F(1, 29) = 57.54$, $MSE = .33$, $p < .001$, with *must* rules showing a stronger preference for the $p \& q$ possibility over the $p \& \neg q$ possibility than *may* rules. The ANOVA also revealed a near significant main effect of Content, $F(1, 29) = 3.30$, $MSE = .17$, $p = .08$, which goes some way toward supporting the finding from Experiment 1

that rules embodying H&S content are associated with a stronger preference for the $p \& q$ over the $p \& \neg q$ possibility than rules without H&S content. The interaction was not significant, $F(1, 29) = 0.35$, $MSE = 0.26$, $p = .56$, and thus, for the present experiment, the preference effects do not entirely parallel the rule-goodness effects. Notwithstanding this slight anomaly, however, the experiment again provides reasonable support for the CEU account of deontic rule acceptability. Such rules appear acceptable to the extent that the $p \& q$ situation is preferred over the $p \& \neg q$ situation.

Table 4. Mean difference in the rating of the $p \& q$ situation and the rating of the $p \& \neg q$ situation in Experiment 2.

Content	Deontic Rule Form			
	Obligation		Permission	
	M	SD	M	SD
H&S Present	1.92	0.50	1.07	0.56
H&S Absent	1.73	0.52	0.98	0.67

Note. All of the notes associated with Table 2 apply here.

We finally tested the prediction that irrespective of rule content (H&S present/absent) or rule form (*must/may*), deontic conditionals will be judged as good to the extent that $p \& q$ is preferred to $p \& \neg q$. As in Experiment 1, we correlated each participant's goodness ratings for a rule and their $(p \& \neg q) - (p \& q)$ preference ratings for situations associated with the rule (collapsing across Deontic Rule Form and Content). Of the 30 correlations, 24 were in the predicted (positive) direction, which was reliable with a binomial test, $p = .002$, two-tailed.

General Discussion

This paper aimed to replicate and extend Over et al.'s (2004) research on the issue of why some deontic rules are more acceptable than others. Like Over et al., we focused on conditional prudential obligations and permissions and adopted a decision-theoretic stance, whereby we predicted that: (1) people would judge a conditional obligation (*if p then must q*) as a good rule when they preferred the situation in which the antecedent and consequent co-occur ($p \& q$) to the situation in which the antecedent occurs but the consequent does not ($p \& \neg q$); and (2) people would judge a conditional permission (*if p then may q*) as a good rule when either they preferred $p \& q$ to $p \& \neg q$, or when they were indifferent between these two situations.

The key aspect of our experiments that went beyond Over et al.'s research was the manipulation of rule content such that H&S issues were either present or absent. The personal costs associated with our H&S-related rules included risks of injury or fatality. Whilst Experiment 1 used rules where there was *no* financial or legal penalty for rule violation, Experiment 2 used rules linked to such penalties. These rules are of the type that authorities hold as especially worthy of being followed for optimal functioning of

organisations and societies, and—in the case of H&S rules—for people’s personal protection and well-being. Such rules should generally be acceptable to people, who should be familiar with following them and well aware of the associated benefits that derive from their adoption.

The predictions of the CEU theory were borne out across both experiments. We observed an association between rule-goodness evaluations and preference ratings for $p \& q$ over $p \& \neg q$ situations. Both experiments also replicated the strong propensity for people to evaluate obligation (*must*) rules as better than permission (*may*) rules. Over et al. (2004) express concerns that the differences between *must* and *may* rules that they observed may have arisen because of their sample (psychology students) and their rule contents. The fact that we have generalised this effect to a more representative sample of the adult population as well as across new sets of materials attests to the validity of the difference between *must* and *may* goodness evaluations.

Our data also provided some support for predictions that deontic rules containing H&S content would be judged as better than those not containing H&S content, and that people’s preference ratings for H&S $p \& q$ situations over $p \& \neg q$ situations would mirror H&S rule-goodness judgements. Moreover, the effects of H&S content on the acceptability of rules as well as situation-preference ratings generalised from Experiment 1 (where financial and legal costs for rule violation were absent) to Experiment 2 (where such costs were present)—attesting to the robustness of predicted CEU effects across contents. Financial and legal costs associated with rules (in Experiment 2) led to a general increase in the acceptability of such rules compared to Experiment 1 (where such penalties did not pertain).

Our research derives from a decision-theoretic approach to understanding deontic reasoning, and closely follows the work of Over et al. (2004; Manktelow & Over, 1991, 1995). We acknowledge that there are alternative theoretical accounts of reasoning with deontic conditionals (e.g., Almor & Sloman, 1996; Cheng & Holyoak, 1985; Cosmides, 1989; Johnson-Laird & Byrne, 2002; Oaksford & Chater, 1994). However, such accounts are primarily focused on explaining people’s responses with deontic variants of the Wason selection task, where the key response measure concerns people’s capacity to detect rule violators. The emphasis of the CEU account is, of course, somewhat different to these latter theories, in that it attempts to explain why deontic conditionals are adopted or accepted by people in the first place as appropriate rules for guiding behaviour in organisations, institutions or societies.

The underlying assumption of the CEU account is that in making everyday decisions people are trying to work out which actions they *must* or *may* perform in the light of their preferences (Over et al., 2004). Moreover, preferences will be determined by a range of beliefs and desires that have their basis in moral, legal and prudential value systems acquired during people’s life and work experiences. Explaining how people reason about rules in relation to such preferences is a complex and challenging endeavour, and

one that psychologists are only beginning to grapple with (see also Hilton, Kimmelmeier, & Bonnefon, 2005). Our research represents another small attempt to examine issues that are relevant to this endeavour. Our results suggest that the concept of *conditional expected utility* as proposed by Over et al. (2004) may have a valuable role to play in progressing a theoretical understanding of why people see some deontic rules as more acceptable than others.

Acknowledgments

This research was funded by the EPSRC (EP/CO14677/1).

References

- Almor, A., & Sloman, S.A. (1996). Is deontic reasoning special? *Psychological Review*, *103*, 374-380.
- Hilton, D. J., Kimmelmeier, M., & Bonnefon, J-F. (2005). Putting *If*s to work: Goal-based relevance in conditional reasoning. *Journal of Experimental Psychology: General*, *134*, 388-405.
- Cheng, P.W., & Holyoak, K.J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391-416.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, *31*, 187-276.
- Edgington, D. (1995). On conditionals. *Mind*, *104*, 235-329.
- Evans, J. St. B. T., Handley, S. J., & Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *29*, 321-335.
- Evans, J. St. B. T., & Over, D. E. (2004). *If*. Oxford: OUP.
- Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inferences. *Psychological Review*, *109*, 646-678.
- Manktelow, K. I., & Over, D. E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition*, *39*, 85-105.
- Manktelow, K. I., & Over, D. E. (1995). Deontic reasoning. In S. E. Newstead & J. St. B. T. Evans (Eds.), *Perspectives on thinking and reasoning: Essays in honour of Peter Wason*. Hove: LEA.
- Manktelow, K. I., Sutherland, E. J., & Over, D. E. (1995). Probabilistic factors in deontic reasoning. *Thinking and Reasoning*, *1*, 201-220.
- Oaksford, M. & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*, 608-631.
- Oberauer, K., & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *29*, 680-693.
- Over, D. E., Manktelow, K. I., & Hadjichristidis, C. (2004). Conditions for the acceptance of deontic conditionals. *Canadian Journal of Experimental Psychology*, *58*, 96-105.
- Perham, N., & Oaksford, M. (2005). Deontic reasoning with emotional content: Evolutionary psychology or decision theory? *Cognitive Science*, *29*, 681-718.
- Ramsey, F. P. (1990). General propositions and causality (original publication, 1931). In D. H. Mellor (Ed.), *Philosophical papers*. Cambridge: CUP.