THINKING AND REASONING, 1995, *1* (4), 373–388

Pragmatic Reasoning from Multiple Points of View: A Response

Keith J. Holyoak and Patricia W. Cheng
University of California, USA

Rejoinder to Open Peer Commentaries

In our response we consider several broad sets of issues raised in the commentaries on our target article. We provide an elaboration of the general theory of pragmatic reasoning schemas and of the mechanisms posited to explain perspective effects and other contextual influences on reasoning. We consider the strengths and limitations of a number of alternative proposals offered by the commentators. Finally, we discuss possible links between pragmatic schemas and more elementary "building blocks" for human reasoning.

INTRODUCTION

The commentaries on our target article provide a multiplicity of approaches, which we will only be able to touch on in our response. The impact of perspectives on reasoning about social contracts has clearly attracted a great deal of interest, and it is worth stepping back to consider why this is so. For 25 years it has been evident that everyday reasoning about problems that appear formally equivalent with respect to standard propositional logic can differ enormously depending on the specific content of the material (e.g. Fillenbaum, 1976; Johnson-Laird, Legrenzi, & Legrenzi, 1972; Wason & Shapiro, 1971). For some time it was difficult to perceive any clear overall pattern in the vagaries of subjects' responses to specific problem content (Griggs & Cox, 1982). Ten years ago, we proposed that people's everyday inferences typically depend on content-based schemas at a level of generality intermediate between the concrete level of specific known cases and the full abstraction of formal logic (Cheng & Holyoak, 1985; Cheng, Holyoak, Nisbett, & Oliver, 1986). These pragmatic reasoning schemas correspond to such broad but nonetheless content-based domains as social regulations and causal relations.

Over the past decade the domain of social regulations has received particular attention, both in our own work and that of other investigators concerned with deductive reasoning. The bulk of this research has been based on regulation isomorphs to Wason's (1966) selection task. Problems that people interpret in terms of conditional permissions and obligations yield a distinctive pattern of card selections in terms of conditional permissions and obligations yield a distinctive pattern of card selections in the Wason paradigm; moreover, the dominant pattern

observed in many experiments has corresponded to those selections licensed by standard propositional logic (i.e. the *p & not-q* cases). Thus it appeared that social regulations yield "facilitation" of reasoning, in the sense of promoting inferences in accord with the logic of the material conditional.

In this context, the influence of perspective on selection choices (e.g. Gigerenzer & Hug, 1992; Manktelow & Over, 1991) has attracted great interest for two reasons. First, whereas previous work had established that distinct types of content evoke distinct inference patterns, perspective shifts create distinct inference patterns *within* a single domain, namely contractual varieties of regulations. Second, one of these patterns apparently corresponds to the dominant selection of the *not-p & q* cases—precisely opposite to the pattern licensed by standard logic, and hence a prima facie example of "anti-facilitation" of logical inference by meaningful content. Perspective effects thus have raised basic questions about the nature of the knowledge people use to guide their inferences and selection of actions.

Moreover, other recent work has had a converging impact. Several commentators (Over & Manktelow, Oaksford & Chater, O'Brien) as well as our own target article, noted recent evidence that subjective expected utilities play a role in determining whether a potential violation will be checked (e.g. Kirby, 1994). Subjective expected utilities introduce pragmatic considerations beyond those originally postulated in our theory of pragmatic reasoning schemas. In addition, Johnson-Laird and Byrne, Girotto, and Evans and Clibbens note that facilitation can be observed in the selection task even with non-deontic materials, as long as certain general contextual constraints are established (e.g. Sperber, Cara, & Girotto, in press).

Against this general background, several interrelated questions arise. Is the original conception of pragmatic reasoning schemas too rigid to accommodate the broader varieties of context and content effects that have recently been demonstrated? Do pragmatic schemas need to be supplemented, or replaced, by additional types of inference mechanisms? Can schemas be derived from more elementary elements, and if so what might these be?

SOME THOUGHTS ABOUT CONTENT AND CONTEXT IN REASONING

These questions are deep ones, and in our view each of the commentaries provides some valuable insights that deserve further exploration. Not surprisingly, we do not always agree as to which directions are most likely to prove fruitful (nor is any unanimity of option to be found among the commentators!). We will now consider some of the specific points raised by the commentaries. As several shared themes recur, we will focus on these issues and questions of general concern.

How Do Pragmatic Schemas Work?

Our aim in the target article was to derive and test an account of perspective effects using the permission and obligation schemas as originally formulated by Cheng and Holyoak (1985) and elaborated by Politzer and Nguyen-Xuan (1992). The core of each schema is a set of four conditional rules. Over and Manktelow note that these rules differ from standard production rules in that in each case the consequent is not an action to be performed, but rather a state or action that "may" or "must" hold. Indeed, we should have clarified that a schema rule is like a typical "production rule" in that when its left-hand side is matched by instantiating the variables in it then the right-hand side is executed; but they differ in that the right-hand side of a matched rule yields an assertion about the deontic status of a state or potential act, rather than a direct prescription for an action. That is, the direct inferential import of the rules in regulation schemas is to specify what situations would or would not violate the stated rule, rather than what action to take. In the target article we made the ancillary assumption that people will typically be inclined to check situations in which their own rights (and the duties owed them by others) are at risk, an assumption critical to our account of the influence of perspective on selection performance.

Our use of the term "production rules", which is often associated with conscious reasoning (although such linkage is questionable), seems to have fostered the main reservation expressed by Evans and Clibbens. They consider the application of schema rules (considered as production rules) to be a conscious and serial process. Based on Evans' (in press) finding that subjects inspect the cards they will eventually select much longer than those they will not select, Evans and Clibbens argue that subjects use some unconscious mechanism to determine which cases are relevant, and afterwards apply conscious processes to those cases that "pass" the relevance check. If use of the schema rules corresponded to the conscious decision phase, it would seem that the rules apply too late (as the cases to be selected have already been tacitly determined). However, the application of the schemas rules need not be conscious. The violating cases may be selected by an implicit process that matches cases to schema rules. Those cases that match rules with a definite consequent (i.e. "must") will become the focus of subsequent decision and confirmation processes, such as assessing the subjective expected utilities of possible selections. (Other mechanisms are presumably involved in the determination of relevance for non-schematic versions of the selection task.)

O'Brien suggests that some details of our theory remain to be worked out, which is indeed the case. The spirit of our approach is simply that whatever knowledge is typically used to "make sense" of a regulation is included in the relevant schemas. Much work remains to specify what knowledge is in fact used

to "make sense". Our quasi-formalisations of pragmatic schemas have focused on the rules, as these are most directly pertinent to the choices in the selection task. However, the content of schemas is presumably more general than rules. Essentially, a schema is defined by a set of relational roles. For the deontic schemas, major elements would include an authority who establishes a regulation, the parties subject to it (one of which may be the authority), and the actions and states involved in the regulation. We would thus consider it entirely natural to postulate enriched schemas that include such knowledge as the utilities of potential outcomes with respect to the goals of the relevant parties (i.e. costs and benefits, as discussed by Cosmides, 1989). Indeed, it is this type of contextual information that figures prominently in the "rationales" that have been used to increase the likelihood that subjects will map a stated rule into a pragmatic schema (Cheng & Holyoak, 1985, 1989; Gigerenzer & Hug, 1992).

A basic computational problem to be solved in specifying how pragmatic schemas are used is to specify a mechanism that could accomplish the mapping of the specific elements of the stated problem into the roles and relations in the schema (given that the mapping will often be imperfect). A prerequisite of schema use is the mapping process, coupled with the subsequent propagation of systematic role bindings across the schema rules (the "racasting" of rules to which O'Brien refers). In keeping with a suggestion of Evans and Clibbens concerning the possible role of connectionist networks, we would envisage eventual development of a schema-mapping model similar to current computational models of analogical mapping by "soft" constraint satisfaction, based on structured connectionist-style representations (e.g. Holyoak & Thagard, 1989, 1995; Hummel, Burns, & Holyoak, 1994).

Girotto calls attention to apparent variations in response patterns that have been observed across experiments using highly similar materials. Although the basic reversal of the modal selection pattern triggered by perspective shifts appears to be a reliable phenomenon, it is by no means a clean, deterministic result. We agree with Girotto that although caution is warranted (as some of the empirical variations may be attributable to noise in the data), there is reason to suspect that some differences reflect the interplay of additional pragmatic factors, such as perceived probability of various kinds of violations. Although it is difficult at present to make strong predictions, it should be noted that a connectionist-style constraint-satisfaction model of schema mapping would have the flexibility potentially to accommodate such nuances. For example, in the experiment described in our target article, the predicted *p & not-q* selection for the employee-O1 condition was unexpectedly infrequent, because a considerable number of subjects added a third card selection, the *q* case. This case (in which the employee takes a day off) represents a potential violation from the point of view of the employer (rather than the employee). It is possible, given a constraint-satisfaction model of mapping, that some subjects may have generated a "mixed" mapping in which the non-focal point of view was partially

activated in addition to the focal view. (Politzer & Nguyen-Xuan, 1992, have shown that it is possible, although apparently not easy, for the context to guide a subject to consider two different perspectives on a social contract.)

Our conception of pragmatic schemas, like other uses of the schema concept (Rumelhart, 1980), assumes that schemas provide default inferences. For example, the authority who establishes a regulation is by default assumed to have an interest in enforcing it, and the person subject to a regulation is assumed to have a motive for obeying the rule. It is in this sense that we characterise the rules in the permission and obligation schemas as "pragmatically biconditional" at the deontic level (e.g. the precondition in a conditional permission by default is presumed to be a necessary and sufficient condition for establishing the right to take an action). Because schemas support default reasoning, rather than strict deductive reasoning, we see no conflict with Over and Manktelow's emphasis on the importance of non-monotonic inference in everyday reasoning. Defaults can potentially be overridden by more specific information. Thus if a mother has told her son he may go out if he tidies his room, it may be that an apparent violation (the son going out without tidying his room) will be excused if it is covered by some additional rule (e.g. the mother also has said the boy can go out if he does his homework). In this case, the second conditional permission effectively blocks the (default) biconditional interpretation of the first (i.e. providing an alternative precondition removes the default assumption that any one precondition is necessary to establish the right).

The broader conception of pragmatic schemas sketched have suggests how the theory can account for all of the six "unsatisfactory" situations that Over and Manktelow point out may arise in social contracts. Besides the two cases of violations, where one side carries out their side of the agreement and the other does not, it is possible for a party to fail either to enforce duties owed to them or to exercise their own rights (e.g. the mother lets her son go out even though he doesn't tidy his room; the son does tidy his room but doesn't bother to go out). Roughly, the former case is a "defective" contract because it was not enforced (a responsibility typically attached to the authority), and the latter is defective because a party to the contract neglects to attain their apparent goal even though they have earned the right to do so. The deviations from the default inferences generated by the schema provide an understanding of how each situation is defective.

The final two "unsatisfactory" cases discussed by Over and Manktelow are those in which the rule constitutes an unaccepted offer, such that "nothing happens": neither side fulfils their part of the bargain (e.g. the son doesn't tidy his room but instead just sits in it sulking). Unaccepted offers can be readily understood in terms of our (informal) schemas for social contracts. Pragmatically, each party to a contract must perceive their cost as being outweighed by the contingent benefit before they are willing to act. Suppose I have a house to sell and you wish to buy one, but I ask for a higher price than

you are willing to pay. Then there is no sale, as the offer is pragmatically defective; I linger in possession of my unwanted house, and you continue to lack one. As in the more standard cases discussed in connection with the selection task, an enriched set of social-contract schemas can provide an understanding of the various defective situations, and may suggest a range of options for correcting the deficiencies (without directly specifying what action to take). For example, I may lower the asking price for my house; you may turn your attention to other properties on the market.

In summary, many of the questions raised in the commentaries call attention to the need to conceive of pragmatic schemas more broadly as systems of roles and relations, to which specific situations can be mapped by constraint satisfaction. Regulation schemas can provide a default inferences not only about potential violations, but also about the wider range of contractual situations, including both successful and pragmatically defective cases. Such schemas are considerably richer than the small sets of conditional rules that have borne most of the explanatory burden in the target article and other previous work on pragmatic schemas.

Can Pragmatic Schemas be Replaced?

None of the commentaries (with the possible exception of Oaksford & Chater) seem to altogether dismiss the potential usefulness of pragmatic schemas as a theoretical concept; even Johnson-Laird and Byrne offer faint praise when they note that "Our argument does not rule out the existence of pragmatic reasoning schemas . . . ." Most of the commentaries suggest additional theoretical machinery that is arguably needed to augment pragmatic schemas in some way; and in many cases, we find ourselves quite sympathetic to these suggestions. Here we will consider three approaches that might appear to offer replacements (in whole or in part) for pragmatic schemas.

Johnson-Laird and Byrne, proponents of mental-model theory, object to our characterisation of their approach as being based on "content-independent reasoning procedures", and hence incapable of explaining influences of content on reasoning. They counter with examples of how content can influence the construction of counterexamples, a process that in turn appears to govern inferences in many tasks. We agree that counterexample reasoning is a psychologically plausible process. However, it should be noted that the mental-model account of such reasoning has to grapple with an awkward paradox. A conditional rule of the sort used in the selection task admits of a single type of counterexample: a case of p in the absence of q. But the most basic phenomenon that emerges with arbitrary selection rules reveals that this single counterexample is psychologically asymmetrical: people readily see that p implies q, but must struggle mightily to comprehend that the absence of q implies the absence of p. The theory of mental models, which assiduously

avoids postulating rule-like mental entities, requires considerable ingenuity to explain how a single "mental token", the *p & not-q* case, can be psychologically asymmetrical. Indeed, it is not clear how the mental-model theory distinguishes between the meaning of "if", which is asymmetrical with respect to the propositions it connects, versus the meaning of "and", which is symmetrical. The basic asymmetry of the connective "if" seems quite mysterious when described in terms of this theory. (Johnson-Laird & Byrne distinguish "if" from "and" by a contentless "footnote", notated by triple dots; this "footnote" is attached to an initial model of *p* paired with *q* that is identical to that for "and".)

In contrast, the selection-task asymmetry for arbitrary conditional rules can readily be characterised in terms of inference rules of a natural logic: people seem to have modus ponens but lack modus tollens. As O'Brien observes, the theory of pragmatic schemas explains the facilitative impact of deontic content by postulating that people *do* possess content-specific specialisations of tollens (i.e. schema rules P4 and O4). In fact, most of Johnson-Laird and Byrne's discussion appears to miss the question that lies at the heart of the theory of pragmatic reasoning schemas: What knowledge in memory that can be evoked by problem content generates reliable inference patterns? The central point of pragmatic schema theory is to specify the knowledge that allows broad types of situations to "make sense" in a way that allows contact with goal-oriented reasoning procedures (including specification of relevant counterexamples). We do not, of course, pretend to offer a theory of language comprehension; as O'Brien correctly notes, our theory (like his natural-logic theory) proposes that "inference schemas are applied not to the surface-structure forms of linguistic input, but to propositions as interpreted." And as our discussion in the previous section indicates, the formalisation of schematic knowledge and the mechanisms that map specific situations onto schemas remains a project in progress. However, our approach has already generated quite specific and testable predictions concerning such phenomena as the impact of rationales on reasoning, and facilitation of reasoning for extremely abstract versions of permission and obligation rules (Cheng & Holyoak, 1985, 1989; Kroger, Cheng, & Holyoak, 1993)—manipulations designed to reveal the basic content of the knowledge that guides inference patterns. Mental-model theory as developed by Johnson-Laird and Byrne and their colleagues remains silent about such phenomena.

Oaksford and Chater argue that our processing algorithm is a three-stage model that is unnecessarily complex, and could be readily replaced by a two-stage model based more directly on computation of expected subjective utilities. Although we agree that expected utility considerations are important, an examination of the specific proposal made by Oaksford and Chater indicates that they too have missed the central point of an explanation of performance on the selection task. They agree with us that in a first stage people map the given situation into a schema that determines their perspective; however, these critics

wish to do away with the second stage of drawing deontic inferences (i.e. inferring which cases are governed by "must" in the consequent of a matched schema rule), instead proceeding directly to a stage of calculating expected utilities "based on background knowledge" and the output of the first stage. But when they actually illustrate their proposal, it appears that Oaksford and Chater simply assert that certain utility values arise under each perspective, without providing any theoretical account of where these values come from. The answer, from our point of view, is that the utilities are provided in large part by inferences from the matched schema rules, which constitute the "background knowledge" that Oaksford and Chater assume but do not model. In other words, the schema-based inferences (stage 2) serve to inform the reasoner as to what cases count as potential violations of relevant rights and duties. A model that simply assumes the relevant utilities cannot thereby claim greater parsimony.

As we noted in the target article, more work needs to be done to determine how expected utilities may interact with pragmatic reasoning schemas. To a first approximation, our model describes how people decide what situations involve rights and duties. Such knowledge, as indicated earlier, is at a different level than actually deciding what to *do*. Our theory predicts, for example, that people will be able to say which cases could constitute violations, even for cases they deem unnecessary to check due to their low probability (e.g. the 4-year-old customer in a bar; Kirby, 1994). The pragmatic schemas theory also suggests how people are able (at least sometimes) to take the perspective of others (thereby being able to predict, for example, whether a contemplated action will be perceived as violating a right of someone else). Such knowledge can be used to generate a variety of expected utilities that will enter into a decision process to select an action.

As an aside, it would be quite interesting to apply the methodology of Evans (in press) to selection tasks of the sort used by Kirby (1994), in which some perceived violations are not selected due to the low subjective expected utility of doing so. Perhaps some cases of this nature would require relatively long processing times, thereby providing exceptions to the generalisation made by Evans and Clibbens (that only cases to be selected elicit long processing times). Experiments of this sort might help determine whether the components of subjective expected utility—probabilities of the outcomes and the utilities of these outcomes based on knowledge of violations—are more basic than subjective expected utility itself.

A number of commentators (Girotto, Evans, & Clibbens; Johnson-Laird & Byrne) note the recent application to relevance theory to the selection task by Sperber et al. (in press). The latter investigators have provided a "recipe" for facilitation in the selection task (i.e. a high proportion of *p & not-q* choices). The essential ingredients in the recipe are to create a context in which *p & not-q* case is easily represented, and in which knowing whether there are *p & not-q* cases

would have greater "cognitive effects" than knowing whether there are *p* & *q* cases. Sperber et al. (in press) demonstrate that facilitation can be obtained for non-deontic rules that meet these relevance conditions. These results raise the possibility that pragmatic schemas for regulations are too narrow to explain the basis for facilitation in the selection task (although Sperber et al. acknowledge that there may be independent reasons for postulating domain-specific schemas).

We view Sperber et al.'s (in press) work as a significant advance in understanding conditional reasoning. It is indeed the case that deontic schemas can usefully be viewed as a special case of more general conditions for creating explicit representations of potential violations or falsifications of a conditional rule. At the same time, we would emphasise the possible contributions of pragmatic schemas to the establishment of relevance conditions, which extent beyond the deontic cases. (Because we view pragmatic schemas as a broader concept than the deontic schemas alone, we are unperturbed by O'Brien's observation that deontic conditionals were relatively scarce in the *New York Times* on the day he checked.) We noted earlier, in agreement with O'Brien, that the deontic schemas provide domain-specific versions of modus tollens (if *not-q*, then *not-p*). It is reasonable to suppose that any scenario that serves to construct an isomorph of tollens for the subject will thereby highlight the relevance of the *not-q* case. In some cases the resulting "rule" may be ultra-specific, constructed by the context so as to apply only to the exact content of the target problem. But in other cases schemas may encourage construction of tollens-like rules, and hence help to establish the relevance conditions described by Sperber et al. (in press).

The deontic domain may not be the only one in which schemas can provide specialised analogs of tollens. Of the non-deontic rules studied by Sperber et al., several appear to have a causal basis of a deterministic sort, and hence may evoke causal schemas (Cheng & Nisbett, 1993; Kelley, 1972). For example, in their "virgin-mothers problem", the rule to be evaluated (presented in a fairly elaborate context) is, "If a woman has a child, she has had sex." First, the content of this rule is likely to be interpreted in terms of a necessary cause: sex is generally presumed to be a necessary cause of having a child (the effect). The stated rule therefore instantiates a schema rule of the form, "If < effect >, then < cause >," where the cause is presumed to be *necessary* for the effect. This causal schema may encourage instantiation of a tollens-like rule of the form, "If < no cause >, then < no effect >." Second, in the context provided to subjects by Sperber et al., the presumptive necessity of the cause for the effect is questioned (because of the possibility that virgin girls have been artificially inseminated). This feature of the context brings the "no cause" case (a woman who has not had sex) to the focus of attention, providing a match to the tollens-like rule, which will in turn trigger checking for the required "no effect" (i.e. no child).

Other rules used by Sperber et al. (in press) appear to have the form, "If < cause >, then < effect >," with a context that establishes (1) a presumption that the cause is *sufficient* to produce the effect, and (2) reason to doubt this presumption. Such cases may evoke a schema for sufficient causes, encouraging formation of the tollens-like rule, "If < no effect >, then < no cause >." The latter rule would then encourage selection of the *not-q* (here, "no effect" case), as Sperber et al. found.

More generally, we would offer a slight recasting of Sperber et al.'s recipe for success on the selection task. The key ingredients appear to be (1) creating a mental representation of a tollens-isomorph for the stated conditional rule, and (2) providing a reason to expect that the stated rule may be violated or false. The deontic versions of the selection task clearly meet these requirements, as the permission and obligation schemas provide tollens-isomorphs, and the task requirements focus on checking for violations. It is instructive to contrast the rather elaborate scenarios that Sperber et al. (in press) used to obtain facilitation for non-deontic rules with the minimal context required to obtain facilitation for an abstract permission rule, "If one is to take action 'A', then one must first satisfy precondition 'P' " (Cheng & Holyoak, 1985, Girotto, Mazzocco, & Cherubini, 1992; Griggs & Cox, 1993; Kroger et al., 1993). Evoking the permission schema does not require much contextual support, and its evocation directly satisfies the relevance requirements.

But as these examples illustrate, the relevance requirements can also be met by certain types of causal contexts, which may also evoke pragmatic schemas. It is important to note, however, that these causal cases are severely restricted, depending on the causal relation being hypothesised to be deterministic (either a necessary cause stated as "If < effect >, then < cause >", or a sufficient cause stated as "If < cause >, then < effect >"), coupled with a context that makes a violation of the relation appear plausible. Contrary to these restrictions, most causal relations outside the physical domain are understood to be non-deterministic. Moreover, when causal relations are understood to be deterministic, they are often assumed never to be violated. In the more general probabilistic case, a causal schema will actually block generation of a tollens-isomorph. For example, the causal conditional "If one smokes, then one gets cancer" cannot be falsified by any single counterexample (i.e. a smoker who does not get cancer), and does not license the tollens-isomorph "If one does not get cancer, then one does not smoke." Not only does the tollens-isomorph fail to apply because the stated causal rule is interpreted probabilistically, but it is pragmatically anomalous because the "if–then" frame imposes a forward temporal directionality (Evans & Newstead, 1977), thereby suggesting that not getting cancer somehow causes non-smoking (Cheng & Holyoak, 1985).

Except in very special cases, then, causal schemas will not "facilitate" performance in the selection task, in the sense of encouraging selection of the

create or otherwise regulate rights and duties, whereas a liability implies being at risk of having a duty imposed by someone else. The "authority" in a regulation schema has the power to create a regulation, to which various parties (often including the authority) are then liable. For example, it is natural to assume that it is the employer who has the power to actually create the day-off rule, which constitutes an "offer" to the employee. (Outside the context of collective bargaining, an individual employee presumably lacks the power to create a rule concerning days off.) Once the regulation is created, however, both parties are interlocked by power and liability relations. The employee, for example, may at any time accept the employer's offer by electing to work on the weekend (thereby exercising a power created by the regulation); by the same token, the employer has a liability flowing from the regulation (namely, the risk of incurring a duty to provide the employee with a day off during the week, which will be triggered whenever the employee chooses to work on the weekend).

The concepts of deontic power and liability (and their respective opposites, which Hohfeld termed *disability* and *immunity*) appear useful in characterising when a contract is well-formed and enforceable. For example, someone who possesses stolen goods (even unknowingly) lacks the power to sell them; because of this disability, any sales agreement made in such circumstances is unenforceable. Everyday contractual disputes are based not only on claims of violation of rights and neglect of duties, but also on claims related to the prerequisite conditions of power and liability. Suppose that the mother discussed by Manktelow and Over (1991) had the audacity to say, "If you go out, you have to tidy your room first", to her teenage son. Besides all the six "unsatisfactory" outcomes described by Over and Manktelow, yet another type of rule failure would be exemplified by the response, "You don't run my life!" (a denial of the mother's power to make the rule, tantamount to a claim of immunity from it). To which the mother might retort, "You don't pay rent for your room!", thereby challenging the claim of immunity, and so on. A complete theory of social contracts will need to provide analyses of the felicity conditions for establishing rules, as well as of the conditions under which rules are followed or violated.

Several commentaries argue that the elements of deontic schemas need to be defined in terms of more basic concepts. Both Over and Manktelow and Johnson-Laird and Byrne note that the rules in our permission and obligation schemas make use of undefined modal terms, "may" and "must", and suggest that deeper definitions can be based on the logical concepts of "possibility" and "necessity". As Johnson-Laird and Byrne point out, the latter concepts have both epistemic and deontic senses, which appear to be systematically related. Over and Manktelow argue that we were premature in dismissing deontic logic as a theoretical basis for reasoning about regulations. O'Brien argues that pragmatic schemas for regulations are very limited in scope (and that regulations are ecologically infrequent content for conditionals), and that a natural logic for conditionals is more basic.

We agree that the identification of foundational concepts that may underlie human reasoning would be extremely valuable. We do have reservations, however, regarding whether contextual influences on reasoning will be finally understood in terms of some variant of formal logic. In the case of deontic logic, our objection is not to the use of the concepts of possibility and necessity, but rather to various casual suggestions in the literature to the effect that some already formulated deontic logic readily explains human deontic reasoning. As we argued in our target article, no existing deontic logic provides an account of the contextual constraints on regulations and social contracts that seem to be central to the permission and obligation schemas. Over and Manktelow, responding to our observation that certain deontic logics preclude the possibility of conflicting obligations, note that "minimal" deontic logics are compatible with such conflicts. This is indeed the case, but even these weaker deontic logics axiomatically exclude the possibility of an obligatory act that is epistemically impossible to perform (Chellas, 1980, p.202). However, we can readily imagine "Catch-22" situations in which an authority imposes a duty that is impossible to fulfil. It seems that such situations are perceived as unreasonable rather than illogical.

In any case, as Over and Manktelow point out, deontic logics only provide information about the validity of inferences, rather than the selection of actions. Furthermore, the list of psychologically plausible "natural" deontic inferences seems quite short. Thus in addition to the biconditional converse implication between a right and a duty that we stated in our target paper, we are happy to acknowledge such plausible deontic inferences as that having a duty to perform an action implies having a right to do so (with respect to the other party to the relevant contract, but not necessarily with respect to everyone); that having a duty to perform an action implies not having the right not to perform it; and that having a right to perform an action implies not having a duty not to perform it. If deontic logic is to have more to offer as an account of human understanding of permission and obligation, someone will have to develop a model of "natural deontic logic" with greater psychological relevance than anything so far proposed in the psychological literature.

In contrast to the rather ephemeral proposals to harness deontic logic as a psychological model, O'Brien and his colleagues have made highly specific and testable proposals about natural propositional logic. For conditionals, Braine and O'Brien's natural logic includes two basic components: modus ponens and the schema for conditional proof. O'Brien takes the position that pragmatic schemas and other content-based inference mechanisms operate to supplement, but never to replace, content-free natural logic. We have noted that the theory of pragmatic schemas is compatible with a core natural logic (Cheng et al., 1986). However, we would also repeat our point that the basic pragmatic constraint of relevance permeates human reasoning so deeply that it is doubtful that a psychologically plausible natural logic can be encapsulated from pragmatic considerations. For

example, O'Brien defines the schema for conditional proof as follows: "CP holds that to derive or infer *if p then* . . . , first suppose *p*; for any proposition *q* that follows from the supposition *p* taken together with other information assumed, one may assert *if p then q*." It follows that people ought to readily assent to the proposition that "If humans have three heads, then 13 is a prime number" (because if we suppose as *p* that humans have three heads, then the stated *q* surely follows from *p* "taken together with other information assumed"). This seems to us a dubious psychological prediction. The general problem for O'Brien's schema for conditional proof, and for all formal treatments of the conditional, is that people seem to demand that *p* appear *relevant* to *q* in some way before they accept the truth of *if p then q*. In so far as natural logics fail to provide an account of the relevance requirement, it seems that pragmatic considerations can override, as well as supplement, the inferences that a natural logic licenses.

Finally, both O'Brien and Gigerenzer draw attention to perhaps the most difficult question facing theories of human reasoning: What is the origin of inference rules at different levels of abstraction? As O'Brien notes, we have tended to view pragmatic schemas as emerging by inductive generalisation from types of goal-related experiences. O'Brien raises the alternative view that the rules of a domain-independent natural logic are first acquired (perhaps due to maturation), with content-specific schemas being added later as specialisations. Gigerenzer emphasises that innate constraints may govern the acquisition of "modules" for reasoning about particular content domains.

These are intriguing suggestions, but it seems that the origins of reasoning mechanisms remain largely unexplored. The most general problem facing schema theories in psychology, including our own, is that we lack adequate models of how schemas based on relational predicates can be induced from experience. One might content that reasoning schemas are really innate modules, but it is hardly plausible to suppose that the schemas acquired by expert physicists, for example, are innate. However we lack good models of the induction of problem schemas as well as of reasoning schemas. If we assume that people in fact use inductive mechanisms to form problem schemas (while acknowledging we do not yet understand these mechanisms well), it remains reasonable to consider how such inductive mechanisms might also operate to produce content-specific reasoning schemas.

At the same time, there are good reasons to suppose that constraints at various levels of generality (some innate, some the currently available products of prior experience) play important roles in human induction (Holland, Holyoak, Nisbett, & Thagard, 1986). In fact, Gigerenzer's conception of modules as hierarchically organised sets of procedures seems quite consistent with Holland et al.'s conception of rule-based default hierarchies. Given his analysis of deontic situations in terms of elementary dimensions of similarity and difference, Gigerenzer's "modules" seem much more like overlapping schemas than like the

encapsulated Fodorian notion of modules. But however reasoning schemas or modules are acquired, O'Brien, Gigerenzer, Johnson-Laird and Byrne, and Over and Manktelow are quite correct that some representational elements must serve as primitives. Concepts such as "cause" and "effect", "necessity" and "possibility", "rights" and "duties" may provide some of these primitives, or they may be derived in turn from yet more fundamental elements. These basic questions remain open.

Despite the huge gaps that remain in our understanding of everyday reasoning—indeed, to some extent because of them—we view the field as a whole as entering an exciting period for new empirical and theoretical developments. The recent flurry of findings on perspective effects, the influence of subjective utilities, and general prerequisites for triggering specific selection patterns, are providing valuable stimulation and guidance for further research. The ideas in the commentaries we have discussed provide an excellent sample of some of the theoretical directions that deserve to be pursued vigorously. We echo Gigerenzer's closing admonition: "Let the work begin."

ACKNOWLEDGEMENTS

Preparation of this paper was supported by NSF Grants SBR-9310614 to Keith Holyoak and DBS-9121298 to Patricia Cheng. Requests for reprints may be directed to either author at Department of Psychology, University of California, Los Angeles, California 90095-1563, USA.

REFERENCES

Chellas, B.F. (1980). *Modal logic: An introduction*. Cambridge: Cambridge University Press.

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

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

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

Cheng, P.W., & Nisbett, R.E. (1993). Pragmatic constraints on causal deduction. In R.E. Nisbett (Ed.), *Rules for reasoning* (pp.207–227). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Cheng, P.W., & Novick, L.R. (1992). Covariation in natural causal induction. *Psychological Review, 99*, 365–382.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition, 31*, 187–276.

Evans, J.St.B.T. (in press). Deciding before you think: Relevance and reasoning in the selection task. *British Journal of Psychology*.

Evans, J.St.B.T., & Newstead, S.E. (1977). Languge and reasoning: A study of temporal factors. *Cognition, 8*, 387–390.

Fillenbaum, S. (1976). Inducements: On phrasing and logic of conditional promises, threats and warnings. *Psychological Research, 38*, 231–250.

Gigerenzer, G., & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition, 43*, 127–171.

Girotto, V., Mazzocco, A., & Cherubini, P. (1992). Judgements of deontic relevance in reasoning: A reply to Jackson and Griggs. *Quarterly Journal of Experimental Psychology, 44A*, 469–482.

Griggs, R.A., & Cox, J.R. (1982). The elusive thematic materials effect in Wason's selection task. *British Journal of Psychology, 73*, 407–420.

Griggs, R.A., & Cox, J.R. (1993). Permission schemas and the selection task. *Quarterly Journal of Experimental Psychology, 46A*, 637–651.

Hohfeld, W.N. (1919). Some fundamental legal conceptions as applied in judicial reasoning. In W.W. Cook (Ed.), *Some fundamental legal conceptions as applied in judicial reasoning and other legal essays, by Wesley Newcomb Hohfeld* (pp.23–64). New Haven, CT: Yale University Press.

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

Holyoak, K.J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295–355.

Holyoak, K.J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.

Hummel, J.E., Burns, B., & Holyoak, K.J. (1994). Analogical mapping by dynamic binding: Preliminary investigations. In K.J. Holyoak & J.A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2: Analogical connections* (pp.416–445). Norwood, NJ: Ablex.

Johnson-Laird, P.N., Legrenzi, P., & Legrenzi, S.M. (1972). Reasoning and a sense of reality. *British Journal of Psychology, 63*, 395–400.

Kelley, H.H. (1972). Causal schemata and the attribution process. In E.E. Jones et al. (eds), *Attribution: Perceiving the cause of behavior* (pp.151–174). Morristown, NJ: General Learning Press.

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

Kroger, K.K., Cheng, P.W., & Holyoak, K.J. (1993). Evoking the permission schema: The impact of explicit negation and a violation-checking context. *Quarterly Journal of Experimental Psychology, 46A*, 615–635.

Legrenzi, P., Politzer, G., & Girotto, V. (in press). Contract proposals: A sketch of a grammar. *Theory and Psychology*.

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

Politzer, G., & Nguyen-Xuan, A. (1992). Reasoning about promises and warnings: Darwinian algorithms, mental models, relevance judgments or pragmatic schemas? *Quarterly Journal of Experimental Psychology, 44A*, 402–421.

Rumelhart, D.E. (1980). Schemata: The building blocks of cognition. In R.J. Spiro, B.C. Bruce & W.F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp.33–58). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Sperber, D., Cara, F., & Girotto, V. (in press). Relevance theory explains the selection task. *Cognition*.

Tversky, A., & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp.49–72). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Waldmann, M.R., & Holyoak, K.J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*, 222–236.

Waldmann, M.R., Holyoak, K.J., & Fratianne, A. (1995). Causal models and the acquisition of category structure. *Journal of Experimental Psychology: General, 124*, 181–206.

Wason, P.C. (1966). Reasoning. In B.M. Foss (Ed.), *New horizons in psychology* (Vol. 1). Harmondsworth, UK: Penguin.

Wason, P.C., & Shapiro, D. (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology, 23*, 63–71.