The Taming of Content: Some Thoughts About Domains and Modules

Gerd Gigerenzer, Max Planck Institute for Psychological Research, Munich

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

Peter Wason invented the Selection Task in 1966. Thirty years and many, many experiments later, two results are evident for me. First, the view that sound reasoning can be reduced to propositional logic (or first order logic) is myopic. Human thought operates in more dimensions than entailment and contradiction (Strawson, 1952). We need to work out how the mind infers the meaning of a conditional statement from the content of the Ps and Qs and the social context in which the statement is uttered, rather than exclaiming "Cognitive illusion! Hurrah! Error!" whenever human reasoning cannot be reduced to propositional logic. Second, the hope that these "errors" or their flip-side, the "facilitation" of logical reasoning, would be the royal road to discovering the laws of human reasoning did not materialise. This hope was fueled by the (misleading) analogy with visual illusions (Gigerenzer, 1991). What were seen as "errors" were attributed to deeper cognitive deficits such as confirmation biases, to crude heuristics such as availability, or simply to "the subjects' incorrigible conviction that they are right when they are, in fact, wrong" (Wason, 1983, p.356). Unfortunately, this programme of research has brought little progress on the theoretical front.

In the last decade, a few adventurous researchers freed themselves from the straitjacket of propositional logic and looked at dimensions of reasoning beyond entailment and contradiction. The content rather than the logical structure of the conditional statement moved into the foreground. As early as 1972, Wason and Johnson-Laird pointed out (p.245) that, contrary to their expectation, "content is crucial" to reasoning and that "any general theory of human reasoning must include an important semantic component". But neither they nor the others who studied human reasoning at that time found a way to include semantics in a theory of reasoning. Instead, content remained but a decorative element in reasoning problems—which either might "confuse" subjects or "facilitate" their logical reasoning. In the 1980s, Patricia Cheng (Cheng & Holyoak, 1985; Cheng, Holyoak, Nisbett, & Oliver, 1986) and Leda Cosmides (1989; Cosmides & Tooby, 1989, 1992) dared to introduce the content of the Ps and Qs into their theories. Cheng and her collaborators postulated a set of pragmatic reasoning schemas, such as permission and obligation schemas. Cosmides and Tooby postulated Darwinian algorithms, with social contracts and threats as examples. Content, the *terra incognita* where no established researcher on reasoning dared or cared to venture, became a legitimate topic of study. Cheng and her colleagues made a significant first proposal: permission and obligation schemas, each defined by four production rules. This was an important move away from propositional logic, but the production rules still resembled the four rules of the truth table, with Ps and Qs replaced by "actions" and "preconditions". Cosmides and Tooby made a bolder and theoretically richer leap, connecting information search with pragmatic goals such as cheater detection, with cost–benefit analyses, and with the broader evolutionary theory of reciprocal altruism. These researchers set the stage for the discovery of a genuinely pragmatic dimension of reasoning unknown to logic: a person's perspective (Gigerenzer & Hug, 1992; Light, Girotto, & Legrenzi, 1990; Manktelow & Over, 1991).

The semantic and pragmatic approach to reasoning, however, is still in its infancy: researchers often continue to focus on the "logically correct answer", to report their results in terms of "logical facilitation", or to criticise a competing pragmatic theory by saying that its semantic conditionals are simply logical biconditionals (e.g. Politzer & Nguyen-Xuan, 1992). Too much energy has been spent in the debate on asserting who is wrong, and too little on actually developing the theories. But whether a theory is "right" or "wrong" cannot be decided unless it is well specified in the first place. For instance, one group now needs to work out exactly what the Darwinian algorithms are, and another group needs to put pragmatic contents—a person's point of view, their goals, their cost–benefit computations among others—explicitly into the production rules that define permission and obligation schemes. Unless this theoretical specification is accomplished, one can only make plausible arguments that one's theory *could* be consistent with the perspective effect. For instance, Holyoak and Cheng's account, in this issue, of the perspective effect along the lines of Politzer and Nguyen-Xuan (1992) is plausible, but the pragmatic reasoning schema theory is not richly specified enough to allow for a deduction of this effect. The challenge is not to slip content and perspective effects in through the back door, but to devise richer theories that allow semantic and pragmatic dimensions in through the front door.

So, how can we make progress in the study of reasoning? How can we improve on the first ideas suggested by Cheng and Holyoak, and Cosmides and Tooby? I will not join the battle over who is right. Instead, I will try to work out some of the unresolved theoretical issues and suggest how these could be solved.

BEYOND "LOGICAL FACILITATION": WHERE TO GO FROM HERE?

First we need to get clear what we were studying when we discovered perspective effects. What I saw as the issue was reasoning about conditional statements in natural languages—not deductive or logical reasoning. If one wanted to study logical reasoning, one would have to instruct subjects: "Forget the content and just treat this problem as an exercise in propositional logic." But this is not what has been done in the last 30 years. Given the way this research has been conducted, the question to ask is: How does a mind understand a natural language conditional, and what cognitive processes and behaviours are controlled by that understanding?

I start with two assumptions. First, understanding a conditional statement is impossible without analysing its semantic content and the pragmatic context in which the statement is uttered. Or, as Fillenbaum (1977) put it, "Mind your Ps and Qs". The second assumption is that minds reason about conditional statements neither at the level of the particular nor at the abstract level of propositional logic, but at an intermediate level. This intermediate level of abstraction retains enough of semantics and pragmatics but at the same time discards enough particulars to allow for both rich and fast cognitive processes and behaviour. What I call domain-specific reasoning is carried out at this intermediate level of abstraction.

Here is an example. Suppose you say to me: "If you touch me, I'll kill you". I infer that if I don't touch you, you won't kill me. Thus, from "if P then Q", I infer "if not-P then not-Q". This is logically invalid, but it may save my life. Threats, warnings, social contracts, bribes, and many other classes of conditional statement do not follow propositional logic. They have their own "natural logics". Note the plural. The term "natural logics" refers to how people reason with or in natural language (Fillenbaum, 1976, 1977). For instance, linguists study paraphrases and "invited inferences" (as opposed to truth-preserving logical inferences). A conditional threat can be paraphrased with "P and Q" as well as with "not-P or Q". You could as well have said "Touch me *and* I kill you" or "Don't touch me *or* I kill you". These paraphrases bypass propositional logic in that the English "or" and "and" do not map onto the logical OR and AND. Now take a social contract: "If you mow my lawn, I'll give you $50". The social contract can be paraphrased by "P and Q", but not by "not-P or Q", as was the case with the threat. "Don't mow my lawn *or* I'll give you $50" is not an acceptable paraphrase. These examples illustrate that invited inferences and paraphrases systematically vary between domains, such as conditional threats, warnings, bribes, and social contracts (Fillenbaum, 1976, 1986).

It is not just paraphrases that are specific to domains. What would a research programme that studied domain-specific reasoning look like?

MENTAL MODULES: A RESEARCH PROGRAMME

Here I propose a two-step programme for studying reasoning about conditional statements (and beyond). The first step is to model the mapping algorithm that recognises a particular statement as falling into the range of a domain; the second step is to work out the subsequent activity of the module: how it deploys attention, makes inferences, and executes other processes.

1. Search for the Mapping Algorithm

Here is the first problem to be solved: There is a potentially unlimited number of conditional statements C_i (in English or any other language) and a smaller number of domains D_j. What is the algorithm that people use to map C_i onto D_j? Note that we have simplified the problem by only considering linguistic information. Once we have solved this problem, we may go on to tackle the issue of how nonverbal information—such as facial cues that signal a threat (Ekman, 1982) or motion cues signalling a self-propelled actor (Premack, 1990)—is used in parallel with verbal information.

Our problem is like solving one equation with two unknowns. We know the C_i but not the mapping algorithm or the domains. One strategy is to assume a value for one of the two unknowns and try to solve for the other. Let us assume a few domains D_j and solve for the algorithm. I suggest proceeding in this way because we have several converging proposals for candidate domains, but I do not know of any specified proposals for mapping algorithms.

What are plausible candidates for domains? With respect to conditional statements, linguists have distinguished inducements, such as conditional promises and bribes, and deterrents, such as conditional threats and warnings (Fillenbaum, 1977, 1986). Evolutionary psychologists focusing on important adaptive problems in the history of humans have proposed strikingly similar candidates; social contracts and social exchange, precautions, and threats (Cosmides & Tooby, 1992). Cognitive psychologists have postulated permissions and obligations (Cheng & Holyoak, 1985) as candidate domains, which include social exchange but also other forms of social regulations and deontic reasoning (Over & Manktelow, 1993). No complete list of domains is known, nor is one needed; it is sufficient to start analysing a few.

If we provisionally commit ourselves to a group of domains then we can ask: What is the mapping algorithm? That is, how does a mind infer that a particular semantic conditional is a threat and not a social contract or something else? Here is an idea. I conjecture that only a small number of dimensions are needed to make fairly reliable classifications of individual English conditionals into domains. For instance, all statements of the kind "if P then Q" in which you issue a threat to me seem to have in common that (i) P is an action of mine, (ii) Q is an action of yours, (iii) both actions are possible in the future, (iv) P is first and Q second, and (v) Q has negative consequences for me. This is most likely not the fullest possible characterisation of a conditional threat, but with a few hundred examples of conditional threats we could find the dimensions I have overlooked. Now consider the social contract: "If you mow my lawn, I'll give you $50". The values on the dimensions (i), (ii), and (iii) are as characteristic of a social contract as of a threat and would not distinguish between them. The strict temporal order in (iv) is not characteristic of a social contract because it may or may not hold: you may pay me in advance or after the work. The decisive difference is in feature (v). If your future action Q is a benefit for me, then I can infer that the statement is a social contract rather than a conditional threat. As a last example, imagine you issue a warning to me: "If you give him your finger, he'll take your hand". Warnings share with threats the values (i), (iii), (iv), and (v), but differ on (ii). Q is typically not an action of yourself but rather an action of a third party (as in the last example) or a negative consequence for me ("If you ski in this snow, you may break a leg"). These are some candidate dimensions from which a richer account of pragmatic reasoning schemas might be constructed.

If this view is correct, the algorithm for mapping conditional statements C_i onto domains D_j could be modelled analogously to a key, where the teeth are the values on the dimensions. Each threat, each social contract, and so on has the same characteristic profile, and the mapping algorithm uses the dimensions on which they differ to discriminate between them. Of course, there will be some residual uncertainty, which may be further reduced by non-verbal information. The challenge is to examine hundreds of conditional statements to improve on the preliminary sketch I have given here. Working models of the mapping algorithm can subsequently be tested in the form of computer programs together with a knowledge component, e.g. for defining positive and negative consequences. My intuition is that the mapping algorithm uses no more than a half-dozen content dimensions of the Ps and Qs to infer a domain.

2. The Control Structure of Cognitive Modules

Once the mapping algorithm has inferred a domain (e.g. "threat!"), a "cognitive module" that controls the processes necessary for coping with this type of situation is activated. The advantages of cognitive modules that respond to the social and physical environment at the intermediate level of a domain (rather than at the level of the particular or of logical abstraction) are the following. First, attention can be focused. For instance, if the situation is identified as a social contract, then attention can be directed to information that could reveal that one is being cheated (Cosmides, 1989; Gigerenzer & Hug, 1992). If, however, a threat is identified, information that could reveal being cheated needs no attention, but information that can reveal being bluffed or double-crossed does. Thus, modules can help to reduce one fundamental problem of induction: What to observe? (Popper, 1959). Second, inferences can be made more efficiently. Modules that contain semantic and pragmatic structures enable the organism to react quickly and to reduce the problem of computational explosion when making inferences. For instance, John Garcia (e.g. Garcia & Koelling, 1966; Garcia y Robertson & Garcia, 1985) showed that when the taste of flavoured water was paired with experimentally induced nausea, rats could learn in just one trial to avoid the flavoured water, even if the nausea occurred two hours later. In contrast, when the taste of flavoured water was repeatedly paired with electric shock, rats had great difficulty learning to avoid the water. Thus, rats seem to have a built-in specific mechanism for food avoidance that *enables* them to develop some associations rapidly but not others. Note that such specific mechanisms are often phrased in terms of "constraints" on inference or learning. The more appropriate statement seems to me that these mechanisms enable, not constrain, inference. Semantic relations built into mechanisms can enable what "unconstrained", that is, content-independent, algorithms could not do in the first place because combinatorial explosion might paralyse any system that is truly domain-general (Cosmides & Tooby, 1994). Third, learning could hardly proceed without domain-specific devices—an argument that Chomsky made against Skinner's view of language learning. Garcia's experiments are further examples in support of this argument.

The second problem that needs to be solved, then, is how to model the control structure of a module? Here are some thoughts.

In order to design the structure of a module, we need to define more clearly the notion of a module. (In the literature, almost everything has been suggested as a domain to which some module is attached; see Hirschfeld & Gelman, 1994). Fodor (1983) distinguished "horizontal" from "vertical" faculties. Current curricula and textbooks typically organise cognitive psychology according to the doctrine of horizontal faculty psychology, assuming that the mind is divided up into general-purpose processes: memory, attention, thinking, judgment, perception, volition, and so forth. All memories are in the same place; they may depend on time and rehearsal, but not on the *content* of the memory. Similarly, in this view all thoughts, judgments, and so forth are of one kind across content domains. For the "vertical" view of faculties, however, there are no such things as memory, attention, thinking, judgment, perception, volition, and so forth. Instead there are domain-specific capacities, each with different mechanisms in which the horizontal faculties are like parts of an engine. The modules we are searching for are domain-specific and are thus vertically, not horizontally, organised.

For Fodor there exist about six input systems, one for each of the five senses and one more for language. He envisions modules as functionally more specific than these six systems, including, for instance, modules designed for colour perception, analysis of shape, recognition of faces, and recognition of voices. Fodor's focus on input systems makes him believe that modules are informationally encapsulated.

Where I depart from Fodor is when he restricts modules to informationally encapsulated input systems, arguing that central cognitive processes such as thinking tend to be domain-general. The mind in Fodor's *Modularity of mind* (1983) is decidedly anti-modular at its centre. The stronger modularity thesis is that central cognitive processes like thoughts are modular to some important degree, too. Evidence for this stronger thesis stems from studies of cognition and development (e.g. Barkow, Cosmides, & Tooby, 1992; Goldstein & Weber, in press; Leslie, 1994; Premack & Premack, 1994; Sperber, 1994; Todd & Miller, 1993) and from selective impairments following brain damage (e.g. Caramazzo, Hillis, Leek, & Miozzo, 1994). This descriptive evidence is complemented by the normative argument that content-independent formulations of principles of rational reasoning and decision making, such as consistency, are essentially confused, in the sense that there is no way to determine whether choice is consistent or rational without referring to something external to choice behaviour, such as a person's motives and values (e.g. Gigerenzer, in press; Sen, 1993).

Furthermore, if we think of modules as hierarchically organised, then modules can activate and inhibit one another and the notion of encapsulation is misleading. Take for instance a sequential, tree-like arrangement of modules and mapping algorithms. Assume you are out in the woods at night, it is windy, and you notice at some distance a large, dark, moving object. The postulated mapping algorithm in your brain would analyse the motion pattern to classify the object as either "self-propelled" (animal or human) or not self-propelled (plant or physical object), to use David Premack's terms. If the algorithm infers that there is a self-propelled object, a module for unrecognised self-propelled objects may be activated, which initiates physiological, emotional, and behavioural reactions that alert the organism and direct attention to information that can reveal whether it is human or animal. A second, more specialised mapping algorithm may infer from shape and motion information that the object is human. This inference may in turn activate a module lower in the hierarchy that initiates reactions appropriate for unidentified humans (warnings, threats, and the like) and guides attention to information that could reveal whether that human is friend or enemy (or predator or prey, in the case of an animal), and so on. Note that such hierarchically organised, specialised modules can act quickly, as only a few branches of the hierarchical tree are travelled. The organism thereby avoids combinatorial explosion. For instance, if the first mapping algorithm had inferred that the object is not self-propelled, then any information that could reveal whether it is a human or animal, and subsequently, friend or enemy, or predator or prey, would not need to be searched for.

There need not be a one-to-one correspondence between domain and module. The domain of a module may shift over evolutionary time. Sperber (1994) tries to capture this phenomenon with his distinction between the "proper" and the "actual" domain of a module. The proper domain of a module designed for social contracts may once have been exclusively the small hunter-gatherer group in which trust had to be established, forgiveness granted, and repeat-cheaters expelled. The "actual" domain today, in contrast, may range from the task of managing a large company to board games like *Diplomacy*.

THE TAMING OF CONTENT

In his seminal book, *The taming of chance* (1990), Ian Hacking described how chance was tamed by statistical laws in the sense that chance became the very stuff of the fundamental processes of nature and society. Statistical mechanics, quantum theory, and evolutionary theory epitomise that revolution. Chance was no longer the essence of the lawless and unpredictable. In theories of reasoning, content has played a role similar to the one chance once played in theories of nature: something to ignore or banish. The taming of content by a new class of theories of reasoning can fundamentally change our understanding of the mind, just as the probabilistic revolution once changed our understanding of nature and society (Gigerenzer et al., 1989).

But we are only beginning to build theories that model the role content plays in reasoning: domains and modules mark only a first step. Cognition is still studied by most psychologists in terms of what Fodor calls horizontal faculties: deductive thinking, probabilistic reasoning, problem solving, and so on. And the preferred models are propositional logic, probability theory, or variants thereof that ignore the content and pragmatics of thought. In this article, in contrast, I have sketched a two-step programme for studying modular thought. I have not provided an answer to the questions I proposed to study, but have at least tried to define these questions and to outline possible research strategies. Let the work begin.

ACKNOWLEDGEMENTS

I am grateful to Valerie Chase and Lorraine Daston for many helpful comments on an earlier version of this paper.

REFERENCES

Barkow, J.H., Cosmides, L., & Tooby, J. (Eds.), (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford: Oxford University Press.

Caramazza, A., Hillis, A., Leek, E.C., & Miozzo, M. (1994). The organization of lexical knowledge in the brain: Evidence from category- and modality-specific deficits. In L.A. Hirschfeld & S.A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture*. 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., Nisbett, R.E., & Oliver, L.M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology, 18*, 293–328.

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

Cosmides, L. & Tooby, J. (1989). Evolutionary psychology and the generation of culture, Part II. A computational theory of social exchange. *Ethology and Sociobiology, 10*, 51–97.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J.H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163–228). Oxford: Oxford University Press.

Cosmides, L., & Tooby, J. (1994). Origins of domain specificity: The evolution of functional organization. In L.A. Hirschfeld & S.A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture*. Cambridge: Cambridge University Press.

Ekman, P. (Ed.) (1982). *Emotions in the human face* (2nd edn.). Cambridge: Cambridge University Press.

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

Fillenbaum, S. (1977). Mind your *p's* and *q's*: The role of content and context in some uses of *and, or*, and *if*. *The Psychology of Learning and Motivation, 11*, 41–100.

Fillenbaum, S. (1986). The use of conditionals in inducements and deterrents. In E.C. Traugott, A. ter Meulen, J.S. Reilly, & C.A. Ferguson (Eds.), *On conditionals*. Cambridge: Cambridge University Press.

Fodor, J.A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Garcia, J., & Koelling, R.A. (1966). The relation of cue to consequence in avoidance learning. *Psychonomic Science, 4*, 123–124.

Garcia y Robertson, R., & Garcia, J. (1985). X-rays and learned taste aversions: Historical and psychological ramifications. In T.G. Burish, S.M. Levy, & B.E. Meyerowitz (Eds.), *Cancer, nutrition and eating behavior: A biobehavioral perspective* (pp. 11–41). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Gigerenzer, G. (1991). On cognitive illusions and rationality. *Poznan Studies in the Philosophy of the Sciences and the Humanities, 21*, 225–249.

Gigerenzer, G. (in press). Rationality: Why social context matters. In P. Baltes & U.M. Staudinger (Eds.), *Interactive minds. Life-span perspectives on the social foundations of cognition*. Cambridge: Cambridge University Press.

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

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L.J., Beatty, J., & Krüger, L. (1989). *The empire of chance: How probability changed science and everyday life*. Cambridge: Cambridge University Press.

Goldstein, W.M., & Weber, E. (in press). Content and discontent: Indications and implications of domain specificity in preferential decision making. In J.R. Busemeyer, R. Hastie, & D.L. Medin (Eds.), *The psychology of learning and motivation*. New York: Academic Press.

Hacking. I. (1990). *The taming of chance*. Cambridge: Cambridge University Press.

Hirschfeld, L.A., & Gelman, S.A. (Eds.). (1994). *Mapping the mind; Domain specificity in cognition and culture*. Cambridge: Cambridge University Press.

Leslie, A.M. (1994). ToMM, ToBY, and agency: Core architecture and domain specificity. In L.A. Hirschfeld & S.A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture*. Cambridge: Cambridge University Press.

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

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

Over, D.E., & Manktelow, K.I. (1993). Rationality, utility and deontic reasoning. In K.I. Manktelow & D.E. Over (Eds.), *Rationality: Psychological and philosophical perspectives* (pp. 231–259). London: Routledge.

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.

Popper, K.R. (1959). *The logic of scientific discovery*. London: Hutchinson.

Premack, D. (1990). The infant's theory of self-propelled objects. *Cognition, 36*, 1–16.
Premack, D., & Premack, A.J. (1994). Moral belief: Form versus content. In L.A. Hirschfeld & S.A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture*. Cambridge: Cambridge University Press.
Sen, A. (1993). Internal consistency of choice. *Econometrica, 61*, 495–521.
Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp.39–67). Cambridge: Cambridge University Press.
Strawson, P.F. (1952). *Introduction to logical theory*. New York: Wiley.
Todd, P.M., & Miller, G.F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior, 2*, 5–47.
Wason, P.C. (1983). The importance of cognitive illusions. *Behavioral and Brain Sciences, 4*, 356.
Wason, P.C., & Johnson-Laird, P.N. (1972). *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.

Contextual Factors in Deontic Reasoning

Vittorio Girotto, University of Trieste and CREPCO, Aix-en-Provence

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

The discovery of perspective effects is one of the most interesting findings in the recent literature on the selection task. Several investigators have found that people's selections can vary as a function of the specific perspective that they are asked to take (cf. Gigerenzer & Hug, 1992; Light, Girotto, & Legrenzi, 1990; Manktelow & Over, 1991; Politzer & Nguyen-Xuan, 1992). Consider a conditional contractual regulation of the form:

If you perform action A1 for me, I'll perform action A2 for you

in which the two actions increase the positive utility of the two target actors[1]. Given the task of checking for the possible violation of the rule from the hearer's perspective, people will tend to select the "action A1 performed" (*p*) and "action A2 not performed" (*not-q*) cases. In other words, (using the terminology proposed by Holyoak & Cheng, this issue; henceforth H&C), they will tend to select the cases corresponding to the "duty of the hearer towards the speaker fulfilled" and the "right of the hearer against the speaker acquired". By contrast, if the same rule is checked from the speaker's point of view, people will

[1]For an analysis of the entire set of contract proposals, including those expressing negative utilities, cf Legrenzi, Politzer, & Girotto (in press).