

Putting *If*s to Work: Goal-Based Relevance in Conditional Directives

Denis J. Hilton
Université Toulouse 2

Markus Kimmelmeier
University of Nevada

Jean-François Bonnefon
Université Toulouse 2

Conditional directives are used by speakers to instruct hearers which actions are to be taken should certain events occur. The authors demonstrate that conditional directives are distinct from indicative conditionals in which speakers predict what is likely to be observed should certain events occur. The 1st set of experiments shows that goal structure determines what information speakers will select to test whether conditional directives have been followed but that these selections do not reflect their interpretations of the deontic necessity and sufficiency of the conditional relation. The 2nd set of experiments shows that formulations of conditional directives differ in how clearly speakers consider them to express their situation-specific intentions and that hearers accurately perceive what speakers intend them to do as a result of these formulations. The authors' findings illustrate a form of social rationality common in everyday interaction, which broadens normative conceptions of conditionals.

Keywords: reasoning, conditionals, pragmatics, conversation, rationality

As Austin (1962) observed, people use words to get things done—and “*if*” seems to be a very useful communicative device to use when making and communicating plans under conditions of uncertainty. The varied uses of “*if*” to coordinate work can be illustrated by a passage from Philip Roth’s (2000) novel, *The Human Stain*. It describes a scene in which Louie, the leader of a Vietnam veterans’ self-help group, is giving directives to Les, who has been suffering for many years from posttraumatic stress disorder from the war. Louie wants to help Les overcome his flashbacks by desensitizing him to Asian people. Using a mixture of instructions and promises, Louie tries to cajole Les into staying for as long as possible in a Chinese restaurant and eating a meal with other Vietnam veterans. Louie must carefully manage the conflicting goals of keeping Les long enough to at least partially desensitize him to the restaurant’s Asian staff, without provoking in Les the onset of a full-blown anxiety attack. To do this, Louie de-

scribes a kind of decision-tree with options for Les that maximize the chances of attaining this goal while minimizing the risks of collateral damage:

“Breathe,” Louie said. “That is it. Breathe, Les. If you cannot make it after the soup, we’ll go. But you have to make it through the first course. If you cannot make it through the double-sautéed pork, that is okay. But you have to make it through the soup. Let us make a code word if you have to get out. A code word that you can give me when there’s just no two ways about it. How about ‘tea leaf’ for the code word? That is all you have to say and we’re out of here. Tea leaf. If you need it, there it is. But *only* if you need it.” (pp. 219–220)

This passage illustrates some key points about the use of conditional directives to coordinate action. First, Louie’s uses of “*if*” in this passage are pragmatic in that they are being used to make plans to achieve a goal and to maximize benefits while minimizing (or avoiding) costs. Second, they are conversationally pragmatic in that achievement of these goals has to occur through social coordination in communication, in particular through successful recognition and execution of the speaker’s intentions by the hearer (Austin, 1962; Grice, 1975; Levinson, 1983; Sperber & Wilson, 1986; Wilson & Sperber, 2004). Third, the passage illustrates how nuances of the speaker’s intention are conveyed through use of different formulations of the conditional relation (*if*, *if not*, *only if*).

The pragmatic use of conditionals to make things happen rather than to describe states of the world poses a fundamental challenge for psychological theories of conditional use and inference. Although a number of theorists have acknowledged the importance of pragmatic influences on the interpretation of conditional statements (Evans et al., 2003; Johnson-Laird & Byrne, 2002; Oaksford, Chater, & Larkin, 2000), they are chiefly concerned with the extent to which reasoners in the role of *observers* will use world knowledge to infer how likely it is that Q will be true under the assumption that P is true. For example, Evans and Over (2004)

Denis J. Hilton and Jean-François Bonnefon, Laboratoire Dynamiques Sociocognitives et Vie Politique (D.S.V.P.), Université Toulouse 2, Toulouse, France; Markus Kimmelmeier, Interdisciplinary Ph.D. Program in Social Psychology, University of Nevada.

Portions of these data were reported in “Pragmatics at Work: Formulation and Interpretation of Conditional Instructions,” published in the *Proceedings of the XXIIIrd Annual Conference of the Cognitive Science Society* (pp. 417–422), Edinburgh, Scotland, August 2001. We thank Pierre Chandon, Fouad El-Ouardighi, and Benoît Monin-Morandat for their help in data collection and Heather Gillespie for help in the preparation of the manuscript. We thank Vittorio Girotto, David Mandel, David Over, Guy Politzer, Christine Sorsana, Steven Sloman, Valérie Tartas, and Gaëlle Villejoubert for helpful comments on drafts of this article.

Correspondence concerning this article should be addressed to Denis J. Hilton, Laboratoire D.S.V.P., Maison de la Recherche, Université Toulouse 2, 5, Allées Antonio Machado, 31058 Toulouse Cedex 9, France. E-mail: hilton@univ-tlse2.fr

argued that the conditional probability of Q given P plays a central role in the interpretation of a conditional “if P, then Q,” but their analysis rests on the assumption that people use language (and conditionals) mainly to describe the world to each other. Such analysis thus targets the constative use of language (Austin, 1962), in which the value of an utterance will indeed depend in large part on how *true* or *likely to be true* it is. For example, we may use “if” to suggest conclusions to be drawn from observations (e.g., “If it is a three-star restaurant, then the food there is certainly divine”). Such indicative conditionals are primarily statements about the world, which are true or false regardless of how a hearer responds to them.

However, we argue that this quasi-scientific *observer* perspective is irrelevant for *actors* such as Louie and Les who are communicating through conditional performatives (e.g., directives, promises, threats), which are not aimed at describing states of the world but at influencing the actions of other individuals. For performatives, the likelihood of Q given P is often beside the point. What matters for conditional performatives is whether the recipient of a directive responds to it in a way that is appropriate in the circumstances. For example, if Louie says to Les, “If you cannot make it after the soup (P), then we’ll go (Q),” both Louie and Les will know that the conditional probability of Q given P is being fixed for present purposes at 1. Indeed, given standard conversational assumptions (Grice, 1975; Hilton, 1995), Les would have no reason to doubt that P(Q|P) is 1 if Louie’s sincerity and ability to make good his promise are not in question.¹ What matters here is Louie’s expectation about the effect his conditional will have on Les’s behavior, an expectation that is actually made explicit in his next sentence: “But you have to make it through the soup.”

In this article, we outline a goal-based relevance model of use, interpretation, and testing of conditional directives, which we apply to conditional instructions. Conditional instructions presuppose the truth of the conditional relation (so Q must follow P), but establish that an action Q, to be taken by the hearer, is to be recognized by him or her as required by the speaker. Conditional instructions may be formulated in various ways: One can give others instructions to follow should a specific situation occur (e.g., “If a customer buys two make-up products, offer her a sample of this perfume”) and direct people on what *not* to do if a condition is *not* fulfilled (e.g., “If you do not get an invitation to the party, then do not go”). In addition, one can, if he or she wishes, be more specific, by telling people what conclusions to draw and actions to take *if and only if* or *only if* a condition is fulfilled.

The observation that there are variations in how conditional directives are formulated and used draws attention to certain lacunae in research on conditional reasoning. Much research has focused on the “if P, then Q” formulation of indicative conditionals (see Evans, Newstead, & Byrne, 1993, for a review). Further, numerous studies have examined how different formulations of conditional statements (e.g., “if and only if P, then Q”, “Q only if P,” and “if not P, then not Q”) may be paraphrased, reasoned about, and tested (Evans, 1972; Evans, Clibbens, & Rood, 1995; Evans & Newstead, 1977; Fillenbaum, 1978; Johnson-Laird & Byrne, 2002; McKenzie, Ferreira, Mikkelsen, McDermott, & Skrable, 2001; Ormerod, Manktelow, & Jones, 1993; Schaeken & Schroyens, 2000; Schroyens, Schaeken, & d’Ydewalle, 2001; Thompson & Mann, 1995; see Evans & Over, 2004, for a review). However, to our knowledge, research has not yet systematically

studied what might be the communicative functions of these different formulations of conditional relations. In like vein, research in cognitive ergonomics has examined the ease of processing various formulations of conditional instructions (e.g., Dixon, 1982; Wright & Hull, 1986; Wright & Wilcox, 1979) without developing a conversational pragmatic analysis of why speakers would prefer one formulation over another in the first place. Therefore, one innovation of our research is to systematically examine how goal structure influences the perceived appropriateness of different formulations of conditional directives.

Our central argument is that conditional directives aim to achieve effective social coordination in planning actions. We propose that a speaker’s formulation of a conditional rule (e.g., using *if* rather than *only if*) depends on the effect that he or she seeks to have on the hearer’s actions. Rational speakers will formulate a conditional directive in such a way that it results in effective social coordination through successful illocutionary uptake, that is, the correct perception of the speaker’s intended meaning by the hearer. Specifically, we hypothesize that the speaker’s goal structure (e.g., promotion vs. prevention focus; cf. Friedrich, 1993; Higgins, 2000; Lewicka, 1989) will determine how a conditional directive is to be formulated and how one should go about finding out whether a particular directive was followed.

Conditional Directives: Rationality Through Successful Social Coordination

Conditional directives seem to be pervasive in human life. For Luria (1959), they represent an important way in which language is used to control the actions of others. For instance, parents may use conditional directives to influence how their children respond to specific stimuli, and this device tends to be successful with toddlers as young as 3 years old. *If* invites the hearer to suppose the antecedent (Braine & O’Brien, 1991; Levinson, 2000; Strawson, 1952); therefore, parents can present their children with an imaginary situation and can give children instructions about what to do in such cases. Children tend to internalize such directives in inner speech and thus can be expected to obey their parents or caregivers’ directives and wishes even when they are absent. In adult life, conditional directives are widely used in social and organizational settings where an expert or a social superior will instruct neophytes or social inferiors on how to behave in identified situations (cf. Manktelow & Over, 1991). Whether with children or adults, conditional directives only work in the presence of a status difference, when the speaker occupies an authority role with respect to the hearer. For example, in our introductory example, Louie as group leader has the right to be the director who gives instructions to Les. Given their importance for parental and social control, it seems plausible that conditional directives are among the first kinds of

¹ We thank Gaëlle Villejoubert for pointing this out. Confirmation of the importance of the hearer’s perceptions of a speaker in interpreting his or her conditional utterances is provided by Evans and Twyman-Musgrove (1998), who found that hearers are more likely to assume Q given P if the speaker is likely to be able to keep his or her promise. See also Stevenson and Over’s (2001) work for related evidence.

conditional utterances to be understood by children and among the most widely used in adult life.²

In terms of the functions of language identified by Austin (1962), conditional directives are performatives that act felicitously or infelicitously on the social world, that is, they constitute a class of utterance whose success or failure is to be judged in terms of whether or not they bring about the intended behavior in others. Specifically, conditional directives seek to bring about states of the world that are *desired* by the speaker by telling the hearer what he or she *should* do in a hypothetical situation. Conditional directives thus do not make assertions about the current or future state of the world but rather aim to bring about future states of the world through directing the actions of others. In contrast, the indicative conditionals that are most frequently studied in the literature are constative in nature, that is, they represent descriptions of the world. Whether an indicative conditional is true or false depends on whether a particular conditional relationship *is* the case or not, that is, whether it does or does not exist in the world of discourse. For instance, the indicative conditional statement “If you heat water to 100° centigrade, then it will boil” is true if on observing the antecedent (the water is at 100 °C), the consequent indeed follows (the water boils) and is false if the consequent is not observed (the water does not boil).

Conditional directives achieve their objectives through the establishment of a social arrangement, which prescribes what should be done in a given case. Because of their inherently social nature, the value of conditional directives cannot be measured in terms of truth-value or probability but rather in terms of how well the hearer’s behavior satisfies the goals that the speaker had in mind when issuing the conditional directive (Kenny, 1966). Conditional directives can therefore be evaluated for their effectiveness in achieving a goal but not for their accuracy in describing reality. Imagine a parent whose goal is to prevent his or her child from being bullied at school. She instructs the child that “if another kid is mean to you, always tell the teacher.” The obedient child does exactly as he is told and ends up being beaten every day for being such a rat. The conditional is an exact description of the reality (every time the child is beaten, he tells the teacher), but this is of no relevance to its evaluation as a directive. The point that matters here is that it fails to achieve the parent’s goal and is thus a bad directive.

Directives have certain specific characteristics. They are like fiats (“Would that it rains!,” “If only he would behave himself!,” “Hopefully there’ll be someone at their post to fire a missile if we are attacked!”), in that they have to describe a desired state of affairs. However, unlike fiats, they are commonly stated in the second person and only in the future tense, and their contradictory is stated as “Do not do P” rather than “Do not P” (Kenny, 1966). Second,

a directive is a fiat for utterance to an agent: its point is to give the agent to understand that he is to realize the fiat . . . a directive cannot be satisfied unless the state of affairs is brought about by the recipient of the directive. Moreover, it must be brought about by him *because* of the utterance of the directive. (Kenny, 1966, pp. 68–69)

Thus, successful social coordination rather than truth-values determines the satisfactoriness of directives (Grice, 1975; Strawson, 1964). For this reason, to evaluate their effectiveness, we need to evaluate the extent to which they secure illocutionary uptake

from the hearer (Austin, 1962), that is, their success in getting the hearer to realize what the speaker wants him to do.

Directive and Indicative Attitudes to Deontic Rules

Deontic conditionals address the question of what should, must, may, or ought to be done in a given situation. An extensive body of research using the Wason selection task (Wason & Johnson-Laird, 1972) has shown that people are more likely to make “correct” (P and not-Q) selections with deontic conditionals than with indicative conditionals (Evans & Over, 2004). Our goal in this section is therefore to clarify the relationship between deontic conditionals and conditional directives: Specifically, we argue that whereas all conditional directives are deontic conditionals, not all deontic conditionals are conditional directives. We show that considering whether one takes an indicative or a directive attitude to a deontic conditional clarifies how deontic statements are asserted and used and how they should be evaluated.

Directives given to others, by their very nature, constitute deontic rules that are to be followed by that other. For example, a police chief in a county can dictate to a subordinate that “If workers are repairing a road, then assign a traffic policeman to direct the traffic.” Directives may therefore take the form of instructions that a speaker wants a hearer to follow in order to achieve the speaker’s goals. To be felicitous, the hearer needs to perceive that the speaker has the appropriate authority to issue the directive, such as being a parent, boss, expert, and so forth (Hilpinen, 2001). Assuming cooperative hearers, the effectiveness of a directive depends on whether its addressees perform appropriately or not according to the given rule.

However, it is also possible to speak of deontic rules in an indicative mood by focusing on them as descriptive statements of the world. This can happen in one of two ways. First, one might want to verify whether a deontic rule holds in a certain jurisdiction or not (indicatives about the existence of a rule). Second, if the deontic rule exists in that jurisdiction, one can verify whether it is in fact being followed or not (indicatives about the practice of a rule). In an illustration of the first case, imagine two national police superintendents are surveying what is going on in a local county. To them, the statement “If workers are repairing a road, then a traffic policeman must be directing the traffic” is the indicative description of an existing deontic rule prescribed by another authority, and they are not using their own authority to establish a new deontic rule to regulate police behavior in this county. It simply asserts a fact about what rules already hold in this county and can be true or false. Presumably, the existence of this rule can be verified by checking whether it is in the county police rule book or not. Having satisfied themselves that the rule indeed exists in this county, the two police superintendents can then pose

² Three- and four-year-old children understand indicative conditionals asserting the existence of deontic rules that apply to a third person (e.g. Cummins, 1996; Harris & Núñez, 1996). However, there appears to have been no research that systematically compares young children’s understanding deontics addressed to them rather than to others. Given that three-year-old children’s language indicates that they understand others’ desires (Harris, 1996), it seems possible that performatives (instructions, permissions, obligations, etc.) that constitute deontic rules are understood at least as early.

the second question of whether the rule is being followed or not. To do this, they can inspect traffic works through the county to see whether policemen on the spot are indeed guiding traffic or not.

In this regard, it is important to note that many of the deontic conditionals studied by psychologists are treated as indicative descriptions and not as directives. For example, much previous research (e.g., Gigerenzer & Hug, 1992; Politzer & Nguyen-Xuan, 1992) has often focused on whether a statement describing a deontic rule accurately describes practice or not (e.g., whether operatives such as postal workers are following instructions about how to stamp letters or not). A second point to note is that many deontic conditionals in previous research are expressed ambiguously in that it is not clear whether they are directive or indicative in mood because they are often stated in the third person, for example, "If workers are repairing a road, then a traffic policeman must be directing the traffic" (e.g., Cheng & Holyoak, 1985; Fiddick, 2000). To make the directive mood clear, one should express statements in the second person, for example, "If you know that workers are repairing the road, then assign a traffic policeman" (cf. Manktelow & Over, 1991). To avoid this ambiguity about the status and interpretation of deontic rules, in our own research we focus exclusively on the directive function by using elementary second-person conditional instructions that an organizational superior (the speaker) gives to an inferior (the hearer) to get the organization's work done properly.

In sum, indicatives answer the question "If P is the case, what will happen?" whereas directives answer the question "If P is the case, what should I do?" We therefore hold that probabilistic approaches of how people use and test conditionals are appropriate for indicatives but not for directives, whether these probabilistic approaches be disinterested (Evans, Handley, & Over, 2003; Evans & Over, 1996; McKenzie et al., 2001; and the model of the abstract selection task in Oaksford & Chater, 1994) or decision-theoretic (e.g., Klauer, 1999, or the model of the thematic selection task by Oaksford & Chater, 1994). This is because the point of the conditional directive for the speaker is to bring Q about through correctly instructing the hearer what to do when P is present, not to evaluate the probability of P happening or how likely Q is given P. For this reason, in directives it seems incongruous to qualify the probability of Q given P by saying "If workers are repairing the road, then probably assign a traffic policeman" but natural to do so for indicatives (e.g., "If workers are repairing the road, then there will probably be a traffic policeman").

That is not to say that conditional probabilities will not play any role in issuing or following conditional directives. Speakers are certainly concerned about the expected utility of issuing a directive, and hearers are certainly concerned about the expected cost of being caught violating the directive. Yet those two expected utilities do not feature the conditional probability of the consequent of the directive, given its antecedent. Instead, they feature the conditional probability of some good consequence, given that the directive has been issued or the conditional probability of some bad consequence, given that the directive has been violated. So although we believe that probabilistic analyses have yielded invaluable insights about the use and testing of indicative conditionals, we consider that their application to understanding how conditional directives achieve relevance is necessarily limited. What is needed instead is a goal-based calculus for evaluating the desir-

ability of courses of action, and for this we turn to an analogy with signal detection theory in decision-making.

Goal-Based Relevance in Conversational Planning: Signal Detection and Taking Action

From the speaker's point of view, uttering a conditional directive forms part of a plan to achieve a desired effect. Cooperative hearers will interpret conditional statements in light of what they know or what they infer the speaker's goals to be and will thus determine which actions are most relevant to fulfilling the speaker's conditional directive. Speakers will signal to hearers what they want them to do, and the choice of a particular form of conditional directive will be rational to the extent that it will predictably succeed in conveying what is intended to the hearer and in getting him or her to do what the speaker wants.

Work and organization are particularly suitable domains in which to study social coordination in reasoning. Typically, supervisors instruct employees on what to do on the basis of well-defined organizational goals. Perfect illocutionary uptake of the speaker's intentions by the hearer is achieved when the hearer (instructee) does exactly as the speaker (instructor) intended. The efficiency of the organization will indeed often depend on how well hearers pick up and execute speakers' expressed intentions. We assume that such intention recognition and action planning is achieved by the kinds of generalized, goal-based inference rules posited by models of language comprehension (e.g., Schank & Abelson, 1977; Wilensky, 1983). The speaker's goals will determine which kind of error is most to be avoided in a work situation, a situation that we consider below through an analogy between complying with a conditional instruction and being engaged in a signal detection task.

In a signal detection task, the observer has to judge whether a signal in fact indicates its putative referent or not. Does a certain kind of blip on a sonar screen indicate the presence of an enemy submarine (that has to be sunk) or not? The two types of error conceivable here are (a) an enemy submarine may really be out there, but its sonar "signature" has been distorted by underground wave patterns or rock formations in such a way that the operator fails to recognize it, leading to a *miss* (MS), and (b) other kinds of objects may have caused the blip, for example, whales or friendly submarines, causing an incorrect identification, thus leading to a *false alarm* (FA). Consequently, there are two kinds of mistake to be avoided in planning (and communicating about) actions to be taken: errors of omission (MSs, when one does not do something one should do) and errors of commission (FAs, when one does something one should not do). A real-life example of a miss was when the British Navy lost a warship, *H.M.S. Sheffield*, because of a failure to respond to a 19-min radar warning of an incoming Exocet missile during the Falklands conflict in 1982 (Ezard, 2000). An example of a costly FA was when the *U.S.S. Vincennes* mistook a civilian airliner for an attacking military aircraft and shot it down with a missile while patrolling the Persian Gulf in 1988.

Whether an MS or an FA has the highest expected cost may depend on the situation. For example, when under attack in war, an MS is likely to prove costly: If you do not sink the enemy submarine or aircraft first, then it will sink you. However, an FA may also prove costly: If you mistakenly sink a neutral country's

aircraft, submarine, or ship, you may provoke that country to declare war on you. This could prove especially costly if the neutral country happens to be, say, the United States. (Think of the *Lusitania*. . .) Now, imagine a warship commander wishing to give his weapons operators a conditional instruction linking the observing of an unusual “blip” (P) and the launching of depth charges (Q); for example, “If you see an unusual ‘blip’ on the screen (P), then launch depth charges (Q).” In Context A, the commander knows that enemy submarines’ unusual blips usually indicate enemy submarines, which are lethal if allowed within range and must be destroyed at first sighting. Clearly, what the commander should fear are MSs, that is, enemy submarines that are not attacked (cases of P and not Q). In Context B, the commander knows that enemy submarines are outside range at first sighting and that there is a considerable risk of destroying the commander’s own submarines or those of a neutral superpower that are also lurking in the area. What the commander should fear here is FAs, that is, nonenemy submarines that are attacked (cases of not P and Q).³

We present our experimental work in two parts. The first part shows that contextual goal structure affects the speaker’s information preferences for testing whether an *if-then* conditional rule has been followed without influencing the speaker’s intended logical interpretations of the conditional in terms of the perceived necessity or sufficiency of the antecedent for the outcome. In Experiment 1, we use a version of the Wason selection task (Wason & Johnson-Laird, 1972) to assess how context influences participants placed in the position of instructors who wish to verify the quality of the instructee’s work—that is, how the instructor checks that the instructee did indeed follow an *if-then* rule correctly in the context. Experiment 2 uses a truth-table task to show that the speaker’s information selection in the previous experiment cannot be explained by the logical interpretation he or she gives to the *if-then* rule, thus suggesting that rule-testing behavior in these examples is driven by perceived utility and not by logical interpretation.

The second part of our experimental work addresses how goal structure affects the communication of intentions through the formulation and interpretation of different forms of conditional instruction. Experiment 3 shows that speakers have very clear intuitions about which conditional formulations are appropriate for expressing their intentions when it is important to avoid an MS rather than an FA, and vice versa. Experiment 4 suggests that different formulations of the conditional do indeed achieve effective social coordination and successful illocutionary pick-up by the hearer, even though this does not seem to require correspondence between the speaker’s and the hearer’s logical interpretation of the conditional formulations in terms of perceived necessity and sufficiency.

Part 1: Checking for Mistakes in Following a Rule— Goal-Based Relevance or Logic?

We begin by testing our goal-based relevance approach on the Wason selection task.⁴ Various authors concur that participants’ information selections are determined by intuitions of relevance (Evans, 1996; Oaksford & Chater, 1995; Sperber, Cara, & Girotto, 1995), but thus far research has not yet systematically examined how goal structure can influence relevance (but see Perham & Oaksford, 2003, for an exception). We argue that participants in

versions of the Wason selection task in which they test indicative conditional statements about whether operatives are in fact following a rule can be thought of as testing whether that operative has correctly reacted (“performed”) in response to a signal. We propose that the relevant utility matrix determines the appropriate response, which defines the relative cost of MSs and FAs. This, in turn, decides the perceived relevance of information in the selection task: Specifically, we expect that the P and not-Q cards will be judged as more relevant in the avoid-MS condition and the not-P and Q cards as more relevant in the avoid-FA condition.

Explanation of Perspective Effects on Information Selection: Goal-Based Relevance Versus Switching Social Roles

Initial studies of the selection task sought to explain what factors prompt people to select the P and the not-Q cards, which would enable them to observe any P and not-Q pairings required to falsify the sufficiency of P for Q. To select P and not-Q is rational or “logical,” given an interpretation of the *if-then* conditional of P as sufficient but not necessary for Q. Following this traditional normative interpretation, theorists are interested in what facilitates

³ We wish to make clear that there is an important difference between our signal detection analogy and that of Kirby (1994). In our approach, the signal that is to be detected denotes the occurrence of P. As a consequence, a miss is defined as not doing Q when P is the case (a “P and not-Q situation”), and, accordingly, the cost of a miss is the cost of not doing Q when P is the case. On the contrary, in Kirby’s (1994) approach, the signal denotes the occurrence of a “P and not-Q” situation, rather than the occurrence of a P situation. Thus, a miss is defined as the *failure to detect a P and not-Q situation*, rather than the occurrence of the situation itself. Accordingly, the cost of a miss is the cost attached to the nondetection of a “P and not-Q situation,” rather than the cost of the “P and not-Q situation” itself. As an example of this difference, consider the situation in which a student has to take the following examination: A list of statements about experimental psychology is provided, and the student must underline (Q) every statement that is true (P). In the present article, a miss is defined as the student not underlining a true statement. In Kirby’s approach, a miss would be defined as failing to detect that the student has not underlined a true statement. Thus, the cost of a miss will be different in the two approaches. Within our framework, a miss is costly to the student, as the mistake will impinge on his or her mark. Within Kirby’s approach, a miss benefits the student, given that a miss would mean that the mistake of the student was not detected!

⁴ The selection task was created by Peter Wason as a tool for investigating hypothetical thinking with conditional statements and has achieved immense popularity among cognitive scientists. A typical experiment using the selection task will present a conditional statement and ask participants which information they will seek in order to decide whether the statement is true or false. The instructions will be along the following lines: “There are four cards lying on a table. Each is known to have a capital letter on one side and a single figure number on the other side. The following statement applies to these cards and may be true or false: “If there is an A on one side of the card, then there is a 3 on the other side of the card. The four cards respectively show on their exposed sides the symbols: A, D, 3, and 7. Your task is to choose those cards and only those cards that need to be turned over in order to decide whether the statement is true or false.” The reason why the selection task has been so intensely investigated is that usually less than 10% of reasoners can find the normatively correct answer that can falsify the rule (in this case, A and 7).

such rational selections (Almor & Sloman, 1996; Cheng & Holyoak, 1985; Kirby, 1994; Manktelow & Over, 1991; Oaksford & Chater, 1994, 1995; Sperber et al., 1995). However, other studies demonstrated consistent not-P and Q selections in the Wason selection task especially when “switched” social contracts are used (Cosmides, 1989; Gigerenzer & Hug, 1992; Manktelow & Over, 1991; Politzer & Nguyen-Xuan, 1992). For example, Politzer and Nguyen-Xuan (1992) used a department store scenario to examine how the unique perspectives of a social actor would influence choices in the Wason selection task. Focusing on the rule “If the purchase exceeds 10,000 francs (P), then the salesman must stick a voucher gift on the back of the receipt (Q),” they found that a customer will check whether the store is cheating on the rule by checking P and not-Q cards, thus avoiding an MS for the customer.⁵ By contrast, managers would give the instruction to check for customers cheating by examining more not-P and Q cards, thus avoiding an MS for the manager. Politzer and Nguyen-Xuan (1992) attributed this shift in selection to the fact that different perspectives give rise to different concerns. Likewise, Gigerenzer and Hug (1992) argued that these “switching social contract” effects support Cosmides’ claim that people have specialized “cheater detectors,” with each party to the contract checking that the other is not cheating on them.

A natural way for us to explain these well-established between-subjects perspective effects is to say that the utility distribution (costs and benefits of hits, MSs, FAs, and correct rejections) is different for the two parties. However, “anomalous” (not-P & Q) selections can also be obtained through emphasizing the probabilities or costs of an FA to an actor, without taking another party’s perspective or switching social contracts. For instance, Fairley, Manktelow, and Over (1999) showed that skepticism about the truth of an indicative conditional rule makes relevant counterexamples salient (see also Staller, Sloman, & Ben-Zeev, 2000). Fairley et al. (1999) asked participants to test the rule, “If you wear rubber gloves (P) when clearing up spilt blood, you are protected from infection (Q).” Participants tended to select P and not-Q cards if they were cued into the role of a hospital manager in charge who questions the sufficiency of rubber gloves to protect his staff against infection. However, when put in the role of a hospital finance manager who questions the necessity of rubber gloves because of the extra financial burden created by them, participants tended to select not-P and Q cards. Fairley et al. argued that they obtained this pattern of results because they successfully sensitized participants to the causal sufficiency or necessity of P for Q.

It is important, however, to note that Fairley et al.’s (1999) procedure simultaneously manipulates expectancy and utility because these authors varied both beliefs about the effectiveness of gloves *and* the perspective of the different managers. Without doubt, their paradigm is an appropriate manipulation of expected utility. However, from a goal-based relevance perspective, one cannot be sure whether its effects on information selection are due to expectancy, utility, or a combination of both. In addition, it remains unclear whether Fairley et al.’s findings were indeed caused by changing perceptions of the sufficiency and necessity of the rule in question because the authors did not directly assess participants’ logical interpretations of the conditional statements. Without an independent assessment of conditional interpretations, we cannot be sure that these effects come about through changing participants’ perceptions of the sufficiency and necessity of the

conditional relation. It is quite possible that varying utilities affect information selection on the Wason selection task independently of necessity and sufficiency. Consequently, Fairley et al.’s account of their own findings should be viewed as suggestive, at best.

Utility can also be manipulated independently of the perceived chances of success of an action, namely, by simply changing the perceived desirability of action’s consequences. Such a manipulation is likely to change the perceived *deontic* sufficiency or necessity of a condition for an action to be taken (cf. Over, Manktelow & Hadjichristidis, 2004). Rules where the case of P and Q is more desirable than P and not-Q (our avoid-miss case) can be called *deontically sufficient* (if the Condition P is present, a hearer *may* perform the Action Q). Conversely, an action in which the case of not-P and not-Q is more desirable than not-P and Q (our avoid-FA case) can be termed *deontically necessary* (if the Condition P is present, a hearer *must* perform the Action Q). This implies that a manipulation of utilities associated with different goal structures should not influence perceptions of the strength of the conditional relation (i.e., the likelihood that the action will have an intended effect; cf. Fairley et al., 1999) but rather its desirability. Because of their nonprobabilistic nature, conditional instructions provide a paradigmatic test case for assessing whether utility will affect information selection on the Wason task through deontic necessity and sufficiency or whether it will do so independently. This is important because Fairley et al. (1999) have argued that using deontic rules facilitates performance in the selection task because the context changes the perceived necessity and sufficiency of the causal relation. Consequently, a major innovation of our research will be to evaluate whether participants’ tests of whether a conditional instruction was followed in the Wason task can be predicted from their reasoning about the same rules on a deontic variant of the truth-table task.

We therefore assessed the effect of utilities by changing the contextually specified utility distribution that dictates which action errors (MSs or FAs) are most costly in a given case. Our paradigm thus allows us to manipulate the different perspectives within the same actor (cf. Light, Girotto & Legrenzi, 1990; Perham & Oaksford, 2003), whereas much previous research on perspective effects has manipulated utilities through taking the perspective of different actors (e.g., Fairley et al., 1999; Fiddick, Cosmides, & Tooby, 2000; Gigerenzer & Hug, 1992; Holyoak & Cheng, 1995; Manktelow & Over, 1991; Politzer & Nguyen-Xuan, 1992; see Girotto, 1995, for discussion). We constructed scenarios in which we varied goal structure to determine whether an error of omission (MS) or an error of commission (FA) would be the most costly error for a conditional instruction given by an organizational superior (the speaker) to a subordinate (the hearer). For example, we constructed a situation in which a warship commander wants to test how well his or her weapons operators follow the rule in the

⁵ In this example, we insert “(P)” and “(Q)” in the rule that was presented by Politzer and Nguyen-Xuan (1992) to their experimental participants in order to facilitate mapping from the rule onto the experimental results. We adopt the same expository procedure throughout this article, for example, when mapping the rules we presented to our participants to our predictions and results. However, like Politzer and Nguyen-Xuan, we did *not* present experimental participants with the bracketed “P”s and “Q”s in the rules used in the experimental procedure itself.

following form: "If you see an unusual 'blip,' then launch depth charges by testing their performance in a computer simulation under 'virtual reality' battle conditions." In the avoid-MS context, the commander knows that enemy submarines are lethal if allowed within range and must be destroyed at first sighting. Clearly, what the commander should fear are MSs (errors of omission, P and not-Q), that is, enemy submarines that are not attacked. This should lead the commander to select P and not-Q cards to ensure that sensitivity is maximized (i.e., all positives are detected and acted on). In the avoid-FA context, the commander knows that there is a considerable risk of misidentification and of destroying one's own submarines or those of allies. Here what the commander should fear are FAs (errors of commission, not-P and Q), that is, nonenemy submarines, which are attacked. This should then lead the commander to select not-P and Q cards to ensure that specificity is maximized (that only positives are identified and acted on).

Experiments 1A and 1B: Effect of Goal-Based Utilities on Perceived Relevance of Information for Rule-Checking

Four rules were created in avoid-MS and avoid-FA settings. The first was "If you see an unusual 'blip,' then launch the depth charges." Participants learned that enemy submarines are particularly likely to emit the unusual blips. The avoid-MS context explained that enemy submarines are lethal and have to be destroyed as soon as possible, whereas the avoid-FA context explained that the weapons operators may be worryingly "trigger happy" and liable to destroy an allied submarine. The second was an airport security rule, "If you detect a suspect object, then take the bag out," in which operators were instructed either to aim for accuracy of detecting suspect baggage (avoid MS) or to avoid unnecessary delays to passengers (avoid FA). The third was a border speeding rule, "If the speed of the vehicle exceeds 130 km per hour, then stop it," in which a French prefect seeks either to enforce rigor in the application of the rule (avoid MS) or to avoid causing unnecessary annoyance to motorists (avoid FA) in the context of cars coming from Germany (where there is no speed limit on motorways) into France (where there is a speed limit of 130 km per hour). The fourth and final rule was a variant of Gigerenzer and Hug's (1992) mail scenario, in which an office worker was told "If the letter is over 20 g in weight, put 2 marks in stamps on it," knowing that understamping (i.e., putting 1 mark stamps on letters over 20 g) would damage the firm's public image (avoid MS), whereas overstamp (i.e., putting 2 marks stamps on letters under 20 g) would be costly to the firm's finances (avoid FA). To illustrate the nature of our material, we provide the submarine story in the Appendix.

Method

Participants. A total of 100 psychology and humanities students at the University of Heidelberg, Germany, completed the questionnaires (roughly 50% women; age $M = 24.22$ years, $SD = 3.58$; range = 19–40 years). The questionnaire was administered in German.

Procedure. Participants read a first page of instructions and then, as a warm-up, worked on an abstract form of the Wason four-card selection task. On the next four pages, they read the four target scenarios: two on the avoid-MS version and two in the avoid-FA version. Two experimental groups were created with different pairings of scenario and version such

that each group did two avoid-MS and two avoid-FA tasks, with each task being presented in an avoid-MS form in one experimental group and avoid-FA form in the other. Specifically, roughly one half of the participants received the avoid-MS version of the submarine and the airport, as well as the avoid-FA version of the speeding and the postal scenario, whereas the versions were switched for the other half of the participants. Within each of these two groups, questionnaires were presented in two orders: one randomly generated and the other the reverse of that order. Participants were asked to imagine that they were wondering whether the operators were indeed following the rule in question. The 59 participants of Experiment 1A were asked to evaluate the relevance (on a 7-point scale anchored with *not at all relevant* and *extremely relevant*) of turning over each card (P, not-P, Q, not-Q) to aid them in their task. The remaining 41 participants of Experiment 1B were required to select those and only those items of information necessary to check whether the rule in question had been followed. In each of the scenarios, the "cards" (P, not-P, Q, not-Q) were presented in a fixed order (reading left to right) at the bottom of the page describing the task across experimental condition. The cards were presented in the order of P, not-P, not-Q, and Q in the submarine and border police tasks; P, not-P, Q, and not-Q in the surface-to-air missile task; and not-P, P, Q, and not-Q in the airport security task. The questions about the relevance of turning over each card were presented on a facing page in the same order as the cards.

Results and Discussion

Data are summarized in Table 1. Because for either response format there were no significant differences between the four scenarios, we collapsed across them.

First, it has to be noted that both response formats, that is, rated relevance versus information selection, resulted in remarkably similar patterns of results. This finding speaks to the robustness of the hypothesized effects of goal structure on reasoning about conditionals.

More important, though, results obtained in a series of 2 (goal focus) \times 2 (order of presentation) analyses of variance offer substantial confirmation for our predictions (see Table 1). Judged relevance of the P and not-Q cards was significantly higher in the avoid-MS condition, consistent with the assumption that participants here were most concerned with searching for an MS (i.e., a case of P and not-Q; e.g., not firing a depth charge when an enemy submarine was present). The converse pattern was observed for not-P and Q, whose relevance participants judged to be higher in the avoid-FA condition, even though across both response formats the difference only reached statistical significance for Q. This finding is consistent with the assumption that participants here were searching for an FA (i.e., a case of not-P and Q; e.g., firing a depth charge on a nonenemy submarine). The medium to large effects found across the four scenarios attest to the significance of our goal-based relevance analysis. For either response format, there were no effects of order of presentation, either as a main effect or interaction with goal focus (avoid MS vs. avoid FA), for the rating response format, all $F_s(1, 55) < 2.55$, $p > .12$, and the selection response format, all $F_s(1, 37) < 2.51$, $p > .12$.

It is noteworthy that, although the relevance of the P and not-Q cards is lower in the avoid-FA context than in the avoid-MS context, it remains high: In particular, the relevance of the P card is greater than the relevance of the not-P card even in the avoid-FA context (e.g., rated $M = 4.76$, $SD = 1.31$, vs. $M = 4.13$, $SD = 1.47$). It thus appears that the avoid-FA context cannot completely override a tendency to avoid MSs. Similarly, the relevance of the

Table 1
Experiment 1A: Effect of Utilities on Relevance of Information for Checking Rule, Averaged Across Four Scenarios

Goal structure	P		not-P		Q		not-Q	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Ratings of relevance (7-point scale; <i>N</i> = 59)								
Avoid MS	5.86	1.19	3.67	1.69	4.31	1.66	5.27	1.46
Avoid FA	4.76	1.31	4.13	1.47	5.95	1.19	3.69	1.65
<i>F</i> (1, 55)	10.32**		0.78		16.95***		13.84***	
Cohen's <i>d</i>	0.87		-0.28		-1.10		1.01	
Proportion of selection of relevant information (<i>N</i> = 41)								
Avoid MS	.79	.25	.30	.35	.44	.31	.65	.31
Avoid FA	.58	.35	.36	.34	.65	.29	.40	.38
<i>F</i> (1, 37)	5.04*		0.40		6.41*		5.65*	
Cohen's <i>d</i>	0.67		-0.17		-0.69		0.71	

Note. *F* tests pertain to the difference between the avoid-MS and avoid-FA conditions with their respective panels. MS = miss; FA = false alarm.

* $p < .05$. ** $p < .01$. *** $p < .001$.

not-P card was quite high even among participants in an avoid-MS context: In the subsample that rated the perceived relevance, the mean was slightly below the midpoint of the scale ($M = 3.67$, $SD = 1.69$, on a 7-point scale), and in the subsample that selected relevant pieces of information, almost a third of all participants chose the not-P card. Although these results seem to contrast with the spectacular reversals obtained in perspective shift studies, it should be noted that we had no reason to expect such a spectacular reversal to take place here. Indeed, in our scenarios, the goals of both avoiding MSs and avoiding FAs are bound to be in play at least to some degree, and the experimental manipulation primarily serves to alter the emphasis between the two goals. However, the high ratings or high proportions of selections of P and not-Q in the avoid-FA context and of not-P and Q in the avoid-MS context could of course be justified if participants are interpreting the conditional relation as a “deontic biconditional” (where Q must be performed when P is the case and Q must *not* be performed when P is not the case), an issue we examine in our next experiment.

The results of Experiment 1 provide corroborating evidence that goal-driven intuitions of relevance drive information selection on the Wason task. They are consistent with earlier results that show that explicit mention of a goal facilitates P and not-Q selections (Fiddick et al., 2000, Experiments 1 and 2; George, 1991), while, however, generalizing this result to “anomalous” not-P and Q selections. As such, our results support our goal-based relevance model and complement informativeness-based relevance approaches (e.g., Evans, 1996; Klauer, 1999; Oaksford & Chater, 1995; Sperber et al., 1995). However, as we show below, our results are difficult to reconcile with some other theoretical perspectives.

Comparison to Other Work on the Wason Selection Task

As we noted above, the rules used in our experiment are not social contracts involving exchanges in the way Cosmides (1989) has defined it: A conditional instruction does not relate perceived

benefits to perceived costs, and it does not express a social exchange in which an individual is required to pay a cost (or meet a requirement) to another individual in order to be eligible to receive a benefit from that individual. Because conditional instructions do not involve exchange, they do not leave room for cheating, that is, obtaining the benefit without paying the cost. Our pragmatic analysis does suggest that not-P and Q selections may be obtained by changing the perspective of a single party, by using instructions that contain (at most) only unilateral obligations in which a superior (the speaker) requires a subordinate (the hearer) to respond appropriately to a condition (signal). Specifically, if the context suggests an avoid-FA situation, then we predict that not-P and Q selections should be favored. This is because we argue that changes in the perspective of the instructor will be sufficient in themselves to change card selections through changing the perceived payoffs from the point of view of one party alone (the instructor). In this view, it is not necessary to obtain perspective shifts from one party to a social contract to another to obtain not-P and Q selections. Hence, theories that seek to explain not-P and Q selections through changes of perspective from one party to another, such as bilateral social contracts (Gigerenzer & Hug, 1992) or mutual obligations (Holyoak & Cheng, 1995; Politzer & Nguyen-Xuan, 1992), cannot account for our data.

Likewise, our analysis seems to preclude the need to posit a new category of “precaution schemas” that have been hypothesized to facilitate not-P and Q selections (Fiddick et al., 2000, Experiment 3). Fiddick et al. (2000) used a conditional rule-testing task that, like ours, varies perspective within the same actor. Thus, a Western visitor to a remote tribe who brings a limited supply of protective rubber gloves says to a tribe member, “If you make poison darts, then you may wear rubber gloves.” They put experimental participants in the role of the visitor who either tests in a “precaution” orientation whether tribespeople were risking their lives or who, in a “social contract” orientation, tests whether tribespeople were abusing their privileges. Fiddick et al. observed

a strong tendency to observe P and not-Q selections in the first case, and not-P and Q selections in the second case. In our scheme, we could of course attribute this first result (P and not-Q selections) to an avoid-MS perspective (not taking action against a potential threat) and the second result (not-P and Q selections) to an avoid-FA perspective (taking unnecessary and wasteful action), without there being a need to attribute this pattern to different “ecologically rational, content-dependent notions of violation” (Fiddick et al., 2000, pp. 59–60). Rather, our data suggest that the same underlying content-independent logic of utility can explain both P and not-Q selections, in which precautions are being taken to avoid unintended and undesirable side effects to plans, and also not-P and Q selections, in which attainment of goals is being promoted (through social contracts or other means). Our manipulation of the contextual goal structure is consistent with the utility theory of conditional reasoning (Manktelow & Over, 1991), though we have extended the utility analysis through the promotion versus prevention focus implicit in our use of the signal detection analogy. Thus, as well as focusing on “hits” (P and Q) and MSs (P and not-Q), like Kirby (1994) we use a signal detection analogy to consider the cases of “correct rejection” (in our model, not-P and not-Q) and FAs (not-P and Q). However, we go further than Kirby in showing that utilities should and do affect selections of not-P and Q cases. Our goal structure approach thus allows a more comprehensive analysis of the way that utilities will influence information selection when testing a conditional rule and shows how both “correct” (P and not-Q) and “anomalous” (not-P and Q) selections can be rational for the same individual.

Finally, in deriving our theoretical predictions from the structure of goal-driven action, our approach has some fundamental orientations in common with the pragmatic schemas approach of Cheng and Holyoak (1985), especially when Holyoak and Cheng (1995) wrote that “such schemas represent knowledge about the types of regulations typically constructed by those in authority to constrain human voluntary action in a manner consistent with their goals” (p. 291). Nevertheless, our approach differs from theirs in important ways. First, Cheng and Holyoak did not distinguish directive and indicative attitudes to deontic conditionals as we do. Second, it seems to us that pragmatic schemas theory cannot explain the patterns of use of conditionals that we observe. In particular, by changing goal context, we obtained different patterns of information selection to test the same rule content, a finding that is impossible to explain by cuing of a pragmatic permission or obligation schema based on the content of the rule. One way of revising pragmatic schemas theory to account for our results might be to argue that the avoid-MS context activates an obligation schema (e.g., “If you see an enemy blip, then you must launch a missile”), whereas the avoid-FA context activates a permission schema (e.g., “If you see an enemy blip, then you may launch a missile”). However, even here, one problem for the pragmatic schemas approach becomes immediately clear: Given the life-or-death nature of this situation, in the avoid-FA case, it would seem rather casual of a commander to simply give permission to a weapons operator to fire, implying that the subordinate has the right to decide not to fire even if there is indeed an incoming missile.

A final problem for pragmatic schemas theory is that permission and obligation schemas imply different patterns of conditional reasoning; therefore, if our context manipulations have differen-

tially activated these schemas, then we should observe the corresponding patterns of conditional reasoning on a truth-table task—or, more precisely, on a deontic version of the truth-table task. We turn to this question below.

Experiment 2: Does Goal Structure Affect the Logical Interpretation of the Rule?

The preceding experiments have shown that goal structure has a strong and robust effect on information selection in the Wason selection task. Using the same research materials as in Experiment 1, we now ask whether this result is obtained because goal structure directs the hearer’s attention to what is important in the situation or whether it is because goal structure affects the perceived deontic necessity and deontic sufficiency of the conditional relation. Is it the case, for example, that the avoid-FA context prompts a “deontic biconditional” interpretation, in which P is both deontically necessary and sufficient for Q, making not-P and Q selections relevant in addition to P and not-Q selections, whereas the avoid-MS context prompts an interpretation in which P is deontically sufficient but not deontically necessary to Q, making only P and not-Q selections relevant? So, if context influences the logical interpretation of the rule, we expect different patterns of response in the avoid-MS and avoid-FA conditions. Specifically, if participants use a deontic sufficiency interpretation in the avoid-MS condition, we predict that they will judge that the hearer will have followed the rule when P and Q is observed but not when P and not-Q is observed, with the cases of not-P and Q and not-P and not-Q perhaps being judged as irrelevant. However, if they shift to a deontic biconditional (sufficiency *and* necessity) interpretation in the avoid-FA condition, we expect them to judge cases of not-P and not-Q as cases in which the hearer has followed the rule and cases of not-P and Q as cases in which they have not followed the rule.

Because we are interested in the *speaker’s* judgment of the intended meaning of his utterance, participants were placed in the role of the speaker who wishes to test whether the hearer has correctly interpreted his statement.

Method

Participants. Participants were 46 University of Heidelberg students (17 men, 27 women; 2 missing data) from the same population as the participants of Experiment 1. The mean age was 23.57 years ($SD = 2.40$; range = 20–29 years). Again, the experiment was conducted in German.

Procedure. We used the same four scenarios used in Experiment 1. Again, participants were presented with all four scenarios composed of two avoid-MS and two avoid-FA contexts in a yoked design. There were two different pairings of scenarios, which were presented in two different orders. Participants were placed in the position of a superior (e.g., warship commander) who wished to check that his or her subordinate had followed the directive (e.g., “If you see an unusual blip, then launch depth charges”). They were presented with a deontic variant of a truth-table task (e.g., Staudenmayer, 1975), in which the presence and absence of the signal (e.g., the presence of an “enemy blip” or “allied blip” on the screen) was paired with the presence or absence of the action to be taken (e.g., launch the depth charges). Participants responded to four different combinations of events: P and Q, P and not-Q, not-P and Q, and not-P and not-Q. Participants indicated whether in each combination the subordinate had “followed” or “not followed” the rule or whether the pair of events was

Table 2
Experiment 2: Truth-Table Task ($N = 46$)

Outcome	Response options			Outcome	Response options		
	“Followed”	“Didn’t follow”	“Irrelevant”		“Followed”	“Didn’t follow”	“Irrelevant”
P and Q				Not-P and Q			
Abstract				Abstract			
Avoid MS	87	10	4	Avoid MS	5	50	45
Avoid FA	95	0	5	Avoid FA	18	41	41
Submarine				Submarine			
Avoid MS	75	4	21	Avoid MS	21	54	25
Avoid FA	62	0	38	Avoid FA	9	36	55
Airport security				Airport security			
Avoid MS	88	4	8	Avoid MS	17	42	42
Avoid FA	100	0	0	Avoid FA	5	50	46
Border speeding				Border speeding			
Avoid MS	92	4	4	Avoid MS	4	58	38
Avoid FA	95	0	5	Avoid FA	5	50	45
Mail				Mail			
Avoid MS	92	4	4	Avoid MS	25	25	50
Avoid FA	100	0	0	Avoid FA	9	46	45
P and not-Q				Not-P and not-Q			
Abstract				Abstract			
Avoid MS	4	92	4	Avoid MS	39	4	56
Avoid FA	0	100	0	Avoid FA	48	5	48
Submarine				Submarine			
Avoid MS	8	63	29	Avoid MS	63	8	29
Avoid FA	5	57	38	Avoid FA	33	5	62
Airport security				Airport security			
Avoid MS	4	88	8	Avoid MS	58	8	33
Avoid FA	0	95	5	Avoid FA	52	5	43
Border speeding				Border speeding			
Avoid MS	12	83	4	Avoid MS	54	4	42
Avoid FA	5	86	10	Avoid FA	57	0	43
Mail				Mail			
Avoid MS	8	83	8	Avoid MS	46	4	50
Avoid FA	0	91	9	Avoid FA	45	0	55

Note. Numbers reflect proportions of categorical judgments for each combination of scenario and goal condition. Within each row, percentages add up to 100% (give or take rounding error). $n = 24$ for avoid-MS context, and $n = 22$ for avoid-FA context. MS = miss; FA = false alarm.

irrelevant for an assessment of whether the subordinate had followed the directive.

Results and Discussion

A summary of the results is presented in Table 2. Various statistical analyses did not yield any evidence of changing conditional interpretations as a function of context. Specifically, when we collapsed across all scenarios, an analysis of variance on the frequency with which participants had responded to each of the four combinations of events showed that “followed,” “did not follow,” and “irrelevant” responses did not vary as a function of goal focus, all $F_s(1, 42) < 2.38$, $p > .13$. Similarly, order of presentation had no effect for all main and interaction effects, $F_s(1, 42) < 2.04$, $p > .16$. Instead, for all four scenarios, in both avoid-MS and avoid-FA contexts, participants showed a strong tendency to interpret the presence of the antecedent in the *if-then* directive as sufficient but not necessary for the action to be taken. Consequently, the effects of context on information selection we observed in Experiment 1 cannot be attributed to changes in the

perceived deontic necessity or sufficiency of P in regard to Q; they can only be attributed to differences in the perceived relevance of the information because of changes in the perceived desirability of the outcome.⁶

Our findings that goal structure influences how participants check whether the rule has been followed on the Wason selection task (see Experiment 1), but not how the rule has been logically interpreted in the indicative form of the truth-table task concerning whether the rule has been followed or not, show that utility can affect information selection on the Wason selection task without changing the interpretation of the conditional. This contradicts the assumptions of Fairley et al. (1999) but is consistent with Thomp-

⁶ We also ran two informal studies in French using the same options of “followed” and “not followed” as in the German study, as well as the options of “true” and “false” that have been more traditionally used in previous research using truth tables, and obtained essentially the same pattern of results—a strong preference for conditional interpretations of P as sufficient but not necessary for Q regardless of context.

son's (1994, 2000) findings that whereas conditional necessity and sufficiency relations predict performance on the truth-table and conditional argument (inference) tasks, different considerations (such as content) may determine information selection on the Wason task. For example, Thompson (1994) found that variations in necessity and sufficiency of conditional relations had far more effect on performance on a truth-table task (as on a conditional arguments task) than did type of content (permission schema, obligation schema, causal relation, definition).

Experiment 2 thus offers further support for the position that performance using conditional instructions on the Wason selection task is determined by goal-based intuitions of relevance rather than by logical reasoning (cf. Evans, 1996; Girotto, Kimmelmeyer, Sperber, & van der Henst, 2001; Sperber et al., 1995). In addition, it seems implausible that the intuitions of relevance on this particular task can be explained through optimal data selection based on disinterested probabilistic principles. For example, even if one argued that low-desirability events are also likely to have low probabilities of occurrence (cf. Weber & Hilton, 1990), following Oaksford et al. (2000) we should expect our context manipulation to induce differences in propensity to deny the antecedent (i.e., given not-P, infer not-Q) on the indicative version of the task concerning whether the rule has been followed or not. This was not the case. Hence, these results support our claim that unlike probabilities, utility determines perceptions of relevance of cards on the Wason task without affecting the perceived deontic necessity and sufficiency of P in regard to Q.

Our theoretical approach is consistent with the relevance theory approach of Sperber et al. (1995) to conditionals, in that in reasoning about the potential consequences of (not) performing an action, people do indeed seem to be following Sperber et al.'s effect principle of relevance. However, our approach allows us to go further in certain respects than Sperber et al. did. Specifically, our goal-based relevance model allows principled predictions to be made regarding a single individual's intuitions of relevance about conditional instructions that will motivate both "correct" (P and not-Q) and "anomalous" (not-P and Q) selections in a Wason task. Although Sperber et al. and others (e.g., Fiddick et al., 2000; Girotto et al., 2001; Liberman & Klar, 1996) have also demonstrated similar contextual effects on perceived relevance in the Wason task, such effects have not been motivated by an analysis in terms of a single individual's utility distributions. Even when Oaksford, Chater, and Grainger (1999) suggested that Sperber et al.'s manipulations of effect could be regarded as manipulations of subjective utility, in line with Kirby's (1994) signal detection analogy, they saw the utility being manipulated as the cost of *failing to detect a violation*, rather than the cost of a violation itself (see Footnote 3).

A further motivation for making the distinction between kinds of errors of omission (MS) and commission (FA) in the way that we do would be that it enables discovery and explanation of new phenomena. We argue that the approach explains facts about how language gets used: in particular, that goal structure specifies felicity conditions for different formulations of conditional relations that may be either successful or unsuccessful in transmitting the speaker's intentions (avoid MS vs. avoid FA) to the hearer. It thus promises a systematic explanation of use, interpretation, and testing of conditionals that is a priori consistent with the kind of

goal-based reasoning that underlies a wide variety of pragmatic phenomena (Levinson, 2000).

Part 2: Conditional Directives as Social Coordination Devices—Goals and Illocutionary Uptake

Why is it more natural to phrase a conditional directive one way rather than another? For example, in the avoid-FA situations we studied above, did our participants find it odd that we stated the rule as "If P, then do Q"? Anecdotal evidence suggests that this was indeed the case. Several weeks after participating in a replication of Experiment 1, a participant accosted Denis J. Hilton to let him know that the avoid-FA condition "had given him a headache," as his commonsense intuitions about what information he should select directly contradicted what he had been trained to think of as the logically correct interpretation.

We, of course, chose the *if-then* formulation of the conditional rule to maintain continuity with other work on the Wason task. However, in doing so, it is very possible that we put our participants in the avoid-FA condition into a kind of double bind—because according to the canons of logical inference, selections of not-P and Q are manifestly irrational even if it is dictated by the instinct for self-preservation. Here, what is normative reasoning for a logician diverges from what is normative utility analysis for an evolutionary psychologist or economist. In our view, the resolution of this paradox is to be found in explicating the role of goal structure in determining the felicity conditions for uttering different formulations of conditional instructions. We argue that if speakers follow their intuitions about which formulation best expresses their intention, then hearers will be spared the fate that befell the students in our first experiment, because formulations that are inappropriate for the context will be screened out. Further, when the context is either neutral or opaque, these formulations can substitute for explicit context through implicitly conveying what the speaker's chief concerns are.

As argued earlier, in conversational action planning, the real utility of the speaker using a conditional directive is to be evaluated in terms of its effect, namely, its success in eliciting the intended response from the hearer. As such, rationality is evaluated in terms of the successful coordination of the speaker and hearer around the goals of the task in hand, implicit in Grice's (1975) cooperative principle and, in particular, in his maxim of relation (Levinson, 2000). What is important is that the superior successfully communicates his or her hopes (goals to be attained) and fears (pitfalls to be avoided), which leads to successful illocutionary uptake by the hearer. If both speaker and hearer succeed in coordinating their actions successfully around promoting intended goals while avoiding unwanted side effects, then both will stand to gain (e.g., the supervisor will have an efficient organization and the subordinate will have the feeling of a job well done). Likewise, if the communication misfires for some reason (e.g., if the superior gives the wrong instruction or is misunderstood), then both parties stand to lose through feeling frustrated, apart from the losses to the organization.

Felicity Conditions for Formulations of Indicative Conditionals

What little work has been done on the felicity conditions for making conditional assertions has focused on indicative condition-

als that make propositions that can be either true or false. For example, Stevenson and Over (1995) cited influential work in philosophy (e.g., collected in Jackson, 1991) to support the claim that “if P, then Q” statements become “pragmatically assertable” (Stevenson & Over, 1995, p. 617) when people have grounds to believe that Q is highly probable given P. Consistent with this reasoning, research has shown that modus ponens inferences from indicative conditionals (i.e., given that “if P, then Q” holds, and P exists, then inferring that Q also exists) are suppressed when circumstances make it unlikely that Q will follow P, either because a relevant contextual precondition fails to hold (Bonnefon & Hilton, 2002) or because the action in question will lead to an undesirable consequence (Bonnefon & Hilton, 2004).

In some relevant experimental work, McKenzie et al. (2001) showed that people tend to encode the rare event when phrasing a conditional hypothesis to be tested (cf. the rarity principle of Oaksford & Chater, 1994). Thus, if given data suggesting that successful entrants into a school are rare, participants are more likely to phrase a hypothesis as follows: “If applicants have high SATs, then they will be accepted” (“If P, then Q”). However, about a nonselective school, where few are rejected, people will formulate a hypothesis to be tested as follows: “If applicants have low SATs, then they will be rejected” (akin to “if not P, then not Q”).

Sperber et al. (1995, p. 73) have argued that the naturalness of conditional expressions is conditioned by what they implicitly deny in a given context. Thus, utterances of the form “if P, then Q” (e.g., “If workers are repairing a road, then a traffic policeman must be directing the traffic”) characteristically achieve relevance when P and not-Q has been asserted (e.g., someone said they saw men working in the street without there being traffic policemen). The conditional statement would contradict this observation through casting doubt on its accuracy, in much the same way a counterfactual conditional achieves relevance by speculating on what would have happened if a fact were undone (cf. Fillenbaum, 1978). This kind of “implicit denial” approach (cf. Wason, 1965) can be extended to other conditional formulations. For example, what if someone asserts that she saw traffic policemen (Q) directing traffic in this county even though there were no men working in the road (not-P)? Here, we suggest that it might be felicitous to contradict this observation by saying “Q only if P,” that is, “Well, that cannot be true because traffic policemen must be directing the traffic only if workers are repairing a road.”

The implicit denial approach can be adapted to conditional directives through explicating of the role of goal-based reasoning in recognition of their relevance, even if the appropriate counterexample has not been explicitly mentioned in a preceding turn in the conversation. For example, a superior may establish a rule of the form “Q only if P” to be followed by a subordinate (e.g., “Assign a traffic policeman to direct traffic only if there are men working in the road”). The form “Q only if P” can achieve relevance for the hearer if he or she infers that there is a *risk* in the superior’s mind of not-P and Q (the hazard being absent but a costly precaution being taken anyway) that his or her boss wishes to avoid. For example, the hearer may infer that his or her superior may chiefly be worried that valuable police resources could be wasted by assigning traffic policemen when they were not needed (e.g., at school crossings where there were already independent supervisors) even if this state of affairs has in fact not been

previously mentioned or observed. For relevance to be achieved, the hearer has therefore to correctly perceive and execute the beliefs and intentions that the speaker conveys by a particular formulation of the conditional relation.

Expressing Intentions Clearly: The Effects of Contextual Goal Structure on Formulations of Conditional Instructions

We suggest that felicity conditions for formulations of conditional instructions are based on goal-based reasoning, which drives intuitions of relevance about what is most of concern in a given situation. Over et al. (2004) have shown that deontic conditionals expressing obligations (e.g., “If you wash the dishes, then you must wear rubber gloves”) and permissions (e.g., “If you work over the weekend, you may take a weekday off”) are most acceptable when the P and Q state of affairs is preferred to the P and not-Q state of affairs. For our conditional instructions, the *if-then* formulation should therefore be most appropriate in avoid-MS situations, where the speaker wishes to emphasize the need to attain the goal in question (P and Q) due to the costs of missing out (P and not Q). Conversely, we expect that the conditional statement of “Q only if P” is typically used when the speaker is most worried about cases of not-P and Q and is most likely to be appropriate to avoid-FA situations, in which the speaker wishes to emphasize the importance of taking this action only in the signaled case.

The explicit biconditional “if and only if P, then Q” will convey that the speaker is worried about both not-P and Q and P and not-Q combinations and should be most appropriate for cases when *both* MSs and FAs should be avoided. Finally, our goal-based adaptation of Sperber et al.’s (1995) implicit denial approach is that “if not-P, then not-Q” suggests that the speaker will be most concerned by not-P and Q occurrences. We predict that it will be perceived as more appropriate for the avoid-FA case, in which the speaker wishes to emphasize the importance of not taking this action except in the signaled case.

In the following experiment, we investigate how context influences the perceived felicity of conditional formulations. What would be the best way for the commander to frame his conditional instruction in Context A (avoid MS)?

If you see an enemy blip, then launch the depth charges? If and only if you see an enemy blip, then launch the depth charges? If you do not see an enemy blip, then do not launch the depth charges? Launch the depth charges only if you see an enemy blip?

In addition, what would be the best choice in Context B (avoid FA)? Our overall prediction is that people can and do perceive differences in the naturalness of these formulations as a function of what they perceive to be the goal structure of the context, that is, an avoid-MS goal versus an avoid-FA goal. In the avoid-MS context, one would like to stress the importance given P (observing an enemy blip) of doing Q (launching the depth charges), whereas in the avoid-FA context, stress should be on the importance, given not P, of not doing Q. Hence, we expect the ideal formulation in the avoid-MS context to direct attention to deontic sufficiency would be “if P, then Q,” whereas the other three formulations would be more appropriate in the avoid-FA context. In this context, a natural formulation will be one that directs the attention of

the hearer to deontic necessity: Is it important, given not-P, not to do Q? The experiment presented below explores these questions.

Experiment 3: The Effect of Contextual Goal Structure on Conditional Formulations

Method

Participants. Forty-seven students at the University of Heidelberg took part in this study. Students majored in a variety of different academic disciplines, with the largest group majoring in law (26%). The average age was 22.84 years ($SD = 2.55$, range = 19–40 years). All experimental materials were presented in German. One participant was excluded because of excessive missing responses.

Materials. We used the same four scenarios and experimental design as had been used in Experiments 1 and 2, and each scenario came either in an avoid-FA context or in an avoid-MS context. Participants were cued into the role of a superior who wanted to give a directive to his or her underling that assured that a specific rule was followed. Participants were asked to evaluate the clarity of different formulations for communicating the rule to their subordinates. Following each scenario, four conditional instructions were introduced (“if P, then Q,” “if and only if P, then Q,” “Q only if P,” “if not-P, then not-Q”), with participants having to rate on a 7-point scale “how well each formulation expresses the speaker’s intentions in these circumstances.” Because participants responded to all four conditional instructions, formulation was thus treated as a repeated-measures factor.

Each questionnaire featured the four scenarios, all of them in their avoid-MS or avoid-FA version (context was thus a two-level, between-groups factor), rotated over two experimental blocks, with two orders within each experimental block. That is, as in our previous experiments, each participant saw two avoid-MS and two avoid-FA contexts paired with different content scenarios.

Results and Discussion

Data were analyzed using a 2 (goal focus) \times 2 (order of presentation) factorial design. Because the observed pattern of results was remarkably stable across scenarios, we collapsed across the four scenarios and display the mean clarity of expression ratings in Table 3. With the exception of “if and only if P, then Q,” goal focus had a significant effect on clarity ratings. First of all, our primary predictions were confirmed. “If P, then Q” was indeed perceived as more natural in the avoid-MS context, whereas “Q only if P” and “if not P, then not Q” were perceived as more natural in the avoid-FA context. Second, “if and only if P, then Q” was perceived as fairly natural in both avoid-MS and avoid-FA contexts but less so than “if P, then Q” in the avoid-FA context. This may reflect the conversational “overkill” achieved by the “if and only if” formulation in contexts in which there is either an avoid-MS or avoid-FA concern but not both.

Finally, a closer examination of our results suggests that there may be some fine-grained differences of meaning in conditional formulations that our analysis did not predict. In particular, it seems that “if not-P, then not-Q” was judged as even less appropriate than “Q only if P” in the avoid-MS context. In addition, our results suggest that “if not P, then not Q” is not simply the mirror image of “if P, then Q,” given that it is overall less natural than the latter, even though it is judged the most natural expression of all in the avoid-FA context. These results are consistent with those of Fillenbaum (1978), who found that rephrasings of negative conditionals showed systematic differences to those of positive conditionals.

Overall, the results clearly demonstrate that goal structure is a strong determinant of preferences for formulating conditional directives. Where the cost of not taking the action in response to the signal is salient, as in the avoid-MS context, the “if P, then Q” formulation is most favored. Conversely, where the cost of taking the action in the absence of the signal is salient, as in the avoid-FA context, the other formulations (“Q only if P,” “if not P, then not Q”) are favored, with use of “if and only if P, then Q” being intermediately favored in both avoid-MS and avoid-FA contexts. Consequently, these results show, among other things, that the temporal antecedence of P (Evans & Newstead, 1977) is not a necessary condition for finding “Q only if P” more natural than “if P, then Q,” just as rarity is not a necessary condition for preferring the “if P, then Q” formulation to the “if not P, then not Q” formulation (McKenzie et al., 2001). These results suggest that conditional directives do indeed have different felicity conditions to those of indicative conditionals.

The choice of these formulations will only be rational for the speaker if they indeed induce the hearer to do what the speaker wants. In the next experiment, we evaluate whether the different formulations of the conditional do indeed achieve successful illocutionary uptake.

Experiment 4: Illocutionary Uptake of Conditional Directives—Do Hearers Understand What They Are Meant to Do?

The experiments we have reported have all taken the speaker’s perspective, and Experiment 3 has shown that contextual goal structure has a considerable effect on what the speaker considers to be a clear formulation of his or her intention. However, in many cases, a hearer will not be aware of the context that motivated a conditional directive. A superior may not let an inferior know the goals that led him or her to formulate an instruction in a certain way. In some cases, he or she may not wish to make such information public; in others, he or she might not have the opportunity to do so. For example, a ship commander may not wish to give explicit information about the policy and intelligence considerations that lead him or her to consider some submarine threats to be more dangerous than others or may not have time to do so in the heat of the moment. Nevertheless, the commander will normally wish for subordinates to follow the commander’s instructions *as if* they knew explicitly his or her goals. Consequently, the success of the commander’s policy will depend on whether his or her conditional instructions are successful performatives: Do they achieve successful illocutionary uptake and execution because subordinates successfully detect what is expected of them and execute their commander’s policy as intended? In other words, does each of our four conditional formulations (“if P, then Q,” “if and only if P, then Q,” “Q only if P,” and “if not-P, then not-Q”) succeed in inducing subordinates to avoid MSs and FAs appropriately? Specifically, do “if P, then Q” formulations succeed in orienting hearers to the importance of avoiding MSs (avoid-MS condition: i.e., “if P is present, then it is important to do Q”), whereas do the other formulations orient hearers to the importance of avoiding FAs (avoid-FA condition: “if P is absent, then it is important not to do Q”).

Table 3
 Experiment 3: Perceived Clarity of Expression of Intention by Conditional Formulations as a Function of Context (N = 46)

Formulation	Avoid MS (n = 22)		Avoid FA (n = 24)		F(1, 42)	Cohen's d
	M	SD	M	SD		
If P, then Q	5.74 _a	1.21	3.10 _c	1.21	51.93***	2.14
If and only if P, then Q	4.03 _b	1.48	4.47 _b	1.20	1.47	-0.32
Q only if P	3.19 _c	1.21	4.27 _b	1.50	6.28*	-0.85
If not-P, then not-Q	2.42 _d	0.87	3.91 _b	1.33	18.69***	1.77

Note. All ratings were made on a 7-point scale. Means that do not share the same subscript differ at $p < .05$.

MS = miss; FA = false alarm.

* $p < .05$. *** $p < .001$.

Method

Participants. A total of 39 students (approximately 50% female) of the Ecole Supérieure des Sciences Economiques et Commerciales (ESSEC) at Cergy-Pontoise, France, participated in this study. All participants were in their early 20s. The experiment was conducted in French.

Materials. Two scenarios were constructed, the *shop scenario* and the *restaurant scenario*. In the shop scenario, participants were told that they were selling clothes in a shop; they had to decide whether they would offer a customer some help, knowing that there was an instruction to be strictly followed (e.g., "If a customer is touching some clothes, offer him some help"). In the restaurant scenario, participants were told they were establishing a list of providers for the chef; they had to decide whether a provider should be put on the list, again knowing that there was an instruction to be strictly followed (e.g., "If a provider does not offer you a reduced price, do not put him on the list").

Each questionnaire featured both the shop scenario and the restaurant scenario, with the order of the scenarios being counterbalanced across questionnaires. Within each scenario, the four formulations of the conditional instruction were introduced in turn such that each participant responded to all formulations for each of the two scenarios. (For instance, for the shop scenario, the four formulations went as follows: "If a customer is touching some clothes, offer him your help"; "If and only if a customer is

touching some clothes, offer him your help"; "Offer a customer your help only if he is touching some clothes"; and "If a customer is not touching any clothes, do not offer him your help.") For any given formulation, participants were asked to suppose P and to choose one of the three following answers: "I must do Q," "I must not do Q," or "I am free to decide what to do." Subsequently, participants were asked to suppose not-P and again to choose one of the three following answers: "I must do Q," "I must not do Q," or "I am free to decide what to do."

Results and Discussion

As in Experiment 3, each rule formulation differs clearly from the others on at least one dimension, suggesting that these formulations form part of a paradigmatic contrast set for speakers to choose from to convey different nuances of meaning (Lyons, 1977). A first way to look at the results is to consider the most frequent pattern of answers elicited by the participants for each formulation (see Table 4). Regarding the formulations "if P, then Q," "if and only if P, then Q," and "if not-P, then not-Q," there is a clear dominance of a single pattern for each rule (eliciting 82%–90% of answers), whereas the formulation "Q only if P"

Table 4
 Experiment 4: Most Frequent Patterns Associated With Each Formulation of the Conditional Instruction (N = 39)

Formulation	Most frequent pattern	Percentage	
		Frequency with Shop Scenario	Frequency with Restaurant Scenario
If P, then Q	<i>Must do Q</i> when P is the case but <i>free to decide</i> when not-P is the case	82	82
If and only if P, then Q	<i>Must do Q</i> when P is the case but <i>must not do Q</i> when not-P is the case	85	82
If not-P, then not-Q	<i>Free to decide</i> when P is the case but <i>must not do Q</i> when not-P is the case	85	90
Q only if P	<i>Free to decide</i> when P is the case but <i>must not do Q</i> when not-P is the case	56	46
	<i>Must do Q</i> when P is the case but <i>must not do Q</i> when not-P is the case	31	46

elicits two main patterns. (Whatever the formulation, no other pattern elicited more than 13% of answers.)

The dominant pattern elicited by “if P, then Q” is that one must do Q when P is the case but is free to decide what to do when P is not the case. Hence, this formulation succeeds in avoiding MSs yet leaves open the possibility of an FA. The dominant pattern elicited by “if and only if P, then Q” is that one must do Q when P is the case and one must not do Q when P is not the case. This formulation is thus interpreted in a way that prevents both MSs and FAs. The dominant pattern elicited by “if not-P, then not-Q” is that one is free to decide what to do when P is the case but one must not do Q when P is not the case. This formulation thus prevents FAs but leaves open the possibility of an MS. Finally, the formulation “Q only if P” elicits both patterns attached to “if not-P, then not-Q” and “if and only if P, then Q.” Thus, this formulation always prevents FAs but only sometimes prevents MSs.

The corresponding “deontic-logical” way to look at the results is to consider for each formulation of the instruction the frequency with which participants answered *as if* P was deontically necessary or deontically sufficient for Q. To compute the percentages in Tables 5 and 6, we considered participants as (a) answering *as if* P was deontically necessary for Q if they answered that they would have to avoid doing Q in the not-P situation and (b) answering *as if* P was deontically sufficient for Q if they answered that they would have to do Q in the P situation. Whatever the scenario, P was overwhelmingly considered to be deontically necessary for Q with all formulations except “if P, then Q.” Turning to the deontic sufficiency of P in regard to Q, results are unambiguous for the formulations “if P, then Q”; “if and only if P, then Q”; and “if not-P, then not-Q”: This time, P is overwhelmingly deemed as deontically sufficient for Q. With respect to the formulation “Q only if P,” P is deemed deontically sufficient for Q by 36%–51% of participants.

The data shown in Tables 4, 5, and 6 confirm that even in the absence of any contextual hints to the goal structure of the situation, the recipients of a conditional instruction can pick up the concerns of the instructor through his or her choice of a formulation. Indeed, those formulations, which are perceived to be natural when the goal is to avoid MSs (or, alternately, FAs), do lead instructees to avoid MSs (or, alternately, FAs) when the goal structure of the situation is not explicit.

Table 5
Experiment 4: Sufficiency of P in Regard to Q

Formulation	Shop Scenario (%)	Restaurant Scenario (%)
If P, then Q	87 _a	92 _a
If and only if P, then Q	87 _a	85 _a
Q only if P	36 _b	51 _b
If not-P, then not-Q	8 _c	8 _c
$\chi^2(3, N = 39)$	33.45***	29.48***

Note. Percentages reflect proportion of “must do Q” responses when P was given. Proportions that do not share the same subscript differ at $p < .05$.
*** $p < .001$.

Table 6
Experiment 4: Necessity of P in Regard to Q

Formulation	Shop Scenario (%)	Restaurant Scenario (%)
If P, then Q	8 _a	13 _a
If and only if P, then Q	95 _b	95 _b
Q only if P	87 _b	92 _b
If not-P, then not-Q	92 _b	97 _b
$\chi^2(3, N = 39)$	38.52***	26.55***

Note. Percentages reflect proportion of “must not do Q” responses when not-P was given. Proportions that do not share the same subscript differ at $p < .05$.
*** $p < .001$.

Conclusion

The focus of this article has been the formulation and interpretation of conditional instructions, that is, conditionals that relate the occurrence of some event to the undertaking of some action. Drawing an analogy from signal detection theory, we labeled as MS the situation in which the event is occurring but the action is not taken and as FA the situation in which the action is taken without the event occurring. We proposed that the relative expected costs of errors of omission (MSs) and of commission (FAs) determine the goal structure of the situation. This utility distribution determines what problem is in question for a rational hearer: What is to be avoided in the situation? MSs? FAs? Both?

Accordingly, in the initial set of experiments, we demonstrated that goal structure determines the perceived relevance of information on the Wason selection task and that these changes in relevance predict selection patterns better than deontic-logical interpretations of the conditional statement, which did not change at all according to goal structure. Experiments 1A and 1B used the selection task to examine goal structure effects on what information the superior checks to see whether an *if-then* rule has been followed by his subordinates. Experiment 2 revealed a lack of correspondence between performance on the selection task and on a deontic variant of the truth-table task, suggesting that relevance is not communicated through deontic-logical interpretations of the necessity and sufficiency of the conditional relation. In other words, contextual manipulations of goal structure affect the perceived relevance of information in the Wason selection task without influencing the deontic-logical interpretation of the conditionals.

In Experiments 3 and 4, we examined the question of whether different conditional formulations (a) are perceived as more or less felicitous depending on the contextual goal structure and (b) succeed in securing appropriate illocutionary uptake in the absence of explicit information about speaker goals. Experiment 3 indeed showed that the goal structure of the situation caused participants placed in the role of the speaker to prefer different syntactic formulations of the conditional instruction. For example, the usual conditional formulation “if P, then Q” will be perfectly appropriate for situations in which MSs must be avoided but will be infelicitous in situations in which FAs must be avoided.

We then turned to the question of the communicative effectiveness of the different formulations of the conditional that we examined. In Experiment 4, we took the hearer’s perspective and addressed the question as to whether the conditional formulations

studied in Experiment 3 will be correctly understood by hearers and lead them to interpret instructions in line with the speaker's intentions. Our results showed that people have a very clear understanding of conditional directives and that different formulations of the conditional are very effective devices in securing successful illocutionary uptake. We conclude that conditional directives can be viewed as very effective devices for social coordination of action.

How might this social coordination be achieved? Our data seem more consistent with the view that this happens through a shared understanding of goal-based reasoning that specifies utilities in a given context. When the context is not explicitly given, speakers can encode their concerns through different formulations of the conditional directive, which implicitly communicate to the hearer what the speaker wants the hearer to pay attention to. In other words, they successfully communicate what the speaker considers to be relevant in a given situation, and the cooperative hearer simply tries to achieve the speaker's goals as best as possible.

Taken together, these four studies suggest that speakers' perception of the felicity of different kinds of conditional expressions is strongly determined by goal structure (avoid-MS vs. avoid-FA goals) and that hearers' reactions to these conditionals are well aligned with this goal structure, even if hearers have no explicit knowledge of these goals. The results, therefore, suggest that the function of these different formulations of conditional instructions is to direct the hearer's attention to potential consequences of his actions that the speaker considers important. That hearers so successfully detect the speaker's intentions not only suggests—in the language of Austin (1962)—high illocutionary uptake but also successful coordination of action by the speaker and the hearer. For the rational speaker to get what he or she wants done with words, the speaker should choose a form of the conditional that encodes the goal structure implicit in the context, with the knowledge that the hearer should react in a way that will fulfill the speaker's intention. Rationality, here, is thus social and pragmatic, determined by the successful coordination of the speaker and the hearer to achieve shared organizational goals.

Recognizing the social rationality of conditionals allows a new perspective to be taken regarding their development and adaptive functions. Evolution should favor communication and reasoning skills that enable evocation and discussion of joint projects that yield collective benefits (Gärdenfors, 2003). In addition, internalization of conditional directives in inner speech (Luria, 1959) should form a developmental platform for domain-independent planning and reasoning (Luria & la Yudovich, 1956; Vygotsky, 1962; see also Sperber, 2001). Although evolutionary psychologists find it implausible that the human mind would have evolved some domain-general reasoning ability (e.g., Cosmides, 1989; Gigerenzer & Hug, 1992), it does seem to us that one plausible evolved general ability would be the skills needed for speakers to get others to understand one's goals and, conversely, to understand how one's own actions can serve others' goals. Indeed, if these skills are integrated with general goal-based planning and reasoning strategies, then it would not be surprising that people reason especially well about propositions when they can see the point of them.

References

- Almor, A., & Sloman, S. (1996). Is deontic reasoning special? *Psychological Review*, *103*, 374–380.
- Austin, J. L. (1962). *How to do things with words*. Oxford, England: Clarendon Press.
- Bonnefon, J. F., & Hilton, D. J. (2002). The suppression of modus ponens as a case of pragmatic preconditional reasoning. *Thinking & Reasoning*, *8*(1), 21–40.
- Bonnefon, J. F., & Hilton, D. J. (2004). Consequential conditionals: Invited and suppressed inferences from valued outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 28–37.
- Braine, M. S., & O'Brien, D. P. (1991). A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review*, *98*, 182–203.
- 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 we reason? Studies with the Wason selection task. *Cognition*, *31*, 187–276.
- Cummins, D. D. (1996). Evidence for deontic reasoning in 3- and 4-year old children. *Memory & Cognition*, *24*, 823–829.
- Dixon, P. (1982). Plans and written directions for complex tasks. *Journal of Verbal Learning and Verbal Behavior*, *21*, 70–84.
- Evans, J. St. B. T. (1972). Interpretation and matching bias in a reasoning task. *British Journal of Psychology*, *24*, 193–199.
- Evans, J. St. B. T. (1996). Deciding before you think: Relevance and reasoning in the selection task. *British Journal of Psychology*, *87*, 223–240.
- Evans, J. St. B. T., Clibbens, J., & Rood, B. (1995). Bias in conditional inference: Implications for mental models and mental logic. *Quarterly Journal of Experimental Psychology*, *48A*, 644–670.
- Evans, J. St. B. T., Handley, S. J., & Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 321–335.
- Evans, J. St. B. T., & Newstead, S. E. (1977). Language and reasoning: A study of temporal factors. *Cognition*, *8*, 265–283.
- Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). *Human reasoning: The psychology of deduction*. Hove, England: Erlbaum.
- Evans, J. St. B. T., & Over, D. E. (1996). Rationality in the selection task: Epistemic utility versus uncertainty reduction. *Psychological Review*, *103*, 356–363.
- Evans, J. St. B. T. & Over, D. E. (2004). *If*. Oxford, England: Oxford University Press.
- Evans, J. St. B. T., & Twyman-Musgrove, J. (1998). Conditional reasoning with inducements and advice. *Cognition*, *69*, B11–B16.
- Ezard, J. (2000, September 26). Officers dismissed radar warning of Exocet attack on H. M. S. Sheffield. *The Guardian*, p. 43.
- Fairley, N., Manktelow, K. L., & Over, D. E. (1999). Necessity, sufficiency, and perspective effects in causal conditional reasoning. *Quarterly Journal of Experimental Psychology*, *52A*, 771–790.
- Fiddick, L. (2000). *A word of precaution: Not all deontic rules support perspective effects on the Wason selection task*. Unpublished manuscript.
- Fiddick, L., Cosmides, L., & Tooby, J. (2000). No interpretation without representation: The role of domain-specific representations in the Wason selection task. *Cognition*, *77*, 1–79.
- Fillenbaum, S. (1978). How to do some things with IF. In J. W. Cotton & R. L. Klatzky (Eds.), *Semantic factors in cognition* (pp. 169–214). Hillsdale, NJ: Erlbaum.
- Friedrich, J. (1993). Primary error detection and minimization (PEDMIN) strategies in social cognition: A reinterpretation of confirmation bias phenomena. *Psychological Review*, *100*, 298–319.
- Gärdenfors, P. (2003). *How homo became sapiens*. Oxford, England: Oxford University Press.

- George, C. (1991). Facilitation in the Wason selection task with a consequent referring to an unsatisfactory outcome. *British Journal of Psychology*, 82, 463–472.
- Gigerenzer, G., & Hug, K. (1992). Domain specific reasoning, social contracts, and perspective change. *Cognition*, 43, 127–171.
- Giroto, V. (1995). Contextual factors in deontic reasoning. *Thinking and Reasoning*, 1, 333–338.
- Giroto, V., Kemmelmeier, M., Sperber, D., & van der Henst, J. B. (2001). Inept reasoners or pragmatic virtuosos? Relevance and the deontic selection task. *Cognition*, 81, B69–B76.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: 3. Speech acts* (pp. 41–58). New York: Academic Press.
- Harris, P. L. (1996). Beliefs, desires and language. In P. Carruthers & P. K. Smith (Eds.), *Theories of theories of mind* (pp. 200–220). Cambridge, England: Cambridge University Press.
- Harris, P. L., & Núñez, M. (1996). Understanding of permission rules by preschool children. *Child Development*, 67, 1572–1591.
- Higgins, E. T. (2000). Making a good decision: Value from fit. *American Psychologist*, 55, 1217–1230.
- Hilpinen, R. (2001). Deontic logic. In L. Goble (Ed.), *The Blackwell guide to philosophical logic* (pp. 159–182). Oxford, England: Blackwell.
- Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin*, 118, 248–271.
- Holyoak, K. J., & Cheng, P. W. (1995). Pragmatic reasoning with a point of view. *Thinking and Reasoning*, 1, 289–313.
- Jackson, F. (Ed.). (1991). *Conditionals*. Oxford, England: Oxford University Press.
- Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, 109, 646–678.
- Kenny, A. (1966). Practical inference. *Analysis*, 26, 65–75.
- Kirby, K. N. (1994). Probabilities and utilities of fictional outcomes in Wason's four-card selection task. *Cognition*, 51, 1–28.
- Klauer, K. C. (1999). On the normative justification for information gain in Wason's selection task. *Psychological Review*, 106, 215–222.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge, England: Cambridge University Press.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge, MA: MIT Press.
- Lewicka, M. (1989). Towards a pragmatic perspective of cognition: Does evaluative meaning influence rationality on lay inferences? *Polish Psychological Bulletin*, 20, 267–285.
- Lieberman, N., & Klar, Y. (1996). Hypothesis-testing in Wason's selection task: Social exchange, cheating detection or task understanding. *Cognition*, 58, 127–156.
- Light, P., Giroto, V., & Legrenzi, P. (1990). Children's reasoning on conditional promises and permissions. *Cognitive Development*, 5, 369–383.
- Luria, A. R. (1959). The directive function of speech in development and dissolution: Part I. *Word*, 15, 341–352.
- Luria, A. R., & la Yudovich, F. I. (1956). *Speech and the development of mental processes in the child*. Harmondsworth, England: Penguin.
- Lyons, J. (1977). *Semantics* (Vol. 1). Cambridge, England: Cambridge University Press.
- Manktelow, K. I., & Over, D. E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition*, 39, 85–105.
- McKenzie, C. R. M., Ferreira, V. S., Mikkelsen, L. A., McDermott, K. J., & Skrabble, R. P. (2001). Do conditional hypotheses target rare events? *Organizational Behavior and Human Decision Processes*, 35, 291–309.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608–631.
- Oaksford, M., & Chater, N. (1995). Information gain explains relevance, which explains the selection task. *Cognition*, 57, 97–108.
- Oaksford, M., Chater, N., & Grainger, B. (1999). Probabilistic effects in data selection. *Thinking and Reasoning*, 5, 193–243.
- Oaksford, M., Chater, N., & Larkin, J. (2000). Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 883–899.
- Ormerod, T. C., Manktelow, K. I., & Jones, G. V. (1993). Reasoning with three types of conditional: Biases and mental models. *Quarterly Journal of Experimental Psychology*, 46A, 653–677.
- 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. (2003). *Adaptation or adaptability in deontic reasoning*. Unpublished manuscript, University of Wales at Cardiff, United Kingdom.
- Politzer, G., & Nguyen-Xuan, A. (1992). Reasoning about conditional promises and warnings: Darwinian algorithms, mental models, relevance judgments or pragmatic schemas? *Quarterly Journal of Experimental Psychology*, 44A, 401–421.
- Roth, P. (2002). *The human stain*. London: Vintage Books.
- Schaeken, W., & Schroyens, W. (2000). The effect of explicit negatives and of different contrast classes on conditional syllogisms. *British Journal of Psychology*, 91, 533–550.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An enquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.
- Schroyens, W. J., Schaeken, W., & d'Ydewalle, G. (2001). The processing of negations in conditional reasoning: A meta-analytic study in mental model and/or mental logic theory. *Thinking and Reasoning*, 7, 121–172.
- Sperber, D. (2001). An evolutionary perspective on testimony and argumentation. *Philosophical Topics*, 29, 401–413.
- Sperber, D., Cara, F., & Giroto, V. (1995). Relevance theory explains the selection task. *Cognition*, 57, 31–95.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford, England: Blackwell.
- Staller, A., Sloman, S. A., & Ben-Zeev, T. (2000). Perspective effects in nondeontic versions of the Wason selection task. *Memory & Cognition*, 28, 396–405.
- Staudenmayer, H. (1975). Understanding reasoning with meaningful propositions. In R. J. Falmagne (Ed.), *Reasoning: Representation and process* (pp. 55–79). Hillsdale, NJ: Erlbaum.
- Stevenson, R. J., & Over, D. E. (1995). Deduction from uncertain premises. *Quarterly Journal of Experimental Psychology*, 48A, 613–643.
- Stevenson, R. J., & Over, D. E. (2001). Reasoning from uncertain premises: Effects of expertise and conversational context. *Thinking and Reasoning*, 7, 367–390.
- Strawson, P. F. (1952). *Introduction to logical theory*. London: Methuen.
- Strawson, P. F. (1964). Intention and convention in speech acts. *Philosophical Review*, 73, 439–460.
- Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory & Cognition*, 22, 742–758.
- Thompson, V. A. (2000). The task-specific nature of domain-general reasoning. *Cognition*, 76, 209–268.
- Thompson, V. A., & Mann, J. M. (1995). Perceived necessity explains the dissociation between logic and meaning: The case of "only if." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 1554–1567.
- Vygotsky, L. S. (1962). *Thought and language*. Cambridge, MA: MIT Press.
- Wason, P. C. (1965). The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior*, 4, 7–11.
- Wason, P. C., & Johnson-Laird, P. N. (1972). *The psychology of reasoning*. London: Batsford.
- Weber, E. U., & Hilton, D. J. (1990). Contextual effects in the interpretation of probability words: Perceived base-rate and severity of events.

- Journal of Experimental Psychology: Human Perception and Performance*, 16, 781–789.
- Wilensky, R. C. (1983). *Planning and understanding: A computational approach to human reasoning*. Reading, MA: Addison Wesley.
- Wilson, D., & Sperber, D. (2004). Relevance theory. In G. Ward & L. Horn (Eds.), *Handbook of pragmatics* (pp. 607–632). Oxford, England: Blackwell.
- Wright, P., & Hull, A. J. (1986). Answering questions about negative conditionals. *Journal of Memory and Language*, 25, 691–709.
- Wright, P., & Wilcox, P. (1979). When two no's nearly make a yes: A study of conditional imperatives. In P. A. Kollers, M. E. Wrolstad, & H. Bouma (Eds.), *Processing of visible language 1* (pp. 413–430). New York: Plenum Press.

Appendix

Example Task Used in Experiment 1

For clarity, the sections used in the avoid-FA conditions are italicized in their entirety.


Ship-Submarine Story

You are the weapons officer aboard a warship. You know that the task of the sonar operators who are looking for submarines is very difficult. Often unusual “blips” are due to rock formations or shoals of fish.

In addition, even though allied submarines produce special kinds of “blip” that your operators have learned to identify, it is not always possible to differentiate them from the blips produced by enemy submarines. An unusual blip could come from an allied submarine in an unusual position.

In the same way, an apparently familiar “blip” could be produced by an enemy submarine in an unusual position, or one whose sonar trace has been deformed by the state of the sea.

At the moment, your operators are performing a simulation exercise, which puts them in exactly the same conditions as they will experience in reality through producing the classic sonar traces of allied submarines, as well as unusual “blips.” You have four cards in front of you. On one side, one kind of case is written (allied or enemy submarine), and on the other side, the operator’s response (did nothing or fired).



You have decided to test the speed and alertness of your operators. No enemy submarine should get close enough to be able to attack your ship without being stopped (**avoid-MS version**)/ *You are worried that your operators might be trigger-happy and that they might one day destroy an allied submarine (avoid-FA version)*. What will you do to check that your operators are indeed following the rule “If you see an unusual ‘blip,’ then launch the depth charges”?

Judge the relevance of turning over each card for helping you in your task.

Received July 22, 2003
 Revision received May 16, 2005
 Accepted May 24, 2005 ■