Psychological Review
2009, Vol. 116, No. 4, 888–907

0033-295X/09/$12.00 DOI: 10.1037/a0017186

A Theory of Utility Conditionals: Paralogical Reasoning From Decision-Theoretic Leakage

Jean-François Bonnefon
Centre national de la recherche scientifique and Université de Toulouse

Many "if *p*, then *q*" conditionals have decision-theoretic features, such as antecedents or consequents that relate to the utility functions of various agents. These decision-theoretic features leak into reasoning processes, resulting in various paralogical conclusions. The theory of utility conditionals offers a unified account of the various forms that this phenomenon can take. The theory is built on 2 main components: (1) a representational tool (the utility grid), which summarizes in compact form the decision-theoretic features of a conditional, and (2) a set of folk axioms of decision, which reflect reasoners' beliefs about the way most agents make their decisions. Applying the folk axioms to the utility grid of a conditional allows for the systematic prediction of the paralogical conclusions invited by the utility grid's decision-theoretic features. The theory of utility conditionals significantly extends the scope of current theories of conditional inference and moves reasoning research toward a greater integration with decision-making research.

Keywords: reasoning, conditionals, inference, decision, utility

Conditionals play a pivotal role in rational thinking. The cognitive manipulation of "if . . . then" relations is central to reasoning processes, but it is also a critical aspect of decision making: Careful decisions reflect the assessment of what would happen if a given choice were made and what would be the utility of this outcome to relevant agents. The central argument of this article is that the way we manipulate conditionals for decision making leaks into the way we manipulate them for reasoning. The central contribution of the article is then to offer a theory of the paralogical inferences invited by the decision-theoretic features of conditional statements.

The article runs as follows: I start by emphasizing that conditionals allow for conclusions to be reached and decisions to be made—two activities that traditionally fall within the scope of two different subfields of psychology. I then raise the concern that these two activities cannot be plausibly considered insulated from each other. In particular, the way individuals use conditionals for decision making is bound to influence the inferences they make from conditionals, as soon as these conditionals show decision-theoretic features.

I then remark that this fact has long been neglected due to the disinterested view toward conditional statements that used to dominate the psychology of reasoning. Although this disinterested view was tenable as long as the psychology of reasoning focused on abstract, knowledge-lean contents, the increased interest in knowledge-rich contents has allowed the decision-theoretic features of conditionals to leak into reasoning experiments. This has led to a variety of empirical phenomena, calling for a unified theoretical account.

I then move on to the central contribution of the article: the theory of utility conditionals. The theory of utility conditionals is built on two main components: a representational tool (the utility grid), which summarizes in compact form the decision-theoretic features of a conditional statement, and a set of folk axioms of decision, which captures reasoners' likely beliefs about the way most agents make their decisions. Applying the folk axioms to the utility grid of a conditional allows for the prediction of paralogical conclusions invited by its decision-theoretic features. I review the empirical findings supporting the predictions of the theory and contribute some new empirical data supporting other predictions of the theory. Finally, I explore several possible extensions and refinements of the basic theory, and I consider how it might be integrated with three broader theories of conditional inference: the Bayesian theory, the suppositional theory, and the model theory.

Conditionals in Decision and Reasoning

Conditionals are the natural mode of expression of *hypothetical thinking,* that is, the ability to reach conclusions and make decisions on the basis of the mental simulation of states of affairs (Evans, 2007). As noted by Evans and Over (2004), this is what makes *if* one of the most important and interesting words in the language. Reflecting the twofold nature of hypothetical thinking, conditionals are involved in the production of beliefs as well as in the production of actions. Compare 1a and 1b:

1a. If it's a cow, then it goes "moo."

1b. If I move my knight there, then I'll win the game.

Jean-François Bonnefon, Cognition, Langues, Langage et Ergonomie (CLLE), Centre national de la recherche scientifique, Paris, France, and CLLE, Université de Toulouse, Toulouse, France.

Wim De Neys, Jonathan Evans, Mike Oaksford, David Over, and Guy Politzer all made helpful comments on drafts of this article.

Correspondence concerning this article should be addressed to Jean-François Bonnefon, Cognition, Langues, Langage et Ergonomie, Maison de la recherche, 5 allées A. Machado, 31058 Toulouse Cedex 9, France. E-mail: bonnefon@univ-tlse2.fr

Conditional 1a is oriented toward the production of a belief. It expresses that, under the supposition that an animal is a cow, the conclusion is warranted that this animal goes "moo." In contrast, 1b is oriented toward the production of an action. It expresses that, given its beneficial consequences, the action of moving the knight is warranted.

Following Baron (2000), I define *utility* as a measure of goal achievement. For example, an action that achieves the goals of the decision maker to a great extent has a large positive utility, whereas an action that hinders the goals of the decision maker to some small extent has a small negative utility. I use the term *utility conditional* to denote a conditional whose antecedent or consequent bears on the utility function of some agent, as in 1b. Other conditionals, such as 1a, are not known to have this property. Borrowing a term from Oaksford and Chater (1994), I call them *disinterested* conditionals.

It is important to note right away that utility conditionals can be used in a disinterested fashion. For example, 1b can be used without any reference to the goals of any decision maker, in other words, as just any disinterested conditional. It would then simply express that under the supposition that I move my knight there, the conclusion that I will win the game is warranted—independent of any assumption about whether I want to win the game.

Now, it is tempting to define *decision making* as the assessment of potential actions through the manipulation of utility conditionals and to draw a clear separation with conditional *reasoning,* defined as the production of new beliefs through the manipulation of (a) disinterested conditionals and (b) utility conditionals qua disinterested conditionals. Adopting this perspective amounts to considering that the role utility conditionals play in decision making has no specific impact on the conclusions that people draw from them in reasoning contexts—or, in other terms, that the processes involved in conditional reasoning are insulated from the processes involved in conditional decision making. This is, broadly defined, the perspective adopted in the deduction paradigm in the study of human reasoning (see for reviews, Evans, 2002; Politzer, 2007).

How tenable is that assumption, though? In recent years, arguments have been made that the deduction paradigm and its logicist grounds are outdated (Evans, 2002), that the processes underlying reasoning and decision making are essentially the same (Evans, 2007; Evans & Over, 1996; Stanovich, 1999), and that progress in the study of conditional reasoning will require that it be integrated with research on decision making (Evans & Over, 2004). This realization is the end point of a long process that I briefly recapitulate in the next section. The process started with the introduction of thematic materials in conditional reasoning experiments and was once largely confined to specific tasks or thematic contents, but then it was revived and generalized to a larger array of tasks, up to the point where it called for redefining the scope of reasoning research.

Beyond Disinterested Conditionals

For a long while, the question of whether decision-making processes might leak into conditional reasoning experiments was irrelevant, and doubly so. First, because the psychology of conditional reasoning was, de facto, limited to the study of disinterested conditionals. Second, because experiments framed conditional reasoning itself as a contemplative activity, disconnected from any goal other than pondering the microworld featured in the experiment. These two constraints, however, were progressively relaxed, and experimenters started to use richer materials and richer contexts.

Opening the Thematic Gates

The psychology of conditional reasoning originated in the conviction that standard deductive logic was the norm for correct reasoning against which human responses ought to be compared in order to assess people's rationality (Evans, 2002). Because the objective was to assess people's formal deductive competence, it made sense to present them with abstract, algebraic arguments such as 2a–2c below or arguments with neutral contents and arbitrary relations such as 3a–3c.

2a. If A, then B.

2b. A is true.

2c. Therefore, B is true.

3a. If there is a circle, then there is a triangle.

3b. There is not a triangle.

3c. Therefore, there is not a circle.

Conditionals 2a and 3a are disinterested conditionals whose antecedents and consequents are not known to bear on anyone's utility function. As long as reasoning experiments involved conditionals of that sort and that sort only, then the observation that other conditionals could be used in decision making was irrelevant.

However, experimenters gradually started to consider the kind of realistic conditionals that were more likely to be encountered in daily life. One reason for doing so was to investigate the effect of background knowledge on conditional reasoning: When they reason about thematic contents, reasoners are likely to take into account their personal beliefs about whether (and to what extent) the antecedent p is a necessary or sufficient condition of the consequent q (Byrne, 1989; Cummins, Lubart, Alksnis, & Rist, 1991; De Neys, Schaeken, & d'Ydewalle, 2002, 2003; Politzer & Bonnefon, 2006; Thompson, 1994). Oftentimes, these realistic conditionals received labels reflecting the flavor they gave to an interaction, for example, *threat, intention,* or *prediction* (e.g., Dieussaert, Schaeken, & d'Ydewalle, 2002; Evans & Twyman-Musgrove, 1998; Fillenbaum, 1975, 1976, 1978; Newstead, Ellis, Evans, & Dennis, 1997; Staudenmayer, 1975). What is interesting for my current purpose is not to analyze in detail these early typologies but to point out that, alongside disinterested conditionals, they included obvious as well as less obvious instances of utility conditionals. Consider the following examples:

4a. If you mow the lawn, then I will pay you $5.

4b. If you stand by the pillar, then you'll be served immediately.

4c. If you turn the radio on one more time, then I will hit you.

Statements 4a–4c are utility conditionals, which unambiguously feature contents that bear on the utility function of a decision maker. This utility takes the form of money for the first conditional (a *promise* from Fillenbaum, 1975), of being served faster for the second (a *tip* from Newstead et al., 1997), and of getting hurt for the third (i.e., a negative utility; this is a *threat* from Dieussaert et al., 2002).

Several authors did notice that there was something special with promises, threats, tips, and warnings, which were accordingly targeted for specific study (e.g., Beller, Bender, & Kuhnmünch, 2005; Evans, Neilens, Handley, & Over, 2008; Evans & Twyman-Musgrove, 1998; Ohm & Thompson, 2004, 2006; Verbrugge, Dieussaert, Schaeken, & Van Belle, 2004). I address this line of research later in the article; for the time being, I wish to emphasize that other realistic conditionals such as the following, which have attracted less attention, can also be construed as utility conditionals:

5a. If I have well-trained muscles, then I will be lucky with women.

5b. If he goes to Washington, then he will get drunk.

Conditional 5a was labeled a *prediction* by Dieussaert et al. (2002). According to these authors, the main characteristic of a prediction conditional is its expression of a weak causal relation that is based on loose ideas rather than evidence. This analysis focuses on the question of whether the truth of the antecedent in Conditional 5a warrants the belief in its consequent—in other words, it treats 5a as a disinterested conditional. However, 5a clearly has the characteristics of a utility conditional, which can be used for decision making: Assuming that being lucky with women has positive utility for the speaker, 5a is an argument for having well-trained muscles, that is, for working out.

Fillenbaum (1975) considered 5b an example of a mere *causal* conditional: Knowing that the person will go to Washington warrants the belief that he will get drunk.

However, it is conceivable that the consequent of 5b has positive or negative utility for various agents. For example, the agent denoted as "he" might be looking for a chance to get drunk, in which case 5b would be an argument for going to Washington. The speaker asserting 5b, though, might have a different perspective. For example, the speaker might be the agent's parent, who dislikes the possibility that the agent might get drunk; in that case, 5b would be an argument for not letting the agent go to Washington.

Thus, the shift toward everyday, realistic contents has served as a Trojan horse for utility conditionals to disseminate into reasoning experiments. Before I consider the consequences of this dissemination, I briefly address another utilitarian shift that prepared a redefinition of the scope of reasoning research.

Goal-Directed Reasoning in the Selection Task

As experimenters started to use everyday statements in conditional argument tasks, an important development occurred within a kindred paradigm, the Wason (1966) selection task. In the original form of the selection task, participants are presented with the pictures of four cards and told that the cards all have a letter on one side and a number on the other side. Two of the cards display their lettered side (e.g., *A* and *B*), and the two other cards display their numbered side (e.g., *2* and *5*). Participants are instructed to choose those cards and only those cards that need to be turned over in order to decide whether the following rule is true: "If there is a vowel on one side, then there is an even number on the other side."

The finding that the vast majority of participants fail to select the *A* and *5* cards (which are the only cards that can yield a vowel–odd number falsifying combination) generated a massive literature (see Klauer, Stahl, & Erdfelder, 2007, for a review). One turning point in this literature was the discovery that some everyday conditionals could lead to a dramatic change in participants' responses. More specifically, Griggs and Cox (1982) instructed participants to imagine that they were police officers observing people drinking in a bar and that they were checking whether the following rule was followed: "If a person is drinking beer, then that person must be over 19 years of age." With these instructions and this material, a large majority of participants selected the "drinking beer" and "under 19 years of age" cards—precisely the combination that is almost never selected in abstract versions of the task.

Manktelow and Over (1991) eventually offered an elegant explanation for this result by observing that participants in the Griggs and Cox (1982) drinking rule version of the task were not reasoning just for the sake of reasoning—rather, their reasoning was aimed at making the best possible decision with respect to a particular goal, namely, discovering contraveners of the drinking legislation. Consequently, they selected the cards that were the most likely to help them achieve this goal—that is, the cards with the highest expected utility. This account correctly predicts that any manipulation that affects the expected utilities of the four cards also affects participants' choices of cards, even when the conditional rule itself is kept constant (Gigerenzer & Hug, 1992; Hilton, Kemmelmeier, & Bonnefon, 2005; Manktelow & Over, 1991; Perham & Oaksford, 2005; Politzer & Nguyen-Xuan, 1992).

This line of research emphasized the fact that conditional reasoning ought not to be construed as a purely contemplative activity, detached from practical purposes. This development, together with the observation that reasoning experiments gradually opened up to conditionals featuring actions and their utility, called for a large redefinition of the scope of reasoning research.

Reasoning Research Redefined

For historical reasons I have briefly considered, the psychology of conditional reasoning was long committed to a narrow definition of its object: Conditional reasoning was to be construed as the contemplative manipulation of disinterested conditionals.

As long as reasoning experiments took place in contemplative contexts and involved only disinterested conditionals, it was a tenable assumption that the processes involved in these experiments were insulated from the processes involved in decision making. Methodological shifts, however, extended this traditional paradigm into two broadly orthogonal directions: Some studies moved from contemplative contexts to goal-directed contexts, and others moved from disinterested contents to utility contents.

As a consequence, current research on conditional reasoning broadly falls into one of four categories. The first category corresponds to the traditional, contemplative–disinterested approach. The second category corresponds to reasoning on disinterested conditionals in goal-directed contexts. Proving a theorem, checking whether a procedure was applied, and detecting a contradiction in a body of laws are examples of this kind of reasoning. Studies of perspective shifts in the selection task typically belong to this category.

More relevant to my current purpose are the last two categories, which involve utility conditionals. Just as disinterested conditionals can be used in either goal-directed or contemplative contexts, so, too, can utility conditionals. Arguing for a policy or influencing the actions of other people are examples of goal-directed uses of utility conditionals (Corner & Hahn, 2007; Hahn & Oaksford, 2006, 2007; Thompson, Evans, & Handley, 2005). Predicting the behavior and attitude of other people are examples of contemplative uses of utility conditionals (Beller et al., 2005; Bonnefon & Hilton, 2004; Evans et al., 2008; Evans & Twyman-Musgrove, 1998; Ohm & Thompson, 2004, 2006; Verbrugge et al., 2004). These studies, which provide empirical support for the theory I offer in this article, are reviewed after I introduce the theory of utility conditionals.

Utility Grids

So far, I have opposed utility and disinterested conditionals in a loose way, noting that utility conditionals are special in that they connect a potential action (the antecedent) with its valued consequences for some agents. I now offer a more systematic examination of what characterizes utility conditionals, by formalizing the key notion of the *utility grid* of a conditional statement. I argue that all conditionals have a utility grid, although that of disinterested conditionals is a rather special one.

My notion of a utility grid draws on various typologies of everyday conditional statements (Amgoud, Bonnefon, & Prade, 2007; Bonnefon & Hilton, 2004, 2005; Evans, 2005; Legrenzi, Politzer, & Girotto, 1996; López-Rousseau & Ketelaar, 2004, 2006). Later, in the Remarkable Utility Grids section, I show how the new notion elegantly encompasses these previous typologies and extends them to new, uncharted territories.

The Utility Grid of Conditional Statements

Let us consider that p and q in any conditional "if p, then q" are two actions undertaken by agents x and x', respectively. For this to be true of all conditionals, we need a special convention when p or q cannot be construed as the action of an identified human agent. Consider, for example,

6a. If there is a financial crisis, then Obama will be elected president.

6b. If there is a storm, then my pear trees will be damaged.

In both conditionals, neither p nor q can be easily construed as an action undertaken by a given agent. The financial crisis, the storm, Obama being elected, the trees being damaged—none of this is really something for somebody to do. The convention that I adopt in such a case is to consider that p (or q) is still an action, but one undertaken by a special, neutral agent ω. The agent ω can be thought of as "the world," or the body of laws that govern the world.

Equipped with this convention, let us repeat the basic assumption of this section: For all conditionals "if p, then q," p and q are two actions undertaken by some agents x and x', respectively. The two agents, x and x', belong to A, the set of all agents, which includes ω, the neutral agent.

Let us now expand the initial conditional "if p, then q" into a developed form: "If agent x takes action p that has utility u to agent y, then agent x' will take action q that has utility u' to agent y'."

The decision-theoretic information contained in this expanded formulation can be represented in the following utility grid of the conditional:

$$\left\{\begin{array}{cccc} \text{If} & x & u & y \\ \text{Then} & x' & u' & y' \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\},$$

The first row of the grid contains the information related to the *if* clause of the conditional. That is, it displays the agent who can potentially take action p (left column) and the utility of this action (central column) for a given target (right column). The second row of the grid contains the corresponding information with respect to the *then* clause of the conditional. The utility grid can also be written in this abbreviated form:

$$\left\{\begin{array}{ccc} x & u & y \\ x' & u' & y' \end{array}\right\}.$$

In the rest of this section, for familiarization purposes, I use the full notation of utility grids. Beginning in the next section, and for the remainder of the article, I use the abbreviated notation.

Disinterested conditionals do have a utility grid, that is, one where $u = u' = 0$. Note that $u = 0$ means that action p is not known to have any utility for any agent. By convention, an action that has no known utility for any agent has zero utility and A as a target (A, as previously noted, is the whole set of agents).

Some agents in A will prove to be especially relevant to the analysis of utility conditionals, namely the agent asserting the conditional and the agent to whom it is addressed. These agents will be denoted s (for "speaker") and h (for "hearer"), respectively. It is also convenient to use the notation e (for "someone else") as a generic label for identified agents who are not the speaker, the hearer, or the neutral agent (i.e., agents in A who are not s, h, or ω).

Although utility may take any positive or negative value, it is enough for the time being to simply consider its sign. Thus, I consider that u and u' take their values from $\{-, 0, +\}$, where $-$ and $+$, respectively, stand for any significantly negative and positive values. Finally, I write of *proximal utility* for u (the utility involved in the *if* clause of the conditional) and of *distal utility* for u' (the utility involved in the *then* clause of the conditional). Proximal utility refers to the immediate foreseeable consequences of action p, whereas distal utility refers to the more long-term consequences of action p brought about by virtue of producing q.

Now that we are equipped with these notations, let us consider a few examples of utility grids, starting from the purely disinterested conditional "if it is a triangle, then it is not a square." Being a triangle, just as not being a square, is an action of the neutral agent, which has no known utility for the whole set of agents. Thus, this conditional has the following utility grid:

$$\left\{\begin{array}{cccc} \text{If} & \omega & 0 & A \\ \text{Then} & \omega & 0 & A \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

Consider now the conditional "if you move just two inches, then I'll be able to see the screen." Moving two inches is an action of the hearer that has no known proximal utility for anyone (let us assume in particular that it has no significantly negative utility for the hearer). However, if that action is taken, a state of affairs will occur that has positive distal utility for the speaker. The utility grid of the conditional is thus

$$\left\{\begin{array}{cccc} \text{If} & h & 0 & A \\ \text{Then} & \omega & + & s \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

Finally, the conditional "if I let you get away with it, then my boss will fire me" has the utility grid

$$\left\{\begin{array}{cccc} \text{If} & s & + & h \\ \text{Then} & e & - & s \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\},$$

where *e* is an agent (the boss) who is neither the speaker nor the hearer nor the neutral agent.

A single utility grid will not be enough to describe some conditionals. Consider again Conditional 5b, "If he goes to Washington, then he will get drunk." As previously explained, different agents may have different takes on the consequences featured in this conditional. Imagine that the speaker is the mother of the teenager referred to as "he." He has plans to spend a wild night in Washington, and his mother disapproves of it. In that example, although the action of going to Washington, per se, does not have any utility to any agent, its consequences that "he" (some agent *e*) will get drunk has positive utility for *e* but negative utility for the speaker, *s*. These different utility functions are reflected in separate subgrids, leading to the following complex utility grid of Conditional 5b:

$$\left\{\begin{array}{cccc} \text{If} & e & 0 & A \\ \text{Then} & e & + & e \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}$$

$$\left\{\begin{array}{cccc} \text{If} & e & 0 & A \\ \text{Then} & e & - & s \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

More intricate situations may require a greater number of subgrids, as appropriate. There is no guarantee that all possible subgrids of a given conditional are considered by a given reasoner, and contextual cues or perspective manipulations may lead reasoners to prefer one subgrid in their reasoning to the detriment of the others.

Remarkable Utility Grids

Some specific subclasses of utility conditionals have already been identified and scrutinized, and various typologies or characterizations were offered for these subclasses. Of course, because utility conditionals were not yet identified as a group, these efforts at organization and characterization were neither concerted nor systematic. Nevertheless, all these classification schemes are captured by the utility grid framework (see Table 1 for a selection of examples).

López-Rousseau and Ketelaar (2006) offered a simple scheme for identifying conditionals as promises, threats, warnings, or tips, building on Evans (2005) and López-Rousseau and Ketelaar (2004). Their algorithm operates in situations where an "if *p*, then *q*" conditional is asserted by a speaker to a listener and where *p* is a potential action of this listener. When these conditions are met, the classification algorithm applies in two steps. If *q* is a potential action of the speaker, then the conditional is a promise when *q* is desirable for the listener and a threat when *q* is undesirable for the listener. If *q* is not a potential action of the speaker, then the conditional is a tip when *q* is desirable for the listener and a warning when *q* is undesirable for the listener. This simple algorithm (for further formalization, see Amgoud et al., 2007) proved quite successful, capturing over 90% of the classifications made by human subjects.

Table 1
Examples of Remarkable Utility Grids (in Abbreviated Notation)

Promise	Warning	Persuasion
$\left\{\begin{array}{ccc} h & \bullet & \bullet \\ s & + & h \end{array}\right\}$	$\left\{\begin{array}{ccc} h & \bullet & \bullet \\ \omega & - & h \end{array}\right\}$	$\left\{\begin{array}{ccc} e & \bullet & \bullet \\ \bullet & + & A \end{array}\right\}$
Consequential	**Request**	**Veiled threat**
$\left\{\begin{array}{ccc} e & \bullet & \bullet \\ \bullet & - & e \end{array}\right\}$	$\left\{\begin{array}{ccc} h & 0 & A \\ \omega & + & s \end{array}\right\}$	$\left\{\begin{array}{ccc} h & - & s \\ \omega & - & h \end{array}\right\}$

Note. Some grids (promise, warning, persuasion, consequential) capture previously identified subclasses of utility conditionals; other grids (request, veiled threat) point to new subclasses obtained by studying the permutations of the utility grid. *h* = hearer; ● = any legitimate value of the parameter; *s* = speaker; + = any significantly positive value; ω = a neutral agent; – = any significantly negative value; *e* = someone else; 0 = zero utility of an action that has no known utility for any agent; A = the whole set of agents.

This algorithm for classifying promises, threats, tips, and warnings can readily be translated into the utility grid framework. First, note that the algorithm assumes that *p* is an action of the hearer. This means that all four types of conditionals have a grid featuring *h* in the upper left cell. In all four conditionals, *q* is expected to be of some value to the listener; therefore, all grids will feature *h* in the lower right cell. Whether *q* is an action of the speaker is represented by having or not having *s* in the lower left cell of the grid. Finally, whether *q* is desirable or undesirable for the listener is represented by a + or a – sign in the lower central cell of the grid.

Thus, a promise (in the sense used in the algorithm) has the following utility grid:

$$\left\{\begin{array}{cccc} \text{If} & h & \bullet & \bullet \\ \text{Then} & s & + & h \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\},$$

where the black dot stands for a blank parameter, which is allowed to take any legitimate value. A threat has the same utility grid except that the + sign becomes a – sign. Tips have the same utility grid except that the value in the lower left cell is no longer *s* but any agent in *A* who is not *s*.

Beller et al. (2005; see also Legrenzi et al., 1996) introduced an additional requirement for an "if *p*, then *q*" conditional to be considered as a promise (or a prototypical social contract). They argued that not only should *q* benefit the hearer, but *p* should benefit the speaker. In terms of utility grids, they suggested that the grid above should be replaced with this more specific grid:

$$\left\{\begin{array}{cccc} \text{If} & h & + & s \\ \text{Then} & s & + & h \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

This sounds sensible, because we usually think of a promise as an exchange of favors: The speaker gives something in exchange for receiving something. Intuitively, we would not call a "promise" a conditional where the speaker gives something in exchange for giving something. It is quite hard, however, to check this intuition by finding a compelling example of a conditional with this critical utility grid:

$$\left\{\begin{array}{cccc} \text{If} & h & - & s \\ \text{Then} & s & + & h \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

Consider the following attempts:

7a. If you slap me in the face, then I will pay you 20 euros.

7b. If you steal my computer, then I will invite you to dinner.

It seems that we spontaneously adjust our perceptions of the speaker's utility function in both examples. Most people would assume that the speaker of 7a has masochistic tendencies and actually finds pleasure in being slapped. Conditional 7b might not so readily make sense, but we also seem to be looking for a reason why the stealing of a computer might increase the speaker's utility, and some of us eventually come to suspect an insurance scam.

That is not to say that the typical grid of a promise is not that suggested by Beller et al. (2005) and Legrenzi et al. (1996). Rather, it is to say that this grid is an attractor for all the grids that fit the Evans (2005) and López-Rousseau and Ketelaar (2006) schema for promises. When we encounter one of these, we seem to fill in the blanks in order to reconstitute the prototypical grid identified by Beller et al. (2005) and Legrenzi et al. (1996). This phenomenon is nicely illustrated by the following example, as noted by Van Canegem-Ardijns and Van Belle (2008):

8. If you date my daughter, then I will give you a C.

The C could be a good grade or a bad grade, depending on the usual performance of the student. The beauty of this example is that, depending on the usual performance of the student and thus the utility of getting a C, we seem to switch from a threat grid to a promise grid. If C is a bad grade for the student, then we understand that the professor does not want this student around her daughter. If C is an improvement over the usual grades, we understand that the professor is encouraging the student to date her daughter, for reasons we cannot fathom but that we nevertheless assume to exist.

Promises—and, more generally, inducement conditionals—are not the only subclass of utility conditionals that have appeared in previous classification schemes. For example, Bonnefon and Hilton (2004) gave the following definition of *consequential conditionals.*

> [An] "if *p,* then *q*" statement is a consequential conditional when (a) *p* is an action of a third party, the agent, who is neither the speaker nor the hearer; (b) *q* is a consequence of taking action *p;* (c) *q* is valued: it is either a good (positive, desirable) or a bad (negative, undesirable) outcome to the agent; (d) to take or not to take action *p* is a matter of choice to the agent; and (e) the agent knows *q* to be a consequence of *p.* (p. 29)

The statement is called a positive consequential when *q* is desirable, as in Conditional 9a below, and a negative consequential when *q* is undesirable, as in Conditional 9b.

9a. If Sophie takes this drug, then she will be entirely cured.

9b. If the CEO admits the fraud, then he'll be sent to jail.

We can easily translate the definition of a positive consequential into the following utility grid:

$$\left\{\begin{array}{cccc} \text{If} & e & \bullet & \bullet \\ \text{Then} & \bullet & + & e \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

For a negative consequential, the + sign is replaced with a – sign. All the examples offered by Bonnefon and Hilton (2004) reflect various instantiations of this utility grid. Finally, Thompson et al. (2005) gave the following definition of *persuasion* and *dissuasion* conditionals:

> (a) *p* refers to a hypothetical action of a third party or agent, who is neither the speaker nor listener of the conditional and whose agency may be implied rather than stated explicitly; (b) the action *p* is as yet hypothetical, and is under the control of the third party; (c) *q* is a consequential outcome of taking action *p;* (d) *q* is valenced, representing either a desirable or undesirable outcome of action *p;* and (e) the consequences of this outcome do not accrue specifically to the listener, the speaker, or the agent of the conditional. (p. 240)

The statement is called a persuasion conditional when *q* is desirable, as in Conditional 10a below, and a dissuasion conditional when *q* is undesirable, as in Conditional 10b.

10a. If the railway line is closed, then the taxpayers will save millions of dollars.

10b. If the Kyoto accord is ratified, then there will be a severe downturn in the economy.

Again, we can easily translate the definition of a persuasion conditional into the following utility grid:

$$\left\{\begin{array}{cccc} \text{If} & e & \bullet & \bullet \\ \text{Then} & \bullet & + & A \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\}.$$

Thus far, I have shown how utility grids can capture previous characterizations of known subclasses of utility conditionals (e.g., threats, tips, consequentials). But the expressive power of utility grids goes far beyond this limited subset, and there resides part of their appeal. Consider, for example, the following grid:

$$\left\{\begin{array}{cccc} \text{If} & h & 0 & A \\ \text{Then} & \omega & + & s \\ & \text{Actor} & \text{Utility} & \text{Target} \end{array}\right\},$$

which reflects Conditional 11 below.

11. If you move just two inches, then I'll be able to see the screen.

This class of utility conditionals, which we may call requests, is probably as frequent as promises or warnings, but to my knowledge it has never been investigated empirically. This is an example of the way we can identify brand new subclasses of utility conditionals by studying the different permutations of the utility grid. Some of these permutations also allow for the nuancing of existing classifications. For example, we can express the missing link between threats and warnings with the following utility grid:

$$\left\{\begin{array}{cccc} \text{If} & h & - & s \\ \text{Then} & \omega & - & h \\ & \text{Actor} & \text{Utility} & \text{Target}\end{array}\right\},$$

which is reflected in Conditional 12 below.

12. If you denounce me to the police, then you will have an accident.

The classifications I have mentioned would consider Conditional 12 a warning rather than a threat, but anybody who has seen a Mafia movie knows better than that. Finally, some utility conditionals seem to entirely elude any labeling. Consider, for instance, this utility grid:

$$\left\{\begin{array}{cccc} \text{If} & e & - & e \\ \text{Then} & \omega & + & h \\ & \text{Actor} & \text{Utility} & \text{Target}\end{array}\right\},$$

illustrated by the following conditional:

13. If she gives up buying herself a new office computer on the department budget, then there will be funding for your conference trip.

I cannot think of any satisfying label for the speech act that this conditional is performing. The label, however, is irrelevant from the perspective of the utility grid framework. We need not study the properties of requests (Conditional 11), veiled threats (Conditional 12), or whatever we may call Conditional 13; what we need to study are the properties of their utility grids. We need a theoretical framework that predicts the paralogical inferences invited by utility grids, one that can explain why, upon reading Conditional 13, we feel it unlikely that the conference trip will be funded and why, upon reading Conditional 12, we feel it unlikely that the listener will denounce the speaker to the police. This is the topic of the next section.

Folk Axioms of Decision

In the previous section, I suggested that reasoners, when they process a conditional statement, include in their representation some information about what is at stake for whom in the situation described by the conditional. Although I did not make any hypothesis about the way this information is represented in the mind, I introduced the utility grid of a conditional as a theoretical tool reflecting the sort of information people retrieve.

Now that we are equipped with a theoretical tool to represent the decision-theoretic features of a conditional statement, it is time to address the question that drives this article: How do decision-making processes take advantage of the decision-theoretic features of a conditional to leak into reasoning processes?

Clearly, people have a folk theory about how other people make decisions. Now, this folk theory is quite probably not as complex as the choice theories developed by economists and psychologists. It may rather consist of a series of rules of thumb, of simple heuristics. For example, people are highly likely to entertain some folk version of consequentialism, according to which other people generally do what they think is good for themselves (Miller, 1999).

I suggest that when they process a conditional statement, people retrieve the information that can be represented in its utility grid, then apply to this information the folk axioms that form their folk theory of decision. It is the application of these folk axioms to this information that yields paralogical conclusions, as soon as the conditional is not purely disinterested.

I now go through several possible folk axioms of the folk theory of decision and consider in detail the utility grids from which they deliver paralogical inferences (see Table 2 for a summary). I do not claim my list of folk axioms to be exhaustive; rather, I believe that the exciting task before us is to formulate additional folk axioms of decision, predict the paralogical inferences that they should yield, and put these predictions to the test.

Table 2
Three Folk Axioms of Decision and the Utility Grids They Typically Apply To

Folk axiom of decision	Definition	Utility grid
Self-interested behavior	*People take actions that increase their own personal utility.* (And they do not take actions that decrease their own personal utility.)	$\begin{Bmatrix} x & + & x \\ \bullet & \bullet & \bullet \end{Bmatrix}$ $\begin{Bmatrix} x & \bullet & \bullet \\ \bullet & + & x \end{Bmatrix}$
Self-interested attitude	*People think that actions that increase their own personal utility should be taken by others when these other agents can take these actions.* (And they think that actions that decrease their own personal utility should not be taken by others when these other agents can take these actions.)	$\begin{Bmatrix} x & + & y \\ \bullet & \bullet & \bullet \end{Bmatrix}$ $\begin{Bmatrix} x & \bullet & \bullet \\ \bullet & + & y \end{Bmatrix}$
Limited altruism	*People take actions that increase the utility of others insofar as doing so does not decrease their own personal utility.* (And they do not take actions that decrease the utility of others insofar as these actions do not increase their own personal utility.)	$\begin{Bmatrix} x & + & y \\ \bullet & u & z \end{Bmatrix}$ $\begin{Bmatrix} x & u & z \\ \bullet & + & y \end{Bmatrix}$

Note. The + sign could indifferently be a – sign. x = an agent; + = any significantly positive value; ● = any legitimate value of the parameter; y = an agent; u = the utility of an action; z = an agent.

One remark about the language that these folk axioms ought to be couched in is also in order. It would be easy enough to give the folk axioms a standard Bayesian decision-theoretic formulation. Indeed, simple ideas like "people do what increases their utility" have immediate translations in formal decision-theoretic language. Although this is a feature of the theory that might prove useful in the future, in this article I do not attempt to translate the folk axioms into standard, formal decision-theoretic language. One immediate reason is that the framework I am offering consists of applying the folk axioms of decision to the utility grids, in order to predict utility-based inferences. As a consequence, the language of the folk axioms needs to be the same as that for the utility grids, and the folk axioms' Bayesian decision-theoretic formulation can only be construed as an extra feature of the framework, without an immediate use.

In the future, though, it might be relevant to give the folk axioms a fully formal translation, in a language that might or might not be that of standard Bayesian decision theory. There are indeed alternative (qualitative) formal languages for decision theory that might prove to be as appropriate as the standard quantitative language for psychological modeling. My second reason for not explicitly translating the folk axioms into classic decision-theoretic language in this article is twofold: (1) in order not to commit too early to one set of formal assumptions and (2) in order not to rule out the use of alternative qualitative formal translations.

The Folk Axiom of Self-Interested Behavior

The first folk axiom of decision states that people believe others to be, broadly speaking, consequentialist decision makers who do what achieves their goals and do not do what hurts their goals. This folk axiom is consistent with findings in social psychology (e.g., Ratner & Miller, 1998), demonstrating that laypersons predict the actions of other agents on the basis of these agents' perceived self-interest and that they even tend to overestimate the predictive power of self-interest when making these predictions.

> *Folk Axiom 1 (Self-Interested Behavior):* People take actions that increase their own personal utility, and they do not take actions that decrease their own personal utility.

Consider an "if p, then q" conditional with one of the following two utility grids (as always, the black dot stands for a blank parameter, which can take any legitimate value):

$$\begin{Bmatrix} x & + & x \\ \bullet & \bullet & \bullet \end{Bmatrix},$$

$$\begin{Bmatrix} x & \bullet & \bullet \\ \bullet & + & x \end{Bmatrix}.$$

The first grid expresses that p, the taking of some action by agent x, has positive proximal utility for this same agent x. The second grid expresses that p, the taking of some action by agent x, has positive distal utility for x. Applying the folk axiom of self-interested behavior, we conclude in both cases that x is going to take the action. Thus, both of these grids invite the inference that agent x is going to take the action—that is, they invite an inference to the truth of p. Mutatis mutandis, if the + sign were replaced with a – sign, then the grids would invite an inference to the falsity of p.

Note that agent x can be any agent in A except ω. For example, consider that $x = e$, that is, that agent x is someone who is neither the speaker nor the hearer of the conditional. Consider the following example:

14. If Alice takes the new job, then her life will improve in every respect.

Assume that taking the new job is an action made by Alice that, per se, has no known proximal utility for any agent and that improvement in Alice's life in every respect is an action of the neutral agent with positive utility for Alice, which can be represented by

$$\begin{Bmatrix} e & 0 & A \\ \omega & + & e \end{Bmatrix}.$$

The folk axiom of self-interested behavior, applied to this utility grid, predicts that Conditional 14 will invite an inference to the truth of its antecedent: Reasoners should conclude that Alice will take the new job. This is exactly the finding reported in Bonnefon and Hilton (2004). Conversely, reasoners presented with the conditional

15. If Alice takes the new job, then she will be paid less and be less happy

concluded that Alice would not take the new job. This is, again, exactly the prediction of the first axiom of folk decision theory, applied to this utility grid:

$$\begin{Bmatrix} e & 0 & A \\ \omega & - & e \end{Bmatrix}.$$

In another experiment, Bonnefon and Hilton (2004) compared responses to simple problems such as Conditional 16 below and augmented problems such as Conditional 17.

16a. If Emma needs to take the train, then Cédric drives her to the station.

16b. Emma needs to take the train.

17a. If Emma needs to take the train, then Cédric drives her to the station.

17b. If Cédric drives Emma to the station, then he will miss his job interview.

17c. Emma needs to take the train.

Bonnefon and Hilton (2004) observed that reasoners accepted the conclusion "Cédric drives Emma to the station" for the simple problem (16) but not for the augmented problem (17). Logically speaking, both problems entail this conclusion, since their formal structure is "if p, then q; p" and "if p, then q; if q, then r; p," respectively. However, the two conditionals in the augmented problem do not have the same utility grid.

Let us consider the conditional featured in 16a and 17a. Emma needing to take the train is not an action under the control of Emma but an action of the neutral agent that has no known significant utility for anyone; driving Emma to the station is an action by

Cédric that has positive utility for Emma. This leads to the following utility grid:

$$\begin{Bmatrix} \omega & 0 & A \\ e' & + & e \end{Bmatrix}.$$

The folk axiom of self-interested behavior does not apply to this utility grid. Thus, the conditional in 16a and 17a does not trigger extralogical, decision-based inferences. Things are different with Conditional 17b. Driving Emma to the station is an action by Cédric that has positive utility for Emma, but missing the job interview is an action by the neutral agent that has negative utility for Cédric, as represented in the following:

$$\begin{Bmatrix} e' & + & e \\ \omega & - & e' \end{Bmatrix}.$$

From this utility grid, the folk axiom of self-interested behavior invites the inference that Cédric (a.k.a. agent e') is not going to drive Emma to the station—hence reasoners' reluctance to accept that conclusion, despite the fact that it is a legitimate, logical conclusion from the conditional (17b) and the categorical premise (17c).

We have considered so far examples where the folk axiom of self-interested behavior applies to the goals of an agent e who is neither the speaker nor the hearer of the conditional statement. The first principle, however, can apply to the goals of these special agents. Consider the following examples:

18a. If you mow the lawn, then I will pay you $5.

18b. If you stand by the pillar, then you'll be served immediately.

Conditional 18a has the characteristic utility grid of a promise, where p is an action of the hearer that has positive utility for the speaker and q is an action of the speaker that has positive utility for the hearer, as represented by

$$\begin{Bmatrix} h & + & s \\ s & + & h \end{Bmatrix}.$$

Similarly, Conditional 18b has the characteristic utility grid of a tip, where p is an action of the hearer that has no known utility to anyone and q has positive utility for the hearer. Note that q in Conditional 18b is an action of an agent e (a waiter), which is not necessarily true of all tips (for some tips, q is an action of the neutral agent), as represented by

$$\begin{Bmatrix} h & 0 & A \\ e & + & h \end{Bmatrix}.$$

The folk axiom of self-interested behavior applies to both grids and predicts that both conditionals 18a and 18b invite the inference that agent h will make p true, by mowing the lawn and by standing next to the pillar, respectively. Support for that prediction is provided in Ohm and Thompson (2004, Experiment 1). Reasoners were presented with conditionals such as 18a and 18b and had to assess what the probability of p would be if the conditional were asserted and what it would be if the conditional were not asserted. The assertion of the conditional increased this probability by about .30.

The results of Evans et al. (2008) offered independent support for the folk axiom of self-interested behavior. Reasoners were presented with tips and promises wherein the utility of q to h was either high (e.g., having a healthy pregnancy) or low (e.g., saving 5 pence a week) and had to rate on a 5-point scale the likelihood that h would undertake action p. Under appropriate conditions (i.e., p is not itself overly costly to h, and "if p, then q" is a reasonably credible prediction), the folk axiom of self-interest predicts that a high-utility q (as compared with a low-utility q) strongly invites the conclusion that h will undertake p. Reasoners rated, in line with this prediction, as very likely (about 4.8) the likelihood that h would undertake p when q had high utility and expressed their ignorance (with a rating of about 3.3) as to whether h would undertake p when q had low utility.

The folk axiom of self-interest also applies to the negative utility counterparts of promises and tips (i.e., threats and warnings), as represented by the following:

19a. If you tell the boss I was late, then I will tell her about your affair with the intern.

19b. If you click on that link, then your computer will be infected by a virus.

Conditional 19a has the characteristic utility grid of a threat, as seen in

$$\begin{Bmatrix} h & - & s \\ s & - & h \end{Bmatrix},$$

and Conditional 19b has the characteristic utility grid of a warning (note that q is an action of the neutral agent but could also be an action of some agent e), as seen in

$$\begin{Bmatrix} h & 0 & A \\ \omega & - & h \end{Bmatrix}.$$

The folk axiom of self-interested behavior, applied to these utility grids, predicts that conditionals 19a and 19b will invite the inferences that the hearer is not going to tell the boss and not going to click on the link, respectively. Again, support for this prediction is given in Ohm and Thompson (2004, Experiment 1). Reasoners were presented with conditionals such as 19a and 19b and had to assess what the probability of p would be if the conditional were asserted and what it would be if the conditional were not asserted. The assertion of the conditional decreased this probability by about .20.

Evans et al. (2008) also presented reasoners with threats and warnings in which the (negative) utility of q for h was either high (e.g., crashing a friend's brand new Porsche) or low (e.g., waiting 30 s longer to get served). Under appropriate conditions (i.e., p is not itself very beneficial to h and "if p, then q" is a reasonably credible prediction), the folk axiom of self-interested behavior predicts that a high-cost q (as compared with a low-cost q) will strongly invite the conclusion that h will not undertake p. Reasoners rated, in line with this prediction, as rather unlikely (about 2.2) the likelihood that h would undertake p when q had a high cost and expressed their ignorance (with a rating of about 3.1) as to whether h would undertake p when q had a low cost.

The folk axiom of self-interested behavior, conjoined to the use of utility grids, thus accounts for a significant range of prior results related to threats, promises, tips, warnings, and consequential conditionals. Because the scope of the folk axiom is much broader than just these few categories of utility conditionals, it also allows for the derivation of a vast number of new predictions. The two grids that it typically applies to (as featured in Table 2) allow for 428 distinct, legitimate combinations of parameters from which the folk axiom of self-interest makes a straightforward prediction. That is, there are 428 different well-formed utility grids (i.e., if u or u' is zero, then y or y', respectively, is A), from which the folk axiom invites consistent conclusions (more on this later; in this case, it means that x must be different from x' as soon as u and u' have opposite signs). Most of these grids have never been considered in previous theoretical or empirical work. The initial definitions of threats, promises, tips, warnings, and consequential conditionals that I have considered in this section potentially account for 136 of the 428 combinations, and previous empirical work has been limited to only a subset of these 136 combinations. For example, although the definition of consequential conditionals alone allows for 94 combinations, the Bonnefon and Hilton (2004) article used only a dozen different statements as experimental material.

The examples I have considered so far are simple in the sense that, for all of them, the folk axiom of self-interested behavior invites only one conclusion. More complex situations are possible, however, wherein this folk axiom yields conflicting conclusions from a single utility grid, or conflicting conclusions from several utility subgrids of a single conditional. I examine these situations in a later section, after introducing additional tentative folk axioms of decision.

The Folk Axiom of Self-Interested Attitude

The folk axiom of self-interested behavior addressed situations wherein agents had the capacity to take action relevant to their goals. In other situations, though, agents may be aware that someone else has the capacity to take an action that can help them or hurt them. The folk axiom of self-interested attitude addresses these situations. It states that agents think that actions that achieve their goals ought to be taken by the agent who can take them. Again, this folk axiom is consistent with findings in social psychology (Miller, 1999; Ratner & Miller, 1998) that demonstrate that laypersons predict the attitudes of other agents on the basis of their perceived self-interest and that they even tend to overestimate the predictive power of self-interest when making these predictions.

> *Folk Axiom 2 (Self-Interested Attitude):* People think that actions that increase their own personal utility should be taken by others, when these other agents can take these actions; and they think that actions that decrease their own personal utility should not be taken by others, when these other agents can take these actions.

Consider an "if p, then q" conditional with one of the following utility grids, where $x \neq y$ and where the black dot stands for any legitimate value of the appropriate parameter:

$$\begin{Bmatrix} x & \bullet & \bullet \\ \bullet & + & y \end{Bmatrix},$$

$$\begin{Bmatrix} x & + & y \\ \bullet & \bullet & \bullet \end{Bmatrix}.$$

The folk axiom of self-interested attitude predicts that conditionals with one of these utility grids will invite the inference that y thinks that x should take action p. Mutatis mutandis, when one replaces the + sign with a – sign, then the folk axiom of self-interested attitude predicts that agent y will think that x should not take action p.

Thompson et al. (2005) investigated specific examples of the kind of conditionals that the second principle may apply to, such as

20a. If the Kyoto accord is ratified, then greenhouse gas emissions will be reduced.

20b. If the Kyoto accord is ratified, then there will be a downturn in the economy.

Ratifying the Kyoto Protocol is unlikely to be an action under the control of the speaker or the hearer of Conditionals 20a and 20b, but it is an action under the control of some other agent e. The mere action of ratifying the protocol does not have any known proximal utility. However, this action will have distal consequences in the form of an action of the neutral agent ω. Not only does Conditional 20a frame the consequence as positive, but it arguably makes the consequence beneficial to all agents in A, as represented in

$$\begin{Bmatrix} e & 0 & A \\ \omega & + & A \end{Bmatrix}.$$

Thus, from the folk axiom of self-interested attitude, Conditional 20a invites the inference that all agents, as long as they believe 20a to be correct, will endorse the deontic stance that p should be taken (i.e., that the protocol should be ratified). Now, the speaker of the conditional, in particular, can be safely assumed to believe that her conditional is correct. The folk axiom of self-interested attitude thus predicts that Conditional 20a will invite the inference that the speaker believes that p should be taken. Mutatis mutandis, this same folk axiom predicts that Conditional 20b will invite the inference that the speaker believes that p should not be taken.

This is indeed what Thompson et al. (2005) observed when they presented reasoners with conditionals such as 20a and 20b and asked them to rate the extent to which the speaker believed that p should be taken on a 7-point scale ranging from –3 (*p should be taken*) to +3 (*p should not be taken*). Reasoners unambiguously judged that the speaker of 20a-like conditionals wanted p to be taken (with an average rating of –2.2) and that the speaker of 20b-like conditionals did not want p to be taken (with an average rating of +2.2).

Thompson et al. (2005) also reported that whether reasoners believed themselves that p should be taken was not much affected by whether they were presented with 20a- or 20b-like conditionals. Note that the folk axiom of self-interested attitude does not make any prediction on that issue. The folk axiom accounts for how reasoners *predict* the attitude of other people from conditionals that bear on these other people's utility functions; it does not account for the way reasoners form their own attitudes on the basis of the same conditionals.

The folk axiom of self-interested attitude accounts for prior results related to persuasion and dissuasion conditionals but also allows for a vast number of other predictions. Indeed, just as with the folk axiom of self-interested behavior, the folk axiom of self-interested attitude applies to hundreds of utility grids for which no predictions have been available so far.

The Folk Axiom of Limited Altruism

The folk axioms of self-interested behavior and of self-interested attitude assume that people care only about their own goals and disregard the consequences that any action may have for others. The folk axiom of limited altruism tentatively injects a dose of benevolence into the folk theory of decision. It states that people do what achieves the goals of others as long as it does not hurt their own goals and, mutatis mutandis, that people avoid actions that hurt the goals of others, as long as those actions would not achieve their own goals.

> *Folk Axiom 3 (Limited Altruism):* People take actions that increase the utility of others insofar as doing so does not decrease their own personal utility, and they do not take actions that decrease the utility of others insofar as these actions do not increase their own personal utility.

Consider an "if p, then q" conditional with one of the following utility grids:

$$\begin{Bmatrix} x & + & y \\ \bullet & u & z \end{Bmatrix},$$

$$\begin{Bmatrix} x & u & z \\ \bullet & + & y \end{Bmatrix}.$$

The folk axiom of limited altruism predicts that such a conditional invites an inference to the truth of p, as long as $(u, z) \neq (-, x)$. Mutatis mutandis, if the $+$ sign in the grid is replaced by a – sign, then the folk axiom of limited altruism predicts that the conditional invites an inference to the falsity of p, as long as $(u, z) \neq (+, x)$.

Consider again Conditional 11, "If you move just two inches, then I'll be able to see the screen." Moving 2 in. is an action of the hearer that, presumably, has no proximal utility for anyone (including the hearer). Being able to see the screen is a state of affairs (i.e., an action of the neutral agent) that has positive utility for the speaker. The utility grid of the conditional is thus

$$\begin{Bmatrix} h & 0 & A \\ \omega & + & s \end{Bmatrix}.$$

Applied to this utility grid, the folk axiom of limited altruism predicts that the conditional "If you move just two inches, then I'll be able to see the screen" invites the inference that the hearer will move 2 in. Consider another example, featuring a negative distal utility. Louis, who has never been interested in ecological or humanitarian issues, has recently seen a documentary and has come to the conclusion that

21. If I eat Nile perch fish, then I'll make life harder in some African countries.

Louis has no personal stake in that consequence. In addition, Louis neither likes nor dislikes eating Nile perch. The utility grid of the conditional is thus

$$\begin{Bmatrix} s & 0 & A \\ \omega & - & e \end{Bmatrix}.$$

Applied to this utility grid, the folk axiom of limited altruism predicts that Conditional 21 will invite the inference that Louis is not going to eat Nile perch anymore.

The extant evidence of inferences invited by the folk axiom of limited altruism is indirect only. Vartanian, Mandel, and Blackler (2008) presented reasoners with conditionals describing government policies such as

22a. If policy A is adopted, then spending on social programs will increase.

22b. If policy B is adopted, then spending on defense programs will increase.

Participants then rated the probability that these policies would be implemented, as well as the value they saw in the outcome of each policy. Results showed that the perceived probability that the policy would be implemented was a positive function of the value that individual participants attached to the outcome of the policy. This finding is consistent with the folk axiom of limited altruism, provided that we construe the value attached to the outcome of a policy as the distal utility of the policy for the whole set of agents. Assuming that an individual participant thinks approvingly of spending on social programs, then Conditional 22a would be represented by the following utility grid:

$$\begin{Bmatrix} e & 0 & A \\ \omega & + & A \end{Bmatrix},$$

where e is the government. From this utility grid, the folk axiom of limited altruism invites the inference that the policy will be implemented. This evidence, however, is indirect at best (if only because the folk axiom of self-interest already invites the same inference from the same grid). Just as with the other folk axioms, the folk axiom of limited altruism applies to hundreds of utility grids for which no predictions have been available so far. Because no evidence was currently available for the most basic of these predictions, I conducted the following experiment.

Experiment: Limited Altruism

Method

Sixty participants read three conditional reasoning problems. All problems had the following same simple context (the names of the characters changed in each problem):

> [One character *X*] and [another character *Y*] are employed by the same company, at the same level of responsibility, but in different sectors. [Character *X*] tells [character *Y*]: . . .

This contextual information was followed by one of these three conditionals:

23a. If you walk by my office, then you'll see a plant.

23b. If you walk by my office, then that would be good for me.

23c. If you decline the promotion you have been offered, then that would be good for me.

In each case, participants rated the likelihood that the antecedent action would be undertaken on a 5-point scale ranging from –2 (*certainly not going to happen*) to +2 (*certainly going to happen*). The folk axiom of limited altruism predicts that this rating will be close to zero for 23a, greater than zero for 23b, and lower than zero for 23c. These predictions can be checked from the following utility grids of the three conditionals, respectively:

$$\begin{Bmatrix} h & 0 & A \\ \omega & 0 & A \end{Bmatrix}, \begin{Bmatrix} h & 0 & A \\ \omega & + & s \end{Bmatrix}, \begin{Bmatrix} h & - & h \\ \omega & + & s \end{Bmatrix}.$$

None of the folk axioms apply to the disinterested conditional (23a), which should not invite any utility-based conclusion. The folk axiom of limited altruism applies to 23b, inviting the conclusion that the antecedent action is going to be undertaken. Finally, 23c is an example in which the folk axiom of self-interested behavior will override the folk axiom of limited altruism and should invite the conclusion that the antecedent action is not going to be undertaken, regardless of its benefits to another agent.

Results

Participants' ratings agreed with the predictions. The average rating was 0.0 (SD = 0.7) for 23a, +1.1 (SD = 0.9) for 23b, and –0.6 (SD = 0.9) for 23c. A repeated-measure analysis of variance with the type of grid as a three-level factor revealed a large overall effect, $F(2, 58) = 51.4$, $p < .001$, $\eta_p^2 = .51$. The follow-up contrast analysis confirmed that the rating associated to 23b was greater than the rating associated to 23a, $F(1, 59) = 101.9$, $p < .001$, $\eta_p^2 = .63$, and that the rating associated to 23c was lower than the rating associated to 23a, $F(1, 59) = 17.7$, $p < .001$, $\eta_p^2 = .23$.

Multiple Conclusions From a Single Utility Conditional

So far, I have illustrated only the three folk axioms with simple examples, which yielded only one invited conclusion. However, in many cases, multiple conclusions can be invited from a given conditional: (a) One folk axiom may invite several conclusions from a single utility grid; (b) several folk axioms may apply to a single utility grid; (c) one folk axiom may invite different conclusions from the different subgrids of a conditional; and (d) several folk axioms may apply to the different subgrids of a conditional.

Furthermore, two situations may arise when different conclusions are invited by the application of one or several folk axioms to one or several subgrids. The various conclusions may be consistent (in which case they are likely to simply coexist), or some conclusions may be inconsistent with each other. In this section, I begin with the simpler situation of nonconflicting multiple conclusions. Then I turn to the complex situations of multiple conflicting conclusions and the refinements required to arbitrate these cases.

Nonconflicting Conclusions

One grid. In some situations, a conclusion can be invited twice from the same grid, because a given folk axiom applies to two different aspects of the grid. Consider, for example, the following grid:

$$\begin{Bmatrix} x & + & x \\ \bullet & + & x \end{Bmatrix}.$$

The folk axiom of self-interest doubly applies to this grid, since action p has both positive proximal utility and positive distal utility for agent x. Such a grid should strongly invite conclusion p, more so in any case than a grid where action p would have only proximal or distal utility for agent x. For example, the conditional

24. If Alain goes partying with his dear friends, then he will meet up with the woman he loves

should strongly invite the conclusion that Alain will go partying. Mutatis mutandis, if the + signs in the grid are replaced with – signs, then the conditional should strongly invite the conclusion $\neg p$.

Similarly, the folk axiom of self-interested attitude doubly applies to the following grid and doubly invites the conclusion that agent y will adopt the deontic stance that action p should be taken, presumably making this invitation stronger:

$$\begin{Bmatrix} x & + & y \\ \bullet & + & y \end{Bmatrix}.$$

Note that the folk axiom of limited altruism also doubly applies to this grid, thus strongly inviting the conclusion that agent x is going to take action p, since this action doubly benefits agent y without any negative utility to agent x. For example, the conditional

25. If I call my mother on her birthday, then all her friends will compliment her about how good a son I am

should strongly invite the conclusion that the speaker is going to call his mother on her birthday. In parallel, from the perspective of the folk axiom of self-interested attitude, Conditional 25 invites the conclusion that the speaker's mother is likely to insist that he call her on her birthday. A generalization of this situation occurs when the grid above features two different agents, y and y', as represented by

$$\begin{Bmatrix} x & + & y \\ \bullet & + & y' \end{Bmatrix}.$$

The folk axiom of limited altruism predicts that the conditional strongly invites the conclusion p, since this action of agent x benefits two different agents (proximally for y, distally for y'). In parallel, the folk axiom of self-interested attitude applies twice to this grid, yielding two separate yet nonconflicting conclusions: that y thinks p should be taken and that y' thinks p should be taken. Consider, for example, the conditional

26. If I call my grandmother on her birthday, then my mother's friends will compliment her about how well she raised me.

The folk axiom of limited altruism predicts that Conditional 26 will strongly invite the conclusion that the speaker is going to call her grandmother, for this action will benefit both her grandmother (proximally) and her mother (distally). The folk axiom of deontic self-interest predicts that Conditional 26 will invite the conclusions that the mother and the grandmother both think the call should be made, albeit for different motives.

In other situations, the same conclusion can be invited by the application of different folk axioms to the same grid. Consider, for example, the following grid:

$$\begin{Bmatrix} x & + & x \\ \bullet & + & y \end{Bmatrix}.$$

The taking of action p by agent x has positive proximal utility for x and positive distal utility for another agent y. Therefore, both the folk axiom of self-interested behavior and that of limited altruism predict that the conclusion p will be invited—and as a consequence, we may expect this invitation to be stronger than if it were invited by one folk axiom only. For example, the conditional

27. If you wear your favorite suit, then I'll be proud to be seen with you

should strongly invite the conclusion that the hearer will wear his favorite suit—not only does this action intrinsically please him (self-interest), but it also has positive utility for the speaker (limited altruism). Note that, in parallel, the folk axiom of self-interested attitude applies to the grid above and predicts that Conditional 27 will invite the conclusion that the speaker thinks the suit should be worn.

Two subgrids. A single action can bear on the goals of several agents. When an action has nonnull proximal utility for several agents, or when it has nonnull distal utility for several agents, then these different aspects of the utilities at stake are represented in different utility subgrids of the conditional. Sometimes, the same folk axiom will apply to the different subgrids, yielding each time a similar conclusion. Consider, for example, the conditional

28. If she invites both Pierre and Luc, then they will be upset.

Assuming for simplicity's sake that the mere action of inviting both Pierre and Luc has no proximal utility to any agent and that this action has no significant distal utility to the female agent, then Conditional 28 receives the following complex utility grid, consisting of two subgrids:

$$\begin{Bmatrix} e & 0 & A \\ \omega & - & e' \end{Bmatrix} \begin{Bmatrix} e & 0 & A \\ \omega & - & e'' \end{Bmatrix}.$$

The folk axiom of limited altruism applied to the left subgrid predicts that she will not invite both Pierre and Luc—because this would upset Pierre. Applied to the right subgrid, this same folk axiom predicts again that she will not invite both Pierre and Luc—because this would upset Luc. Thus, the conclusion that she will not invite both Pierre and Luc should be strongly invited by Conditional 28.

It can also be the case that the same folk axiom applies to different subgrids, yielding different yet nonconflicting conclusions. Consider, for example, that Alice is indifferent about whether she takes her annual vacation in July or August. Bastien, her coworker, tells her that

29. If you take your vacation in July, then both I and Cécile will have to stay at the office at that time and won't be able to spend our vacation with our families.

This conditional has the following complex utility grid:

$$\begin{Bmatrix} h & 0 & A \\ \omega & - & s \end{Bmatrix} \begin{Bmatrix} h & 0 & A \\ \omega & - & e \end{Bmatrix}.$$

The folk axiom of self-interested attitude applied to the left subgrid predicts that Conditional 29 will invite the conclusion that the speaker (Bastien) objects to the taking of action p (Alice taking her vacation in July). Applied to the right subgrid, this same folk axiom predicts that Conditional 29 will invite the conclusion that another agent (Cécile) also objects to the taking of action p. Furthermore, as in Conditional 28, the folk axiom of limited altruism doubly predicts that 29 will invite the conclusion that the hearer (Alice) will not take action p, because it has negative utility for two other agents but no positive utility for herself.

Finally, different folk axioms may apply to different subgrids, yielding different yet nonconflicting conclusions. Consider the following example, in which the owner of a factory is considering a strategic move:

30. If I relocate the factory, then I will make greater profits, all the current workers will lose their jobs, and they will be replaced by workers from the new location.

This conditional can be represented by a complex utility grid consisting of these three subgrids:

$$\begin{Bmatrix} s & 0 & A \\ \omega & + & s \end{Bmatrix} \begin{Bmatrix} s & 0 & A \\ \omega & - & e \end{Bmatrix} \begin{Bmatrix} s & 0 & A \\ \omega & + & e' \end{Bmatrix}.$$

The folk axiom of self-interested behavior, applied to the left subgrid, predicts that Conditional 30 will invite the conclusion that the speaker is going to relocate the factory. This folk axiom does not directly apply to the other subgrids, but, as per the definition of the folk axiom of limited altruism, it prevents any use of limited altruism that would invite the opposite conclusion. In other terms, the fact that action p (relocating the factory) has negative utility (central subgrid) or positive utility (right subgrid) for other agents is irrelevant as soon as it has positive utility for the agent himself.

The folk axiom of self-interested attitude, however, applies to both the central and the right subgrids. It predicts that the current workers will endorse the deontic stance that p should not be undertaken and that the future, potential workers, if they know about it at all, will endorse the deontic stance that action p should be undertaken.

Conflicting Conclusions

Conditional 30 serves as a reminder of a built-in arbitration rule between the folk axiom of self-interested behavior and that of limited altruism: The folk axiom of self-interested behavior is assumed to systematically override the folk axiom of limited altruism when they suggest conflicting conclusions. This "charity begins at home" rule was already explicit in the definition of the folk axiom of limited altruism. However, other conflicts can occur that cannot readily be solved.

All these conflicts are the result of a single folk axiom suggesting contradictory conclusions from a single grid or suggesting contradictory conclusions from separate subgrids. This is a structural consequence of the starting list of folk axioms that I am considering in this article: (a) No conflict can derive from self-interested behavior and limited altruism, because of the built-in arbitration rule that charity begins at home and (b) no conflict can derive from self-interested attitude and the other folk axioms, because the conclusions it invites are deontic, whereas the conclusions of the two other folk axioms are factual.

Conflicts based on the folk axiom of self-interested behavior. This folk axiom can easily generate conflicts from a single grid as soon as the proximal and distal utilities of action p have different signs for the focal agent of the conditional, as in either of the following two grids:

$$\left\{\begin{matrix} x & - & x \\ \bullet & + & x \end{matrix}\right\},$$

$$\left\{\begin{matrix} x & + & x \\ \bullet & - & x \end{matrix}\right\}.$$

These grids denote a trade-off between a proximal loss and a distal benefit (see Conditional 31a below) or between a proximal benefit and a distal loss (see 31b).

31a. If I stay home and study tonight instead of partying, then I'll be ready for the test tomorrow.

31b. If you have an affair with her, then it will destroy your marriage.

Consider the following utility grid for Conditional 31a:

$$\left\{\begin{matrix} s & - & s \\ \omega & + & s \end{matrix}\right\}.$$

Staying home and studying, rather than partying, is (presumably) an action that has negative utility for the speaker. The folk axiom of self-interested behavior thus predicts that Conditional 31a will invite the conclusion that the speaker is not going to stay home. However, being ready for the test is (again, presumably) a consequence that has positive utility for the speaker. The folk axiom of self-interested behavior thus predicts that 31a will invite the conclusion that the speaker will stay home.

It is tempting to consider that the final decision of the agent depends on a comparison between the positive utility of being ready for the test and the negative utility of staying home. Rational agents, indeed, are expected to take the action that maximizes their net utility. However, the coarse granularity that I have adopted to describe utility does not enable such a comparison. We may thus want to let utility vary on a refined scale, allowing for more precise comparisons between the various utilities at stake.

I say more on this issue in the Extensions and Future Directions section. Note, though, that a coarse granularity seems appropriate for simple conditional reasoning problems, wherein the personal preferences of the agents are usually poorly defined. In the absence of any specific information, it is anybody's guess whether the agent in 31a prefers partying to passing the test or whether the agent in 31b values his marriage more than his affair. It would seem especially ludicrous and arbitrary to assign a precise quantitative utility to these consequences for this agent, say +28 utility points for the affair and –76 utility points for destroying the marriage. Because there is no obvious answer to the question of what the agent will value more, it is reasonable to accept that the inconsistent conclusions simply cancel out and that conditionals such as 31a and 31b do not invite, in fine, any utility-based conclusion.

This simple analysis seems broadly in line with results obtained by Evans et al. (2008), who presented reasoners with inducements and advice featuring a significant intrapersonal conflict of utility for the hearer, for example, warnings that a high-cost consequence q would follow a high-utility action p, as in Conditional 32a below, or promises that a high-benefit consequence q would follow a high-cost action p, as in Conditional 32b.

32a. If you go and see your favorite band tonight, then you'll fail the exam.

32b. If you stop smoking, then I'll take you to Paris for the weekend.

Participants rated on a 5-point scale the likelihood that the hearer would undertake action p. Overall, these ratings did not deviate much from the midpoint of the scale, suggesting that reasoners were not willing to commit to a prediction as to how the agent would arbitrate his or her intrapersonal conflict of utility.

Conflicts based on the folk axiom of self-interested attitude. The folk axiom of self-interested attitude can make conflicting predictions both from a single grid and from separate subgrids of the same conditional. A single grid may feature an action p by an agent x whose proximal and distal utilities for another agent y are of different polarities, as in either of the following two grids:

$$\left\{\begin{matrix} x & - & y \\ \bullet & + & y \end{matrix}\right\},$$

$$\left\{\begin{matrix} x & + & y \\ \bullet & - & y \end{matrix}\right\}.$$

Conditional 33 below is an example of the first grid above.

33. If the coach trains the players much harder, then the team will win the high school championship.

The folk axiom of self-interested attitude predicts that Conditional 33 will invite the conclusion that the players will encourage harder training by the coach (because they presumably want to win the championship) but also that the players will discourage harder training by the coach (because harder training is, after all, harder on them).

Just as in the case of conflicts based on the folk axiom of self-interested behavior, we may be happy with considering that in such a situation, and in the absence of any specific information about the specific preferences of the team, the two invited conclusions just cancel out, eventually yielding no invited conclusion about whether the players think p or $\neg p$ should be taken.

The second principle can also lead to conflicting predictions when the proximal or distal utility of a given action has two different aspects of different polarities for the same agents, as in the following example:

34. If the new regulation is passed, then employees will work more hours but be paid more.

Conditional 34 is represented by these two subgrids:

$$\left\{\begin{matrix} e & 0 & A \\ \omega & - & e' \end{matrix}\right\} \left\{\begin{matrix} e & 0 & A \\ \omega & + & e' \end{matrix}\right\}.$$

Applied to the left-hand grid, the folk axiom of self-interested attitude predicts that Conditional 34 will invite the conclusion that the employees think p should be taken. However, applied to the right-hand grid, the same folk axiom predicts that 34 will invite the opposite conclusion. Just as in the previous example, and in the absence of any specific information about the specific preferences of the employees, we may be happy to consider that the two invited conclusions just cancel out, eventually yielding no invited conclusion about whether the employees think p or $\neg p$ should be taken.

Conflicts based on the folk axiom of limited altruism. The folk axiom of limited altruism can make conflicting predictions both from a single grid and from separate subgrids of the same conditional. A single grid may feature an action p by an agent x whose proximal and distal utilities are of different polarities and concern different agents, as in either of the following two grids:

$$\begin{Bmatrix} x & + & y \\ \bullet & - & y' \end{Bmatrix},$$

$$\begin{Bmatrix} x & - & y \\ \bullet & + & y' \end{Bmatrix}.$$

Conditional 35 below is an example of the first grid above.

35. If I give Simon the best seat, then I will upset Theo.

The folk axiom of limited altruism predicts that Conditional 35 will invite the conclusion that the speaker will give Simon the best seat (for it would probably please Simon without harming the speaker) but also that the speaker will not give the best seat to Simon (for it would upset Theo without benefiting the speaker).

Similarly, the folk axiom of limited altruism can make conflicting predictions when the proximal or distal utility of a given action has different aspects of different polarities for different agents, as in the following example:

36. If we take the kids to see the snakes at the vivarium, then Nina will enjoy seeing her favorite animals, but Olivier will have nightmares.

Whereas the conflicts based on the two self-interest folk axioms derive from intrapersonal conflicts of utility, conflicts based on the folk axiom of limited altruism derive from interpersonal conflicts of utility. Because interpersonal conflicts of utility are involved in a wide diversity of hard problems (from moral dilemmas to social justice), conflicts based on the folk axiom of limited altruism are likely to require specific refinements. One attractive possibility is to assume that the folk theory of decision includes some variant of the *primum non nocere* (above all, do no harm) principle; that is, that the negative utility incurred by an agent y as a consequence of a possible action by x trumps the positive utility of this same action for another agent y'. With this refinement, we can predict that Conditional 36, or any conditional with a comparable utility grid, will invite the inference that the antecedent action will not be taken. This would capture the common intuition that a parent asserting 36 would lean toward not taking the kids to the vivarium.

Extensions and Future Directions

The theory of utility conditionals that I offer in this article is complex enough to encompass most previous research about special conditional reasoning (based on inducements, persuasion, or otherwise consequential conditionals) and to account for most empirical findings in this literature. It is also simple enough to allow for straightforward new predictions about which paralogical inference is invited by which utility grid. This simplicity nevertheless leaves room for extensions and refinements of the theory. In this section, I consider three natural extensions of the theory that would endow it with more utility degrees, more probability degrees, and more folk axioms.

More Utility Degrees

So far, the utility grid framework uses only a coarse measure of utility. Utility can be significantly positive (+), significantly negative (–), or none of the above (0). As a consequence, the framework deals rather summarily with conflicts of utilities such as one posed by the following grid:

$$\begin{Bmatrix} h & + & h \\ \omega & - & h \end{Bmatrix},$$

of which warning Conditional 37 below is an illustration.

37. If you go on a romantic date tonight, then you will fail the exam tomorrow.

I have suggested that in such a case, the conflicting conclusions invited by the folk axiom of self-interested behavior simply cancel out, leaving reasoners undecided as to whether the hearer will or will not go out on a date. In the absence of any specific information about how important the exam is, or how infatuated the hearer is, this seems a reasonable solution. It is not difficult, though, to find conditionals that make this solution counterintuitive. Consider the following, for example:

38. If you eat this delicious fugu fish, then you'll die twenty minutes later.

Clearly, Conditional 38 invites the conclusion that the hearer is not going to eat the fish. However, the utility grid of 38 is the same as that for 37. Thus, the utility grid framework would currently predict that 38 will not invite any inference as to whether the hearer will or will not eat the fish.

One appropriate solution to account for conditionals such as 38 may be to use a scale of intermediate granularity, which would capture obvious differences in utility without committing to any numerical framework. For example, utility may take its value from {– –, –, 0, +, + +}, where + denotes a utility value that is significantly positive but one order of magnitude smaller than + + (Bonnefon, Dubois, Fargier, & Leblois, 2008; Dubois, Fargier, & Bonnefon, 2008). Then a conditional such as 38 could receive the following utility grid:

$$\begin{Bmatrix} h & + & h \\ \omega & -\,- & h \end{Bmatrix}.$$

The folk axiom of self-interested behavior still applies to this grid and still predicts conflicting invited conclusions. However, the negative distal utility of dying clearly outweighs the positive proximal utility of eating delicious fish, and the invited conclusion from distal utility accordingly overrides the invited conclusion from proximal utility.

More Probability Degrees

In their current form, utility grids miss out on one decision-theoretic aspect of utility conditionals: that of their probability. Not only does the probability of conditionals play an important role in reasoning (Evans, Handley, & Over, 2003; Evans & Over, 2004; Oaksford & Chater, 2001, 2007; Oberauer & Wilhelm, 2003; Over, Hadjichristidis, Evans, Handley, & Sloman, 2007), but it is also instrumental to decision making, because it allows for consideration of the expected utility of courses of action rather than their raw utility. Consider, for example, a child being told the following:

39. If you eat candy at school, then we will return all your Christmas presents to the store.

The huge gap between the modest proximal utility of eating candy and the momentous distal cost of foregoing all Christmas presents should make Conditional 39 an excellent deterrent. A smart child, however, might realize that the probability that her parents will act upon their disproportionate threat is really quite low (Verbrugge et al., 2004). In other words, a smart child might realize that the expected cost of foregoing Christmas presents is low and maybe not on par with the sure gain of eating candy at school. Accordingly, this child may well decide to call her parents' bluff and eat candy at school.

So far, the utility grid framework cannot model this complex inference about what the child is going to do. Let us assume that we allow utility to take its values from the $\{--, -, 0, +, ++\}$, ordinal, bipolar scale. Conditional 39 can then receive the following utility grid:

$$\begin{Bmatrix} h & + & h \\ s & -- & h \end{Bmatrix}.$$

Applied to this grid, the folk axiom of self-interested behavior straightforwardly invites the inference that the child is not going to eat candy. This might be judged an unsatisfactory result: We would want to predict that 39 will not invite this inference, because 39 itself has a low subjective probability.

More generally, we would want to capture the fact that inferences based on the utility u' of q are affected by the subjective probability of the utility conditional "if p, then q" (Evans et al., 2008). For many utility conditionals (e.g., inducements and pieces of advice), this subjective probability will itself be affected by the control that the speaker has on the consequent event and its associated utility (Evans & Twyman-Musgrove, 1998; Newstead et al., 1997).

There are several solutions we may consider in order to factor in the subjective probability of utility conditionals. One option is to replace distal utility with expected distal utility in the grid (proximal utility is not concerned because it is attached to p, not to q). This would be a solution of choice if utility grids used a full-blown, numeric utility measure. It would be quite simple, then, to decide that the grid should not display the raw distal utility u' but rather the product $u' \times \Pr(q|p)$. Such a fine-grained scale of utility, however, is descriptively implausible. It seems reasonable and realistic to assume that we have only a gross perception of the utility function of other agents, best described by a qualitative scale consisting of a handful of values. The problem, then, is to aggregate this qualitative measure of utility with the uncertainty measure, which may or may not be itself a full-blown probability measure (see Amgoud, Bonnefon, & Prade, 2005, for an example and Fargier & Sabbadin, 2005, for a review). Deciding on such an aggregation rule involves multiple formal assumptions, none of which I am willing to commit to at this stage of theoretical development.

Another option is to consider that the utility grid of a conditional and its subjective probability are two independent components of its broad *epistemic mental model* (Evans, 2007; Evans et al., 2008). The broad epistemic mental model of a conditional would encode the subjective probability of q given p as well as the utility of p and q to various agents. Then, we may assume that the probability of the conditional moderates paralogical inferences based on distal utility: That is, the strength with which paralogical inferences involving u' are invited is an increasing function of the probability of "if p, then q." I am aware that this is a purely operational treatment of the problem of aggregation, which eschews its formal difficulty. I believe, though, that such a treatment can be acceptable when dealing with a psychological phenomenon whose formalization would require multiple arbitrary assumptions.

More Folk Axioms

As for now, the theory of utility conditionals includes three folk axioms of decision: the folk axiom of self-interested behavior, the folk axiom of self-interested attitude, and the folk axiom of limited altruism. These three folk axioms already account for many paralogical inferences invited by utility conditionals, but future research is likely to consider additional folk axioms accounting for yet-undocumented paralogical inferences.

The opportunity to formulate new folk axioms may arise in particular from a more expressive representation of utility. Moving from a $\{-, 0, +\}$ scale to a $\{--, -, 0, +, ++\}$ scale would already allow for tentatively introducing new folk axioms. Consider, for example, this extended version of the folk axiom of limited altruism:

> *Folk Axiom 4 (Extended Altruism):* People take actions that increase the utility of others insofar as doing so does not decrease their own personal utility by at least the same order of magnitude, and they do not take actions that decrease the utility of others insofar as these actions do not increase their own personal utility by at least the same order of magnitude.

The folk axiom of extended altruism articulates the idea that people believe other people to be willing to sacrifice a small measure of their own utility in order to ensure a much larger gain of utility for another agent. Consider, for example, this utility grid:

$$\begin{Bmatrix} s & - & s \\ \omega & ++ & e \end{Bmatrix}.$$

Applied to this grid, the folk axiom of extended altruism invites the conclusion that s is going to do p, because the personal loss of utility incurred by s is one order of magnitude lower than the gain in utility enjoyed by e. This prediction fits our intuitions regarding Conditional 40 below.

40. If I donate 50 euros to this charity, then the life of a starving child will be saved.

Conditional 40 can be represented by the utility grid above, and it does seem to invite the conclusion that the speaker is going to donate the 50 euros, for the reasons articulated by the folk axiom of extended altruism. This intuition cannot be captured by any of the three folk axioms considered so far, which pleads for the inclusion of some stronger version of the folk axiom of limited altruism, albeit perhaps not as strong as the folk axiom of extended altruism.

Integration With Current Theories of Conditionals

The theory of utility conditionals aims at predicting the paralogical inferences invited by conditionals whose antecedents and consequents relate to the utility functions of various agents, but it does not constitute a theory of conditional inference per se. As a consequence, it cannot stand alone but needs to be integrated with a broader theory of conditional reasoning. In this section, I consider how the theory of utility conditionals might be integrated with three current theories of conditional reasoning: the Bayesian theory, the suppositional theory, and the model theory.

The Bayesian Theory

The Bayesian theory of conditional reasoning (Oaksford & Chater, 2001, 2007) postulates that reasoners apply to conditional reasoning tasks the general probabilistic strategies that they use to handle everyday uncertain information. To solve a conditional reasoning problem, reasoners are assumed to engage mental processes that are computationally equivalent to probabilistic calculus. For example, reasoners are assumed to endorse the conclusion of a modus tollens argument (if p, then q; $\neg q$; therefore, $\neg p$) in proportion to the probability of the conclusion $\neg p$ given the premise $\neg q$; this probability is itself a function of the prior probabilities $\Pr(p)$ and $\Pr(q)$ and of the conditional probability $\Pr(q|p)$, which are all retrieved from background knowledge about p and q.

The decision-theoretic notion of utility was introduced early on in the Bayesian theory, albeit not in relation to utility conditionals. Rather, utility was called upon to account for situations wherein reasoners manipulate disinterested conditionals in order to achieve a goal. This situation is most often studied with the deontic selection paradigm, which was discussed in the section Goal-Directed Reasoning in the Selection Task.

More recently, though, the Bayesian theory has started to address the question of utility in a manner that can be more directly related to the theory of utility conditionals. Hahn and colleagues (Corner & Hahn, 2007; Hahn & Oaksford, 2006, 2007) offered in particular a Bayesian account of *slippery slope* arguments such as Conditional 41a below, which can naturally be rephrased using a conditional statement, as in 41b.

41a. Legalizing cannabis will ultimately lead to increased use of cocaine or heroin. Hence, it should remain banned.

41b. If cannabis is legalized, then it will ultimately lead to increased use of cocaine or heroin. Hence, it should remain banned.

As observed by Hahn and colleagues (Corner & Hahn, 2007; Hahn & Oaksford, 2006, 2007), the subjective force of Conditional 41 depends not solely on the probability that the outcome (increased use of heroin) will follow the action (legalizing cannabis) but also on the utility of this outcome. Hence, the Bayesian theory needs to incorporate such considerations of utility to account for arguments such as Conditional 41. Indeed, Hahn and colleagues convincingly argued that this constitutes a natural extension of the theory. (Conversely, a natural and necessary extension of the theory of utility conditionals is to incorporate the probabilistic information emphasized in the Bayesian theory.)

Now, of course, the conditional in 41b is a utility conditional as defined in the present article. It can be represented by a utility grid, and folk axioms of decision can be applied to this grid to predict the paralogical inferences that it invites, including that cannabis should remain banned. The theory of utility conditionals also applies to other arguments that, like the slippery slope, require one to augment a purely probabilistic account with a notion of decision-theoretic utility. I thus believe in the possibility of augmenting the Bayesian theory with all the insights and predictive power of the theory of utility conditionals. As illustrated at length in the present article, many ordinary conditionals have utilitarian features, which trigger inferences that cannot be predicted on the sole basis of the probability of the conditional. Slippery slope arguments are, in that respect, the tip of the iceberg. The theory of utility conditionals can help the Bayesian theory to identify and address all the utility arguments that have not attracted as much attention as has the slippery slope argument.

The Suppositional Theory

The suppositional theory of conditionals (Evans, 2007; Evans & Over, 2004) assumes that *if* is a linguistic device whose purpose is to trigger a process of hypothetical thinking. According to the theory, the mental representation of an "if p, then q" conditional includes a subjective connection between the representation of p and that of q. This subjective connection indicates the degree of belief in q given p, and the degree of belief in the conditional accordingly corresponds to the conditional probability $\Pr(q|p)$.

This mental representation of the conditional is an epistemic mental model, which represents what a reasoner believes about possible states of the world. The epistemic mental model of a conditional features possible states of the world, propositional attitudes about them (it encodes the extent to which one *believes q* when *supposing p*), and some probabilistic information related to the subjective probability $\Pr(q|p)$. Other aspects of the conditional may be considered for inclusion in its epistemic mental model. For example, Evans et al. (2008) suggested that epistemic mental models ought to be augmented to represent the utilities associated with p and q for various agents. Utility grids precisely offer such a representational extension, and augmenting the current epistemic mental models with utility grids could deliver the full-fledged epistemic mental model sought by the suppositional theory. In any case, the theory of utility conditionals must be developed in order to incorporate the probabilistic links emphasized in the suppositional theory.

Another important aspect of the suppositional theory is its dual-process approach to conditional reasoning. Dual-process accounts of conditional reasoning (see, e.g., Bonnefon, Eid, Vautier, & Jmel, 2008; Klaczynski & Daniel, 2005; Oberauer, 2006; Verschueren, Schaeken, & d'Ydewalle, 2005; or see Evans, 2007, 2008a, for broader surveys) assume that inferences can reflect, at different times, the operation of one type of mental processes (Type 1) or the other (Type 2). Type 1 processes are fast, automatic, and undemanding of cognitive resources. Type 2 processes are slow, deliberately engaged and controlled, and demanding of capacity.

Whereas Type 2 processes depend on the decontextualization of the reasoning task and on the activation of abstract rules of inference, Type 1 processes take into account the whole context of the reasoning task. In conversational contexts in particular, many factors—ranging from the relative expertise of the interlocutors (Stevenson & Over, 2001) to their personality or affective bond (Demeure, Bonnefon, & Raufaste, 2009) to their relative power (Kilpatrick, Manktelow, & Over, 2007)—can affect the processing of a conditional. It would be tempting to consider that utilities are yet another contextual factor in conditional reasoning, processed by fast and automatic Type 1 computations. Indeed, Evans et al. (2008) would appear to make that suggestion when commenting that

> Our results only make sense if it is assumed that the participants processing conditionals in context, in some sense simulate (*however, rapidly and intuitively*) an action scenario taking into account the utility of both actions and consequences and the strength of the link between the two. (p. 115, emphasis added)

Note, though, that we need to distinguish the representational processes that retrieve the utility grid of the conditional from the inferential processes that result in paralogical conclusions (invited by the folk axioms of decision). The simulation described by Evans et al. (2008), which allows for retrieval of the utility grid of a conditional, is likely to be a fast, automatic Type 1 process. But the application of (at least some) folk axioms of decision to the utility grid, once it is retrieved, might well require some conscious, deliberate reasoning (see also Thompson et al., 2005, for a related suggestion). If really the application of the folk axioms of decision requires deliberate thinking, then a conflict between the conclusions invited by two folk axioms would reflect a conflict between two Type 2 conclusions, unlike the more commonly investigated case of a conflict between Type 1 and Type 2 conclusions (De Neys & Glumicic, 2008; Evans, 2008b).

The Model Theory

The mental model theory of conditionals (Johnson-Laird & Byrne, 2002) assumes that an "if *p,* then *q*" conditional is mentally represented as the set of possibilities that make it true (i.e., $< pq, \neg pq, \neg p \neg q >$). Coreferential links between *p* and *q,* as well as contextual information, can add possibilities or prevent the construction of some possibilities. These two influences are called semantic modulation and pragmatic modulation, respectively. Due to semantic and pragmatic modulation, a conditional can be represented by 10 distinct combinations of the possibilities pq, $p\neg q$, $\neg pq$, $\neg p \neg q$ (for general and local comments on the model theory of conditionals, see, e.g., Bonnefon, 2004; Bonnefon & Vautier, 2008; Evans, Over, & Handley, 2005; Politzer, 2007).

We can assume that the utility of *p* and *q* for various agents is part of the contextual information surrounding the conditional, and accordingly this affects the conditional through pragmatic modulation. The theory of utility conditionals thus allows for the unpacking of the workings of pragmatic modulation (as far as utility is concerned) into a two-step process.

First, reasoners retrieve contextual information about utility in the form of the utility grid of the conditional. Second, folk axioms of decision determine which possibilities are blocked during the construction of the mental model of the conditional. Consider, for example, Conditional 42 below.

42. If Bruce goes to the meeting, then he will see the woman he loves.

Reasoners first retrieve the following utility grid of 42:

$$\begin{Bmatrix} e & 0 & A \\ \omega & + & e \end{Bmatrix}.$$

Then, the folk axiom of self-interested behavior leads reasoners to discard the possibilities featuring $\neg p$, because $\neg p$ is an action that would decrease the utility of Bruce (Bonnefon, 2007; Bonnefon & Hilton, 2004). As a consequence, the mental model of 42 is reduced to the single possibility $< pq >$, from which participants endorse the conclusion *p.*

This rudimentary analysis suggests that the theory of utility conditionals could provide the mental model theory with one of its mostly required add-ons: that is, an unpacked account of the processes underlying (part of) pragmatic modulation. Note, however, that some phenomena predicted by the theory of utility conditionals would appear to fall beyond the predictive power of mental model theory. For example, the predictions derived from the folk axiom of self-interested attitude, or the reinforcement of an invited conclusion by multiple folk axioms, do not seem amenable to an account in terms of mental models.

Conclusion

As the primary expressions of hypothetical thinking, conditionals are involved in decision making as well as reasoning. The logicist paradigm long assumed that the decision-theoretic features of conditionals were irrelevant to the conclusions that reasoners would derive from these conditionals. This assumption, however, was gradually abandoned when students of reasoning started to investigate goal-directed reasoning, as well as reasoning based on special types of conditionals (e.g., inducements, persuasion, consequentials). In this article, I have organized this substantial body of literature into an integrated theoretical framework, the theory of utility conditionals.

The theory of utility conditionals describes how the decision-theoretic features of conditional statements leak into reasoning contexts, accounting for the paralogical inferences invited by conditionals whose antecedents and consequents relate to the utility functions of various agents. The theory of utility conditionals rests on two main components: a representational tool (the utility grid), which summarizes in compact form the decision-theoretic features of a conditional statement, and a set of folk axioms of decision, which captures reasoners' likely beliefs about the way most agents make their decisions. Applying the folk axioms of decision to the utility grid of a conditional allows for prediction of the paralogical conclusions invited by its decision-theoretic features.

Not only does the theory account for previously identified subtypes of utility conditionals, but it also allows for the identification of brand-new classes of such conditionals by studying the permutations of the utility grid. The inferential properties of these novel classes of utility conditionals are then straightforwardly predicted by the theory, by applying the folk axioms of decision to their utility grid. The predictive power of the theory of utility conditionals can be refined by introducing finer grained variations of utility and probability and by investigating additional folk axioms of decision. Finally, the theory of utility conditionals can be integrated with broader theories of conditional inference, supplying them with an add-on they are explicitly looking for. It is my hope that the theory of utility conditionals will constitute a milestone in the gradual integration of reasoning research with decision-making research.

References

Amgoud, L., Bonnefon, J.-F., & Prade, H. (2005). An argumentation-based approach for multiple criteria decision. *Lecture Notes in Computer Science, 3571,* 269–280.

Amgoud, L., Bonnefon, J.-F., & Prade, H. (2007). The logical handling of threats, rewards, tips, and warnings. *Lecture Notes in Artificial Intelligence, 4724,* 235–246.

Baron, J. (2000). *Thinking and deciding* (3rd ed.). Cambridge, England: Cambridge University Press.

Beller, S., Bender, A., & Kuhnmünch, G. (2005). Understanding conditional promises and threats. *Thinking and Reasoning, 11,* 209–238.

Bonnefon, J.-F. (2004). Reinstatement, floating conclusions, and the credulity of mental model reasoning. *Cognitive Science, 28,* 621–631.

Bonnefon, J.-F. (2007). Reasons to act and the mental representation of consequentialist aberrations. *Behavioral and Brain Sciences, 30,* 454–455.

Bonnefon, J.-F., Dubois, D., Fargier, H., & Leblois, S. (2008). Qualitative heuristics for balancing the pros and cons. *Theory and Decision, 65,* 71–95.

Bonnefon, J.-F., Eid, M., Vautier, S., & Jmel, S. (2008). A mixed Rasch model of dual-process conditional reasoning. *Quarterly Journal of Experimental Psychology, 61,* 809–824.

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.

Bonnefon, J.-F., & Hilton, D. J. (2005). Getting the point of conditionals: An argumentative approach to the psychological interpretation of conditional premises. In G. Kern-Isberner, W. Rodder, & F. Kulmann (Eds.), *Conditionals, information, and inference* (pp. 59–64). Berlin, Germany: Springer Verlag.

Bonnefon, J.-F., & Vautier, S. (2008). Defective truth tables and falsifying cards: Two measurement models yield no evidence of an underlying fleshing-out propensity. *Thinking and Reasoning, 14,* 231–243.

Byrne, R. M. J. (1989). Suppressing valid inferences with conditionals. *Cognition, 31,* 61–83.

Corner, A., & Hahn, U. (2007). Evaluating the meta-slope: Is there a slippery slope argument against slippery slope arguments? *Argumentation, 21,* 349–359.

Cummins, D. D., Lubart, T., Alksnis, O., & Rist, R. (1991). Conditional reasoning and causation. *Memory & Cognition, 19,* 274–282.

Demeure, V., Bonnefon, J.-F., & Raufaste, E. (2009). Politeness and conditional reasoning: Interpersonal cues to the indirect suppression of deductive inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35,* 260–266.

De Neys, W., & Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. *Cognition, 106,* 1248–1299.

De Neys, W., Schaeken, W., & d'Ydewalle, G. (2002). Causal conditional reasoning and semantic memory retrieval: A test of the semantic memory framework. *Memory & Cognition, 30,* 908–920.

De Neys, W., Schaeken, W., & d'Ydewalle, G. (2003). Inference suppression and semantic memory retrieval: Every counterexample counts. *Memory & Cognition, 31,* 581–595.

Dieussaert, K., Schaeken, W., & d'Ydewalle, G. (2002). The relative contribution of content and context factors on the interpretation of conditionals. *Experimental Psychology, 49,* 181–195.

Dubois, D., Fargier, H., & Bonnefon, J.-F. (2008). On the qualitative comparison of decisions having positive and negative features. *Journal of Artificial Intelligence Research, 32,* 385–417.

Evans, J. St. B. T. (2002). Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin, 128,* 978–996.

Evans, J. St. B. T. (2005). The social and communicative function of conditional statements. *Mind & Society, 4,* 97–113.

Evans, J. St. B. T. (2007). *Hypothetical thinking: Dual processes in reasoning and judgment.* Hove, England: Psychology Press.

Evans, J. St. B. T. (2008a). Dual-processing accounts of reasoning. *Annual Review of Psychology, 59,* 255–278.

Evans, J. St. B. T. (2008b). On the resolution of conflict in dual-process theories of reasoning. *Thinking and Reasoning, 13,* 321–329.

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., Neilens, H., Handley, S. J., & Over, D. E. (2008). When can we say "if"? *Cognition, 108,* 100–116.

Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning.* Hove, England: Psychology Press.

Evans, J. St. B. T., & Over, D. E. (2004). *If.* Oxford, England: Oxford University Press.

Evans, J. St. B. T., Over, D. E., & Handley, S. J. (2005). Suppositions, extensionality, and conditionals: A critique of the mental model theory of Johnson-Laird and Byrne (2002). *Psychological Review, 112,* 1040–1052.

Evans, J. St. B. T., & Twyman-Musgrove, J. (1998). Conditional reasoning with inducements and advice. *Cognition, 69,* B11–B16.

Fargier, H., & Sabbadin, W. (2005). Qualitative decision under uncertainty: Back to expected utility. *Artificial Intelligence, 164,* 245–280.

Fillenbaum, S. (1975). If: Some uses. *Psychological Research, 37,* 245–260.

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

Fillenbaum, S. (1978). How to do some things with IF. In J. W. Cotton & R. L. Klatzky (Eds.), *Semantic factors in cognition* (pp. 169–231). Hillsdale, NJ: Erlbaum.

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

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

Hahn, U., & Oaksford, M. (2006). A Bayesian approach to informal argument fallacies. *Synthese, 152,* 207–236.

Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to argument fallacies. *Psychological Review, 114,* 704–732.

Hilton, D. J., Kemmelmeier, M., & Bonnefon, J.-F. (2005). Putting ifs to work: Goal-based relevance in conditional directives. *Journal of Experimental Psychology: General, 135,* 388–405.

Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review, 109,* 646–678.

Kilpatrick, S. G., Manktelow, K. I., & Over, D. E. (2007). Power of source as a factor in deontic inference. *Thinking and Reasoning, 13,* 295–317.

Klaczynski, P. A., & Daniel, D. B. (2005). Individual differences in conditional reasoning: A dual-process account. *Thinking and Reasoning, 11,* 305–325.

Klauer, K. C., Stahl, C., & Erdfelder, E. (2007). The abstract selection task: New data and an almost comprehensive model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33,* 680–703.

Legrenzi, P., Politzer, G., & Girotto, V. (1996). Contract proposals: A sketch of a grammar. *Theory & Psychology, 6,* 247–265.

López-Rousseau, A., & Ketelaar, T. (2004). "If . . .": Satisficing algorithms for mapping conditional statements onto social domains. *European Journal of Cognitive Psychology, 16,* 807–823.

López-Rousseau, A., & Ketelaar, T. (2006). Juliet: If they do see thee, they will murder thee: A satisficing algorithm for pragmatic conditionals. *Mind & Society, 5,* 71–77.

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

Miller, D. T. (1999). The norm of self-interest. *American Psychologist, 54,* 1053–1060.

Newstead, S. E., Ellis, C., Evans, J. St. B. T., & Dennis, I. (1997). Conditional reasoning with realistic material. *Thinking and Reasoning, 3,* 49–76.

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. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences, 5,* 349–357.

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning.* Oxford, England: Oxford University Press.

Oberauer, K. (2006). Reasoning with conditionals: A test of formal models of four theories. *Cognitive Psychology, 53,* 238–283.

Oberauer, K., & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 680–693.

Ohm, E., & Thompson, V. (2004). Everyday reasoning with inducements and advice. *Thinking and Reasoning, 10,* 241–272.

Ohm, E., & Thompson, V. (2006). Conditional probability and pragmatic conditionals: Dissociating truth and effectiveness. *Thinking and Reasoning, 12,* 257–280.

Over, D. E., Hadjichristidis, C., Evans, J. St. B. T., Handley, S. J., & Sloman, S. A. (2007). The probability of causal conditionals. *Cognitive Psychology, 54,* 62–97.

Perham, N. R., & Oaksford, M. (2005). Deontic reasoning with emotional content: Evolutionary psychology or decision theory? *Cognitive Science, 29,* 681–718.

Politzer, G. (2007). Reasoning with conditionals. *Topoi, 26,* 79–95.

Politzer, G., & Bonnefon, J.-F. (2006). Two varieties of conditionals and two kinds of defeaters help reveal two fundamental types of reasoning. *Mind and Language, 21,* 484–503.

Politzer, G., & Nguyen-Xuan, A. (1992). Reasoning about conditional promises and warnings: Darwinian algorithms, mental models, relevance judgements or pragmatic schemas? *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 44*(A), 401–412.

Ratner, R. K., & Miller, D. T. (1998). The disparity between the actual and assumed power of self-interest. *Journal of Personality and Social Psychology, 74,* 53–62.

Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning.* Mahwah, NJ: Erlbaum.

Staudenmayer, H. (1975). Understanding conditional reasoning with meaningful propositions. In R. J. Falmagne (Ed.), *Reasoning: Representation and process in children and adults* (pp. 55–79). Hillsdale, NJ: Erlbaum.

Stevenson, R. J., & Over, D. E. (2001). Reasoning from uncertain premises: Effects of expertise and conversational context. *Thinking and Reasoning, 7,* 367–390.

Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory & Cognition, 22,* 742–758.

Thompson, V. A., Evans, J. St. B. T., & Handley, S. J. (2005). Persuading and dissuading by conditional argument. *Journal of Memory and Language, 53,* 238–257.

Van Canegem-Ardijns, I., & Van Belle, W. (2008). Conditionals and types of conditional perfection. *Journal of Pragmatics, 40,* 349–376.

Vartanian, O., Mandel, D. R., & Blackler, K. (2008, August). *When liking is believing: Consequential conditionals as a framework for communicating government policy.* Paper presented at the Sixth International Conference on Thinking, Venice, Italy.

Verbrugge, S., Dieussaert, K., Schaeken, W., & Van Belle, W. (2004). Promise is debt, threat another matter: The effect of credibility on the interpretation of conditional promises and threats. *Canadian Journal of Experimental Psychology, 58,* 106–112.

Verschueren, N., Schaeken, W., & d'Ydewalle, G. (2005). A dual-process specification of causal conditional reasoning. *Thinking and Reasoning, 11,* 239–278.

Wason, P. (1966). Reasoning. In B. M. Foss (Ed.), *New horizons in psychology* (pp. 106–137). Harmandsworth, England: Penguin.

Received March 31, 2009
Revision received July 6, 2009
Accepted July 7, 2009 ■