

Is the Basic Conditional Probabilistic?

Geoffrey P. Goodwin
University of Pennsylvania

Nine experiments examined whether individuals treat the meaning of basic conditional assertions as deterministic or probabilistic. In Experiments 1–4, participants were presented with either probabilistic or deterministic relations, which they had to describe with a conditional. These experiments consistently showed that people tend only to use the basic *if p then q* construction to describe deterministic relations between antecedent and consequent, whereas they use a probabilistically qualified construction, *if p then probably q*, to describe probabilistic relations—suggesting that the default interpretation of the conditional is deterministic. Experiments 5 and 6 showed that when directly asked, individuals typically report that conditional assertions admit no exceptions (i.e., they are seen as deterministic). Experiments 7–9 showed that individuals judge the truth of conditional assertions in accordance with this deterministic interpretation. Together, these results pose a challenge to probabilistic accounts of the meaning of conditionals and support mental models, formal rules, and suppositional accounts.

Keywords: conditionals, conditional meaning, conditional reasoning, conditional probability, pragmatics

If you do not brush your teeth, you will get cavities. What does this conditional assertion mean? Does it mean that getting cavities is an inevitable (i.e., deterministic) consequence of not brushing your teeth? Or does it mean that getting cavities is merely a highly probable though not inevitable consequence of not brushing your teeth? Despite more than 40 years of psychological research on conditionals, this question about the basic meaning of conditional assertions (*if p then q*) has not been resolved. Recent, vigorous debate has focused on processes of conditional inference—how people reason from conditional assertions—and, more recently, on how people judge the probability of conditionals. Yet few studies have directly addressed the question of what conditionals mean, even though the various theories of conditional inference all make assumptions about their meaning. The aim of the present article is therefore to fill this gap.

Conditionals are ubiquitous as a means of representing knowledge and belief. Reasoning from conditionals is thus pervasive in everyday, legal, and scientific contexts (see, e.g., Johnson-Laird & Byrne, 2002; Oberauer, 2006). Yet the various extant theories of conditionals are sharply divided on whether the meaning of basic conditionals is probabilistic or deterministic. Mental model, formal rules, and suppositional accounts treat the conditional as having a deterministic meaning. In contrast, a variety of probabilistic accounts treat the conditional as having a probabilistic meaning.

Thus, resolving whether ordinary individuals treat conditionals as probabilistic or deterministic would help adjudicate between these competing accounts.

Despite fundamental disagreement concerning the processes of conditional inference, formal rule theories (see, e.g., Braine & O'Brien, 1991, 1998; Rips, 1994, 2002) and the mental model theory (see, e.g., Johnson-Laird, 2006; Johnson-Laird & Byrne, 2002) take conditional assertions to express a deterministic relation between antecedent and consequent. According to formal rule theories, a conditional assertion of the form *if p then q* expresses a deterministic relation between its antecedent, *p*, and its consequent, *q*. If *p* happens to hold, then *q* automatically follows, as a matter of valid deduction (a modus ponens inference). Such a conditional tolerates no exceptions. The mental model theory similarly holds that conditional assertions express a deterministic relation between antecedent and consequent. According to the model theory, the core meaning of a conditional assertion, *if p then q*, is consistent with three separate possibilities (see also Table 1):

<i>p</i>	<i>q</i>
not <i>p</i>	<i>q</i>
not <i>p</i>	not <i>q</i>

As these possibilities indicate, *p* only ever occurs with *q*. In other words, whenever *p* occurs, *q* invariably follows. This interpretation corresponds to that of material implication in logic. According to the model theory, this interpretation will be made for *basic conditionals*. A basic conditional is one “for which general knowledge, the meaning or reference of its clauses, or knowledge of its context, has no effect on the interpretation of the relation between its if-clause and its then-clause” (Johnson-Laird, 2011, p. 121; see also Evans, Handley, Neilens, & Over, 2007, pp. 1772–1773, who define basic conditionals as concerning abstract relations and as being indicative in form).

However, the model theory allows that interpretations of the conditional may sometimes diverge from the material implication

This article was published Online First November 4, 2013.

For their insightful comments and support, I thank Jeremy Boyd, Adele Goldberg, Sam Glucksberg, Dena Gromet, Phil Johnson-Laird, Sangeet Khemlani, Paula Rubio, Joe Simmons, and Uri Simonsohn. I also thank David Over for his helpful and stimulating critiques. Alison Newman, Sarah Patrick, Yael Diamond Schlenger, and Christina Torres provided valuable research assistance.

Correspondence concerning this article should be addressed to Geoffrey P. Goodwin, Department of Psychology, University of Pennsylvania, 3720 Walnut Street, Solomon Lab Building, Philadelphia, PA 19104. E-mail: ggoodwin@psych.upenn.edu

Table 1
Representation of the Possibilities Compatible With the Material Implication and Defective Interpretations of Conditionals

Possibilities		Material implication conditional	Defective conditional
p	q	True	True
p	$\neg q$	False	False
$\neg p$	q	True	Irrelevant
$\neg p$	$\neg q$	True	Irrelevant

interpretation. Meaning, general knowledge, and context can sometimes modulate the meaning of a conditional assertion, usually by adding temporal, spatial, or other sorts of relation between antecedent and consequent. In addition to adding such relations, modulation can sometimes eliminate possibilities from the above set (see Johnson-Laird & Byrne, 2002, for more details; see also Quelhas, Johnson-Laird, & Juhos, 2010, for empirical corroboration). There is, therefore, no single meaning for conditional assertions (Johnson-Laird & Byrne, 2002). Nonetheless, for basic conditionals, the model theory posits an interpretation corresponding to that for material implication. This interpretation expresses a deterministic relation between antecedent and consequent, unless the content of the conditional includes an explicit reference to probabilistic information, for example, “If it rains, then it’s likely that he’ll get wet.”

Contrasting with both model- and rule-based theories is a variety of more recent theories that are generically grouped together as *probabilistic* theories (see, e.g., Evans & Over, 2012; Fugard, Pfeifer, Mayerhofer, & Kleiter, 2011). Each of these theories stems from Adams’s seminal work on conditionals (Adams, 1975, 1998). Yet important differences between these theories exist with respect to both their explanatory focus and their core assumptions about the meaning of conditionals.

The most extensively investigated of these theories is based on a suppositional account of conditionals (see, e.g., Evans et al., 2007; Evans, Handley, Neilens, & Over, 2008; Evans, Handley, & Over, 2003; Evans, Neilens, Handley, & Over, 2008; Evans & Over, 2004; Evans, Over, & Handley, 2005; Over & Evans, 2003; Over, Evans, & Elqayam, 2010; Politzer, Over, & Baratgin, 2010). The main impetus of this theory has been to understand how individuals judge the probability of conditionals (which, according to this theory, refers to how individuals establish their subjective degree of belief in a conditional). The theory relies on three interrelated assumptions. First, according to Evans and Over (2004), individuals judge their confidence or belief in a conditional assertion by relying on the *Ramsey test*, that is, by hypothetically adding its antecedent to their stock of knowledge and judging the likelihood of the consequent in the resulting state of affairs. Thus, as a consequence, individuals judge the probability of a conditional assertion in accordance with the conditional probability of the consequent given the antecedent, that is, $P(\text{if } p \text{ then } q) = P(q|p)$, sometimes referred to simply as *The Equation*. The third assumption, which follows from the first two, is that individuals represent the conditional in terms of a defective truth table (see Table 1). According to this account, the conditional is true in the case in which antecedent and consequent are both true but false when the antecedent is true and the consequent false. The false antecedent

cases are treated as irrelevant to the truth of the conditional, which has no truth value (i.e., is indeterminate) for such cases (hence the *defective* truth table). This defective interpretation of the conditional is regarded as closely linked to an individual’s reliance on The Equation, and recent evidence supports this connection (Evans et al., 2007; Politzer et al., 2010).

While, on the surface, this suppositional account might seem to yield a probabilistic interpretation of the meaning of conditionals, in fact, it arguably leads to the very same conclusion that formal rule and mental model theories give rise to, namely, that the conditional expresses a deterministic relation between antecedent and consequent. To see this, it is important to recognize that the suppositional theory applies only to conditionals that pertain to single events (e.g., *If the card is yellow, then it has a circle printed on it; If Sally does not brush her teeth, then she will get cavities*), rather than more generally across a range of events (e.g., *If a card is yellow, then it has a circle printed on it; If people do not brush their teeth, then they will get cavities*; see Evans et al., 2003, p. 333). Accordingly, the conditional can be seen as making a deterministic statement about the single event or case in question, such that the conditional is true when both antecedent and consequent are true and is false when the antecedent is true and the consequent false. Nonetheless, information pertaining to the distribution of pertinent cases may still be relied upon to judge the probability of a specific conditional statement being true of the case in question—as has been corroborated empirically (see, e.g., Evans et al., 2003, 2007; Geiger & Oberauer, 2007; Handley, Evans, & Thompson, 2006; Oberauer, Geiger, Fischer, & Weidenfeld, 2007; Oberauer & Wilhelm, 2003; Over, Hadjichristidis, Evans, Handley, & Sloman, 2007). Thus, because, in the defective truth table, p also only ever occurs with q and never occurs with *not* q , the consequent’s following from the antecedent is just as invariable according to the defective interpretation as it is on an interpretation that corresponds to material implication (David Over, personal communication, February 13, 2012). This important consequence of the suppositional theory has not been made explicit previously.

At the same time, however, there are some indications from the proponents of the suppositional theory that run contrary to this deterministic interpretation and that are more consistent with the idea that the conditional has a probabilistic interpretation. For instance, the proponents of this theory have remarked that the conditional “expresses a subjective degree of conditional confidence, which equals the conditional probability” (Evans & Over, 2004, p. 31), where the conditional probability refers to the probability of the consequent holding given that the antecedent is true. In a similar vein, these theorists have asserted that “an ordinary conditional assertion *if* p *then* q is interpreted as q is probable given p ” (Evans et al., 2003, p. 322) and that “an indicative conditional tends to be asserted and accepted when its consequent seems highly probable given its antecedent” (Over, 1993, p. 361). Nonetheless, setting these relatively isolated conflicting remarks aside, it seems that the suppositional theory holds the same position as do rule and model theories, namely, that the conditional expresses a deterministic relation between antecedent and consequent.

However, other theories within the broad suite of probabilistic theories clearly endorse the weaker claim that the conditional expresses merely a probabilistic relation between antecedent and consequent. On these theories, the conditional conveys merely that

the conditional probability of the consequent given the antecedent is high, but not necessarily certain. The line demarcating true (or acceptable) conditionals from false (or unacceptable) conditionals, even if fuzzily drawn, thus represents what might be referred to as a *semantic threshold* (David Over, personal communication, February 13, 2012). These probabilistic theories have typically had a different explanatory focus than has the suppositional theory and have been more concerned with processes of conditional inference than with judgments of the probability of conditionals. The earliest of these theories is due to Oaksford and Chater (1994), who developed a probabilistic theory that was designed originally to account for performance on the Wason selection task and that was later updated to handle conditional inference more generally (Oaksford, Chater, & Larkin, 2000). The processes of conditional inference were rendered in a probabilistic way, such that the likelihood of basic conditional inferences such as modus ponens was a function of the strength of the relation between antecedent and consequent. Oaksford and Chater (1998b) formally built in an *exceptions parameter* to represent the occurrence of exceptions to the dependency between antecedent and consequent (see also Oaksford & Chater, 2001; Oaksford et al., 2000). This exceptions parameter was designed to reflect the assumption that such exceptions are included in the ordinary meaning of conditional assertions. Several illustrative remarks by these authors make this clear. For instance, Oaksford and Chater (1998b, p. 377) wrote that:

We have argued elsewhere that people do not interpret everyday conditional rules as exceptionless (Chater & Oaksford, 1990, 1993; Oaksford & Chater, 1991, 1992, 1993, 1995a; Oaksford, 1995). Any everyday rule, such as *birds fly, if you put money in the coke machine, you get a coke*, and so on, succumb to indefinitely many exceptions, such as ostriches, penguins and broken or empty coke machines.

Oaksford and Chater (2001, pp. 350–351) similarly wrote that “the contingency table represents a conditional rule, *if p then q*, where there is a dependency between the *p* and *q* that may admit exceptions (ϵ).” More recently, Oaksford and Chater (2009, p. 78) elaborated this account, writing that:

The idea again relies on interpreting a conditional in terms of conditional probability. For example, the hypothesis, *if swan (p) then white (q)*, is interpreted as making the claim that the probability of a bird being white given that it is a swan, $P(q|p)$, is high, certainly higher than the base rate of being a white bird, $P(q)$. This hypothesis is called the dependence hypothesis (HD).

The theory is clear, and there is thus no doubt that it takes the conditional to have a probabilistic meaning (on this theory, the conditional is usually referred to as a *probability conditional*; see Evans & Over, 2012).

This probabilistic assumption is also made by several other theorists, who, like Oaksford and Chater, have tended to focus on conditional inference. Liu, Lo, and Wu (1996) developed a theory according to which the conditional expresses a probabilistic relation between antecedent and consequent. They wrote:

However, most people may understand the first premise of the present argument [which was: *If it rains, then there is enough water for the mowers*] as “*If it rains, then probably there will be enough water for the mowers*”. In other words, judging from our experiences of the extent to which *p* predicts *q*, we tend to interpret “*if p then q*” to mean “*in case p occurs, q is likely to occur to that extent*”. Thus, “*if-then*”

is interpreted probabilistically, depending on the perceived sufficiency of *p* for *q*. (Liu et al., 1996, p. 830, bracketed remark not in original)

Similarly, Oberauer and Wilhelm (2003, p. 685), who focused on both conditional inference and judgments of the probability of conditionals, expressed this same idea as follows:

If people have a continuous degree of belief in a conditional “*if p then q*” and are forced to choose between “*true*” or “*false*,” they must judge whether their degree of belief is sufficiently high to accept the conditional as true. We assume that they set a threshold such that degrees of belief larger than the threshold are regarded as sufficiently high to warrant a “*true*” judgment, and degrees lower than the threshold result in a “*false*” judgment. . . . The threshold can be assumed to vary between individuals. We expected that on average it would be high, because a conditional expresses the belief that $P(q|p)$ is high or very high.

Finally, a somewhat more recent theory, based on conditional probability logics, also shares the assumption of probabilistic meaning and regards conditionals as expressing a *conditional event* (see, e.g., Pfeifer & Kleiter, 2005, 2006, 2009). According to this theory, as Pfeifer and Kleiter (2005) wrote, “In everyday contexts it seems to be more plausible to interpret conditionals not by material implications, but by much weaker conditional probabilities” (p. 72). Similarly, they wrote that “the semantics of an indicative ‘*if A, then B*’ is given by the conditional probability, $P(B/A)$ ” (Fugard et al., 2011, p. 636). And, perhaps most clearly:

In common sense reasoning, conditionals are inherently uncertain, as they hold only “*probably*”, “*normally*”, or “*usually*”. A few exceptions do not invalidate the conditional. Nonmonotonic conditionals express uncertain relations between the *if* and the *then* part of a conditional assertion. The nonmonotonic conditional is interpreted as a “*high*” conditional probability assertion. (Pfeifer & Kleiter, 2010, p. 154)

Thus, although Pfeifer and Kleiter’s theory differs from other probabilistic theories in important respects (e.g., it treats conditional inference as deductive, rather than probabilistic, unlike Oaksford and Chater’s theory), for present purposes, its key claim is that conditional assertions express probabilistic relations.

As these quotes illustrate, these probabilistic theories all appear to be intended as theories of the core semantics of conditionals, rather than simply as theories of the pragmatics of conditional assertion and interpretation in everyday contexts (although, of course, patterns of pragmatic usage should be closely connected with the core semantics). The focus on semantics is explicit in many of the illustrative quotes just reviewed, which either directly refer to the notion of semantics or instead refer to how conditionals will be understood and interpreted. None of the probabilistic theories described has made a distinction between semantic and pragmatic interpretations of conditionals, nor have any stressed the role of pragmatics in their accounts of conditionals. It thus seems most plausible to interpret these theories as making semantic claims, and I will treat them as such in what follows.

Nonetheless, the notion of a probabilistic threshold, described in the theories above, could potentially be reconstrued as referring not to a semantic threshold but rather to a more pragmatic threshold that governs the everyday patterns of usage and interpretation of conditionals. One version of this weaker claim is that while individuals may recognize that conditionals have a core meaning

that is strict and exceptionless, they are nonetheless willing to accept and assert conditionals in everyday contexts notwithstanding that some evidence conflicts with this exceptionless reading. This version thus treats probabilistic evidence as sometimes being sufficient to make a nonprobabilistic claim.¹ A second, closely related view is that conditional assertions generate pragmatic implicatures that license a weaker, probabilistic interpretation. Accordingly, individuals will typically use and interpret conditionals as expressing a loose, exception-tolerating relation between antecedent and consequent in accordance with this pragmatic meaning while nonetheless still recognizing that the core meaning of a conditional does not tolerate exceptions. Both of these reinterpretations do not accord with how these theories have been described in the literature. However, I elaborate on how both of these possibilities might square with the evidence of the present experiments in the General Discussion.

To briefly summarize thus far, mental model and formal rule theories claim that the meaning of basic conditionals is deterministic (or exceptionless). In the case of the model theory, this conclusion results from the analysis of possibilities that the conditional gives rise to. In the case of rule theories, it results from the deterministic rule that corresponds to a modus ponens inference. There is a rich variety of probabilistic accounts of the conditional, which appear to differ in their interpretations. The suppositional account is the only one of the probabilistic accounts presented that allows that the conditional can have truth values. And owing to its reliance on the defective truth table, it appears to make the same deterministic interpretation of conditionals as do rule and model theories. But all of the other existing probabilistic accounts explicitly favor a probabilistic reading of the conditional (and also deny that it can have truth values). Moreover, all proponents of these probabilistic theories of the conditional have expressed clear adherence to this probabilistic reading.

There is thus a major theoretical divide among psychologists when it comes to the interpretation of conditional assertions. This issue has also divided the opinions of philosophers. Some philosophers have taken the conditional to express a deterministic or exceptionless relation between antecedent and consequent (see, e.g., [Dudman, 1987, 1992](#); [Lowe, 1996](#)), whereas other philosophers have explicitly argued that conditionals tolerate exceptions (see, e.g., [Adams, 1975, 1998](#); [Appiah, 1985, 1987](#); [Douven, 2008](#); [Edgington, 1996](#); [Jackson, 1979](#)). Thus, resolving this issue should be of theoretical interest to philosophers as well. Moreover, separate from this theoretical divide, resolving the meaning of conditionals is an important task in its own right, given the ubiquity of conditional knowledge and conditional reasoning in daily life.

Curiously, however, the question of whether the conditional expresses a deterministic or probabilistic meaning has been largely neglected in the existing empirical literature. Prior studies have tended to focus either on conditional inference tasks (which in general seem to support the mental model theory; see [Oberauer, 2006](#), for a comprehensive review) or, alternatively, on how people judge the probability of conditional statements (which in general seem to support conditional probability accounts). However, neither of these sorts of studies directly answers the question of what conditionals mean. And in particular, they do not answer the main question of interest for the present article, which is whether people judge conditionals to have a deterministic or a probabilistic meaning.

Some prior studies do bear indirectly on whether conditionals have a probabilistic or deterministic reading. For instance, studies that ask participants to describe what is possible given a basic conditional assertion (see, e.g., [Barrouillet, Gauffroy, & Lecas, 2008](#); [Barrouillet, Grosset, & Lecas, 2000](#)) have shown that individuals tend to produce the three possibilities that a material implication interpretation of the conditional gives rise to, which is consistent with the mental model theory's deterministic account of basic conditionals. Other studies, which have asked participants to indicate what would falsify a conditional, have shown that they routinely select *p*, not *q* cases as falsifying, which is again consistent with the deterministic view (e.g., [Johnson-Laird & Tagart, 1969](#); [Oaksford & Stenning, 1992](#)).

However, this evidence does not directly address the issue presently at hand. Moreover, other evidence, from tasks in which participants have been asked to judge the truth of conditionals, has been interpreted as indicating a probabilistic, exception-tolerating reading of conditionals (see, e.g., [Evans, Ellis, & Newstead, 1996](#); [Liu et al., 1996](#); [Oberauer & Wilhelm, 2003](#)). I defer further discussion of this evidence, including some problems with the standard interpretation of these findings, until the General Discussion. For now, however, it seems clear that the meaning of conditional assertions is not settled and that a new form of evidence is called for.

Overview of the Present Experiments

The present article concerns only what have been referred to as *basic* conditionals. A basic conditional is one “for which general knowledge, the meaning or reference of its clauses, or knowledge of its context, has no effect on the interpretation of the relation between its if-clause and its then-clause” ([Johnson-Laird, 2011](#), p. 121; see also [Evans et al., 2007](#)). It also exclusively concerns conditionals for which some distribution of relevant evidence can be brought to bear regarding the strength of the relation between antecedent and consequent, as has been the typical practice in the existing psychological literature (see, e.g., [Evans et al., 2003, 2007](#); [Fugard et al., 2011](#); [Geiger & Oberauer, 2007](#); [Oaksford et al., 2000](#); [Oberauer et al., 2007](#); [Oberauer & Wilhelm, 2003](#); [Over et al., 2010](#), among others). The investigation includes conditionals that function as generalizations (e.g., *If a person does not brush their teeth, then they will get cavities*), as well as those that pertain to single events (e.g., *If this person does not brush their teeth, then they will get cavities*), but it excludes conditionals that pertain to truly unique single events for which no past evidence can be brought to bear (e.g., *If aliens land tonight, then John will not brush his teeth*). It aims to bring new forms of evidence to bear on the meaning of conditionals by investigating how individuals use conditionals to describe various states of affairs, how they judge what is possible given a conditional assertion, and how they judge the truth or falsity of conditionals.

¹ I thank David Over for bringing this issue to my attention.

In Experiments 1–4, I investigated how individuals use conditionals to express both probabilistic² and deterministic patterns of data. If the meaning of a basic conditional, *if p then q*, is probabilistic, then it should be perfectly adequate to use such a conditional to describe probabilistic patterns of data. Alternatively, if the meaning of a conditional is deterministic and not probabilistic, then an unqualified conditional is not adequate to describe probabilistic data. Instead, speakers should frequently use some sort of probabilistic qualification when they describe probabilistic data.

Consider a conditional, such as the following:

If an individual smokes, then he or she will get lung cancer.

General knowledge implies that the relation between the antecedent and consequent in this conditional is not deterministic because we know that there are cases in which individuals smoke and yet do not end up with lung cancer. Granted that there are individuals who smoke and who do not get lung cancer, the assertion is thus false according to deterministic theories. Accordingly, to make the assertion true, according to such theories, it needs to be qualified probabilistically, that is,

If an individual smokes, then he or she will probably get lung cancer.

or

If an individual smokes, then he or she will most likely get lung cancer.

or, alternatively,

If an individual smokes, then he or she stands a greater chance of getting lung cancer.

Inserting probabilistic qualification of this sort is vital to convey the truly probabilistic nature of the relation between antecedent and consequent (see also [Dudman, 1992](#)).³

In the first four experiments, individuals were presented with information about the strength of the relation between two different events. These events could function as the antecedent and consequent of a conditional describing the relation. The chief experimental manipulation was whether the relation between antecedent and consequent was probabilistic (i.e., the probability of the consequent given the antecedent being high, but less than 1) or deterministic (i.e., the conditional probability being equal to 1). Deterministic accounts of the conditional predict that when describing probabilistic data, individuals will strongly prefer probabilistically qualified conditionals—and this should be reflected in their selection of conditional descriptions (Experiment 1) and in their generation of conditional descriptions (Experiments 2–4). For deterministic patterns of data, however, individuals will prefer the unqualified, *if p then q*, formulation.

In contrast, probabilistic accounts treat the meaning of such general conditionals as being inherently probabilistic, that is, they treat the meaning of *if p then q* as being *if p then probably q*.⁴ As a consequence, when individuals describe probabilistic relations between antecedent and consequent, they should have no preference between conditionals that insert probabilistic qualification and those that do not. In fact, according to these accounts, this sort of extra qualification seems not only unnecessary but also redundant. Since conditional assertions are already inherently probabilistic, there is no added meaning conveyed by inserting the prob-

abilistic qualification. Indeed, as [Pfeifer and Kleiter \(2010, pp. 163–164\)](#) noted, the fact that conditionals tolerate exceptions is supposed to be an ordinary part of their meaning and is not often stated explicitly. Thus, if anything, adding this extra qualification is an extra burden on speakers, since the effort to insert redundant information is contrary to the Gricean maxim of quantity: Be informative, but not more informative than necessary ([Grice, 1975](#)). In sum, when describing probabilistic data, probabilistic accounts imply that, at the very least, individuals should be indifferent as to whether conditionals are probabilistically qualified. On a stronger reading, which takes into account the Gricean maxim of nonredundancy, conditional probability accounts suggest that individuals should in fact prefer *not* to include such redundant probabilistic qualification.

Following these first four experiments, five further experiments investigated alternative ways of assessing the meaning of conditional assertions. Experiments 5 and 6 asked individuals to indicate what should hold given a conditional assertion. If individuals represent *if p then q* as deterministic, they should think that all cases of *p* are also *q*, and they should not allow any exceptions of the form *p, not q*. But if individuals instead represent *if p then q* as probabilistic, they should happily allow at least some *p, not q* exceptions. Experiments 7–9 investigated how individuals judge the truth or falsity of conditional assertions. Following the prediction just outlined, if individuals regard conditionals as deterministic, they should regard even a small number of *p, not q* counterexamples as rendering the conditional false, whereas if they regard conditionals as probabilistic, a few counterexamples should not be sufficient for a judgment of falsity.⁵

These three strategies—investigating how individual use conditionals, how they judge their meaning, and how they judge their truth—provide convergent methods to address the central theoretical issue at stake, which is whether individuals treat conditionals as probabilistic or deterministic in meaning.

Experiment 1: Selecting Conditional Descriptions

Experiment 1 presented participants with different patterns of data and asked them to select which of four alternative descriptions was most accurate in describing the data. The data patterns were either probabilistic or deterministic, and the question of interest was whether individuals would be more likely to select probabilistically qualified conditionals for probabilistic as opposed to deterministic patterns of data.

² Henceforth, by *probabilistic data*, I mean data reflecting a conditional probability that is strong relative to some salient reference point (usually, but not necessarily 0.5) but less than 1.

³ In this way, conditionals should be interpreted quite differently than generic statements, such as *Smokers get lung cancer*, which are typically interpreted as tolerating exceptions (see, e.g., [Cimpian, Brandone, & Gelman, 2010](#); [Leslie, Khemlani, & Glucksberg, 2011](#)).

⁴ Arguably, this phrase, *if p then probably q*, can refer to cases in which the conditional probability of *q*, given *p*, is less than 1, as well as cases where it is equal to 1. The main point for present purposes is that it can indeed refer to cases in which this conditional probability is less than 1.

⁵ Indeed, [Oaksford and Chater \(1998b\)](#) treated conditionals as being able to tolerate conditional probabilities of *q* given *p* as low as .875 (and possibly lower).

Materials and Method

Forty-one undergraduate students (gender not recorded) participated in the experiment for course credit. They each responded to six different problems. Each problem provided a short summary of some pattern of data. The data pertained to conditional probability relations holding between two variables. For instance, Problem 1 was of the following form (see Table 2), with the experimental manipulation indicated in square parentheses:

Say that you are told the following facts: A particularly nasty virus has been sweeping through Africa. After some investigation it turns out that [all/close to all (approximately 95%)] of the countries in which the virus has appeared have an annual per person income of less than \$500.

The nature of the conditional probability relation (probabilistic or deterministic) was varied within subjects. For half the participants, the first three problems contained a deterministic relation between the two variables in which the chance of the consequent holding given the antecedent was always 100%. That is, in the example above, *all* the countries in which the virus has appeared have an annual per person income of less than \$500. Correspondingly, for these participants, the last three problems contained a probabilistic relation. That is, in the example above, *close to all (approximately 95%)* of the coun-

tries in which the virus has appeared have an annual per person income of less than \$500. For the remaining participants, this order was reversed—the first three problems contained a probabilistic relation between the variables, and the last three contained a deterministic relation. For the probabilistic problems, the chance of the consequent holding given the antecedent varied in 5% increments between 70% and 95%, as Table 2 shows.

Following the initial description, participants were asked to select from a set of four conditional descriptions the one that best captured the relation. For Problem 1, they were instructed as follows:

Please indicate which of the following statements best describes this set of facts. You may only select one of the statements. Please read each of them carefully.

1. If a country has the virus, then it has an annual per person income of less than \$500.
2. If a country has the virus, then it probably has an annual per person income of less than \$500.
3. If a country has an annual per person income of less than \$500, then it has the virus.

Table 2

Descriptions of Data Patterns Used in Experiments 1 and 2, and the Percentages of Probabilistically Qualified Conditionals That the Participants Selected in Each Condition (Experiment 1), or Generated in Each Condition (Experiment 2)

Problem number	Problem contents	Experiment 1		Experiment 2	
		Probabilistic condition	Deterministic condition	Probabilistic condition	Deterministic condition
1	Say that you are told the following facts: A particularly nasty virus has been sweeping through Africa. After some investigation it turns out that [all/close to all (approximately 95%)] of the countries in which the virus has appeared have an annual per person income of less than \$500.	100	61	75	8
2	Say that you are told the following facts: A large study has been performed on the most effective therapies for depression. The results show that [all/a very large majority (approximately 90%)] of the most effective therapies for depression include some element of drug treatment.	100	50	82	0
3	You and a close friend are planning to go on vacation together. Your friend asks you to make a list of the places you would like to visit. It turns out that [all/a large majority (approximately 85%)] of the vacation spots you list are ones that your friend has already visited.	95	14	92	0
4	Each day when you go to work you park your car illegally in a spot that is closer to where you need to go. Most of the time, you do not get ticketed, however occasionally you do. You would like to see whether there is any pattern to when you get ticketed so you collect some data. It turns out that [all/a large majority (approximately 80%)] of the days you get ticketed are ones when it is raining.	100	36	100	10
5	Imagine that you are a smoker who is trying to give up. Although you try hard, there are still some occasions when you can't help having a cigarette. You want to try to find out what might cause you to have lapses, so you start observing your behavior carefully. It turns out that [all/a majority (approximately 75%)] of the times when you smoke are ones where you are anticipating giving a public talk.	100	25	100	22
6	Imagine that you are a researcher who is trying to work out why people commit various crimes. You commission a large national study which reveals that [all/a majority (approximately 70%)] of the individuals in your sample who committed violent crimes did not complete high school.	100	25	100	45

Note. In both cases, the means reflect the percentages only for conditionals that were expressed in the correct direction. The experimental manipulation is bolded and italicized here, but was displayed in normal font for participants.

4. If a country has an annual per person income of less than \$500, then it probably has the virus.

The four different response options were structured so as to capture four different sorts of conditional. As the example shows, one pair of conditionals expressed the correctly inferable relation between the variables (in the example above, Conditionals 1 and 2: if virus, then income less than \$500), whereas the other pair expressed the incorrect *reverse* relation (Conditionals 3 and 4: if income less than \$500, then virus). For each participant, the order of these pairs was counterbalanced, such that a participant would receive one order (i.e., the pair of conditionals expressing the correct direction first, and the pair expressing the incorrect direction second) for the entire set of problems within a particular block (probabilistic or deterministic) and the reverse order for the entire set of problems in the other block. As the example shows, within each of these pairs of conditionals, the first conditional was always unqualified, whereas the second one explicitly included the qualifier *probably*. The counterbalancing of the order of the response options within each block (correct vs. incorrect first) was fully crossed with the counterbalancing of the order of the deterministic and probabilistic blocks. Finally, the order of presentation of the six different problem contents was fixed and followed the order shown in Table 2. As in Table 2, the actual probabilities used for the probabilistic relations (varying between 70% and 95% in 5% increments) were fixed to the specific problem contents.

The main prediction was that when presented with probabilistic patterns of data, participants would select conditionals that explicitly included the qualifier *probably*—Option 2 above—whereas when presented with deterministic patterns of data, participants would select conditionals that did not explicitly include the word *probably*—Option 1 above. In contrast, conditional probability accounts suggest that there should be no preference between Options 1 and 2 regardless of the data pattern. According to these accounts, the meaning of Option 1 is inherently probabilistic—a country's having the virus only probabilistically implies that it has an annual per person income under \$500—and so its meaning does not differ from that of Option 2. Indeed, since Option 1 is a more economical description, it might even be preferred.

Results

Table 2 presents the means for the six problems for each of the two sorts of data pattern. As predicted, participants tended to select probabilistically qualified conditionals for probabilistic relations and unqualified conditionals for deterministic relations. They selected probabilistically qualified conditionals 99% of the time for probabilistic relations and only 41% of the time for deterministic relations (Wilcoxon test, $z = 5.23$, $p < .00001$). Moreover, the selection of probabilistically qualified conditionals for probabilistic conditionals was much greater than chance (Wilcoxon test, $z = 6.30$, $p < .000001$). Participants did not always select those conditionals that expressed the correct direction of the relation—they selected the *reverse* direction conditionals 19% of the time. Such selections were no more or less frequent for probabilistic than for deterministic data patterns, although their likelihood did differ depending on problem content. They were most frequent for Problems 4 (32%; parking tickets) and 5 (42%; smoking), which may be because the reverse conditionals in these cases would allow a useful prediction—for instance, a prediction about when

parking tickets and behavioral lapses—both events to be avoided—are most likely. The means in Table 2 include the selections of conditionals in the reverse direction, but the key results are unchanged when they are excluded (99% of probabilistically qualified conditionals for probabilistic data vs. 37% of probabilistically qualified conditionals for deterministic data, Wilcoxon test, $z = 5.08$, $p < .000001$).

The percentages of probabilistic descriptions chosen for deterministic problems were surprisingly high overall, but these percentages have a straightforward explanation. Although participants were instructed to select conditionals that applied only to the data in question, some participants may have interpreted their task as being to formulate a general description of the relation in question, including one that held with regard to unsampled instances. For some patterns of data, therefore, participants may have been skeptical that the deterministic relation that was asserted would hold up with more sampling. As Table 2 indicates, describing deterministic data in probabilistic terms was especially prevalent for Problems 1 and 2. These two problems refer to abstract data patterns pertaining to events outside of the scope of the participant's own personal life (viruses in Africa, treatments for depression), with the data being ostensibly provided to the participant (rather than having been collected by the participant themselves). In contrast, Problems 3–6 refer to data that, hypothetically, pertained directly to the participant's own life, with the data having ostensibly been collected by the participant him- or herself. This difference in the abstractness of the conditionals may account in part for the observed difference in probabilistic descriptions.

One potential concern with these results stems from the within-subjects design. Because participants were exposed to data sets that were both probabilistic and deterministic, they may have been cued to adjust their responses in line with the main hypothesis. However, this objection can be circumvented by examining only the first three problems that participants responded to, since these problems exclusively described either probabilistic or deterministic relations. The results for these first three problems mirrored the overall results. Participants selected probabilistically qualified conditionals 98% of the time for the probabilistic relations and 47% of the time for the deterministic relations (Wilcoxon test, $z = 5.23$, $p < .00001$), and this result also held when excluding reverse conditional selections (98% vs. 44%, respectively; Wilcoxon test, $z = 5.14$, $p < .00001$).

Discussion

These results suggest that individuals do indeed prefer explicitly qualified conditionals when describing probabilistic relations. In contrast, they prefer conditionals that contain no probabilistic qualification when describing deterministic relations (although there was some inconsistency across problem contents in this respect). Consistent with the accounts offered by the mental model theory, formal rules theories, and the suppositional theory and contrary to explicitly probabilistic theories, these results suggest that the basic conditional is not treated as probabilistic.

Two potential problems arise with this interpretation, however. First, the four response options differed in whether they explicitly included the word *probably*. This difference may have cued participants to select the response option that best matched the description of the data. For the probabilistic relations, participants

may have been quite satisfied with the *if p then q* framing but nevertheless have selected the *if p then probably q* framing in order to best match the initial description. Second, it is possible that, if given the opportunity, participants would have further qualified the conditionals that described the deterministic relations. For instance, they may have preferred conditionals with modal terms such as *must*, *always*, *necessarily*, or some other term (e.g., *strictly speaking*) that explicitly conveyed the conditional probability of 1.

Experiment 2: Generating Conditional Descriptions

Experiment 2 was designed to deal with both of these potential problems and to investigate whether the main effect would replicate using a different methodology. Participants were presented with the same materials as in the previous experiment, but this time, they were asked to generate a conditional description of the data in their own words.

Materials and Method

Thirty-nine undergraduate students (gender not recorded) participated in the experiment for course credit. They were presented with the same six problems as in the first experiment, but instead of selecting a conditional description, this time they were instructed, "How would you describe this set of facts using an IF-THEN statement?" and were asked to respond in an open-ended way. As in the first experiment, for half the participants, the first three problems were deterministic and the second three probabilistic, whereas, for the other half of the participants, this assignment was reversed. The probabilities varied and were fixed to the different statement contents in the same way as in Experiment 1. The problems were presented in the same fixed order as in Experiment 1.

Results

Table 2 presents the means for the six different problems. Participants' responses were coded in terms of whether they explicitly included probabilistic terminology such as *probably*, *usually*, and *likely that* and in terms of whether the conditionals were framed in the correct direction. Five descriptions were excluded from the data analysis (2% of the total number of descriptions) because they introduced additional information and as a result were not easily coded. Replicating the first experiment, participants wrote descriptions that explicitly included probabilistic qualification much more often for the probabilistic data patterns than for the deterministic data patterns (83% vs. 25%, respectively; Wilcoxon test, $z = 4.82$, $p < .00001$). And they were more likely to insert probabilistic qualification for the probabilistic data patterns than chance would dictate (Wilcoxon test, $z = 4.99$, $p < .00001$). There was, again, some tendency to frame conditional descriptions in the incorrect, *reverse* direction, which was quite high in this experiment: Thirty-nine percent of all conditional descriptions were framed in the incorrect, *reverse* direction. However, the previous results held when excluding these incorrect framings: Probabilistically qualified conditionals were much more common for probabilistic data than for deterministic data (89% vs. 20%, respectively; Wilcoxon test, $z = 3.68$, $p < .001$) and were more likely than chance (Wilcoxon test, $z = 4.41$, $p < .00001$).

The within-subjects design again poses a potential problem—because participants were exposed to both probabilistic and deterministic data patterns, they may have been cued to adjust their responses in the direction of the hypothesis. However, this possibility cannot explain the results, which held just as strongly when considering only the first three problems. Probabilistically qualified conditionals were more common for probabilistic data than for deterministic data (76% vs. 10%, Mann-Whitney U, $z = 5.11$, $p < .001$), they were more common than chance (Wilcoxon test, $z = 3.01$, $p < .01$), and both of these results held when excluding the incorrect, *reverse* conditionals (81%, probabilistic data, vs. 6%, deterministic data; Mann-Whitney U, $z = 4.60$, $p < .001$; probabilistic qualification more likely than chance, Wilcoxon test, $z = 2.80$, $p < .01$).

One advantage of the present method over that used in Experiment 1 is that it allows an examination of whether participants qualified their responses with vocabulary specifically designed to indicate a deterministic relation. For deterministic data, individuals might choose to insert modal or arithmetic terms such as *always*, *must*, *necessarily*, *all instances*, *100%*, *probability of 1*, *without exception*, or other language designed to convey an exceptionless reading, for example, *strictly speaking*. This explicit qualification would count against the hypothesis that the basic conditional expresses a deterministic relation—if the conditional is treated as deterministic, there should be no need to explicitly qualify conditionals with deterministic or modal vocabulary. However, consistent with the deterministic meaning hypothesis, this sort of qualification virtually never occurred. Out of 229 total responses to deterministic data patterns across the six problems, only four (i.e., less than 2%) explicitly included any modal or deterministic vocabulary, strongly suggesting that the default meaning of a basic conditional is in fact taken to be deterministic.

Discussion

These results further support the idea that the meaning of the basic conditional is not probabilistic. In an open-ended response format, when describing probabilistic patterns of data, participants explicitly qualified conditional descriptions with probabilistic terminology.

One further possible objection to these findings is that participants were responding to an experimental demand to provide maximal accuracy in their descriptions (Nikki Pfeifer, personal communication, August 19, 2008). Thus, when presented with a probabilistic data pattern, participants may have felt that their task was to provide as much information about the relation as possible, such that they chose to qualify their description probabilistically, notwithstanding the fact that they would have found the basic *if p then q* construction a perfectly adequate description. However, this objection is undermined by the fact that participants did not add extra qualifications for the deterministic data patterns. According to this maximal accuracy hypothesis, they ought equally to have qualified their descriptions of the deterministic data patterns with some extra modal or deterministic terminology. But, as the results showed, they did not do this. Instead, the straightforward *if p then q* construction was used the great majority of the time to capture deterministic relations.

Perhaps, however, conditionals that expressed deterministic data patterns were not explicitly qualified because specific numeric

values were used only for the probabilistic data patterns, in contrast with the deterministic data patterns, for which only the word *all* was used (see Table 2). This objection seems implausible, but the data from Experiment 2 cannot rule it out conclusively.

Experiment 3: Conditional Descriptions of Sets of Possibilities

Experiment 3 aimed to extend the effects of the previous experiments to a new sort of task format in which participants were presented with sets of raw data and asked to describe them. In so doing, the experiment aimed to rule out the linguistic objection just described—here, the data sets for the deterministic and probabilistic relations did not differ in how they were linguistically framed.

Materials and Method

Thirty-seven undergraduate participants (13 male, 24 female) participated in the experiment for course credit. They each carried out 20 problems in total, which differed in content from those used in the first two experiments and which are shown in Appendix A.

The basic task was the same as in Experiment 2, but the nature of the stimuli that participants described was different. In this experiment, participants were presented with sets of instances, or cases, of a probabilistic or deterministic relation between two events and were asked to describe this relation using a conditional assertion. For instance, they were presented with a description of the following sort:

You are interested why a certain beach gets closed. In particular, you are interested in whether the days when the beach is closed are bad weather days. You have examined a series of days when the beach was closed, shown below. Each row represents a different day.

Beach closed	Bad weather
Beach closed	Not bad weather

Participants were then asked to complete the following task:

On the basis of this information, you can now make an inference about the relationship between the beach being closed, and bad weather. Please complete the following IF-THEN statement, which describes what you know about this relationship. If the beach is closed, then:

This example conveys a probabilistic relation between beach closures and bad weather—the beach being closed tends to coincide with bad weather, but this relation is not invariable. Deterministic relations differed simply by making all of the cases identical, such that the relation was invariable. As the example shows, in coming up with their descriptions, participants were presented with the antecedent of the desired conditional and asked to provide only the consequent, which prevented conditionals being described in the reverse direction. The question of interest was whether individuals would choose to explicitly qualify the consequents of the conditionals—either probabilistically for probabilistic data patterns or deterministically for deterministic data patterns.

As in the previous experiment, participants carried out two separate blocks of problems. For half the participants, the first block contained 10 probabilistic problems, and the second block contained 10 deterministic problems. For the remaining participants, this ordering of blocks was reversed. In the interests of generalizability, within each block, half the items were designed so that participants would have no strong prior beliefs about the conditional relation between the events (e.g., “If there is a mop in the closet, then there is a wrench in the closet”), and the remaining half were designed so that participants would have some prior belief about the conditional (e.g., “If the beach is closed, then it is a bad weather day”).

The experiment was computerized, using the Eprime program. Twenty different problem contents were designed, as shown in Appendix A. The 20 different contents were presented in four overall versions of the experiment. The different versions counterbalanced the order of the materials presented and the assignment of probabilistic relations to materials. Version 1 started with the materials ordered from 1 to 20, as in Appendix A. Different probabilities were assigned to these materials, as shown in the Version 1 column in Appendix A. The first 10 problems were probabilistic. The number of instances used ranged from five to 12, and the actual probabilities assigned ranged from 3/5 (Problem 1) to 11/12 (Problem 10). The aim in doing this was to vary both the probabilities and the number of instances each problem consisted of. The second 10 problems were deterministic. For this set of problems, the data were all changed so that the conditional probabilities were 1, but the actual number of instances used for each problem was the same as in the first set of 10 problems. Version 2 differed by swapping the set of assigned probabilities from the first block to the second block and vice versa. Versions 3 and 4 simply reversed the order of assignment of probabilities to problem contents. The problems in the first block of Version 3 were probabilistic, and the problems in the second block were deterministic, with this order reversed for Version 4. Within each block, problem order was randomly generated by the computer program.

Results

Participants’ responses were coded in the same way as in the previous experiment. That is, they were coded as probabilistic if they explicitly included terms such as *probably*, *usually*, *likely*, *most of the time*, and variants on these phrasings. Responses were also coded as probabilistic if they were explicitly disjunctive, that is, *if p then q or r* (since such responses were typically used to indicate some uncertainty about the occurrence of the consequent); if they included a specific probability value less than 1; or if they used the word *can*. Responses were coded as deterministic if they explicitly included terms such as *always*, *must*, *all of the time*, *necessarily*, *should*, *strictly speaking*, or the explicit probability value of 1.

A preliminary analysis showed that there was no difference in the percentage of probabilistically qualified conditionals expressed for the arbitrary and realistic problem contents, and so, this variable was collapsed for the remaining analyses. As in the previous experiments, probabilistically qualified conditionals were expressed far more often for probabilistic than for deterministic relations (96% vs. 38%, respectively; Wilcoxon test, $z = 4.96$, $p < .00001$). For probabilistic relations, this percentage was reliably

greater than chance (Wilcoxon test, $z = 5.67$, $p < .00001$). This percentage might have been artificially inflated, owing to the within-subjects design. Again, however, this possibility is ruled out by an analysis of the first block of problems. For these problems, probabilistically qualified conditionals were again far more frequent for probabilistic than deterministic relations (96% vs. 33%, respectively; Mann-Whitney U-test, $z = 4.57$, $p < .00001$) and were again more frequent than chance (Wilcoxon test, $z = 4.57$, $p < .00001$).

In contrast, there were very few instances of explicit deterministic or modal qualification when participants described deterministic relations. For deterministic relations, these sorts of qualification occurred only 4% of the time. The most pertinent comparison is between the percentage of deterministic qualifications for deterministic relations and the percentage of probabilistic qualifications for probabilistic relations. The data show clearly that deterministic qualification was markedly less frequent than probabilistic qualification (4% vs. 96%, respectively; Wilcoxon test, $z = 5.50$, $p < .00001$), which strongly suggests that participants' default interpretation of the conditional was deterministic.

Discussion

Consistent with the results of the previous experiments, these results strongly support the idea that the basic conditional is taken to express a deterministic relation. When describing probabilistic data, participants almost always took the extra step of qualifying their descriptions with probabilistic terminology. It is hard to understand why they would do this if the *if p then q* construction adequately expresses a probabilistic relation. On the other hand, when describing deterministic relations, participants virtually never took the extra step of qualifying their descriptions with deterministic or modal terminology. This is also hard to reconcile with a probabilistic reading of the conditional. If the *if p then q* construction is capable of expressing probabilistic relations, there might seem to be an extra need to qualify the use of this construction for deterministic relations, which involve a special kind of probability. Instead, it seems that people treated *if p then q* as adequate only to describe deterministic relations and to be able to do so without any further adornments.

These data go beyond those from the previous experiments in three ways. First, they show that the results generalize to a different kind of task, in which participants formulated a conditional description from instances or cases rather than paraphrasing a description of data—a task that is somewhat more inferential than that used in the earlier experiments. Second, these data refute the objection that the difference seen in previous experiments owed to the fact that explicit, numeric descriptions were used for probabilistic relations, whereas linguistic descriptions (e.g., *all*) were used for deterministic relations. No objection of this sort can explain the large differences in the present study because the data were always presented as instances. Third, these data also generalize the previous findings to more realistic (rather than purely arbitrary) conditional relations.

Experiment 4: Conditional Descriptions of Both General and Specific Conditionals

The conditionals that subjects described in the previous experiment were all, arguably, general conditionals. That is, participants

were asked to formulate a conditional that expressed a relation holding over an entire set of events or objects, as in the following example, which applies to the entire set of cards in a specified set:

If a card contains a letter on one side, then it contains a number on the other side.

This sort of conditional differs from one that expresses a relation that holds only with respect to a single entity or event (which I refer to as a *specific conditional*). For instance, a person might choose to select a specific card from a set, and assert with respect to this specific card:

If the card contains a letter on one side, then it contains a number on the other side.

This more specific conditional applies only to a specific entity rather than an entire set of entities, and it uses the definite article *the* to pick out that entity (rather than the indefinite article *a*). This difference between general and specific conditionals is potentially important and has been emphasized by some theorists (see, e.g., Evans et al., 2003; Sperber, Cara, & Girotto, 1995). And, as mentioned in the introduction, the suppositional account pertains only to specific conditionals (see, e.g., Evans et al., 2003, p. 333; David Over, personal communication, March 13, 2009). It therefore seemed pertinent to extend the present investigation to both general and specific conditionals, in order to examine how participants interpret both sorts of conditional. The experiment again varied the specific methodology used for the sake of generality.

Materials and Method

Sixty-nine undergraduate students (28 male, 41 female) participated in the experiment for course credit. They each carried out 16 problems in total, which differed in content from those used in the first three experiments and which are shown in Appendix B. The basic task was the same as in the previous experiments, but the nature of the stimuli that participants described was different. This time, individuals were presented with summary statistics on the frequency of each of the four possibilities relevant to a conditional of the form *if p then q*, that is, the following four possibilities—*p and q*, *p and not q*, *not p and q*, and *not p and not q*—and were asked to formulate the consequent of a conditional assertion. For instance, one problem had the following form:

Imagine that you are looking at a wine list that has 100 different wines that differ in terms of their country of origin and their type (red or white). The wine list includes the following numbers of different kinds of wines (each row represents a different kind of wine):

Italian	red	9
Italian	white	1
French	red	40
French	white	10

Suppose that you want to describe to someone the overall relation between the country of origin and the type (red or white) of wines on this wine list. Please complete the following IF-THEN statement, which describes what you know about this relation. If a wine is Italian, then _____.

Participants typed their responses into a space provided in a program implemented on Eprime.

For each problem, there were 50 *p* cases (i.e., Italian wines in the example above) and 50 *not p* cases (i.e., French wines in the

example above). The key manipulation was the ratio of *p* and *q* cases to *p* and *not q* cases. For one block of eight problems, the conditional probability of *q* given *p* (i.e., in the case above, the probability of red given Italian) was high (above .8 in each case) but not 1. For a second block of problems with the same content, the conditional probability of *q* given *p* was always 1. The main question of interest was how participants chose to describe the consequent for the probabilistic problems as compared with the deterministic problems. The order of the deterministic and probabilistic blocks was counterbalanced within subjects, and the order of the problems within each block was randomized. The frequencies of the *not p* and *q* and the *not p* and *not q* cases were also varied from problem to problem, although this was merely an incidental feature.

A further between-subjects variable was whether the conditionals were framed as general or specific. The example above is general since it talks about describing the *overall relation* and uses the indefinite article *a* in the antecedent of the conditional. Half the participants received conditionals of this sort in both blocks. The remaining participants received specific conditionals, which had the same introductory wording but which differed in the wording of the main part of the problem. The wording of the specific conditionals focused on a specific card drawn at random and used the definite article *the*; for example:

Suppose that you randomly select a single wine from this wine list, and that you want to describe to someone the relation between its country of origin and its type (red or white). Please complete the following IF-THEN statement, which describes what you know about this relation. If the wine is Italian, then:

The actual frequencies used were identical across the general and specific blocks.

Results and Discussion

Responses were coded in the same way as previously. For both general conditionals and more specific conditionals, participants tended to insert probabilistic qualification into the consequents of conditionals that described probabilistic data, and they did so at a rate that was much higher than chance (general conditionals, 89%; specific conditionals, 90%; Wilcoxon tests, $z_s = 4.47, 4.79$, $p_s < .00001$, respectively). By contrast, they tended not to insert probabilistic qualification into conditionals that described deterministic data (10% of the time on average). Nor did they tend to insert modal or other deterministic qualification into their descriptions of deterministic data patterns. Such insertions happened only 12% of the time for general conditionals and 11% of the time for specific conditionals. Overall, the tendency to insert probabilistic qualification into conditionals that described probabilistic data was much greater than the tendency to insert modal or deterministic qualification into conditionals that described deterministic data (Wilcoxon test, $z = 7.12$, $p < .00001$).

The same pattern held when looking only at the first block of problems that subjects were exposed to. Individuals inserted probabilistic qualification into conditionals describing probabilistic data at a rate that was much greater than chance (88%, Wilcoxon test, $z = 4.11$, $p < .0001$) but tended not to insert modal or deterministic language into conditionals describing deterministic data (7%). Nor did participants typically insert probabilistic qual-

ification into conditionals describing deterministic data (14%). The tendency to insert probabilistic qualification into conditionals that described probabilistic data was thus much more pronounced than the tendency to insert deterministic qualification into conditionals describing deterministic data (Wilcoxon test, $z = 6.81$, $p < .00001$).

The present experiment is thus, once again, consistent with the previous experiments. It supports the idea that the basic meaning of the conditional is deterministic and that extra qualification is required to convey a probabilistic meaning. And it demonstrates this point with regard to both general and specific conditionals.

Experiment 5: Judging the Meaning of Conditional Assertions

All of the experiments thus far examined how individuals use conditionals to describe various patterns of data. These experiments showed that individuals use an unqualified *if p then q* construction to express deterministic data but that they use probabilistic qualification to express probabilistic data. Together, this evidence suggests that individuals' default representation of conditional assertions is deterministic. If this is correct, it should have consequences for how individuals make judgments about the meaning of a conditional assertion—that is, judgments about what possibilities it gives rise to. Specifically, given a conditional, *if p then q*, a deterministic interpretation would mean that no *p* and *not q* cases are possible, whereas a probabilistic interpretation would allow for the possibility of such cases. Because little existing evidence appears to have directly addressed this question, Experiment 5 aimed to do so.

The use of this more direct method also helps address a potential criticism of the initial studies. According to this criticism, participants may in fact regard the *if p then q* construction as adequate to capture probabilistic relations, yet in their normal linguistic usage, they tend to reserve this construction for special cases only—that is, cases in which the conditional probability of the consequent given the antecedent is 1.⁶ This sort of interpretation would be consistent with all of the prior data on how people use conditionals to describe relations. And, arguably, the only way to refute it would be to ask participants directly about the strength of the relation between antecedent and consequent, which is what Experiment 5 did.

Method

Fifty-five undergraduate students (24 male, 31 female) participated in the experiment for course credit. Individuals were presented with conditionals and were asked to indicate the proportion of cases in which the consequent, or its negation, occurred given that the antecedent occurred. The same problem set as was used in Experiment 4 was also used here, except that participants were not shown any data pattern. Instead, they were presented with a conditional description of a possible data pattern, either in an unqualified *if p then q* form or alternatively in the probabilistically qualified *if p then probably q* form. Following the statement of each conditional, participants were then asked one of two ques-

⁶ I thank an anonymous reviewer for this suggestion.

tions. For some problems, they were asked to indicate, given that the conditional is true, what proportion of the p cases are q , while for other problems they were asked to indicate what proportion of the p cases are *not* q . They responded on a scale that included the following five response options: *none, a small proportion, a moderate proportion, a large proportion, and all*.

Assuming that participants treat the conditional as deterministic, the main prediction was that when presented with an unqualified *if p then q* conditional, participants would indicate that all of the p cases are q and, correspondingly, that none of the p cases are *not* q . On the other hand, there is no reason why participants should make this response if they treat the conditional as probabilistic. In that case, a more accurate (and more conservative) response would be to indicate that most (or at least some reasonable proportion) of the p cases are q and that few of the p cases (but not necessarily none) are *not* q . Both general conditionals and specific conditionals were used in this experiment.

An illustrative general conditional problem was as follows:

Imagine that another person is presented with a box filled with 100 balls that differ in terms of their size and their color. The balls in the box are either large or small, and they are either green or blue. Suppose that this person describes the overall relation between the size and the color of balls within the box with the following statement:

Following this description, the probabilistically qualified conditional was:

If a ball is large, then it is probably green.

whereas the unqualified conditional was:

If a ball is large, then it is green.

In the consistent (p , what proportion of q) cases, participants were then asked:

Given that this statement is TRUE, what proportion of LARGE BALLS in the box are GREEN?

whereas in the inconsistent (p , what proportion of *not* q) cases, participants were asked:

Given that this statement is TRUE, what proportion of LARGE BALLS in the box are BLUE?

The corresponding specific version of this problem differed in its preamble, which described the selection of a specific entity from the box, and in terms of the conditional itself, which used the definite article *the*, as follows:

Suppose that this person is going to randomly select a single ball from this box. Before selecting the ball, this person describes the relation between its size and its color with the following statement [the terms in brackets were not shown to participants]:

[Probabilistically qualified conditional] If the ball is large, then it is probably green.

[Unqualified conditional] If the ball is large, then it is green.

The questions that were asked of participants for the specific conditionals were the same as those asked for the general conditionals.

Each participant received 32 problems in total, 16 of which used conditionals with probabilistic qualification inserted and 16 of which used unqualified conditionals. Within each of those subsets, there were eight distinct problem contents. Each problem content was presented twice within each block—once where participants were asked to indicate what proportion of p cases are q and once where they were asked what proportion of p cases are *not* q . Approximately half of the participants ($n = 28$) were randomly assigned to receive only general conditionals, whereas the other half of participants ($n = 27$) received the corresponding set of specific conditionals.

Results and Discussion

There were very few differences as a function of problem content, and so, the main analyses below aggregate across content. Across all of the analyses, whether or not the conditionals were specific or general had no effect on participants' responses and did not interact with whether or not the conditionals were probabilistically qualified. I therefore collapsed across this variable (general vs. specific) as well. As Figure 1 illustrates, when asked what proportion of p cases are q , participants provided higher responses on the 5-point scale (1 = none, 5 = all) when interpreting unqualified conditionals than when interpreting probabilistically qualified conditionals (4.62 vs. 3.93, Wilcoxon test, $z = 6.10$, $p < .00001$). Similarly, participants provided lower responses for unqualified conditionals when asked about the proportion of p cases that are *not* q (1.42 vs. 2.09, $z = 5.98$, $p < .00001$). Neither of these differences is predicted by probabilistic accounts of the meaning of conditionals, but they are directly predicted by deterministic accounts. Moreover, both of these differences held when considering only the first block of problems that participants encountered (p , q cases, unqualified: 4.79 vs. probabilistically qualified: 3.9, Mann-Whitney test, $z = 5.70$, $p < .00001$; p , *not* q cases, unqualified: 1.24 vs. probabilistically qualified: 2.11, Mann-Whitney test, $z = 5.36$, $p < .00001$).

A more stringent test of the deterministic hypothesis is to consider only those responses in which participants indicated that all of the p cases are q and, conversely, that none of the p cases are *not* q . Indeed, for unqualified conditionals, such responses were highly prevalent, as Figures 1c and 1d illustrate. The modal response (68%) for unqualified conditionals was that all of the p cases are q . This percentage is substantially higher than the corresponding percentage of *all* responses for probabilistically qualified conditionals (1%, $z = 5.83$, $p < .00001$) and is reliably higher than the most conservative estimate of chance (50%, $z = 2.94$, $p < .01$). Similarly, for the complementary, p and *not* q question, the modal response (68%) for unqualified conditionals was that none of the p cases are *not* q . This was again much higher than the corresponding percentage of *none* responses for probabilistically qualified conditionals (1%, Wilcoxon test, $z = 5.76$, $p < .00001$) and was also higher than 50% (Wilcoxon test, $z = 2.72$, $p < .01$). This tendency was even more pronounced when considering participants' responses to just the first block of problems that they encountered. Participants reported that all p cases were q 86% of the time for unqualified conditionals compared with 0% for the probabilistically qualified conditionals (Wilcoxon test, $z = 7.16$, $p < .00001$), and they reported that no p cases were *not* q 88% of the time for unqualified conditionals compared with .03% for

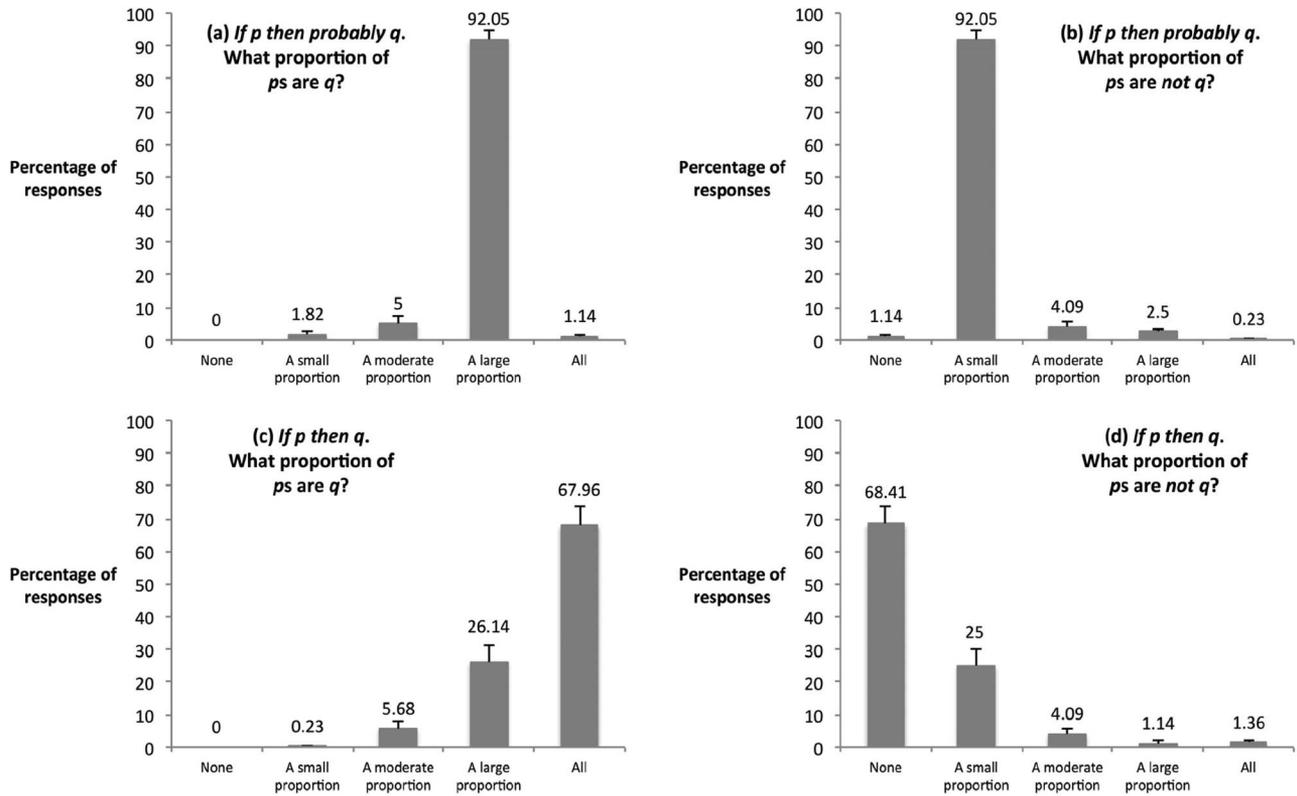


Figure 1. Responses to the four different questions in Experiment 5, shown separately for probabilistically qualified conditionals (Panels a and b) and for unqualified conditionals (Panels c and d). Error bars represent standard errors.

probabilistically qualified conditionals (Wilcoxon test, $z = 7.05$, $p < .00001$).

Thus, Study 5 shows that although the tendency to regard conditionals as deterministic was not universal, it was clearly the dominant response tendency, contrary to probabilistic accounts of the meaning of conditionals. This evidence constitutes the most direct evidence that people interpret the conditional deterministically. And it appears to refute pragmatic accounts of the earlier data (i.e., that *if p then q* can be used to express probabilistic relations, even if it is not typically used in this way).

A skeptic might perhaps object that people are still relying on some special default sense of the *if p then q* construction, such that it is reserved especially for cases in which the conditional probability of the consequent given the antecedent is 1, even though it can still adequately describe probabilistic relations. Given that the question in the present study was simply to indicate what percentages of p and q or p and not q actually obtained, perhaps people simply relied on this default usage when responding.⁷ A still more stringent test would be to ask people questions that have a wider scope, that is, about what percentages of p and q must obtain and what percentages of p and not q can obtain given a conditional, *if p then q*. If people once again tend to provide a deterministic reading of the conditional even in response to this wider scope question, it would rule out the idea that the *if p then q* construction is usually reserved for deterministic relations but can just as well

describe probabilistic relations. Accordingly, Experiment 6 examined this.

Experiment 6: Judging What Is Possible and What Must Obtain Given a Conditional Assertion

Method

Four hundred and eight participants (286 male, 121 female, one unspecified) completed the experiment on Amazon.com's Mechanical Turk system, in exchange for a small payment. An additional 15 participants began the survey but did not complete it, and so, their partial data were removed from the analyses.

The materials from Experiment 5 were once again used here, although on this occasion the design of the experiment was between-subjects. Each participant received eight problems corresponding to the content displayed in Appendix B. Half of the participants ($n = 204$) received the unqualified *if p then q* construction for all problems, while the other half of the participants ($n = 204$) received a probabilistically qualified formulation, *if p then probably q*. Orthogonal from this manipulation, for each problem, approximately half of the participants ($n = 209$) judged the proportion of ps that must be q , given that the conditional is

⁷ I thank an anonymous reviewer for raising this objection.

true. In its general form, the question was “Given that this statement is true, what proportion of ps must be q (in order for the statement to remain true)?” The remaining participants ($n = 199$) judged the proportion of ps that could be *not* q with the conditional still remaining true. In its general form, the question was “Given that this statement is true, what is the largest proportion of ps that can be *not* q (with the statement still remaining true)?” Finally, orthogonally crossed with these two manipulations was a manipulation that concerned the scope (general vs. specific) of the conditional assertions—whether they applied generally to the overall relation between antecedent and consequent ($n = 205$) or instead to a specific entity that would be selected at random from the overall set ($n = 203$). Thus, eight between-subjects conditions resulted from the crossing of these three independent variables, which participants were assigned to at random.

Results

As in Experiment 5, there were very few differences as a function of problem content, and so, the main analyses below aggregate across content. Similarly, whether or not the conditionals were framed in general or specific terms had very little effect on responses—neither a main effect nor any interactive effect—and so the analyses collapse across this variable as well.

Despite the change in the wording of the questions, the same general picture emerged from this experiment, such that participants once again interpreted the conditional as deterministic. As

Figure 2 shows, when asked what proportion of p cases must be q , participants provided higher responses on the 5-point scale (1 = none, 5 = all) when interpreting unqualified conditionals than when interpreting probabilistically qualified conditionals (4.52 vs. 3.72, Wilcoxon test, $z = 8.68$, $p < .00001$). Similarly, participants provided lower responses for unqualified conditionals when asked about the proportion of p cases that can be *not* q (1.40 vs. 2.24, $z = 8.79$, $p < .00001$).

The more stringent test is to examine the percentage of responses for which people thought that *all* of the ps must be q and, conversely, that *none* of the ps can be *not* q . These analyses provide even more decisive support for the idea that people interpret conditionals as deterministic. For general conditionals, when asked what proportion of ps must be q , participants gave the deterministic response—that all ps must be q —81% of the time for the unqualified *if p then q* construction, compared with only 5% of the time for the qualified *if p then probably q* construction (Wilcoxon test, $z = 7.88$, $p < .00001$). The deterministic response for *if p then q* was given more often than a conservative 50% estimate of chance (Wilcoxon test, $z = 4.69$, $p < .00001$). Conversely, when asked what proportion of ps could be *not* q , participants gave the deterministic response—that no ps could be *not* q —81% of the time for *if p then q* , compared with 5% of the time for *if p then probably q* (Wilcoxon test, $z = 8.14$, $p < .00001$). Once again, the deterministic response was given more than 50% of the time (Wilcoxon test, $z = 4.74$, $p < .00001$).

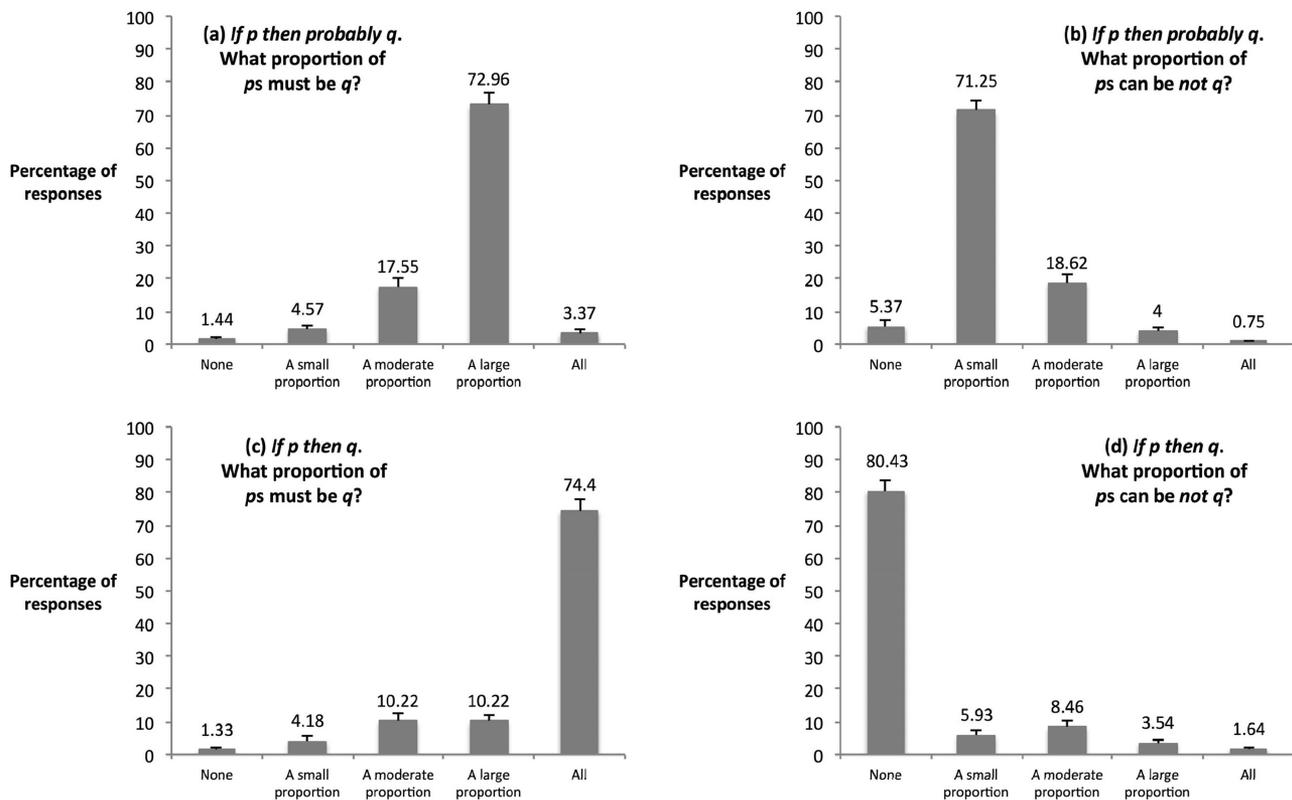


Figure 2. Responses to the four different questions in Experiment 6, shown separately for probabilistically qualified conditionals (Panels a and b) and for unqualified conditionals (Panels c and d). Error bars represent standard errors.

For specific conditionals, the rate of deterministic responses was slightly lower but still very high. When asked what proportion of *ps* must be *q*, participants said that all *ps* must be *q* 68% of the time for the unqualified construction, *if p then q*, compared with only 2% of the time for the qualified construction, *if p then probably q* (Wilcoxon text, $z = 7.37, p < .00001$). The deterministic response was given more than half the time (Wilcoxon text, $z = 2.66, p = .008$). And when asked what proportion of *ps* could be *not q*, participants said that no *ps* could be *not q* 80% of the time for *if p then q*, compared with 6% of the time for *if p then probably q* (Wilcoxon text, $z = 8.00, p < .00001$). Once again, the deterministic response was given more than 50% of the time (Wilcoxon text, $z = 4.57, p < .00001$). Figure 2 displays these results aggregated across general and specific conditionals.

These results thus provide further support for the deterministic reading of conditional assertions. Moreover, because of the wide scope of the questions asked in this study—regarding what is possible and what must be the case, given that a conditional is true—it rules out an objection to the interpretation of the results of Experiment 5, namely, that *if p then q* is conventionally used for deterministic relations but also adequately captures probabilistic relations. If that were true, participants should have been quite happy to admit the possibility of exceptions in Experiment 6 when explicitly asked about them. Yet, as the results show, they were not.

Experiment 7: Judgments of the Truth of General Conditionals

The final three experiments examined a related way of assessing the meaning of conditionals. If conditional assertions are interpreted as deterministic, this should have consequences for how individuals judge the truth of such assertions. To investigate this, I focused on conditionals where the probability of the consequent obtaining given the antecedent was high but not certain. Will individuals judge such conditionals as true, false, or neither? On the one hand, if it is indeed the case that conditionals are treated as deterministic and not probabilistic, a single counterexample should suffice to render them false, and so deterministic theories predict that in cases of high but not certain conditional probability, the *false* response should be most likely. It is less clear exactly what probabilistic accounts predict—they might either predict the *true* response or, alternatively, the *neither* response—but it is at least clear that they do not predict the *false* response. A single counterexample, or a few counterexamples, is not sufficient to render a conditional false according to such probabilistic accounts. Experiment 7 focused solely on general conditionals.

Method

Thirty-two participants (15 male, 17 female) were recruited using Amazon.com's Mechanical Turk system. They each performed three problems. For each one, they were asked to suppose that they were in a restaurant looking at a wine list, which contained either Italian or French wines, which could be either red or white. The task was to judge whether an assertion that described this wine list was either true, false, or neither true nor false. The key instructions were as follows:

The wine list that you are reading lists 100 different wines. Each wine is either from Italy or France, and is either red or white. The following numbers represent the overall frequencies of the different kinds of wine that are available:

Italian	red	5
Italian	white	5
French	red	20
French	white	30
TOTAL		100

Imagine that someone describes the selection of wines that are available with the following statement:

"If a wine is Italian, then it is red".

Given the selection of available wines, please indicate whether you think this statement is true, false, or neither true nor false.

Subjects received each of three different versions of the problem, which varied only in terms of the frequencies of Italian wines. The high conditional probability version, which was always presented first, exactly matched the preceding description. The deterministic version, which was always presented second, had 50 Italian red wines and zero Italian white wines, and the low conditional probability version, presented last, had 30 Italian red wines and 20 Italian white wines. According to the deterministic interpretation hypothesis, the prediction is straightforward: For both the high and low conditional probability versions, participants should judge the conditional to be false, whereas they should judge it to be true in the deterministic version. In contrast, at least some conditional probability accounts predict that participants will judge both the high and low conditional probability versions of the conditional as being either true or as neither true nor false.

Results and Discussion

The deterministic prediction was strongly reflected in the results. Twenty-four participants (75%) judged the high conditional probability version as false, three (9%) judged it as true, and five (16%) judged it as neither true nor false, which is a reliably nonuniform distribution, $\chi^2(2) = 25.19, p < .0001$. Similarly, 25 participants (78%) judged the low conditional probability version as false, one (3%) judged it as true, and six (19%) judged it as neither true nor false, which is also reliably nonuniform, $\chi^2(2) = 30.06, p < .0001$. Notably, the number of false responses did not differ between the high and low conditional probability conditionals (Sign test, $z = .58, p > .5$), which is at odds with the predictions of probabilistic accounts. In contrast to the high and low probability conditionals, the deterministic conditional was generally judged true: Twenty-three participants (72%) judged it as true, eight (25%) judged it as false, and one (3%) judged it as neither true nor false, again a reliably nonuniform distribution, $\chi^2(2) = 23.69, p < .0001$.

Thus, even in a case where the conditional expresses a very high conditional probability, the majority of participants judged the conditional as false. It is hard to understand why they would do this if the meaning of the conditional is probabilistic. But it accords with the view that individuals treat the meaning of a basic conditional as deterministic.

Experiment 8: Judgments of Truth and Falsity for General and Specific Conditionals

Experiment 8 tested the same prediction for both general and specific conditionals.

Method

One hundred seventy-three undergraduate students (55 male, 118 female) participated for course credit as part of a larger experimental session. Each participant received two problems. One of the problems was the same as in the previous experiment and described a set of 100 wines on a wine list in the same fashion as previously. Frequency information was provided for each of the four relevant possibilities. The second problem described a set of 100 meals that were available on a restaurant's menu. The meals either contained meat or were vegetarian, and they were either spicy or mild. Participants were again presented with frequency information on each of the four relevant possibilities.

All participants received each of two different sorts of conditional, a general one and a specific one. The general one had the same form as in the previous experiments, such that the conditional described the overall set of entities (wines or dishes) available. In contrast, the specific conditional pertained only to one specific entity that would be picked out at random. Thus, for the meals problem, participants were instructed as follows:

As part of a game with friends, you decide that one of them will choose your dish completely at random from the menu. Imagine that before this selection takes place, someone describes the dish you will end up with using the following statement:

"If the dish is meat, then it will be spicy".

The full wording of the problems is available upon request. For both sorts of problem, participants were asked to indicate whether the stated conditional was true, false, or neither true nor false. The assignment of the problem contents (wine vs. meals) was fully crossed with the scope of the conditionals (general vs. specific). The order in which the problems were received was also fully counterbalanced. This gave rise to four overall versions of the experiment to which participants were assigned at random. The actual frequencies presented in each problem were uniform, such that the probability of the consequent following the antecedent was always .9 (45/100).

Results and Discussion

Collapsing across the two different problem contents (for which the results did not differ), individuals had a reliable tendency to judge both the specific and general conditionals as false, which was the most frequent response for both types—specific: true 12%, neither 37%, false 51%, $\chi^2(2) = 40.79$, $p < .0001$; general: true 9%, neither 31%, false 60%, $\chi^2(2) = 67.17$, $p < .0001$. The same general pattern held when examining just the very first problems that participants received, although for specific problems the *neither* response was also highly prevalent on the first problem—specific: true 13%, neither 46%, false 42%, $\chi^2(2) = 18.15$, $p < .0001$; general: true 9%, neither 19%, false 72%, $\chi^2(2) = 54.79$, $p < .0001$.

Thus, as in Experiment 7, even in cases where the conditional expresses a very high conditional probability, participants tended

to judge it as false, and this was true for both general and specific conditionals. These results once again challenge probabilistic accounts of the conditional yet are directly predicted by deterministic accounts.

Experiment 9: Judgments of the Truth and Falsity of Conditionals and Universally Quantified Statements

Experiment 9 comprised one final investigation of the interpretation of conditionals by comparing judgments of the truth of conditional statements, both general and specific, with judgments of the truth of universally quantified statements (i.e., statements with the quantifier *all*). There is no major debate in the literature concerning such universally quantified statements—the statement that *All X are Y* is treated as unambiguously exceptionless and as logically equivalent to the statement that *No X are not Y*. It therefore represents a pertinent point of comparison for the interpretation of conditionals. I was chiefly interested in whether, once again, the predominant interpretation of conditional statements would be deterministic (exceptionless) and, if so, whether this pattern would be shown to the same extent for conditionals as for universally quantified statements.

Method

Three hundred and thirty-two participants (219 male, 113 female) completed the experiment on Amazon.com's Mechanical Turk system, in exchange for a small payment. One participant began the survey but did not complete it, and so this participant's data were removed from the analyses.

Each participant received a single problem, modeled after the wine list task used in Experiments 7 and 8. The frequencies of the respective wines in Experiment 9 were Italian red (48), Italian white (2), French red (25), French white (25). Thus, the probability of the consequent given the antecedent was very high: .96. The full wording of the problems was very similar to that used in Experiment 8 and is available upon request. Approximately one third ($n = 108$) of the participants received a general conditional statement about the overall relation that held between the wines on the wine list (*If a wine is Italian, then it is [probably] red*), another third ($n = 113$) received a specific conditional statement about a single wine that would be randomly selected from the list (*If the wine is Italian, then it is [probably] red*), and the remaining participants ($n = 111$) received a universally quantified statement about the overall relation (*[Almost] all of the Italian wines are red*). Separate from this manipulation and as indicated by the square brackets above, approximately half of the participants ($n = 167$) received a basic, unqualified conditional or universally quantified statement, whereas the remaining half ($n = 165$) received probabilistically qualified statements (*If p then probably q*; *Almost all ps are q*). Thus, the design comprised six, between-subjects cells. In each case, participants were simply asked to indicate whether the statement they received was true, false, or neither true nor false.

Results

Consistent with the previous results, participants generally regarded the unqualified statements as false, and this was true for all

three sorts of statement. The percentages of each sort of response were as follows: general conditional, false 67%, true 13%, neither 19%, $\chi^2(2) = 27.27$, $p < .0001$; specific conditional, false 65%, true 17%, neither 29%, $\chi^2(2) = 17.62$, $p < .0001$; universally quantified (all) statement, false 93%, true 5%, neither 2%, $\chi^2(2) = 91.37$, $p < .0001$. In contrast, the predominant response for the probabilistically qualified statements was *true*: general conditional, true 89%, false 5%, neither 5%, $\chi^2(2) = 78.89$, $p < .0001$; specific conditional, true 80%, false 7%, neither 13%, $\chi^2(2) = 52.33$, $p < .0001$; universally quantified (all) statement, true 91%, false 4%, neither 5%, $\chi^2(2) = 82.07$, $p < .0001$. Thus, once again, these results provide strong evidence that the predominant interpretation of the basic, unqualified conditional is deterministic (exceptionless). Interestingly, for the unqualified statements, the proportion of *false* responses (indicating a deterministic interpretation) was significantly greater for universally quantified (*all*) statements than for the general and specific conditional statements, $\chi^2(4) = 21.52$, $p < .001$. This reflects the fact that although the deterministic interpretation is predominant for conditionals, it is not unanimous, and there is a minority of participants who seem not to adopt the deterministic interpretation for conditionals (see the General Discussion for further discussion on this point).

General Discussion

Summary of the Present Findings, Alternative Interpretations, and Clarifying Remarks

The present investigation aimed to resolve a controversy underlying the psychology of conditional assertions. Various theories of conditionals have recently emerged in the literature, several of which clearly argue that conditionals have a probabilistic meaning (Liu et al., 1996; Oaksford & Chater, 1998a, 1998b, 2001, 2007; Oaksford et al., 2000; Oberauer & Wilhelm, 2003; Pfeifer & Kleiter, 2006, 2009, 2010). These accounts conflict with the classical logical interpretation of conditionals, which is deterministic, and as such, they also conflict with the interpretation assigned to conditionals by the mental model theory (Johnson-Laird & Byrne, 2002) and by formal rule theories (e.g., Rips, 1994, 2002). And perhaps most interestingly, they also clash with the suppositional theory's account of conditionals. Based on its adherence to the defective truth table, the suppositional theory also adheres to a deterministic interpretation (David Over, personal communication, February 13, 2012; see also Evans & Over, 2004; Evans et al., 2003, 2005). This may strike some readers as a surprise, given that the suppositional theory has often been grouped with other probabilistic theories, in opposition to the mental model theory (see, e.g., Evans & Over, 2012; Fugard et al., 2011; Oaksford & Chater, 2009).

There is therefore a sharp conflict in views represented in the present psychological literature, which is echoed by a similar clash of views among philosophers (see, e.g., Dudman, 1987, 1992; Lowe, 1996, who advocate a nonprobabilistic interpretation of conditionals, as compared with, e.g., Adams, 1975, 1998; Appiah, 1985, 1987; Douven, 2008; Edgington, 1996; Jackson, 1979, who advocate a probabilistic interpretation). Moreover, independent of any theoretical conflict, understanding how people represent the meaning of conditionals is important because so much of everyday

knowledge is conditional (Johnson-Laird & Byrne, 2002) and because conditional reasoning is ubiquitous in both everyday and scientific contexts (Oberauer, 2006).

Despite the importance of this problem, very little existing research has directly examined people's interpretation of the meaning of conditional assertions. Instead, this semantic issue has had a more auxiliary role, with research focusing chiefly on how individuals reason from conditional assertions or judge their probability. In the present article, my aim has been to investigate people's interpretation of the meaning of basic conditional assertions (i.e., abstract, indicative conditionals) using a variety of methods.

Experiments 1–4 examined how individuals use conditional assertions to describe patterns of data, Experiments 5 and 6 examined what individuals think is possible given a conditional assertion, and Experiments 7–9 examined how individuals judge the truth of conditional assertions. With no exceptions, these experiments each strongly corroborated the view that the default interpretation of basic conditionals is deterministic. This evidence thus provides a coherent overall picture of the ordinary, default meaning of basic conditionals. Below, I briefly review this evidence, consider some objections, and discuss its implications.

Experiments 1–4 examined how individuals used conditional assertions to describe data patterns that were either deterministic or probabilistic. The underlying rationale for these experiments was that if the default interpretation of a conditional assertion, *if p then q*, is in fact probabilistic, then individuals should see no need to further qualify such an assertion with probabilistic terminology (e.g., *if p then probably q*) when describing probabilistic data. Conversely, if the basic interpretation is deterministic, then individuals should see no need to further qualify such an assertion with modal or deterministic language (e.g., *if p then necessarily q* or *strictly speaking, if p then q*) when describing deterministic data. Indeed, the results of these experiments consistently showed that individuals preferred to describe probabilistic relations (relations where the probability of the consequent, *q*, given the antecedent, *p*, is less than 1) with conditionals that were explicitly qualified in a probabilistic fashion. Conversely, when describing deterministic relations (relations where the probability of the consequent, *q*, given the antecedent, *p*, is equal to 1), participants preferred the unqualified, *if p then q* construction and only very rarely qualified it further with modal or deterministic language. Experiment 1 demonstrated this effect for a task in which participants had to select conditional descriptions. Experiments 2–4 demonstrated the effect just as strongly on tasks where participants had to generate conditional descriptions, which relied on a variety of different ways of presenting the critical data to participants.

One potential alternative explanation of these results is that participants were aiming to be maximally informative in their responses, even at the price of redundancy (cf. Grice, 1975). Indeed, it has been shown recently that individuals will sometimes overdescribe certain states of affairs, contrary to the Gricean maxim of quantity (Engelhardt, Bailey, & Ferreira, 2006). But the trouble with this explanation is that it cannot explain why participants did not add any explicit qualification of their conditional descriptions of deterministic patterns of data. If the meaning of basic conditionals is in fact probabilistic, the tendency to overdescribe should be especially prevalent for deterministic conditionals, which are special cases of a probabilistic relation. Yet the

tendency to add extra qualification was very rare for deterministic conditionals and not nearly as prevalent as it was for probabilistic conditionals.

Alternatively, perhaps when individuals probabilistically qualify a conditional assertion, they are simply trying to convey a *weakening* of the strength of the conditional probability relation rather than to express a more categorical difference between deterministic and probabilistic relations. According to this argument, the presence of probabilistic qualification should increase as the probabilistic relation between antecedent and consequent weakens. There are two problems with this account, however. First, across Experiments 1–4, there was no discernible increase in the likelihood of probabilistic qualification as the strength of the conditional probability relation weakened. The percentages of such qualified conditionals were very high regardless of the strength of the conditional probability relation. Second, in all three experiments, there was a very abrupt difference in the percentages of probabilistically qualified conditionals when comparing conditional probabilities of 1 with the very highest conditional probabilities that were less than 1. In Experiment 1, probabilistically qualified conditionals were selected for deterministic relations only 41% of the time, but for conditional probabilities of .95, they were selected 100% of the time. The corresponding percentages of probabilistically qualified conditionals in Experiments 2, 3, and 4, in which participants generated conditionals, were as follows: Experiment 2, deterministic relation: 25%, conditional probability of .95: 78%; Experiment 3, deterministic relation: 38%, conditional probability of .92: 95%; Experiment 4, deterministic relation: 10%, conditional probability of .98: 88%. Such large differences are not easily explained on the view that explicit probabilistic qualification merely serves to convey a weakening in the strength of the relation. They are more easily explained by positing a categorical difference between unqualified conditionals, which describe deterministic relations, and probabilistically qualified conditionals, which describe probabilistic relations.

A final aspect of the data from Experiments 1–4 that may seem puzzling is the high preponderance of probabilistic conditionals used to express deterministic relations. Why would participants choose to probabilistically qualify conditionals that express deterministic relations? The most obvious explanation for this result is that participants took sampling issues and uncertainty into account. They most likely surmised that since only a sample of the relevant population of data had been examined, then for all they knew, the deterministic pattern that had been observed up to the present point might not hold in future samples.

If this explanation is true, it raises the question whether participants truly believed that the deterministic patterns of data were deterministic. And if they did not, one might be concerned that the comparison between the insertion of probabilistic language for probabilistic data and the insertion of modal or deterministic language for deterministic data is not informative. However, while this argument has some merit, it cannot explain the entire pattern of data in Experiments 1–4. In each experiment, the instructions did tell participants to focus their conditional descriptions only on the present sample of data. And while it seems reasonable to assume that this may have been forgotten or neglected by some participants, it does not seem likely that this was the case for a large number of participants. Experiment 4 had the clearest instructions to focus only on the particular sample of data in ques-

tion. For instance, for the wine problem, participants were told, “Suppose that you want to describe to someone the overall relation between the country of origin and the type (red or white) of wines on *this wine list*” (emphasis not in the original). In this experiment, the percentage of probabilistic insertions for the probabilistic data was 88%, while the percentage of modal or deterministic insertions for the deterministic data was 7%. The large difference between these percentages is not readily explained by the idea that some substantial percentage of participants treated the deterministic data as in fact being only probabilistic, since the percentage of probabilistic insertions for the deterministic data was also very low in this experiment (only 14%). A much larger percentage of probabilistic insertions for the deterministic data should have been generated if many participants were genuinely skeptical about the exceptionlessness of the relation in such cases.

Experiment 5 used an even more direct method to examine the meaning of conditionals. Individuals were asked to indicate what follows given a conditional assertion. They were given unqualified conditional descriptions, *if p then q*, and were then asked what proportion of *ps* are *qs* (consistent cases), as well as what proportion of *ps* are not *qs* (inconsistent cases). The deterministic interpretation of conditionals is that no exceptions are possible. And indeed, this was the dominant response. More than two thirds of the time, participants said that all *ps* are *qs* and that no *ps* are not *qs*. This result lends direct support to the deterministic account of conditional meaning.

One further objection to the results provided by Experiments 1–5 is that perhaps people do think that the *if p then q* construction is capable of conveying a probabilistic relation between antecedent and consequent, it is just that they typically reserve this construction for special cases in which the conditional probability of the consequent given the antecedent is 1. This argument conflicts with some of the claims made by probabilistic theorists, who have noted that the exception-tolerating nature of conditionals is usually not marked explicitly (Pfeifer & Kleiter, 2010, pp. 163–164). Nonetheless, it remains a viable account of the data.

However, this idea cannot account for the results of Experiment 6, in which participants were asked, given a conditional assertion, what percentage of *ps* must be *qs* and similarly what percentage of *ps* can be not *qs*. If, in their normal linguistic usage, people typically prefer the *if p then q* construction for cases where the conditional probability of the consequent given the antecedent is 1 but nonetheless *allow* that this construction can also be used to express probabilistic relations, this should have been revealed by this question that asked about what is possible (rather than what actually obtains). This was not found, however. The results were virtually identical to those in Experiment 5, showing instead that *if p then q* is predominantly used to express deterministic relations.

Nor can this alternative, preferential use account explain the data provided by Experiments 7–9. These experiments further buttressed the deterministic account by showing that individuals judge the truth and falsity of conditionals in line with a deterministic interpretation—a very small number of counterexamples being judged as sufficient to falsify a conditional.

In sum, across all three methods—conditional descriptions, direct judgments of conditional meaning, and truth assignments—the deterministic interpretation of conditionals was supported. Moreover, the data revealed this interpretation both for general conditionals, which range over a set of events or objects and which

are typically demarcated with the indefinite article (e.g., *If a ball is green, then it is large*), and for specific conditionals, which apply only to a single event or object and which are typically demarcated with the definite article (e.g., *If the ball is green, then it is large*; see Evans et al., 2003, p. 333; Sperber et al., 1995, for emphases of this distinction). The present results show that both sorts of conditional are generally not interpreted probabilistically.

There were some deviations from this general story. The results most clearly supported the deterministic interpretation for tasks in which participants were required to use conditional descriptions (Experiments 1–4). On these tasks, typically greater than 90% of responses supported the deterministic interpretation. The results were somewhat less clear-cut for the meaning and truth judgment tasks (Experiments 5–9), for which typically around two thirds to four fifths of the responses were in line with the deterministic interpretation. Indeed, deterministic responses were significantly less prevalent for conditional statements than they were for universally quantified statements (Experiment 9). There is thus a reasonably large set of minority responses to explain.

One possible explanation for these responses is that the core meaning of *if* is susceptible to individual differences, such that while the majority response is deterministic, at least some individuals genuinely (and consistently) treat *if* as probabilistic. A second and perhaps more plausible possibility is that there are different senses of *if*. On the one hand, the basic, default interpretation is deterministic, which is strongly suggested by the present experiments. But on the other hand, perhaps there is also a weaker, more colloquial or perhaps more pragmatic sense that is probabilistic and that can be activated in certain contexts. Indeed, the idea that there may be at least two different senses of *if* has been suggested previously (see, e.g., Oberauer et al., 2007; Verschueren, Schaeken, & d'Ydewalle, 2005).

If this two senses account is correct and there is some such looser interpretation of *if*, it is clearly not the predominant interpretation, as the present results show. However, a limitation of the present research is that it used relatively neutral (and impoverished) contexts. It therefore remains possible that the extent to which conditionals generate deterministic interpretations is moderated by variations in the context in which they are uttered. Take, for instance, the claim *If John went to the garage, then he passed by the coffee shop (and its security cameras)*, based on data regarding John's previous movements around the neighborhood. If uttered in a formal setting or in a setting where the stakes are particularly high, for example, in a court of law, in which John is a suspect in a murder case (with the murder occurring in the garage mentioned), this conditional might take on an even more deterministic meaning than it would in more neutral contexts; whereas, if uttered informally or conversationally, perhaps in connection with a query about whether John might bring back coffees, this conditional might be interpreted in a somewhat more probabilistic sense. Given the strength and consistency of the present evidence, it seems unlikely that contextual differences would completely reverse the default, deterministic interpretation of conditionals. Yet revealing whether and the extent to which different pragmatic contexts yield different interpretations remains an important task for future research.

Might recourse to pragmatics offer a more profound way to salvage a probabilistic account of conditionals? As mentioned in the introduction, none of the most prominent probabilistic theories

have made a semantic–pragmatic distinction in their accounts of conditional assertion and interpretation, and so, they seem most naturally interpreted as making semantic claims. Nonetheless, there seem to be at least two potential reconstructions of the extant probabilistic theories in pragmatic terms. One such reconstruction is that while individuals recognize that the core meaning of a conditional is exceptionless, they nonetheless adopt a probabilistic, pragmatic threshold when it comes to the evidence required to assert and accept a conditional. That is, individuals are willing to accept imperfect, probabilistic evidence when making and interpreting (nonprobabilistic) conditional claims. A second, closely related possibility is that conditional utterances generate pragmatic implicatures, which license this more probabilistic reading. On both of these reconstructions, existing probabilistic accounts are conceived as offering a view about the everyday, pragmatic use and interpretation of conditionals, such that they are typically interpreted probabilistically, at deviance with their core semantics.

However, while perhaps attractive at first, neither of these possibilities survives a confrontation with the present evidence, which shows clearly that participants are not willing to assert and accept conditionals when the evidence for the relation between antecedent and consequent is clearly probabilistic. Experiments 2–4, which examined the actual use and assertion of conditionals, show this especially clearly. Assuming that the present tasks engaged participants' ordinary pragmatic practices of assertion and interpretation—which seems reasonable—it is clear that participants did not adhere to a probabilistic pragmatic threshold for conditionals, let alone a semantic threshold.

The only way a pragmatic recasting of this sort might work is to posit that participants in the present experiments provided only their strict, nonprobabilistic interpretation of the core semantics of conditionals while nonetheless still having access to a looser (i.e., probabilistic) pragmatic reading that would have been revealed by a different sort of task. However, while this possibility cannot be entirely ruled out, the tasks used in the present investigation did not explicitly ask participants only to provide a strict interpretation of the conditionals, nor did they recruit contexts that seem especially likely to yield only a strict interpretation. Being relatively free of context, it seems more plausible to suppose that the present tasks induced participants to use and interpret conditionals as they ordinarily would in daily life. It thus seems incumbent on the critic at this point to specify just what sort of task would have yielded this looser, more pragmatic interpretation and why it is that the present tasks failed to access this interpretation. In sum, the evidence of the present experiments seems relevant both to participants' core semantic conceptions and to their pragmatic conceptions of basic conditionals, which can be presumed to converge at least for the conditionals I investigated.

These remarks are not meant to deny the possibility that pragmatic processes might play a moderating role in guiding individuals' interpretations of conditionals—they are meant merely to show that a wholly pragmatic recasting of probabilistic theories would not square those theories with the present evidence.

With respect to this notion of a pragmatic threshold, one aspect of the present account requires clarification, pertaining to the difference between meaning and certainty. The notion that the conditional has an exceptionless reading does not necessarily imply that one must be absolutely certain that the conditional holds in order to assert it. That is, what the conditional means and how

confident or certain one must be in the relevant evidence pertaining to it are separate issues. This can most clearly be seen by way of analogy with universally quantified assertions. If a person were to assert that “All the balls in the urn are white,” there is no doubt that a single black ball in the urn falsifies this assertion. This follows directly from the meaning of *all*. But the exceptionlessness of a universally quantified claim does not necessarily mean that a person must be absolutely 100% certain that all of the urns in the ball are white before asserting it. Perhaps the person is only 90% confident in this fact, having only briefly looked inside the urn and seen no black balls. Or perhaps the person has simply heard someone else’s possibly unreliable testimony about the urn without having checked it carefully him- or herself. What the standards are for assertion is a related but ultimately separate issue from what the assertion means. Nonetheless, irrespective of the speaker’s degree of certainty, the basic meaning of the claim is clear—no black balls can be tolerated, and the definite fact that a single black ball is present falsifies the assertion. The present results suggest a similar picture for conditionals. The assertion that “If you pick a ball from the urn, it will be white,” seems to preclude any black balls. This does not necessarily mean that the asserter must be absolutely certain that none of the balls are black. But if one of the balls is in fact black, the assertion will generally be judged false. Thus, just as with universally quantified assertions, the notion that a conditional has an exceptionless meaning does not necessitate that one must be absolutely certain in order to assert it.⁸

Boundary Conditions

Three boundary conditions may affect the scope of the present results. First, while I endeavored to investigate high but not certain conditional probabilities of the consequent given the antecedent (in many cases, higher than .95), it might be that investigating even higher conditional probabilities would provide a different pattern of results. It remains to be seen, for instance, whether participants would regard the truth of a conditional as being impinged when the probability of a falsifying counterexample is less than one in 1,000, one in 10,000, or some even more minuscule number.

Second, the present research focused almost exclusively on basic conditionals, which are relatively abstract and for which knowledge does not affect the interpretation of the relation between the antecedent and consequent (the one exception being the more realistic relations used in Experiment 3). It therefore speaks less directly to how people interpret non-basic conditionals, such as those that involve causal relations between antecedent and consequent (e.g., Bonnefon & Sloman, 2013; Evans, Handley, et al., 2008; Over et al., 2007; Weidenfeld, Oberauer, & Hörnig, 2005) or some other form of evidentiary support relation (e.g., Douven, 2008; Douven & Verbrugge, 2012). Experiment 3 provides the most pertinent evidence relevant to this issue, showing that deterministic interpretations were just as prevalent for non-basic conditionals that did involve some evidentiary relation between antecedent and consequent (e.g., *If a person votes Republican, then they are conservative*). However, the more realistic conditionals used in this study were all of general scope, and they were not tightly controlled in terms of the strength of the evidentiary relation between antecedent and consequent. It thus remains

an open question whether this form of non-basic conditional will typically be interpreted probabilistically or deterministically.

Third, the present investigation did not examine conditionals that pertain to truly unique events, for which no prior distribution of evidence is directly relevant, for instance, *If the economy improves, Hillary Clinton will win the next election*. For this sort of conditional, unlike the specific conditionals used in the present experiments, there is no clear-cut way to determine the conditional probability of the consequent given the antecedent and no directly relevant frequency data that can be brought to bear to determine this conditional probability. Perhaps largely as a result of this, past research that has examined how people judge the probability of conditionals has eschewed examining such conditionals about truly unique events, favoring conditionals that are amenable to the presentation of relevant probabilistic evidence (see, e.g., Evans et al., 2003, 2007; Geiger & Oberauer, 2007; Handley et al., 2006; Oberauer et al., 2007; Oberauer & Wilhelm, 2003; although see Douven & Verbrugge, 2012; Over et al., 2007, for relevant exceptions).

The present research, in following this trend in the literature, cannot speak directly to whether conditionals about truly unique events are also interpreted deterministically. It would perhaps seem unlikely that only this sort of conditional is interpreted probabilistically, whereas conditionals for which objective evidence can be brought to bear are interpreted deterministically. Nonetheless, caution seems warranted here given how intuitions sometimes diverge in similar cases. For instance, Kyburg’s (1961) famous lottery paradox captures the intuition that it is not acceptable baldly to assert that “The ticket lost” (or “The ticket will lose”) when confronted with evidence that there is only a minuscule chance of your ticket winning a lottery but without yet knowing the outcome of the lottery (see, e.g., DeRose, 1996; Williamson, 1996, 2000). Yet, at the same time, a familiar extension of this example suggests that it is acceptable to assert that “The ticket lost” when an identical probability of error exists stemming from a different source, for instance, when the information about the winning ticket is read about in the newspaper and there is a correspondingly minuscule chance that there has been a misprint and that the ticket in fact won (see, e.g., Harman, 1968). Indeed, this prediction has recently been corroborated psychologically (Turri & Friedman, *in press*). The discrepancy between these cases is puzzling because it cannot rest entirely in the difference in the objective probabilities, and so, various alternative accounts of it exist (see, e.g., DeRose, 1996; Harman, 1968; Turri & Friedman, *in press*). For present purposes, this example prompts caution in generalizing a nonprobabilistic reading of conditional assertions to all circumstances and to all kinds of conditionals. If the assertion that the ticket lost is seen as acceptable when it is read about in the newspaper, notwithstanding the small probability that is wrong, it may also be that single-event conditionals about truly unique

⁸ By contrast, however, if the conditional were instead to pertain to a unique event about which no prior evidence can be brought to bear, for example, “If I pick a white ball from the urn, I will have a sleepless night,” it is not as obvious what the conditional probability of the consequent given the antecedent must be in order for it to be judged true (or acceptable). And for this sort of conditional, it seems that issues of meaning and certainty become harder to separate cleanly. This strikes me as being part of what makes such conditionals interesting and difficult to study (see later).

events that are not based on any objective distribution of relevant evidence are also seen as inherently more resistant to probabilistic counterevidence than those that are grounded in an objective distribution of prior evidence. In sum, conditionals about truly unique events may be distinct, which makes them interesting to investigate further.

Relation to Existing Research

How do the present findings accord with other relevant evidence, particularly evidence that seems contrary to a deterministic interpretation of conditionals? As mentioned in the introduction, some prior research has been interpreted as providing support for a probabilistic interpretation of conditionals. The tasks used in these prior investigations involved either a conditional inference task (Liu et al., 1996) or a truth judgment task (Evans et al., 1996; Oberauer & Wilhelm, 2003) similar to the tasks used in Experiments 7–9 of the present article. In each case, however, the conclusion that participants made a probabilistic reading of the conditional seems unjustified given the data. Liu et al. (1996), for instance, manipulated the perceived sufficiency of the antecedent for the consequent (essentially, the conditional probability of the consequent given the antecedent) across three levels, low, medium, high, using realistic materials. Participants (in Study 1) were somewhat sensitive to this and endorsed modus ponens inferences at differential rates. Participants were asked to suppose that premises of the form, *if p then q*, and *p* were true, they indicated whether the putative conclusion, *q*, was also true. They endorsed this modus ponens inference at the following rates—72% (low), 86% (medium), and 94% (high)—which was taken to support the probabilistic account. However, these inference rates substantially outstripped the judged sufficiency of the antecedent for the consequent, as indicated by a separate group of participants (Study 2). These participants responded to the question “Knowing *p*, what is the probability of *q*?” for the materials used in Study 1 (i.e., they provided direct conditional probability assessments). Across the three levels, these conditional probabilities were judged to be 37% (low), 64% (medium), and 86% (high). Thus, contrary to the authors’ interpretation, the main finding of this research seems to be that people drew such modus ponens inferences at rates that were much higher than would be predicted by probabilistic inference accounts and that are better accounted for by deterministic accounts (though, of course, the small amount of probabilistic modulation is consistent with probabilistic accounts).

Evans et al. (1996) presented participants with frequency information about the four truth table rows that are relevant to a conditional assertion, *if p then q*, manipulating in particular the frequency of the *p, q* (true antecedent, true consequent) and *p, not q* (true antecedent, false consequent) cases. The frequency of these cases ranged over five levels, which they presented as follows (see Evans et al., 1996, p. 1093):

“absolutely true”	<i>p, q</i>	120;	<i>p, not q</i>	0
“fuzzily true”	<i>p, q</i>	116;	<i>p, not q</i>	4
(intermediate)	<i>p, q</i>	60;	<i>p, not q</i>	60
“fuzzily false”	<i>p, q</i>	4;	<i>p, not q</i>	116
“absolutely false”	<i>p, q</i>	0;	<i>p, not q</i>	120

Participants were asked to indicate on a –10 (absolutely false) to 10 (absolutely true) scale whether they thought the conditional was true in each case (this task is very similar to the tasks used in

Experiments 7 and 8 of the present investigation). The proportion of counterexamples did moderate judgments, and participants did not judge conditionals with only a few counterexamples to be absolutely false, which the authors took to support a probabilistic reading of the conditional. However, this conclusion ignores an important result, which is perhaps the most important finding: Only the exceptionless conditional was rated above the 0 midpoint of the scale (i.e., as leaning more toward truth than falsity). Even the conditional with only four counterexamples out of 120 was judged more false than true ($M = -1.15$). Thus, this result again provides more support for a deterministic interpretation than for a probabilistic interpretation.⁹

Evans et al.’s (1996) later experiments also furnish further support for this view. When participants were asked to construct truth table cases that would make conditionals either true or false, they sometimes included seemingly falsifying *p, not q* cases, which the authors took to provide evidence for a fuzzy (i.e., probabilistic) interpretation of the conditional. However, in Experiment 2, only 29% of participants ever did this—the clear majority did not (similar percentages would appear to hold in Experiments 3 and 4, though they are not provided). And while it is not provided in the study, the percentage of trials on which such *p, not q* counterexamples were provided was presumably considerably lower than 29% (which represents the percentage of individuals who responded this way at least once across all trials). Thus, despite the authors’ readiness to interpret such results as signaling a probabilistic interpretation of the conditional, a more accurate reading is that participants’ predominant reading was in fact deterministic.

Finally, Oberauer and Wilhelm (2003) relied on a similar method, in which they manipulated the presented frequencies of the true, *p, q*, and false, *p, not q*, cases relevant to a conditional statement across four different levels: HH (i.e., proportion of both true and false cases high; in subsequent levels, L stands for low), HL, LH, and LL. In each case, there was at least some number of falsifying, *p, not q*, cases. The authors were impressed by the fact that participants’ judgments were moderated by this frequency information to a small extent, that is, judgments of the truth of conditionals were HH (21%), HL (10%), LH (20%), and LL (12%). But what in fact is more striking is that in all cases, the overwhelming majority of participants regarded the conditional as false (there were only two response options, true and false). Thus,

⁹ It is also worth noting that this result has sometimes been misdescribed in the literature. For instance, Evans et al. (2003, p. 322) wrote that:

Evans, Ellis, and Newstead (1996) asked people to evaluate or construct arrays of colored shapes relating to rules such as “If it is a triangle then it is blue.” People rated such rules as true even when a small number of exceptions—say red triangles—were included in an array, and they included such exceptions when constructing their own arrays. This also suggests that an ordinary conditional assertion *if p then q* is interpreted as *q is probable given p*.

Similarly, Oberauer and Wilhelm (2003, p. 695) wrote that “Experiments 2 and 4 also provide a replication of a finding reported by Evans et al. (1996): People are ready to accept a conditional statement as true even if there are a few cases of *p –q*.” Both of these interpretations are incorrect, however, as the main text makes clear. In Evans et al. (1996), people generally treated conditionals that had even very few exceptions as more false than true (though, admittedly, not absolutely false).

once again, this evidence points toward a predominantly deterministic reading of the conditional, despite an original interpretation to the contrary.

In sum, it seems hard to escape the conclusion that this past data, along with the data presented in the present article, strongly point toward a deterministic reading of the basic conditional.

How do the present findings accord with the wealth of recent studies that have examined how individuals judge the probability of conditional assertions? Studies that have asked for assessments of degree of belief in conditional assertions or of the probability of conditional assertions appear to have generally favored Evans et al.'s (2003) conditional probability hypothesis by showing that individuals' judgment of the probability that a conditional statement is true declines roughly linearly as a function of the conditional probability of the consequent given the antecedent (e.g., Evans et al., 2003, 2007; Geiger & Oberauer, 2007; Handley et al., 2006; Oberauer et al., 2007; Oberauer & Wilhelm, 2003; Over et al., 2007; although for countervailing argument and evidence, see also Girotto & Johnson-Laird, 2004; Johnson-Laird, 2011).

At first blush, these findings may seem difficult for deterministic accounts of conditional meaning to explain—if it is true that conditionals express deterministic relations between antecedent and consequent, should they not be judged as simply false (and assigned a probability of 0) when a single counterexample exists, that is, when the conditional probability of the consequent given the antecedent is less than 1? While this might seem plausible at first, there are ways to reconcile these conditional probability findings with the deterministic findings reported here.

First, as previously noted, proponents of the suppositional account have distinguished conditionals that have general scope ranging across a class of events from those that apply to a single event. Evans et al. (2003, p. 333) noted that general scope conditionals are falsified by single counterexamples—a position that is consistent with the defective truth table and with the positions adopted by both formal rule and mental model theories. Thus, such general scope conditionals should be simply judged as false when even a single counterexample is present in the class of events that they range over (and also, presumably, as having a probability of 0). Such general conditionals thus do not provide a fertile testing ground for testing different theories of how people judge the probability of conditionals.

As a consequence, most of the existing evidence pertaining to judgments of the probability of conditionals pertains to specific conditionals. The standard task involves the selection of a single entity from an extensional array, much like those used in the present experiments (see Experiments 4 and 7–9), with participants being asked to judge the probability that a conditional will be true of the specific entity selected from that array. For such specific conditionals, the suppositional theory holds that the probability of the conditional corresponds to the conditional probability of the consequent given the antecedent. Most pertinently for present purposes, the suppositional theory appears to hold that the mere possibility that a counterexample *might* arise does not automatically falsify the conditional. Instead, the conditional is only falsified (with respect to the specific case in question) should such a counterexample actually be selected. Thus, the difference in scope of conditional assertions (general vs. specific) is critical. Notably, the mental model theory seems to adopt a similar general position, despite differing in its details regarding how probability is judged.

Although the model theory has not explicitly distinguished general and specific conditionals, the way it has been implemented indicates that, just like the suppositional theory, it also holds that a specific conditional is not falsified by the mere possibility that a counterexample will arise but that it is only falsified should such a falsifying, *p*, *not q* case actually be selected (see, e.g., Girotto & Johnson-Laird, 2004). (Of course, the model theory does differ from the suppositional theory in how it suggests people will judge the probability of conditionals, i.e., the model theory posits that people will rely on an interpretation corresponding to the material implication interpretation rather than a defective truth table interpretation.) This therefore suggests that the existing evidence regarding judgments of the probability of conditionals is consistent with the notion that people interpret conditionals in a deterministic, exceptionless fashion. General or wide-scope conditionals are falsified by the presence of any *p*, *not q* counterexamples in an extensional array. Specific or narrow-scope conditionals are falsified only if such a *p*, *not q* case should actually arise when selected from such an array and they can be assigned probabilities prior to that selection having taken place (or being known about).

The one major complication that the present evidence poses for these accounts is that people seem to judge specific conditionals as not even allowing the possibility, let alone the actuality, of specific counterexamples. Individuals tended to judge such specific conditionals as false (Experiments 7–9) even when a specific, falsifying example had not been selected but simply could have been selected. They also tended to think that such specific conditionals were true only if no *p*, *not q* cases were in the array from which a specific entity would be selected (Experiments 5 and 6). This result is thus hard to accommodate for both suppositional and mental model theories. One, albeit post hoc way to accommodate these findings is to suggest that questions about a conditional's truth invite a general reading of them (even when the conditional is framed specifically), whereas questions about their probability invite a more specific reading. To investigate this possibility, I ran a separate experiment (not reported in the main body of the article) in which some participants were asked to judge the truth of specific conditionals like those used in Experiments 7–9 (*If the wine is Italian, then it is red*), whereas other participants judged their probability ("What are the chances that this statement is true?"). The conditional probability of the consequent (the wine is red) given the antecedent (the wine is Italian) was .9. An interesting discrepancy emerged to these two questions. While 71% of participants who judged truth regarded the conditional as outright false, 57% who judged the probability of the very same conditional gave the well-documented conditional probability response (that there was a 90% chance that the conditional was true). This result is consistent with the notion that truth questions invite a general reading of the conditional (such that it is taken to range over a set of cases), whereas probability questions invite a more specific reading (pertaining only to a single case), though it is hardly definitive. It also suggests that there may be important, unexplored differences in how people interpret truth versus probability questions for conditionals. Future research that aims to further explore and understand these potential differences is an important task.

Conclusion

The psychology of conditionals is so complex that no current theory encompassing how individuals reason from and to conditionals, how they judge the probability of conditionals, and how they interpret their meaning may be able to fully account for all the existing data about conditionals. Nonetheless, the present results constitute strong support for an assumption shared by mental model and formal rules accounts of the conditional. Nine experiments, using a variety of methodologies, all consistently showed that the basic conditional, *if p then q*, is generally interpreted in a deterministic fashion, whereas only a more elaborated, qualified construction, *if p then probably q*, is taken accurately to reflect probabilistic relations.

One major account that is usually classified as probabilistic—the suppositional theory—also predicts the present findings. This theory postulates the defective truth table in which *q* follows from *p* just as invariably as it does on the mental model theory’s interpretation. However, all of the other probabilistic theories that have populated the recent literature are challenged by the present results because they explicitly claim that the conditional has a probabilistic meaning. And unlike the suppositional theory, these theories build this assumption into their core theories of inference (see, e.g., Liu et al., 1996; Oaksford & Chater, 1998a, 1998b, 2001, 2007, 2009; Oaksford et al., 2000; Oberauer & Wilhelm, 2003; Pfeifer & Kleiter, 2005, 2006, 2009, 2010). It is hard to see how such probabilistic theories can accommodate the present results. Thus, one of the main contributions of the present research is to elucidate a hitherto previously unacknowledged difference between the various probabilistic theories of the conditional. A second contribution is to refute a core assumption of theories that claim that the basic meaning of the conditional is probabilistic. While the present results may not count as decisive refutation of these theories in their entirety, they pose a distinct and novel challenge to them. Proponents of probabilistic accounts of the conditional may be able to accommodate these results but seemingly not without substantial modification to their existing theories. In contrast, the results are directly predicted by the mental model theory, formal rules, and suppositional accounts of the conditional.

References

- Adams, E. W. (1975). *The logic of conditionals*. Dordrecht, the Netherlands: Reidel. doi:10.1007/978-94-015-7622-2
- Adams, E. W. (1998). *A primer of probability logic*. Stanford, CA: CLSI Publications.
- Appiah, A. (1985). *Assertion and conditionals*. Cambridge, England: Cambridge University Press. doi:10.1017/CBO9780511895975
- Appiah, A. (1987). “If” again. *Analysis*, 47, 193–199. doi:10.1093/analysis/47.4.193
- Barrouillet, P., Gauffroy, C., & Lecas, J. F. (2008). Mental models and suppositional accounts of conditionals. *Psychological Review*, 115, 760–771. doi:10.1037/0033-295X.115.3.760
- Barrouillet, P., Grosset, N., & Lecas, J. F. (2000). Conditional reasoning by mental models: Chronometric and developmental evidence. *Cognition*, 75, 237–266. doi:10.1016/S0010-0277(00)00066-4
- Bonnefon, J.-F., & Sloman, S. A. (2013). The causal structure of utility conditionals. *Cognitive Science*, 37, 193–209. doi:10.1111/cogs.12014
- Braine, M. D. S., & O’Brien, D. P. (1991). A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review*, 98, 182–203. doi:10.1037/0033-295X.98.2.182
- Braine, M. D. S., & O’Brien, D. P. (Eds.). (1998). *Mental logic*. Mahwah, NJ: Erlbaum.
- Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, 34, 1452–1482. doi:10.1111/j.1551-6709.2010.01126.x
- DeRose, K. (1996). Knowledge, assertion, and lotteries. *Australasian Journal of Philosophy*, 74, 568–580. doi:10.1080/00048409612347531
- Douven, I. (2008). The evidential support theory of conditionals. *Synthese*, 164, 19–44. doi:10.1007/s11229-007-9214-5
- Douven, I., & Verbrugge, S. (2012). Indicatives, concessives, and evidential support. *Thinking & Reasoning*, 18, 480–499. doi:10.1080/13546783.2012.716009
- Dudman, V. H. (1987). Appiah on “if”. *Analysis*, 47, 74–79. doi:10.1093/analysis/47.2.74
- Dudman, V. H. (1992). Probability and assertion. *Analysis*, 52, 204–211. doi:10.1093/analysis/52.4.204
- Edgington, D. (1996). Lowe on conditional probability. *Mind*, 105, 617–630. doi:10.1093/mind/105.420.617
- Engelhardt, P. E., Bailey, K. G. D., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean maxim of quantity? *Journal of Memory and Language*, 54, 554–573. doi:10.1016/j.jml.2005.12.009
- Evans, J. St. B. T., Ellis, C. E., & Newstead, S. E. (1996). On the mental representation of conditional sentences. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 49(A), 1086–1114.
- Evans, J. St. B. T., Handley, S. J., Neilens, H., & Over, D. E. (2007). Thinking about conditionals: A study of individual differences. *Memory & Cognition*, 35, 1772–1784. doi:10.3758/BF03193509
- Evans, J. St. B. T., Handley, S. J., Neilens, H., & Over, D. (2008). Understanding causal conditionals: A study of individual differences. *Quarterly Journal of Experimental Psychology*, 61, 1291–1297. doi:10.1080/17470210802027961
- 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. doi:10.1037/0278-7393.29.2.321
- Evans, J. St. B. T., Neilens, H., Handley, S. J., & Over, D. E. (2008). When can we say if? *Cognition*, 108, 100–116. doi:10.1016/j.cognition.2008.02.001
- Evans, J. St. B. T., & Over, D. E. (2004). *If*. Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780198525134.001.0001
- Evans, J. St. B. T., & Over, D. E. (2012). Reasoning to and from belief: Deduction and induction are still distinct. *Thinking & Reasoning*. Advance online publication. doi:10.1080/13546783.2012.745450
- 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. doi:10.1037/0033-295X.112.4.1040
- Fugard, A. J. B., Pfeifer, N., Mayerhofer, B., & Kleiter, G. D. (2011). How people interpret conditionals: Shifts toward the conditional event. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 635–648. doi:10.1037/a0022329
- Geiger, S. M., & Oberauer, K. (2007). Reasoning with conditionals: Does every counterexample count? It’s frequency that counts. *Memory & Cognition*, 35, 2060–2074. doi:10.3758/BF03192938
- Giroto, V., & Johnson-Laird, P. N. (2004). The probability of conditionals. *Psychologia*, 47, 207–225. doi:10.2117/psysoc.2004.207
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 64–75). New York, NY: Seminar Press.

- Handley, S. J., Evans, J. St. B. T., & Thompson, V. A. (2006). The negated conditional: A litmus test for the suppositional conditional? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 559–569. doi:10.1037/0278-7393.32.3.559
- Harman, G. (1968). Knowledge, inference, and explanation. *American Philosophical Quarterly*, *5*, 164–173.
- Jackson, F. (1979). On assertion and indicative conditionals. *Philosophical Review*, *88*, 565–589. doi:10.2307/2184845
- Johnson-Laird, P. N. (2006). *How we reason*. New York, NY: Oxford University Press.
- Johnson-Laird, P. N. (2011). The truth about conditionals. In K. Manktelow, D. Over, & S. Elqayam (Eds.), *The science of reason: A Festschrift for Jonathan St. B. T. Evans* (pp. 119–144). Hove, England: Psychology Press.
- Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, *109*, 646–678. doi:10.1037/0033-295X.109.4.646
- Johnson-Laird, P. N., & Tagart, J. (1969). How implication is understood. *American Journal of Psychology*, *82*, 367–373. doi:10.2307/1420752
- Kyburg, H. E. (1961). *Probability and the logic of rational belief*. Middletown, CT: Wesleyan University Press.
- Leslie, S.-J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language*, *65*, 15–31. doi:10.1016/j.jml.2010.12.005
- Liu, I., Lo, K., & Wu, J. (1996). A probabilistic interpretation of “if-then”. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *49(A)*, 828–844.
- Lowe, E. J. (1996). Conditional probability and conditional beliefs. *Mind*, *105*, 603–615. doi:10.1093/mind/105.420.603
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*, 608–631. doi:10.1037/0033-295X.101.4.608
- Oaksford, M., & Chater, N. (1998a). *Rationality in an uncertain world*. Hove, England: Psychology Press. doi:10.4324/9780203345955
- Oaksford, M., & Chater, N. (1998b). A revised rational analysis of the selection task: Exceptions and sequential sampling. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 372–398). Oxford, England: Oxford University Press.
- Oaksford, M., & Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences*, *5*, 349–357. doi:10.1016/S1364-6613(00)01699-5
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780198524496.001.0001
- Oaksford, M., & Chater, N. (2009). Précis of *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. *Behavioral and Brain Sciences*, *32*, 69–84. doi:10.1017/S0140525X09000284
- Oaksford, M., Chater, N., & Larkin, J. (2000). Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 883–899. doi:10.1037/0278-7393.26.4.883
- Oaksford, M., & Stenning, K. (1992). Reasoning with conditionals containing negated constituents. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 835–854. doi:10.1037/0278-7393.18.4.835
- Oberauer, K. (2006). Reasoning with conditionals: A test of formal models of four theories. *Cognitive Psychology*, *53*, 238–283. doi:10.1016/j.cogpsych.2006.04.001
- Oberauer, K., Geiger, S. M., Fischer, K., & Weidenfeld, A. (2007). Two meanings of “if”? Individual differences in the interpretation of conditionals. *Quarterly Journal of Experimental Psychology*, *60*, 790–819. doi:10.1080/17470210600822449
- 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.
- Over, D. E. (1993). Deduction and degrees of belief. *Behavioral and Brain Sciences*, *16*, 361–362. doi:10.1017/S0140525X00030570
- Over, D. E., & Evans, J. St. B. T. (2003). The probability of conditionals: The psychological evidence. *Mind & Language*, *18*, 340–358. doi:10.1111/1468-0017.00231
- Over, D. E., Evans, J. St. B. T., & Elqayam, S. (2010). Conditionals and non-constructive reasoning. In M. Oaksford & N. Chater (Eds.), *Cognition and conditionals: Probability and logic in human thinking* (pp. 135–151). Oxford, England: Oxford University Press.
- Over, D. E., Hadjichristidis, C., Evans, J. St. B. T., Handley, S. J., & Sloman, S. (2007). The probability of causal conditionals. *Cognitive Psychology*, *54*, 62–97. doi:10.1016/j.cogpsych.2006.05.002
- Pfeifer, N., & Kleiter, G. D. (2005). Towards a mental probability logic. *Psychologica Belgica*, *45*, 71–99.
- Pfeifer, N., & Kleiter, G. D. (2006). Inference in conditional probability logic. *Kybernetika*, *42*, 391–404.
- Pfeifer, N., & Kleiter, G. D. (2009). Framing human inference by coherence based probability logic. *Journal of Applied Logic*, *7*, 206–217. doi:10.1016/j.jal.2007.11.005
- Pfeifer, N., & Kleiter, G. D. (2010). The conditional in mental probability logic. In M. Oaksford & N. Chater (Eds.), *Cognition and conditionals: Probability and logic in human thinking* (pp. 153–173). Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780199233298.003.0009
- Politzer, G., Over, D. E., & Baratgin, J. (2010). Betting on conditionals. *Thinking & Reasoning*, *16*, 172–197. doi:10.1080/13546783.2010.504581
- Quelhas, A. C., Johnson-Laird, P. N., & Juhos, C. (2010). The modulation of conditional assertions and its effects on reasoning. *Quarterly Journal of Experimental Psychology*, *63*, 1716–1739. doi:10.1080/17470210903536902
- Rips, L. J. (1994). *The psychology of proof*. Cambridge, MA: MIT Press.
- Rips, L. J. (2002). Reasoning. In D. Medin (Ed.), *Stevens' handbook of experimental psychology: Vol. 2. Memory and cognitive processes* (3rd ed., pp. 317–362). New York, NY: Wiley.
- Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, *57*, 31–95. doi:10.1016/0010-0277(95)00666-M
- Turri, J., & Friedman, O. (in press). Winners and losers in the folk epistemology of lotteries. In J. R. Beebe (Ed.), *Advances in experimental epistemology*. New York, NY: Continuum.
- Verschueren, N., Schaeken, W., & d'Ydewalle, G. (2005). A dual-process specification of causal conditional reasoning. *Thinking & Reasoning*, *11*, 239–278. doi:10.1080/13546780442000178
- Weidenfeld, A., Oberauer, K., & Hörnig, R. (2005). Causal and noncausal conditionals: An integrated model of interpretation and reasoning. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *58(A)*, 1479–1513.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, *105*, 489–523. doi:10.2307/2998423
- Williamson, T. (2000). *Knowledge and its limits*. Oxford, England: Oxford University Press.

Appendix A

Materials Used in Experiment 3

Problem number	Problem contents	<i>p</i>	<i>q</i> ; not <i>q</i>	Version 1 (block; probability)	Version 2 (block; probability)	Version 3 (block; probability)	Version 4 (block; probability)
1	You are interested in how a certain card player plays cards. In particular, you are interested in whether when this player holds the ace card, they also hold a king card. You have examined a series of this player's hands of cards, shown below. Each row represents a different hand of cards.	Ace	King; Not king	Block 1; 3/5	Block 1; 5/5	Block 2; 12/12	Block 2; 11/12
2	You are interested in the drinks on offer in a hospital. In particular, you are interested in whether drinks that contain caffeine are coffee, or some other drink. You have examined a series of drinks in the hospital, shown below. Each row represents a different drink.	Caffeine	Coffee; Not coffee	Block 1; 5/6	Block 1; 6/6	Block 2; 10/10	Block 2; 8/10
3	You are interested in the playing cards used in a board game. In particular, you are interested in whether cards that have a letter on one side also have a number on the other side. You have examined a series of playing cards for this board game, shown below. Each row represents a different card.	Letter	Number; Not number	Block 1; 6/8	Block 1; 8/8	Block 2; 8/8	Block 2; 7/8
4	You are interested in why a certain beach gets closed. In particular, you are interested in whether the days when the beach is closed are bad weather days. You have examined a series of days when the beach was closed, shown below. Each row represents a different day.	Beach closed	Bad weather; Not bad weather	Block 1; 9/10	Block 1; 10/10	Block 2; 6/6	Block 2; 4/6
5	You are interested in how a certain teacher teaches basic geometry. In particular, you are interested in which shapes are written on the board at the same time. You have examined a series of lessons, shown below. Each row represents a different lesson.	Triangle	Square; Not square	Block 1; 10/12	Block 1; 12/12	Block 2; 5/5	Block 2; 4/5
6	You are interested in the relationship between wedding postponements and the mood of the guests, in a particular city. In particular, you are interested in whether weddings that are postponed have unhappy guests. You have examined a series of postponed weddings, shown below. Each row represents a different postponed wedding.	Wedding postponed	Guests unhappy; Guests not unhappy	Block 1; 4/5	Block 1; 5/5	Block 2; 12/12	Block 2; 10/12

(Appendices continue)

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Appendix A (continued)

Problem number	Problem contents	<i>p</i>	<i>q</i> ; not <i>q</i>	Version 1 (block; probability)	Version 2 (block; probability)	Version 3 (block; probability)	Version 4 (block; probability)
7	You are interested in the color of the scarves that people own at a certain University. In particular, you are interested in whether people who own green scarves also own blue scarves. You have examined a series of people's wardrobes, shown below. Each row represents a different wardrobe.	Green scarf	Blue scarf; No blue scarf	Block 1; 4/6	Block 1; 6/6	Block 2; 10/10	Block 2; 9/10
8	You are interested in why people vote Republican in a certain town. In particular, you are interested in whether people who vote Republican are conservatives. You have interviewed a series of people who vote Republican, shown below. Each row represents a different voter.	Votes Republican	Conservative; Not conservative	Block 1; 7/8	Block 1; 8/8	Block 2; 8/8	Block 2; 6/8
9	You are interested in items that people have in their closet in your apartment building. In particular, you are interested in whether people who have a mop in the closet also have a wrench in the closet. You have examined a series of people's closets, shown below. Each row represents a different closet.	Mop	Wrench; No wrench	Block 1; 8/10	Block 1; 10/10	Block 2; 6/6	Block 2; 5/6
10	You are interested in the composition of jazz bands at a music competition. In particular, you are interested in whether the jazz bands have pianists. You have examined a series of different bands at the competition, shown below. Each row represents a different band.	Plays jazz	Pianist; No pianist	Block 1; 11/12	Block 1; 12/12	Block 2; 5/5	Block 2; 3/5
11	You are interested in the dishes at a certain restaurant. In particular, you are interested in whether dishes that contain onion come with a side salad. You have examined a series of dishes at this restaurant, shown below. Each row represents a different dish.	Onion	Side salad; No side salad	Block 2; 5/5	Block 2; 3/5	Block 1; 11/12	Block 1; 12/12
12	You are interested in crime in a particular police jurisdiction. In particular, you are interested in whether districts that experienced more violent crime had less of a police presence. You have examined a series of police districts, shown below. Each row represents a different district.	Violent crime	Less police presence; Not less police presence	Block 2; 6/6	Block 2; 5/6	Block 1; 8/10	Block 1; 10/10
13	You are interested in the available apartments in a certain apartment building. In particular, you are interested in whether apartments that come with a hardwood floor also come with a built-in microwave. You have examined a series of apartments in the building, shown below. Each row represents a different apartment.	Hardwood floor	Built-in microwave; No built-in microwave	Block 2; 8/8	Block 2; 6/8	Block 1; 7/8	Block 1; 8/8

(Appendices continue)

Appendix A (continued)

Problem number	Problem contents	p	q ; not q	Version 1 (block; probability)	Version 2 (block; probability)	Version 3 (block; probability)	Version 4 (block; probability)
14	You are interested in dog attacks over the past year. In particular, you are interested in whether attacks that are vicious are caused by Rottweilers. You have examined a series of dog attacks from the past year, shown below. Each row represents a different attack.	Vicious attack	Inflicted by Rottweiler; Not inflicted by Rottweiler	Block 2; 10/10	Block 2; 9/10	Block 1; 4/6	Block 1; 6/6
15	You are interested in the books on a friend's bookcase. In particular, you are interested in whether their hard-cover books are less than 200 pages long. You have examined a number of books from your friend's bookcase, shown below. Each row represents a different book.	Hard-cover	Less than 200 pages; Not less than 200 pages	Block 2; 12/12	Block 2; 10/12	Block 1; 4/5	Block 1; 5/5
16	You are interested in the drugs developed by a medical research institute. In particular, you are interested in whether drugs that have no side effects are approved for use by the public. You have examined a series of drugs developed by the institute, shown below. Each row represents a different drug.	No side effects	Approved for use; Not approved for use	Block 2; 5/5	Block 2; 4/5	Block 1; 10/12	Block 1; 12/12
17	You are interested in a local high school's basketball team. In particular, you are interested in whether the team wins when a player called Edwards is in the team. You have examined a series of games from the past season, shown below. Each row represents a different game.	Edwards on team	Win; Don't win	Block 2; 6/6	Block 2; 4/6	Block 1; 9/10	Block 1; 10/10
18	You are interested in the plays that a local theater shows. In particular, you are interested in whether plays that are sold out are ones that are well reviewed. You have examined a series of plays that the theater put on, shown below. Each row represents a different play.	Play sold out	Well reviewed; Not well reviewed	Block 2; 8/8	Block 2; 7/8	Block 1; 6/8	Block 1; 8/8
19	You are interested in the items hanging on the walls in a local mansion. In particular, you are interested in whether the items that are framed are photographs. You have examined a series of items in the mansion, shown below. Each row represents a different item.	Item framed	Photograph; Not photograph	Block 2; 10/10	Block 2; 8/10	Block 1; 5/6	Block 1; 6/6
20	You are interested in the recovery from serious knee injuries by players in your favorite football team. In particular, you are interested in whether injuries that are operated on allow the injured player to make a comeback. You have examined a series of knee injuries sustained by players in your team, shown below. Each row represents a different injury.	Injury operated on	Comeback; No comeback	Block 2; 12/12	Block 2; 11/12	Block 1; 3/5	Block 1; 5/5

Note. These items were designed to vary in terms of whether the conditional that described them was either arbitrary or realistic. Item number 1 is arbitrary, item number 2 is nonarbitrary, and this alternating sequence continues throughout the rest of the items.

(Appendices continue)

Appendix B

Materials Used in Experiment 4 and Also Adapted for Use in Experiments 5 and 6

Problem number	Introduction to problem	Frequencies presented		Antecedent	
1	Imagine that you are presented with a box filled with 100 balls that differ in terms of their size and their color. The box contains the following numbers of different kinds of balls (each row represents a different kind of ball):	Large Large Small Small	Green Blue Green Blue	45 (50) 5 (0) 25 25	If a/the ball is large, then:
2	Imagine that you are looking at a wine list that has 100 different wines that differ in terms of their country of origin and their type (red or white). The wine list includes the following numbers of different kinds of wines (each row represents a different kind of wine):	Italian Italian French French	Red White Red White	49 (50) 1 (0) 40 10	If a/the wine is Italian, then:
3	Imagine that you are looking at a menu that has 100 different meals that differ in terms of whether or not they include meat and their level of spiciness. The menu includes the following numbers of different kinds of meals (each row represents a different kind of meal):	Meat Meat Vegetarian Vegetarian	Spicy Mild Spicy Mild	42 (50) 8 (0) 5 45	If a/the meal is meat, then:
4	Imagine that you are fishing in a lake filled with 100 fish that differ in terms of their speed and their size. The lake contains the following numbers of different kinds of fish (each row represents a different kind of fish):	Fast Fast Slow Slow	Large Small Large Small	44 (50) 6 (0) 10 40	If a/the fish is fast, then:
5	Imagine that you are presented with a barrel filled with 100 toys that differ in terms of their outer material and their mode of operation. The barrel contains the following numbers of different kinds of toys (each row represents a different kind of toy):	Furry Furry Plastic Plastic	Battery operated Not battery operated Battery operated Not battery operated	48 (50) 2 (0) 30 20	If a/the toy is furry, then:
6	Imagine that you go to a car yard filled with 100 cars that differ in terms of whether they are new or used and their price. The car yard contains the following numbers of different kinds of cars (each row represents a different kind of car):	New New Used Used	Expensive Cheap Expensive Cheap	46 (50) 4 (0) 25 25	If a/the car is new, then:
7	Imagine that you are presented with an iTunes playlist that contains 100 songs that differ in terms of their genre and their era. The playlist contains the following numbers of different kinds of songs (each row represents a different kind of song):	Jazz Jazz Rock Rock	From the 1970s From the 1980s From the 1970s From the 1980s	43 (50) 7 (0) 35 15	If a/the song is jazz, then:
8	Imagine that you are looking at a bookcase filled with 100 books that differ in terms of their binding and their genre. The bookcase contains the following numbers of different kinds of books (each row represents a different kind of book):	Hardcover Hardcover Paperback Paperback	Fiction Nonfiction Fiction Nonfiction	47 (50) 3 (0) 20 30	If a/the book is a hard-cover, then:

Note. Materials were also adapted for use in Experiments 5 and 6. In the first two rows of each problem, the frequencies for the probabilistic problems are presented first, with the corresponding frequencies for the deterministic problems presented in parentheses. The specific problems referred to a specific entity chosen at random from the set, and used the definite article “the,” whereas the general problems referred to the “overall relation” between the antecedent and consequent in the set as a whole, and used the indefinite article “a.”

Received December 1, 2011
Revision received July 16, 2013
Accepted July 25, 2013 ■