THE LOGIC OF CONDITIONALS[1]

by

Ernest Adams

University of California, Berkeley

The standard use of the propositional calculus ('P.C.') in analyzing the validity of inferences involving conditionals leads to fallacies, and the problem is to determine where P.C. may be 'safely' used. An alternative analysis of criteria of *reasonableness* of inferences in terms of conditions of *justification* rather than *truth* of statements is proposed. It is argued, under certain restrictions, that P. C. may be safely used, except in inferences whose conclusions are conditionals whose antecedents are *incompatible* with the premises in the sense that if the antecedent became known, some of the previously asserted premises would have to be withdrawn.

1. The following nine 'inferences' are all ones which, when symbolized in the ordinary way in the propositional calculus and analyzed truth functionally, are valid in the sense that no combination of truth values of the components makes the premises true but the conclusion false.

F1. John will arrive on the 10 o'clock plane. Therefore, if John does not arrive on the 10 o'clock plane, he will arrive on the 11 o'clock plane.

F2. John will arrive on the 10 o'clock plane. Therefore, if John misses his plane in New York, he will arrive on the 10 o'clock plane.

F3. If Brown wins the election, Smith will retire to private life. Therefore, if Smith dies before the election and Brown wins it, Smith will retire to private life.

F4. If Brown wins the election, Smith will retire to private life. If Smith dies before the election, Brown will win it. Therefore, if Smith dies before the election, then he will retire to private life.

F5. If Brown wins, Smith will retire. If Brown wins, Smith will not retire. Therefore, Brown will not win.

F6. Either Dr. A or Dr. B will attend the patient. Dr. B will not attend the patient. Therefore, if Dr. A does not attend the patient, Dr. B will.

F7. It is not the case that if John passes history, he will graduate. Therefore, John will pass history.

F8. If you throw both switch S and switch T, the motor will start. Therefore, either if you throw switch S the motor will start, or if you throw switch T the motor will start.

F9. If John will graduate only if he passes history, then he won't graduate. Therefore, if John passes history he won't graduate.

We trust that the reader's immediate reaction to these examples agrees with ours in rejecting or at least doubting the validity of these inferences — one would not ordinarily 'draw' the inferences if one were 'given' the premises. If this is so, it poses a problem for the *application* of formal logic: either applied formal logic is wrong (and it is only as an applied theory that formal logic can be said to be right or wrong at all) or, if it is right and these inferences are valid after all, then we have to explain why they 'appear' fallacious. This is the problem which this paper is concerned with. More exactly, we shall be concerned with criteria of validity for inferences involving ordinary conditionals, and with determining to what extent formal logical theory may safely be applied to analyzing and formalizing arguments involving conditionals. As will be seen, this objective is only partially achieved, if at all, and this partial 'solution' itself raises what are probably more profound difficulties.

2. Before passing to a detailed discussion, it may be useful to consider briefly the significance of these examples. If these inferences are truly fallacious, then clearly they call into question some of the principles of formal logic and its application. The use of formal logic in the analysis of the logic of conditionals rests on two principles, one general, the other specific to conditionals. These are:

Principle 1 (general)

An inference is deductively valid if and only if it is logically impossible for the premises to be true and the conclusion false.

Principle 2

At least for purposes of formal analysis, 'if then' statements can be treated as truth functional — true if either their antecedents are false or their consequents true, false only if their antecedents are true and the consequents false.

Thus baldly stated, of course, Principle 2 appears very dubious, and one might immediately be led to expect that fallacies would arise in applications of the formal theory which tacitly make the assumption. Examples F1–F9 might be regarded simply as confirmation of this expectation. F1 and F2 are in fact counterparts of the paradoxes of material implication.

It should be noted, however, that the principles of material implication (i.e. the truth conditions for material conditionals) are not themselves arbitrary, but can be 'deduced' from still more fundamental principles which must in turn be questioned if Principle 2 is questioned. For instance, the 'law' that a conditional is true if its consequent is true follows from Principle 1 together with the principle of conditionalization. The argument goes thus. *p* and *q* logically imply *p* (Principle 1). If *p* and *q* imply *p*, then *p* implies 'if *q* then *p*' (principle of conditionalization), hence *p* implies 'if *q* then *p*'. If *p* implies 'if *q* then *p*', and *p* is true, then 'if *q* then *p*' is true (Principle 1), hence, if *p* is true, so is 'if *q* then *p*'. Therefore, if we reject Principle 2, we must reject either the principle of conditionalization or Principle 1, or both, as well.

Similarly, the truth conditions for the material conditional follow from the apparent logical equivalence of 'not *p* or *q*' and 'if *p* then *q*' plus Principle 1. Therefore, rejection of Principle 2 requires either that Principle 1 be rejected or that the supposition that 'not *p* or *q*' is equivalent to 'if *p* then *q*' be re-examined. A closely related supposition is in fact at the heart of example F6.

Looked at in another way, example F6 seems to bring into question a principle which follows from Principle 1 alone: namely, that if a conclusion follows logically from a single premise, then it also follows logically from two premises one of which is the same as the original premise. One might be willing to grant that 'If Dr. A does not attend the patient, Dr. B will' follows logically from 'Either Dr. A or Dr. B will attend the patient', but if one rejected the inference in F6, this would be a counterinstance to the above consequence of Principle 1.

Finally, examples F3 and F4 call into question laws of conditionals not specifically associated with material implication. F4, if it is fallacious, is a counterinstance to the hypothetical syllogism (if 'if *q* then *r*' and 'if *p* then *q*' then 'if *p* then *r*'). A closely related principle is taken as a postulate in C. I. Lewis' [4, p. 125] theory of strict implication. It is unlikely, therefore, that fallacies of the kind given here can be entirely avoided by going over to formal analysis in terms of strict implication or related systems.

3. Are the inferences in examples F1–F9 fallacious? We may concentrate for the present on F1 and ask whether

C If John does not arrive on the 10 o'clock plane, then he will arrive on the 11 o'clock plane.

is a logical consequence of

P John will arrive on the 10 o'clock plane.

To avoid irrelevant misunderstandings, let us specify that both statements are about the same person, and that both P and the antecedent of C are about the same event — John's arrival on the 10 o'clock plane. If we attempt to answer this question by applying the criterion of validity of formal logic, formulated in Principle 1, we must ask: Is it logically possible for P to be true, but C false?

To attempt to answer this question is to see that it has no clear answer. The reason is that the term 'true' has no clear ordinary sense as applied to conditionals, particularly to those whose antecedents prove to be false, as the antecedent of C must if P is true. This is not to say that conditional statements with false antecedents are not sometimes called 'true' and sometimes 'false', but that there are no clear criteria for the applications of those terms in such cases. This is, of course, an assertion about the ordinary usage of the terms 'true' and 'false', and it can be verified, if at all, only by examining that usage. We shall not conduct such an examination here, but leave it to the reader to verify by observation of how people dispute about the correctness of conditional statements whose antecedents prove false, that precise criteria are lacking.

Formal logicians might perhaps be inclined to dismiss the vagueness of the ordinary usage of 'true', and provide a new and more precise definition which is appropriate to their needs. This is in fact what Tarski has done with very fruitful results in the *Wahrheitsbegriff*. One might even regard the truth conditions of material conditionals not as defining a new sentential connective, but rather as defining the term 'true' precisely as applied to ordinary conditionals. The danger in this approach lies in the likelihood of begging the question raised at the beginning of this section. If the application of 'true' is defined arbitrarily for those cases in which it has no clear ordinary sense, then this arbitrariness will extend to the determination as to whether statement C is a logical consequence of P or not.

It may, of course, be argued that the imprecision of the truth conditions for conditional statements shows that these statements are vague, and not that the term 'true' is. This argument would, it seems to us, prove too much. Conditional statements are ubiquitous in argument, including the most rigorous, and it seems doubtful to us that even in the case of conditional statements which occur in rigorous mathematical demonstration the word 'true' has any clear ordinary sense.

In view of the foregoing remarks, it seems to us to be a mistake to attempt to analyze the logical properties of conditional statements in terms of their truth conditions. But, if the notion of truth has no clear sense as applied to conditionals, then it no longer makes sense to apply the criterion of deductive validity formulated in Principle 1 to inferences involving conditionals. That is, the question as to whether C can be validly inferred from P cannot be decided by asking whether it is possible for P to be true and C false, and some other way of characterizing the validity of inferences involving conditionals must be sought.

The foregoing observations do not answer the original question, which was: Do the conclusions in the examples follow logically from their premises? In one sense, the conclusions do follow logically from the premises, because it is possible to give logical arguments that the truth conditions of ordinary 'if then' statements are the same as those of the material conditional, and once this is granted, the conclusions in the examples follow directly from the premises. Aren't these examples, then, merely cases — like countless others in mathematics — in which a conclusion which really does follow from assumptions is not 'seen' to follow because the argument which proves it is not given? In fact, the missing 'proof' of the conclusion in F1 is very easily supplied:

Assume P:	John will arrive on the 10 o'clock plane.
Therefore:	Either John will arrive on the 10 o'clock plane or he will arrive on the 11 o'clock plane.
Therefore C:	If John does not arrive on the 10 o'clock plane, he will arrive on the 11 o'clock plane.

Q. E. D.

This settles the question as to whether C follows logically from P unless we are inclined to suspect that the argument is fallacious — and why should we suspect that? It seems to us that most people would have doubts about the soundness of this argument (even if they could not verbalize them), and the reason for this suspicion could only be in the fact that there can be independent, non-deductive grounds for rejecting the conclusion, even while accepting the premise. As we have argued, these grounds cannot be that the premise may be true but the conclusion false. In the following sections we shall inquire into what these grounds might be. It will then remain, of course, to re-examine the above and like arguments to 'prove' the conclusions in the examples, to determine where, if anywhere, they go astray.

4. Let us return again to example F1 and ask what grounds there might be for accepting P (that John will arrive on the 10 o'clock plane) and rejecting C (that if John does not arrive on the 10 o'clock plane, he will arrive on the 11 o'clock plane), other than the knowledge that P is true and C is false. It seems to us that in a very general sense of 'grounds', the following telegram from John might do just this:

> I will arrive on the 10 o'clock plane. If I don't arrive on that plane, I will arrive on the 2 o'clock plane.

The receipt of the telegram does not, of course, logically imply P nor contradict C. We hope, though, that the reader will grant, without necessarily being able to explain why, that the telegram would give its recipient reason for accepting P, for *asserting* it and for acting accordingly (for instance, by meeting the 10 o'clock plane in case he wanted to meet John). It would also give reason for rejecting C, for *denying* it, and acting accordingly (for instance, by not meeting the 11 o'clock plane in case John proved not to be on the 10 o'clock plane). We might say that example F1 violates a condition for reasonable inferences. *Condition for reasonableness of inferences*: If an inference is reasonable, it should not be the case that on some occasion the assertion of its premises would be justified, but the denial of its conclusion also justified.

That the foregoing condition, vague as it is, cannot be maintained in full generality is shown from the example of complicated arguments such as occur frequently in science and mathematics. In such arguments it often occurs that the premises (accepted principles) are justified, but in the absence of the argument justifying the conclusion, one would be justified in denying it. The condition seems more compelling as applied to *immediate* inferences, though to clarify this idea and the application of the principle to it would be hard. In any case, we are concerned here with the question of whether such inferences as those in examples F1–F9 can be justified at all and not just whether they are reasonable immediate inferences (which they obviously are not). For the present we shall take the above condition as a rough heuristic guide in inquiring into the justification of such specific inferences as those in F1–F9, and the justification of the general rules of inference of which they are particular applications. The cash value of this will be to direct inquiry to conditions of 'justified assertability' rather than conditions of truth in assessing inferences. Later we will give the condition a precise mathematical formulation within the calculus of probability, which will provide a systematic means of exploring the reasonableness of formal inference schemata which represent many simple rules of inference.

5. In the following section we will be discussing justifications for the assertion and denial of conditionals and their components. What we will be concerned with is not what *ought* to be required in order to justify the assertion of statements, but rather what is expected in practice. It is important to note that what is expected in practice depends to an extent on the context of the assertion. In certain 'critical' contexts, the required justification for the assertion of a mathematical statement is a rigorous proof, which leads to 'logical' certainty in the thing asserted. The justification demanded for the assertion of generalizations in empirical science is clearly different, and seldom yields as complete certainty as does mathematical proof. The assertions in the examples which presently concern us are neither typically mathematical nor scientific. As in example F1, one is usually justified in asserting such statements even where there is a distinct possibility of their later being 'contradicted by the facts'.

The foregoing observations suggest that reasonable inferences involving conditionals asserted in 'everyday' situations should be analyzed in terms of requirements of justification or 'assertability' which yield something less than certainty in the statement asserted. As we shall see, it can happen that it is reasonable to infer certain consequences from a statement if its assertion is very strictly justified, which would not be reasonable inferences from the same statement made with less strict justification. We shall argue, in fact, that it is this difference in the strictness of the requirements of assertability which explains why such fallacies as F1–F9 arise in 'everyday speech', but do not (at least, they are hard to find) in mathematical argument. Also looking ahead, we shall suggest that where strict requirements of assertability are satisfied, conditional statements can safely be analyzed as though they were material conditionals, but they cannot if they are less strictly justified.

6. We shall not be able to explain in detail what are the ordinary requirements of 'assertability' for the sorts of statements which occur in the examples. What we shall do instead is to suggest a hypothesis which connects 'justification of assertion' with another sort of justification (and makes a similar connection for 'justification of denial'), which we hope will be illuminating. Actually, this hypothesis applies only to a rather limited class of statements, and not, unfortunately, to most of those of mathematics and science which have traditionally been the center of the logician's interests. We shall state the hypothesis first, and then discuss its scope and significance.

Hypothesis 1

> The assertion of a statement is justified if and only if a bet on it is justified in the same situation; the denial of a statement in a situation is justified if and only if a bet against it is justified in that situation.

Of course, we are under no illusions as to the precision of this requirement. What it does is to link the justified assertion of a statement with a particular kind of 'action in accordance' with a belief in what it says. The particular action in this case is betting, and in speaking of a bet's being justified, we shall pay attention only to what might be called the 'game-theoretic' aspects of the situation; that is, ethical justifications, even such long-range strategic considerations as 'it is foolish to bet if you can't afford to lose', etc. are to be left out of account.

The hypothesis applies only to those statements for which there are clear *ordinary* procedures for settling bets on them.[2] The conditions for settling bets are just the conditions under which they win or lose. For such particular 'empirical' statements as 'John will arrive on the 10 o'clock plane' these conditions are well defined, and our hypothesis applies to that sort of statement. On the other hand, many meaningful statements do not have such clear conditions: we would argue that all 'infinite' general statements (e.g. 'There is no highest prime') have less clear conditions, if any, for settling bets.[3] We shall speak of statements for which the ordinary conditions of settling bets on them are clear as 'bettable'.

We can now inquire further into the conditions which justify bets on bettable statements (and which therefore justify their assertion, if Hypothesis 1 is correct). Consider first particular unconditional statements, such as that John will arrive on the 10 o'clock plane, uttered on a particular occasion. It seems to us intuitively obvious that a bet that John will arrive on the 10 o'clock plane is justified (in a situation) just in case there is reason to believe that John will arrive on the 10 o'clock plane — i.e. that the statement 'John will arrive on the 10 o'clock plane' will prove true. Such reasons may be better or worse. Very good reasons would yield certain knowledge that the statement would prove true, and that the corresponding bet would win; weaker but still frequently 'adequate' reasons would only yield a high probability that the statement would prove true and the bet would win. We shall speak of the bet and the corresponding assertion as 'strictly justified' in the former case, as 'probabilistically justified' in the latter. We formulate these requirements in a second hypothesis:

Hypothesis 2

a. The assertion of a particular bettable statement is strictly justified on an occasion if what is known on that occasion makes it certain to prove true; its denial is strictly justified if it is certain to prove false.

b. The assertion of a particular bettable statement on an occasion is probabilistically justified if what is known on that occasion makes it probable that the statement will prove true; its denial is probabilistically justified if it is probable that the statement will prove false.

Again, of course, we have a vague hypothesis which more than anything else can be useful only as pointing in the direction of an analysis of the logic of conditional statements. We are not yet able to go much deeper than this, however, and can only hope that as it stands (and in the mathematical form we shall give it later) it will yield some insight into the problems we have been discussing, and in particular 'explain' what is paradoxical about examples F1–F9. However, we must first attempt to formulate a similar hypothesis about conditional statements.

For many conditional statements, the conditions of settling bets on them are just as clear as they are on particular non-conditional statements. Normally, bets on conditional statements are considered to be themselves conditional: for instance, a bet on the statement 'If John does not arrive on the 10 o'clock plane, he will arrive on the 11 o'clock plane' is conditional on John's not arriving on the 10 o'clock plane. More generally, a bet that 'if *p* then *q*' is conditional — in force only if *p* proves true, and in that case winning if *q* is true, and losing if *q* is false.[4,5]

Under what circumstances would a bet that 'if *p* then *q*' be justified (and therefore, if our Hypothesis 1 is correct, the conditional statement 'if *p* then *q*' be justified)? We may expect, as before, differing degrees of strictness in the requirements of assertability. However, here the situation proves somewhat more complicated than in the unconditional case. It would be too much to require to justify betting that 'if *p* then *q*' that the bet be certain or even very likely to win, since this would mean that 'if *p* then *q*' would be justified only if both *p* and *q* were (because the bet that 'if *p* then *q*' wins only if both *p* and *q* prove true). It would be too little to require merely that the bet be very likely not to lose, since it might be much more likely not to win (thus, a bet that 'If the Chinese land first on the moon, they will plant rice paddies on it' would be irrational because, though it would probably not be lost it would far more certainly not be won). Tentatively we shall suggest that a conditional bet is justified in case it is much more likely to be won than lost (though it may be more likely still to be called off). If the bet on the conditional were *certain* not to be lost, then, of course, it could not be still less likely to win, and we may regard the bet in this case as 'strictly' justified. Correspondingly, we may say that the assertion of a conditional statement is strictly justified in case it is certain that either the antecedent is false or the consequent true (that is, that the corresponding conditional bet would not lose), and that it is probabilistically justified if the likelihood that both the antecedent and consequent are true is much greater than that the antecedent is true and the consequent false.

Before formulating the foregoing criteria as another 'hypothesis', though, we must briefly consider a difficulty. What about the case of a conditional asserted in the certainty that its antecedent is false? In this case the corresponding bet would be certain to be called off, and it makes no sense to speak of it as being either more or less likely to be won than lost (hence our criterion of 'probabilistic justification' breaks down). On the other hand, the bet would be certain not to lose and therefore, according to the 'strict' criterion just outlined, the bet on it would be justified, and so would the assertion. Worse, the bet *against* the same statement would be equally justified, and hence, by parity of reasoning, we should say that under these circumstances both the assertion and the denial of the conditional are strictly justified.

It is significant that indicative conditional statements are seldom made in the *knowledge* that their antecedents are false. The reason is clear: such an assertion would be misleading. Where conditionals are stated in such knowledge, they are usually put in the subjunctive mood. It is doubtful that *ordinary* usage prescribes criteria either for the rationality of bets made on such conditionals, or for the justification of the making of such conditional assertions (grammatically they are unjustified, but this is a different sense of justification, we believe, from that which we have so far been discussing). Nonetheless, requirements of theoretical completeness demand that we give a place to such 'vacuous' conditionals, and that we say something about their justification. We choose to say that they are *vacuously* both strictly and probabilistically justified. In doing this, though, we are simply making an arbitrary stipulation in the interests of theoretical completeness and simplicity, and not one which is in any sense based on an analysis of the concepts of 'justification' or 'rationality'. In this case also, the same argument leads us to denominate the denials of these vacuous conditionals (which are equally vacuous) as justified in both senses, vacuously. Thus, we have a special category of statements which, according to our arbitrary stipulation, may be justifiably both asserted and denied. This, of course, leads to some awkward complications in the logical laws which these statements obey, but we have reason to think that any attempt to bring vacuous conditionals into a logical theory must be paid for at a price at least as high.

Now we are ready to formulate these hypotheses, and give a definition of vacuous justification.

Hypothesis 3

a. The assertion of a bettable conditional 'if p then q' is strictly justified on an occasion if what is known on that occasion makes it certain that either p is false or q is true; its denial is strictly justified if it is certain that either p is false or q is false.

b. The assertion of a bettable conditional 'if p then q' is probabilistically justified on an occasion if what is known on that occasion makes it much more likely that p and q are both true than that p is true and q is false; its denial is probabilistically justified if it is much more likely that p is true and q is false than that p and q are both true.

c. (Definition) The assertion and denial of a bettable conditional 'if p then q' are both vacuously probabilistically and strictly justified on an occasion if what is known on that occasion makes it certain that p is false.

One effect of making this arbitrary stipulation in 3c is to make strict assertability and deniability special cases of probabilistic assertability and deniability respectively. In a later section (section 13) we will formulate these hypotheses mathematically within the framework of the calculus of conditional probability, and on that basis derive some general results. However, it is possible to apply these criteria of justified assertion and denial, together with the necessary condition for reasonable inference formulated in section 4, directly to the examples F1–F9, which we now proceed to do.

7. Before taking up the examples F1–F9 one by one, we must first modify the condition for reasonable inference slightly to take into account the 'odd' case of the vacuous conditional. The modification is self-explanatory.

Modified condition for reasonableness of inferences:

If an inference is reasonable, it cannot be the case that on any occasion the assertion of its premises is justified, and the non-vacuous denial of its conclusion is justified.

Now for the examples. F1 has already been discussed (in fact, it is the one whose discussion led to the formulation of the criterion), and F2 is so similar that no discussion seems necessary. F3 and F4 are slightly more complicated, but can be dealt with in the same way. To take F3, it is clear that in an election campaign between Smith and Brown, in which Smith had intimated his intention to retire to private life in the event of Brown's winning, a conditional bet that if Brown won Smith would retire would be justified, but the conditional bet that in the event of Smith's death prior to the election and Brown's winning, Smith would retire would not be justified. According to our 'betting' criterion of justification, therefore, the assertion of the premise would be justified, and the denial of the conclusion would be too (the conditional would not be vacuous since it would be certain that Smith would not die or Brown would not win). The same kind of analysis applies to F4. Later, of course, we must attempt to explain why it is that, in spite of counterexamples such as F3 and F4, we are inclined to look on the general logical laws of which they are special cases (e.g. the hypothetical syllogism) as valid, whereas the 'paradoxical' laws of material implication which lead to F1 and F2 are almost universally regarded with suspicion.

8. F5 represents a somewhat more complex situation. According to our criterion, the only situation in which two 'inconsistent' conditionals 'if *p* then *q*' and 'if *p* then not *q*' may justifiably be asserted is that in which both are vacuous — i.e. in which *p* is certain to be false. Thus, in one sense, the inference in F5 satisfies our necessary condition for reasonable inference since in this case the assertion of not *p* is justified and its denial is not. However, as we previously suggested, one of the requirements of proper usage demands that conditionals not be asserted in the knowledge that their antecedents are false (though subjunctive conditionals are not subject to this restriction). Therefore a pair of conditional statements of the forms 'if *p* then *q*' and 'if *p* then not *q*' are seldom if ever justifiably asserted on the same occasion. When such a pair of statements are made on the same occasion, it is usually the case that one is asserted in contradiction to the other, and this carries the implication that the contradicted statement is false or at least that it may justifiably be denied (and non-vacuously). Thus, in the case of a contradiction, it is usually the case that at least one of the two contradicting statements is *not* justified.[6] And, if one of the premises of the inference is not justifiably asserted, then it does not follow that the conclusion can be justifiably asserted. That is, we would be quite right in *not* concluding that Brown would not win, upon hearing it said by one person that if Brown won, Smith would retire, and by another that if Brown won, Smith would not retire; rather, we would be justified in concluding that one of the statements we heard was not justified.

9. Example F6 raises what are perhaps the most interesting issues of all of F1–F9. For purposes of ready reference, let us symbolize its three assertions (letting '$\rightarrow$' stand for the 'if then' of English):

$A \vee B$	Either Dr. A or Dr. B will attend the patient.
$-B$	Dr. B will not attend the patient.
$-A \rightarrow B$	If Dr. A does not attend the patient, Dr. B will.

As noted, this seems to be a case in which, though the consequence $-A \rightarrow B$ follows from premise $A \vee B$ alone, it does not from premises $A \vee B$ and $-B$ together. Nor is this a case (like F5) in which one would refrain from 'inferring' $-A \rightarrow B$ on hearing $A \vee B$ and $-B$ asserted, on the grounds that they appeared to contradict each other and therefore that not both of the assertions were justified.

According to our criterion of justification, the inference from $A \vee B$ to $-A \rightarrow B$ does not satisfy the condition for reasonable inference. The reason is that there can be a situation in which, though the assertion of $A \vee B$ is justified, the non-vacuous denial of $-A \rightarrow B$ is justified. This is that in which the grounds for asserting the disjunction $A \vee B$ are just the probable knowledge of its first member; that is, the knowledge that Dr. A would attend the patient. Knowing that Dr. A would attend the patient would be sufficient reason for making a bet that either Dr. A or Dr. B would attend the patient, and hence would justify asserting that one of them would attend the patient, if our Hypothesis 2 is correct. On the other hand, this same knowledge would not justify the making of a conditional bet that if Dr. A did *not* attend the patient, Dr. B would. Thus, F6 appears to be simply a case in which the premises and conclusion do not satisfy a criterion of reasonable inference.

Unfortunately, we are here left with a difficulty: why, if the inference from the disjunction 'either p or q' to the conditional 'if not p then q' is not always reasonable, do we normally accept it without hesitation? Part of the answer is to be found in a further condition of correct assertability for disjunctions (beyond simply grounds for believing that they will prove true). This is a requirement that disjunctions not be asserted in the knowledge (either certain or probable) that one of the parts is true. Though the assertion of disjunctions violating this requirement is justified in the 'betting' sense we have been concentrating on, such a statement is misleading, and therefore probably runs against standards (at least implicit ones) of correct communication, and hence is seldom made. Now, as we shall see later, on any occasion on which the assertion of a disjunction 'either p or q' is justified, but neither p nor q can be justifiably asserted, the assertion of 'if not p then q' is justified. Therefore, the inference of 'if not p then q' from the correct assertion of the disjunction 'either p or q' *is* justified in the betting sense, and this may help to explain why we do not hesitate to make the inference.

But this leaves us just where we started: if the inference from 'either

p or q', correctly and justifiably asserted, to 'if not p then q' is justified, why is it apparently not in F6? There is, in fact, no prima facie reason to suppose that premise A v B was not either justifiably or correctly asserted. It seems to us that the answer is this. The two statements A v B and -B in F6 have the look of ones uttered on the same occasion by different speakers, A v B being asserted first. We shall suppose that this in fact was the situation. If both speakers were justified, and the first speaker spoke correctly, it could only be that the second speaker knew something the first one did not: namely that Dr. A would attend the patient. Now consider the inferences which a third person, listening to the conversation, would be entitled to draw. Having heard statements A v B alone, he would be justified in asserting $-A \rightarrow B$, but having heard the second statement -B as well, he would no longer be justified in doing this. This can be seen more clearly, perhaps, in considering what bets it would be reasonable to make in the circumstances. Having heard that either Dr. A or Dr. B would attend the patient, it would be reasonable to bet that if Dr. A did not, then Dr. B would. Later, having received the additional information that Dr. B would not attend the patient, it would no longer be reasonable to make this bet.

What the foregoing shows is that increasing the amount of information (which corresponds to adding further premises) may not just add to the number of statements which may be reasonably inferred from them; it can also subtract. Of course, there should be no surprise at this in view of the fact that we are dealing with only probabilistic requirements of assertability. If a statement can be justifiably asserted under circumstances in which it may be later 'contradicted by the facts', it is clear that there will be occasions on which a statement whose assertion is justified at one time is not justified (must be withdrawn) at another. Of course, this cannot occur where strict requirements of assertability are demanded (which may partially explain why the logic of mathematical inference is comparatively simpler than that of everyday life).

10. The most interesting feature of the next two fallacies is that they both involve what are normally regarded as 'truth-functional' compounds of conditionals — a denial in F7 and a disjunction in F8. If, as we have argued, the concept of truth has no clear sense as applied to conditionals, these 'truth-functional' compounds of them cannot be defined in the usual way: by giving a truth table in which for each

combination of truth values of the components, the truth value of the compound statement is listed. Thus, we should expect the logical laws of these compounds to be different from those which they obey when applied to components for which the notion of truth has a clear sense. This expectation is fulfilled.

As has already been suggested, the ordinary meaning of the denial 'it is not the case that if p then q' is just to assert 'if p then not q', and clearly the antecedent of the denied conditional is not a logical consequence of this. This is probably obvious, and the only significance it has in this context is to lend indirect support to our contention that the concept of truth is not clear as applied to conditionals.

11. Disjunctions of conditionals present a special and interesting problem in their own right. This is that disjunctions of most of the conditionals occurring in the examples F1–F7 have no clear ordinary sense. Consider, for example, the grammatically correct disjunction: 'Either if John passes history he will graduate, or, if John does not pass history he will not graduate.' It is difficult to see what the quoted sentence means because it is hard to imagine what anyone uttering it might mean by it, what would justify its assertion, etc. The most likely interpretation (the one most commonly given the author on occasions on which he has asked people how they would interpret this statement) is very close to a 'meta-assertion' to the effect that at least one of the two conditionals could be justifiably asserted, but it was not known which. Under such an interpretation, of course, the meaning of the disjunction is clearly different from that of the corresponding disjunction of material conditionals. (If the conditionals disjoined in the sentence quoted were material, that disjunction would be equivalent to the tautology 'If John passes and doesn't pass history, then he will either graduate or not graduate'.)

As further evidence that disjunctions including conditionals sometimes have this meta-assertive character, consider the following:

F10. Either the restricted continuum hypothesis is provable, or, if the restricted continuum hypothesis is provable, then so is the general continuum hypothesis.

It seems to us that this disjunction would normally be taken to mean that either the restricted continuum hypothesis had been proved, or that it had been proved that if the restricted hypothesis could be proved, then so could the general. Analyzed as a material conditional, of course, the foregoing statement is one of the simplest of tautologies.

The problem which these usages present to logic as a *descriptive* theory is to explain why disjunctions of conditionals have no clear ordinary sense in many cases, and not to *give* a meaning to such expressions. However, we are only able at present to raise the problem, and not to give a solution to it.

The conditionals disjoined in example F8 are somewhat different from those occurring in the previous examples, in that their antecedents refer to *repeatable* conditions, whereas the antecedents of the earlier conditionals could only be fulfilled once. Thus, the antecedent of the conditional 'If you throw switch S and switch T, the motor will start' can be fulfilled several times by throwing the switches repeatedly, whereas in the conditional 'If John does not arrive on the 10 o'clock plane, he will arrive on the 11 o'clock plane' the antecedent can only be fulfilled once, since John can only not arrive once on the 10 o'clock plane. We suggest, without attempting to pursue it in detail here, that such conditionals as these demand a rather different analysis from the 'unrepeatable' conditionals (quantifying over occasions on which the antecedent is fulfilled, and at the same time making the standard model-theoretical analysis more nearly applicable).[7]

12. The last example, F9, presents two problems: (1) the meaning and conditions of assertability of 'only if' constructions, and (2) the conditions of assertability of conditionals whose antecedents are not themselves particular statements. We shall make only very brief remarks on both, chiefly with the object of giving an understanding of what is wrong with the inference in the example.

One rather striking difference between 'if then' and 'only if' statements is that the latter are never bettable. Thus, for example, there are no ordinary conditions for settling a bet that 'John will graduate only if he passes history'. In one sense, therefore, the quoted statement is not equivalent to either of the two conditionals 'If John graduates, then he will have passed history', and 'If John does not pass history, then he will not graduate', which are alternative ways of analyzing the 'only if' statement customarily used in formal logic (as we shall see later, the two conditionals are not themselves probabilistically equivalent).

One reason why 'only if's are frequently not bettable is that they assert, in effect, a conjunction of conditionals. If, for instance, some-one asserts of a player X in a card game 'X will win only if he gets an ace on his last card', he asserts in effect that for anything other than an ace as a last card, if X gets that he will not win. Such a conjunction of conditionals is closely related to, but is not probabilistically equivalent to a single conditional with a disjunctive antecedent (which is bettable): i.e. the conditional 'If X gets either a 2 or a 3, or anything besides an ace, he will not win'.[8]

Conjunctions of conditionals are never bettable in the sense that there are standard conditions for the settling of a *single* bet on the conjunction (e.g. a bet that 'If John passes history he will graduate, and if he doesn't pass history he won't graduate'). On the other hand, the assertion of a conjunction is justified on an occasion just in case the assertion of each of its members is justified. Hence, if 'only if's can be analyzed as conjunctions of conditionals, the requirement of assertability for conditionals can be extended to include 'only if's as well.

The analysis of the conditions of assertability for bettable conditionals does not extend to those whose antecedents are not particular bettable statements, and therefore it is not possible to apply the criterion of justified assertability stated in Hypothesis 3 to showing that there may be situations in which the premise of example F9 may be justifiably asserted, whereas the conclusion may not. It is, however, possible to apply the general criterion of reasonable inference directly to show that the inference in this example is not reasonable. That is, it is possible to describe a circumstance in which (at least intuitively) we would say that the assertion of 'If John will graduate only if he passes history, then he won't graduate' is justified, but the statement 'If John passes history, he won't graduate' can be justifiably denied. This is one in which there is good reason to believe: (1) that John won't pass history, and (2) that if John passes history he *will* graduate (perhaps he has been promised this by the dean).

The general problem of describing conditions of justified assertability for non-bettable statements is one which we are not prepared to discuss at present. Hence we cannot go further here into the analysis of statements such as the premise of F9. With this we end our discussion of the examples.

13. In view of the foregoing discussion, one might well wonder whether there are any general rules of inference involving conditionals which are reasonable in the sense we have defined. In order to investi-

gate this question we have formalized the criterion of reasonableness using the calculus of probability, and studied its consequences. In this and the next section we will describe the basic assumptions of this formulation and some of its more significant consequences. Except for giving precise statements of the assumptions, the discussion will be informal, although the results described have been rigorously proved.

We begin with a formal language for representing truth-functional compounds of atomic statements, and conditionals whose antecedents and consequents are such truth-functional compounds. We here use the '&', 'v' and '-' in their usual senses as the truth-functional 'and' 'or' and 'not', respectively, and use the arrow '→' to symbolize the *non*-material 'if-then'. For the sake of formal simplicity it is also useful to include two atomic sentences '*T*' and '*F*', representing logically true and logically false statements, respectively. The formal language is then defined recursively in the obvious way:

The formal language

A formula is any expression of one of the following three kinds:

i) Atomic — '*T*', '*F*', and lower case letters (e.g. '*p*', '*q*', etc.) possibly with numeral subscripts;
ii) Truth functional — atomic formulas and expressions constructed from them using the binary connectives '&' and 'v' and the unary connective '-';
iii) *Conditional* — all formulas of the form $\varphi \rightarrow \psi$, where φ and ψ are truth functional.[9]

In view of the fact that we are able to characterize the conditions of assertability for denials of conditionals, we might have included them too in our formalism. However, there would be little point in doing so since, if what we have previously argued is correct, 'it is not the case that if *p* then *q*' is equivalent to 'if *p* then not *q*', and hence can be replaced by it.

The notions of *tautology* and *tautological consequence* will be assumed to be applicable in the usual way to truth-functional formulas of the formal language. Later we shall also speak of conditional formulas as tautological consequences of other formulas ('premises', conditional or truth functional), in the sense that if the conditionals were replaced by material conditionals (i.e. all formulas of the form $\varphi \rightarrow \psi$ were replaced by $-\varphi \vee \psi$), then the resulting formula would be a tautological consequence of its premises.

Next we define the notion of a *conditional probability function* for the formal language, modifying slightly the definition given in Kolmogorov [3, p. 47]. The modification consists in defining the probability of a conditional $\varphi \rightarrow \psi$ to be 1, in case the probability of its antecedent, φ, is 0. This is in accord with our convention to regard a conditional as assertable in case its antecedent is known with certainty to be false.

Conditional probability functions

A real-valued function Pr whose domain is the class of formulas is a conditional probability function if and only if:

i) for all formulas A, $0 \leqq Pr(A) \leqq 1$ and $Pr(F) = 0$ and $Pr(T) = 1$;

ii) for all truth-functional formulas φ and ψ, if ψ is a tautological consequence of φ then $Pr(\varphi) \leqq Pr(\psi)$;

iii) for all truth-functional formulas φ and ψ, if φ tautologically implies $-\psi$ then $Pr(\varphi \vee \psi) = Pr(\varphi) + Pr(\psi)$;

iv) for all truth-functional formulas φ and ψ, if $Pr(\varphi) \neq 0$ then $Pr(\varphi \rightarrow \psi) = Pr(\varphi \& \psi)/Pr(\varphi)$, and if $Pr(\varphi) = 0$ then $Pr(\varphi \rightarrow \psi) = 1$.

It remains now to formulate the condition for a formula A to be reasonably inferable from a *set* of formulas, S. The general requirement is that it should not be the case that there be a situation under which all of the formulas of S are assertable, but A is non-vacuously deniable. We will actually formulate a slightly stronger requirement: namely, that in any situation in which all of the formulas of S are assertable, A must also be *assertable* (not just not deniable), though it can be shown that this apparently stronger requirement is actually equivalent to the weaker one. There are two degrees of assertability, strict and probabilistic, and two corresponding requirements of reasonableness for inferences. Thus, we shall call the inference of A from S *strictly reasonable* in case the strict assertability (corresponding to probability 1) of all formulas of S guarantees the strict assertability of A. Roughly, we may say that the inference of A from S is *probabilistically reasonable* in case the fact that all formulas of S have high probability guarantees that the probability of A is high. The latter idea is vague, and so there is necessarily some arbitrariness about the precise definition given below.

Reasonable inferences

Let S be a set of formulas and let A be a formula. Then the inference of A from S is *strictly reasonable* if and only if for all conditional probability functions Pr such that $Pr(B) = 1$ for all B is S, $Pr(A) = 1$. The inference of A from S is *probabilistically reasonable* if and only if for all $\varepsilon > 0$ there exists $\delta > 0$ such that for all conditional probability functions Pr such that $Pr(B) \geqq 1 - \delta$ for all B in S, $Pr(A) \geqq 1 - \varepsilon$.

According to the above definition, the inference of A from S is probabilistically reasonable in case an arbitrarily high probability (at least $1 - \varepsilon$) can be guaranteed for A by requiring the probabilities of all the members B of S to be at least $1 - \delta$. Under the circumstance that the inference of A from S is strictly or probabilistically reasonable, we shall also say that A is a strict or probabilistic *consequence* of S.

We list now some consequences of the fundamental definitions, and discuss their significance.

Consequence 1

Probabilistic reasonableness implies strict reasonableness, and strict reasonableness is equivalent to tautological implication.

Thus, according to the above, the requirement that the inference of A from S be probabilistically reasonable is, in a sense, the strictest requirement, since it entails both that this inference is strictly reasonable, and that A is a tautological consequence of the formulas of S (in the generalized sense of tautological consequence applying to conditional formulas). The second part of Consequence 1 is important in that it guarantees that symbolizing 'if-then' by the material conditional and applying truth-functional inference cannot lead from assertable premises to deniable conclusions so long as the premises are logically certain. It is also important in that it shows that the only formal problem remaining to investigate is the nature of probabilistic reasonableness.

Concerning the properties of the relation of probabilistically reasonable inferability which are of primarily formal significance, it can be shown that the relation of probabilistic consequence has most of the general properties of deducibility relations. For instance, every element of a set is a probabilistic consequence of it, and probabilistic consequences of probabilistic consequences are probabilistic consequences. It is also the case that any probabilistic consequence of a subset is a probabilistic consequence of the whole set. This shows that no more than the standard truth-functional logic does our calculus admit the possibility that adding premises may actually detract from what can be inferred. The reason for this is implicit in the use of the calculus of probability: it, like the truth-functional logic, treats probabilities (analogous to truth values) as things fixed, and not susceptible to change in the light of evidence. One rather interesting respect in which our calculus differs from the standard one is that ours does not satisfy the *compactness* requirement — there are infinite sets S of formulas and formulas A such that A is a probabilistic consequence of S, but A is not a probabilistic consequence of any finite subset of S.

Coming now to the more detailed properties of probabilistic consequence, these are most simply explained in terms of the auxiliary notions of a truth assignment's *verifying* or *falsifying* a formula.

Verification and falsification

> Let A be a formula and let f be a function with domain the class of truth-functional formulas and range $\{0,1\}$. f is a *truth-assignment* if and only if $f(F) = 0, f(T) = 1$, and for all truth-functional formulas φ and ψ, $f(\varphi \ \& \ \psi) = f(\varphi) \cdot f(\psi), f(\varphi \vee \psi) = f(\varphi) + f(\psi) - f(\varphi) \cdot f(\psi)$, and $f(-\varphi) = 1 - f(\varphi)$. If A is truth functional, then A is *verified* or *falsified under* f according as $f(A) = 1$ or $f(A) = 0$. If $A = \varphi \rightarrow \psi$ then A is verified under f if $f(\varphi) = f(\psi) = 1$, and A is falsified under f if $f(\varphi) = 1$ and $f(\psi) = 0$.

Going back to our analysis of justification in terms of bets, we may say that a statement is verified just in case a bet on it wins, and is falsified in case the bet on it loses. Note that in the conditional case, the statement $\varphi \rightarrow \psi$ is neither verified nor falsified under a truth-assignment f for which $f(\varphi) = 0$ (i.e. φ is false). The notions of verification and falsification are used to define a new pair of 'entailment' relations:

Strong and weak entailment

> Let S be a set of formulas and A be a formula. Then S *weakly entails* A if and only if all truth-assignments f which falsify A falsify at least one member of S. S *strongly entails* A if and only if it weakly entails A, and all truth-assignments f which falsify no members of S and verify at least one of them also verify A.

Intuitively, we might say that S weakly entails A in case not losing on any bet on a member of S guarantees not losing on a bet on A, and S strongly entails A in case not losing on any bet on S, and winning on at least one of them entails *winning* on a bet on A. In case S is empty, then strong and weak entailment coincide. We are now ready to state our most important result:

Consequence 2

Let S be a finite set of formulas and let A be a formula. Then A is a strict consequence of S if and only if S weakly entails A, and A is a probabilistic consequence of S if and only if there is a subset S' of S such that S' strongly entails A.

The first part of Consequence 1 simply states the not surprising fact that weak entailment coincides with strict consequence and hence with tautological consequence. The second part is the more significant: intuitively, the inference of A from a (finite) set S is probabilistically reasonable in case there is a subset S' of S such that (1) not losing on bets on any member of S' guarantees not losing on A, and (2) winning on at least one bet on a member of S' and not losing on the others guarantees winning on A. It is easy to check the inferences in examples F1–F4 to see that they don't satisfy the requirement of strong entailment. In the case of the hypothetical syllogism, for example, any truth-assignment f such that $f(p) = 0$ and $f(q) = f(r) = 1$ verifies $q \rightarrow r$, does not falsify $p \rightarrow q$, but fails to verify $p \rightarrow r$, which it should do if $q \rightarrow r$ and $p \rightarrow q$ strongly entailed $p \rightarrow r$.

Some significant corollaries follow immediately from Consequence 2.

Corollaries to Consequence 2

Let S be a finite set of formulas and let A be a formula. Then:

2.1 If S is empty or A is truth functional, then A is a probabilistic consequence of S if and only if A is a strict consequence of S.

2.2 If $A = \varphi \rightarrow \psi$ and all of the formulas of S are truth functional, then A is a probabilistic consequence of S if and only if either both φ and ψ are tautological consequences of S, or ψ is a tautological consequence of φ.

2.3 If $A = p \rightarrow \psi$, where p is an atomic formula not occurring in any formula of S, then A is a probabilistic consequence of S if and only if p tautologically implies ψ, or F is a strict consequence of S.

The importance of 2.1 is obvious. Truth-functional logic is 'safe' as applied to inferences whose conclusions are truth functional (even if some of the premises are not). To the extent that what ultimately interests us in reasoning are unconditional conclusions (and we reason only with simple conditionals), this suggests that it can do no harm to treat conditionals as truth functional. Corollary 2.2 is cited only to show agreement between consequences of our calculus and intuition: essentially it says that unconditional premises can only probabilistically entail a conditional conclusion if in fact they entail both its antecedent and its consequent (in which case one would not normally assert the conditional anyway). Corollary 2.3 is cited for the same reason: it says in effect that if the antecedent of a conditional is not mentioned among a set of premises, then the conclusion can only be a probabilistic consequence of the premises if either the premises are themselves inconsistent or the conclusion is tautological.

There are some particular inferences with conditional consequences which are probabilistically reasonable, among which we list the following (let us abbreviate 'probabilistic consequence' by 'p.c.'):

Some probabilistically reasonable inferences

Let φ, ψ and η be truth-functional formulas.

1. If φ is tautologically equivalent to ψ then $\psi \rightarrow \eta$ is a p.c. of $\varphi \rightarrow \eta$.
2. φ is a p.c. of $T \rightarrow \varphi$ and conversely.
3. If φ tautologically implies ψ then $\varphi \rightarrow \psi$ is a p.c. of the empty set.
4. $\varphi \vee \psi \rightarrow \eta$ is a p.c. of $\varphi \rightarrow \eta$ and $\psi \rightarrow \eta$.
5. $\varphi \rightarrow \eta$ is a p.c. of $\varphi \vee \psi \rightarrow \eta$ and $\psi \rightarrow -\eta$.
6. $\varphi \rightarrow \psi$ is a p.c. of $\varphi \rightarrow \psi$ & η.
7. $\varphi \rightarrow \psi$ & η is a p.c. of $\varphi \rightarrow \psi$ and $\varphi \rightarrow \eta$.
8. $\varphi \rightarrow \eta$ is a p.c. of $\varphi \rightarrow \psi$ and φ & $\psi \rightarrow \eta$.

Of course, these schemata of inference are all valid in the truth-functional calculus. It will be noted that schemata 5 and 8 are weakened versions of the frequently used rules that $p \vee q \rightarrow r$ entails $p \rightarrow r$, and $p \rightarrow q$ and $q \rightarrow r$ entails $p \rightarrow r$ (hypothetical syllogism). What is significant about this particular set of probabilistically reasonable inference schemata is that it is *complete* in the sense that using these as rules of inference in a natural deduction system is sufficient to yield all probabilistically reasonable consequences of any finite set of premises.

Let us note, in concluding this section, some inferences which are not probabilistically reasonable. These are easily determined, using the result stated in Consequence 2. We have already noted the hypothetical syllogism, and the law that $p \vee q \rightarrow r$ entails $p \rightarrow r$. Clearly, in view of our earlier discussion, $-p \rightarrow q$ is not a p.c. of p, or of the disjunction $p \vee q$ (any situation under which p is true verifies $p \vee q$ but not $-p \rightarrow q$). Likewise, the law of contraposition, inferring $-q \rightarrow -p$ from $p \rightarrow q$ is not probabilistically reasonable, since verifying $p \rightarrow q$ actually precludes verifying $-q \rightarrow -p$. We will attempt to explain in the next section why most of these inferences might be said to be 'normally reasonable' in spite of the fact that it can happen that the probabilities of the premises can be high but those of the conclusions low. For the present, we record the following remarkable fact: if any of the above laws are added to those which have been listed as probabilistically reasonable, as rules of inference, then the inferences which they allow include all tautological consequences. The same, it might be added, applies to the rule of conditionalization: if ψ can be inferred from S and φ, then $\varphi\rightarrow\psi$ can be inferred from S. Thus, if S tautologically implied a conditional, say $\varphi\rightarrow\psi$, then S and φ would tautologically imply ψ, and, in virtue of Corollary 2.1, S and φ would *probabilistically* imply ψ, so, if the principle of conditionalization held, S would probabilistically imply $\varphi\rightarrow\psi$; hence probabilistic consequence would coincide with tautological consequence. What this shows, we think, is how very strong the 'logical forces' are which lead to the identification of 'if-then' with the material conditional.

14. Probably the most useful employment of the calculus of probability is in determining the nature of the situations in which rules of inference, like the hypothetical syllogism, which are valid in truth-functional logic, can lead from premises with high probability to conclusions with low probability. To carry out this kind of inquiry we have developed a routine procedure which can be used to determine certain essential aspects of these situations, and which can be used as a guide in the search for actual 'counterexamples' to the rules of inference in question. This method was used in fact to obtain examples F3, F4 and F6. The detailed description of the method would be too tedious to justify inclusion here, but we will describe the intuitive idea upon which it is based, as it applies to a particular example.

Consider contraposition as a rule of inference: i.e. that $-q\rightarrow -p$ can be inferred from $p\rightarrow q$. The atomic sentences p and q can be associated with sets (the sets of states of affairs in which these sentences would be true), and these sets represented by a Venn Diagram in the usual way (see diagram at right). The probabilities of p and q and their sentential combinations can be regarded as being represented by the areas of the corresponding regions in the diagram. Thus, the probability of $p\&q$ is proportional to area c and that of $p\&-q$ is proportional to area b in the diagram. We

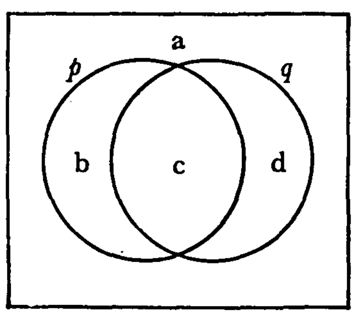

can represent the fact that $p \rightarrow q$ has a high probability by requiring that the probability of p&-q be much smaller than that of p&q: i.e. by requiring that area b be much smaller than area c (which we will indicate by writing 'b$\ll$c'). The fact that -$q \rightarrow$ -p has high probability corresponds to the condition that the probability of -q&p is much smaller than that of -q&-p: i.e. that b$\ll$a (b is much smaller than a). If the fact that $p \rightarrow q$ had high probability guaranteed that -$q \rightarrow$ -p also had, it would also have to be the case that b$\ll$c entailed b$\ll$a. But clearly this is not the case, since it is easy to describe a situation in which b$\ll$c but not b$\ll$a. What would such a situation be like? It would clearly have to be one in which both a and b were much smaller than c, and therefore were smaller than c+d: i.e. in this situation, the probability of -q (corresponding to a and b) would have to be much smaller than that of q. Hence, we can assert: in any situation in which $p \rightarrow q$ has a high probability and -$q \rightarrow$ -p has a low probability, q must have a high probability.

Do such situations as described above arise in practice? The answer is yes. They are ones in which one would assert both $p \rightarrow q$, and also q outright. Under such circumstances, however, it seems that correct usage demands the employment of the locution 'even if', rather than 'if then'. Thus, in asserting:

The game will be played, even if it rains.

one implies both

The game will be played.

and

If it rains, the game will be played.

Clearly the inference of the contrapositive

If the game is not played, it will not have rained.

would not be reasonable from such a pair of premises.

The observation that standards of correct usage demand the employment of 'even if' in situations in which a speaker is in a position to assert both q and $p \rightarrow q$ explains, we think, why the contrapositive inference -$q \rightarrow$ -p from $p \rightarrow q$ is normally warranted: i.e. if saying 'if p then q' is not correct when one is in a position to assert q outright, then one can legitimately infer -$q \rightarrow$ -p from $p \rightarrow q$ correctly and justifiably asserted. Note, however, that this standard of correctness is one we have not represented in our analysis of the conditions of justification of assertions. To attempt to represent it formally would lead to a far more radical revision of the traditional laws of logic than we have here considered. The reason is that standards of correctness such as that one should not assert 'if p then q' *simpliciter* when one is in a position to assert q outright are ones which demand the withdrawal of assertions in the light of evidence, even where the evidence in one sense does not conflict with the original assertion.[10]

Very similar remarks apply to the inference of the conditional $-p \rightarrow q$ from the disjunction pvq (already briefly discussed in section 9), and also show what is 'wrong' in the argument given in section 3 to derive $-p \rightarrow q$ from p. The Venn Diagram analysis applied to the inference of $-p \rightarrow q$ from pvq shows that the only situation under which pvq has high probability but $-p \rightarrow q$ has low probability is that in which p also has a high probability: i.e. the inference of $-p \rightarrow q$ from pvq is only unwarranted in the case in which the first member of the disjunction, p, can be asserted outright. As we noted in section 9, if standards of correct usage demand that disjunctions not be asserted when one is in a position to assert one of their members outright, then the inference of $-p \rightarrow q$ from pvq, correctly and justifiably asserted, is reasonable.

The derivation of $-p \rightarrow q$ from p given in section 3 consisted in two steps: (1) to derive pvq from p, and (2) to derive $-p \rightarrow q$ from pvq. What our analysis has shown is that the 'error' lies in step 2, but this is an error of a rather odd sort. Under normal circumstances of correct usage the inference of $-p \rightarrow q$ from pvq is reasonable, but pvq is here derived from premise p which justifies the assertion of pvq in the betting sense, but in fact makes the 'correct' use of pvq unjustified.

It is possible to analyze the other ordinarily accepted rules of inference for conditionals, such as the hypothetical syllogism, in a manner similar to the analysis of contraposition and the inference of conditionals from disjunctions. The analysis is quite complicated, and in any case leads to rather similar conclusions as those reached about contraposition. The upshot is to yield an explanation as to why, though these rules are not reasonable in our 'betting' sense, they are nevertheless usually accepted without question. The rule of conditionalization deserves special comment.

15. As noted at the end of section 13, the rule of conditionalization does not hold for reasonable inference: if S is a set of premises and φ and ψ are truth-functional formulas, then the fact that ψ is a probabilistic consequence of S and φ does not guarantee that $\varphi \rightarrow \psi$ is a probabil-

istic consequence of S. There is, however, an informal argument to show that the rule of conditionalization *ought* to hold, even for probabilistic inference, which we shall examine. The argument goes as follows. Suppose that ψ is a probabilistic consequence of S and φ: i.e. suppose that on all occasions in which all the members of S and φ have high probability, so does ψ. To show that $\varphi \rightarrow \psi$ is a probabilistic consequence of S it should be sufficient to show that on all occasions on which the assertions of the members of S are all justified, the conditional bet that if φ then ψ is a good one. Now, consider any occasion on which all the members of S could be justifiably asserted: i.e. they all have high probabilities. The bet that if φ then ψ would be a good one in that situation in case one were justified in supposing that, in the event of φ's proving true, ψ would be very likely to prove true. But the fact that ψ is a probabilistic consequence of S and φ would appear to provide this justification, since, in the event of φ's proving true, all of the members of S *and* φ would have high probability, and therefore so would ψ.

It seems to us that what is wrong with the foregoing argument is its tacit assumption that, though all of the members of the set S were assumed to have a high probability, they would continue to have these high probabilities in the event of φ's proving true. That this need not be so becomes evident if one considers particular examples, and it should not be surprising in view of the fact that statements made with less than complete certainty are always liable to being withdrawn. If this analysis is correct, it is significant in two ways. First, by emphasizing once again the importance of the fact that probabilities may change, it brings out more sharply a basic limitation in the kind of analysis we have used up to now, which treats conditional probabilities as static. Indeed, in view of this limitation, it may seem surprising that our calculus is somehow able to account for the fact that the principle of conditionalization sometimes fails. That is, our analysis of the informal argument to justify the principle of conditionalization seemed to suggest that its error lay in the tacit assumption that probabilities don't change — but our calculus does not represent probabilities as changing, so it ought to imply the validity of conditionalization. We have as yet no solution to this little puzzle which satisfies us, so we will just leave it here, with the comment that we think that to analyze it clearly should involve consideration of difficult questions concerning the 'intended interpretation' of the calculus of probability.

The second inference to be drawn from the argument for conditionalization and its analysis has to do with the nature of situations under which, though an inference involving conditionals is truth-functionally valid, nevertheless the assertions of its premises are probabilistically justified, but its conclusion is probabilistically unjustified.

Our discussion above leads to the following:

Consequence 3 (informal)

> Let S be a set of bettable particular or conditional statements, and let p and q be particular bettable statements such that $p \rightarrow q$ is a truth-functional consequence of S. If it is the case that all statements of S have high probability but $p \rightarrow q$ has low probability, then p is incompatible with S in the sense that in the event of p's proving true, not all statements of S would continue to have high probability.

The foregoing rule is satisfied in examples F1–F4 and F6, as well as in the counterexample to the principle of contraposition given in section 14. Note that it does not say that if an inference of a conditional from premises is truth-functionally but not probabilistically valid, then the premises *entail* the falsity of the antecedent of the conditional. What it suggests rather is that if it is *evident* (as it must be in convincing counterexamples) that the premises have high probability but the conclusion does not, then it must also be evident that some of the premises would have to be withdrawn in the event of the antecedent of the conditional's being found true.

16. We conclude this paper with a brief comment on some of its limitations. Two critical limitations have already been noted, which ought to be re-emphasized. The first is that our formal calculus applies only to particular bettable statements and to conditionals with such statements as antecedents and consequents. The conclusions of this analysis therefore apply only to such statements, and in particular not to compounds involving conditionals or other statements of non-bettable types. This is made clear in examples F7–F9. The second limitation lies in the fact that the use of the calculus of probability is not entirely appropriate to the representation of conditions of assertability even for bettable statements, especially in view of the fact that it treats probabilities as static. Its employment here was necessitated by the fact that it is the only existing calculus of probability with sufficient precision and generality to help us with the problem at hand.

The final limitation to be noted is the fact that, in dealing only with the question of what can reasonably be inferred from *assertions*, we have left aside uses of reasoning in *evaluation*. In applying logic in evaluation, one typically considers what follows from a supposition (i.e. the supposition that the proposition in question is true). Our analysis does not deal with what follows from suppositions, and, indeed, our argument in section 3 suggests that there is no clear sense in the supposition that a conditional is *true*. It seems to us significant that conditionals are seldom put as suppositions, and we hypothesize that when this does occur, in effect what is being supposed is that the conditional is *assertable*. Similarly, we suggest that *questions* concerning conditionals can be understood as properly answered in the affirmative if the questioned conditional is assertable: e.g. the question 'Is it the case that if Brown wins, Smith will retire?' is to be answered 'yes' just in case the assertion 'If Brown wins, Smith will retire' is justified.

Though conditionals are seldom put as suppositions, they are ubiquitous as consequences of suppositions — in which case they are no more asserted than are the suppositions themselves. The crucial lack in our analysis is, we believe, its failure to apply to inferences involving conditionals derived from suppositions. At the present we can do no more than note this lack, as showing the need for further study. We close with the suggestion that conditionals employed in this way probably obey rather different logical laws than do asserted ones, and that this, in fact, is what is at the heart of Lewis Carroll's 'Barbershop Paradox' [1].

NOTES

[1] This research was supported in part by a grant from the National Science Foundation to the Institute for Human Learning. The author would like to express his appreciation to many colleagues and students for helpful comments and criticisms, including Professors Benson Mates, Wallace Matson, David Rynin, John Searle and David Shwayder, and especially Dr. William S. Cooper, whose own theory of conditionals [2] disagrees with the one presented here, but whose own ideas on the subject have been very illuminating to the author.

[2] Talk of 'betting on a statement' (rather than on an event) is a slight perversion of ordinary usage which we think has no harmful consequences. We may, if we wish, simply define a bet on a statement as being a bet on the event which the statement says will take place. As will be seen, though, this must be modified in the case of bets on conditional statements.

[3] Of course, one can define 'extraordinary' ways of settling such bets at will: for example, we might agree that a bet on such a statement as 'there is no highest prime' won, in case the person making it could convince his opponent of its truth by an argument (i.e. a proof). The 'ordinary' way of settling a bet on a finite generalization (e.g. that all of the men in a room on an occasion had on red ties) is to examine all of its instances, and not to argue about its truth.

[4] In connecting conditional statements with conditional bets we are doing something akin to what Quine [5, p. 12] has suggested: that is to regard conditional statements as having truth values only in case their antecedents are true. On this view, conditional statements might more properly be described as statements made conditionally, the condition being that the antecedent be true, and the consequent being what is asserted subject to that condition. Seen in this way, nothing is asserted in case the antecedent is not 'fulfilled', just as nothing is bet (i.e. the bet is called off) in case its conditions are not fulfilled in a conditional bet, and in a conditional promise nothing is promised in case its conditions are not met.

[5] It is important to note that there are many conditional statements for which there are no ordinary conditions for settling bets on them. One very important class of these non-bettable conditionals consists of the subjunctive ones, and hence our analysis does not apply to them.

[6] A problem which will arise more acutely in connection with F6 is also implicit here. It may happen that each of two people who make a pair of contradictory assertions are justified, from their own points of view, or even that a single person *is justified within a short space of time in saying contradictory things: when he* gets new information and 'changes his mind'. We can do no more than note the problem here, though, and not solve it.

[7] It is worth noting that it is this kind of repeatable conditional which occurs usually in the standard definitions of so-called 'dispositional' terms: e.g. 'x is malleable $=_{df}$ if x is subject to moderate stress, it is permanently deformed'.

[8] This is a little over simple. There is a rather interesting logical phenomenon connected with conditionals whose antecedents are disjunctions. This is that such conditionals are not equivalent in meaning to the conditional statement obtained when the disjunction is replaced by another expression which is probabilistically and strictly equivalent to it. The two conditionals 'If that animal is a tiger it is dangerous' and 'If that animal is either a harmless tiger or not a harmless tiger, it is dangerous' furnish an example. What this shows is that often conditionals of the form 'if p or q, then r' are asserted as a conjunction of conditionals 'if p then r' and 'if q then r', and not as a single conditional. The foregoing example suggests that the single conditional is not equivalent to the conjunction, and we shall show later that they are not probabilistically equivalent since it can happen that, though the assertion of the single conditional may be justified, the assertion of one (but not both) of the conjuncts may not.

[9] Henceforth, we will ordinarily use Greek letters as variables ranging over truth-functional formulas, and capital roman letters as variables ranging over formulas.

[10] A theory of conditionals which does formally represent the situation in which statements have to be withdrawn in this way is presented in the Ph.D. dissertation of W. Cooper [2].

REFERENCES

1. Carroll, Lewis, 'A Logical Paradox', *Mind* N.S. 3, 1894, pp. 436–8.
2. Cooper, William S., *A Formalization of the Logic of Ordinary Conditionals*, Ph. D. Thesis. University of California 1964.
3. Kolmogorov, A., *Foundations of Probability*, trans. M. Morrison, Chelsea, New York 1950.
4. Lewis, C. I. and Langord, C. H., *Symbolic Logic*, Century, New York 1932.
5. Quine, W. V. O., *Methods of Logic*, Revised Edition, Holt, Rinehart and Winston, Inc., New York 1959.