

VII

INDICATIVE CONDITIONALS*

ROBERT STALNAKER

'EITHER the butler or the gardener did it. Therefore, if the butler didn't do it, the gardener did.' This piece of reasoning—call it the *direct argument*—may seem tedious, but it is surely compelling. Yet if it is a valid inference, then the indicative conditional conclusion must be logically equivalent to the truth-functional material conditional,¹ and *this* conclusion has consequences that are notoriously paradoxical. The problem is that if one accepts the validity of the intuitively reasonable direct argument from the material conditional to the ordinary indicative conditional, then one must accept as well the validity of many arguments that are intuitively absurd. Consider, for example, 'the butler did it; therefore, if he didn't, the gardener did.' The premiss of this argument entails the premiss of the direct argument, and their conclusions are the same. Therefore, if the direct argument is valid, so is this one. But this argument has no trace of intuitive plausibility. Or consider what may be inferred from the *denial* of a conditional. Surely I may deny that if the butler didn't do it, the gardener did without affirming the butler's guilt. Yet if the conditional is material, its negation entails the truth of its antecedent. It is easy to multiply paradoxes of the material

Robert Stalnaker, 'Indicative Conditionals' first published in *Philosophia* 5 (1975): 269–86, reprinted in *Ifs*, ed. W. L. Harper, R. Stalnaker and, G. Pearce, copyright © 1976 by D. Reidel Publishing Company. Reprinted by permission of Kluwer Academic Publishers.

* The ideas in this paper were developed over a number of years. During part of this time my research was supported by the National Science Foundation, grant number GS-2574; more recently it was supported by the John Simon Guggenheim Memorial Foundation.

¹ The argument in the opposite direction—from the indicative conditional to the material conditional—is uncontroversially valid.

conditional in this way—paradoxes that must be explained away by anyone who wants to defend the thesis that the direct argument is valid. Yet anyone who denies the validity of that argument must explain how an invalid argument can be as compelling as this one seems to be.

There are thus two strategies that one may adopt to respond to this puzzle: defend the material conditional analysis and explain away the paradoxes of material implication, or reject the material conditional analysis and explain away the force of the direct argument.² H. P. Grice, in his William James lectures,³ pursued the first of these strategies, using principles of conversation to explain facts about the use of conditionals that seem to conflict with the truth-functional analysis of the ordinary indicative conditional. I will follow the second strategy, defending an alternative semantic analysis of conditionals according to which the conditional entails, but is not entailed by, the corresponding material conditional. I will argue that, although the premiss of the direct argument does not *semantically entail* its conclusion, the inference is nevertheless a *reasonable inference*. My main task will be to define and explain a concept of reasonable inference which diverges from semantic entailment, and which justifies this claim.

Grice's strategy and mine have this in common: both locate the source of the problem in the mistaken attempt to explain the facts about assertion and inference solely in terms of the semantic content, or truth conditions, of the propositions asserted and

² This does not exhaust the options. Three other possible strategies might be mentioned. (1) Defend the direct argument, not by accepting the truth-functional analysis of the conditional, but by rejecting the truth-functional analysis of the disjunction. (2) Give a three-valued interpretation of the indicative conditional, assigning the neutral value when the antecedent is false. (3) Interpret the indicative conditional as a conditional assertion rather than the assertion of a conditional proposition. Alternative (1) might disarm this particular puzzle, but it seems *ad hoc* and would not help with other persuasive arguments for the material conditional analysis. Alternative (2) would conflict with some basic and otherwise plausible pragmatic generalizations such as that one should not make an assertion unless one has good reason to think that it is true. Alternative (3) seems to me the most promising and plausible alternative to the account I will develop, but to make it precise, I think one needs much of the framework of a pragmatic theory that I shall use in my account.

³ Photo-copies have been widely circulated; part of it has been published in: D. Davidson and G. Harman (eds.), *The Logic of Grammar* (Encino, Calif., 1975): 64–75.

inferred. Both attempt to explain the facts partly in terms of the semantic analysis of the relevant notions, but partly in terms of pragmatic principles governing discourse. Both recognize that since assertion aims at more than truth, and inference at more than preserving truth, it is a mistake to reason too quickly from facts about assertion and inference to conclusions about semantic content and semantic entailment.

My plan will be this: first, I will try to explain, in general terms, the concept of reasonable inference and to show intuitively how there can be reasonable inferences which are not entailments. Second, I will describe a formal framework in which semantic concepts like content and entailment as well as pragmatic concepts like assertion and inference can be made precise. Third, within this framework, I will sketch the specific semantic analysis of conditionals, and state and defend some principles relating conditional sentences to the contexts in which they are used. Fourth, I will show that, according to these analyses, the direct argument is a reasonable inference. Finally, I will look at another puzzling argument involving reasoning with conditionals—an argument for fatalism—from the point of view of this framework.

1

Reasonable inference, as I shall define it, is a pragmatic relation: it relates speech acts rather than the propositions which are the contents of speech acts. Thus it contrasts with entailment which is a purely semantic relation. Here are rough informal definitions of the two notions: first, reasonable inference: an inference from a sequence of assertions or suppositions (the premisses) to an assertion or hypothetical assertion (the conclusion) is *reasonable* just in case, in every context in which the premisses could appropriately be asserted or supposed, it is impossible for anyone to accept the premisses without committing himself to the conclusion; second, entailment: a set of propositions (the premisses) *entails* a proposition (the conclusion) just in case it is impossible for the premisses to be true without the conclusion being true as well. The two relations are obviously different since they relate different things, but one might expect them to be

equivalent in the sense that an inference would be reasonable if and only if the set of propositions expressed in the premisses entailed the proposition expressed in the conclusion. If this equivalence held, then the pragmatic concept of inference would of course have no interest. I shall argue that, and try to show why, the equivalence does not hold. Before discussing the specific framework in which this will be shown, let me try to explain in general terms how it is possible for an inference to be reasonable, in the sense defined, even when the premisses do not entail the conclusion.

The basic idea is this: many sentences are context-dependent; that is, their semantic content depends not just on the meanings of the words in them, but also on the situations in which they are uttered. Examples are familiar: quantified sentences are interpreted in terms of a domain of discourse, and the domain of discourse depends on the context; the referents of first- and second-person pronouns depend on who is speaking, and to whom; the content of a tensed sentence depends on when it is uttered. Thus context constrains content in systematic ways. But also, the fact that a certain sentence is uttered, and a certain proposition expressed, may in turn constrain or alter the context. There are two ways this may happen: first, since particular utterances are appropriate only in certain contexts, one can infer something about a context from the fact that a particular utterance is made (together with the assumption that the utterance is appropriate); second, the expression of a proposition alters the context, at the very least by changing it into a context in which that proposition has just been expressed. At any given time in a conversation, the context will depend in part on what utterances have been made, and what propositions expressed, previously in the conversation. There is thus a two-way interaction between contexts of utterance and the contents of utterances. If there are general rules governing this interaction, these rules may give rise to systematic relations between propositions expressed at different points in a conversation, relations which are mediated by the context. Such relations may become lost if one ignores the context and considers propositions in abstraction from their place in a discourse. It is because entailment relates propositions

independently of their being asserted, supposed, or accepted, while reasonable inference concerns propositions which are expressed and accepted that the two relations may diverge.

These general remarks are not an attempt to show that the notions of entailment and reasonable inference do in fact diverge, but only an attempt to point to the source of the divergence that will be shown. To show the divergence, I must say what contexts are, or how they are to be represented formally. I must say, for some specific construction (here, conditionals) how semantic content is a function of context. And I must state and defend some rules which relate contexts to the propositions expressed in them.

2

The framework I will use begins with, and takes for granted, the concept of a possible world. While model theory based on possible worlds is generally agreed to be a powerful and mathematically elegant tool, its intuitive content and explanatory power are disputed. It is argued that a theory committed to the existence of such implausible entities as possible worlds must be false. Or at least the theory cannot do any philosophical work unless it can provide some kind of substantive answer to the question, what is a possible world? Possible worlds are certainly in need of philosophical explanation and defence, but for the present I will make just a brief remark which will perhaps indicate how I understand this basic notion.⁴

It is a common and essential feature of such activities as enquiring, deliberating, exchanging information, predicting the future, giving advice, debating, negotiating, explaining, and justifying behaviour, that the participants in the activities seek to distinguish, in one way or another, among alternative situations that may arise, or might have arisen. Possible-worlds theory, as an explanatory theory of rational activity, begins with the notion of an alternative way that things may be or might have been (which is all that a possible world is) not because it takes this notion to be unproblematic, but because it takes it to be fundamental to the

⁴ See David Lewis, *Counterfactuals* (Cambridge, Mass., 1973): 84–91 for a defence of realism about possible worlds.

different activities that a theory of rationality seeks to characterize and relate to each other. The notion will get its content, not from any direct answer to the question, what is a possible world? or from any reduction of that notion to something more basic or familiar, but from its role in the explanations of such a theory. Thus it may be that the best philosophical defence that one can give for possible worlds is to use them in the development of substantive theory.

Taking possible worlds for granted, we can define a *proposition* as a function from possible worlds into truth values.⁵ Since there are two truth values, this means that a proposition is any way of dividing a set of possible worlds into two parts—those for which the function yields the value true, and those for which it yields the value false. The motivation for this representation of propositions is that, as mentioned above, it is an essential part of various rational activities to distinguish among alternative possible situations, and it is by expressing and adopting attitudes toward propositions that such distinctions are made.

How should a context be defined? This depends on what elements of the situations in which discourse takes place are relevant to determining what propositions are expressed by context-dependent sentences and to explaining the effects of various kinds of speech acts. The most important element of a context, I suggest, is the common knowledge, or presumed common knowledge and common assumption of the participants in the discourse.⁶ A speaker inevitably takes certain information for granted when he speaks as the common ground of the participants in the conversation. It is this information which he can use as a resource for the communication of further information, and against which he will expect his speech acts to be understood. The presumed common ground in the sense intended—the *presuppositions* of the speaker—need not be the beliefs which are really common to the speaker and his audience; in fact, they need

⁵ See M. J. Cresswell, *Logics and Languages* (London, 1973): 23–4, and Stalnaker, 'Pragmatics,' in G. Harman and D. Davidson (eds.), *Semantics of Natural Languages* (Dordrecht, 1972): 381–2 for brief discussions of the intuitive motivation of this definition of proposition.

⁶ For a fuller discussion and defence of this concept, see Stalnaker, 'Presuppositions', *Journal of Philosophical Logic*, 2 (1973): 447–57.

not be beliefs at all. The presuppositions will include whatever the speaker finds it convenient to take for granted, or to pretend to take for granted, to facilitate his communication. What is essential is not that the propositions presupposed in this sense be believed by the speaker, but rather that the speaker believe that the presuppositions are common to himself and his audience. This is essential since they provide the context in which the speaker intends his statements to be received.

In the possible-worlds framework, we can represent this background information by a set of possible worlds—the possible worlds not ruled out by the presupposed background information. I will call this set of possible worlds the *context set*.⁷ Possible worlds within the set are situations among which the speaker intends his speech acts to distinguish. I will sometimes talk of propositions being *compatible with* or *entailed by* a context. This means, in the first case, that the proposition is true in some of the worlds in the context set, and in the second case that the proposition is true in all of the worlds in the context set. Intuitively, it means, in the first case, that it is at least an open question in the context whether or not the proposition is true, and in the second case, that the proposition is presupposed, or accepted, in the context.

Propositions, then, are ways of distinguishing among any set of possible worlds, while context sets are the sets of possible worlds among which a speaker means to distinguish when he expresses a proposition.

3

The semantic analysis of conditionals that I will summarize here is developed and defended more fully elsewhere.⁸ The analysis was constructed primarily to account for counterfactual conditionals—

⁷ Elsewhere, I have called this set the *presupposition set*, but this terminology proved misleading since it suggested a set of presuppositions—propositions presupposed—rather than a set of possible worlds. The terminology adopted here was suggested by Lauri Karttunen.

⁸ Stalnaker, 'A Theory of Conditionals', in N. Rescher (ed.), *Studies in Logical Theory* (Blackwell: Oxford, 1968): 98–112, Essay II in this volume, and Stalnaker and R. H. Thomason, 'A Semantic Analysis of Conditional Logic', *Theoria*, 36 (1970): 23–42. See also Lewis, *Counterfactuals*. The formal differences between Lewis's theory and mine are irrelevant to the present time.

conditionals whose antecedents are assumed by the speaker to be false—but the analysis was intended to fit conditional sentences generally, without regard to the attitudes taken by the speaker to antecedent or consequent or his purpose in uttering them, and without regard to the grammatical mood in which the conditional is expressed.

The idea of the analysis is this: a conditional statement, *if A, then B*, is an assertion that the consequent is true, not necessarily in the world as it is, but in the world as it would be if the antecedent were true. To express this idea formally in a semantic rule for the conditional, we need a function which takes a proposition (the antecedent) and a possible world (the world as it is) into a possible world (the world as it would be if the antecedent were true). Intuitively, the *value* of the function should be that world in which the antecedent is true which is most similar, in relevant respects, to the actual world (the world which is one of the *arguments* of the function). In terms of such a function—call it '*f*'—the semantic rule for the conditional may be stated as follows: a conditional *if A, then B*, is true in a possible world *i* just in case *B* is true in possible world $f(A,i)$.⁹

It may seem that little has been accomplished by this analysis, since it just exchanges the problem of analysing the conditional for the problem of analysing a semantic function which is equally problematic, if not more so. In one sense this is correct: the analysis is not intended as a reduction of the conditional to something more familiar or less problematic, and it should not satisfy one who comes to the problem of analysing conditionals with the epistemological scruples of a Hume or a Goodman. The aim of the analysis is to give a perspicuous representation of the formal structure of conditionals—to give the *form* of their truth conditions. Even if nothing substantive is said about how antecedents select counterfactual possible worlds, the analysis still has non-trivial, and in some cases surprising, consequences for the logic of conditionals.

⁹ If *A* is the impossible proposition—the one true in *no* possible world—then there will be no possible world which can be the value of the function, $f(A,i)$, and so the function is left undefined for this case. To take care of this special case, the theory stipulates that all conditionals with impossible antecedents are true.

But what more can be said about this selection function? If it is to be based on *similarity* in some respect or other, then it must have certain formal properties. It must be a function that determines a coherent ordering of the possible worlds that are selected. And, since whatever the respects of similarity are that are relevant, it will always be true that something is more similar to itself than to anything else, the selection function must be one that selects the actual world whenever possible, which means whenever the antecedent is true in the actual world. Can anything more substantive be said about the relevant respects of similarity on which the selection is based? Not, I think, in the *semantic* theory of conditionals. Relevant respects of similarity are determined by the context, and the semantics abstracts away from the context by taking it as an unexplained given. But we can, I think, say something in a pragmatic theory of conditional statements about how the context constrains the truth conditions for conditionals, at least for indicative conditionals.

I cannot *define* the selection function in terms of the context set, but the following constraint imposed by the context on the selection function seems plausible: if the conditional is being evaluated at a world in the context set, then the world selected must, if possible, be within the context set as well (where C is the context set, if $i \in C$, then $f(A, i) \in C$). In other words, all worlds within the context set are closer to each other than any worlds outside it. The idea is that when a speaker says 'If A ', then everything he is presupposing to hold in the actual situation is presupposed to hold in the hypothetical situation in which A is true. Suppose it is an open question whether the butler did it or not, but it is established and accepted that whoever did it, he or she did it with an ice-pick. Then it may be taken as accepted and established that if the butler did it, he did it with an ice-pick.

The motivation of the principle is this: normally a speaker is concerned only with possible worlds within the context set, since this set is defined as the set of possible worlds among which the speaker wishes to distinguish. So it is at least a normal expectation that the selection function should turn first to these worlds before considering *counterfactual* worlds—those presupposed to be non-actual. Conditional statements can be directly relevant to their

primary uses—deliberation, contingency planning, making hedged predictions—only if they conform to this principle.

Nevertheless, this principle is only a defeasible presumption and not a universal generalization. For some special purposes a speaker may want to make use of a selection function which reaches outside of the context set, which is to say he may want to suspend temporarily some of the presuppositions made in that context. He may do so provided that he indicates in some way that his selection function is an exception to the presumption. Semantic determinants like domains and selection functions are a function of the speaker's intentions; that is why we must allow for exceptions to such pragmatic generalizations. But they are a function of the speaker's intention *to communicate* something, and that is why it is essential that it be conveyed to the audience that an exception is being made.

I take it that the subjunctive mood in English and some other languages is a conventional device for indicating that presuppositions are being suspended, which means in the case of subjunctive *conditional* statements, that the selection function is one that may reach outside of the context set. Given this conventional device, I would expect that the pragmatic principle stated above should hold without exception for indicative conditionals.

In what kinds of cases would a speaker want to use a selection function that might reach outside of the context set? The most obvious case would be one where the antecedent of the conditional statement was counterfactual, or incompatible with the presuppositions of the context. In that case one is forced to go outside the context set, since there are no possible worlds in it which are eligible to be selected. But there are non-counterfactual cases as well.¹⁰ Consider the argument, 'The murderer used an ice-pick. But if the butler had done it, he wouldn't have used an ice-pick. So the murderer must have been someone else.'¹¹ The subjunctive conditional premiss in this *modus tollens* argument

¹⁰ I was slow to see this despite the existence of clear examples in the literature. Comments by John Watling in a discussion of an earlier version of this essay helped me to see the point.

¹¹ This is Watling's example.

cannot be counterfactual since if it were the speaker would be blatantly begging the question by presupposing, in giving his argument, that his conclusion was true. But that premiss does not conform to the constraint on selection functions, since the consequent denies the first premiss of the argument, which presumably is accepted when the second premiss is given.

Notice that if the argument is restated with the conditional premiss in the indicative mood, it is anomalous.

My second example of a subjunctive non-counterfactual conditional which violates the constraint is adapted from an example given by Alan Anderson many years ago.¹² 'If the butler had done it, we would have found just the clues which we in fact found.' Here a conditional is presented as evidence for the truth of its antecedent. The conditional cannot be counterfactual, since it would be self-defeating to presuppose false what one is trying to show true. And it cannot conform to the constraint on selection functions since if it did, it would be trivially true, and so no evidence for the truth of the antecedent. Notice, again that when recast into the indicative mood, the conditional seems trivial, and does not look like evidence for anything.

The generalization that all indicative conditionals conform to the pragmatic constraint on selection functions has the following consequence about appropriateness-conditions for indicative conditionals: *It is appropriate to make an indicative conditional statement or supposition only in a context which is compatible with the antecedent.* In effect, this says that *counterfactual* conditionals must be expressed in the subjunctive. This follows since indicative conditionals are those which must conform to the constraint, while counterfactuals are, by definition, those which cannot.

I need just one more assumption in order to show that the direct argument is a reasonable inference—an assumption about conditions of appropriateness for making assertions. The generalization that I will state is a quite specific one concerning disjunctive statements. I am sure it is derivable from more general conversational principles of the kind that Grice has discussed, but since I am not sure exactly what form such general principles

¹² 'A Note on Subjunctive and Counterfactual Conditionals', *Analysis*, 12 (1951): 35–8.

should take, I will confine myself here to a generalization which has narrow application, but which can be clearly stated and easily defended. The generalization is this: *a disjunctive statement is appropriately made only in a context which allows either disjunct to be true without the other.* That is, one may say *A or B* only in a situation in which both *A and not-B* and *B and not-A* are open possibilities. The point is that each disjunct must be making some contribution to determining what is said. If the context did not satisfy this condition, then the assertion of the disjunction would be equivalent to the assertion of one of the disjuncts alone. So the disjunctive assertion would be pointless, hence misleading, and therefore inappropriate.¹³

4

All of the ingredients of the solution to the puzzle are now assembled and ready to put together. It may seem that this is a rather elaborate apparatus for such a simple puzzle, but each of the elements—propositions and contexts, the semantic analysis of conditionals, the pragmatic constraint on conditions, and the generalization about appropriateness—is independently motivated. It is not that this apparatus has been assembled just to solve the little puzzle; it is rather that the puzzle is being used to illustrate, in a small way, the explanatory capacity of the apparatus.

¹³ As with the pragmatic constraint on selection functions, there may be exceptions to this generalization. One exception is a statement of the form *A or B or both*. (I assume that the meaning of 'or' is given by the truth table for inclusive disjunction.) But statements which conflict with the principle must satisfy two conditions if they are to be appropriate. First, the statement must wear on its face that it is an exception so that it cannot be misleading. Second, there must be some explanation available of the purpose of violating the generalization, so that it will not be pointless. In the case of the statement *A or B or both*, it is clear from the logical relation between the last disjunct and the others that it must be an exception, so it satisfies the first condition. The explanation of the point of adding the redundant third disjunct is this: the disjunctive statement, *A or B*, requires that *A and not-B* and *B and not-A* be compatible with the context, but leaves open whether *A and B* is compatible with the context. The addition of the third disjunct, while adding nothing to the *assertive content* of the statement, does change the appropriateness-conditions of the statement, and thus serves to indicate something about the context, or about the presuppositions of the speaker.

The argument we began with has the form *A or B, therefore, if not-A, then B*. This inference form is a reasonable inference form just in case every context in which a premiss of that form could appropriately be asserted or explicitly supposed, and in which it is accepted, is a context which entails the proposition expressed by the corresponding conclusion. Now suppose the premiss, *A or B*, is assertable and accepted. By the constraint on the appropriateness of disjunctive statements, it follows that the context is compatible with the conjunction of *not-A* with *B*. Hence the antecedent of the conditional conclusion, *not-A*, is compatible with the context. Now it follows from the pragmatic constraint on selection functions that if a proposition *P* is compatible with the context, and another proposition *Q* is accepted in it, or entailed by it, then the conditional, *if P, then Q*, is entailed by it as well. So, since *not-A* is compatible with the context, and the premiss *A or B* is accepted, the conditional, *if not-A, then A or B*, must be accepted as well. But this conditional proposition entails the conclusion of the argument, *if not-A, then B*. So the inference is a reasonable one.

Since the argument works the other way as well, it follows that the indicative conditional and the material conditional are equivalent in the following sense: in any context where either might appropriately be asserted, the one is accepted, or entailed by the context, if and only if the other is accepted, or entailed by the context. This equivalence explains the plausibility of the truth-functional analysis of indicative conditionals, but it does not justify that analysis since the two propositions coincide only in their assertion and acceptance conditions, and not in their truth conditions. The difference between the truth conditions of the two propositions will show itself if one looks at acts and attitudes other than assertion and acceptance. To take the simplest case, it may be reasonable to deny a conditional, even when not denying the corresponding material conditional. For example, I know I didn't do it, so I know that it is false that if the butler didn't do it, I did. But since I don't know whether the butler did it or not, I am in no position to deny the material conditional, which is equivalent to the disjunction, either the butler did it or I did. I may even think that that disjunction is very probably true.

There are two other familiar inference forms involving conditionals which are judged to be reasonable, although invalid, by this analysis: contraposition and the hypothetical syllogism. It was one of the surprising consequences of the semantic analysis sketched above that these inferences are, in general, invalid. Nevertheless, these consequences count in favour of the semantic analysis rather than against it since there are clear counter-examples to both inference forms. But all the counter-examples involve subjunctive conditionals which are counterfactual conditionals whose antecedents are presupposed to be false. Now we can explain why there are no purely indicative counter-examples, and also why the arguments have the appearance of validity which they have. Both argument forms can be shown to be reasonable inferences, given that all conditionals involved are indicative, and given the assumption that indicative conditionals always conform to the pragmatic constraint on selection functions.¹⁴

5

I want to conclude by looking at a notorious argument involving indicative conditionals. The argument for fatalism is, I will argue, unreasonable as well as invalid. But it gains its appearance of force from the fact that it is an artful sequence of steps, each one of which has the form of a reasonable or of a valid inference. The trick of the argument, according to the diagnosis I will give, is that it exploits the changing context in an illegitimate way. Subordinate conclusions, legitimately drawn within their own subordinate contexts, are illegitimately detached from those contexts and combined outside of them. To make clear what I mean, let me sketch the argument. The specific form it takes, and the example used to present it, are taken from Michael Dummett's discussion of fatalism in his paper, 'Bringing about the Past'.¹⁵ The setting of the example is wartime Britain during an air raid. I reason as follows: 'Either I will be killed in this raid or I will not be killed.

¹⁴ Strictly, the inference to the contrapositive is reasonable only relative to the further assumption that the indicative conclusion is not inappropriate.

¹⁵ *Philosophical Review*, 73 (1964): 338-59.

Suppose that I will. Then even if I take precautions I will be killed, so any precautions I take will be ineffective. But suppose I am not going to be killed. Then I won't be killed even if I neglect all precautions; so, on this assumption, no precautions are necessary to avoid being killed. Either way, any precautions I take will be either ineffective or unnecessary, and so pointless.'

To give an abstract representation of the argument, I will let K mean 'I will be killed', P mean 'I take precautions', Q mean 'precautions are ineffective', and R mean 'precautions are unnecessary'. The argument, reduced to essentials, is this:

1. K or not- K
2. | K
3. | If P , K
4. | Q
5. | not- K
6. | if not- P , not- K
7. | R
8. Q or R

Now I take it that the main problem posed by this argument is not to say what is wrong with it, but rather to explain its illusion of force. That is, it is not enough to say that step x is invalid and leave it as that, even if that claim is correct. One must explain why anyone should have thought that it was valid. Judged by this criterion, Dummett's analysis of the argument does not solve the problem, even though, I think, what he says about the argument is roughly correct. Dummett argues that any sense of the conditional which will validate the inference from 2 to 3 (and 5 to 6) must be too weak to validate the inference from 3 to 4 (and 6 to 7). Hence, however the conditional is analysed, the argument as a whole cannot be valid. Dummett's argument to this conclusion is convincing, but it would be a full solution to the problem only if he supplemented it by showing that there *are* in our language distinct senses of the conditional that validate each of those steps. This I do not think he can do, since I do not think the force of the argument rests on an equivocation between two senses of the conditional.

According to the semantic and pragmatic analyses sketched

above, there is *one* sense of the conditional according to which the inference from 2 to 3 is a *reasonable inference*,¹⁶ and which is also strong enough to justify the inference from 3 to 4. The fallacy, according to the diagnosis, is thus in neither of the steps that Dummett questions. Both of the sub-arguments are good arguments in the sense that anyone who was in a position to accept the premiss, while it remained an open question whether or not the antecedent of the conditional was true, would be in a position to accept the conclusion. That is, if I were in a position to accept that I were going to be killed even though I hadn't yet decided whether or not to take precautions, then I would surely be reasonable to conclude that taking precautions would be pointless. Likewise if I knew or had reason to accept that I would not be killed.

The problem with the argument is in the final step, an inference which seems to be an instance of an unproblematically valid form—constructive dilemma which has nothing essential to do with conditionals. The argument form that justifies step 8 is this: A or B ; C follows from A ; D follows from B ; therefore, C or D . It is correct that the conclusion follows *validly* from the premiss provided that the sub-arguments are *valid*. But it is not correct that the conclusion is a *reasonable inference* from the premiss, provided that the sub-arguments are *reasonable inferences*. In the fatalism argument, the sub-arguments are reasonable, but not valid, and this is why the argument fails. So it is a confusion of validity with reasonable inference on which the force of the argument rests.

6

One final remark: my specific motivation for developing this account of indicative conditionals is of course to solve a puzzle, and to defend a particular semantic analysis of conditionals. But I have a broader motivation which is perhaps more important. That

¹⁶ As with contraposition, the inference from 2 to 3 is reasonable only relative to the further assumption that the conclusion of the inference is appropriate, which means in this case, only relative to the assumption that P , the antecedent of the conditional, is compatible with the context. This assumption is obviously satisfied since the setting of the argument is a deliberation about whether or not to make P true.

is to defend, by example, the claim that the concepts of pragmatics (the study of linguistic contexts) can be made as mathematically precise as any of the concepts of syntax and formal semantics; to show that one can recognize and incorporate into abstract theory the extreme context-dependence which is obviously present in natural language without any sacrifice of standards of rigour.¹⁷ I am anxious to put this claim across because it is my impression that semantic theorists have tended to ignore or abstract away from context-dependence at the cost of some distortion of the phenomena, and that this practice is motivated not by ignorance or misperception of the phenomenon of context-dependence, but rather by the belief that the phenomenon is not appropriately treated in a formal theory. I hope that the analysis of indicative conditionals that I have given, even if not correct in its details, will help to show that this belief is not true.

APPENDIX

Entailment and reasonable inference relate propositions and speech acts, respectively, but in both cases, given an appropriate language, one can define corresponding logical notions— notions of entailment and reasonable inference which relate formulas, or sentences independently of their specific interpretations.

Let L be a language which contains sentences. A *semantic interpretation* of the language will consist of a set of possible worlds and a function which assigns *propositions* (functions from possible worlds into truth-values) to the sentences, relative to *contexts*. The formal semantics for the language will define the class of legitimate interpretations by saying, in the usual way, how the interpretation of complex expressions relates to the interpretation of their parts. A *context* is an n -tuple, the first term of which is a *context set* (a set

¹⁷ I recognize, of course, that the definitions and generalizations presented here are nothing like a rigorous formal theory. But some parts of the apparatus (in particular, the semantics for conditionals) have been more carefully developed elsewhere, and I believe it is a relatively routine matter to state most of the definitions and generalizations which are new in precise model theoretic terms. Just to show how it might go, I will give in an appendix a very abstract definition of a logical concept of reasonable inference.

of possible worlds). The other terms are whatever else, if anything, is necessary to determine the propositions expressed by the sentences.

Notation: I will use P, P_1, P_2 , etc. as meta-variables for sentences, ϕ, ϕ_1, ϕ_2 , etc. as meta-variables for propositions (for convenience, I will identify a proposition with the set of possible worlds for which it takes the value true); k, k_1, k_2 , etc. will be variables ranging over contexts. $S(k)$ will denote the context set of the context k . $\|P\|_k$ will denote the proposition expressed by P in context k under the interpretation in question. (Reference to the interpretation is suppressed in the notation.)

Entailment. One may define several notions of entailment. The basic notion is a language-independent relation between propositions: ϕ_1 entails ϕ_2 if and only if ϕ_2 includes ϕ_1 . The *logical* concept of entailment, entailment-in- L , is a relation between sentences of L : P_1 entails P_2 if and only if for all interpretations and all contexts k , $\|P_1\|_k$ entails $\|P_2\|_k$. Logical entailment is entailment in virtue of the logical structure of the sentences. Similarly, the logical concept of reasonable inference will identify the inferences which are reasonable in virtue of the logical structure of the sentences.

Pragmatic interpretations. To define the logical notion of reasonable inference, we need to expand the concept of an interpretation. A *pragmatic interpretation* of L will consist of a semantic interpretation, an *appropriateness relation*, and a *change function*. The appropriateness relation A is a two-place relation whose arguments are a sentence of L and a context. $A(P, k)$ says that the assertive utterance of P in context k is appropriate. The change function g is a two-place function taking a sentence of L and a context into a context. Intuitively, $g(P, k)$ denotes the context that results from the assertive utterance of P in context k .

Since L is unspecific here, I leave these notions almost completely unconstrained, but it is easy to see how the generalizations about disjunctive and conditional statements would be stated as postulates which give some substance to these notions as applied to a language containing these kinds of statements. Just as the semantics for a specific language will include semantic rules specifying the elements of the context and placing constraints on the allowable semantic interpretations, so

the pragmatic theory for a specific language will include rules constraining the two distinctively pragmatic elements of a pragmatic interpretation, as well as the relations among the elements of the context.

I will give here just two constraints which will apply to any language intended to model a practice of assertion.

1. $A(P, k)$ only if $\|P\|_k \cap S(k) \neq 0$.

One cannot appropriately assert a proposition in a context incompatible with it.

2. $S(g(P, k)) = S(k) \cap \|P\|_k$.

Any assertion changes the context by becoming an additional presupposition of subsequent conversation. (In a more careful formulation the second of these would be qualified, since assertions can be rejected or contradicted. But in the absence of rejection, I think it is reasonable to impose this constraint.)

Both the appropriateness relation and the change function can be generalized to apply to finite sequences of sentences in the following way: Let σ be a finite sequence of sentences of L , P_1, P_2, \dots, P_n . Let k_1, k_2, \dots, k_n be a sequence of contexts defined in terms of σ and a context k as follows: $k_1 = k$; $k_{i+1} = g(k_i, P_i)$. Then $A(\sigma, k)$ if and only if, for all i from 1 to n , $A(P_i, k_i) \cdot g(\sigma, k) =_{at} k_n$.

Reasonable inference. The inference from a sequence of sentences of L , σ , to a sentence of L , P is *reasonable-in-L* if and only if for all interpretations and all contexts k such that $A(\sigma, k)$, $S(g(\sigma, k))$ entails $\|P\|_{g(\sigma, k)}$.

Note that there is no language-independent concept of reasonable inference analogous to the language-independent notion of entailment. The reason is that, while we have in the theory a notion of proposition that can be characterized independently of any language in which propositions are expressed, we have no corresponding non-linguistic concept of statement, or assertion. One could perhaps be defined, but it would not be a simple matter to do so, since the identity conditions for assertion types will be finer than those for propositions. The reason for this is that different sentences may have different appropriateness-conditions even when they express the same proposition.

VIII

LOGIC AND CONVERSATION

H. P. GRICE

IT is a commonplace of philosophical logic that there are, or appear to be, divergences in meaning between, on the one hand, at least some of what I shall call the FORMAL devices— $\sim, \wedge, \vee, \supset, (x) \exists (x) \int x$ (when these are given a standard two-valued interpretation)—and, on the other, what are taken to be their analogues or counterparts in natural language—such expressions as *not, and, or, if, all, some (or at least one), the*. Some logicians may at some time have wanted to claim that there are in fact no such divergences; but such claims, if made at all, have been somewhat rashly made, and those suspected of making them have been subjected to some pretty rough handling.

Those who concede that such divergences exist adhere, in the main, to one or the other of two rival groups, which for the purposes of this essay I shall call the formalist and the informalist groups. An outline of a not uncharacteristic formalist position may be given as follows: In so far as logicians are concerned with the formulation of very general patterns of valid inference, the formal devices possess a decisive advantage over their natural counterparts. For it will be possible to construct in terms of the formal devices a system of very general formulas, a considerable number of which can be regarded as, or are closely related to, patterns of inferences the expression of which involves some or all of the devices: Such a system may consist of a certain set of simple formulas that must be acceptable if the devices have the meaning that has been assigned to them, and an indefinite number of further formulas, many of them less obviously acceptable, each of which can be shown to be acceptable if the members of the original

H. P. Grice, 'Logic and Conversation', first published in *Syntax and Semantics*, Vol. 3, ed. Peter Cole and Jerry L. Morgan (Academic Press, 1975). Copyright Kathleen Grice and used with permission.