NOTES ON CONDITIONAL SEMANTICS*

Robert Stalnaker
Department of Linguistics and Philosophy
MIT
Cambridge, MA 02139

1. Introduction

Philosophers have been puzzling about conditional sentences, and conditional reasoning, at least since the time of the ancient Stoics. In this century, the problem was first raised by logicians trying to give a plausible account of the logic of conditionals, and by philosophers of science in the empiricist tradition trying to understand the relation between counterfactual conditionals and scientific laws, and the role of such conditionals in the formation of dispositional and theoretical concepts. Conditional sentences are problematic, first from an abstract point of view: they are, or at least seem to be, non-truth-functional, and so are not analyzable with the resources of extensional semantics, the only kind of semantics that is and unproblematic. One can, of course, define a truth-functional connective that has some of the properties of the conditional -the so-called material conditional - but it neither gives an intuitively plausible account of the logic and semantics of the conditional sentences of natural language that we find ourselves using, nor does it have the promise to do the conceptual work that we would like to use conditionals to do.

There are also more substantive philosophical reasons for finding conditionals problematic: On the face of it, counterfactual conditional sentences seem to be about unactualized possibilities, yet they also seem to be making contingent, factual claims - claims that must be true or false in virtue of actual facts about the world as it is, and that must be evaluated on the basis of empirical evidence gathered in the actual world. Since Hume first articulated these worries, philosophers in the empiricist tradition have been skeptical about the intelligibility of concepts of natural necessity and possibility, about claims about potentialities, propensities, and connections that seem to be statements about how things could or must be, and not just about how they are. Possibility and necessity might be all right so long as they concern purely conceptual relations or derive from linguistic conventions, but the empiricist argues that there cannot be necessity in the world, and that is what counterfactual conditionals seem to commit us to.

But while conditionals are problematic, they are also a central, and apparently ineliminable part of the conceptual resources we use to describe and find our way about in the world. We make contingency plans - what to do *if* certain situations arise - and we deliberate by considering what would be the best thing to do under various alternative hypotheses. And in justifying and evaluating actions, we may appeal to hypothetical possibilities that we know are not realized: I went around the pond rather than across it because the ice was thin and I would have fallen through. The Fed was justified in lowering interest rates, because

---

*This paper was the basis for a survey talk at the third TARK conference, March, 1990.

if it had not, there would have been a recession. In forming beliefs about the world, we prepare to receive evidence that we may or may not receive. Our dispositions to change what we believe in response to potential evidence are expressed in conditional beliefs: if he doesn't call by ten, then we'll know that he didn't miss the train; if the test result is negative, then we can rule out AIDS. And we appeal to counterfactual possibilities in reasoning about how to interpret evidence: If he had missed the train, he would have called, so since he didn't call, he must have made the train; If the gardener had done it, there would have been muddy footprints in the parlor, but there were not. In theorizing about the world, we find more powerful generalizations and deeper ways of categorizing things by considering not only how things actually behave, but how they would behave under normal or ideal conditions, or when subjected to certain conditions that they may or not in fact be subjected to. As predicted by Aristotelian physics, bodies in motion in the world as it is slow down and come to a stop unless a force continues to be applied to them. But contrary to what Aristotelian physics implies, they *would* continue in a straight line at constant velocity if acted on by no force at all. Things that never dissolve in water form a heterogeneous class, difficult to generalize about. But if we group the things that dissolved with those that didn't, but would have if they had been put to the test then we can get somewhere.

2. Conditional semantics

Empiricists were skeptical about counterfactual possibilities, but committed to making sense of scientific theory, and of practical and inductive reasoning. So they tried to analyze counterfactuals away - to find a way to understand them as only apparently about unrealized possibilities. But these attempts were unsuccessful, and were acknowledged to be unsuccessful by empiricists such as Nelson Goodman, whose penetrating criticisms of his own project (in 1955) showed why it was doomed to fail.[1] The developers of conditional semantics (starting about a decade later) adopted a different, more modest strategy: don't try to analyze away the problematic references to alternative possibilities; instead, take conditionals at face value and try to give a clear, precise and rigorous statement of what they seem, on the surface, to be saying. Ask what conceptual resources we need to postulate in order to give a straightforward interpretation of conditionals, and then help yourself to those resources, and use them to clarify the abstract structure of conditionals, and the relation between conditionals and the other problematic concepts that they are so deeply interconnected with. Perhaps this way we can get clearer about the role of conditionals in describing the world and reasoning about it even if we can't explain them away.

On the face of it, a statement such as "If Cuomo had run in 1988 he would have won" seems to be making a claim about a counterfactual possible situation in which Mario Cuomo runs for president in 1988. The claim it is making about this counterfactual situation is that in it, Cuomo is elected. To have a model theory to give a precise statement of this straightforward account, we need a set of possible situations or possible worlds, and a selection function that picks the situation in which the proposition expressed in the consequent of the conditional is said to be true. This selection function must have two arguments: first, the proposition expressed in the antecedent, and second the possible sit-

[1]Goodman, 1954. This is the classic statement of the problem of counterfactuals.

uation relative to which the conditional as a whole is being evaluated. The selection is dependent on the first argument - the antecedent - since the world selected represents the way things would have been if the antecedent were true. The selection is dependent on the second argument - the possible world from which the conditional is being evaluated - since conditionals may be contingent - true in some possible worlds and false in others - which implies that the possible world selected from one possible world may be different from the possible world selected from another.

So suppose we have a set of possible worlds and a function that takes a possible world i and a proposition P into a possible world j = f(i,P). If i is the actual world, then j is the world that would be actual if P were true. Now in terms of these resources, we can state a rule that gives the semantic value - the proposition expressed - by a conditional as a function of the propositions expressed by the antecedent and consequent. (For present purposes propositions may be identified with sets of possible worlds - intuitively, the worlds in which the proposition is true.) The rule is this: "If P, the Q" is true in possible world i if and only if Q is true in possible world f(i,P). This is pretty much all there is to the semantics. All that remains to be done is to make explicit the formal constraints on the selection function. Here are three plausible constraints: first, the function must always selects a possible world in which the antecedent is true (if there is one; if the antecedent is necessarily false, then the function is undefined, and a clause is added to the rule making the conditional vacuously true in that case). Second, the actual world (more generally, the world we are selecting from) must be selected if it is eligible (that is, if the antecedent is true there). The third condition is motivated by the following intuitive considerations: intuitively, it seems reasonable to take the counterfactual possibility one is talking about when one says "if P" to be a situation that differs from the actual situation only in ways that are required to make P true. If this is right, we can assume that the selection is based on an ordering of possible worlds with respect to the extent to which they differ from the actual world: the function selects the closest or most similar world from among those that are eligible. The abstract theory can say nothing substantive about the basis for the ordering, but the selection will be based in this way on some such ordering if and only if it conforms to the following condition: if Q is true in f(i,P) and P is true in f(i,Q), then f(i,P) = f(i,Q).

That is the whole of one semantic theory of conditionals, a theory developed in the late 1960's by Rich Thomason and myself.[2] There are a number of closely related theories, most notably one developed independently by David Lewis that takes the relevant ordering to be one that permits ties and infinite sequences of more similar worlds without limit.[3] Many find it more plausible to take the selection function to select a set of possible situations rather than a single situation. John Pollock has argued that the ordering should be assumed

---

[2]Stalnaker, 1968, is an informal exposition and philosophical defense of the theory. Stalnaker and Thomason, 1970, is a formulation of the semantics for a quantified conditional language and a completeness proof for an axiomitization of it. Stalnaker, 1984, ch. 7-8 is a later philosophical discussion of the theory.

[3]The fullest exposition of Lewis's theory is in his book, Lewis, 1973.

only to be a partial ordering that permits incomparabilities as well as ties.[4] Bas van Fraassen likes a theory that selects a single situation, but rejects the assumption that selection is based on any ordering.[5] Brian Chellas spelled out the semantics in its most general form.[6] But these other theories have very similar properties, and most of what I say will not depend on which of the variations one chooses.

This kind of theory does not solve the philosophical problems that motivated philosophers to try to analyze conditionals, but it does give those problems a new form and a sharper focus. It justifies a logic for conditionals, one that has some surprising properties. And it is a rich theory that raises new problems of interest, and provides a framework for exploring the relations between conditionals and other concepts such as knowledge and time, probability, both subjective and objective, causation, deliberation, explanation. Many of these other concepts have been analyzed within the possible worlds framework, which is therefore an ideal setting to explore the relationships between these concepts and conditionals.

3. The logic of conditionals

The logic of conditionals that is validated by this semantic theory contrasts with the logics for both the material conditional and the strict conditional. More generally, the logic is different from the logic of any conditional that can be analyzed as some monadic modal operator applied to a material conditional. Some inferences that are common to the logics of all such concepts, such as contraposition and the hypothetical syllogism, are invalid in conditional logic. The invalidity of these inferences is confirmed by intuitive counterexamples, as well as by general arguments that connect the inferences to each other. Consider the following counterexample to the hypothetical syllogism (if A then B; if B then C, therefore if A then C): if Bush had lost in 1988, Dukakis would have become president; if the Communist party candidate had won, Bush would have lost; therefore, if the Communist party candidate had won, Dukakis would have become president. The failure of the hypothetical syllogism is closely connected with the following nonmonotonicity property exhibited by conditional logic: From *if A, then C*, it does not follow that *if A and B, then C*. Intuitive examples showing that ordinary conditionals have this property are easy to construct, and the abstract semantics explains why one should expect conditionals to have it: the closest A-world may not be the same as the closest A&B world, since one may have to look farther to make both A and B true. Since the conditionals we accept are connected with our inductive practices and policies - our ways of responding to evidence - this nonmonotonicity property in conditional logic is connected with the nonmonotonic character of inductive inference, and conditional semantics should help to clarify that feature of reasoning.

The failure of monotonicity implies the invalidity of another pattern of inference ordinarily associated with conditionals: the inference of contraposition, and this feature of conditional logic also finds intuitive support. Con-

---

[4]Pollock, 1976. cf. Ginsberg, 1986.

[5]van Fraassen, 1976b.

[6]Chellas, 1975.

sider the following counterexample: If Dukakis had carried California, he (still) would have lost the election. Therefore, if Dukakis had won the election, he (still) would not have carried California. Again, the abstract semantics helps explain the failure: that B is true in the closest A world says nothing about what must be true in the nearest world in which B is false.

These properties - nonmonotonicity and the failure of contraposition and hypothetical syllogism - are common to all of the logics validated by the variations of conditional semantics mentioned above. Other patterns of inference distinguish the different theories, and remain controversial within the context of conditional semantics: for example, a theory that assumes a selection function that selects a single world validates the principle of conditional excluded middle, *if A, then B, or if A, then not-B*, whereas in a theory with a set selection function this principle is invalid. Arguments about this difference concern both the interpretation of particular examples, and the plausibility of the different assumptions needed to interpret the examples.[7]

4. Subjunctive vs. indicative

The semantic analyses of conditionals were directed primarily at counterfactual conditionals - conditional statements whose antecedents are incompatible with what is presupposed in the context, which is to say statements about a possible situation that is presupposed to be nonactual. Or sometimes the kind of conditional that is the primary concern is characterized grammatically: the subjunctive conditional. There is some controversy about just how to make the distinction[8] - the grammar is complicated and not very well understood - but there is a clear contrast between the paradigm cases of what have been called subjunctive and indicative conditionals, and it is clear from examples first proposed by Ernest Adams[9] that the contrast is semantically significant. Here is one such example of a contrasting pair:

If Shakespeare didn't write Hamlet, someone else did.

If Shakespeare hadn't written Hamlet, someone else would have.

The first seems obviously true, but the second would be affirmed only by someone who had a bizarre theory of the identity and destiny of works of literature.

Assuming that the possible worlds semantic analysis is appropriate for subjunctive conditionals, what kind of account is appropriate for the indicatives, and how are the two kinds of conditionals related? At least three kinds of answers to the first of these questions have been proposed. First, some theorists have defended the material conditional analysis for indicative conditionals[10],

---

[7]Stalnaker, 1981.

[8]see Dudman, 1984 and Bennett, 1989.

[9]see Adams, 1970.

[10]Grice, 1989, ch. 4 and Thomson, 1990 are defenses of this thesis written in the 1960's. Lewis, 1976 defends it in the context of a defense of the thesis that conditionals are assertible when the conditional probability is high.

using pragmatic theory - theory of conversation - to explain why such conditionals are often not appropriately asserted, even though they are true, according to the analysis. Second, some have denied that indicative conditionals express propositions at all[11], arguing that a conditional assertion is a hedged assertion of the consequent; one version of this kind of theory uses conditional probability to explain the assertability conditions for indicative conditionals. Third, some have argued that the possible worlds semantic analysis applies to both kinds of conditionals.[12] The last answer is compatible with a semantic difference between the two kinds of conditionals, since the grammatical difference might indicate different constraints on the selection function relative to which the conditional is interpreted.

Whichever kind of account one favors, it seems clear that indicative conditional sentences are normally used to express something about the speaker's epistemic situation rather than to state a fact that is independent of what the speaker does and does not know. The possible worlds analysis accommodates this fact by tying the selection function used to interpret indicative conditionals to the speaker's state of knowledge or belief. For an indicative conditional to be appropriate, the antecedent must be compatible with the speaker's knowledge, or at least with what the speaker presupposes, and the possible situation selected must be compatible with that state of knowledge, or presupposition. I can say "if Shakespeare didn't write Hamlet" only in a context in which it is an open question whether he did or not, and the relevant situation in which he didn't must be compatible with everything we know, or take for granted, in the context. Because even in a context where it is an open question whether Shakespeare wrote Hamlet, we take for granted that the play Hamlet was written by someone, it follows that if Shakespeare didn't write it someone else did. But in the case of the subjunctive, the situation envisioned is counterfactual, and so need not be compatible with what is known or presupposed.

Because of the epistemic role of indicative conditionals, their interpretation will be highly context sensitive. If they express propositions in the way proposed in the semantic theory of conditionals, they will express different propositions at different times and in different settings. Conditionals that appear to contradict each other may be true together if they are taken from different contexts. Consider this example, used by Kenneth Warmbrod[13] to illustrate the context dependence of indicative conditionals: An electrical circuit contains a power source, a lightbulb, and two switches wired in series. Two observers know this, and know that everything is in working order. The first knows also that switch b is closed, but cannot see switch a, or the light. The second observer sees that the light is off, but has no information about either switch. The first observer will assent to the conditional claim that *if switch a is closed, then the light is on*, and will dissent from the claim that *if switch a is closed, then the light is off*. The second observer will agree that *if switch a is closed, then switch b is open*, but will deny that *if switch a is closed then b is also closed.*

Jackson, 1987, is a defense of a modified version of the thesis.

[11]See Adams, 1975, Appiah, 1985, and Gibbard, 1981.

[12]See Stalnaker, 1975.

[13]Warmbord, 1983.

So the first observer seems committed to saying that if a is closed, b is closed and the light is on, while the second observer is committed to denying this. The two seem to disagree, but neither's statement rests on any mistake or false belief. Both, it seems, are right. The contradiction is only apparent, resolved by a contextual difference - a difference in the epistemic states relative to which the conditionals are interpreted. So it is like the apparent contradiction in the magician's true statement, "now you see it, now you don't."

5. Conditionals and probability

One kind of theory of indicative conditionals - developed in detail by Ernest Adams - uses conditional probability (where probability is interpreted as degree of belief) to give assertability conditions for conditional assertions. Ordinary statements are assertible when they have high probability; conditionals are assertible when the conditional probability of consequent on antecedent is high. Adams did not take such conditionals to make statements with truth conditions, but one might ask what kind of truth conditions conditionals would have to have in order to make the probability of the truth of a conditional coincide with the conditional probability of consequent on antecedent. There are some striking similarities between the logic that Adams extracts from his assumption, and the logic validated by the possible worlds semantics for conditionals, and the assumptions about the epistemic role of conditionals that were used to motivate the possible worlds theory are closely related to Adams's assumption. So there was some reason to hope that a probability conditional - a conditional for which $Pr(A>B) = Pr(B/A)$ - might coincide with a possible worlds conditional. But David Lewis showed, in a series of elegant and clarifying arguments, that there could not be any conditional at all that met this requirement for any class of probability functions meeting certain natural conditions. However conditionals are interpreted, if they express propositions and have truth conditions at all, then the relation between conditional statements and conditional probabilities will have to be more complicated than we might like to suppose. I will state Lewis's main result more precisely, and then try to explain the intuitive idea behind his proof.

A conditional is a *probability conditional*, relative to a probability function, or to a class of probability functions, if the probability identity, $Pr(A>B) = Pr(B/A)$ holds in that probability function, or in all the probability functions in the class, for all A and B for which $Pr(B/A)$ is defined (that is, for all A such that $Pr(A) \neq 0$). What Lewis proved was that there cannot be a probability conditional for any class of probability functions that is (1) closed under conditionalization, and (2) nontrivial in the sense that for at least one probability function in the class there are at least three disjoint propositions that have nonzero probability.

The reason there cannot be a probability conditional satisfying these constraints is that the constraints require that the same conditional be tied to different probability functions - to a probability function and also to all those defined by conditionalizing on it - in a way that is impossible. The reason this is impossible is that a conditional probability, $Pr(C/A)$, depends only on the measure on that part of the space in which A is true, but since the probability of the conditional, $Pr(A>C)$, is an absolute probability, it depends on the measure on the whole space. The strategy of Lewis's argument, roughly speaking, is to define, by conditionalization, different probability functions for which a conditional probability must be the same, but the corresponding probability of a conditional must be different. If you define a new probability function, Pr', by conditionalizing on some condition that is entailed by A, then the conditional probability Pr'(C/A), must remain the same as the original conditional probability. But it cannot be that the probability of the conditional is the same in *all* new probability functions defined in this way. One easy way to see that this is impossible is to conditionalize on the disjunction, A ∨ ~(A>C).

So the argument depends crucially on a comparison of the probabilities of a conditional, A>C, in different probability functions - a prior probability function and the posterior probability functions defined by conditionalizing on it. The argument assumes, implicitly, that the conditional sentence expresses the same proposition in the context of these different probability functions. This assumption is implicit in the identification of the conditional probability, Pr(A>C/X), where "Pr" is the prior probability function, with Pr'(A>C), where "Pr'" is the posterior probability function defined by conditionalizing on X. But the conditionals for which the probability identity is appropriate are indicative conditionals, and these conditionals we have independent reason to believe are context dependent - more specifically, dependent for their interpretation on the beliefs and presuppositions of the person using the conditional sentence. The probability identity can be seen as one way of trying to spell out the way in which the interpretation of the conditional is constrained by beliefs, since the probability function represents a belief state. What Lewis's arguments show is that we cannot have a conditional that is constrained by the agent's beliefs in the way required by the probability identity and that also has truth conditions that are independent of the epistemic context in which it is interpreted. But if one allows the conditional to express different propositions relative to different probability functions, then there can be a probability conditional for nontrivial probability functions. In fact Bas van Fraassen has shown that any probability function defined on the nonconditional sentences of a language can be extended to the conditional sentences in a way that validates the probability identity.[14] Van Fraassen's construction makes use of a possible worlds-selection function semantics for the conditional.

There are qualitative analogues of Lewis's triviality arguments. Peter Gardenfors has used such an argument, in the context of his theory of belief revision, to cast doubt on the *Ramsey test*, a qualitative analogue of the probability identity.[15] The Ramsey test is the thesis that a conditional is accepted in a given belief state if and only if the consequent would be accepted in the revised belief state that would be induced by receiving as total new information the antecedent. Gardenfors shows that the Ramsey test implies an untenable monotonicity condition for belief revision, but his argument, like Lewis's, implicitly assumes that the interpretation of the conditional does not change when beliefs are revised - an assumption that is not plausible for indicative conditionals.

6. Connections and applications

The pure abstract semantic theory of conditionals, with possible worlds represented as simple unstructured points, can say little about the basis for the selection or ordering of possible worlds relative to which conditionals are in-

[14]van Fraassen, 1976a.

[15]Gardenfors, 1988.

terpreted. But in the context of a possible worlds framework in which other concepts are analyzed as well, or in applications of such a framework in which possible worlds are defined in terms of more primitive elements of some structure (for example, paths through a branching tree of temporal nodes, runs of a distributed system, complete plays of a game), one may have the resources to say something more substantive about the truth conditions for conditionals, and about the relations between conditionals and the other concepts with which they interact. In conclusion I will mention some of the work of this kind that has been done, and some possibilities that I think are worth exploring.

The context dependence of indicative or epistemic conditionals that I have been discussing can be made precise in a semantic framework that provides a representation of context - of the beliefs, presumptions and presuppositions of the participants in a conversation - and of the way context changes in response to the events that take place in an extended discourse. In such a theory one can separate the semantic from the pragmatic elements in an explanation of inferential relations between conditional statements made at different points in a conversation.

The interaction between counterfactual conditionals and temporal concepts is explored in the setting of a semantic theory of time dependent necessity and possibility by Rich Thomason and Anil Gupta.[16] The branching time structure, in which possible worlds are partly overlapping paths through the tree, suggests some very natural constraints on the way possible worlds should be ordered for the purposes of selection, and these constraints have consequences for the semantic relations between conditionals and statements of time dependent possibility.

The branching time framework is the natural setting for a representation of chance - objective probability. Bas van Fraassen examines the relations between conditionals and chance in such a framework.[17] Brian Skyrms has explored the relations between conditionals and both objective and subjective probability.[18] Among other things, he connects the truth conditional semantic theory to Ernest Adams's probabilistic theory by constructing a more general framework in which both are special cases.

Conditionals - both epistemic and counterfactual - obviously play an important role in deliberation and decision making. A Bayesian framework in which conditionals are among the propositions that are believed to various degrees, and in which the probabilities of conditionals may diverge from conditional probabilities, helps to explain the role of causal beliefs in practical reasoning. Allan Gibbard and Bill Harper used such a framework in the first formulation of causal decision theory, a theory formulated in other terms by Brian Skyrms, and by David Lewis.[19] The basic assumption of causal decision theory is that one should decide what to do on the basis of one's beliefs about the causal consequences of one's alternative actions, and so one needs to separate causal from merely evi-

---

[16]Thomason and Gupta, 1980.

[17]van Fraassen, 1980.

[18]Skyrms, 1980 and 1984.

[19]Gibbard and Harper, 1978, Skyrms 1980, and Lewis, 1981.

dential components of one's degrees of belief. Suppose that the choosing of a certain action (say refraining from smoking) would be positively relevant to a certain desirable future state of affairs (say one's not getting lung cancer) according to your degrees of belief. But suppose this evidential relevance is explained, not by a belief that smoking tends to cause lung cancer, but by a belief that smoking is a symptom of a genetic predisposition to get lung cancer. In this case, according to the causal decision theorist, it would be unreasonable to decide not to smoke on the grounds that this decision would strengthen your belief that you would avoid lung cancer. The relevant probability to use in your Bayesian reasoning is not the conditional probability of the outcome, conditional on the hypothesis that you choose the action, but the probability of the (subjunctive) conditional, if you were to choose the action, then the outcome would obtain. Even where smoking is evidentially relevant, if it is believed to be causally independent, then you will believe the conditional, *if I were to smoke, I would get lung cancer* to the same degree as you believe the consequent, *I will get lung cancer*. The distinctions that causal decision theory uses conditionals to make are obviously relevant to game theoretic reasoning, where one's choice of strategy is known to be causally independent of the choices of one's partners or opponents, but where one's own choices may still be evidentially relevant to their choices.[20]

In conclusion, I want to mention one final potential application of conditional semantics that is relevant to the theoretical aspects of reasoning about knowledge: to the general foundations of nonmonotonic logic. Any semantic theory - any theory that defines a set of models and gives rules saying which sentences of some language are true in which models - will determine a relation of semantic consequence defined as follows: a sentence A is a consequence of a set of sentences S if and only if all models of S are also models of A. No matter what the models are, and no matter what the semantic rules interpreting the language are like, this consequence relation will be monotonic. One can give an equally abstract and general definition of a nonmonotonic consequence relation: A sentence A is a nonmonotonic consequence of a set of sentences S if and only if all *preferred* models of S are models of A.[21] The particular nonmonotonic semantics will spell out the relevant relation of preference, but however it is done, it will determine some kind of selection function selecting as a function of S some, but not necessarily all, of the models satisfying S. There are obvious parallels between such a function and the selection functions of conditional semantics, and they explain the parallels that have been noted between the abstract structure of nonmonotonic consequence relations and the structure of conditional logic. The preference relation might or might not determine an ordering of models, just as the selection function in a conditional semantic theory might or might not determine an ordering of possible worlds. And in both cases, the ordering might be partial or total. Different decisions about the constraints on the selection function in conditional semantics will correspond to different kinds of preference concepts in the semantics of nonmonotonic logic, and will validate different principles of nonmonotonic consequence. For example, consider the following principle of nonmonotonic consequence, which has been called "rational monotonicity" ("$\vdash$" represents nonmonotonic consequence):

[20]See Harper, 1988.

[21]See Shoham, 1987 and Lehmann, 1989.

If $\phi \vdash \psi$ and $\phi \nvdash \sim\chi$, then $\phi\wedge\chi \vdash \psi$

This principle holds when the preference relation determines a total order, and fails when the preference relation determines only a partial order.

The relations between conditional logic and nonmonotonic logic are, I suspect, more complex than has sometimes been supposed, but I think they are worth exploring further.

References

Adams, E. 1970. Subjunctive and indicative conditionals. *Foundations of Language*, *6*.

Adams, E. 1975. *The Logic of Conditionals*. Dordrecht, Holland: D. Reidel.

Appiah, A. 1985. *Assertion and Conditionals*. Cambridge: Cambridge University Press.

Bennett, J. 1989. Farewell to the phlogiston theory of conditionals. *Mind*, *97*.

Chellas, B. 1975. Basic conditional logic. *Journal of Philosophical Logic*, *4*.

Dudman, V. H. 1984. Conditional interpretations of 'if' sentences. *Australian Journal of Linguistics*, *4*.

Gardenfors, P. 1988. *Knowledge in Flux*. Cambridge, MA: MIT Press.

Gibbard, A. 1981. Two recent theories of conditionals. in Harper, Stalnaker and Pearce.

Gibbard, A. and W. L. Harper. 1978. Counterfactuals and two kinds of expected utility. In C. A. Hooker, J. J. Leach and E. F. McClennen (eds.) *Foundations and applications of Decision Theory*, *1*, Dordrecht, Holland: D. Reidel. Reprinted in Harper, Stalnaker and Pearce.

Ginsberg, M. 1986. Counterfactuals. *Artificial Intelligence*, *30*.

Goodman, N. 1954. *Fact, Fiction and Forecast*. Cambridge, MA: Harvard University Press,

Grice, H. P. 1989. Logic and conversation, in Grice, *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.

Harper, W. L. 1988. Causal decision theory and game theory. In W. L. Harper and B. Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics*, Dordrecht, Holland: D. Reidel.

Jackson, F. 1987. *Conditionals*. Oxford: Basil Blackwell.

Lehmann, D. 1989. What does a conditional knowledge base entail? In R. Brackman and Hector Levesque (eds.), *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto: Morgan Kaufman.

Lewis, D. K. 1981. Causal decision theory. *Australasian Journal of Philosophy, 59*. Reprinted with added postscript in Lewis, 1986.

Lewis, D. K. *Counterfactuals*. 1973. Cambridge, MA: Harvard University Press.

Lewis, D. K. 1986. *Philosophical Papers*. New York: Oxford University Press.

Lewis, D. K. 1976. Probabilities of conditionals and conditional probabilities. *Philosophical Review, 85*. Reprinted with added postscript in Lewis, 1986.

Pollock, J. 1976. *Subjunctive Reasoning*. Dordrecht Holland: D. Reidel.

Shoham, Y. 1987. A semantical approach to nonmonotonic logics. In *Proceedings, Logics in Computer Science*, Ithaca, NY.

Skyrms, B. 1980. *Causal Necessity*. New Haven: Yale University Press.

Skyrms, B. 1984. *Pragmatism and Empiricism*. New Haven: Yale University Press.

Stalnaker, R. 1981. A defense of conditional excluded middle, in Harper, Stalnaker and Pearce.

Stalnaker, R. 1968. A theory of conditionals. *Studies in Logical Theory*, American Philosophical Quarterly Monograph Series, #2 (1968). Reprinted in Harper, Stalnaker and Pearce.

Stalnaker, R. 1975. Indicative conditionals. *Philosophia, 5*. Reprinted in Harper, Stalnaker and Pearce.

Stalnaker, R. 1984. *Inquiry*, Cambridge, MA: MIT Press.

Stalnaker, R. and R. Thomason. 1970. A semantical analysis of conditional logic. *Theoria, 36*.

Thomason, R. and A. Gupta. 1980. A theory of conditionals in the context of branching time. *Philosophical Review, 89*. Reprinted in Harper, Stalnaker and Pearce.

Thomson, J. F. 1990. A Defense of '$\supset$'. *Journal of Philosophy, 87*.

van Benthem, J. 1984. Foundations of conditional logic. *Journal of Philosophical Logic, 13*.

van Fraassen, B. 1980. A temporal framework for conditionals and chance. *Philosophical Review, 89*, reprinted in Harper, Stalnaker and Pearce.

van Fraassen, B. 1976a. Probabilities of conditionals, in W. L. Harper and C. A. Hooker(eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, Dordrecht, Holland: D. Reidel.

van Fraassen, B. 1976b. Report on conditionals. *Teorema, 6*.

Warmbrod, K. 1983. Epistemic conditionals. *Pacific Philosophical Quarterly, 64*.