Hamilton College
Russell Marcus
rmarcus1@hamilton.edu

Class 4 - September 3
Philosophy Friday #1: Conditionals
Nelson Goodman, "The Problem of Counterfactual Conditionals"

I. The Material Interpretation of the Natural-Language Conditional

The standard truth table for the material conditional is used to interpret the indicative conditional in natural language.
The material conditional is true unless the antecedent is true and the consequent is false.

α	⊃	β
⊤	⊤	⊤
⊤	⊥	⊥
⊥	⊤	⊤
⊥	⊤	⊥

Since the truth table is only false in the second line, we can think of the material conditional as saying, 'If α then β', is equivalent to 'Not (α and not-β)'
We can distinguish indicative conditionals from other kinds of conditionals:

A. Indicative conditionals: If the Mets lost, then the Cubs won.
B. Conditional questions: If I like logic, what class should I take next?
C. Conditional commands: If you want to pass this class, do the homework.
D. Conditional prescriptions: If you want a good life, you ought to act virtuously.
E. Cookie conditionals: There are cookies in the jar, if you want them.
F. Subjunctive conditionals: If Rod were offered the bribe, he would take it.

B, C, and D are not propositions; as they stand, they lack truth values.
We can, if we wish, parse them truth-functionally by turning them into indicatives, perhaps as follows:

B'. If you like logic, then you take linear algebra next.
C'. If you want to pass the class, you do the homework.
D'. If you want a good life, you act virtuously

We can regiment B'-D' as material conditionals, just as we did for A.
E is just a fraud; it's not really a conditional despite the use of 'if'.
F is really the problem with which Goodman is concerned.
The material conditional is probably the best truth-functional option for representing the conditional as it appears in English and other natural languages.
But, the natural-language conditional is more complex than the material interpretation.
Thinking about a proper treatment of conditionals quickly leads to important questions regarding the nature of scientific laws, and the ways in which they are confirmed or disconfirmed.

Indeed, discussion of the proper treatment of conditionals is a central topic in the philosophy of science.
Recent work in the logic of conditionals has also led to sophisticated modal extensions of classical logic, called conditional logics.
Conditional logics are beyond the scope of this course, but are worth your time, if you are ambitious.
We will discuss a few of the subtleties of conditionals, and the challenges facing those who wish to pursue their proper logical treatment.

II. Logical Truths and the Paradoxes of Material Implication

The material conditional creates what are called the paradoxes of material implication.
To understand the paradoxes of material implication, one has first to understand the nature and importance of logical truth.
A logical truth is a special, or privileged, sentence of a logical system.
The logical truths are the theorems of a system of logic.
In an axiomatic system, like Euclidean geometry or a formal treatment of Newtonian mechanics, we choose a small set of privileged sentences that we call axioms.
The axioms define the system.
In many cases, we insist that the axioms be obvious and uncontroversial.
The theorems of a formal system are the statements that are provable from the axioms.

We are using a system of propositional logic which has no axioms.
We could, if we wished, adopt an axiomatic system.
Still, even without axioms, some sentences of propositional logic will be theorems.
Such statements are laws of logic, or logical truths.
In propositional logic, the theorems are also called tautologies.
A tautology is a sentence that is true in every row of the truth table.
Logical truths are thus statements that can not be false.

We identify a system of logic, indeed any formal system, with its theorems.
Competing theories have different theorems.
Two theories with different axioms, or assumptions, can turn out to be equivalent, if they yield the same theorems.

> Demarcate the totality of logical truths, in whatever terms, and you have in those terms specified the logic (Quine, *Philosophy of Logic*, p 80.)

In order to know which system of logic we are using, which assumptions that system makes, we look at the theorems, at the logical truths of the system.
We will talk more about the logical truths later in the term, and show two different methods of proving that a wff is a logical truth.
To get a feel for logical truths, now, consider two.

G. $P \supset P$
H. $[(P \supset (Q \supset R)] \supset [(P \supset Q) \supset (P \supset R)]$

G and H have a natural obviousness that properly characterizes a theorem of logic, which is supposed to be the most obvious of disciplines.

Many other tautologies are also obvious.
The paradoxes of material implication are that statements of the following forms, among others, turn out to be logical truths even though they are not obvious.

I. $\alpha \supset (\beta \supset \alpha)$
J. $\sim\alpha \supset (\alpha \supset \beta)$
K. $(\alpha \supset \beta) \vee (\beta \supset \alpha)$

The paradoxes of material implication are both unobvious, and have awkward consequences.
I says, approximately, that if a statement is true, then anything implies it.
For, the truth table for the material conditional is true on every line in which the consequent is true.
So, L is true, on the material interpretation.

L. If Martians have infra-red vision, then Obama is president

J says that if a statement is false, its opposite entails any other statement.
So M is true on the material interpretation.

M. If Bush is still president, then Venusians have a colony on the dark side of Mercury.

Lastly, K says that for any statement, β, either any other statement entails it, or it entails any statement.
Every statement must be either true or false.
If a given statement is true, then, as in I, any statement entails it.
If a given statement is false, then, as in J, it entails any statement.
So N is true, according to the material interpretation of the conditional.

N. Either 'Neptunians love to wassail' entails 'Saturnians love to foxtrot' or 'Saturnians love to foxtrot' entails 'Neptunians love to wassail'.

Indeed, N is not only true, but a law of logic.
Further, either O or P is true.

O. 'It is raining' implies 'Chickens are robot spies from Pluto'.
P. 'Chickens are robot spies from Pluto' implies 'It is raining'.

For, if it is raining, then P is true; if it is not raining, then O is true.
According to a further law of logic, called excluded middle, either it is raining or it is not raining.

In sum, the paradoxes of the material conditional are two kinds of awkward results.
First, specific statements of the form of I, J, and K (like N) are laws of logic that are not obviously true.
Second, statements like L, M, and either O or P are true, given the truth values of their component propositions, even we do not intuitively see them as true.

III. Dependent and Independent Conditionals

To begin to diagnose the problems with the material conditional, let's distinguish between two types of conditional statements:

A *dependent conditional* has a connection between its antecedent and consequent.

Q-U are all dependent conditionals.

Q. If I run a red light, then I break the law.
R. If you paint my house, I will give you five thousand dollars.
S. If it is raining, then I will get wet.
T. If the car is running, then it has fuel in the tank.
U. If I were to jump out of the window right now, I would fall to the ground.

The material interpretation seems acceptable for dependent conditionals like Q-U.
Even when the antecedents are false, connections between the antecedents and consequents hold.
If the antecedent of, say, R is not true, if you do not paint my house, we can take the conditional to be true as a standing offer.
Whether or not it is actually raining, I will get wet if it is.
Whether or not I run a red light, the connection between doing so and breaking the law remains.

In contrast, consider some independent conditionals.

An *independent conditional* lacks the connection we find in a dependent conditional

V-Y are independent conditionals.

V. If 2+2=4, then cats are animals.
W. If 2+2=4, then cats are robots.
X. If pigs fly, then Utica near Rome.
Y. If pigs fly, then Utica is the capital of Canada.

The material interpretation is awkward for independent conditionals.
Since '2+2=4' is true and 'cats are animals' is true, V is true.
Since '2+2=4' is true and 'cats are robots' is false, W is false.
Since 'pigs fly' is false, X and Y are both true.
All of these results seem counter-intuitive for the natural-language conditional.
I am hesitant to pronounce at all on their truth values.

The material interpretation of the conditional thus seems acceptable for dependent conditionals.
But, it is only uncomfortably applied to independent conditionals.
The paradoxes of material implication are awkward because they hold for any values of the propositional variables, whether the relation is dependent or independent.

Still, we might accept the material analysis merely for the benefits it yields to the dependent conditional.
The material interpretation of the conditional returns a truth value for any conditional combination of propositions.

It allows us to maintain the truth functionality of our logic: the truth value of any complex sentence is completely dependent on the truth value of its component parts.
Imagine you were using logic to program a computer, or a robot.
We do not want the program to stall on an empty truth value.
We want it to have rules for how to proceed in any case.

The dependent conditional is much more common than the independent one, anyway, and I don't have strong feelings about the truth values of sentences like V-Y.
The paradoxes of material implication may thus just be seen as the price we have to pay to maintain truth-functionality.

One response you might be considering is finding a different truth table for the conditional.
It's worth a moment to see that this response is not productive.
We will do so in two stages, first looking at the first two rows of the truth table for the material conditional, and then at the second two rows.

IV. Nicod's Criterion and the First Two Rows of the Truth Table

The first two lines of the truth table for the material conditional, which are not counterfactual, look fine, especially in dependent conditionals like Q-U.
The first two lines of the truth table represent what is known as Nicod's criterion for confirmation of a scientific claim.
Many scientific laws are conditional in form.
Nicod's criterion says that evidence will confirm a law if it satisfies both the antecedent and consequent of such a law.
It also says that evidence will disconfirm a law if it satisfies the antecedent, but fails to satisfy the consequent.

Consider Coulomb's Law: $F = k\ |q_1q_2| / r^2$.
We analyze it as a claim that *if* two particles have a certain amount of charge and a certain distance between them, *then* they have a certain, calculable force between them.
We take evidence to confirm the law if it satisfies the antecedent and the consequent of that conditional.
We take evidence to disconfirm the law if it were to satisfy the antecedent and falsify the consequent.
If we were to find two particles which did not have the force between them that the formula on the right side of Coulomb's Law says should hold, and we could not find over-riding laws to explain this discrepancy, we would seek a revision of Coulomb's Law.

To take a simpler example, consider the claim that all swans are white.
We may analyze that claim as, 'if something is a swan, then it is white'.
When we find a white swan, which satisfies the antecedent and the consequent, it confirms the claim.
If we were to find a black swan, which satisfies the antecedent but falsifies the consequent, then it would disconfirm the claim.

According to Nicod's criterion, instances which do not satisfy the antecedent are irrelevant to confirmation or disconfirmation.
A white dog and a black dog and a blue pen have no effect on our confidence in the claim that all swans are white.

Call a conditional in which the antecedent is false a counterfactual conditional.
Nicod's criterion thus says nothing about counterfactual conditionals.
We are considering alternatives to the material interpretation of the conditional.
The point of mentioning Nicod's criterion was to say that we should leave the first two lines of the truth table alone.

V. The Immutability of the Last Two Rows of the Truth Table for the Material Conditional

Given the first two rows, there are three possibilities for the third and fourth lines of the truth table for the conditional that are different from the material interpretation.

Option A

α	⊃	β
⊤	⊤	⊤
⊤	⊥	⊥
⊥	⊤	⊤
⊥	⊥	⊥

Option B

α	⊃	β
⊤	⊤	⊤
⊤	⊥	⊥
⊥	⊥	⊤
⊥	⊤	⊥

Option C

α	⊃	β
⊤	⊤	⊤
⊤	⊥	⊥
⊥	⊥	⊤
⊥	⊥	⊥

Option A gives the conditional the same truth-values as the consequent.
It thus makes the antecedent irrelevant.
Option B gives the conditional the same truth-values as a biconditional.
Option C gives the conditional the same truth-values as the conjunction.
The conditional seems to have a different role in natural language from either the biconditional or the conjunction.
Thus, the truth table for the material conditional is the only one possible with those first two lines that doesn't merely replicate a truth table we already have.

To see the problem more intuitively, consider again a good counterfactual dependent conditional like S.

T. If I were to jump out of the window right now, I would fall to the ground.

Option A says that T is falsified when I don't jump out the window and I don't fall to the ground.
Options B and C say that T is falsified when I don't jump out of the window and I do fall to the ground.
But, neither case seems to falsify T, as it is intended.

The only time that T is falsified, as on Nicod's criterion, is in the second line of the truth table, when I jump out of the window and, like the Coyote for just a moment after he races off a cliff, I hang in the air.

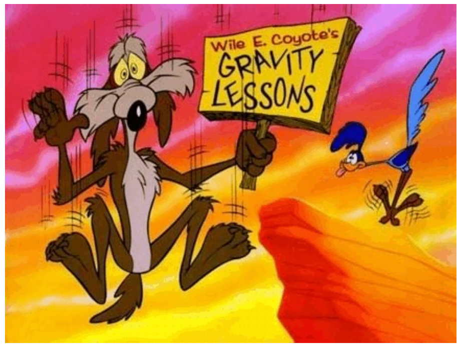

It looks like we have to stick with the original truth table, if we want the conditional to be truth-functional.

VI. Subjunctive and Counterfactual Conditionals

We have been considering whether and to what extent the material conditional stands for the natural-language conditional.
Our worries about independent indicative conditionals were allayed by two considerations.
First, we don't have strong feelings about the truth values of many of those sentences, like V-Y, which seem rare and deviant.
Second, our desire to maintain truth-functionality entails that we have to give all complex propositions truth values on the basis of the truth values of their component sentences, and the material conditional is the only option that respects our strong feelings about the first two rows of the table.
The first two rows of the truth table for the material conditional capture our intuitions about dependent conditionals, as well.
Despite worries about the paradoxes of the material implication and the oddities of material interpretations of independent conditionals, the material interpretation seems to be the best way to maintain truth-functionality.

Unfortunately, there are more problems besetting the material conditional.
Subjunctive conditionals, especially in their counterfactual interpretations, raise further problems.
Consider again F.

F. If Rod were offered the bribe, he would take it.

If its antecedent is true, we know how to evaluate F.
If Rod takes the bribe, then F is true; if he refuses the bribe, then F is false.
According to the material interpretation of the conditional, if Rod is never offered the bribe, then F is true.
So far, so good.
But, there are cases in which a conditional with a false antecedent should be taken as false.

Compare T, which is a subjunctive conditional similar to F, with T'.

T: If I were to jump out of the window right now, I would fall to the ground.
T': If I were to jump out of the window right now, I would flutter to the moon.

According to the material interpretation, T and T' are true, since I am not jumping out of the window.
But, T is true and T' is false.

The difference between T and T', and its inconsistency with the material interpretation of the conditional, is what Goodman calls the problem of counterfactual conditionals.
Goodman contrasts the following pair of propositions.

Z. If that piece of butter had been heated to 150°F, it would have melted.
Z'. If that piece of butter had been heated to 150°F, it would not have melted.

Let's imagine that we never heat that piece of butter, so that Z and Z' both contain false antecedents.
According to the material interpretation, since the antecedents of Z and Z' are both false, or counterfactual, both sentences come out as true.
But, it seems that they should be taken as contraries.
They can't both be true.
Indeed, just as we want to call T true and T' false, we want to call Z true and Z' false.

We have already seen that there are no good options for alternative truth tables.
If we want to distinguish among counterfactual conditionals, we may have to quit thinking of the natural-language conditional as a truth-functional operator.

VII. Non-Truth-Functional Operators and the Two-State Solution

We might resolve the tension between Z and Z' by claiming that 'if...then...' has two meanings.
The, let's say, logical aspects of the natural-language conditional may be expressed by the truth-functional '⊃' , encapsulated by the truth table for the material conditional.
Other aspects of the natural-language conditional might not be truth-functional at all.

We could introduce a new operator, strict implication, ⇒, to regiment conditionals whose meaning is not captured by the material interpretation.
Statements of the form '$\alpha \supset \beta$' could continue to be truth-functional, even though statements of the form '$\alpha \Rightarrow \beta$' would be non-truth-functional.
So, consider again, sentence V.

V. If 2+2=4, then cats are animals.

We could regiment it as '$T \supset C$' in the standard way, or we could regiment it as '$T \Rightarrow C$'.
'$T \Rightarrow C$' would lack a standard truth-value in the third and fourth rows.
We could leave the third and fourth of the truth table rows blank, neither true nor false.
Or, we could add a third truth value, often called undetermined, or indeterminate.
This solution would leave many conditionals, especially counterfactuals, without truth values.
We will look at three-valued logics later in the term; it's a good paper topic.

Another option, deriving from the early-twentieth-century logician C.I. Lewis and popular currently, is to interpret strict implication modally.
Lewis defined '$\alpha \Rightarrow \beta$' as '$\Box(\alpha \supset \beta)$'.
The '$\Box$' is a modal operator.
There are many interpretations of modal operators.
They can be used to construct formal theories of knowledge, moral properties, tenses, or knowledge.
For Lewis's suggestion, called strict implication, we use an alethic interpretation of the modal operator, taking the '$\Box$' as 'necessarily'.
So, on the modal interpretation of conditionals, a statement of the form '$\alpha \Rightarrow \beta$' will be true if it is necessarily the case that the consequent is true whenever the antecedent is.
Modal logics, while currently popular, are controversial.
Some philosophers believe that matters of necessity and contingency are not properly logical topics.
Other philosophers worry that our ability to know which events or properties are necessary and which are contingent is severely limited.
I will not pursue strict implication or modal logic here, but they would each make good paper topics.

One advantage of introducing a modal operator to express implication is that it connects conditional statements and scientific laws.
A scientific law is naturally taken as describing a necessary, causal relation.
When we say that event A causes event B, we imply that A necessitates B, that B could not fail to occur, given A.
To say that lighting the stove causes the water to boil is to say that, given the stability of background conditions, the water has no choice but to boil.
Thus, we might distinguish the two senses of the conditional by saying that material implication represents logical connections, where strict implication attempts to regiment causal connections.

VIII. Counterfactual Conditionals and Causal Laws

As we have seen, the natural-language conditional often indicates a connection between its antecedent and its consequent.
Consider again the dependent conditionals Q-U.

> Q. If I run a red light, then I break the law.
> R. If you paint my house, I will give you five thousand dollars.
> S. If it is raining, then I will get wet.
> T. If the car is running, then it has fuel in the tank.
> U. If I were to jump out of the window right now, I would fall to the ground.

The connection in Q is mainly conventional.
R refers to an offer or promise.
But, S, T, and U are fundamentally causal, depending on most basic scientific laws.
Our investigation of the logic of the conditional has taken us into questions about causation.

Consider:

> AA: If this salt had been placed in water, it would have dissolved.

AA indicates a dispositional property of salt, one that we use to characterize the substance.
Other dispositional properties, like irritability, flammability, and flexibility, refer to properties interesting to scientists.
Psychological properties, like believing that it is cold outside, are often explained as dispositions to behave, like the disposition to put on a coat or say, "It's cold."
Contrast AA with:

> BB: This marble counter is soluble in water.

If we never place the counter in water, then BB comes out true on the material interpretation.
To be flammable is just, by definition, to have certain counterfactual properties.
These pajamas are flammable just in case they would burn if subjected to certain conditions.
The laws of science depend essentially on precisely the counterfactual conditionals that the logic of the material conditional gets wrong.

Goodman argues that the problem with counterfactual conditionals is that they are not merely logical relations.
The problem of giving an analysis of the logic of conditionals is intimately related to the problem of distinguishing laws from accidental generalizations.
Compare:

> CC: There are no balls of uranium one mile in diameter.
> DD: There are no balls of gold one mile in diameter.

The explanation of CC refers to scientific laws about critical mass.
If you gather too much uranium in one spot, it explodes.
The explanation of DD, in contrast, is merely accidental.
It is entirely possible that we could gather that much gold together, while it is impossible to gather the same amount of uranium.
In order to know that difference, though, you must know the laws which govern the universe.
Goodman's claim is that the problem of distinguishing CC from DD, the problem of knowing the laws of nature, is inextricably linked to the problem of understanding the logic of the natural-language conditional.
We may use conditionals as truth-functional connectives, sometimes.
More commonly, especially in counterfactual cases, we use them to state connections between antecedents and consequents.
So, a conditional will be true if the relevant connections hold among the antecedent and the consequent.
It is false if such connections do not.

> A counterfactual is true if a certain connection obtains between the antecedent and the consequent. But...the consequent seldom follows from the antecedent by logic alone (7-8).

Consider AA, or Goodman's sentence:

> EE. If that match had been scratched, it would have lighted.

For either AA or EE to be true, there have to be a host of other conditions satisfied.

> We mean that conditions are such - i.e. the match is well made, is dry enough, oxygen enough is present, etc. - that "That match lights" can be inferred from "That match is scratched." Thus the connection we affirm may be regarded as joining the consequent with the conjunction of the antecedent *and other statements that truly describe relevant conditions* (8-9, emphasis added).

When we assert a conditional, we commit ourselves to the relevant related claims.
But, to understand the claims to which we are committed, we must understand the relevant connections between the antecedent and the consequent.
We must understand the general laws connecting them.

> The principle that permits inference of 'That match lights' from 'That match is scratched. That match is dry enough. Enough oxygen is present. Etc.' is not a law of logic but what we call a natural or physical or causal law (8-9).

In order to infer the consequent of AA, for example, from its antecedent, we need to presume causal laws governing the dissolution of salt in water.
In order to infer the consequent of ε from its antecedent, we need to presume causal laws about the lighting of matches.

Goodman thus argues that a proper analysis of counterfactual conditionals would include two elements.

1. A definition of the conditions that are relevant to the inference of a consequent from an antecedent.
2. A characterization of causal laws.

We have gone far from just understanding the logic of our language.
We are now engaged in a pursuit of the most fundamental features of scientific discourse.
Distinguishing between T and T', or between Z and Z', in contrast to the material interpretation, would import extra-logical features into our logic.
While we believe that T' and Z' are false, our reasons for that belief do not seem, now, to be a matter of logic.
The reasons we think that they are false are due to the laws of physics.
If we were living on a planet with very little gravitational force, but on which buildings had limited force fields that kept us tethered to the ground inside, it might indeed be the case that if I jumped out of the window, I would fly to the moon, rather than fall to the ground.
We really want our logic to be independent of all the extra-logical facts.
We don't want to import the physical facts into our logic, since we want our logic to be completely independent of the facts about the world.

Thus, we rest with the material interpretation of the natural-language conditional, giving up hope for a truth-functional analysis of the causal conditional, and remembering that the natural-language conditional represents a strictly logical relation.

Paper Topics

1. Contrast the following pair of counterfactual conditionals.

> If bin Laden didn't plan the 9-11 attacks, then someone else did.
> If bin Laden hadn't planned the 9-11 attacks, then someone else would have.

The antecedents and consequents of these statements are nearly identical, but, our estimations of the truth values and semantics of U and U' are different. Discuss the similarities and differences among these sentences. Can we use the material conditional for this example? Are there other options? See Bennett and Jackson for discussions of a relevantly similar pair of sentences.

2. Consider the following inference.

> If this is gold, then it is not water-soluble.
> So, it is not the case that if this is gold then it is water-soluble.

Intuitively, this argument seems valid. But, if we regiment the argument in a standard way, we get an invalid argument. Discuss this problem in the light of the discussion of the material conditional. For possible solutions, you might look at Lewis and Langford 1932; Priest 2008; Goodman's work, or

3. Another good topic would be to explore what are known as relevance logics.
In relevance logic, we insist that for a conditional to be true, its antecedent and consequent must be appropriately related.
As I briefly mentioned in class, last time we met, there are folks working on relevance logics, mostly following C.I. Lewis's suggestion concerning strict implication. See Priest 2008.

4. Lewis on strict implication

5. The philosopher Paul Grice, responding in part to the problems of the conditional, distinguished between the logical content of language, and other, pragmatic, features of language. In addition to Grice's paper, Fisher, Priest, and Bennett all have useful discussions of Grice's suggestion.

6. Connections to three-valued logics

7. Lewis Carroll's paper, "A Logical Paradox"

8. Goodman, and the relation between conditionals and scientific laws. Hempel.

9. Frank Jackson and David Lewis have extended treatments of conditionals

Suggested Readings

Bennett, Jonathan. *A Philosophical Guide to Conditionals*. Oxford University Press, 2003.
Carroll, Lewis. "A Logical Paradox." *Mind* 3:436-8, 1894.
Edgington, Dorothy. "Do Conditionals Have Truth Conditions?" In Hughes, R.I.G. *A Philosophical Companion to First-Order Logic*. Hackett, 1992. A selection of very good, advanced papers.
Fisher, Jennifer. *On the Philosophy of Logics*. Chapter 8.
Goodman, Nelson. "The Problem of Counterfactual Conditionals." In *Fact Fiction and Forecast*, Harvard University Press, 1983.
Grice, H.P. "Logic and Conversation." Reprinted in Jackson 1991.
Harper, William, Robert Stalnaker, and Glenn Pearce. *Ifs*. Reidel, 1981. Lots of advanced papers, some quite technical.
Jackson, Frank. *Conditionals*. Oxford University Press, 1991.
Lewis, David. *Counterfactuals*. Blackwell, 1973.
Lewis, C.I., and Langford, C.H. *Symbolic Logic*. New York: Century Company, 1932. Reprinted by Dover Publications (New York), 1959.
Priest, Chapters 1, 4, and 5.
Read, Chapter 3.
Weiner, Joan. "Counterfactual Conundrum." *Nous* 13.4: 499-509, 1979.