

HOW GOOD ARE YOUR LOGICAL INTUITIONS?

Raymond D. Bradley

[Extract from an unfinished elementary textbook, *Reasoning: the 4th "R"*]

Some children seem blessed, almost from birth, with a capacity for critical thinking. They won't let a fallacious argument pass unnoticed or unscathed. And some are fortunate enough to be exposed at an early age to fine examples of good reasoning. In their listening and their reading they learn, by intellectual osmosis as it were, to think logically. Yet even these fortunate ones, like the rest of us, can benefit by having their logical intuitions and reasoning skills sharpened by precept and practice.

So, no matter how logical you already are, you may benefit from working your way through this chapter and picking up on some of the logical lessons it offers.

1. HOW GOOD ARE YOUR LOGICAL INTUITIONS?

Let me invite you to test your logical intuitions on a few examples. In each case, I'd like you to think about it a bit and then answer the question that follows as honestly as you can. Don't give an answer merely because you think it is the one expected of you. Give an answer only if you genuinely believe it to be the right one. Don't "hedge" or "sit on the fence". Think. Then, before going on to the next example, commit yourself by writing your answer on a piece of paper so that you don't later lose track of what you originally thought the answer to be. After all, you might want to re-examine it.

None of the questions I'll ask involves any purely "verbal" tricks. I'll use words in their ordinary, hum-drum, senses. Hence you won't need to be on the look-out for syntactic ambiguities, such as you find in the sentence "Flying planes can be dangerous". Nor will you have to watch for semantic ambiguities, such as you find in the question "Which is healthier: Mr Chretien or the Canadian

economy?" What you will need to do, though, is get your powers of imagination and reasoning working together.

The square house.

Is it possible for a square house to have all four walls facing due south?

Answer "Yes" or "No".

Brothers.

A is a brother of B, and B is a brother of C.

Does it follow from the above information that A and C are brothers?

Answer "Yes" or "No".

If you find it difficult to think in the abstract about people called "A", "B", and "C", try giving them corresponding names - "Al", "Bobby", and "Chris", for instance. Then ask, and try to answer, the question again.

Rain and wet pavements.

For each of the following cases, ask yourself whether, from the given premises, the suggested conclusion follows:

Example (1): If it is raining, then the pavement is wet.

It is not raining.

Can you correctly infer that the pavement isn't wet? Answer "Yes" or "No".

Example (2): If it is raining, then the pavement is wet.

The pavement is wet.

Can you correctly infer that it's raining? Answer "Yes" or "No".

Example (3): If it is raining, then the pavement is wet.

It is raining.

Can you correctly infer that the pavement is wet? Answer "Yes" or "No".

Example (4): If it is raining, then the pavement is wet.

The pavement is not wet.

Can you correctly infer that it's not raining? Answer "Yes" or "No".

Some apples in a bag.

You have been handed a bag of apples and told, truthfully, that some of them are rotten.

Does it follow that some of the apples are not rotten? [To put the question more abstractly: Does the statement that some Xs are Ys imply the statement that some Xs are not Ys?]

Answer "Yes" or "No".

The barber.

Consider the following story:

Once upon a time, back in 1961 it was, I went on a ski holiday to the little Austrian village of Lech. Unfortunately, I'd been so preoccupied with my duties at Oxford University that I'd let my hair grow to a length which, though acceptable in those academic precincts, led the villagers to look askance at me. So, after I'd been living in the village for just over a week, I decided I'd better get a hair cut. But where? My ski instructor told me that the barber lived in the center of the village and that he'd certainly be happy to cut my hair provided I hadn't already tried to cut it myself. It seems, you see, that if someone living in the village cut his own hair, the barber refused to cut it. But, by the same token, he always cut the hair of anyone living in the village who didn't cut his own. Since I certainly hadn't cut my own hair, or even made an attempt at it, I went to this barber and had my hair cut to a socially acceptable length.

Most of the above story is true autobiography. Could *all* of it be true? Answer "Yes" if you can't see anything wrong with it. Remember to be honest with yourself. Answer "No" only if you think you it involves a contradiction, and are

prepared to say exactly how you think that contradiction arises.

Now let's examine each of these examples in turn and see what we can learn from them about how we ought to reason.

2. THE CONCEPT OF LOGICAL POSSIBILITY

The square house

Before supplying the answer to our first question, I want you to think about it again. Note that our question does not ask whether any *actual* house has all four walls facing south. Our question is not one of architecture or of geography or of any other subject that requires you to indulge in fact-gathering about how the world happens to be.¹ Rather, it is a question about *how the world might be or have been*. Roughly, we're asking: Can you CONCEIVE of any situation in which a square house might be located so as to have all its four walls facing south? So let your IMAGINATION run wild, as it were, allowing yourself - in the way that science-fiction writers do - to dream up POSSIBLE WORLDS in which such a house might exist.

What's the answer, then?

Most people to whom I've put the question think it's "No". They think that if a house is square and one of its walls faces due south, then it follows logically that the others must face due north, due west, and due east. But their reasoning is flawed. They ignore a relevant possibility: that a square house could be situated right on the North Pole! The fact that there isn't a house so located (at least so far as I know), is irrelevant since we are dealing with what is possible not what is actual. So the answer is: "Yes."

¹ Philosophers like to say that a question that can be answered without indulging in fact finding about how the world happens to be is a question that can be answered A PRIORI. The term "a priori", then, means roughly "without appeal to experience". Questions that can be answered only by appeal to experience are said to be EMPIRICAL ones.

Reflections on the concept of possibility

Let's pause briefly and reflect about this example. In saying that it is perfectly POSSIBLE for there to be such a house, we are using the term "possible" in an extremely wide sense. We are saying that such an envisaged situation is LOGICALLY POSSIBLE, that there is no contradiction or inconsistency involved.

Needless to say, not all described situations are logically possible. It isn't logically possible, for instance, that there should be a square house with only three sides. The reason is obvious. In calling it square, we are implying that it has four sides; and something with four sides can't at the same time have only three. The very supposition that a square house might simultaneously have only three sides therefore contains a contradiction within itself. It is self-contradictory. And self-contradictory statements, remember, cannot be true. No POSSIBLE WORLD contains a square house with only three sides.

There is, then, a boundary, between what is logically possible and what is logically impossible.

Not all possibility and impossibility, however, are logical in nature. Indeed, the concepts of possibility and impossibility are generic in character, admitting - as it were - different species. That which is logically possible (or impossible) needs to be distinguished from that which is physically possible (or impossible), and that in turn from what is technologically possible (or impossible). Nor do these three species exhaust the whole range of senses in which the terms "possible" and "impossible" are used. We can further distinguish, for instance, between what is legally possible (or impossible) and what is morally possible (or impossible), i.e., what is morally permissible (or morally forbidden). And still other senses of these expressions can be distinguished.

All talk of what is possible or impossible is said to be MODAL talk. [Roughly speaking, it is so-called because it has to do with modes of existence.] So, too, is talk of what is necessary (or nonnecessary), contingent (or noncontingent). And the point we have just been making is that modal talk in general needs to be disambiguated so that we are clear, in any given context, about the precise sense in which an occurrence of a modal word like "possible",

"impossible", "necessary", "non-necessary", "contingent", or "non-contingent" should be understood.

Let me illustrate by asking you to consider this question:

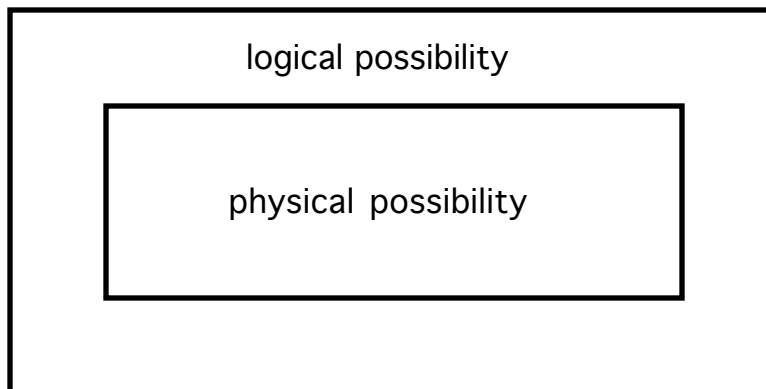
Is it possible that the earth should, without the intervention of any other physical body, instantaneously stop spinning on its axis?

How should you construe the question? Mindful of the distinction between logical and physical possibility, a thoughtful person will say something like the following:

It all depends on what you mean by "possible". If you are asking whether it is logically possible, then the answer is "Yes". There certainly is no contradiction involved in the supposition that the earth should behave in this erratic, indeed catastrophic, way. One could, without any inconsistency, imagine an omnipotent (all-powerful) god performing just such a miracle, i.e., just such a violation of a law of nature. But if you are asking whether it is physically possible, then the answer is "No". For such behavior would be entirely contrary to the laws of physics; and anything contrary to the laws of physics is physically impossible. In short, the answer is that the envisaged occurrence is logically possible but physically impossible.

Here, in this little reply, we see an illustration of the kind of intellectual acuity that does not allow itself to be bamboozled by ambiguous questions. Such an ability to perceive, and make use of, fine semantic/conceptual distinctions is frequently the mark of someone well-trained in philosophy. Not surprisingly, then, philosophers have spent a good deal of effort on drawing such distinctions and seeing how the concepts distinguished relate to one another.

How are the various modal concepts related to one another? So far as the distinction between LOGICAL POSSIBILITY and PHYSICAL POSSIBILITY is concerned, there is broad agreement that they relate to one another in the manner depicted on the following CONCEPTUAL MAP:



Our conceptual map should be read thus: The larger rectangle depicts the realm of what is logically possible. Anything outside that rectangle, therefore, is to be thought of as logically impossible. The smaller rectangle depicts the realm of what is physically possible. And anything outside that rectangle, once again, is to be thought of as physically impossible. Our conceptual map clearly shows, then, that anything which is physically possible is also logically possible. It also shows that the converse doesn't hold. Something can be logically possible yet physically impossible.

For the most part, our concern in this text will be with logical uses of the modal expressions "possible", "impossible", "necessary", "non-necessary", "contingent", and "non-contingent", and with other modal expressions (such as "implication", "deductive validity") that are definable in terms of them.²

Exercises:

1. How is the concept of technological possibility related to those of physical and logical possibility. How would you depict it on the above conceptual map?
2. Try to draw a conceptual map (using rectangles, as above) depicting the relationship between the concepts of legal possibility (what is legally permissible) and moral possibility (what is morally permissible). Will one rectangle be located within the other? If so, which is inside which? If neither sort of nesting seems right, will the rectangles overlap? Try to explain your reasoning as simply and clearly as you can.

² Recall the section "Concepts related to that of Deductive Validity" in Chapter 1, pages 19-23.

3. INFERENCES INVOLVING RELATIONAL CONCEPTS

Brothers

In this example we're asked to consider a hypothesis which, for all I know, might well in fact be true: that there are three persons, A, B, and C - Al, Bobby, and Chris, if you prefer - such that A is a brother of B and B a brother of C. It doesn't matter whether or not the hypothesis IS true, i.e., whether or not the actual world contains three such persons so related. The question is only whether, IF some possible world were to contain three such persons, what would follow, i.e, what else would HAVE to be true. Would A and C have to be brothers? Or, if you like, would Al and Chris have to be brothers?

Another way of asking the question is: Is it LOGICALLY POSSIBLE - in the sense discussed above - for the premises:

Al is Bobby's brother

Bobby is Chris's brother

to be true, and yet the conclusion:

Al and Chris are brothers

to be false?

Let's take it in stages. One conclusion we can validly infer is:

Al is a brother of Chris.

For if Al is a brother of Bobby, then Al must be one of Bobby's siblings. And if Bobby is a brother of Chris, then Bobby must be one of Chris's siblings.

Therefore from the premises we can infer that Bobby is a sibling both of Al and of Chris. Now since any two persons who have a sibling in common must be siblings of one another, we can infer that Al is a sibling of Chris. Moreover, since we know that Al is a brother of Bobby, we know that Al is a male. Therefore, Al is not just a sibling of Chris but a male sibling, to boot. And that is just what it means to say that Al is a brother of Chris. So Al's being a brother of Chris certainly does follow logically from the information we've been given.

So far so good. We've correctly deduced that Al is a brother of Chris. But we still haven't answered the question: Does it follow that Al and Chris are

brothers? From the given premises, can we infer the truth of:

Al and Chris are brothers.?

Well, think about it. In order for Al and Chris to be brothers, it would have to be the case not only that the statement:

Al is a brother of Chris

is true - something we've already established. It would also have to be the case that the statement:

Chris is a brother of Al

is true. But what grounds have we for inferring this further statement?

None. The claim that Chris is a brother of Al doesn't follow from any of the information given us in the premises. If we think it does, this can only be because we have jumped to the conclusion that Chris is male. But the information in the premises gives us no warrant for that. Anyone who concluded that Chris is a brother of Al, has ignored a POSSIBILITY that would make the premises true and the conclusion false, viz., the possibility that while Al is a brother of Chris, Chris is not Al's brother but his sister.³ It is entirely consistent with everything stated in the premises, that Chris is Al's sister rather than brother. Hence it is possible for the statement that Chris and Al are brothers to be false, even in circumstances in which both the premises happen to be true. But a valid argument, remember, is one in which it is not possible for the conclusion to be false when the premises are true. A valid argument, a valid inference, is one which *guarantees* you against deriving falsehood from truth. So anyone who concluded that Chris and Al are brothers was arguing invalidly.

Transitivity and symmetry

We can generalize from this example by pointing out that the relation (relationship) *being a brother of*, though transitive, is not symmetrical. Let's use the variable "R" to stand for any relation (relationship) whatever. Now a relation R is classified as transitive iff when it holds between one thing, A, and a second, B, and also between B and a third thing, C, then it must hold between A and C.

³ There's no "verbal trick" here. Anyone who jumped to the conclusion that "Chris" was a boy's name also ignored a possibility, viz., the possibility - indeed the actuality - of its being a girl's name as well. The names "Al" and "Bobby", too, it should be noted, may be used to refer to females as well as males. The conclusion that Al is a male can be inferred validly from the premise that Al is Bobby's brother; but it can't be inferred validly from the use of the name "Al" alone.

The relation *being a brother of*, we have seen, is transitive. However, this relation is not symmetrical. For a relation is classified as symmetrical iff it is a "two-way relation", i.e., iff when it holds between a first thing, A, and a second, B, it must also hold between B and A. Yet Al can be a brother of Chris without Chris being a brother of Al. Transitivity and symmetry sometimes go together, as in the cases of *being a sibling of* and *being the same height as*. But they don't always go together, as the cases of *being a brother of* and *being brothers of each other* both demonstrate.

Exercises

1. Think of two more relations (other than the ones just discussed) one of which is both transitive and symmetrical, the other of which is transitive but not symmetrical.
2. Is the relation of implication symmetrical? Is it transitive? [Try to justify your answers.]

An important logical lesson:

There is an important logical lesson to be learned from this example. It can be expressed by the following PRINCIPLE OF CONTAINMENT:

The conclusion of a valid argument can't go beyond or contain more than what is contained in its premises.⁴

In a fairly obvious sense, it is just because it violates this principle that the argument from Al's being a brother of Bobby, and Bobby 's being a brother of Chris, to Chris's being a brother of Al, goes wrong. The conclusion here does go beyond what is contained in the premises in so far as it assumes something that the premises give us no warrant for assuming, viz., that Chris is male.

4. INFERENCES INVOLVING CONDITIONALS

⁴ We will give a more precise account of this talk of "containment" later.

Rain and wet pavements

In each of the four cases you were given certain premises and asked whether a certain conclusion could validly be inferred from them.

Now the question of validity, let me remind you, amounts to this:

Are there any POSSIBLE circumstances in which the premises would be true and the conclusion would be false? If there are, then the inference is invalid. If there are not, then the inference is valid.

Think about the examples once again. Maybe by now, you'll want to rethink the answer you gave first time around. So, DO IT! NOW!

Now let's see what the correct answers are, and - more importantly - *why* they are correct.

Example (1): You are given the premises:

If it's raining, then the pavement is wet.

It is not raining.

And the question is whether you can validly draw the conclusion:

The pavement isn't wet.

The answer is "No".

The reason is, of course, that someone who infers that the pavement isn't wet is overlooking all sorts of logical possibilities that would leave the premises true but make the conclusion false: e.g., the possibility that the pavement is wet from a recent hosing, that the pavement hasn't had time to dry off since the last shower, and so on. In each of these possible situations the conclusion would have been false - i.e., the pavement would have been wet - even if the premises were true. The truth of the premises, then, hasn't given us a guarantee of the truth of the conclusion. So the argument is invalid.

Another logical lesson:

The foregoing explanation invokes a quite general principle which can be formulated as follows:

An argument is invalid iff there is at least one statement which is consistent with its premises (i.e., would leave them true) but inconsistent with its conclusion (i.e., would make the conclusion false).

Thus the statement:

The pavement is wet from a recent hosing
is consistent with both the premises

If it's raining, then the pavement is wet

and

It is not raining

but is inconsistent with the conclusion

The pavement isn't wet.

That is precisely *why* the premises of this invalid argument could be true and the conclusion false.

Example (2): You are given the premises:

If it's raining, then the pavement is wet.

The pavement is wet.

And the question is whether you can validly draw the conclusion:

It is raining

Once more the answer is "No".

By now it should be easy to see why one can't correctly conclude that it is raining. Anyone who draws the conclusion that it's raining is again ignoring all sorts of possibilities that would leave the premises true but make the conclusion false, viz., the very same set of logical possibilities, in fact, as were involved in example (1). The wetness of the pavement might have resulted a previous shower, from someone spraying it with a hose, or the like. Once more, then, the truth of the premises is perfectly consistent with the falsity of the conclusion. Hence the conclusion doesn't follow from the premises.

Note on "if" and "only if"

This brings me to a minor, but nevertheless significant, semantic point: the need to distinguish between "if" and "only if". I suspect that many who gave incorrect answers to (1) and (2) did so because they read the first premise

If it's raining, then the pavement is wet
 as if it were equivalent in meaning to
 Only if it's raining is the pavement wet.

But the two certainly don't mean the same. Nor are we indulging in mere verbal pickiness in pointing out the difference. The claim

Only if it's raining is pavement wet
 allows for only one possible explanation of the pavement's wetness. It says that it must have rained if the pavement is wet. That is to say, it states that having been rained upon is a NECESSARY condition of the pavement's being wet.

But the claim

If it's raining, then the pavement is wet
 doesn't say this at all. It allows for other possible explanations as well. It merely states that one SUFFICIENT condition of the pavement's being wet is its having been rained upon.⁵ It is quite compatible with there being other sufficient conditions as well, e.g., hosing it down. So it doesn't say that having been rained upon is necessary in order for the pavement to be wet. Herein lies the difference.

All too often, fallacious reasoning has its roots in sloppy reading or careless listening, in a failure to think about and digest precisely what is being said and precisely what it means.

Example (3): You are given the premises:

If it is raining, then the pavement is wet.
 It is raining.

And the question is whether you can validly draw the conclusion:

The pavement is wet.

The answer, as you've probably already worked out, is "Yes".

But even if the answer was obvious to you all along, it will prove instructive to through the reasoning in fairly explicit detail.

⁵ If you are unclear about this, you might want to review the discussion of necessary and sufficient conditions given in Chapter 1, page 19.

Our first premise

If it's raining, then the pavement is wet

is a conditional statement which states that if its ANTECEDENT (the if-clause, "it is raining") is true, then its CONSEQUENT (the then-clause, "the pavement is wet") will also be true. Note that it shouldn't be construed as saying that its raining is logically sufficient for the pavement's being wet; it shouldn't be construed, that is, as saying that the antecedent logically implies the consequent. For if it were so construed, the first premise would be false⁶; and you were, after all, told to assume it to be true. All that our first premise says is simply that the rain's falling is CAUSALLY SUFFICIENT to bring about the wetness of the pavement.

Now suppose that this first premise is true, and further, that the second premise

It is raining

is also true. This means that the second premise simply AFFIRMS THE ANTECEDENT of the first premise. It asserts that these causally sufficient conditions of the pavement's being wet are fulfilled (satisfied) - a matter of physics. But this means that IF both premises were true, then the consequent would also HAVE to be true - a matter of logic. It means that there are no logically possible circumstances in which the premises would be true and yet the conclusion false. In short, it means that the truth of the premises is LOGICALLY SUFFICIENT for the truth of the conclusion. The premises, whether or not they ARE true, imply the conclusion. The argument is valid. Hence the conclusion would be true if the premises were.

But, some might object, haven't you here ignored a possibility, viz., the possibility that there are sheets of plastic or some other such shield over the pavement such that in its presence the pavement wouldn't get wet, and hence the conclusion wouldn't be true?

Think for a moment about how you would reply to this objection. Try to see what is wrong with it.

And now that you've tried, see whether your reply is along the same lines

⁶ There is certainly no contradiction involved in asserting the antecedent of this conditional and denying its consequent.

as mine.

My reply would be that we have indeed ignored the possibility that the pavement is shielded from the falling rain, whether by plastic, an awning, or an act of God. But we've ignored such possibilities for good reason. For these possibilities are irrelevant to the question whether IF the premises were true the conclusion would also have to be true. Just think about it. Suppose the possibility that the pavement was shielded somehow or other had been actualized. Would this possibility show that the premises of our argument could be true and the conclusion false? Obviously not. Had this possibility been actualized, our premises would have been false, not true. So the fact that our conclusion would have been false in circumstances in which the premises are also *false* is totally irrelevant to the question whether our argument is valid. It is irrelevant, that is, to the question whether our conclusion would have been false in circumstances in which the premises are *true*.

Example (4): You are given the premises:

If it's raining, then the pavement is wet.

The pavement isn't wet.

And the question is whether you can validly draw the conclusion:

It isn't raining.

Once again, as you've probably concluded, the answer is "Yes".

Again, let's see why. As we've seen, the first premise

If it's raining, then the pavement is wet

tells us that a sufficient condition of the pavement's being wet is that it should have rained. But if the falling of the rain is a sufficient condition of the pavement's being wet, then the pavement's being wet is a necessary condition of its having been raining, i.e., the pavement's being wet is a causally necessary condition (more usually called a causal consequence) of its having been raining. And that means that were the pavement not to have been wet, it couldn't have been raining. But now think. Our second premise is the statement

The pavement isn't wet.

Our second premise, then, DENIES THE CONSEQUENT of the first premise. It tells us that this necessary consequence of it having rained is not satisfied. But

this means that IF both premises were true, then the antecedent of the first premise would have to be false. It means, in other words, that the conclusion

It isn't raining

(which denies the antecedent) would have to be true. The truth of the premises, once more, is logically sufficient for the truth of the conclusion. There are no possible circumstances in which both premises are true and yet the conclusion is false. Hence the premises imply the conclusion and the argument is valid.

FIRST INTERLUDE

VALIDITY, INVALIDITY, AND THE FORMS OF ARGUMENTS

Consider the argument featured as Example (3) again:

If it is raining, then the pavement is wet.

It is raining.

Therefore, the pavement is wet.

Now compare it with this argument:

If God is in heaven above, then all's well with the world.

God is in heaven above.

Therefore, all's well with the world.

There is a very obvious sense in which these two arguments have the same FORM. We can obtain the second from the first by substituting "God is in heaven above" for all occurrences of "It is raining", and substituting "All is well with the world" for all occurrences of "The pavement is wet". For that matter, we could substitute any statement we pleased for "It is raining" and "It is wet" as they occur in Example (3), and - provided we did so uniformly - we would obtain another argument having the same form as our original. All arguments having this particular form are said to be instances of **AFFIRMING THE ANTECEDENT** since, in each of them, one of the premises affirms the antecedent of the other.

Ancient logicians, way back in the time of Aristotle (384-322 BC) and Chrysippus (c.280-207 BC), recognized that countless arguments can all share the same form. But how, in general terms, is one going to describe the **COMMON FORM** of the above antecedent-affirming arguments? Here is the description that Chrysippus and his stoic followers came up with: "If the first, then the second; but the first; therefore the second." Obviously, by "the first" they

meant to refer to whichever statement features as the antecedent of the first premise of any of these arguments, and then again as the second premise. And by "the second" they meant to refer to whichever statement features as the consequent of the first premise, and then again as the conclusion. But if we try to confine ourselves to ordinary language in giving general descriptions of the forms of arguments, we'll find that some of our descriptions will be very complicated and potentially confusing. Try to give such a description for an argument containing half a dozen or more different statements!

These days, we can do a bit better. Availing ourselves of the idea of a STATEMENT-VARIABLE, a symbol which is used as a "stand in" for any given statement or statements, we can use a letter such as "P" instead of the expression "the first", and a letter such as "Q" instead of the expression "the second", and so on. We can then say that both of the above arguments, and countless others, have the form:

If P then Q; P; therefore Q

Next, consider the argument of Example (4):

If it is raining, then the pavement is wet.

The pavement isn't wet.

Therefore, it isn't raining.

It, too, has a form which it shares in common with countless others, arguments such as:

If God is in heaven above, then all's well with the world.

All is not well with the world.

Therefore, it is not the case that God is in heaven above.

That common form can be expressed as:

If P then Q; not Q; therefore not P

All arguments of this form are said to be instances of DENYING THE CONSEQUENT, since, in each of them, one of the premises denies the consequent of the other .

Now what is really significant about all this is that all instances (examples) of these two argument forms are VALID. The reasoning that we went through in order to demonstrate the validity of examples (3) and (4) could be repeated, mutatis mutandis (with the necessary changes having been made), for all the

countlessly many other arguments that involve Asserting the Antecedent or Denying the Consequent. Both, we may say, are VALID ARGUMENT FORMS, and any arguments that are instances of them - no matter what their subject-matter might be - will also be valid.

Here, in these two valid argument-forms, are illustrations of the power that formal logic has as a result of its abstractness. Just as the validity of the simple mathematical truth that $2 + 2 = 4$ doesn't depend on whether we are talking about two apples plus two apples, or two oranges plus two oranges, but has application to any pair of object-pairs whatever, so the validity of the argument-forms, "If P then Q; P; therefore Q" and "If P then Q; not Q; therefore not P", doesn't depend on whether we are talking about falling rain and wet pavements, or God in heaven and the well-being of the world, but has application to any statement-pairs whatever. Test this out for yourself by substituting any statements whatever for the variables "P" and "Q" in these two valid argument-forms. All the resulting arguments will be such that it is impossible for their premises to be true without their conclusions being true also.

Can we generalize analogously for the cases of (1) and (2)? Can we conclude that any arguments having the forms of (1) and (2) will be invalid?

Example (1) has the form:

If P then Q; not P; therefore not Q

It is said to be an instance of the FALLACY OF DENYING THE ANTECEDENT.

Likewise (2) has the form

If P then Q; Q; therefore P

It is said to be an instance of the FALLACY OF AFFIRMING THE CONSEQUENT.

Both these forms of argument are INVALID ARGUMENT-FORMS.

But are all their INSTANCES also invalid? Are all instances of Denying the Antecedent invalid? Are all instances of Affirming the Consequent invalid?

Strictly speaking, the answer is: No. The reason is that there may be OTHER RELEVANT LOGICAL CONSIDERATIONS, not reflected in these argument-forms, by virtue of which an argument that is an instance of one of these invalid argument-forms is valid after all. For example, if P were the statement "John is aged three-score years and ten", and Q were the statement "John is seventy", then both the resulting arguments

If John is three-score years and ten, then John is seventy years old

John is seventy years old

Therefore, John is three-score years and ten

and

If John is three-score years and ten, then John is seventy years old

John is not aged three-score years and ten

Therefore, John is not seventy years old

would be perfectly valid, despite being instances of the fallacies of Denying the Antecedent and Affirming the Consequent, respectively. In such cases, you see, there *are* other relevant logical considerations that make the resulting arguments valid. After all "John is aged three-score years and ten" and "John is seventy" are logically equivalent to one another; and this logical fact makes an enormous difference.

An important logical lesson.

There's a simple but important lesson to be derived from the last few paragraphs, one that is seldom mentioned in logic texts, and is often forgotten even by those who have been alerted to it. Putting it in its most general terms, it amounts to this:

Having a valid argument-form is a sufficient, but not a necessary, condition of the validity of an argument that is an instance of that form.

More particularly, although the fact that an argument has a valid argument-form provides us with a guarantee that it is valid, the fact that an argument has an invalid argument-form does not provide us with a guarantee that it is invalid. We will be justified in concluding that an argument having an invalid argument-form is

invalid only if we are assured that no other relevant logical considerations show the contrary.

Those who assert that possession of an invalid form makes an argument invalid can be hoist on their own logical petard, as it were. For I suspect that their reasoning is of the following form:

If argument X has a valid argument-form, then X is valid.
X does not have a valid argument-form.
Therefore X is not valid.

But this argument itself has an invalid form: it involves the fallacy of Denying the Antecedent. Hence, this argument, for saying that having an invalid form suffices to make an argument invalid, is itself invalid.

In any case it is easy to cite examples of valid arguments which are not instances of valid argument-forms: the inference from A's being a brother of B, and B's being a brother of C, to A's being a brother of C, is a case in point. Hence, to repeat the point:

Having an invalid argument-form does not suffice to make an argument invalid, though having an invalid argument-form in the absence of what I have called "other relevant logical considerations" does suffice.

5. INFERENCES INVOLVING PARTICULAR AND GENERAL STATEMENTS

Some apples in a bag

Here's the example again, in case you've forgotten it. You have been handed a bag of apples and told, truthfully, that some of them are rotten. The question was: Does it follow that some are not rotten? Does the statement that some of the apples are rotten imply the conclusion that some of the apples are not rotten? Putting the point more generally, I am asking whether from a PARTICULAR AFFIRMATIVE statement of the form

Some of the Xs are Ys

we can validly infer the corresponding PARTICULAR NEGATIVE statement of the form

Some of the Xs are not Ys.

Note: In traditional formal logic of the kind we have inherited from Aristotle, one of the ancient founders of the science of logic, these particular statement-forms are contrasted with the GENERAL or UNIVERSAL statement-forms "All of the Xs are Ys" (UNIVERSAL AFFIRMATIVE) and "None of the Xs are Ys" (UNIVERSAL NEGATIVE). We will say more about these later.

Before you go on to read what follows, I'd like you to think about the question again and see whether, now that you've learned a little of what's involved in valid reasoning, you're still inclined to give the answer you gave before. Remember:

(a) that the example should be viewed as one in which you're given certain information, in the premises, the *actual* truth or falsity of which you don't need to know in order to determine what else would have to be the case if the premises were true;

and (b) that the conclusion will follow from (be implied by) the premises only if there is no *possible* situation in which the premises would be true and the conclusion false.

So what is the answer to our question?

It is, "No". Such an inference would be invalid since "Some of the apples are not rotten" does *not* follow from "Some of the apples are rotten."

Those of you who thought it did follow may be comforted to know that you're not alone. The vast majority (usually about 80%) of people to whom I've put this question, frequently in lectures to logic students, and occasionally at dinner parties with friends, agree with you.

Unfortunately, however, logical correctness isn't determined by counting heads. So let me first try to explain why "Some of the apples are not rotten" *doesn't* follow from "Some of the apples are rotten". And then let me offer some diagnoses of why so many people mistakenly think that it does. (Note that, even if you got the answer right, you'll probably benefit from learning how best to justify your position when challenged. So read on.)

A counter-example

It isn't at all difficult to think of possible, indeed plausible, examples of situations in which the premise (that you've been told truthfully that some of the apples are rotten) would be true, and yet the conclusion (that some of them aren't rotten) would be false.

Suppose, for instance, that the gift-bearer had peeked into the bag and seen two or three rotten apples. She hasn't looked at them all, of course; so she doesn't know the condition of the other apples in the bag. One thing she does know, however, is that at least some of them are rotten; and she truthfully passes on this information. You then open up the bag and find that not just some but all of the apples are rotten. Can you fairly accuse her of making a mistake or of lying? Hardly. She had in fact told you the truth. The error would have been yours, not hers, if you had mistakenly inferred that there would be some good apples as well as bad ones.

The fact is, of course, that since "Some of the apples are rotten" is compatible with all of them being rotten, you can't infer, from the fact that some are, that some are not. It just doesn't follow; the one doesn't imply the other.⁷

The apples-in-the-bag case, of course, is only one instance out of countlessly many in which people make inferences from statements of the form "Some of the Xs are Ys". Here's another.

Another counter-example: the fishing expedition

Suppose that you're worried about the possible effects of a forestry company's recent herbicide-spraying program in the area surrounding your favorite fishing-pool. Could the fish in that pool have been affected? Might they all have developed cancerous lesions, you wonder; or might some of them still be fit for your evening meal?

On your next fishing expedition, you decide to find out. The first fish you catch is cancerous. So is the second. And the third. So what do you know so

⁷ Recall the logical lesson that we learned above, viz., that an argument is invalid if there is any statement which is consistent with its premises (would leave them true) but inconsistent with its conclusion (would make the conclusion false).

far? Obviously, you know that some of the fish have cancer.

Suppose you then infer: "Ah-hah, that's great. Now I can infer that some of the fish in the pool *don't* have cancer. So if I continue fishing long enough perhaps I'll catch some that will be good to eat." Would your reasoning be sound?

The absurdity of such reasoning should be obvious. From the fact that some of the fish are inedible, it by no means follows that some of them are perfectly edible. After all, if you subsequently were to throw a drag-net into the pool and inspect all the remaining fish, you might well find that every single fish in the pool has been affected. Not only would your earlier reasoning be logically unsound; it would be out of whack with reality.

The points made so far can be generalized. The particular affirmative statement

Some of the Xs are Ys

is consistent with the universal affirmative (general) statement

All of the Xs are Ys.

Hence it can't imply the particular negative statement

Some of the Xs are not Ys.

The logical intuitions of anyone who thinks the contrary are mistaken and in need of a little education.

A general refutation: Relations between "All", "Some" and "Some are not"

The erroneousness of inferring "Some of the Xs are not Ys" from "Some of the Xs are Ys" can be demonstrated in another way, too. For I think it will be agreed by almost everyone that "Some of the Xs are Ys" is not only consistent with "All of the Xs are Ys" but is implied by it.⁸ That is to say, almost everyone

⁸ [Note to logicians who may be worried, at this point, about problems of so-called "existential import". Although such problems may arise regarding "All Xs are Ys", they don't regarding "All of the Xs are Ys". Thus although statements of the form "All Xs are Ys" may sometimes be true in circumstances in which there aren't any Xs, statements of the form "All of the Xs are Ys" never are. On a Russellian analysis, statements of the form "The X is a Y" imply: (i) there is at least one X; (ii) there is at most one X; (iii) if anything is an X then it is a Y. It is surely in accord with this analysis to say that statements of the form "All of the Xs are Ys" are true just when the uniqueness condition (ii) is dropped but conditions (i) and (iii) are satisfied.]

will agree that - subject to the previously-noted proviso that the terms "X" and "Y" aren't empty - the universal affirmative statement:

All of the Xs are Ys

implies the particular affirmative statement:

Some of the Xs are Ys.

But now - and here's the rub - implication is a transitive relation. That is to say, if one statement P implies a second statement Q, and Q implies a third statement R, then it follows that P implies R.

Suppose, then, that you agree (as surely you must) that "All of the Xs are Ys" implies "Some of the Xs are Ys". Then, if you persist in your belief that "Some of the Xs are Ys" implies "Some of the Xs are not Ys", you are committed - by the transitivity of implication - to saying that "All of the Xs are Ys" implies "Some of the Xs are not Ys"!

But that is obviously absurd. It is impossible for it to be true that all the Xs are Ys, and for it also to be true that some of them aren't! The intuitions of someone who holds that "Some of the Xs are Ys" implies "Some of the Xs are not Ys" are not only mistaken; they need to be brought into harmony with their correct intuition that "All of the Xs are Ys" implies "Some of the Xs are Ys". Otherwise their intuitions will be inconsistent with one another and any persons holding them will be contradicting themselves.

Diagnoses of some confusions about "some".

Why are so many well-educated speakers of English so confused about the meaning of "some"? I suspect that one or more of the following diagnoses will fit a lot of the cases.

1. One reason is the failure to distinguish between

Some of the Xs are Ys

and

Only some of the Xs are Ys.

To be sure, the latter does imply that not all of the Xs are Ys; that is, it does imply "Some of the Xs are not Ys." But the former, as we've seen, doesn't. This error is somewhat akin to that of people who fail to distinguish between "If" and "Only if"; and, like the latter, it calls for a more discerning eye, ear, and mind.

2. A second, and much more prevalent, source of confusion about the meaning and implications of "Some of the Xs are Ys" is hasty generalization. In many of the possible situations in which it is true that some of the Xs are Ys it also just *happens* to be true that some of them are not. Some dogs have long hair; some do not. Some people I know understand the meanings of "some are" and "some are not"; some do not. One could go on, more or less indefinitely, with examples of cases in which both "Some of the Xs are Ys" and "Some of the X's are not Ys" happen in fact to be true. But it doesn't follow that in *all* cases in which one is true, so is the other. There are hosts of counter-examples. Such a generalization is quite unwarranted.

3. Moreover, there's a crucial distinction that needs to be emphasized here: the distinction between what happens to be true as a mere matter of fact, and what must be true as a matter of logic. It may be true as a mere matter of fact - that is in some *actual* circumstances - that some of the apples are rotten and that some of them are not. But it certainly isn't true as a matter of logic. It would be true as a matter of logic only if in all *possible* circumstances when some apples are rotten, some would also not be.

SECOND INTERLUDE

SUBJECTIVE VERSUS OBJECTIVE THEORIES OF MEANING

An even deeper source of confusion about the meaning and implications of statements of the form "Some of the Xs are Ys", I think, is to be found in a widely-held SUBJECTIVIST theory of meaning.

Meaning as a subjective matter

The notion that meaning is a subjective matter is afloat in many people's minds. Sometimes its plausibility derives from the fact that the term "meaning" itself has many meanings, being used by some to encompass both denotation and the fairly strict logical sense of "connotation", by others to encompass only the latter, and by still others to encompass the loose psychological sense of that word as well. (See the discussion of these different senses at the end of the Second Interlude in Chapter 1, pages 13-14.) Identify meaning with connotation, and connotation with mental associations, and understandably (even logically!)

one comes to think of questions about meaning and implication as if they were to be settled introspectively.

Such views can have even deeper roots, roots in the thinking of philosophers themselves! I'll describe two such philosophical sources of subjectivist accounts of meaning.

1. **The Circle of Ideas.**

During the seventeenth century, many influential philosophers, including Rene Descartes (1596-1650), Thomas Hobbes (1588-1679), and John Locke (1632-1704), put forward a doctrine that is still accepted by many: the doctrine that words stand for mental ideas. According to their view, a word like "lead" stands for our idea of lead, and our idea of lead, in turn, somehow "represents" a certain kind of substance. Ideas of things are, as it were, the intermediaries between words and things in the world.

Despite its initial plausibility, there are some pretty obvious flaws in this account of how language works.

One is that, if correct, this doctrine of meaning leads to an absurd account of the meanings of subject-predicate sentences. In general, if one utters a singular subject-predicate sentence of the form "S is P", where "S" is a singular subject-term (a subject-term that refers to a single thing, kind of thing, or set of things), and "P" is the predicate-term, then one intends to ascribe the property that the predicate connotes to the thing that the subject-term denotes. For instance, if one utters a sentence such as "Lead is heavy", one intends to attribute the property of being heavy to that very thing that the word "lead" stands for. But according to the doctrine we are criticizing, what the word "lead" stands for is an idea! Hence the doctrine implies that a statement like "Lead is heavy" ascribes heaviness not to lead but to our idea of lead! But, of course, my idea of lead isn't heavy (or light, for that matter). More generally, our ideas and thoughts aren't the kind of thing that have any weight at all - not in the literal sense, anyway. It is lead itself, not my idea of it, that has the property of being heavy.

A second objection - one well-known to students of the history of

philosophy - is that this psychologistic theory of meaning leaves us trapped within what has aptly been called "The Circle of Ideas". If words are always about ideas, not things, then we will be unable to talk about the external world or things within it. Indeed, we won't even be able to justify our belief that there is an external world beyond that of our mental states. Not surprisingly, Bishop Berkeley (1685-1753), who worked out, and cheerfully accepted, the implications of this way of thinking, was led to espouse a radical kind of Subjective Idealism according to which he knew of the existence of nothing but the contents of his own mind (Solipsism). Most of us, on reflection, are likely to realize that something must be wrong with a theory which leads to such solipsistic consequences - even if we can't say exactly where the error lies.

Yet the principal source of this error isn't too hard to detect. It lies within the mentalistic theory of meaning which Berkeley took over from his philosophical predecessors.

Unfortunately, however, these consequences are seemingly unknown to the many authors of contemporary textbooks on education, psychology, communications, and language who continue to tell their readers that the meanings of words, quite generally, are the ideas they stand for. Nor are dictionaries exempt from this sort of error. Webster's *New Collegiate Dictionary*, for instance, tells us that meaning is "the idea that something conveys to the mind." Evidently, those who are ignorant of mistakes in the history of ideas seem doomed to blithely repeat them.

2. Phenomenological methodology.

Phenomenology is a general philosophical movement, very influential in Europe and in certain literary circles elsewhere, whose exponents are characterized by a certain method of inquiry rather than by any distinctive doctrine. According to the phenomenological method, in order to determine what a word or other linguistic expression means all one needs to do is examine the contents of one's own mind to see what it means to *you*. On this account, all considerations having to do with the external world, objective world are to be, as they put it, "bracketed", i.e., excluded from consideration. As understood by many of its present-day adherents, the phenomenological method for determining

matters of meaning and implication involves introspectively exploring the "inner" world of consciousness to see what ideas are associated with what.

Something like this method for determining meanings can be traced back to works by the German philosophers Franz Brentano (1838-1917) and the early writings of Edmund Husserl (1859-1938), though Husserl later took some pains to disavow the seeming subjectivism of his views.⁹ It has more recently come to the forefront of public attention in some of the works of Martin Heidegger, (1889-1976) and the French existentialists Maurice Merleau-Ponty (1908-61) and Jean-Paul Sartre (1905-80). It is a method to which many people explicitly, and even more people implicitly, seem to subscribe.

Fairly obviously, it is some such account of meaning that is being invoked when people try to justify their claim that "Some of the Xs are Ys" means, or implies, "Some of the Xs are not Ys" by saying such things as:

"That is what 'some' means to me"

"I wouldn't say 'some' if I meant 'all'"

or

mean "If someone else told me that some Xs are Ys, I'd take it to
that some Xs are not Ys"

be "If someone tells me that some Xs are Ys, I'd take them to
implying that some Xs are not Ys."

Such persons are reporting how their own minds work. They are introspecting, and telling us what *they* mean, what *they* would say, and what *they* would infer. They are "bracketing" external considerations to do with situations outside their own minds, ignoring considerations to do with the objectively ascertainable circumstances in which sentences of the form "Some of the Xs are Ys" are actually used. Yet the fact of the matter is that the meanings and implications of our words and sentences are not in general determined by us as individuals. Rather, they are determined by socially accepted conventions for their proper application. Only in the relatively rare case where someone coins a

⁹ Indeed, a major part of Husserl's later philosophical mission consisted in trying to rid philosophical and logical enquiries of any traces of what he called Psychologism. Unfortunately, there is a tension in his work between his avowed aims and his methods for achieving them. And the phenomenological method he espoused has, more often than not, been construed in a subjectivist sense. It is his method, so construed, that I am making the target of criticism here.

neologism or explicitly defines an existing term as having a special, technical meaning, is it up to an individual to determine what a word means. To suppose otherwise is to be guilty of the kind of fallacy that Lewis Carroll revelled in exposing in his philosophically astute "children's books", *Alice in Wonderland* and *Through the Looking Glass*. I refer, in particular, to Humpty Dumpty's fallacy of supposing that he could make words mean exactly what he wanted them to mean - the issue, in his eyes, being solely that of who was to be "boss". To be sure, words don't get their meanings by virtue of some sort of natural tie to things in the world. Nor do they get them by virtue of the dictates of some deity or other. They get their meanings from us and the conventions we adopt for their use. But the question as to what those conventions are cannot be settled simply by inspecting the contents of one's own individual consciousness.

General criticisms of subjectivist doctrines of meaning.

Meaning is not to be found, as it were, "in the head".¹⁰ The introspective, ideational, phenomenological, approach is wrong on two scores.

1. It fails to take into account the fact that the ideas one associates with the use of a word or sentence, may be associated with it for the wrong reasons - because, for example, one has confused "some" with "only some", or because one has been guilty of overhasty generalization of the kind discussed a few paragraphs back.
2. It fails to take into account, also, the fact that introspective reports are notoriously unreliable. Consider, for instance, the case where someone reports that she'd never say "Some of the Xs are Ys" unless she knew "Some of the Xs are not Ys" to be true as well. The fact is that, if faced with the fish-pond situation in which she has caught two or three cancerous fish, she *would* say "Some of these fish are cancerous" without knowing whether or not some fish weren't.
3. The most telling criticism of such subjectivist theories of meaning is that it leads to fallacious thinking. Operate with it, in accordance with its account of how meaning is determined, and - as we've seen - you'll be led to make

¹⁰ This point has been made often, and convincingly, by American philosopher Hilary Putnam. See, for instance, his "The Meaning of 'Meaning'."

inferences that will lead you astray, inferences in which you may lead yourself and/or others from truth to falsity. Operate with it and you may even find yourself accepting contradictory beliefs.

Meaning as an objective matter

What we need, for a self-consistent theory of meaning, is an account that takes into consideration the so-called "external" situations in which language is used, the actual and possible circumstances in which words are correctly applied and sentences are truthfully uttered.

What might such an account look like?

By way of illustration, I invite you to consider a MODEL of the external, objective, possible situations that would have to be considered in order to determine (ascertain) the meanings of sentences that are instances of sentence-forms "All the Xs are Ys", "None of the Xs are Ys", "Some of the Xs are Ys", and "Some of the Xs are not Ys". It is a model which should help you see not only what is MEANT by statements of these various forms, but also what is IMPLIED by each of those statements.

Consider the following five diagrams, invented by the founder of topology, the Swiss mathematician Euler:




Figure 1

Each circle represents the class of things to which a term such as "apples" or "rotten things" applies. In each case we are to presume that the class is NON-EMPTY; that is, we are to presume that there are things to which the terms apply - for instance, that there are at least some apples and at least some rotten things.

These five diagrams exhaust all the topological possibilities: they represent all the LOGICALLY POSSIBLE ways in which two non-empty classes may be related to one another.

In diagram 1, the class Y is a subclass of X. Hence all of the Ys are Xs.

In diagram 2, the class X is a subclass of Y. Hence all of the Xs are Ys.

In diagram 3, the classes overlap. Hence some of the Xs are Ys and some of the Ys are Xs.

In diagram 4, the classes are coextensive. Hence all of the Xs are Ys, and all of the Ys are Xs.

And in diagram 5, the classes are exclusive of one another. Hence none of the Xs are Ys and none of the Ys are Xs.

Now ask yourself these questions:

In which of these possible cases are there Xs that are Ys, i.e., in which cases would "Some of the Xs are Ys" be true?

In which cases are there Xs that aren't Ys, i.e., in which cases would "Some of the Xs are not Ys" be true?

In which cases are there no Xs that aren't Ys, i.e., in which cases would "All the Xs are Ys" be true?

And in which cases are there no Xs that are Ys, i.e., in which cases would "None of the Xs are Ys" be true?

The answers to these questions about the POSSIBLE SITUATIONS in which statements of these four forms are true are given, as Ts and Fs (for "True" and "False", respectively) in the following table:

Statements	Possible situations (as per diagrams)				
	1	2	3	4	5
All the Xs are Ys:	F	T	F	T	F
Some of the Xs are Ys:	T	T	T	T	F
Some of the Xs are not Ys:	T	F	T	F	T
None of the Xs are Ys:	F	F	F	F	T

Table 1

This table provides us with what we shall call the TRUTH-CONDITIONS

for statements of these four forms. It gives us an account of the possible external conditions in which our statements would be true and of the possible external conditions in which they would be false. The truth-conditions for each statement give us an effective grip on what that statement MEANS.

By the same token, the truth-conditions for these statement-forms give us a clear guide to which statements IMPLY which, and which do not imply which. They show, for instance, that since every possible situation in which

All the Xs are Ys

is true (viz., 2 and 4) is one in which

Some of the Xs are Ys

is true (viz., 1, 2, 3, and 4), the former implies the latter. Hence we can never go wrong if we reason that whenever the former is true so is the latter.

Again, they show that since only some of the possible situations in which

Some of the Xs are Ys

is true (1, 2, 3, and 4) are situations in which

Some of the Xs are not Ys

is true (viz., 1, 3, and 5), the former, though CONSISTENT with the latter, does *not* imply it. We might well be led into error were we to reason as though it did.

These are not the only logical relationships that we can, as it were, "read off" our table of truth-conditions. We'll take a look at others in a later chapter. There, too, we'll give other illustrations of this important logical fact:

When it comes to determining matters of meaning and implication, the Truth-conditional Theory gives us the right results whereas subjectivist theories, by way of contrast, yield the wrong results.

One more point before moving on. So far, in this section, we've been

examining cases in which lots of people's logical intuitions let them down. Now think about this:

If otherwise well-educated but logically untutored people can make mistakes about the meaning and implications of sentences involving such simple expressions as "brothers", "if . . . then . . .", and "some", how much more prone to conceptual confusion and inferential error are they likely to be when it comes to dealing with the "big" issues of religion versus science, faith versus reason, romanticism versus logic, freedom versus determinism, capitalism versus socialism, the pros and cons of abortion, capital punishment, and so on?

Not only are the concepts and theories involved in these larger issues much more complex. In our reasoning about them we are even more likely to be carried along by mere association of ideas, uncritically accepted dogma, and emotionally charged rhetoric. Yet here, where the calm voice of reason is perhaps most needed, it is most often absent - and untaught.

6. THE DETECTION OF INCONSISTENCIES

The Barber

This example is far more difficult than the others, and I although I expected many readers to reason that something must be wrong with it (since otherwise I wouldn't have asked whether it could *all* be true), I didn't really expect many to detect exactly what the problem is with my partly autobiographical story.

So why, you may ask, did I include it? Well, let me first tell you about the barber story and why it can't be true. Then I'll explain its significance.

A contradiction.

My story can't be true for the simple reason that it contains a contradiction. Recall that I said that the village barber acts in accordance with two conditions:

- (1) if someone in the village cuts his own hair, the barber doesn't cut it;

(2) if someone in the village doesn't cut his own hair, the barber does cut it.

At first glance, or hearing, my story seems to be a perfectly plausible description of a possible situation. But now ask yourself:

"Who cuts the hair of the village-dwelling barber?"

Either he cuts it or he doesn't cut it. If he does cut it, then it follows (from condition (1)) that since he lives in the village, he doesn't cut his own hair. And if he doesn't cut his own hair, then it follows (from condition (2)) that he does cut it. In other words, the seemingly possible - indeed plausible - description of the village barber yields the conclusion that he cuts his hair if and only if he doesn't cut it. But that is impossible.

My story, then, contains a contradiction. It can't be true that I had my hair cut by such a barber. No such barber can exist. A world satisfying the descriptions given in my story isn't even possible; such a world is an IMPOSSIBLE WORLD.

General point of the story.

My reasons for telling this story, and inviting you to check out its logical credentials, were:

1. to illustrate how easy it is to overlook logical impossibilities in our own and others' statements and beliefs;
 2. to illustrate how the detection of inconsistencies in what we've hitherto taken for granted can have drastic ramifications for whole fields of enquiry, including mathematics and logic.
1. As to the ease with which we overlook inconsistencies, this should already have been evident from the fact that so many of us are prepared (at least at different times, and sometimes even in one breath) to claim both (a) that "All are" implies "Some are", and (b) that "Some are" implies "Some are not", despite the fact that claims (a) and (b) taken together lead to contradiction. If we can overlook inconsistencies in our use of simple concepts such as these, how much

more prone must we be to overlooking them in our thinking about more complex ones?

2. This brings me to the second point. For the fact is that the "Barber Paradox", as it is widely known, doesn't originate with my visit to Lech or merely have the status of a brain-teaser. A simpler version was first told, early this century, as a way of illustrating the kind of paradox discovered in an important branch of mathematic, viz., set-theory, as it was then understood, by the great philosopher-logician, Bertrand Russell (1872-1970).

Intuitively, the following seem obvious.

(1) Things can belong to classes or sets. You and I are members of the set of human beings; Earth and Venus are members of the set of planets; and so on.

(2) Something is a member of a set if and only if it falls under the description that defined membership of that set. You and I are members of the set of human beings iff you and I fall under the description "is a human being"; Earth and Venus are members of the set of planets iff they fall under the description "is a planet".

(3) Some sets, e.g., the set of natural numbers, have infinitely many members. Some, e.g., the set of planets, have finitely many members. Some, e.g., the set of persons writing this book, have just one member. And some, e.g., the set of prime numbers between 7 and 11, have no members at all, i.e., are empty sets.

(4) Some sets are members of themselves. The set of thinkable things, for example, falls under the description "is a thinkable thing" and hence is a member of itself. And the set of non-planets is likewise a member of itself since it is a non-planet.

(5) Some sets are not members of themselves. The set of planets, for instance, since it fails to fall under the description "is a planet", is not a member of itself; that is, the set of planets is not itself a planet. (If it were, then there would be ten of them, not just nine!)

(6) Since there are many sets which are not members of themselves, there must be a set - call it S^* - of all the sets that are not members of themselves.

So far, so good, it would seem. Everything we've said seems plausible, just as did the story of the barber. But now ask yourself the question (analogous to the question, "Who cuts the barber's hair?"):

"Is S^* a member of itself or not?"

On the face of it, S^* must be one or the other. Let's suppose, then, that S^* (defined as the set of all sets that are not members of themselves) is a member of itself. Then, since a set is a member of itself iff it falls under the description characteristic of that set, and S^* (by definition) falls under the description "is not a member of itself", it follows that S^* is one of those sets that are *not* members of themselves. So if S^* is a member of itself, then it isn't.

Obviously, that conclusion can't be accepted. So let's suppose, on the other hand, that S^* (defined as the set of all sets that are not members of themselves) is not a member of itself. Then, since a set is not a member of itself iff it fails to fall under the description that is characteristic of the set, it follows that S^* must fail to fall under the description "is not a member of itself", and hence that S^* must be one of those sets that are members of themselves after all. So if S^* is not a member of itself, then it is.

In short, the set of all sets that are not members of themselves is a member of itself if and only if it isn't! [If you can't follow the reasoning here, don't get hung up on that fact. Take my word for it that it is perfectly valid, and accepted as such by mathematicians and logicians. What is important for our present purposes is the outcome of the reasoning, not the details of how we got there.]

Here, in what is now known as "Russell's paradox", we have another example of how easy it is for us to overlook inconsistencies, especially when dealing with relatively sophisticated concepts.

The great German mathematician, Gottlob Frege (1848-1925), was committed to each of (1) through (6) above and had used the so-called "naive set

theory" of which they are part as the basis for his attempt to derive arithmetic from set theory, and that in turn from logic. Upon hearing from Russell, in 1901, that set theory - as he understood it - was self-inconsistent, he came to the conclusion that his life's work was in ruins.

As things turned out, it wasn't. Nevertheless, certain of the seemingly plausible assumptions of naive set theory did have to be abandoned. Mathematical logicians, these days, usually think of sets as being determined not (as (2) claims) by specifying some description characteristic of all their members, but simply by specifying the members themselves.

The discovery of a contradiction in what we otherwise were inclined to accept as true, can lead to a radical revision in our beliefs. Those who realize that it is self-contradictory to believe both that "All the Xs are Ys" implies "Some of the Xs are Ys", and that "Some of the Xs are Ys" implies "Some of the Xs are not Ys", can be made to see that they must, in all consistency, give up one or the other (preferably the latter). And those who realize that naive set theory is self-contradictory, can be set on the path towards constructing sounder bases for this important branch of mathematics and logic.¹¹

¹¹ Another instructive paradox is the so-called Liar Paradox. It takes many forms. The ancient Cretan philosopher, Epimenedes, reputedly claimed that everything a Cretan says is false.

A more modern version is that of someone who says: "I am lying." Is her statement true or false? If what she is saying is true, then since she is truly saying that she is lying, it is a lie and hence false. If what she is saying is false, then it is false that she is lying, and hence what she is saying must be true.

A third example is that of the sentence: "This sentence says something false." Russell used examples such as these to argue that certain formulations of words, though grammatically correct, are strictly nonsense. Others have concluded only that they "don't express propositions", i.e., that they don't really express anything true or false despite their apparent claims to do so.