Chapter 4: Deduction and Logic

"The supreme task of the physicist is to arrive at those universal elementary laws from which the cosmos can be built up by pure deduction. There is no logical path to these laws; only intuition, resting on sympathetic understanding, can lead to them." [Einstein, 1879-1955]

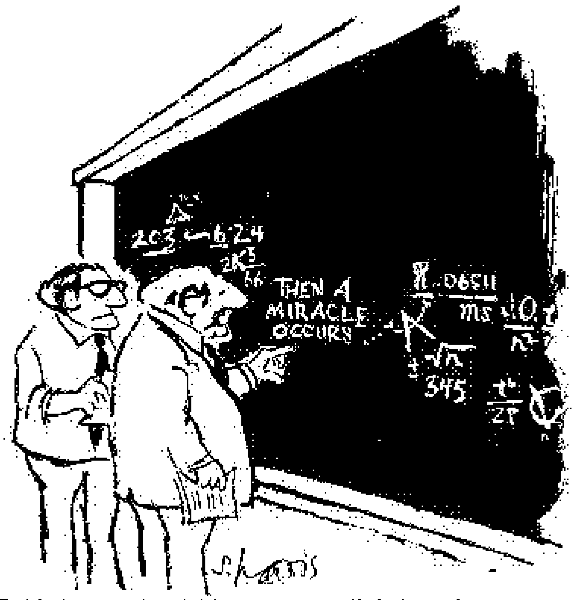

"I think you should be more explicit here in step two."
[Harris, 1970]

"'From a drop of water,' said [Sherlock Holmes], 'a logician could infer the possibility of an Atlantic or a Niagara without having seen or heard of one or the other. So all life is a great chain, the nature of which is known whenever we are shown a single link of it. Like all other arts, the Science of Deduction and Analysis is one which can only be acquired by long and patient study, nor is life long enough to allow any mortal to attain the highest possible perfection in it." [Doyle, 1893b]

*   *   *

Scientific deduction bears little similarity to the mythical conception conveyed by Sherlock Holmes. In science, obvious deductions are ubiquitous, insightful deductions are sporadic, and neither is infallible. We wield our logic with confidence, not noticing our occasional deductive errors. Before declaring that you are immune to such errors and skipping to the next chapter, please take ten minutes to attack the following problem:

> Imagine that four 3"x5" cards are on the table. You can see that each card has a single letter or number on its top: one has the letter 'A', one has 'B', one has the number '4', and one has the number '7'. You may assume that each card contains a single letter on one side and a single numeral on the other side. What cards is it necessary to turn over, to evaluate the validity of this rule: 'If a card has an A on one side, then it has a 4 on the other side'?

This problem, posed by Wason [1966], is considered by many to be a good example of the type of deductive decision-making that scientists face. Only 10% of college students answer the card problem correctly [Kuhn et al., 1988]. I suspect that you, like I, spent only a minute or two on the problem and got the wrong answer. Before proceeding, please consider the problem once more, this time actually using some props such as post-its, sheets of paper, or pencil and pad. Imagine that each card flip will be a major, time-consuming experiment. Will each experiment really be crucial to testing the hypothesis?

The correct answer to the card problem above is the two cards *A* and *7*. Many people answer *A* and *4*. The *B* card is clearly not useful, because it cannot prove or disprove the rule regardless of what is on the other side. Surprisingly, however, the same is true for the *4* card: even if it has an *A* on the other side, it supports but neither proves nor disproves the rule that any card with an *A* on one side has a *4* on the other side. In contrast, flipping the *7* card does test the rule, because the rule would be disproved if the other side is an *A*.

Many philosophers of science interpret the *A* & *4* answer as evidence of a confirmation bias: the chooser of the *4* card is seeking a result that confirms the hypothesis, rather than choosing the *7* card and potentially disproving the hypothesis. Scientists, in contrast, may justify choice of the *4* card as a search for patterns where they are most likely to be found. Not choosing the *7* card, however, is a failure to consider deductively the importance of potential results.

Two problems can involve identical deductive logic yet differ in difficulty. How a deductive problem is posed can affect the likelihood of correct results. Concrete examples are easier to solve than are the same problems expressed in symbols. For example, the success rate on the problem above was increased from 10% to 80% [Kuhn et al., 1988] when the problem was recast: given an envelope that may or may not be sealed and may or may not have a stamp on it, test the hypothesis, 'if an envelope is sealed, then it has a 5-pence stamp on it'.

Our greater facility with the concrete rather than with abstract deductions challenges the very basis of this decision-making. Possibly we do not even make decisions based on learned rules of formal logic [Cheng and Holyoak, 1985], but instead we recognize conceptual links to everyday experience [Kuhn et al., 1988]. The problem must seem real and plausible if there is to be a good chance of a successful solution; thus the postage problem is easier than the 4-card problem. In deductive logic, a similar strategy is often useful: recast the problem so that the logical structure is unchanged but the terms are transformed into more familiar ones. This technique, known as substitution, is one that we shall employ later in this chapter.

The four-card problem illustrates several points:

- prior thought can prevent needless experiments;
- sketches can be valuable in avoiding error;
- the same problem is more likely to be solved correctly if in familiar terms than if in abstract terms;
- confirmation bias is present in science, but to some extent it is a normal consequence of our pervasive search for patterns; and
- many people's 'deductive thinking' may actually be inductive pattern recognition of a familiar deductive form.

*   *   *

Logic

Logic means different things to different people. To Aristotle (384-322 B.C.), the 'Father of Logic', it was a suite of rules for deductive evaluation of syllogisms. To Peter Abelard (1079-1142) and William of Occam (1285-1349), Aristotelian logic was a useful launching point for development of a more comprehensive logic. G. W. Leibniz (1646-1716) sought to subsume all types of arguments within a system of symbolic logic. During the last century, symbolic logic has been the focus of so much study that it almost appeared to be the only type of logic. A notable exception was John Stuart Mill's Canons of inductive logic (Chapter 3).

Logic *is the science of argument evaluation*; it includes methods and criteria for deciding whether arguments are reliable. In this context, the term 'argument' has a meaning quite distinct from its everyday use as a difference of opinion: an **argument** *is a group of statements, consisting of evidence and a conclusion. Evidence statements are called premises,* and the conclusion is claimed to follow from these premises. For example, the following argument consists of three simplified statements, of which the first two are premises and the third is a conclusion:

All *A* are *B*.

All *B* are *C*.

Therefore, all *A* are *C*.

* * *

Deduction vs. Induction

Scientific logic has two distinctive branches: deduction and induction. Surprisingly, most scientists do not know the difference between these two types of inference. I, for example, used the word 'deduced' incorrectly in the title of my first major paper. Sherlock Holmes is indelibly associated with deduction, yet many of his 'deductions' were actually inductive interpretations based on subtle evidence.

To a first approximation, deduction is arguing from the general to the particular, whereas induction is arguing from the particular to the general [Medawer, 1969]. Often scientific induction does involve generalization from the behavior of a sample to that of a population, yet the following inductive argument goes from the general to the particular:

> In spite of many previous experiments, never has a relationship between variables *X* and *Y* been observed. Therefore, this experiment is unlikely to exhibit any relationship between *X* and *Y*.

In a deductive argument, the conclusion follows necessarily from the premises. In an inductive argument, the conclusion follows probably from the premises. Consequently, totally different standards are applied to deductive and inductive arguments. Deductive arguments are judged as valid or invalid by a black-or-white standard: *in a valid deductive argument, if the premises are true, then the conclusion must be true*. Inductive arguments are judged as strong or weak according to the likelihood that true premises imply a correct conclusion. Statistical arguments are always inductive. The following argument is inductively strong but deductively invalid:

> No one has ever lived more than 150 years.
>
> Therefore I will die before age 150.

A mnemonic aid for the difference between deduction and induction is: **de**duction is **de**finite; **in**duction is **in**definite and uncertain.

Both deductive and inductive arguments are evaluated in a two-step procedure:

- Does the conclusion follow from the premises?
- Are the premises true?

The order of attacking the two questions is arbitrary; usually one considers first whichever of the two appears to be dubious. The distinction between induction and deduction lies in the evaluation of whether the conclusion follows from the premises.

Here the focus is on deduction; induction was considered in Chapter 3. Before leaving the deduction/induction dichotomy, however, two common fallacies must be dispelled: ‘scientific deduction is superior to induction,’ and ‘scientific induction is superior to deduction.’ Three centuries ago, great minds battled over whether science *should be* deductive or inductive. René Descartes argued that science should be confined to the deductively certain, whereas Francis Bacon argued that the majority of scientific discoveries were empirical, inductive generalizations. A hallmark of the inception of rapid scientific progress was the realization that both deduction and induction are necessary aspects of science (Chapter 1). Yet the battle continues, fueled by misconceptions. For example, theoretical physicists such as Einstein probably would be outraged by the following statements from Beveridge’s [1955] book on scientific methods:

> “Since deduction consists of applying general principles to further instances, it cannot lead us to new generalisations and so cannot give rise to major advances in science. On the other hand the inductive process is at the same time less trustworthy but more productive.”

Inevitably, theoreticians value deduction and empiricists value induction, but the choice is based on taste rather than inherent superiority.

*   *   *

Scientific deduction uses the science of deduction, but the two do not share the same values or goals. Evaluating the validity of arguments is a primary objective of both, but scientific deduction places more emphasis on the premises. How can they be tested? Can the number of premises, or assumptions, be reduced, and if so what is the impact on the conclusion? How sensitive is the argument to the definition of terms in the premises? Are the premises themselves conclusions based on either deductive or inductive interpretation of other evidence?

Some scientists use a somewhat bootstrap logic that would be abhorrent to logicians. The technique is to tentatively assume an untested premise, and then see where it leads in conjunction with other, more established premises. If the resulting conclusion is one that is independently valued, perhaps on the basis of other deductive paths or perhaps on grounds of elegance or simplicity, then the premise may be tentatively accepted. These other standards of hypothesis evaluation are discussed more fully in Chapter 7.

*   *   *

Deductive Logic

Everyday language provides myriad opportunities for obscuring premises and conclusions, so the first step in evidence evaluation is usually the identification of premises and conclusion. Opinions, examples, descriptions, and many explanations are neither premise nor conclusion and are consequently not integral parts of an argument. Frequently, obvious premises are omitted from an argument:

> “Publish or perish” is an argument of the form:
>
> all A are B,
>
> not B,
>
> ∴ not A.
>
> Here we use the symbol ‘∴’ to indicate ‘therefore’. The premises are ‘all successful scientists are paper publishers’ and ‘consider someone who is not a paper publisher’; the conclusion is ‘that person is not a successful scientist’.

Premises may begin with one of the following flags: because, due to, since, given that, owing to, as indicated by, in that, . . . Likewise, most conclusions have an identifying flag: therefore, consequently, thus, accordingly, hence, so, as a result, it follows that, . . . Usually the conclusion is the first or last statement in an argument. Sometimes, however, one has to search for the conclusion by asking oneself, 'What is the author trying to convince me of?' For example, examine the following argument and identify the premises, conclusion, and any extraneous statements.

> Why should I have to study history? I am a scientist, I have more than enough to do already, I don't like history, and history is irrelevant to science.
>
> If one interprets the conclusion as 'History is irrelevant to me,' then the salient premises are 'History is irrelevant to scientists' and 'I am a scientist.' If one interprets the conclusion as 'History is a waste of time for me,' then the supporting premises are 'History is irrelevant to scientists,' 'I am a scientist,' and 'Doing history would prevent me from doing something more worthwhile.' The logic is valid, but some of the premises are dubious.

*   *   *

With deductive logic, each statement in the argument is either true or false. For the conclusion to be true, two critical preconditions must be met. First, the premises must be true. Second, the form of the argument must be valid. *A valid deductive argument is one in which the conclusion is necessarily true if the premises are true*. Validity or invalidity is totally independent of the correctness of the premises; it depends only on the form of the argument -- thus the term **form**al logic.

The following arguments demonstrate the distinction between the roles of premises and of logical form in determining the correctness of a conclusion:

All dogs are cats. All cats are animals. Therefore, all dogs are animals.	Valid form, but one false premise, so the argument is incorrect (although the conclusion happens to be true).
All dogs are mammals. All mammals are animals. Therefore, all dogs are animals.	Valid form, true premises, so the argument is correct and the conclusion must be true.
All dogs are mammals. All cats are mammals. Therefore, all dogs are cats.	True premises, but invalid form, so the argument is invalid and does not yield this conclusion.

For these three examples, the reader already knows which conclusions are true and which are false without even evaluating the arguments. For scientific arguments, however, it is crucial that one considers separately the two elements -- premise correctness and argument form -- rather than accept or reject the argument based on whether or not the conclusion sounds right. Evaluation of premises requires subjective judgment based on local expertise. Evaluation of argument form, in contrast, is objective. With some practice and a few guidelines, the reader can avoid using invalid argument forms and recognize them in publications. Such is the main goal of this chapter.

*   *   *

Classification Statements

A building block of deductive logic is the classification statement; logicians use the term categorical proposition. The classification statement consists of a subject and predicate, and it states that members of the subject category are or are not included in the predicate category. For example, the statement 'all scientists are people' is a classification statement, in which '**s**cientists' is the **s**ubject and '**p**eople' is the **p**redicate. The four types of classification statement are:

- *All S are P:* The entire subject class lies within the predicate class. Every member of the subject class is also a member of the predicate class.
- *No S are P:* The entire subject class is excluded from, or outside, the predicate class. No member of the subject class is a member of the predicate class.
- *Some S are P:* At least one member of the subject class lies within, and is a member of, the predicate class.
- *Some S are not P:* At least one member of the subject class lies outside, and is not a member of, the predicate class.

Note that 'some' means at least one; it does not mean 'less than all'. Thus it is possible for both statements 'All *S* are *P*' and 'Some *S* are *P*' to be true for the same *S* and *P*; if so, the former statement is more powerful. Similarly, both statements 'Some *S* are *P*' and 'Some *S* are not *P*' may be true for the same *S* and *P*.

The statements 'All *S* are *P*' and 'No *S* are *P*' are sometimes referred to as universal statements because they apply to every member of a class. In contrast, the statements 'Some *S* are *P*' and 'Some *S* are not *P*' apply not to every member but instead to a particular subset; thus they are referred to as particular statements.

Deductive Aids: Venn Diagrams and Substitution

The four classification statements can be illustrated diagrammatically as shown in Figure 17.

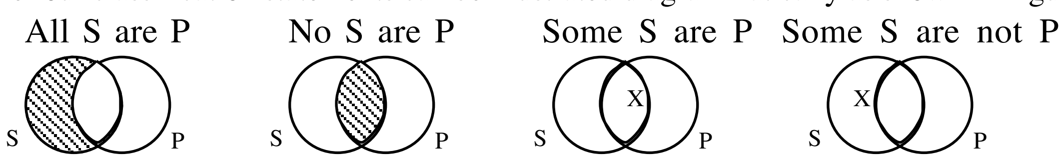

Figure 17. Classification statements, expressed as Venn diagrams.

John Venn, a 19th-century logician, invented this technique of representing the relationship between classes. Each class is represented by a circle; in this case there are only the two classes *S* or *P*. Potential members of the class are within the circle and individuals not belonging to the class are outside the circle. The overlap zone, lying within both circles, represents potential members of both classes. Hatching indicates that a zone contains no members (mathematics texts often use exactly the opposite convention). An *X* indicates that a zone contains at least one ('some') member. Zones that contain neither hatching nor an *X* may or may not contain members. In the next section, we will observe the substantial power of Venn diagrams for enhancing visualization of deductive statements or arguments. For now, it suffices to understand the Venn representations above of the four classification statements:

- *All S are P*: The zone of *S* that is not also *P* is empty (hatched), and the only possible locations of *S* are in the zone that overlaps *P*. Ergo, all *S* are *P*.
- *No S are P*: The zone of *S* that overlaps *P*, *i.e.* that is also *P*, is empty.
- *Some S are P*: The *X* indicates that at least one member lies within the zone that represents members of both *S* and *P*. The remaining members of *S* or *P* may or may not lie within this zone.

• *Some S are not P*: The *X* indicates that at least one member lies within the zone that represents members of *S* but not of *P*. Other members of *S* may or may not lie within *P*.

*   *   *

Substitution is a powerful technique for recognizing valid and invalid deductive arguments. Validity depends only on the form of the argument. Therefore, we can replace any arcane or confusing terms in a deductive argument with familiar terms, then decide whether or not the argument is valid. For example, the following four arguments all have the same invalid form:

> If a star is not a quasar, then it is theoretically impossible for it to be any type of star other than a neutron star. This follows from the fact that no neutron stars are quasars.
>
> No neutron stars are quasars. Therefore, no non-quasars are non-neutron stars.
>
> No S are P. ∴ no non-P are non-S
>
> No cats are dogs. Therefore, no non-dogs are non-cats.

Recognizing that the first three arguments are invalid is easy for some readers and difficult for others. Some of us experience mind-glaze when faced with arguments involving unfamiliar and highly technical terms; others find abstract, symbolic notation even more obscure. Some can analyze arguments easier when the argument is in a standard notation; others prefer their arguments to be couched in everyday language. Everyone can immediately recognize the fallacy of the cats-and-dogs argument, for obviously the world is full of objects that are neither cat nor dog. If this cats-and-dogs argument is invalid, then the other three arguments *must* be invalid because they have the same form.

Substitution relies on four principles that we have encountered in this chapter:

• Validity or invalidity of a deductive argument depends only on the form of the argument, not on its topic (*note*: this is not true for inductive arguments).

• A valid deductive argument is one in which the conclusion is necessarily true if the premises are true (*note*: this is not true for inductive arguments).

• If we know that the premises of an argument are true and yet the conclusion is false, then the argument *must* be invalid.

• Validity or invalidity is much easier to recognize for arguments about familiar objects than for abstract arguments.

To employ substitution, simply identify the elements of the argument and replace each element with a familiar term. In the examples above, the elements are neutron stars and quasars, or *S* and *P*, or cats and dogs, and the structural equivalents are *S*=neutron stars=cats and *P*=quasars=dogs. Formal logic assumes that the premises are true, so it is easiest if one picks substitutions that yield a true initial statement. Then, an absurd result can be attributed correctly to invalid logic.

Substitution may be the main way that most people (logicians excluded) evaluate deductions, but this method seldom is employed consciously. Instead, we unconsciously perceive that an argument is familiar, because it is similar in form to arguments that we use almost every day. Conversely, we may recognize that an argument sounds dubious, because it seems like a distortion of a familiar argument form. With that recognition, we then can deliberately employ substitution to test the argument.

*   *   *

Logically Equivalent Statements

Venn diagrams permit us to identify or remember **logically equivalent statements**. Such statements have exactly the same truth value (whether true or false) as the original. The Venn diagrams in Figure 18 permit us to identify which apparent equivalences are valid (identical Venn diagrams) and which are invalid (different Venn diagrams).

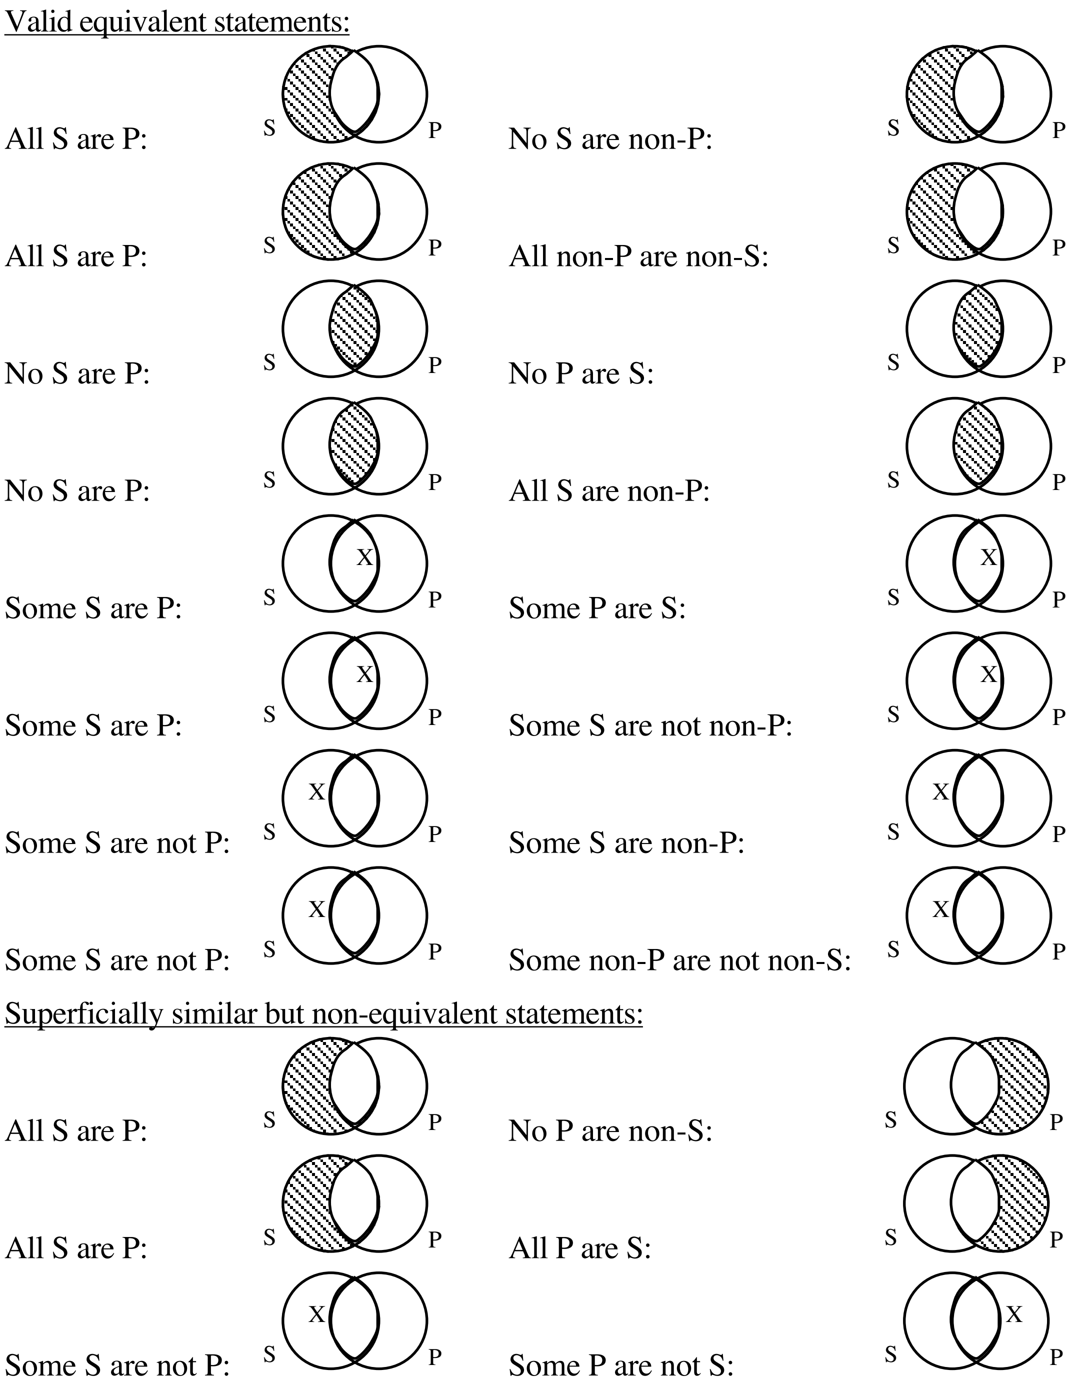

Figure 18. Valid and invalid equivalent statements, and their Venn diagrams.

Logicians use the terms conversion, obversion, and contraposition to define three types of logically equivalent statements, but we will not need to memorize these terms. Below are listed on the right the only logically equivalent statements to those on the left:

Initial statement	Logically equivalent statements	
All S are P.	No S are non-P.	All non-P are non-S.
No S are P.	No P are S.	All S are non-P.
Some S are P.	Some P are S.	Some S are not non-P.
Some S are not P.	Some S are non-P.	Some non-P are not non-S.

Some logically equivalent statements seem cumbersome and overloaded with negatives. That apparent weakness is a strength of the concept of logical equivalence, for we may encounter a statement on the right and want to translate it into a familiar classification statement.

The concept of logical equivalence can also be useful in experimental design. For example, it might be impossible to show that 'some *S* are *P*' but easy to show that 'some *P* are *S*'. In Chapter 7 we will consider the Raven's Paradox: the two statements 'All ravens are black' and 'All non-black things are non-ravens' may be logically equivalent, but testing the latter would involve an inventory of the universe.

*   *   *

For recognizing logically equivalent statements, substitution is an alternative to Venn diagrams. For example, replace *S* with **s**cientists and replace *P* with either **p**eople, **p**hysicists, or **p**oliticians, whichever gives a true initial statement:

Valid equivalent statements:

All Scientists are People.	No Scientists are non-People.
All Scientists are People.	All non-People are non-Scientists.
No Scientists are Politicians.	No Politicians are Scientists.
No Scientists are Politicians.	All Scientists are non-Politicians.
Some Scientists are Physicists.	Some Physicists are Scientists.
Some Scientists are Physicists.	Some Scientists are not non-Physicists.
Some Scientists are not Physicists.	Some Scientists are non-Physicists.
Some Scientists are not Physicists.	Some non-Physicists are not non-Scientists.

Non-equivalent statements:

All Scientists are People.	No People are non-Scientists.
All Scientists are People.	All People are Scientists.
Some Scientists are not Physicists.	Some Physicists are not Scientists.

*   *   *

Relationships Among Statements

The four types of classification statement are formally related in truth value, regardless of the subjects of the statements. The relationships can be summarized in what is called the **square of opposition** (Figure 19).

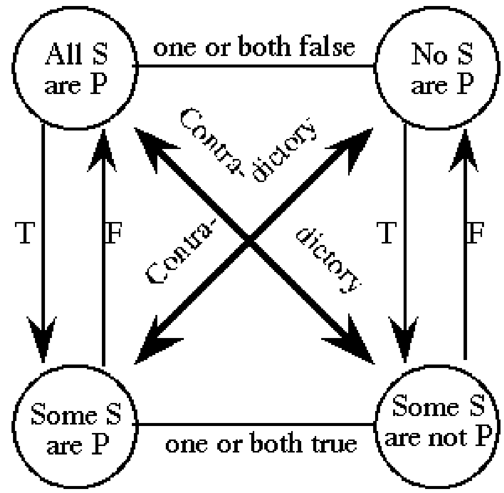

Figure 19. Square of opposition.

The strongest relationship among the statements is that of contradiction along the diagonals: if a statement is true, then its diagonal is false, and vice versa. Without even substituting familiar terms for the subject and predicate, one can recognize readily that:

• 'All *S* are P' contradicts the statement 'Some *S* are not P', and

• 'No *S* are P' contradicts the statement 'Some *S* are P'.

Horizontally along the top, one or both of the statements invariably is false:

- If 'All *S* are *P*' is true, then 'No *S* are *P*' must be false;
- If 'No *S* are *P*' is true, then 'All *S* are *P*' must be false;
- If either 'All *S* are *P*' or 'No *S* are *P*' is false, we cannot infer that the other statement is true; possibly both are false and 'Some *S* are *P*'.

Horizontally along the bottom, one or both of the statements invariably is true:

- If 'Some *S* are *P*' is false, then 'Some *S* are not *P*' must be true;
- If 'Some *S* are not *P*' is false, then 'Some *S* are *P*' must be true;
- Both statements may be true: some *S* are *P* while other *S* are not *P*.

Vertically, the statements lack the perfect symmetry that we saw diagonally and horizontally. Instead, imagine truth flowing downward (from the general to the particular) and falsity flowing upward (from the particular to the general):

• If 'All *S* are *P*' is true, then it is also true that 'Some *S* are *P*'.

The knowledge that 'All *S* are *P*' is false, however, does not constrain whether or not 'Some *S* are *P*'.

• Similarly, if 'No *S* are *P*' is true, then it is also true that 'Some *S* are not *P*'. The knowledge that 'No *S* are *P*' is false, however, does not constrain whether or not 'Some *S* are not *P*'.

• If 'Some *S* are *P*' is false, then 'All *S* are *P*' must also be false. The knowledge that 'Some *S* are *P*' is true, however, does not indicate whether or not 'All *S* are *P*'.

• Similarly, if 'Some *S* are not *P*' is false, then 'No *S* are *P*' must also be false. The knowledge that 'Some *S* are not *P*' is true, however, does not indicate whether or not 'No *S* are *P*'.

These relationships can be visualized more easily with a square of opposition composed of Venn representations of the four types of statement (Figure 20).

For example, the Venn diagrams demonstrate the incompatible, contradictory nature of diagonal statements such as 'All *S* are *P*' and 'Some *S* are not *P*'.

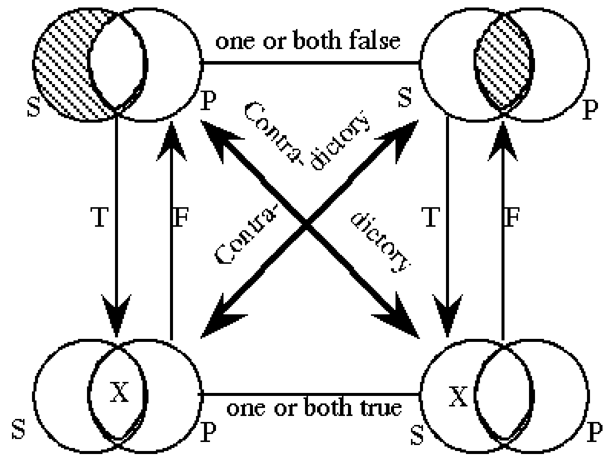

Figure 20. Venn square of opposition.

Table 8 summarizes the relationships that can be determined between any two of the classification statements by examination of the square of opposition.

Table 8. Relationships among classification statements.

		All S are P	No S are P	Some S are P	Some S are not P
If 'All S are P' true,	then		false	true	false
If 'All S are P' false,	then		unknown	unknown	true
If 'No S are P' true,	then	false		false	true
If 'No S are P' false,	then	unknown		true	unknown
If 'Some S are P' true,	then	unknown	false		unknown
If 'Some S are P' false,	then	false	true		true
If 'Some S are not P' true,	then	false	unknown	unknown	
If 'Some S are not P' false,	then	true	false	true	

Finally and most simply (for me at least), one can immediately see the impact of any one statement's truth value on the other three statements through substitution. Again I substitute Scientist for S, and either People, Physicists, or Politicians for P, whichever fits the first statement correctly. For example, if I assume (correctly) that 'Some scientists are physicists' is true, then 'No scientists are physicists' must be false, and I need additional information to say whether 'All scientists are physicists' or 'Some scientists are not physicists'. Some caution is needed to assure that my conclusions are based on the evidence rather than on my independent knowledge. For example, I know that 'All scientists are physicists' is false but I cannot infer so from the statement above that 'Some scientists are physicists'. As another example, if I assume (naïvely) that 'Some scientists are politicians' is false, then it also must be true that 'No scientists are politicians' and that 'Some scientists are not politicians'. Furthermore, the statement that 'All scientists are politicians' must be false.

*   *   *

Syllogisms

Syllogism is the deductive solution of a pervasive scientific problem: what is the relationship between the two classes *A* and *C*, given that I know the relation of both *A* and *C* to the third class *B*?

Aristotle loved syllogisms. He systematized them, developed rules for and patterns among them, and promoted them as the foremost tool for analysis of arguments. But what is a syllogism? Let us examine the syllogism using Aristotle's own example:

> All men are mortal.
>
> Socrates is a man.
>
> Therefore Socrates is mortal.

This argument is recognizable as a syllogism by these characteristics:

- the argument consists of three statements;
- two of the statements (in this case the first and second) are premises and the third is a conclusion that is claimed to follow from the premises.

In so-called standard form such as the Socrates syllogism, the third statement is the conclusion, containing a subject ('Socrates') and predicate ('mortal'), the first statement is a premise dealing with the predicate, and the second statement is a premise dealing with the subject.

Syllogisms are of three types: categorical, hypothetical, and disjunctive. We will consider hypothetical syllogisms briefly later in this chapter. The Socrates syllogism is categorical: three classification statements, each beginning explicitly or implicitly with one of the three words 'all', 'no', or 'some', with two terms in each statement, and with each term used a total of twice in the argument. Each term must be used in exactly the same sense both times. For example, *man* cannot refer to *mankind* in one use and *males* in the second; this is the fallacy of equivocation, described in a later section.

Chambliss [1954] succinctly comments:

> "The syllogism does not discover truth; it merely clarifies, extends, and gives precision to ideas accepted as true. It is, according to Aristotle, 'a mental process in which certain facts being assumed something else differing from these facts results in virtue of them.'"

Aristotle's description that "something else differing from these facts results" is a bit misleading in its hint of getting something for nothing. The conclusion does not really transcend the premises; instead it is really immanent, an implication of the premises that may or may not be obvious. Rather than discover truth, the syllogism reveals the implications of our assumptions. As such, it is a fundamental step in the hypothetico-deductive method (better known as *the* scientific method).

Syllogisms can be difficult to recognize in everyday language. Formal analysis of syllogistic logic requires a translation from everyday language into the so-called standard syllogism form. This translation may involve reorganizing the statements, recognizing that a term can be much longer than one word, using logical equivalences to reduce terms, supplying an omitted (but implied) premise or conclusion, or breaking apart a compound argument into its component syllogisms. This translation is useful to learn but beyond the scope of this book; the reader is encouraged to consult a textbook on logic and practice translation of the many examples therein. Here we focus on the analysis of standard-form syllogisms, because familiarity with standard-form syllogisms has a fringe benefit: invalid syllogisms will sound dubious and invite closer scrutiny, even if they are couched in everyday language.

* * *

Categorical Syllogisms

Categorical syllogisms have 256 varieties; only 24 are valid. Any one of these 256 can occur in scientific arguments or everyday life, and we should be able to recognize whether it is valid or invalid. Simply but brutally put, we cannot always avoid false assumptions, false inductions, or misleading data, but *we must avoid invalid deductions*. A scientist who incorrectly judges the validity of a syllogism may design and undertake an entire experiment based on a fallacious expectation of its potential meaning.

Table 9: Valid categorical syllogisms [Hurley, 1985].

Unconditionally valid:

All M are P.	All S are M.	∴ All S are P.
No M are P.	All S are M.	∴ No S are P.
All M are P.	Some S are M.	∴ Some S are P.
No M are P.	Some S are M.	∴ Some S are not P.
No P are M.	All S are M.	∴ No S are P.
All P are M.	No S are M.	∴ No S are P.
No P are M.	Some S are M.	∴ Some S are not P.
All P are M.	Some S are not M.	∴ Some S are not P.
Some M are P.	All M are S.	∴ Some S are P.
All M are P.	Some M are S.	∴ Some S are P.
Some M are not P.	All M are S.	∴ Some S are not P.
No M are P.	Some M are S.	∴ Some S are not P.
All P are M.	No M are S.	∴ No S are P.
Some P are M.	All M are S.	∴ Some S are P.
No P are M.	Some M are S.	∴ Some S are not P.

Conditionally valid:

All M are P.	All S are M.	∴ Some S are P.	(S must exist)
No M are P.	All S are M.	∴ Some S are not P.	(S must exist)
All P are M.	No S are M.	∴ Some S are not P.	(S must exist)
No P are M.	All S are M.	∴ Some S are not P.	(S must exist)
All P are M.	No M are S.	∴ Some S are not P.	(S must exist)
All M are P.	All M are S.	∴ Some S are P.	(M must exist)
No M are P.	All M are S.	∴ Some S are not P.	(M must exist)
No P are M.	All M are S.	∴ Some S are not P.	(M must exist)
All P are M.	All M are S.	∴ Some S are P.	(P must exist)

Many strategies *could* be employed to distinguish between valid and invalid categorical syllogisms:

- random choice (not a very scientific basis for decision-making at any time, but particularly when the chance of winning is only 24/256);
- memorization, an old, laborious standby;
- knowing where the answer can be found (Table 9);
- recognition that the correct solutions all obey a few rules (only five rules are needed for successful separation of the 24 valid syllogisms from the 232 invalid ones);
- sketching Venn diagrams;
- substitution, in which we recognize that the problem structure is identical to one whose answer is known.

All except for the 'random choice' option are acceptable solutions to the problem, but memorization and substitution have the strong advantage of much greater speed. In the remainder of this section, I list the valid syllogisms for easy reference, and then I describe substitution -- the easiest closed-book technique for evaluating syllogisms.

* * *

Substitution is an easy way to evaluate categorical syllogisms. As with the evaluation of any formal logic, the validity of the form is independent of the actual terms used. If we insert familiar terms into the syllogism, choosing ones that yield true premises, then an untrue conclusion *must* indicate an invalid syllogism. For evaluation of categorical syllogisms, I select substitutions from the following classification tree:

```
                 animals
                 /     \
          mammals       reptiles
          /     \       /      \
    dogs        cats  snakes  turtles
```

The danger of substitution is that a true conclusion does not prove that the logic is valid, as we saw above for the syllogism "Some mammals are dogs; some mammals are cats; therefore no cats are dogs." Substitution can prove that an argument is invalid but, unfortunately, cannot prove that it is valid. If the premises are true, a substitution that yields a true conclusion may or may not be of valid form. In contrast, a substitution with true premises and false conclusion must be of invalid form. Thus one needs to consider several substitutions, to see whether any case can prove invalidity. For example, the following argument is not disproved by the first substitution but is disproved by the second one:

Some physicists are theoreticians.
Some astronomers are theoreticians.
Therefore some physicists are astronomers.

Some dogs are animals.
Some mammals are animals.
Therefore some dogs are mammals.

Some dogs are mammals.
Some cats are mammals.
Therefore some dogs are cats.

Usually, an invalid syllogism couched in familiar terms feels wrong, even if the conclusion is true. Further brief thought then generates a variant that proves its invalidity. Using the 'animal tree' to test syllogisms can generally avoid the juxtaposition of invalid logic and true conclusion: simply confine each statement to adjacent levels in the animal tree, rather than creating statements like 'some dogs are animals' that skip a level.

*   *   *

<u>Hypothetical Syllogisms</u>

Like categorical syllogisms, hypothetical syllogisms consist of two premises and a conclusion. Unlike categorical syllogisms, one or both of the premises in a hypothetical syllogism is a conditional statement: 'if A, then B'.

We can express a conditional, or if/then, statement symbolically as A⇒B. The statement A⇒B can be read as 'A implies B' or as 'if A, then B'; the two are logically equivalent. Both statements state that A is a necessary and sufficient condition for B.

If both premises in a hypothetical syllogism are if/then statements, then only three forms of syllogism are possible:

Valid	Invalid	Invalid
S⇒M.	S⇒M.	M⇒S.
M⇒P.	P⇒M.	M⇒P.
∴ S⇒P.	∴ S⇒P.	∴ S⇒P.

Another type of hypothetical syllogism has one if/then statement, a statement that one of the two conditions is present or absent, and a conclusion about whether the other condition is present or absent. Symbolically, we can indicate presence (or truth) by *S* or *P*, and absence by *-S* or *-P*. If only one premise is an if/then statement, two valid and two invalid forms of syllogism are possible:

Valid	Invalid	Invalid	Valid
S⇒P	S⇒P	S⇒P	S⇒P
S	-S	P	-P
∴ P	∴ -P	∴ S	∴ -S

As with categorical syllogisms, hypothetical syllogisms are readily testable through substitution. The substitution that I use treats *if/then* as a mnemonic for 'if the hen':

A: if the hen lays an egg;
B: we cook omelettes;
C: we eat omelettes.

This substitution readily distinguishes invalid from valid hypothetical syllogisms:

Valid: A⇒B. If the hen lays an egg, then we cook omelettes.
B⇒C. If we cook omelettes, then we eat omelettes.
∴ A⇒C. Therefore, if the hen lays an egg, we eat omelettes.

Invalid: A⇒B. If the hen lays an egg, then we cook omelettes.
C⇒B. If we eat omelettes, then we cook omelettes.
∴ A⇒C. Therefore, if the hen lays an egg, we eat omelettes (invalid; eating omelettes is not necessarily related to the hen's laying).

Invalid: B⇒A. If we cook omelettes, then the hen lays an egg.
B⇒C. If we cook omelettes, then we eat omelettes.
∴ A⇒C. Therefore, if the hen lays an egg, we eat omelettes (invalid, not because the first premise is absurd but because the hen's laying and our omelette eating are not necessarily related).

Valid: A⇒B. If the hen lays an egg, then we cook omelettes.
A. The hen laid an egg.
∴ B. Therefore, we cook omelettes.

Valid: A⇒B. If the hen lays an egg, then we cook omelettes.
-B. We are not cooking omelettes.
∴ -A. Therefore, the hen did not lay an egg.

Invalid: A⇒B. If the hen lays an egg, then we cook omelettes.
-A. The hen did not lay an egg.
∴ B. Therefore, we are not cooking omelettes. (invalid; maybe we can get eggs elsewhere)

Invalid: A⇒B. If the hen lays an egg, then we cook omelettes.
B. We are cooking omelettes.
∴ A. Therefore, the hen laid an egg. (invalid; maybe we can get eggs elsewhere)

The last two fallacies above are so obviously wrong that we might dismiss them as irrelevant to scientists. When couched in technical terms, however, these invalid syllogisms do appear occasionally in print. Both fallacies imply confusion between necessary and sufficient conditions. Both are deductively invalid, but they may have some inductive validity:

Valid: If the hen lays an egg, then we cook omelettes.
The hen did not lay an egg.
Therefore, we *may* not cook omelettes.
(the hen's failure is a setback to our omelette plans, but maybe we can get eggs elsewhere)

Valid: If the hen lays an egg, then we cook omelettes.
We are cooking omelettes.
Therefore, the hen *may* have laid an egg. (true, but maybe we got eggs elsewhere)

This second hypothetical syllogism is a cornerstone of **scientific induction**: "If hypothesis (*H*) entails Evidence (*E*), and *E* is true, then *H* is probably true." It is fallacious to conclude that H is definitely true, but the evidence is relevant to evaluation of the hypothesis.

* * *

Pitfalls: Fallacious Arguments

After a bit of practice, one can readily recognize syllogistic arguments that are expressed in ordinary language, and one can evaluate them by examining their structures. Many arguments can appear to be structurally valid and yet be fallacious; such arguments yield a false conclusion even if the premises are true. These fallacies exhibit an error in execution, such as subtle problems in their premises, use of apparently relevant but logically irrelevant evidence, an incorrect connection of premises to conclusion, and grammatical errors or ambiguities. Many of these fallacies are genuine pitfalls to scientists. Most are deductive pitfalls, but a couple of inductive pitfalls (e.g., hasty generalization) are included here because of their similarity to deductive pitfalls.

The list of fallacies that follows is loosely based on the compilation of Hurley [1985]. Other logicians lump or split these fallacies differently and describe them with different jargon. For our purposes, the names applied to these fallacies have limited usefulness; instead, our goal is to recognize when an argument is fallacious. Practice with a variety of examples is the key, and logic textbooks have a wealth of examples.

Most fallacies fall into one of four types: problems in a premise, extraneous extra evidence, faulty link between premises and conclusion, or case-dependent relationship between parts and whole. Table 10 gives an overview of these different kinds of fallacy, and the remainder of this chapter examines these fallacies in more detail.

* * *

Table 10. Varieties of fallacious argument.

Problems in a premise:

Fallacy	Premises	other 'evidence'	⇒	Conclusion
false dichotomy	2 choices assumed	other choices omitted		
suppressed evidence	weakness ignored			
ambiguity	ambiguity		misinterpreted	
false cause	noncausal, yet assumed causal			
slippery slope	unlikely chain of events		flawed links	

Extraneous other evidence:

Fallacy	Premises	other 'evidence'	⇒	Conclusion
appeal to authority		experts say . . .		
personal attack		fools say . . .		
mob appeal		rest of group says . . .		
might makes right		accept or suffer consequences		
extenuating circumstances		extenuating circumstances		
red herring		smoke-screen distraction		

Faulty link between premises and conclusion:

Fallacy	Premises	other 'evidence'	⇒	Conclusion
missing the point	imply conclusion A			conclusion B drawn
overinterpreting	uncertain			definite
begging the question #1		dubious premise ignored		
begging the question #2	validated by conclusion		circular reasoning	validated by premises
equivocation	one meaning for key word			another meaning for same word
straw man			tested with bad example	

Case-dependent relationship between parts and whole:

Fallacy	Premises	other 'evidence'	⇒	Conclusion
false extrapolation to whole	parts		attribute misapplied	whole
false extrapolation to parts	whole		attribute misapplied	part
false extrapolation to individual	general		attribute misapplied	individual
hasty generalization	nonrepresentative individual		generalized	general

* * *

Fallacies Resulting from Problems in a Premise

> For scientists, few 'victimless' crimes are as outrageous as the burning of the Alexandria library, and with it the destruction of so much ancient knowledge and culture. One legend is that when the Muslim Amrou Ibn el-Ass captured Alexandria, he sought his caliph's guidance on the fate of the library. Caliph Omar responded that the library's books are either inconsistent or consistent with the Koran. If inconsistent, they are heretical; if consistent, they are redundant. In either case they should be burned. [Gould, 1990]
>
> The story is apocryphal and, I suspect, wrong. The library was probably destroyed in 389 A.D., not 642 A.D., and the Muslims embraced other cultures and their science at a time when Christians were suppressing them. As a memorable example of false dichotomy, however, the story is unsurpassed.

A valid deduction does not imply a correct conclusion; accurate premises or assumptions are also essential. When reading a research paper, the scientist must seek and evaluate the premises. Incorrect or overlooked premises are probably the dominant source of incorrect scientific deductions, and these errors can take several forms:

• **False dichotomy** is an incorrectly exclusive 'either . . .or. . .' statement in one of the premises. When one choice is eliminated by another premise, the other choice is accepted incorrectly as the conclusion. The logic is valid, and if there truly are only two choices then the conclusion is valid:

"Either you subscribe to the journal or you don't. Your subscription lapsed, and therefore you don't subscribe to the journal."

The fallacy of false dichotomy is that the either/or premise is false if more than two choices exist. Therefore the conclusion is invalid:

"Either the hypothesis is proved or disproved. This experiment did not prove the hypothesis. Therefore it must have disproved it." Unfortunately, science is almost always less efficient than this. Experiments may support hypotheses, refute them, or disprove them, but never prove them.

False dichotomy is frequent among the general public.

Sometimes one premise and the conclusion are obvious and unstated:

"Either make at least 100 measurements or skip the experiment entirely." The premises (*P*) and conclusion (*C*) are: *P1*: the experiment is worthless if <100 measurements are made; *P2*: surely you want the experiment to be worthwhile; and *C*: therefore you will want to do at least 100 measurements.

• **Suppressed evidence** is the omission of evidence that weakens or fatally undermines one premise. This fallacy is frequent among both lay people and scientists. Few scientists deliberately hide an assumption. Instead, they may suppress evidence passively, by an unconscious 'forgetting' or by a conscious decision that the evidence is too flawed to warrant mention. A different, but related, lapse of objectivity is the ignoring of evidence that leads to a competing conclusion.

• **Ambiguity** creates a fallacious argument, when misinterpretation of an ambiguous premise results in a wrong conclusion. Usually the ambiguity arises from punctuation or grammar and is merely a temporary distraction while reading a publication:

"We analyzed our experiments on monkeys using multivariate statistics." Smart monkeys!

Misinterpretation of someone else's ambiguously stated premise is more serious. People often are unaware of ambiguities in their own statements, because of familiarity with the subject. Others then misinterpret the statement, leading them to incorporate it into an argument that is doomed by the incorrect premise.

A sign on a beach says, "Sharks! No swimming!" [Ennis, 1969]

My colleagues and I have often succumbed to the fallacy of ambiguity in interpreting telexes. The sender cannot foresee the ambiguity that cost-saving brevity has introduced. For example: ". . . STOP MISS YOU STOP LOVE END"

• **False cause** is an argument in which a relationship is incorrectly assumed to be causal. Several types of associations can be misinterpreted as causal: (1) one event may precede another and become misidentified as its cause; (2) the cause may be confused with the effect if the two are nearly simultaneous; (3) a variable may control two others and thereby give those two an indirect association; and (4) the apparent association may be coincidental. Determining causality and dodging the potential pitfall of false cause are fundamental aspects of science. They are discussed in more detail in Chapter 3.

• **Slippery slope** is an argument in which the premises form a chain reaction of assumed causal consequences, beginning with some initial event and culminating with a conclusion. One step onto a slippery slope causes one to slide all the way to an undesirable outcome. The arguer's purpose is usually to prevent that first step. The slippery-slope fallacy is the invalid assumption that a full chain reaction invariably follows the initial event. Almost all chain reactions are invalid, because each step requires a causality that is both necessary and sufficient; only then are alternative paths precluded. Thus chain-reaction arguments are particularly vulnerable to the fallacy of false cause.

Slippery-slope logic is used with mixed success by many fundamentalist preachers. Seldom is it used in science, but sometimes the link between a hypothesis and a testable prediction can involve several steps. If so, one must evaluate whether each step validly involves either pure deduction or a necessary and sufficient causality.

> The most familiar example of a slippery slope, at least to those in my age group, is domino theory. Used successfully in the early justifications of the Vietnam war, domino theory said that if Vietnam were to fall to communism, through chain reaction all of Southeast Asia would eventually become communist. Domino theory was wrong.

> In attempting to refute Galileo's claim that he had discovered satellites of Jupiter, astronomer Francesco Sizi [Holton and Roller, 1958] used a *slippery-slope* argument:
>
> "The satellites are invisible to the naked eye and therefore can have no influence on the earth and therefore would be useless and therefore do not exist."

* * *

Fallacies Employing Extraneous Other Evidence

> When ego is involved, scientific arguments can get personal. This was often the case for Isaac Newton, as the following letter [~1700] illustrates. Note that Newton attempts to demolish an idea without giving a single shred of evidence:
>
> "That gravity should be innate, inherent and essential to matter, so that one body may act upon another at a distance through a *vacuum*, without the mediation of any thing else, by and through which their action and force may be conveyed from one to another, is to me so great an absurdity, that I believe no man who has in philosophical matters a competent faculty of thinking, can ever fall into it."

Unlike Newton's argument, most arguments do involve evidence that can be evaluated in terms of premises and deductive or inductive conclusions. They may also, however, contain a collage of other information that the proponent considers to be relevant but that is extraneous to the core deductive argument. Often this extraneous information is emotionally charged, and the evaluator must cull the deductive argument from among the distractions.

• **Appeal to authority** is the claim that an argument should be accepted because some expert accepts it. Ideally, scientists do not appeal to authority; they evaluate evidence personally. In practice, however, we limit such analyses primarily to our own field, and we tentatively accept the prevailing wisdom of scientists in other fields. The appeal to authority must be considered pragmatically, based on how much more experience the 'authority' has than the arguers have, how mainstream the authority's view is, and how feasible it is for the arguers to evaluate all of the evidence.

> For example, when a biologist considers a physics argument, it is valid to give weight to what physicists believe. Yet when a physicist considers a physics argument, it is a fallacy to accept it merely because some 'great' physicist believes it.

• **Personal attack** is a criticism of the opponent in a debate, rather than refutation of the opponent's arguments. This diversionary attack, like the *red-herring* fallacy discussed later, is a smokescreen that uses emotional impact to draw attention away from the relevant logical arguments. Three types of personal attack are:

• *verbal abuse*, in which one directly attacks the opponent's character, personality, or psychological health, although those factors are irrelevant to the argument being 'refuted'.

> "The so-called discoverers of cold fusion are more interested in glory and a Nobel prize than in well-controlled experiments."

• *challenging the objectivity of the opponent*, in which one argues that the opponents' bias forces them to argue as they do, regardless of the argument's validity.

> "It is not surprising that *A* rejects these experimental data, since they refute his hypothesis."

• *'practice what you preach'*, in which one defends oneself by claiming that the opponent is just as guilty.

> "*A* claims that I have ignored conflicting evidence, but she has ignored . . ."

• **Mob appeal** is the assertion that one should accept an argument in order to join the crowd. Mob appeal is the premise for the emotionally enticing conclusion that 'right thinking' people are a group of winners. Different manifestations of the mob appeal fallacy are:

• *mob psychology*, in which the arguer seeks a simultaneous group response through the bait of inclusion or the threat of exclusion. Politicians and preachers use this technique; scientists do not.

• *bandwagon*, in which it is claimed that the group knows best and an individual is in danger of being left out.

> "Everyone's accepting this new theory and finding applications for their own field."

> "In science the authority embodied in the opinion of thousands is not worth a spark of reason in one man." [Galileo Galilei, 1564-1642]

• *egotistic appeal*, which provides a simple way for an individual to be like someone famous.

> "Nobel prizewinner *A* advocates the hypothesis, so wouldn't you?"

• *status symbol*, in which the individual joins a group composed only of the superior people.

> "Mensa is the most fascinating of clubs, because only those with intelligence in the top 2% can join."

> The last football game that I attended was USC versus UCLA, in about 1969. It was called a 'great game': the sides were evenly matched and the team I favored (USC) came from behind to win in the last couple of minutes. My overriding memory, however, is that the fans on both sides were chanting "Kill! Kill! Kill!" and meaning it. They cheered each time a member of the opposing team was injured or carried from the field, and the game was dirty enough that such incidents were fre-

quent. That day I lost my interest in football, but I gained realization of the need to make my own judgments, rather than accepting mob opinion.

• **Might makes right** is the argument that the listener must accept the arguer's conclusion or suffer the consequences. The threat may be physical or it may involve some other undesirable action. Between parents and children it is usually the former, and among scientists it is the latter. The threat is irrelevant to the validity of the conclusion, and yet it may affect the listener's decision-making.

"Everyone knows that this theory is correct, and if you try to prove otherwise you will destroy your credibility as a scientist."

• **Extenuating circumstances** is the plea for accepting the conclusion out of pity for someone. The arguer claims that acceptance will help someone in trouble, or that rejection of the conclusion will cause undue hardship to someone (usually to the arguer). The extenuating circumstances are irrelevant to the validity of the basic argument.

Students often use the plea of extenuating circumstances on their teachers; lawyers use it on juries and judges. Scientists use it in personal, not scientific, arguments.

• **Red herring** is diversion of attention from an argument's weakness, by creating a distracting smoke-screen. The fallacy is named after a technique used in training hunting dogs: drag a sack of red herring (a strongly scented fish) across the scent trail that the dog is supposed to follow, and train the dog to stick to the main trail without being distracted or diverted to the red-herring trail. The fallacious *red herring* can and usually does consist of valid reasoning, often protracted and sometimes emotional, so the listener is left with the impression that the conclusion is valid. In fact, the red herring is a related issue that is extraneous to the central argument.

This style of misdirection is the secret of many magicians' tricks. It rarely is employed deliberately in scientific arguments. However, a similar smoke-screen – 'straining at a gnat and swallowing a camel' [Matthew 23:24] -- is sometimes adopted: the arguers demolish a minor criticism of their argument, giving the false impression of careful and objective thoroughness, while obscuring a brief mention of a serious weakness in the argument.

When the theory of evolution was proposed by Charles Darwin and advocated by Thomas Huxley, falacious refutations were rampant: Evolution is inconsistent with the <u>Bible</u> (appeal to authority); Darwin is a heretic (personal attack) who should be excluded from our community of scientists (mob appeal) and will certainly burn in Hell (might makes right).

* * *

<u>Faulty Link Between Premises and Conclusion</u>

"The myth of the scientific method keeps the scientific community from recognizing that they must have a humanly developed and enforced professional ethics because there is no impersonal method out there that automatically keeps science the way it ought to be." [Bauer, 1994]

The hardest fallacies to spot are those that lead to a conclusion with which we agree. No scientist would disagree with Bauer's [1994] conclusion that personal ethical responsibility is essential. And yet, his argument is fallacious: we may value the innate checks-and-balances of the scientific method, but no scientist responds by abdicating personal ethics. Between premise and conclusion, the argument has gone astray -- in this case by misrepresenting the relationship between scientist and scientific method. This *straw man* fallacy is one of several ways in which the link between premises and conclusion may be fallacious.

• **Missing the point** is basing a conclusion on an argument in which the premises actually lead to a quite different conclusion. Like *red herring*, much of the argument in *missing the point* is valid but only appears to lead toward the conclusion. *Red herring*, however, is often deliberate whereas *missing the point* is accidental. The fallacy is detectable by deciding what conclusion is actually warranted from the premises, then comparing this valid conclusion to that drawn by the arguer.

> "Hypothesis *A* fails these two tests, and consequently hypothesis *B* is the best explanation."

• **Overinterpreting** is the attempt to claim a firm conclusion although one or more premises is quite uncertain. The fallacy is in asserting a definite, deductive conclusion:

> "Scientists have tried for years to refute this hypothesis and have failed. Thus the hypothesis must be true."

In contrast, the following somewhat similar argument is valid, because it properly considers the evidence as inductive:

> "Many attempts to find N rays have failed, although N rays should be detectable by these tests. Therefore N rays probably do not exist."

• **Begging the question** is an argument in which the logic may be valid, but a dubious premise is either propped up by the conclusion or is ignored entirely. This term is used to describe two different fallacies: ignored dubious premise and circular reasoning.

An *ignored dubious premise*, omitted from but essential to an argument, is a common pitfall. Ignoring a premise is reminiscent of but more extreme than the fallacy of *suppressed evidence*. This fallacy is one of the most serious pitfalls of scientific research, for three reasons. First, everyone is better at noticing characteristics of observed features than at noticing that something is missing. Second, once a premise is overlooked, it will be harder for anyone else to recall. Third, most occurrences of this fallacy could be avoided, if the researcher would just list the premises. Too often, scientists fail to ask themselves what their premises are, or they think about the answers superficially but fail to write them down systematically.

> "You need a different kind of instrument, because the one you have broke down." The dubious premise is that a different type of instrument will not break down.

Circular reasoning is an argument in which the conclusion and premises seem to support each other, but actually they say virtually the same thing in two different ways. The logic is valid if trivial (A, $\therefore A$), yet the repetition lends the illusion of strengthening the conclusion.

> "It is obvious that the instrument is not working reliably because it gives anomalous results. These results must be wrong because the instrument malfunctioned."

• **Equivocation** is use of the same word in subtly different senses; due to ambiguity of word meaning, fallacious logic appears to be structurally valid. Particularly subject to this fallacy are arguments that repeatedly use qualitative descriptions such as large and small, or good and bad:

> "The hypothesis is slightly incorrect, because there is a slight difference between predictions and observations." The first use of 'slight' means 'small but definitely significant,' whereas the second 'slight' may mean 'small and statistically insignificant.'
>
> Proof that -1 is the largest integer [Spencer, 1983]: Listing all the integers . . . -4, -3, -2, -1, 1, 2, 3, 4, . . . where '. . .' extends to infinity, we see that nothing has been omitted. But we also know that the largest integer (*n*) is the only integer for which there is no *n+1*; the only such number is -1.

• **Straw man** is a fallacious strategy for refuting an argument: misinterpret it, refute the misinterpreted version, and then conclude that you have refuted the original argument successfully. The term is a takeoff on the concepts of scarecrows and burning in effigy: imagine that you set up a straw man and easily knock it down, claiming that you have knocked down the real man. The term 'straw man' is sometimes used in a different sense than used here; it can be a 'trial balloon', an idea that is proposed knowing that it will be knocked down but expecting that it will be a productive starting point for further discussions.

Frequently, hypothesis refutations are of the following form: "Let's examine the truth of your hypothesis by seeing how well it fits the following example: . . ." If the argument or hypothesis really should apply to the example, then this technique is compelling. The refutation or confirmation is valid, and any refuted argument must be abandoned or modified. With a *straw man*, in contrast, the original hypothesis was not intended to encompass the example, so the argument is fallacious although the entire logic of the analysis is just as valid. Thus one should evaluate the appropriateness of the example *before* applying it, lest the refutation act as a smoke-screen.

* * *

<u>Case-dependent Relationship Between Parts and Whole</u>

> Cholesterol seems to have surpassed sex as the number-one source of guilt in modern America. Much of this cholesterol consciousness stems from the 1985 National Cholesterol Education Program. All Americans were urged to reduce cholesterol in order to avoid heart disease. Surprisingly, however, there was virtually no direct scientific evidence that cholesterol reduction prevents heart disease in either women or in the elderly, although 75% of heart attacks are in people older than 60 years.
>
> The key studies were all on middle-aged men with high cholesterol. These studies conclusively found that: (1) higher cholesterol level is associated with higher risk of heart disease, and (2) giving cholesterol-lowering drugs to high-cholesterol subjects reduced their risk of heart disease. The first finding established a correlation, and the second result demonstrated causality. Generalization of this pattern to middle-aged women and to the elderly of both sexes is plausible, but neither data nor deduction implies it. [Kolata, 1992b].

The conclusion that everyone should reduce cholesterol is a *hasty generalization*, the extrapolation to an entire population from a possibly nonrepresentative sample. The conclusion *may* be correct -- indeed, it has been demonstrated by subsequent experiments to be correct -- but it does not follow compellingly from *these* data. This inductive fallacy and several deductive fallacies go astray in linking the individual to the general, or parts to the whole, or the universal to the particular. The validity of such arguments is case-dependent: arguments with identical form can be valid or invalid, depending on the specific relationship between parts and whole.

- **False extrapolation to the whole** is the false conclusion that the whole exhibits some characteristic because one or more parts exhibit it. The argument is structurally valid; whether it is correct or not requires careful evaluation of the content, because sometimes extrapolation to the whole is warranted. For example:

> Invalid: "The mistaken belief that technology is applied science . . . implies that any advance in scientific knowledge could be harnessed to useful applications" [Bauer, 1994]. Actually, scientists argue only that *many* scientific advances have valuable practical applications.
>
> Valid: "This prediction of the hypothesis is refuted, and therefore the hypothesis is disproved."
>
> Invalid: "This prediction of the hypothesis is confirmed, and therefore the hypothesis is proved."
>
> Valid: "This premise in the argument is false; thus the argument is false."
>
> Invalid: "Every premise in the argument is true; thus the argument is true." Remember that the truth of a conclusion depends both on the truth of premises and on the validity of the logic.

- **False extrapolation to parts** is the false conclusion that a part exhibits some characteristic because the whole exhibits it. This potential fallacy is the reverse of the previous one. The conclusion may be either correct or incorrect depending on the content:

> Valid: "The argument is correct (valid and true), so every premise must be true."
>
> Invalid: "The argument is incorrect (either invalid or untrue), so every premise must be false."
>
> Valid: "This journal requires peer review for its papers; therefore this article in the journal has been peer reviewed."
>
> Invalid: "That scientific meeting is worth attending, and consequently every talk at the meeting is worth attending."

- **False extrapolation to the individual** is misapplication of a generalization to an individual case. This fallacy is the reverse of *hasty generalization*, and it is somewhat similar to the fallacy of *false extrapolation to parts*. Like these, it may be correct or incorrect depending on the content. The fallacy lies in ignoring evidence that the general rule is inappropriate to this specific case.

> "Publication is an essential part of any research project. Therefore my manuscript should not be refused publication even if the reviews were negative."

- **Hasty generalization** is the inductive extrapolation to an entire population, based on a sample that is nonrepresentative. Often the sample is too small to be representative, but smallness does not necessarily imply the fallacy of *hasty generalization*. A sample of only two or three can be enough if one is dealing with a uniform and representative property: for example, learning not to touch fire. *Hasty generalization* is frequent among non-scientists; it is the origin of many superstitions. A key difference between scientific and popular induction is that the latter usually ignores the need for a representative sample. The consequence is vulnerability to *hasty generalization*.

Hasty generalization is quite similar to the fallacy of *false extrapolation to the whole*. The two differ, however, in the scopes of their conclusions: a general statement about every member of a population (*hasty generalization*) or the collective behavior of a class (*false extrapolation to the whole*).

> Hasty generalization: “Wristwatches with radium dials are safe, so all radium samples are safe.”
>
> False extrapolation to the whole: “The half-life of a radium-226 atom is 1622 years; thus brief exposure to radium poses negligible hazard.”
>
> Which is this? “I seldom detect the effect, so the effect must be rare.”

* * *

H. H. Bauer, a chemist and self-proclaimed expert in STS (science, technology, and society), succeeded in packing a remarkable number of the foregoing fallacy types into a single paragraph:

> “In what sense, then, are the social sciences actually science? They have no unifying paradigm or the intellectual consensus that goes with it. They have not produced distinctive and reliable knowledge that is respected or valued by human society as a whole. Yet those are the very qualities for which the natural sciences are noted and respected; they are the qualities that we associate with something being scientific -- that is, authoritatively trustworthy. The social sciences are simply not, in the accepted meaning of the term, scientific. And that conclusion has been reached by at least a few practicing social scientists -- for example, Ernest Gellner.” [Bauer, 1994]

Bauer’s argument succumbs to at least seven fallacies:

- *Suppressed evidence*: His assertion that none of the social sciences has a paradigm is incorrect.
- *Suppressed evidence*: To claim that the social sciences have not produced reliable knowledge, one must ignore countless concepts such as supply and demand (economics), stimulus and response (psychology), human impacts of environmental change (geography), and human impacts of racial and gender stereotypes (sociology).
- *False dichotomy*: He assumes that the accumulated knowledge of all of the social sciences can be classified as either having or not having a consensus. In actuality, consensus is a continuum and degree of consensus varies tremendously both within and among fields.
- *Mob appeal*: He claims that ‘human society as a whole’ determines the reliability of knowledge.
- *Straw-man*: His definition of ‘scientific’ as ‘authoritatively trustworthy’ is not merely a weak assumption; it is a deliberate misrepresentation.
- *Appeal to authority*: He seeks validation for his stance by quoting one social scientist.
- *False extrapolation to the whole*: He applies a conclusion based mainly on sociology to all social sciences.