A Quick Introduction to Logic

Scott Lehmann

This module was written for use by students in a variety of introductory philosophy courses at The University of Connecticut and by those in more advanced courses who wish to review some of this material. It may also be useful to students in other disciplines. I anticipate that it will be revised from time to time, so please send comments or suggestions to me at lehmann@uconnvm.uconn.edu. (August 2001)

Contents

1. Arguments 1
    Reasoning, premises, conclusion
2. Evaluating arguments 1
    Deductive validity, inductive strength
3. Uses of reasoning 3
    i. Extending belief
    ii. Getting at the truth
    iii. Persuading
    iv. Explaining
4. Evaluation standards compared 6
5. Direct and indirect proofs 8
6. Argument forms 12
7. Extracting arguments from text 14
    Examples, exercises (p.22)
8. Testing hypotheses 23
    Falsification, confirmation, crucial experiments

A QUICK INTRODUCTION TO LOGIC

1. ARGUMENTS. Logic is the study of reasoning. When we reason, we draw conclusions from assumptions. Consider a simple example from Arthur Conan-Doyle's story, "Silver Blaze":

> ... he must have gone to King's Pyland or to Mapleton. He is not at King's Pyland. Therefore, he is at Mapleton.

Here the fictional detective Sherlock Holmes is concluding ("therefore") that Silver Blaze, a missing race horse, is at Mapleton from the assumptions that he went to one of two stables, King's Pyland or Mapleton, and that he is not at King's Pyland.

Holmes' reasoning constitutes an **argument** for the **conclusion** that Silver Blaze is at Mapleton from the assumptions that he went to King's Pyland or to Mapleton and that he is not at King's Pyland. A standard convention for representing arguments is to arrange the assumptions in a vertical list, draw a line (representing "therefore") below them, and write the conclusion underneath it. In such **standard form** Holmes' argument is then:

P1. Silver Blaze went to King's Pyland or to Mapleton.
P2. He is not at King's Pyland.

C. He is at Mapleton.

The assumptions of an argument are termed its **premises** (sometimes written with a double "s" as "premiss" and "premisses"). Here "P1" is short for "premise 1", "P2" for "premise 2", and "C" for "conclusion".

2. EVALUATING REASONING. Has Holmes reasoned well? Is his conclusion well-drawn from his assumptions? To answer this question, we must first settle on a standard of good reasoning.

Deductive logic uses the standard of deductive validity: an argument is **deductively valid** (or simply **valid**) provided its conclusion would *have to* be true if its premises were true. Inductive logic uses the weaker standard of inductive strength: an argument is **inductively strong** (or simply **strong**) provided its conclusion would *probably* be true if its premises were true.

Holmes' argument is not valid. To test for validity, we assume that P1 and P2 are true: that Silver Blaze went to King's Pyland or to Mapleton, and that he is not at King's Pyland. Then we ask: under these assumptions, does C have to be true? The answer is "No": perhaps Silver Blaze is no longer where he originally went. We can see the problem more clearly if we consider a different argument of the same general pattern or form:

P1'. Neil Armstrong went to the moon or to Mars.

P2'. He is not on the moon.

C'. He is on Mars.

Neil Armstrong did go to the moon in 1970, so P1' is true; he is not there now, so P2' is also true; but he is not now on Mars either, so C' is false.

We can make Holmes' argument valid by adding an additional premise (making an extra assumption) to the effect that Silver Blaze is still wherever he originally went:

P1. Silver Blaze went to King's Pyland or to Mapleton.
P2. He is not at King's Pyland.
P3. He is still wherever he went.

C. He is at Mapleton.

With the addition of P3, the argument is valid: if P1, P2, and P3 were true, C would have to be true.

Some arguments that are not (deductively) valid are (inductively) strong, but this standard is more difficult to apply. In the case of Holmes' argument, we really have no idea how likely it is that Silver Blaze is at Mapleton, given that he went to King's Pyland or to Mapleton and isn't at King's Pyland. This is because we have no idea how likely it is that he is still wherever he went, that is, how likely it is that P3 is true. Consequently, we do not know whether Holmes' argument is strong.

In some cases we can apply this standard. Suppose I reason that today is Monday because I teach today and I teach on MWF only. In standard form, my argument is

I teach today.
I teach on MWF only.

Today is Monday.

This argument is neither valid nor strong. It is not valid because if the premises were true, today must be Monday, Wednesday, or Friday, but it need not be Monday. It is not strong because the probability that today is Monday, given that it is Monday, Wednesday, or Friday, is only 1/3 or 33%; it is more likely – 2/3 or 67% – that today is Wednesday or Friday. We can make the argument strong by weakening its conclusion to "Today is Monday or Wednesday":

I teach today.
I teach on MWF only.

Today is Monday or Wednesday.

Here the likelihood that the conclusion is true, given that the premises are true, is 2/3 or 67%, so this argument is strong. More interesting and important applications of the standard of inductive strength will be considered later.

3. USES OF REASONING. We have given two standards for evaluating reasoning, that is, for deciding when a conclusion is well-drawn from premises. The stronger standard is deductive validity: an argument is (deductively) valid just in case its conclusion would *have to* be true if its premises were true. The weaker standard is inductive strength: an argument is (inductively) strong just in case its conclusion would *probably* be true if its premises were true.

We can see that these standards are appropriate if we consider the uses of reasoning. Its most basic use is to **discover what would – or would probably – be the case in a hypothetical situation**. For example, we might try to get evidence for or against a scientific hypothesis (such as the theory of evolution) by asking what we would expect to observe under certain conditions on the assumption that the hypothesis is true. An argument representing this sort of reasoning can be regarded as claiming that its conclusion holds – or probably holds – in the hypothetical situation described by its premises. The standards of deductive validity and inductive strength are just restatements of these claims.

Other uses of reasoning – to extend belief, to get at the truth, to persuade, to explain – build on this basic use.

i. Extending belief. When reasoning is used to extend belief, I begin with assumptions that I believe to be true and ask what else I should believe: what conclusion can I draw from these assumptions? If such an argument is valid, then I should believe its conclusion is true; if its premises are correct, then its conclusion will also be correct. If such an argument is strong, then I should believe its conclusion is probably true; if its premises are correct, then its conclusion will probably be correct as well.

For example, Sherlock Holmes is using reasoning to extend belief. He believes that his assumptions P1 and P2 are correct, and from them draws the conclusion that the missing horse is at Mapleton. Should he believe this? Unfortunately, as we have seen, Holmes' argument is not valid, and we cannot determine whether it is strong. So Holmes' beliefs that P1 and P2 are correct do not justify his believing that C is correct. However, if Holmes also believes that Silver Blaze is still wherever he originally went, then he should believe that he is at Mapleton, because adding P3 as a premise to his argument makes it valid.

ii. Getting at the truth. Reasoning to get at the truth is just like reasoning to extend belief, except that we begin with premises that are not just believed to be true but are in fact true. Believing that p is true can differ from p's being true because – unfortunately! – we all believe things that are false, though we do not know what they are. If an argument from true premises is valid, then its conclusion must also be true. Valid arguments with true premises are said to be **sound**.

In practice, using reasoning to get at the truth does not differ from using reasoning to extend belief. This is because when we use reasoning to extend belief, we assume that our beliefs are correct. If we come to regard some claim that we have accepted (such as: Santa Claus really exists) as false, we cease to believe it.

Mathematical proofs employ reasoning to get at the truth. For example, it is easy to prove that the base angles α and β of an isosceles triangle ABC are equal (in measure), if we appeal to the fact that the corresponding angles of triangles whose corresponding sides are equal (in length) are equal (in measure). The standard proof proceeds by drawing a line from the vertex C to the midpoint D of the base AB and then observing that the corresponding sides of the resulting (right) triangles ACD and BCD are equal in length, so the corresponding angles, in particular α and β, must be equal in measure:

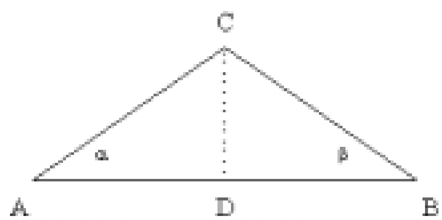

We can represent the assumptions of this proof as the premises of a valid argument for the conclusion that α and β are equal (in measure):

P1. ABC is an isosceles triangle in which side AC and side BC are equal (in length).
P2. D is the midpoint of side AB.
P3. If the corresponding sides of triangles are equal (in length), their corresponding angles are equal (in measure).
P4. In triangles ACD and BCD, α and β are corresponding angles.

C. α and β are equal (in measure).

This argument is valid, but it represents only the assumptions and conclusion of the proof and not the proof itself. The **proof** is the reasoning used to get from the assumptions to the conclusion. We can represent it by expanding the argument into a **chain** of simpler arguments, the conclusion of one link functioning as a premise in the next. See section 5 on proofs below.

When we use reasoning to get at the truth, it is essential to begin with assumptions that are true. A valid argument can be constructed for any conclusion you like, but such ‘proofs’ are worthless unless their premises are true. For example,

George W. Bush is a member of the Green Party.
Each member of the Green Party favors campaign finance reform.

George W. Bush favors campaign finance reform.

is valid, but it does not establish that Mr. Bush favors campaign finance reform, since the first premise (and maybe the second as well) is false.

Note also that if we start from false premises, we needn't end up with a false conclusion. We can see this by replacing "favors campaign finance reform" with "is male". The conclusion of this argument is certainly true. However, the argument proceeds from false premises and therefore does not give us reason to believe that its conclusion is true.

iii. Persuading. To persuade you to accept some claim C, I get you to extend *your* beliefs to include C. To do so, I may try to construct an argument from premises that you accept to the conclusion C. If my argument is deductively valid or inductively strong, then you ought to accept C, because to accept a claim is to believe that it is true. The premises you audience accepts describe a hypothetical situation in which the conclusion of such an argument holds or probably holds; thus, you ought to accept this conclusion. For example, if Anne believes that suicide is wrong because human lives are God's property, I might try to persuade her that executions are also wrong:

> Human lives are God's property, not ours.
> It's wrong to destroy what is not yours.
> Execution destroys a human life.
> ____________________________________
> Execution is wrong.

The first two premises reflect Anne's beliefs; the third is an obvious truth that Anne should accept. The argument is valid: if the premises were true, the conclusion would have to be true as well. So Anne should accept the conclusion.

I need not accept your beliefs in order to persuade you by reasoning in this way – after all, I am trying to extend *your* beliefs, not *mine*! I may, for example, reject the first premise of the argument I have given to Anne, and hold that my body is my property, to do with as I wish. Those who use reasoning to persuade are often more interested in getting people to agree on a position or course of action than they are in getting them to agree on the reasons for accepting it. Indeed, it is often foolish to insist on unanimity of reasons. Legislation, for example, is generally enacted because various interest groups agree that it is desirable, though their reasons for thinking it is desirable may be quite different. The American Civil Liberties Union (ACLU) will be happy to join Catholic Bishops in opposing capital punishment; however, the reasons for their opposition differ and may place these groups in opposition on other issues, such as separation of church and state.

iv. Explaining. Explanations answer questions of the general form "Why X?", where we do not doubt that X is the case: Why is the sky blue? Why did the Soviet Union collapse? Why didn't you buy potatoes yesterday? Why did I get a C+ instead of a B-? When we offer explanations, we generally give an argument for X whose premises are claims about the situation that we believe to be true and which together make X predictable. To make X predictable, the argument for X from these claims must be deductively valid or inductively strong.

For example, perhaps the explanation for your not buying potatoes yesterday is that you just forgot. We may put this in the form of a valid argument as follows:

> You forgot to buy potatoes yesterday.
> One doesn't do what one forgets to do.
> ___
> You didn't buy potatoes yesterday.

Or perhaps the explanation is that the potatoes at Ralph's Pretty Good Grocery were moldy and you didn't have time to look elsewhere. Again, we could cast this in the form of a valid argument:

> The potatoes you inspected at Ralph's yesterday were moldy.
> You never buy moldy potatoes.
> You didn't have time to look elsewhere.
> ___
> You didn't buy potatoes yesterday.

In either case, if I believe the premises, I am justified in concluding that you didn't buy potatoes yesterday.

Which explanation is correct depends upon which set of premises is true. In this case, we would generally interpret the explanations as incompatible, so that at most one of them is correct. However, many instances of human action seem to involve a mixture of motives that are often very difficult to sort out. Suppose a local retailer buys an ad in a community orchestra's concert program. Is it because she thinks the ad will reach an important group of potential customers, or believes it will be good public relations for her business, or wants to support the arts, or didn't want to refuse the little violin student who asked her to buy the ad? All of these factors and more could have influenced her decision, and it is probably a mistake to believe that one of them is its 'true' cause. Indeed, it is hard to make sense of such a claim in this kind of case.

4. EVALUATION STANDARDS COMPARED. Recall that if the premises of a *deductively valid* argument were true the conclusion would have to be true, and that a *sound* argument is one that is deductively valid and has true premises. From these definitions we can derive some important truths that set deductively valid arguments apart from inductively strong arguments:

i. The conclusion of a sound argument is true. (Reason: a sound argument is (a) deductively valid and (b) has true premises; by (a), its conclusion must be true if its premises are true; but, by (b), its premises are true.)

It follows that it is impossible to construct a sound argument for a conclusion C *and* a sound argument for its negation not-C. For if we could do this, both C and not-C would have to be true by (ii), but this is impossible: if C is true, not-C is false, and if not-C is true, C is false. Therefore, if John argues that abortion is wrong, and Mary argues that it is not, at least one of these arguments is unsound (both may be unsound).

By contrast, inductively strong arguments with true premises need not have true conclusions.

For example,

> 1000 raffle tickets were sold.
> Since I really wanted the prize, I bought 990 of them.
> ———————————————————
> I'll win it.

is very strong (the probability that the conclusion is true, given that the premises are true is 99%), and we can imagine situations in which the premises are true. But, alas, we can also imagine all too clearly that in some of these situations the conclusion is false.

ii. If an argument is deductively valid, so is any argument that results from it by adding additional premises. (Reason: if the premises of the expanded argument were true, the premises of the original argument would be true; but this argument is valid by assumption, so its conclusion – which is also the conclusion of the expanded argument – would be true; hence, the expanded argument is valid.)

This is not true of inductively strong arguments. While

A. I teach today.
I teach on MWF only.
———————————
Today is Monday or Wednesday.

is strong,

B. Yesterday was Thursday.
I teach today.
I teach on MWF only.
———————————
Today is Monday or Wednesday.

is not: the additional information that yesterday was Thursday (and, of course, the background assumption that Friday follows Thursday) makes the conclusion unlikely (in fact, it drops the probability from 2/3 to 0).

We can appeal to this difference to see that

C. I teach today.
I teach on MWF only.
———————————
Today is *probably* Monday or Wednesday.

is not deductively valid, contrary to what we might think. If it were valid, it would remain so if "Yesterday was Thursday" were added as a premise. But under these assumptions, "Today is probably Monday or Wednesday" would be false, not true. Argument C is best regarded as claiming, in a somewhat misleading way, that argument A is inductively strong. (Similarly, "I teach MWF only and I teach today, so today *has to be* M, W, or F" is best regarded as asserting that "I teach MWF only and I teach today, so today is M, W, or F" is deductively valid. Can you see why?)

5. DIRECT AND INDIRECT PROOFS. A **proof** of some claim C shows that C follows from accepted assumptions or postulates P. That is, a proof shows that if the assumptions P were true, C would also be true – or, in other words, that the argument with premises P and conclusion C is deductively valid. Roughly speaking, there are two kinds of proofs: direct and indirect. A direct proof reasons from premises to conclusion in small, valid steps; an indirect proof shows (often by direct proof) that a contradiction follows validly from the premises and the assumption that the conclusion is false. These methods rest on two simple consequences of the definition of deductive validity.

i. If the argument from premises P to conclusion C is valid and the argument from premises C and P' to conclusion C' is valid, then the argument from premises P and P' to conclusion C' is valid. (Reason: since the first argument is valid, C must be true under assumptions P; so by (ii), C must be true under assumptions P and P'; but since the second argument is valid, C' must be true under assumptions C and P', so C' must be true under assumptions P and P')

This fact is the basis for direct proofs. In a **direct proof**, the reasoning from P to C is divided into relatively small valid steps, in which the conclusion of a one step functions as a premise in a subsequent step.

The proof sketched in section 2.ii that the base angles of an isosceles triangle are equal in measure is an example of such a proof. Recall the situation of the diagram.

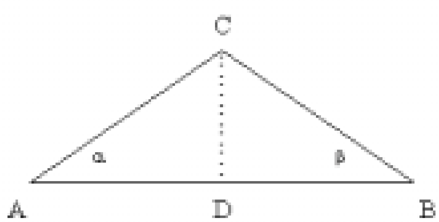

We wish to show that the following argument is valid:

P1. ABC is an isosceles triangle in which side AC and side BC are equal (in length).
P2. D is the midpoint of side AB.
P3. If the corresponding sides of triangles are equal (in length), their corresponding angles are equal (in measure).
P4. In triangles ACD and BCD, α and β are corresponding angles.

C. α and β are equal (in measure)

This we can do by 'chaining together' simpler valid arguments that take us from the premises P1-P4 to the conclusion C:

P2. D is the midpoint of side AB.

C1. AD and DB are equal (in length).

C1. AD and DB are equal (in length).
P1. ABC is an isosceles triangle in which side AC and side BC are equal (in length).

C2. The corresponding sides of ACD and BCD are equal (in length).

C2. The corresponding sides of ACD and BCD are equal (in length).
P3. If the corresponding sides of triangles are equal (in length), their corresponding angles are equal (in measure).

C3. The corresponding angles of ACD and BCD are equal (in measure).

C3. The corresponding angles of ACD and BCD are equal (in measure).
P4. In triangles ACD and BCD, α and β are corresponding angles.

C. α and β are equal (in measure).

Each link in the chain is valid, so by **i** the argument from P1-P4 to C is valid.

We can represent this direct proof as a sequence of claims, starting with the premises P1-P4 and ending with C, such that each claim in the sequence is either a premise or is obtained from earlier claims by a simple valid argument:

1. ABC is an isosceles triangle in which side AC and side BC are equal (in length). {Premise P1}
2. D is the midpoint of side AB. {Premise P2}
3. If the corresponding sides of triangles are equal (in length), their corresponding angles are equal (in measure). {Premise P3}
4. In triangles ACD and BCD, α and β are corresponding angles. {Premise P4}
5. AD and DB are equal (in length). {C1, from 2}
6. The corresponding sides of ACD and BCD are equal (in length). {C2, from 1 & 5}
7. The corresponding angles of ACD and BCD are equal (in measure). {C3, from 3 & 6}
8. α and β are equal (in measure). {Conclusion C, from 4 & 7}

This direct proof establishes that the conclusion C is true provided the premises P1-P4 are true. In this case, P1-P4 are true. P1 and P4 are true by stipulation. An isosceles triangle is one in which two sides are equal in length, so we may identify the equal sides as AC and BC in the diagram; this is P1. Corresponding angles are just those enclosed by corresponding sides, and which sides correspond to which is a matter of stipulation: in this case, we stipulate that AD corresponds to BD and that AC corresponds to BC. Thus we have P4. P2 is true because every line segment has a midpoint, so we may identify the midpoint of AB as D. P3 is true because triangles whose corresponding sides are equal in length can be made to coincide by some combination of translating, rotating, and flipping over one of the triangles.

Flipping suggests an even simpler proof: flip triangle ABC over to obtain triangle BAC, observe that the corresponding sides of ABC and BAC are equal (in length), and conclude that their corresponding angles – including α and β – must therefore be equal (in measure). This shows that we can throw away P2 and still conclude C validly from P1, P3 and P4′, as explained by the following chain:

> P1. ABC is an isosceles triangle in which side AC and side BC are equal (in length).
>
> ---
>
> C2′. The corresponding sides of ABC and BAC are equal (in length).
>
> C2′. The corresponding sides of ABC and BAC are equal (in length).
> P3. If the corresponding sides of triangles are equal (in length), their corresponding angles are equal (in measure).
>
> ---
>
> C3′. The corresponding angles of ABC and BAC are equal (in measure).
>
> C3′. The corresponding angles of ABC and BAC are equal (in measure).
> P4′. In triangles ABC and BAC, α and β are corresponding angles.
>
> ---
>
> C. α and β are equal (in measure).

Exercise: Represent *this* direct proof as a sequence of numbered claims starting with P1, P3, and P4′ and ending with C.

ii. If the conclusion of a valid argument is false, at least one of its premises must also be false. (Reason: if they were all true, the conclusion would have to be true, but it isn't).

This fact is the basis for indirect proofs in mathematics. An **indirect proof** that claim C follows from assumptions P proceeds by showing that a contradiction – a claim of the form "p and not-p" – follows from P and not-C. For suppose this is so. A contradiction must be false, so by **ii** at least one of not-C and the assumptions P must be false. Hence, if the assumptions P were all true, not-C would have to be false and C would have to be true. But then C follows from P: the argument from premises P to conclusion C is valid.

A relatively simple example of an indirect proof is the standard argument that $\sqrt{2}$, the square root of 2, is irrational (cannot be expressed as a fraction). Suppose, on the contrary, that $\sqrt{2}$ is rational. Then we can express $\sqrt{2}$ as a fraction in lowest terms: say, $\sqrt{2} = m/n$, where m and n have no common factor. Then $2n^2 = m^2$ and 2 divides m^2. But 2 is prime, so 2 must also divide m. So we have m = 2p, say, and therefore $n^2 = 2p^2$. By the same reasoning, 2 divides n. But then m and n have a common factor (namely, 2), contrary to supposition. (For Aristotle's slightly different formulation of this famous proof, see Morris Kline, *Mathematical Thought from Ancient to Modern Times*, I, p.33)

In this proof, the supposition that √2 is rational leads validly to the contradiction that m and n both have and lack a common factor. That is, when the assumption that √2 is rational is added to truths about the numbers, we can validly draw the absurd conclusion that certain numbers both have and lack a common factor. The valid argument is:

P1. √2 is rational.
P2. If √2 is rational, then there are numbers m and n such that (a) m and n have no common factor and (b) √2 = m/n.
P3. For any numbers m and n: if √2 = m/n, then $2n^2 = m^2$.
P4. For any numbers m and n: if $2n^2 = m^2$, then 2 divides m.
P5. For any number m: if 2 divides m, then there is a number p such that m = 2p.
P6. For any numbers n and p: if $2n^2 = (2p)^2$, then $n^2 = 2p^2$.
P7. For any numbers m and n: if 2 divides both m and n, then m and n have a common factor.

C. There are numbers that both have and lack a common factor.

The proof that the contradiction C follows validly from P1-P7 is a direct proof and can be represented as above by a sequence of claims:

1. √2 is rational. {P1, the claim we are interested in refuting}.
2. If √2 is rational, then there are numbers m and n such that (a) m and n have no common factor and (b) √2 = m/n. {P2}
3. For any numbers m and n: if √2 = m/n, then $2n^2 = m^2$. {P3}
4. For any numbers m and n: if $2n^2 = m^2$, then 2 divides m. {P4}
5. For any number m: if 2 divides m, then there is a number p such that m = 2p. {P5}
6. For any numbers n and p: if $2n^2 = (2p)^2$, then $n^2 = 2p^2$. {P6}
7. For any numbers m and n: if 2 divides both m and n, then m and n have a common factor. {P7}
8. √2 = m/n, where m and n have no common factor. {From 1 and 2}
9. $2n^2 = m^2$. {From 3 and 8}
10. 2 divides m. {From 4 and 9}
11. m = 2p. {From 5 and 10}
12. $2n^2 = (2p)^2$. {From 9 and 11}
13. $n^2 = 2p^2$. {From 6 and 12}
14. 2 divides n. {From 4 and 13}
15. 2 divides m and n, though m and n have no factor in common. {From 8, 10, and 14}
16. There are numbers that both have and lack a common factor. {C, from 15}

P1 is the claim we wish to deny. P2 - P7 are truths about the numbers. P2(b) and P7 are true in virtue of the meaning of "rational" and "common factor". P4 is true in virtue of the fundamental theorem of arithmetic: each number, m^2 included, can be expressed as the product of primes in one and only one way. This theorem follows validly from more basic truths about the numbers, and so is true. The other premises (P1(a), P3, P5 - P7) also follow validly from more basic truths about the numbers and so are true. The conclusion C cannot be true, yet it is validly drawn from P1 - P7. So if P2 - P2 are true, P1 must be false. Thus, $\sqrt{2}$ is irrational.

6. ARGUMENT FORMS. In showing that Holmes' argument about Silver Blaze is not deductively valid as it stands, we exhibited another argument of the same form with true premises and a false conclusion. This method relies on an analysis of the notion of validity – in particular, the meaning of "have to" – that is beyond the scope of this introduction. Moreover, discerning the forms of arguments requires both learning a theory of forms and considerable practice in applying it.

Here we shall only note that there are argument forms that are valid in the sense that every argument of such a form is valid. If you can learn to discern such forms, you can quickly see that certain arguments are valid. Here are three particularly useful ones, with examples:

Modus ponens:	If p, then q	If it's July, it's summer.
	p	It's July.
	—————	—————————————
	q	It's summer.
Modus tollens:	If p, then q	If it's July, it's summer.
	Not-q	It's not summer.
	—————	—————————————
	Not-p	It's not July.
Disjunctive syllogism:	p or q	The solution is odd or even.
	Not-p	It's not odd.
	—————	—————————————
	q	It's even.

We can also identify some forms that are strong in the same sense: every argument of such a form is strong. The form

All As are Bs
a is an A
—————————
a is a B

is valid. This is really a special case – "all" meaning 100% of – of the form

r% of As are Bs
a is an A
—————————

a is a B

This form is strong when r is greater than 50. It is sometimes called the **statistical syllogism.**

A very important type of argument moves from what is true of a sample to what is true of a population. Political polling organizations, for example, attempt to judge public opinion by determining what a sample of the public thinks. The general pattern or form of this reasoning is

r% of the As in sample S are Bs

Approximately r% of As are Bs

Reasoning of this form is strong when the sample S has been obtained in a way that makes it likely to be representative of the population of As. Random samples, in which the members of S are chosen at random from the population of As, meet this condition. In fact, if we stipulate that "approximately r%" means r±e% for some desired margin of error e, and want the probability that the conclusion is true, given that the premises are true, to be p, then we can *calculate* (from r, e, and p) how big the random sample S has to be to warrant concluding, with probability p, that r±e% of As are Bs.

For example, if r = 50, e = 3, and p = 95%, then sample S must contain about 1000 individuals. If we hold p and the sample size constant, then the margin of error e shrinks as r decreases or increases from 50. This means that if we wish to right about 95% of the time in judging that r±3% of As are Bs on the basis of observing that r% of a random sample S of As are Bs, then S must contain about 1000 As. Suppose a political campaign wishes to gauge support for its candidate C; it might poll 1000 likely voters to ask for whom they would vote if the election were held today. If 450 name candidate C, 400 name her opponent O, and the remaining 150 are undecided, then the campaign might judge that the actual support among likely voters is within 3% of 45% for C, 40% for O, and 15% undecided. This means the race is too close to call: support for C could be as low as 42% and support for O as high as 43%, and we don't know how the undecided voters will split.

The conclusion may also be incorrect. Under the best of circumstances, it will be incorrect about 5% of the time (1 in 20 times). And circumstances may not be the best. The sample may not be random: it is not clear how to identify likely voters, and even if those polled can be randomly selected, some will refuse to answer. Moreover, those who do answer may lie. An exit poll of voters in the 1989 Virginia gubernatorial election predicted that the Democratic candidate, Douglas Wilder, was headed for a 55% win over his Republican opponent. Mr. Wilder did win, but by only 6000 votes (50.01%). The most likely explanation for the disparity is that some voters did not wish to admit that they had voted against Mr. Wilder, who was black.

Exercises:

1. For each of the following conditions, give an argument of the form "All As are Bs, a is an A, therefore a is a B" that satisfies the condition, or explain why this is not possible.

a. The premises are both true, and so is the conclusion.
b. The premises are both true, but the conclusion is false.
c. The premises are both false, and so is the conclusion.
d. The premises are both false, but the conclusion is true.

2. Answer questions a-e about the argument below *in order!*

P1. If this argument is valid, it does not have true premises and a false conclusion.
P2. This argument has true premises and a false conclusion.

C. This argument is not valid

a. Explain why this argument fits the form (say what p and q are in this case):

If p, then not-q
q

---

not-p

b. Explain why this form is valid.
c. Is C true? Why or why not?
d. Is P1 true? Why or why not?
e. Is P2 true? Why or why not?

7. EXTRACTING ARGUMENTS FROM TEXT. To evaluate an argument, we must first know what it is: what conclusion is being drawn from what assumptions? Standard form conveys this information at a glance, but arguments generally do not appear in standard form. They are buried in text and must be dug out.

The job of extracting arguments from text is eased if the text contains **inference indicators**. These are words or phrases that typically point to conclusions or premises. Conclusion indicators include “therefore”, “so”, “thus”, “accordingly”, “hence”, and “it follows that”. Premise indicators include “since”, “because”, “for”, “on account of”, and “the reason is that”.

Let us consider some examples:

Example 1

> Theologians and other moralists have said homosexual acts violate the ‘natural law,’ and that they are therefore immoral and ought to be prohibited by the state. (Burton Leiser, “Is homosexuality unnatural?”, in Rachels, ed., *The right thing to do*, p.165)

In this passage, Leiser is reporting an argument given by “theologians and other moralists”. Its conclusion, indicated by “therefore”, is that homosexual acts are immoral and ought to be prohibited by the state. Only one reason for accepting this claim is given, namely, that homosexual acts violate the ‘natural law’; this is the single premise of the argument. In standard form, then, the argument is:

> P. Homosexual acts violate the ‘natural law’.
> ___
> C. Homosexual acts are immoral and ought to be prohibited by the state.

This argument is not valid as it stands. We can make it valid by adding two premises:

> Whatever violates the ‘natural law’ is immoral.
> Whatever is immoral ought to be prohibited by the state.
> Homosexual acts violate the ‘natural law’.
> ___
> Homosexual acts are immoral and ought to be prohibited by the state.

If the premises of the expanded argument were true, its conclusion would also be true. So if you believe all three premises, you should also believe that homosexual acts are immoral and ought to be prohibited by the state. Whether you should believe the premises is another matter: they are all dubious. For discussion of the first and third, see Leiser (who argues that there is no way to understand ‘natural law’ so that both the first and third premise are true).

Example 2

> Hence, even the fool is convinced that something exists in the understanding, at least, than which nothing greater can be conceived. For, when he hears of this, he understands it. And whatever is understood exists in the understanding. (St. Anselm, *Proslogium*, II)

This passage contains part of a challenging argument for the existence of God, given by St. Anselm in the 12th century. The claim that Anselm is arguing for in this part of the argument is indicated by “hence”: even the fools among us must agree that (C) that, than which nothing greater can be conceived, exists in the understanding. This is the conclusion. The reasons for accepting (C) are indicated by “for”: they are that (P1) we understand the phrase ‘that, than which nothing greater can be conceived’ and (P2) whatever is described by a phrase that we understand exists in the understanding (that is, we can think about it). These are the premises of the argument. In standard form, then, the argument is:

> P1. We understand the phrase ‘that, than which nothing greater can be conceived’
> P2. Whatever is described by a phrase that we understand exists in the understanding.
> ___
> C. That, than which nothing greater can be conceived, exists in the understanding.

This argument is valid: if P1 and P2 were true, C would also be true. Anselm goes on to argue that that, than which nothing greater can be conceived, cannot exist only in the understanding (but must also exist in reality), because what exists in reality is greater than what exists only in the understanding. For discussion, see William L. Rowe, “The ontological argument”, in Feinberg and Shafer-Landau, eds., *Reason and responsibility* (Wadsworth, 1999), pp.8-18.

Example 3

Q: What animal can jump higher than a house?
A: Any animal that can jump at all can jump higher than a house, because houses can't jump.

Q: Which is better, eternal happiness or a ham sandwich?
A: A ham sandwich!! Why? Because nothing is better than eternal happiness, but a ham sandwich is better than nothing.

These jokes depend on **ambiguity**: the fact that certain words and phrases can be used in quite different ways. We are puzzled by the first question because we interpret it as asking "What animal can jump higher than a house *is high*?" whereas the answer interprets it as "What animal can jump higher than a house *can jump*? The second joke exploits different senses of "nothing": the first "nothing" means "there is no thing", while the second means "having nothing at all to eat".

In each case, the argument is presented in the answer: the conclusion (the claim we are supposed to accept) is stated first, and the premises (the reasons for accepting it) are indicated by "because". In standard form, these arguments are:

Houses can't jump.

Any animal that can jump can jump higher than a house.

Nothing is better than eternal happiness.
A ham sandwich is better than nothing.

A ham sandwich is better than eternal happiness.

The first argument is not valid as it stands but can be made so by adding a true premise:

Whatever can jump can jump higher than anything that can't jump.
Houses can't jump.

Any animal that can jump can jump higher than a house.

This argument is sound.

The second argument appears to be valid in virtue of having the form

B is better than C
A is better than B

A is better than C

However, the premises seem true, while the conclusion is ridiculous. So, contrary to appearance, this argument is not valid. The problem is that "nothing" does not mean the same thing in the second premise as it does in the first. The form of the argument is more accurately represented as

> There is no thing that is better than C
> A is better than B
> ___
> A is better than C

The argument of the second joke illustrates what logicians call a **fallacy of ambiguity**: the conclusion follows validly from the premises when ambiguous words are read in one way, but making the premises true requires reading them in another way. Such arguments are unsound (**exercise:** explain why).

Example 4

> If multi-cellular life exists elsewhere in the universe, the earth would by now probably have been colonized by an advanced alien civilization. Since this has not occurred, it is probable that multi-cellular life exists nowhere else in the universe. (Suggested in Ian Crawford, "Where are they?", *Scientific American*, July 2000)

The conclusion of this text seems to be that it's probable that multi-cellular life exists nowhere else in the universe. The reasons for believing this are given by the first sentence and the clause that follows "since". In standard form, the argument is:

> If multi-cellular life exists elsewhere in the universe, the earth would by now probably have been colonized by an advanced alien civilization.
> The earth has not yet been colonized by an advanced alien civilization.
> ___
> Multi-cellular live probably exists nowhere else in the universe.

Is this argument valid? It appears to have the form

> If p, then probably q
> Not-q
> ___
> Probably, not-p

In this case, p is "multi-cellular life exists elsewhere in the universe" and q is "the earth has been colonized by an advanced alien civilization". Some arguments of this form have true premises and a false conclusion; for example,

> If Al Gore is an earthly organism, he's probably a bacterium.
> Al Gore is not a bacterium.

Al Gore is probably not an earthly organism.

Therefore, the form is not valid, and the argument is not valid in virtue of having this form.

This way of looking at the argument takes "probably" in "Multi-cellular life probably exists nowhere else in the universe" as part of the conclusion. However, if we recall the discussion of argument C in Section 4.ii, we see that this is not the only option. Suppose we regard "probably" as weakening the connection between the premises and conclusion from "therefore" to "therefore, probably", and ask about the inductive strength of:

If p, then probably q.
Not-q

Not-p

The first premise claims that Pr(q/p), the probability of q given p, is high. To ask whether the argument is strong is to ask whether the claim that Pr(not-p/not-q), the probability of not-p given not-q, is high follows validly from the assumption that Pr(q/p) is high together with the laws of probability. It does not. Suppose, for example, that we are concerned with randomly selecting marbles from a bag: there are 100 marbles in all, 90 are glass, 70 of these are red, and there are no other red marbles in the bag. Pr(red/glass) = 7/9, which is high, but Pr(not-glass/not-red) = 10/30 = 1/3, which is low.

It is possible to express Pr(not-p/not-q) in terms of Pr(q/p), Pr(p), and Pr(not-q/not-p). However, this does not help much in the present case. Perhaps it could be argued that Pr(not-q/not-p) -- the probability that the earth has *not* been colonized by an advanced alien civilization, given that multi-cellular life exists nowhere else in the universe -- is 1 or close to it. But we do not know Pr(p), the likelihood is that multi-cellular life exists elsewhere in the universe.

The remaining examples are best read as containing more than one argument. Typically, as in a direct mathematical proof, the conclusion of one argument functions as a premise in another argument. Sometimes such 'chained' arguments are called **complex arguments**.

Example 5

> ... during the same days we observed this tailless bird, a chickadee lacking a tail also appeared at the feeder beside the library in the village, two miles farther south along the ridge. Was it the same bird? The chances seem overwhelmingly in favor of that assumption. Inasmuch as we have never seen another tailless chickadee in a decade and a half at Trail Wood, or in our whole lives elsewhere, the coincidence of two chickadees losing their tails at the same time, in the same locality, appears too great to be accepted. (Edwin Way Teale, *A Naturalist Buys an Old Fa*rm, Ch.29)

Here Teale is arguing that the tailless chickadee he and his wife saw at Trail Wood was the same one seen at the library two miles away at about the same time. The reason he gives us for accepting this conclusion is that it is very unlikely that two chickadees will lose their tails at the same time in the same locality. And the reason to believe this is indicated by “inasmuch”: we’ve never seen another tailless chickadee in 15 years at Trail Wood or ever anywhere else. This is evidence that such chickadees are rare. We can lay out Teale’s reasoning in a sequence of three arguments in standard form, the conclusion of one functioning as a premise in the next:

> We’ve never seen another tailless chickadee in 15 years at Trail Wood or ever anywhere else.
> ___
> Tailless chickadees are rare.

> Tailless chickadees are rare.
> ___
> It is very unlikely that two chickadees will lose their tails at the same time in the same locality.

> It is very unlikely that two chickadees will lose their tails at the same time in the same locality.
> ___
> The tailless chickadee we saw at Trail Wood was the same one seen at the library.

Are any of these arguments valid? strong?

Example 6

> ... for a short sale to be legal under securities regulations, an investor must either borrow the shares before selling them or already own them. The borrowing of shares can only be done in a margin account, which Mr. Morse did not have. So, in order to short stocks in his cash accounts at the nine brokerage firms, Mr. Morse should already have owned the shares. But he did not own them, according to the S.E.C. complaint. (“S.E.C. Says an Investor Fleeced Stockbrokers,” *The New York Times*, 28 Feb 95, D5)

This passage explains the Securities and Exchange Commission’s (S.E.C.) case against Morse. The S.E.C. claims that Morse’s short sales were illegal, though this claim is not stated in the passage. Most of the passage is an argument for a premise of this argument, indicated by “so”: if Morse’s short sales were legal, he owned the shares he sold. Here is the argument for the premise in standard form:

> If short sales of stock are legal, one must own the shares or borrow them.
> If one borrows shares of stock, one must have a margin account.
> Morse did not have a margin account.
> ___

If Morse's short sales were legal, Morse owned the shares.

The conclusion of this argument is a premise in the argument for the S.E.C.'s contention that Morse's short sales were illegal:

If Morse's short sales were legal, Morse owned the shares.
Morse did not own the shares.

Morse's short sales were illegal.

Both of these arguments are valid. Therefore, Morse's short sales were in fact illegal provided the premises of the first argument and the second premise of the second argument are true.

Example 7

> "Well, I'll eat it, and if it makes me grow larger, I can reach the key; and if it makes me grow smaller, I can creep under the door; so either way I'll get into the garden, and I don't care which happens!" (Alice speaking in Lewis Carroll, *Alice's Adventures in Wonderland*, Ch.1)

Here Alice is persuading herself to eat a cake that says "Eat Me". She wants to unlock a door into a lovely garden, but drinking from a bottle marked "Drink Me" has shrunk her to about ten inches in height and she cannot reach the key. Her reasoning seems to be:

If I eat it, it'll make me grow larger or smaller.
If I grow larger, I can reach the key and unlock the garden door.
If I grow smaller, I can crawl under the garden door.
If I can unlock or crawl under the door, I'll get into the garden.

If I eat it, I'll get into the garden, though I'll be larger or smaller.

If I eat it, I'll get into the garden, though I'll be larger or smaller.
I want to get into the garden, and I don't care whether I'm larger or smaller.

I ought to eat it.

The "so" in the passage points to the conclusion of the first of these arguments. It is valid. The second argument, whose conclusion is unstated in the passage, is not valid: we need some premise to the effect that I ought to do what will get me what I want. The problem is formulating such a connection in a way that makes it plausible. Sometimes what I want is not good for me (cigarettes, for example), and even when it is, some ways of getting it may be 'costlier' than others.

Example 8

> **“**Let ABΓ be a triangle, and through A let ΔE be drawn parallel to BΓ. Now since BΓ, ΔE are parallel, and the alternate angles are equal, the angle ΔAB is equal to the angle ABΓ, and EAΓ is equal to AΓB. Let BAΓ be added to both. Then the angles ΔAB, BAΓ, ΓAE, that is, the angles ΔAB, BAE, that is, two right angles, are equal to the three angles of the triangle. Therefore the three angles of the triangle are equal to two right angles.” (Proclus, on Euclid, in Newman, ed., *The world of mathematics*, I, p.190)

This is the classic argument that the (degree) measures of the angles of a Euclidean triangle sum to 180. The situation is diagrammed below; to simplify notation, angles ABΓ, AΓB, AΓB have been labeled β, α, γ, and the angles ΔAB, BAE labeled β′, γ′.

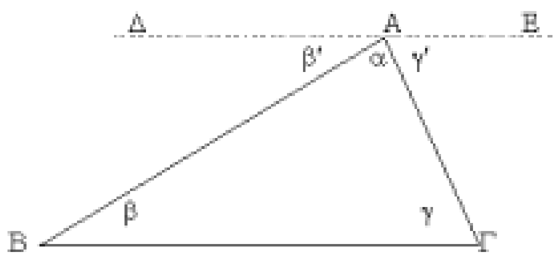

The conclusion of the argument, given in the last sentence, is that the degree measures of β, α, and γ sum to 180 (“the three angles of the triangle are equal to two right angles”). It follows from two assumptions: (1) the degree measures of β′, α, and γ′ sum to 180 (“the angles ΔAB, BAΓ, ΓAE, that is, the angles ΔAB, BAE, that is, two right angles”) and (2) the degree measures of β and β′ are equal, and the degree measures of γ and γ′ are equal (“the angle ΔAB is equal to the angle ABΓ, and EAΓ is equal to AΓB”). The main argument here can be expressed as follows, where m(x) is the (degree) measure of angle x:

P1. $m(\beta') + (m(\alpha) + m(\gamma')) = 180$
P2. $m(\beta') = m(\beta)$, and $m(\gamma') = m(\gamma)$

C. $m(\beta) + (m(\alpha) + m(\gamma)) = 180$

P1 claims that the angles β′, α, γ′ taken together amount to “two right angles.” P2 claims that the “alternate angles” β, β′ and γ, γ′ are “equal”. This argument is valid. Proclus’ text suggests the following direct proof:

1. $m(\beta') + (m(\alpha) + m(\gamma')) = 180$. {P1}
2. $m(\beta') = m(\beta)$, and $m(\gamma') = m(\gamma)$. {P2}
3. When equals are added to equals, the results are equal.
4. $m(\alpha) + m(\gamma') = m(\alpha) + m(\gamma)$ {From 2 and 3}

5. $m(\beta') + (m(\alpha) + m(\gamma')) = m(\beta) + (m(\alpha) + m(\gamma))$ {From 2, 3, and 4}
6. $m(\beta) + (m(\alpha) + m(\gamma)) = 180$ {From 1 and 5}

The premises P1 and P2 follow from facts about Euclidean geometry. The text does not really justify P1, though the move from "angles ΔAB, BAΓ, ΓAE" to "angles ΔAB, BAE" to "two right angles" suggests the following line of argument:

a1. Any point on a line is the vertex of two straight angles, whose rays are the half-lines determined by the point and whose measures are 180.
b1. If an angle is divided by a ray, its measure is the sum of the measures of the resulting angles.
c1. A is a point on line ΔE.
d1. A straight angle of ΔE whose vertex is A is divided by a ray though B into angle β' and an angle BAE, which in turn is divided by a ray through Γ into angles α and γ'.

P1. $m(\beta') + (m(\alpha) + m(\gamma')) = 180$.

Premises a1 and b1 are general truths about Euclidean figures, while c1 and d1 specify details of the situation depicted in the figure. The argument is valid (**exercise:** give a direct proof that shows this), so P1 is true.

The reasoning for P2 is not crystal clear in the text: "the alternate angles are equal" looks like a premise within the scope of "since", though it is really a consequence of "BΓ, ΔE are parallel". The essential premises seem to be a2-d2:

a2. If parallel lines are intersected by a line, alternate interior angles are equal in measure.
b2. ΔE is parallel to BΓ.
c2. Line AB intersects ΔE and BΓ, and β and β' are alternate interior angles.
d2. Line AΓ intersects ΔE and BΓ, and γ and γ' are alternate interior angles.

P2. $m(\beta') = m(\beta)$, and $m(\gamma) = m(\gamma')$.

Here a2 is a general truth about Euclidean figures, while b2, c2, and d2 report particulars of the general situation depicted in the figure. The argument is valid (**exercise:** give a direct proof that shows this), so P2 is true.

Exercises: Put the arguments in standard form and then determine whether they are valid.

1. If I work hard in this course, I'll get a good grade; so if I don't work hard, I won't get a good grade.

2. I ought to work hard in this course, because if I do, I'll get a good grade and I ought to do whatever will get me a good grade.

3. The Bible says that scripture is the word of God, so scripture is the word of God.

4. Scripture is the word of God because that's what the Bible says and whatever the Bible says is so.

5. Whatever the Bible says is so, for the Bible is scripture and the Bible says that scripture is the word of God and whatever God says is so.

6. Nabisco 100% Bran cereal is flavored with two naturally sweet fruit juices, Ann is eating cereal that contains sugar, so Ann isn't eating Nabisco 100% Bran cereal.

7. Every man who didn't dance with Ann danced with Mary's mom, but Ann is Mary's mom, so every man danced with Ann.

8. If infanticide is wrong, so is abortion; but abortion isn't wrong, so neither is infanticide.

9. If abortion is wrong because it destroys a potential person, so is contraception. So if abortion is wrong, it cannot be for this reason.

10. If rights were conferred by the Creator (as the Declaration of Independence alleges), then identical twins could have different rights. But this is absurd, so rights are not conferred by the Creator.

11. Abortion does not violate any rights the fetus may have, even the right to life. For A's rights are not violated if B refuses to let A use something that B owns. But women own their bodies, and aborting an unwanted fetus is refusing it the use of one's body.

12. If people have a right to life, then they also have a right to take their lives. Having a right to something entitles one to do as one wishes with it, provided others' rights are not violated, and taking one's own life violates nobody else's rights.

13. From a bumper sticker: "Jesus Christ or Antichrist: you must choose one." Since Sedat is a devout Muslim, he has not chosen Jesus Christ. Therefore, Sedat has chosen the Antichrist.

8. TESTING HYPOTHESES. Consider the simple argument:

1 If Alice has avian flu, then Alice has a high fever.
2 Alice does not have a high fever.

---

3 Alice does not have avian flu.

This argument is valid: if its premises were true, then its conclusion would have to be true as well. (For suppose, on the contrary, that **1** and **2** are true, while **3** is false. Since **3** is false, Alice *does* have avian flu. Thus, by **1**, Alice has a high fever. But this is impossible, by **2**. So the supposition that **1** and **2** can be true while **3** is false is incorrect.)

If we let “H” represent “Alice has avian flu” and “O” represent “Alice has a high fever”, then the argument fits the pattern (has the form) of *modus tollens:*

MT

If H, then O
not-O

---

not-H

Any argument that fits this pattern is valid. (This can be seen by generalizing the explanation of the argument’s validity. Suppose that both “If H, then O” and “not-O” are true, while “not-H” is false. Since “not-H” is false, H is true. Thus, since “If H, then O” is true, O is also true. But this is impossible if “not-O” is true. So the supposition that the premises can be true while the conclusion is false is incorrect.)

Consider now the argument:

1 If Alice has avian flu, then Alice has a high fever.
2' Alice has a high fever.

3' Alice has avian flu.

This argument is not valid. You can have a high fever without having avian flu, even if **1'** is true: lots of other diseases cause high fevers. This argument fits the pattern (has the form):

AC

If H, then O
O

H

Some arguments that fit this pattern are valid (example: “If not all primes are odd, then some even number is prime, some even number is indeed prime, so not all primes are odd”). But most of the ones you can think of are not. Because not every argument of form AC is valid, this pattern is **fallacious**. It’s labelled “AC” because logicians know it as the fallacy of **affirming the consequent.**

The reason for using “H” and “O” in MT and AC is that the reasoning involved in testing scientific hypotheses roughly follows these simple patterns. Such hypotheses generally make claims about what we cannot directly observe. Like “Alice has avian flu”, they are true or false, but we cannot tell which just by looking. Instead, we must reason *from* the hypothesis H *to* something O that we can observe: we must ask what we’d expect to observe if the hypothesis were true. This reasoning, if deductively valid, gives us the true premise “If H, then O”. Then we proceed to the test: either we see O or we fail to see O. If we see O, then H **passes** the test; if we fail to see O, then H **fails** the test.

How does passing or failing reflect on H? If we fail to see O, then “not-O” is true, so “not-H” is true according to the valid reasoning of MT; in this case, the test **falsifies** H. Unfortunately, the other possible outcome – seeing O – does not prove that H is true, since reasoning of form AC is invalid. Sometimes it is said that O **confirms** H, but this may suggest that the case that O makes for H is stronger than it really is. After all, without more information, we have no idea how likely H is to be true, given that we have seen O.

Moreover, MT is a little too simple to represent real cases of falsification. To reason from H to O (and thus to claim "if H, then O"), we must generally appeal to additional **auxiliary assumptions** A and **conditions** C: we expect O under conditions C, provided that H and A. In the present case, O might be a thermometer reading of at least 102ºF. So C would include a description of taking Alice's temperature, and A would include claims like "avian flu produces a high fever" and "the thermometer is accurate".

If we include auxiliary assumptions A and conditions C, then "If H, then O" becomes "If H and A, then O under C", and the valid reasoning pattern for falsification becomes:

If A and H, then under C, O
C but not-O
____
not-H or not-A

This means that if we fail to see O under conditions C, we cannot be sure that H is incorrect: perhaps one or more of the auxiliary assumptions A is in error. Suppose, in the present case, that the thermometer reads only 99ºF, whereas O is "The thermometer reads at least 102ºF." Should we conclude that Alice does not have avian flu? To do that, we must exclude the other possibilities: that avian flu does not produce a high fever in *all* cases, that this thermometer is inaccurate, that Alice's temperature was taken incorrectly, etc.

The upshot is that neither passing nor failing a test seems conclusive: a hypothesis that passes can be false, and a hypothesis that fails can be true. Assuming the auxiliary assumptions are correct, passing tells us that what we observe is *consistent with* the hypothesis H. But it does not establish that H is true, since this observation will also be consistent with other hypotheses. A high fever, for example, is a symptom of many diseases, not just avian flu. Failing tells us that what we observe is *not consistent with* the hypothesis *and* the extra assumptions we must make to predict O. So either H or at least one of those extra assumptions must be false. But we cannot conclude that it must be H that is in error.

This may look bad for science, but it merely shows is that we cannot be *certain* about the truth or falsity of scientific hypotheses. It does not show that we have *no reason* to believe that some are true and others are false.

Karl Popper believed that we can often arrange things so that failing a test – failing to see O – gives us reason to disbelieve H. This is a matter of appealing to assumptions A that are probably true and conditions C that are probably realized, so that we can move from "If A and H, then O under C" to "If H, then probably O". He also suggests (e.g., in *Conjectures and refutations*, I) that passing a test – seeing O – gives us reason to believe H only when O is something we would not expect unless H were true. If the falsity of H makes O unlikely, then seeing O gives us reason to believe H. However, if O was something we would have expected even without H (if the falsity of H does not make O unlikely), the fact that we see O gives us no reason to believe H. That would certainly be the case here if we already knew that Alice had a high fever. But even if we did not know this, we know that lots of other diseases cause high fevers, so it is false that Alice is unlikely to have a high fever if she doesn't have avian flu.

Ronald Giere ("Testing theoretical hypotheses") has developed this suggestion into the notion of a good test for H. Consider an analogous problem. Suppose we want to develop a lab test for HIV. In addition to being inexpensive, such a test should be *reliable* in the sense of generating few **false negatives** or **false positives**. A false negative here is a test result indicating you are not infected with HIV when in fact you are infected. A false positive is a test result indicating that you are infected when in fact you are not. Both are bad, for obvious reasons.

Analogously, a **good test** for a hypothesis H is one that is unlikely to generate a false negative (H fails, although it is true) or a false positive (H passes, although it is false). So we may characterize a good test is a procedure for generating two possible outcomes, O or not-O, such that both of the following conditions hold: (i) if H, then probably O, and (ii) if not-H, then probably not-O. (i) tells us that false negatives (seeing not-O when H is true) are unlikely; (ii) tells us that false positives (seeing O when H is false) are unlikely. Suppose we accept H if H passes the test (if we see O) and reject H if H fails the test (if we see not-O). Then (i) assures us that we are unlikely to accept H if it is false, while (ii) assures us that we are unlikely to reject H if it is true. In this way, a good test is revealing. The challenge is to devise good tests.

Crucial experiments – where they can be devised – provide good tests of competing hypotheses H1 and H2. Suppose there are conditions C under which if H1 is true, we should see O1, whereas if H2 is true, we should see O2, where O1 and O2 are different. If we have narrowed our hypotheses to H1 and H2, we can find out which is true by setting up conditions C: if we see O1, then H1 is true, whereas if we see O2, then H2 is true. That is because both of the following arguments are deductively valid:

P1. H1 or H2
P2. If H1, then O1; if H2, then O2
P3. If O1, then not O2
P4. O1
____
C. H1

P1. H1 or H2
P2. If H1, then O1; if H2, then O2
P3. If O1, then not O2
P4′. O2
____
C′. H2

A famous example from the history of science concerns the transmission of light. Newton thought that light was transmitted as a stream of particles (call this hypothesis "H1"). Huygens believed that it was transmitted as a wave (call this hypothesis "H2"). For a long time, neither hypothesis enjoyed a clear advantage over the other in accounting for observations. But early in the 19th century, Poisson designed an experiment to decide between them.

Suppose that an opaque circular disk is placed between a light and a screen. If H1 were true, we would expect the disk to cast a circular shadow on the screen; this is O1. This can be depicted in the schematic side-on view below.

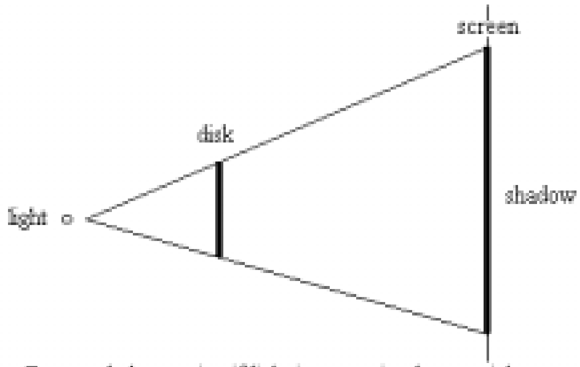

Expected observation if light is transmitted as particles

However, if H2 were true and the dimensions are just right, Poisson calculated that we should see a bright spot at the center of the circular shadow! This is O2:

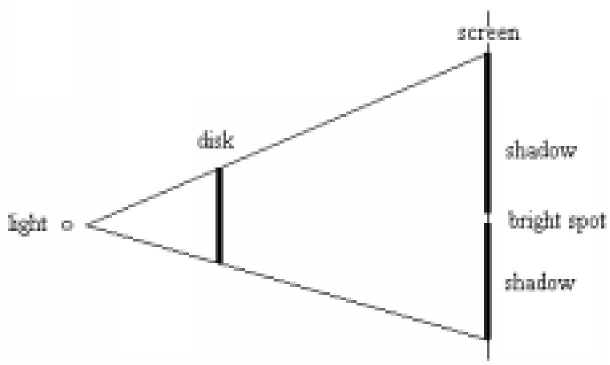

Expected observation if light is transmitted as a wave

When the experiment was performed, the bright spot was observed, much to Poisson's surprise (he, like most physicists of his time, believed H1). This is O2, so H2 is true – provided H1 and H2 exhaust the possibilities.

Unfortunately, we cannot be sure that this is the case for H1 and H2 – or indeed for any pair of hypotheses – unless one of them is just the negation of the other (H2 is not-H1). In the case of light, for example, a third hypothesis H3 is that light is transmitted as a wave in this kind of set-up, but as a stream of particles otherwise.

Now H3 is not a hypothesis that Poisson (or any physicist today) would entertain very seriously. Why not? It is difficult to give a completely satisfactory answer. One popular line of thought is that H2 is simpler (or more general) than H3, and we should accept the simplest (or most general) hypothesis that accounts for the evidence. However, we are interested in accepting *true* hypotheses, so the obvious question is why the simplest (or most general) hypothesis that accounts for the evidence is likely to be true. But any explanation we can dream up – such as:

“God laid down the laws of nature, He prefers economy, hypotheses of this kind claim to express the laws of nature, so simpler (or more general) hypotheses that account for the evidence are more likely to be correct” – is probably going to involve assumptions that are no more obvious than the connection we are trying to make between simplicity and truth.