1. Introduction

1.1. Formal deductive logic

1.1.0. Overview

In this course we will study reasoning, but we will study only certain aspects of reasoning and study them only from one perspective. The special character of our study is indicated by the label formal deductive logic, and we will begin our study by seeing what this label means. Each of the terms formal and logic indicates something about the way in which we will study reasoning while the term deductive indicates the sort of reasoning we will study. In the subsections listed below, we will look at each of these three terms in a little more detail.

1.1.1. Logic
Logic is concerned with features that make reasoning good in certain respects.

1.1.2. Inference
The key form of reasoning that we will consider is inference.

1.1.3. Arguments
The input and output of an inference together form an *argument*.

1.1.4. Deductive *vs.* non-deductive inference
An inference is *deductive* when its conclusion extracts information already present in its premises, and such an inference is risk free.

1.1.5. Bounds on inference
The sentences that constitute risk-free conclusions from given premises form a lower bound on what can be reasonably concluded, and sentences that are absolutely incompatible with those premises form an upper bound.

1.1.6. Entailment and exclusion
Entailment is the relation between the premises and conclusion of a deductive inference, and it will be our main concern. But studying it will involve studying its negative counterpart, *exclusion*.

1.1.7. Inconsistency and exhaustiveness
Other bounds limit the alternatives that can be maintained together—on pain of being *inconsistent*—or assure that there is always one available when a set of alternatives is *exhaustive*.

1.1.8. Formal logic
Many cases of entailment can be captured by generalizations concerning certain linguistic forms, and we will use a quasi-mathematical notation to express these forms.

Several typographical features of the page you are looking at will be reflected throughout the text. A special font (this one) is used to mark language that is being displayed rather than used; the text will frequently use this sort of alternative to quotation marks. Another font (*this one*) is used for special terminology that is being introduced; the index to the text lists these terms and provides links to the points where they are explained. In the list of subsections that appears above, headings have a special formatting (like this) that will be used for links. The links above are links to the subsections themselves, and cross-references in the text with similar formatting will also function as links to portions of the text.

Glen Helman 13 Jul 2012

1.1.1. Logic

Logic is a study of reasoning. However, it does not concern the ways and means by which people actually reason—as psychology does—but rather the sorts of reasoning that count as good. So, while a psychologist is interested as much in cases where people get things wrong as in cases where they get them right, a logician is interested instead in drawing the line between good and bad reasoning without attempting to explain how cases of either sort come about.

Another way of making this distinction between logic and psychology is to say that, in logic, the point of view on reasoning is *internal*: it is a study "from the inside" in a certain sense. As we study reasoning in this way, we will be interested in the norms of reasoning—the rules that reasoners feel bound by, the ideals they strive to reach—rather than the mixed success we observe when we look from outside on their efforts to put norms of reasoning into practice.

This makes logic much like the study of grammar. A linguist studying the grammar of a language will be interested in the sort of things people actually say, but chiefly as evidence of the ways they think words ought to be put together. So, although linguists do not attempt to lay down the rules of grammar for others and see their task as one of description rather than prescription, what they attempt to describe are the (largely unconscious) rules on the basis of which the speakers of a language judge whether utterances are grammatical.

One way of understanding logical norms suggests that there is more than an analogy between logic and the study of language: the norms of thought may be seen to derive from the norms of language, specifically from rules governing certain aspects of meaning. This view is not uncontroversial, but we will see in 1.2 that there is a way of describing the norms of reasoning that makes it quite natural to see them as resting on norms of language.

Glen Helman 11 Jul 2012

1.1.2. Inference

The norms studied in logic can concern many different features of reasoning, and we will consider several of these. The most important one and the one that will receive most of our attention is *inference*, the action of drawing a *conclusion* from certain *premises* or *assumptions*.

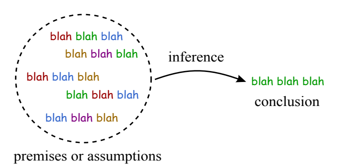

Fig. 1.1.2-1. The action of inference.

This conclusion could be one of the premises, but more often it is different from all of them and draws on several.

Inferences are to be found in science when generalizations are based on data or when a hypothesis is offered to explain some phenomenon. They are also to be found when theorems are proved in mathematics. But the most common case of inference calls less attention to itself. Much of the process of understanding what we hear or read can be seen to involve inference because, when we interpret spoken or written language, our interpretation can usually be formulated as a statement, and we base this statement on the statements we interpret.

The terminology we will use to speak of inference deserves some comment. The terms premise and assumption both to refer to the starting points of inference—whether these be observational data, mathematical axioms, or the statements making up something heard or read. The term premise is most appropriate when we draw a conclusion from a claim or claims that we accept. The term assumption need not carry the suggestion of acceptance (or even acceptability), and we may speak of something being "assumed merely for the sake of argument." In general, we will be far more interested in judging the quality of the transition from the starting point of an inference to its conclusion than in judging the soundness of its starting point, so the distinction between premises and assumptions will not have a crucial role for us. The two terms will serve mainly as alternative expressions for the same idea.

(If it seems strange to consider drawing conclusions from claims that are not accepted, imagine going over a body of data to check for inconsistencies either within the data or with information from other sources. In this sort of case, you may well draw conclusions from data that you do not accept and, indeed, do this as a way of showing that the data is unacceptable—by showing, for example, that it leads to contradictory conclusions.)

Glen Helman 11 Jul 2012

1.1.3. Arguments

It is convenient to have a term for a conclusion taken together with the premises or assumptions on which it is based. We will follow tradition and label such a combination of premises and conclusion an *argument*. A particularly graphic way of writing an argument is to list the premises (in any order) with the conclusion following and marked off by a horizontal line (as shown in Figure 1.1.3-1). The sample argument shown here is a version of a widely used traditional example and has often served as a paradigm of the sort of reasoning studied by deductive logic.

premises	All humans are mortal
	Socrates is human
conclusion	Socrates is mortal

Fig. 1.1.3-1. The components of an argument.

When we need to represent an argument horizontally, we will use / (*virgule* or slash) to divide the premises from the conclusion, so the argument above might also be written as All humans are mortal, Socrates is human / Socrates is mortal.

Notice that the information expressed in the conclusion of this argument is the result of an interaction between the two premises. In its broadest sense, the traditional term *syllogism* (whose etymology might be rendered as 'reckoning together') applies in the first instance to inference that is based on such interaction, and the argument above is a traditional example of a syllogism. Another traditional term, *immediate inference*, applied to arguments with a single premise. The term immediate is not used here in a temporal sense but instead to capture the idea of a conclusion that is inferred from a premise directly and thus without the "mediation" of any further premises.

It is useful to have some abstract notation so that we can state generalizations about reasoning without pointing to specific examples. We will use the lower case Greek letters—most often φ, ψ, and χ—to stand for the individual sentences. And we will use an upper case Greek letter—most often Γ, Σ, and Δ—to stand for a set of sentences, such as the set of premises of an argument. The general form of an argument can then be expressed horizontally as Γ / φ, where Γ is the set of premises and φ is the conclusion.

Although we speak of the premises of an argument as forming a set, in practice what appears above a vertical line or to the left of the sign / will often be a list of sentences, and a symbol like Γ may often be thought of as standing for such a list. The reason for speaking of a set of premises rather than a list is that we will have no interest in the order of the premises or the number of times a premise appears when the premises of an argument are given by a list. We ignore just such features of a list when we move from the list to the set whose members it lists—as we do when we use the notation $\{a_1, a_2, \ldots, a_n\}$ for a set with members $a_1, a_2, \ldots, a_n$. So, although premises will always be listed in concrete examples, we will regard two arguments that share a conclusion as the same when their premises constitute the same set. However, when a symbol for a set appears in a list, we will understand it to abbreviate some list of its members. One common use of this idea will be to employ the notation "Γ, φ" to stand for any list of sentences that includes the sentence φ. That is, we know that φ is included in the list because it is shown; in addition, the list will also include any sentences contributed by Γ, and this may be none at all since Γ might have no members.

Glen Helman 11 Jul 2012

1.1.4. Deductive *vs.* non-deductive inference

Although all good reasoning is of interest to logic, we will focus on reasoning—and, more specifically, on inference—that is good in a special way. To see what this way is, let us begin with a rough distinction between two kinds of reasoning a scientist will typically employ when attempting to account for a body of experimental data.

An example of the first kind of inference is the extraction of information from the data. For instance, the scientist may notice that no one who has had disease A has also had disease B. Even though this conclusion is more than a simple restatement of the data and could well be an important observation, it is closely related to what is already given by the data. It may require perceptiveness to see it, but what is seen does not go beyond the information the data provides. This sort of close tie between a conclusion and the premises on which it is based is characteristic of *deductive reasoning*.

This sort of reasoning appears also in mathematical proof and in some of the inferences we draw in the course of interpreting oral or written language. It is found whenever we draw conclusions that do not go beyond the content of the premises on which they are based and thus introduce no new risk of error. It is this kind of reasoning that we will study, and the traditional name for this study is *deductive logic*.

Science is not limited to the extraction information from data. There usually is some attempt to go beyond data either to make a generalization that applies to other cases or to offer an explanation of the case at hand. A conclusion of either sort brings us closer to the goals of science than does the mere extraction of information, so it is natural to give more attention to an inference that generalizes or explains the data than one that merely extracts information from it. But generalizations and explanations call attention to themselves also because they are risky, and this riskiness distinguishes them from the extraction of information.

There is no very good term—other than *non-deductive*—for the sort of reasoning involved in inferences where we generalize or offer explanations. The term *inductive inference* has been used for some kinds of non-deductive reasoning. But it has often been limited to cases of generalization, and the conclusions of many non-deductive inferences are not naturally stated as generalizations. Although scientific explanations typically employ general laws, they usually employ other sorts of information, too, so they are not just generalizations. And other examples of inferences whose conclusions are the best explanations of some data—for example, the sort of inferences a detective draws from the evidence at a crime scene or that a doctor draws from a patient's symptoms—will often focus on conclusions about particular people, things, or events and are not best thought of as generalizations at all.

Glen Helman 13 Jul 2012

1.1.5. Bounds on inference

Let us now look at the relations between deductive and non-deductive reasoning a little more closely with the aim of distinguishing the role of deductive inference and other aspects of deductive logic.

First notice that there is a close tie between the riskiness of an inference and the question whether it merely extracts information or does something more. The information extracted from data may be no more reliable than the data it is extracted from, but it certainly will be no less reliable. On the other hand, even the generalization or explanatory hypothesis that is most strongly supported by a body of data must go beyond the data if it is to generalize or explain it. And, if this hypothesis goes beyond what the data says, there is a possibility it is wrong even when the data is entirely accurate.

This points to the limits of deductive inference. Still, since the extraction of information can be a first step towards a making a generalization or inferring an explanation. And, by showing how far this first step might take us, deductive logic maps out the territory that we can reach without risking the leap to a generalization or explanatory hypothesis. By distinguishing safe from risky inferences, deductive logic sets a lower bound for inference by marking the range of conclusions that come for free, without risk.

And deductive logic sets bounds for inference also in another respect. One aspect of reasoning is the recognition of tension or incompatibility within collections of sentences, and this, too, has a deductive side. When a incompatibility among sentences is a direct conflict among the claims they make, there is no chance that they could be all be accurate. This sets a sort of upper bound for inference by marking the range of conclusions that could not be supported by any amount of further research. In particular, we know that a generalization can never be supported if our data already provides counterexamples to it, and this sort of constraint is also the concern of deductive logic.

These two bounds are depicted in the following diagram.

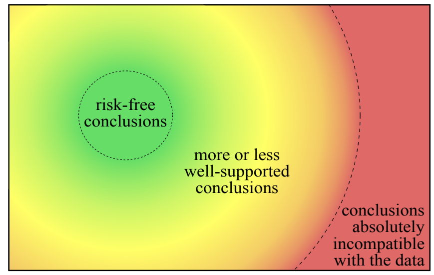

Fig. 1.1.5-1. Deductive bounds on inference.

Sentences in the small circle are the conclusions that are the result of deductive reasoning. They merely extract information and are risk-free and always well-supported. Beyond this circle is a somewhat larger circle with fuzzy boundaries that adds to risk-free conclusions other conclusions that are well supported by the data but go beyond it and are at least somewhat risky. There is large range in the middle of diagram that represents conclusions about which our data tells us nothing. Beyond this, the circle at the right marks the beginning of a region in which we find sentences deductively incompatible with the data. These are claims that are ruled out by the data, that cannot be accurate if the data is accurate. The sentences near this circle but not beyond it are not absolutely incompatible with the data but are in real conflict with it.

The task of deductive logic is to map the sentences within the narrow circle of risk-free conclusions and also to map those that are ruled by our premises. It will turn out that these are not two independent activities: doing one for any substantial range of sentences will involve doing the other.

Glen Helman 13 Jul 2012

1.1.6. Entailment and exclusion

Time has come to provide names for some of the ideas we have been considering. When the conclusion of an argument merely states information extracted from the premises and is therefore risk free, we will say that the conclusion is *entailed by* the premises. If we speak in terms of arguments, entailment is a relation that may or may not hold between given premises and a conclusion, and we will say that an argument is *valid*, or that its has a *valid conclusion*, if the conclusion is entailed by the premises. Figure 1.1.6-1 summarizes these ways of stating the relation of entailment between a set of premises or assumptions Γ and a conclusion φ.

the assumptions Γ *entail* the conclusion φ
the conclusion φ *is entailed by* the assumptions Γ
the conclusion φ *is a valid conclusion from* the assumptions Γ
the argument Γ / φ *is valid*

Fig. 1.1.6-1. Several ways of stating a relation of entailment.

We will use the sign ⊨ (*double right turnstile*) as shorthand for the verb entails, so we add to the English expressions in Figure 1.1.6-1 the claim Γ ⊨ φ as a symbolic way of saying that the assumptions Γ entail the conclusion φ. Using the sign ⊨, we can express the validity of argument in Figure 1.1.2-2 by writing

All humans are mortal, Socrates is human ⊨ Socrates is mortal

Notice that the signs / and ⊨ differ not only in their content but also in their grammatical role. A symbolic expression of the form Γ / φ is a noun phrase since it abbreviates the English expression the argument formed of premises Γ and conclusion φ, so it is comparable in this respect to an expression like x + y (which abbreviates the English the sum of x and y). On the other hand, an expression of the form Γ ⊨ φ is a sentence, and it is thus analogous to an expression like x < y. In short, ⊨ functions as a verb, but the sign / functions as a noun. In Γ / φ, the symbols Γ and φ appear not as subject and object of a verb but as nouns used to specify the reference of a term, much as the names Linden and Crawfordsville do in the term the distance between Linden and Crawfordsville. And the relation between the claims

Γ ⊨ φ
Γ / φ is valid

is analogous to the relation between the claims

Alaska approaches Siberia
The distance from Alaska to Siberia is small

The relation of entailment represents the positive side of deductive reasoning. The negative side is represented by the idea of a statement φ that cannot be accurate when a set Γ of statements are all accurate. In this sort of case, we will say that φ is *excluded by* Γ, and we will say that cases of this sort are characterized by the relation of *exclusion*. We will see later that it is possible to adapt the notation for entailment to express exclusion, so we will not introduce special notation for this relation.

1.1.7. Inconsistency and exhaustiveness

While the bounds on inference marked by entailment and exclusion are the most important topics for deductive logic, it is concerned also with bounds that are not directly bounds on inference. One example comes from looking in a different way at the negative bounds we have just considered. When a possible conclusion is absolutely incompatible with the data, adding it to the data yields a collection of claims that cannot all be true. When a group of sentences are incompatible in this way we are forced to choose among them. So there can be upper bounds on the collections of claims that can be maintained to together. An opposed lower bound comes by trading cannot maintain all for can maintain at least one. And the latter idea lies behind adding None of the above to a list of choices, for we can be sure that, if none of the others can be maintained, it can be, so any list containing it provides at least one acceptable choice. (Notice that, although the relation of this alternative to the others is not one of inference, it can be inferred from the *denials* of the other alternatives, and that will provide the basis for studying this sort of bound by way of bounds on inference.)

Entailment and exclusion are natural opposites, but the nature of the opposition means that the clear distinction between premises and conclusion is no longer found when we consider exclusion. When we say that Γ ⊨ φ, we are saying that there is no chance that φ will fail to be accurate when the members of Γ are all accurate. When we say that Γ excludes φ, we are saying that there is no chance that φ will succeed in being accurate *along with* the members of Γ. In the latter case, we are really saying that a set consisting of sentence consisting of the members of Γ together with φ cannot be wholely accurate, so it is natural to trace the relation of exclusion to a property of *inconsistency* that characterizes such sets: we will say that a set of sentences is *inconsistent* when its members cannot be jointly accurate. Then to say that φ is excluded by Γ is to say that φ is *inconsistent with* (or *given*) Γ in the sense that adding φ to Γ would produce an inconsistent set. The symmetry in the roles of terms in a relation of exclusion is reflected in ordinary ways of expressing this side of deductive reasoning: the difference between saying That hypothesis is inconsistent with our data and Our data is inconsistent with that hypothesis is merely stylistic.

If we turn the idea of inconsistency around—or perhaps inside out—we get the lower bound mentioned above. A set is inconsistent when its members cannot all be true, when at least one must be false. A set is said to be *exhaustive* when its members cannot all be false, when at least one must be true. Such a set exhausts all possibilities in the sense that there is some truth in it no matter what. And the point of adding None of the above is then to provide an exhaustive list of choices, a list that provides at least one option fitting any possibility.

1.1.8. Formal logic

The subject we will study has traditional been given a variety of names. “Deductive logic” is one. Another is *formal logic*, and this term reflects an important aspect of the way we will study deductive reasoning. Even among the inferences that are deductive, we will consider only ones that do not depend on the *subject matter* of the data. This means that these inferences will not depend on the concepts employed to describe particular subjects, and it also means that they will not depend the mathematical structures (systems of numbers, shapes, etc.) that might be employed in such descriptions. This can be expressed by saying that we will limit ourselves to inferences that depend only on the *form* of the claims involved.

The distinction between form and content is a relative one. For example, the use of numerical methods to extract information can be said to depend on content by comparison with the sort of inferences we will study. However, it can count as formal by comparison with other ways of extracting information since all that matters for much of the numerical analysis of data is the numbers that appear in a body of measurements, not the nature of the quantities measured.

Our study is formal in a sense similar to that in which numerical methods are formal, but it is formal to a greater degree. What matters for formal logic is the appearance of certain words or grammatical constructions that can be employed in statements concerning any subject matter. Examples of such logical words are and, not, or, if, is (in the sense of is identical to), every, and some. While this list does not include all the logical words we will consider, it does provide a fair indication of the forms of statements we will study. Indeed, these seven words could serve as titles for chapters 2-8 of this text, respectively. The way in which a statement is put together using words like these (and using logically significant grammatical constructions not directly marked by words) is its *logical form*, and formal logic is a study of reasoning that focuses on the logical forms of statements.

So the subject we will study will be not only deductive logic but formal logic. That means that the norms of deductive reasoning that we will study will be general rules applying to all statements with certain logical forms. It happens that we can give an exhaustive account of such rules in the case of the logical forms that we will consider, so the content of the course can be defined by these forms. *Truth-functional logic*, which will occupy us through chapter 5, is concerned with logical forms that can be expressed using the words and, not, or, and if while *first-order logic (with identity)* is concerned with the full list above, adding to truth-functional logic forms that can be expressed by the words is, every, and some.

Another traditional label for the subject we will study is the term *symbolic logic* that appears in the course title. Most of what this term indicates about the content of our study is captured already by the term *formal logic* because most of the symbols we use will serve to represent logical forms. Certain of the logical forms that appear in the study of truth-functional logic are analogous to patterns appearing in the symbolic statements of algebraic laws. Analogies of this sort were recognized by G. W. Leibniz (1646-1716) and by others after him, but they were first pursued extensively by George Boole (1815-1864), who adopted a notation for logic that was modeled after algebraic notation. The style of symbolic notation that is now standard among logicians owes something to Boole (though the individual symbols are different) and something also to the notation used by Gottlob Frege (1848-1925), who noted analogies between first-order logic and the mathematical theory of functions. This interest in analogies with mathematical theories distinguished logic as studied by Boole and Frege from its more traditional study, and the term symbolic has often been used to capture this distinction. The phrase *mathematical logic* would be equally appropriate, and it has often been used as a label for the subject we will study. However, it has also been used a little more narrowly to speak of an application of logic to mathematical theories that makes these theories objects of mathematical study in their own right. That application of logic in a mathematical style to mathematics itself produces a kind of research that is also known as *metamathematics* (which means, roughly, 'the mathematics of mathematics').

Glen Helman 11 Jul 2012

7 A set of sentences is inconsistent when its members are mutually incompatible. While we can be sure that the members of an inconsistent set are not all true, a set is exhaustive when we can be sure that at least one member is true.

8 We will be interested in the deductive inferences whose validity is a result of the logical form of their premises and conclusions; so our study will be an example of formal logic. The norms of deductive reasoning based on logical form are analogous to some laws of mathematics. The recognition of these analogies (especially by Boole and Frege) has influenced the development of formal deductive logic over the last two centuries, and logic studied from this perspective is often referred to as symbolic logic.

Glen Helman 11 Jul 2012

1.1.s. Summary

The following summarizes this section, looking at it subsection by subsection. Much of the special terminology introduced in the section appears in this summary, and these terms are often links back to the points in the text where they were first introduced and explained.

1 Logic studies reasoning not to explain actual processes of reasoning but instead to describe the norms of good reasoning.

2 The central focus of our study of logic will be inference. We will refer to the starting points of inference as assumptions or premises and its end as a conclusion.

3 These two aspects of a stretch of reasoning can be referred to jointly as an argument. We will mark the distinction by a horizontal line when they are listed vertically and by the sign / when they are listed horizontally. We use the lower case Greek φ, ψ, and χ to stand for individual sentences and upper case Greek Γ, Σ, and Δ to stand for sets of sentences. Our notation for arguments will not distinguish a set from a list of its members; but it is really sets that we focus on because, when considering the norms of inference, we will not distinguish between lists of sentences that determine the same set.

4 Inference that merely extracts information from premises or assumptions and thus brings no risk of new error is deductive inference. Inference that goes beyond the content of the premises to, for example, generalize or explain is then non-deductive. Deductive inference may be distinguished as risk-free in the sense that it adds no further chance of error to the data. The study of the norms of deductive inference is deductive logic, and that is topic of this course.

5 Since deductive inferences are risk free, they provide a lower bound on the inferences that are good. Deductive reasoning also sets an upper bound on good inference by rejecting certain conclusions as absolutely incompatible with given premises.

6 The relation between premises and a conclusion that can be deductively inferred from them is entailment. When the premises and conclusion of an argument are related in this way, the argument is said to be valid. Our symbolic notation for this relation is the sign ⊨, where Γ ⊨ φ says that the premises Γ entail the conclusion φ. To express the upper bounds on inference, we say that a sentence φ is excluded by a set Γ when φ and the members of Γ are mutually incompatible.

1.1.x. Exercise questions

1. Some of the following references to arguments refer to the same argument in different ways (remember that changing the order of premises or the number of times a given premise is referred to does not change the argument being referred to). If Γ stands for the sentences φ, χ, θ, what are the different arguments referred to below? Identify each of the arguments in **a-h** by listing the sentences making up its premises and conclusion and tell which of **a-h** refer to the same argument:

a.	φ, ψ, χ / θ	**f.**	φ, θ, ψ, θ / χ
b.	θ, φ, ψ / χ	**g.**	Γ, φ / ψ
c.	χ, φ, ψ / θ	**h.**	Γ / θ
d.	Γ / ψ	**i.**	χ, θ, φ / ψ
e.	Γ, ζ / ψ	**h.**	Γ, ψ / χ

2. The basis for testing a scientific hypothesis can often be presented as an argument whose conclusion is a prediction about the result of the test and whose premises consist of the hypothesis being tested together with certain assumptions about the test (e.g., about the operation of any apparatus being used to perform the test).

hypothesis to be tested: *hypothesis*
assumptions about the test: { *assumption* ⋮ *assumption* } premises
prediction of the test result: *prediction* conclusion

Suppose that the prediction is entailed by the hypothesis together with the assumptions about the test (i.e., suppose that the argument shown above is valid) and answer the following questions:

a. Can you conclude that the hypothesis is true on the basis of a successful test (i.e., one for which the prediction is true)? Why or why not?

b. Can you conclude that the hypothesis is false on the basis of an unsuccessful test (i.e., one for which the prediction is false)? Why or why not?

Glen Helman 11 Jul 2012

1.1.xa. Exercise answers

1.

	arguments	*references to them*
(1)	φ, χ, ψ / θ	**a**, **c**
(2)	θ, φ, ψ / χ	**b**, **f**
(3)	θ, φ, χ / ψ	**d**, **g**, **i**
(4)	ζ, θ, φ, χ / ψ	**e**
(5)	θ, φ, χ / θ	**h**
(6)	θ, φ, χ, ψ / χ	**j**

2. **a.** Nothing definite can be concluded. The successful test tells you that some true information has been extracted from the hypothesis and auxiliary assumptions. But that can be so even if the hypothesis is not true since a body of information that is not true as a whole can still contain true information. For example, even if the prediction of the result of one test holds true, predictions about other tests may not.

b. You can conclude that the hypothesis is false *provided that the auxiliary assumptions are all true*. The unsuccessful test tells you that a false prediction has been extracted from the hypothesis together with auxiliary assumptions about the test, but this can happen even if the information provided by the hypothesis itself is entirely accurate. The prediction may have failed, for example, because of incorrect assumptions about the way some apparatus would work.