

also function as links.

Glen Helman 01 Aug 2004

1. Introduction

1.1. Formal deductive logic

1.1.0. Overview

The topic of this course is the study of 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 our first task will be to see what this label means. The terms *formal* and *logic* specify the way in which we will study reasoning while the term *deductive* specifies the sort of reasoning we will study. In the course of 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 and arguments

The key form of reasoning that we will consider is inference; the premises and conclusion of an inference make up an **argument**.

1.1.3. Deductive inference

An inference is **deductive** when its conclusion extracts information already present in its premises.

1.1.4. Entailment

Entailment is the relation between the premises and conclusion of a deductive inference.

1.1.5. Formal logic

Many cases of entailment can be captured by generalizations concerning certain linguistic forms.

Several 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. These are links to the subsections themselves, and cross-references in the text with similar formatting will

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 is to say that, in logic, the point of view on reasoning is *internal*, 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 a ***normative discipline***—that is, one whose laws say how things ought to be—rather than a descriptive or explanatory discipline. But there is more than one sense in which a discipline might be seen as normative. Logicians might ask what sorts of reasoning are simply good or, more likely, are good in particular contexts or for particular purposes. But, while they do sometimes ask such questions, logicians more often study features of reasoning that are valued without asking why they are of value.

Now, a study of that sort does not have to be a normative one: chemists may study properties, such as insolubility, that are valuable for certain purposes without that making chemistry a normative discipline. Of course, as a pure science, chemistry does not study properties like insolubility *because* they are valued. But an analogous applied discipline would do this. For example, paint and varnish chemistry will study insolubility because it is valuable property for paints and varnishes to have, and that does not make paint and varnish chemistry a normative discipline.

However, the valuable properties studied in logic are themselves normative in a way that insolubility is not. A comparison with grammar (or the theory of syntax) may help here. A linguist studying the grammar of a language will be interested in the sort of things people actually say, but only 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 utterances as grammatical or ungrammatical. This means that the study of grammar is

normative because its rules prescribe what is required for sentences to be grammatical even if grammarians themselves do not.

One way of putting this would be to say that logic and grammar are disciplines that describe but what they describe are rules or norms. What makes them normative rather than descriptive is that they describe the norms from the inside. Their laws do not generalize about the ways someone who holds such norms will behave but instead state the norms themselves. To cite another example where similar issues arise, someone studying Roman law might do so in the style of a descriptive or explanatory discipline by describing or even trying to explain the operation of the courts or the behavior of the populace, but the study of Roman law has often been normative, attempting to state the content of Roman law as a Roman jurist might have. And it is clear that no one now can really intend to prescribe the law of Rome.

Once the point is seen in this case it can be seen to hold when the law of any state is stated by someone who does not have the legislative or judicial authority to prescribe the content of the law. Logicians and grammarians also lack the authority to prescribe the norms of logic and grammar, not because someone else has it but because no one does. It would be a mistake to place too much emphasis on the contrast between logic and grammar on the one hand and the law on the other because no one has the final authority to prescribe the norms of law either since the authority of legislators and judges depends on the law. The key point is what the study of logic, grammar, and the law share, a normative character that consists in the non-prescriptive statement of rules.

Although such a normative study need not ask the reason for the value of the qualities recognized by the norms it states, one way of understanding the source of logical value suggests that there is more than an analogy between logic and the study of language. However ineffable language itself may sometimes seem, it is vastly more concrete than thought and it has provided logicians with a support for and stimulus to reflection. In the 20th century it acquired an even greater significance because the traditional view of the relation between thought and language (according to which thought is independent of language and language acquires its significance as the expression of thought) came to be reversed, with the significance of thought being seen to derive from the possibility of linguistic expression. As a result, the norms of thought have often been 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 we will study that makes it quite natural to see them as resting on norms of language.

1.1.2. Inference and arguments

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

Inferences are to be found in science when generalizations are based on data or when a hypothesis is offered as the best explanation of 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. We may simply extract information that is provided by the spoken or written text and formulate it as an answer to a question we find of interest, or we may go beyond what has been said or written in a way that clarifies its significance for us. In either case, as in the cases of inference in mathematics and the experimental sciences, we can be understood to formulate a statement that we base on certain other statements. Of course, a reasoner may not formulate an explicit statement of a conclusion or of the data it is based on; but, to the extent that reasoning is articulated sufficiently to apply norms, such statements must be seen to be implicit in it.

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 the claim or claims from which we draw a conclusion are ones that we accept. The term ***assumption*** need not carry this suggestion, and we may speak of something being “assumed for the sake of argument.” But, in general, we will be far more interested in judging 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; and, for the most part, we will use the two terms interchangeably. (If it should seem strange to suppose that you might draw conclusions from claims you do not accept, 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 extract information from data that you do not accept and, indeed, extract this information as a way of showing that the data is unacceptable.)

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 separated off by a horizontal line (as shown in Figure 1.1.2-1). 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	/	<i>All humans are mortal</i>
		<i>Socrates is human</i>
		<i>Socrates is mortal</i>
conclusion		

Fig. 1.1.2-1. The components of an argument.

This example serves to emphasize again that the concepts of inference and argument can be applied not only to reasoning from experimental data or mathematical axioms, but to any reasoning where a conclusion is drawn from certain statements. It also shows that the extraction of information need not be limited to the collection and summary of data. The information expressed in the conclusion is the result of an interaction between the two premises. In its broadest sense, the traditional term **sylllogism** (whose etymology might be rendered as ‘reckoning together’) applies in the first instance to this sort of inference, and the argument above is a traditional example of a syllogism.

It is also useful to have some abstract notation so that we can speak of arguments and their components generally without displaying specific examples. We will use the lower case Greek letters ϕ , ψ , and χ to stand for the individual sentences that may appear as the conclusion of an argument or as its premises. And we will use upper case Greek Γ , Σ , and Δ to stand for sets of sentences, such as the full set of premises of an argument. We will use / (**solidus** or slash) to divide the premises from the conclusion when an argument is represented horizontally, so the argument above might be written horizontally as *All humans are mortal, Socrates is human / Socrates is mortal*. The general form shared by all arguments can then be expressed 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 to the left of the sign / will often be a list of sentences, and a symbol like Γ can often best be thought of as standing for such a list. The reason for speaking of sets at all is that while the items in a list appear in a particular order and can appear more than

once, we have no concern to distinguish arguments on the basis of the order of their premises or the number of times a premise appears; and this means that we regard two arguments that share a conclusion as the same if their premises form the same set. There are other features of sets, however, which are of little use to us. In particular, we have no need to distinguish between a sentence ϕ and the set $\{\phi\}$ that has ϕ as its only member, and we will not attempt to preserve this distinction in our notation for arguments.

If we regard the capital Greek letters as standing for lists of sentences, it makes sense to write $\Gamma, \phi / \psi$ to speak of an argument whose premises consist of the members of Γ together with ϕ —that is, the set of whose premises is the union of Γ and $\{\phi\}$. Since this idea does not exclude the possibility that ϕ is itself a member of Γ , it provides convenient way to refer to any argument whose premises include ϕ .

Glen Helman 01 Aug 2004

1.1.3. Deductive reasoning

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. Let us begin with one example of reasoning: a scientist attempting to account for a body of experimental data. The description of examples of inference in [1.1.2] used a rough distinction between two kinds of reasoning the scientist will typically employ. One kind is the extraction of information from the data. For example, the scientist may notice that nowhere in the data does a certain quantity exceed a certain value. 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. While we might admit that perceptiveness is required to see it in the data, what is seen does not go beyond the information the data provides. The same sort of close relation of conclusion to data can be found in cases where the data is qualitative rather than quantitative in form. For example, someone may notice that two properties are never found together (e.g., no one who has had disease A has also had disease B).

While conclusions like these might attract attention, the process of reasoning used to reach them would probably not be noticed; the extraction of information in such cases is too routine to call attention to itself. Conclusions reached in this way are naturally seen as aspects of the data—albeit aspects that might not be noticed—rather than as results of a process of inference.

The inferences that attract attention are ones that do more than extract information from data. And, at least in science, 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. Either way, we go beyond the data to say something more. A conclusion of one of these sorts will call attention to itself, even when it is easy to reach it from the data, so we will feel that we have done something in forming it. One reason for this sense of accomplishment is that generalization and explanation bring us closer to the goals of science than does the mere extraction of information, so an inference that generalizes or explains the data is more valuable. But another reason why generalizations and explanations call attention to themselves is that they are risky. And this distinguishes them from the extraction of information.

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 completely accurate data can still be wrong. Of course, a scientist will go on to test a generalization or hypothesis, but further testing cannot eliminate the risk. To avoid all risk of error in forming a hypothesis, we would need complete data about the subject matter investigated; and, in the rare cases where a complete set of data can be obtained, there is no further generalization to be made and the best account of the data will merely state information that can be extracted from it. So it seems that, if a generalization or explanation is something other more than the statement of information extracted from an already complete set of data—if it is genuinely hypothetical—it will involve some risk relative to this data. Extraction of information from the data, on the other hand, is completely safe.

Indeed, there is a picture of the matter—a picture that is attached to some of the language we have been using—according to which none of this should be surprising. Information that can be extracted from data is right there in the data, perhaps not for everyone to see but at least implicitly, to be extracted by anyone clever enough. If a conclusion we draw is something other than a statement of such information, it must make claims that are not even implicit in the data, claims that go beyond the data and could prove incorrect however incontrovertible the data might be. This picture also suggests a more positive way of looking at the extraction of information. Extracting information does not merely prepare us to go further; it maps out the territory that we can reach without making the leap to a generalization or explanatory hypothesis.

It is the kind of reasoning exemplified by the extraction of information from data that will be the focus of our study. 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. The traditional name for this study is **deductive logic**. Since the term *deductive* is associated with the features that distinguish the extraction of information from the formation of hypotheses, reasoning that consists in the extraction of information can be labeled **deductive reasoning**.

On the other hand, 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, traditionally, inductive inference was thought of as merely the making of generalizations, but the conclusions of many non-deductive inferences are not naturally stated as generalizations. For example, the sort of

inferences a detective draws will often concern particular people or events (and the interesting examples will be non-deductive in the sense in which we will use the term *deduction*). Because of examples like this, inductive inference would most often now be described as reasoning based on considerations of probability. While this covers inductive generalization and much more besides, it is a matter of controversy whether it covers the full range of reasoning to explanatory hypotheses.

Glen Helman 01 Aug 2004

1.1.4. Entailment

To say that our reasoning is risk-free when we confine ourselves to the extraction of information is to deny the possibility of going wrong when the premises and conclusion of an argument are related in this way. In this sort of case, we will say that the conclusion is **entailed by** the premises. So the extraction of information is characterized by a relation of **entailment** between the initial data and the information extracted from it. 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 can speak of an argument as having the property of **validity** if its premises do entail its conclusion. We will say also that the conclusion of an argument with this property is a **valid conclusion** from its premises. Figure 1.1.4-1 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. 1.1.4-1. Several ways of stating a relation of entailment.

We will use the sign \Rightarrow (**rightwards double arrow**) as shorthand for the verb *entails*, so we add to the English expressions in Figure 1.1.4-1 the symbolic expression $\Gamma \Rightarrow \phi$ as a way of saying that the premises Γ entail the conclusion ϕ . Using this sign, we can express the validity of argument in Figure 1.1.2-1 by writing

All humans are mortal, Socrates is human \Rightarrow Socrates is mortal

Notice that a symbolic expression of the form Γ / ϕ (which amounts to the English expression *the argument formed of premises Γ and conclusion ϕ*) is a noun phrase and is comparable in this respect to the expression $x + y$ (which amounts to the English *the sum of x and y*) while an expression of the form $\Gamma \Rightarrow \phi$ is a sentence (and is comparable to the expression $x < y$).

Entailment and validity are normative concepts since they apply to arguments that are *good* in a certain respect. They might be said to concern the relation that *ought* to hold between the premises and conclusion of an inference if it purports to be risk-free. But this is not to say that all inferences ought to be deductive, that all arguments ought to be valid. There are some contexts, such as mathematical proof, where

the level of security provided by deductive reasoning is required. Still, it often cannot be expected and, in many cases, it would be undesirable. We saw in 1.1.3 that, when our aim is to generalize or explain, deductive inference is not what we want in the end (though it may help along the way).

Glen Helman 01 Aug 2004

1.1.5. Formal logic

Another traditional label for our study is **formal logic**. This term reflects another aspect of our study of 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 they will not depend on the concepts employed to describe particular subjects, but it also means that they will not depend on the mathematical structures (numbers, shapes, etc.) that might be employed in such descriptions. We will limit ourselves to inferences that depend only on the *form* of the claims made in stating the data.

The distinction between form and content is a relative one and, while the use of mathematical methods to extract information will count as a concern with content when it is compared with the sort of inferences we will study, it can count as formal relative to other ways of extracting information. What 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.

In a similar way, what matters for formal logic is the appearance of certain words or grammatical constructions that indicate the kind of claim that is being made by statements expressing the data. Examples of such logical words are *and*, *not*, *or*, *if*, *is* (in the sense of *is identical to*), *every*, and *some*. While this list is by no means exhaustive, it does provide a fair indication of the forms we will study. Indeed, these seven words could serve as title for the seven chapters that will follow this one. The way in which a statement is put together using such words (and using logically significant grammatical constructions not directly marked by words) is its **logical form**, and formal logic is the study of reasoning whose quality rests on the logical forms of statements.

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 forms that can be expressed by the words *is*, *every*, and *some*.

Finally, we will take a quick look at the reason for 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 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 in this way; but this label is also used more specifically for an application of logic to mathematics that takes the theories of mathematics as objects of mathematical study in their own right, a kind of research that is also known as **metamathematics** (which means, roughly, 'the mathematics of mathematics').

Glen Helman 01 Aug 2004

1.1.s. Summary

Logic studies reasoning not to explain actual processes of reasoning but instead to describe valued properties of reasoning by stating norms. It is thus a **normative discipline**.

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**. These two aspects of a stretch of reasoning can be referred to jointly as an **argument**. 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; and we join premises Γ and conclusion ϕ with a **solidus** to indicate the argument Γ / ϕ formed from them.

Considering the difference between extracting information from data and either generalizing from data or offering an explanation of it leads us to a distinction between **deductive** and **non-deductive** inference. Deductive inference may be distinguished as risk free in the sense that it adds no further chance of error to the data. **Deductive logic**, the study of this sort of inference, is our topic in this course.

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 **rightwards double arrow** \Rightarrow , so $\Gamma \Rightarrow \phi$ says that the premises Γ entail the conclusion ϕ .

Among deductive inferences, we can distinguish those that depend on the subject matter of the data and those that depend on the **logical form** of the statements expressing the data; our concern will **only** be with logical form 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 notation for formal deductive logic over the last two centuries, and logic studied from this perspective is often referred to as **symbolic logic**.

Glen Helman 01 Aug 2004

1.1.x. Exercise questions

1. Assume that a statement of entailment $\Gamma \Rightarrow \phi$ holds when the premises Γ listed to the left of the arrow, taken together, contain all the information found in the conclusion ϕ displayed to its right. Using this understanding of entailment, decide for each of the following whether you can be sure that the statement is true (no matter what sentences are put in place of the Greek letters) and briefly explain your reasons. [In some cases a lower case Greek letter (our notation for a single sentence rather than a set) is used on the left of the sign \Rightarrow as shorthand for a set of premises with only a single member.]
 - a. $\phi \Rightarrow \phi$
 - b. if $\phi \Rightarrow \psi$ and $\psi \Rightarrow \chi$, then $\phi \Rightarrow \chi$
 - c. if $\phi \Rightarrow \psi$, then $\psi \Rightarrow \phi$
 - d. if (i) $\Gamma, \phi \Rightarrow \psi$ and (ii) $\Gamma \Rightarrow \phi$, then (iii) $\Gamma \Rightarrow \psi$
[Notice that this says that a premise ϕ of a valid argument $\Gamma, \phi / \psi$ may be dropped without destroying validity provided it is entailed by the remaining premises Γ .]
 - e. if $\chi, \phi \Rightarrow \psi$ and $\chi, \psi \Rightarrow \phi$, then $\phi, \psi \Rightarrow \chi$
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: $\left. \begin{array}{l} \textit{assumption} \\ \dots \\ \textit{assumption} \end{array} \right\} \text{premises}$

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 whose result is as predicted)? Why or why not?
- b. Can you conclude that the hypothesis is false on the basis of an unsuccessful test (i.e., one whose result is not the one predicted)? Why or why not?

Glen Helman 01 Aug 2004

1.1.xa. Exercise answers

1.
 - a. This holds; the premise must contain all the information provided by the conclusion since they are the same sentence.
 - b. If ϕ contains all the information in ψ and ψ contains all the information in χ , then ϕ must contain all the information in χ ; so this claim is true.
 - c. This is not true in general. If ϕ entails ψ then ϕ contains all the information in ψ ; but, if ϕ also contains further information not in ψ , it will not be entailed by ψ .
 - d. This is true; if $\Gamma \Rightarrow \phi$ then all the information provided by the members of Γ together with ϕ is provided by the members of Γ alone, so whenever a sentence is entailed by Γ together with ϕ it will be entailed by Γ .
 - e. This is not true in general. If each of ϕ and ψ is entailed by the other together with χ , we know that χ contains any information that is in one of ϕ and ψ but not the other; but it may also contain further information that is in neither, so it need not be entailed even by the two taken together.
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.

Glen Helman 01 Aug 2004