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 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. Notation for arguments
We will often use some compact ways of stating generalizations about arguments and their components.

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. Deductive 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, exclusion, and inconsistency
Entailment is the relation between the premises and conclusion of a deductive inference, and the terms exclusion and inconsistency are tied to the idea of absolute incompatibility.

1.1.7. 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 topographical 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 03 Aug 2010

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. However ineffable language itself may sometimes seem, it is vastly more concrete than thought, and it has always served logicians as a tool in their study of reasoning. In the 20^{th} 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, and thought was seen to derive its significance from the possibility of linguistic expression. As a result, the norms of thought were 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.

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. 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*.

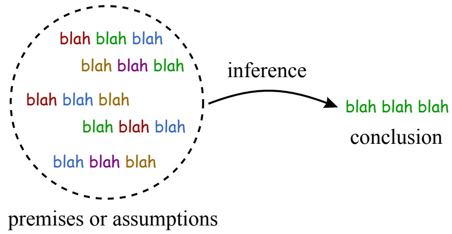

Fig. 1.1.2-1. The action of inference.

This conclusion could be one of the premises, but it is more often formed by drawing on multiple premises.

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 statements in the text 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 should seem strange to consider conclusions inferred 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 draw contradictory conclusions.)

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	All humans are mortal
	Socrates is human
conclusion	Socrates is mortal

Fig. 1.1.2-2. 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.

1.1.3. Notation for arguments

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 φ, ψ, and χ to stand for the individual sentences that may appear as the premises or conclusion of an argument. And we will use upper case Greek Γ, Σ, and Δ to stand for sets 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 basing the idea of an argument on that of a set is that we will have no interest in the order of the premises or the number of times a premise appear if the premises of an argument are listed. 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.

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 {φ} that has φ as its only member, and we will not attempt to preserve the distinction between the two in our notation for arguments. If the capital Greek letters were understood to stand for lists (rather than sets) of sentences, it would make sense to write Γ, φ / ψ to speak of an argument whose premises consisted of the members of Γ together with φ. The *set* of premises of this argument is the *union* Γ ∪ {φ} of the sets Γ and {φ}—i.e., it is the set whose members are the members of Γ and {φ} taken together. Since this idea does not exclude the possibility that φ is already a member of Γ, it provides convenient way to refer to any argument whose premises include the sentence φ. We will understand the notation "Γ, φ" in the same way. That is, imagine the members of Γ are listed, followed by φ. The premises of the argument Γ, φ / ψ are the sentences that appear anywhere in this list. The sentence φ definitely appears, so Γ, φ / ψ is an argument whose premises include φ and whose conclusion is ψ. Since Γ could be any set, this argument may or may not have premises in addition to φ.

We will use an analogous convention in the vertcial notation for arguments.

So, if Γ is the set $\{\varphi, \psi, \chi\}$ (i.e., the set whose members are φ, ψ, and χ) and Σ is the set $\{\psi, \chi\}$, then all of the following refer to the same argument:

horizontal:	Γ / θ	φ, ψ, χ / θ	$\psi, \varphi, \chi, \varphi$ / θ	Σ, φ / θ	Γ, φ / θ	φ, Γ / θ
			ψ			
vertical:		φ	φ			
		ψ	χ	Σ	Γ	φ
	Γ	χ	φ	φ	φ	Γ
	———	———	———	———	———	———
	θ	θ	θ	θ	θ	θ

$\Gamma = \{\varphi, \psi, \chi\}$

$\Sigma = \{\psi, \chi\}$

Fig. 1.1.3-1. Alternative expressions for the same argument (where Γ is the set whose members are ϕ, ψ, and χ and Σ is the set whose members are ψ and χ).

The equivalence of these ways of referring to an argument can be traced to the equivalence among the following ways of referring to the set whose members are φ, ψ, and χ:

$$\{\varphi, \psi, \chi\} = \{\psi, \varphi, \chi, \varphi\} = \{\psi, \chi\} \cup \{\varphi\}$$
$$= \{\varphi, \psi, \chi\} \cup \{\varphi\} = \{\varphi\} \cup \{\psi, \chi\}$$

Glen Helman 03 Aug 2010

1.1.4. Deductive *vs.* non-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. 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 03 Aug 2010

1.1.5. Deductive bounds on reasoning

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.

The extraction of information can be a first step towards a making a generalization or inferring an explanation. We have also seen that extracting information does not merely prepare us to go further: it maps out the territory that we can reach without risking the leap to a generalization or explanatory hypothesis. That is, deductive logic serves to distinguish safe from risky inferences. And this sets a lower bound for inference by marking the range of conclusions that come for free, without risk.

But deductive logic sets bounds for inference in another respect, too. 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. For example, we know that a generalization can never be supported if our data already provides counterexamples to it.

These two bounds are depicted in the following diagram.

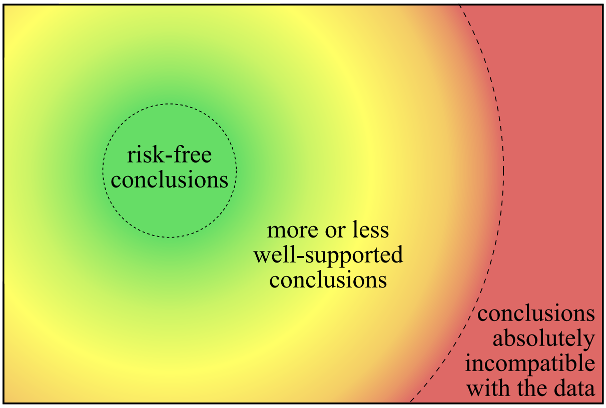

Fig. 1.1.5-1. Deductive bounds on inference.

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 separate activities: doing one for any substantial range of sentences will involve doing the other.

1.1.6. Entailment, exclusion, and inconsistency

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. Using this vocabulary, cases of extraction of information may 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 will say that an argument is *valid* 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 1.1.6-1 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

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.

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.

One aspect of the notation we will use for arguments and entailment deserves a final comment. 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

Linden is close to Crawfordsville
The distance between Linden and Crawfordsville is small

(Of course, there are also many respects in which these pairs of claims are not analogous; for example, the relation expressed by ⊨ has a direction while that expressed by is close to is reversible.)

Glen Helman 03 Aug 2010

1.1.7. 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').

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. These two aspects of a stretch of reasoning can be referred to jointly as an argument. We will separate them by a horizontal line when they are listed vertically and by the sign / when they are listed horizontally.

3 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. A set of sentences is inconsistent when its members are mutually incompatible, and a sentence φ is excluded by a set Γ when φ and the members of Γ are mutually incompatible.

7 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.

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. $\varphi, \psi, \chi / \theta$
b. $\theta, \varphi, \psi / \chi$
c. $\chi, \varphi, \psi / \theta$
d. Γ / ψ
e. $\Gamma, \zeta / \psi$
f. $\varphi, \theta, \psi, \theta / \chi$
g. $\Gamma, \varphi / \psi$
h. Γ / θ
i. $\chi, \theta, \varphi / \psi$
h. $\Gamma, \psi / \chi$

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).

$$\begin{array}{rll}
\text{hypothesis to be tested:} & \textit{hypothesis} & \\
\text{assumptions about the test:} & \left\{\begin{array}{l}\textit{assumption}\\ \vdots \\ \textit{assumption}\end{array}\right. & \text{premises}\\
\hline
\text{prediction of the test result:} & \textit{prediction} & \text{conclusion}
\end{array}$$

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?

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.

1.2. What is said: propositions

1.2.0. Overview

In 1.1.5, we saw the close relation between two properties of a deductive inference: (i) it is a transition from premises to conclusion that is free of any risk of new error, and (ii) the information provided by the conclusion of a deductive inference is already present in its premises. The relation between these properties points to a way of understanding the informational content of a sentence.

1.2.1. Truth values and possible worlds
First we look more closely at the concepts of risk and error involved in the idea of risk-free inference.

1.2.2. Truth conditions and propositions
We can use these ideas to give an account of the content or the meaning of a sentence, an account of what it says.

1.2.3. Ordering by content
When there is a risk-free inference from one sentence to another, the first may say the same thing as a second or it may say more by ruling out some possibility the second leaves open.

1.2.4. Tautologies and absurdities
Two extremes in the ordering of sentences by content are sentences that say nothing and sentences that say too much to distinguish among possibilities.

1.2.5. Logical space and the algebra of propositions
Deductive logic can be seen as the theory of the meanings of sentences in the way that arithmetic is the theory of numbers.

1.2.6. Contrasting content
Other logical relations between sentences concern differences rather than similarities in content. Together with implication, these provide a complete collection of logical relations between two sentences, so sentences related in none of these ways can be described as logically independent.

1.2.1. Truth values and possible worlds

When an inference is deductive, its conclusion cannot be in error unless there is an error somewhere in its premises. The sort of error in question lies in a statement being false, so to know that an argument is valid is to know that its conclusion must be true unless at least one premise is false. Similarly, to know that a set of sentences is inconsistent—to know that it's members are deductively incompatible—is to know that these sentences cannot all be true. This means that the ideas of truth and falsity have a central place in deductive logic, and it will be useful to have some special vocabulary for them.

It is standard to speak of truth and falsity together as *truth values* and to abbreviate their names as **T** and **F**, respectively. So, to say that an argument is valid is to claim that there is no risk of the pattern of truth values for its premises and conclusion shown in Figure 1.2.1-1 occurring. That is (using some of the other terminology we have available), a conclusion is entailed by a set of assumptions if the truth value of the conclusion cannot be **F** when each of the assumptions has the truth value **T**.

$$\begin{array}{rc} & \mathbf{T} \\ \textit{premises:} & \mathbf{T} \\ & \vdots \\ & \mathbf{T} \\ \hline \textit{conclusion:} & \mathbf{F} \end{array}$$

Fig. 1.2.1-1. The pattern of truth values that is not a risk when an argument is valid.

And a set is inconsistent if the truth values of its members cannot all be **T**.

Since to speak of no risk of error is to speak of no possibility of error, it is also useful to have some vocabulary for speaking of possibility and impossibility. The sort of possibility in question in deductive logic is very weak and the corresponding sort of impossibility is very strong. We will refer to this as *logical* possibility and impossibility. A description of a situation that runs counter to the laws of physics (for example, a locomotive floating 10 feet above the earth's surface without any abnormal forces acting on it) might be said to be physically impossible; but it need not be logically impossible, and we must consider many physical impossibilities when deciding whether a conclusion is deductively valid. For, otherwise, anything following from the laws of nature, including the laws themselves, would be a valid conclusion from any premises whatsoever, and these laws would not say anything more than mere descriptions of the facts they were designed to explain. In short, if there is any set of premises such that a sentence φ says something that they do not, then it is logically possible for φ to be false.

We can say that something is impossible by saying that "there is no possibility" of it being true. In saying this, we use a form of words analogous to one we might use to say that there is no photograph of Abraham Lincoln chopping wood. That is, in saying "there is no possibility," we speak of possibilities as if they were things like photographs. This way of speaking about possibilities is convenient, so it is worth spending a moment thinking about what sort of things possibilities might be. The sort of possibility of chief interest to us is a complete state of affairs or state of the world, where this is understood to include facts concerning the full course of history, both past and future. Since Leibniz, philosophers have used the phrase *possible world* as a particularly graphic way of referring to possibilities in this sense. For instance, Leibniz held that the goodness of God implied that the actual world must be the best of all possible worlds, and by this he meant that God made the entire course of history as good as it was logically possible for it to be.

1.2.2. Truth conditions and propositions

When judging the validity of an argument, what we need to know about its premises and conclusion are the truth values of these sentences in various possible worlds. This information about a sentence is an aspect of its meaning that we will call its *truth conditions*. That is, when we are able to tell, no matter what possible world we might be given, whether or not a sentence is true, we know the conditions under which the sentence is true; and, when we know those conditions, we can tell whether or not it is true in a given possible world.

It will also be convenient to be able to speak of this kind of meaning or aspect of meaning as an entity in its own right. We will do this by speaking of the truth conditions of a sentence as encapsulated in the *proposition* expressed by the sentence. This proposition can be thought of as a way of dividing the full range of possible worlds into those in which the sentence is true and those in which it is false—i.e., into the possibilities it leaves open and the ones it rules out. And we can picture a proposition as a division of an area representing the full range of possibilities into two regions.

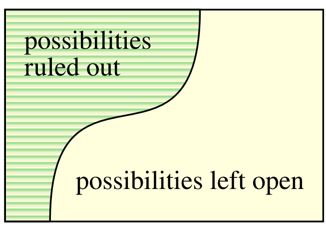

Fig. 1.2.2-1. A proposition dividing the full range of possible worlds into possibilities ruled out and possibilities left open.

Since knowing what possibilities are in one of these regions tells us that the rest are in the other region, we know what proposition is expressed by a sentence when we know what possibilities it rules out—or know what possibilities it leaves open. And focusing on one or the other of these sets of possibilities may be helpful in certain contexts.

1.2.3. Ordering by content

When we judge the validity of an argument we are comparing the content of the conclusion to the contents of the premises, and the ideas of truth values and possible worlds are designed to help us speak about the basis for that comparison. We can see more of what this sort of comparison involves and what similar comparisons are possible by focusing on comparisons of two sentences.

The term *implies* is a more common English synonym of entails, and we will use it often when considering an argument that has only one premise (i.e., an "immediate inference" in traditional terminology noted in 1.1.2). Thus φ implies (or entails) ψ when there is no risk that ψ will be in error without any error in φ—i.e., when there is no logically possible world in which ψ is false even though φ is true. When φ implies ψ, the content of ψ can be extracted from the content of φ, so to say that $\varphi \vDash \psi$ is to say that φ includes the content of ψ. Thus the relation of implication orders sentences according to their content.

If this relation holds in both directions—if both $\varphi \vDash \psi$ and $\psi \vDash \varphi$—then each of the two sentences says everything the other does, so they provide exactly the same information, differing at most in their wording. For example, although one of the sentences Sam lives somewhere in northern Illinois or southern Wisconsin and Sam lives somewhere in southern Wisconsin or northern Illinois might be chosen over the other depending on the circumstances, they allow the same possibilities for Sam's residence and thus provide the same information about it. We will say that sentences that have the same informational content are *(logically) equivalent* (usually dropping the qualification logically since we will not be considering other sorts of equivalence). Our notation for logical equivalence—the sign ≃ (*tilde equal*)—gets used for many different kinds of equivalence, but we will use it only for logical equivalence.

The idea of logical equivalence can also be described directly in terms of truth values and possible worlds. When two sentences say the same thing there is no way for one to be in error when the other is not. That is to say, sentences are equivalent when there is no possible world in which they have different truth values. To put it yet another way, no matter what things are like, a pair of equivalent sentences will both be accurate or both be in error. This means that, when $\varphi \simeq \psi$, we know that in any possible world we might consider, φ and ψ will both have the same truth value. And that means that equivalent sentences have the same truth conditions and express the same proposition.

Since relations of entailment depend only on possibilities of truth and falsity, equivalent sentences entail and are entailed by the same sentences. That means that entailment can be thought of as a relation between the propositions they express. It provides a sort of ordering of propositions by their content that can be compared to the ordering of numbers by $\leq$ and $\geq$. Whether entailment seems more like $\leq$ or $\geq$ depends on whether we think of it as a comparison of possibilities left open or of possibilities ruled out. When a choice needs to be made, we'll general adopt the former perspective. In any case, the analogy is with $\leq$ or $\geq$ rather than $<$ or $>$ because $\varphi \vDash \psi$ tells us that φ says more *or the same as* ψ, that it leaves fewer *or the same* possibilities open.

When φ does say something more than ψ—that is, when $\varphi \vDash \psi$ but $\psi \nvDash \varphi$—the possibilities left open by ψ will include all those left open by φ (because $\varphi \vDash \psi$) but it will leave open some on top of these (because there is some possible world in which ψ is true but φ is false). To see an example of this, consider the following series of successively more specific statements, each implied by the one below it:

The package will arrive sometime
is implied by
The package will arrive next week
is implied by
The package will arrive next Wednesday
is implied by
The package will arrive next Wednesday morning

Each sentence until the last leaves open some possibilities that are ruled out by the sentence below it. And in general, as we add information, we reduce the range of possibilities left open and increase the range that are ruled out. We will often speak of a sentence that rules out more and leaves open less as making a ***stronger*** claim and of one that rules out less and leaves open more as making a ***weaker*** claim. So, in the list above, the sentences closer to the bottom make the stronger claims and those closer to the top make the weaker ones.

We have been employing analogies between implication and numerical ordering and the related sorts of comparison that are associated with terms like stronger and weaker. These analogies rest on two properties that implication shares with many other relations. First of all, it is ***transitive*** in the sense that implication by a premise φ carries over from a valid conclusion ψ to any sentence χ implied by that conclusion: if $\varphi \vDash \psi$ and $\psi \vDash \chi$, then $\varphi \vDash \chi$. That is, we do not count steps in a chain of related items (as is done with parent of, grandparent of, etc., which are not transitive) but simply report the existence of a chain no matter what its length (as is done with ancestor of, which is transitive).

Just about any relation that we would be ready to call an "ordering" is transitive. Implication also shares with certain orderings the more special property of being *reflexive* in the sense that every sentence implies itself. Reflexivity is what distinguishes orderings like $\leq$ and as strong as or stronger than from $<$ and stronger than. In the first two, examples reflexivity is achieved by tacking on a second reflexive relation ($=$ in one case and equally strong as in the other) as an alternative. The analogous relation in the case of implication (i.e., one amounting to "equal in content to") is equivalence, but that is an alternative already built into implication (i.e., one sort of case in which a sentence φ implies a sentence ψ is when they are equivalent), so it does not need to be added.

Relations like $=$, equally strong as, and equivalence are reflexive and transitive, but they are not very effective in ordering things because they have no direction: if they hold between a pair of things in one direction, they hold in the other direction, too. In particular, if $\varphi \simeq \psi$ then $\psi \simeq \varphi$. A relation with this property is said to be *symmetric*. Relations with the three properties of transitivity, reflexivity, and symmetry are said to be *equivalence relations*. Any equivalence relation points to equivalence or equality in some respect, and different relations point to different sorts of equality or equivalence. Logical equivalence points to equivalence in content.

1.2.4. Tautologies and absurdities

There are two extreme examples of truth conditions or propositions. A sentence that is true in all possible worlds says nothing. It has no informational content because it leaves open all possibilities and rules nothing out. For example, the weather “forecast” Either it will rain or it won’t has no chance of being wrong and is, therefore, completely worthless as a prediction. We will say that such a sentence is a *tautology*. Although there are many (indeed, infinitely many) tautologies, all express the same proposition; and the words that they use to express it are beside the point since they all say nothing in the end. In short, any two tautologies are logically equivalent. It will be convenient to establish a particular tautology and mark it by special notation. We will call this sentence *Tautology* and use the sign ⊤ (*down tack*) as our notation for it. Since the logical properties and relations we will consider depend only on the propositions expressed by sentences, any logical property or relation of ⊤ will hold for all tautologies, and we will often simply speak of ⊤ in order to say things about tautologies generally.

At the other extreme of truth conditions from tautologies are sentences that rule out all possibilities. The fact that such a sentence is the opposite of a tautology might suggest that it is maximally informative, but it sets an upper bound on informativeness in a different way: any genuinely informative sentence must say less than it does. The ultimate aim of providing information is to narrow down possibilities until a single one remains, for this would provide a complete description of the history of the universe. To go beyond this would leave us with nothing because there is no way to distinguish among possibilities if all are ruled out. For example, the forecast It will rain, but it won’t is far from non-committal since it stands no chance of being right, but it is no more helpful than a tautologous one.

Sentences that rule out all possibilities make logically impossible claims, and we will refer to them as *absurd*. As was the case with tautologies, any two absurd sentences are logically equivalent. So, as with tautologies, we will introduce a particular example of an absurdity, named *Absurdity*, and we use the special notation ⊥ (the perpendicular sign, or *up tack*) for it.

A tautology is implied by any sentence φ since, as it rules out no possibilities, it cannot rule out any possibility that is left open by φ. The sentence ⊤ is thus the weakest sentence there could be and it can stand at the top of any ordering by logical strength like that depicted in 1.2.3. Analogously, an absurd sentence implies all sentences, and the sentence ⊥ can stand at the bottom of any ordering by logical strength.

Any sentence implying $\bot$ is thus equivalent to it and is itself absurd. More generally, the idea of entailing $\bot$ provides way characterizing inconsistency. That is, we can have $\Gamma \vDash \bot$ only when it is not possible for the premises Γ to all be true, and premises that cannot all be true will entail any conclusion, including $\bot$. This characterization of inconsistency in terms of entailment will help us to keep our focus on entailment. Laws governing inconsistency—and, by way of it, principles governing related ideas like exclusion—will appear as principles governing valid arguments with the conclusion $\bot$. In fact, we are not really dispensing with the idea of inconsistency since an absurdity amounts to a sentence that forms an inconsistent set all by itself. The role of entailment will be to enable us to study the full range of inconsistent sets by way of this one very special example.

1.2.5. Logical space and the algebra of propositions

Logic is concerned with propositions in the way mathematics is concerned with numbers, but propositions are not numbers. While numbers can be ordered in a linear way, the collection of propositions has a more complex structure. The series of examples of increasing strength we looked at in 1.2.3 did form a single chain, but it should be clear that we could have gone in many different directions to find stronger or weaker claims propositions. For example, The package will arrive next Wednesday is implied by The package will arrive next Wednesday morning but also by The package will arrive next Wednesday afternoon, and neither of the latter sentences implies the other. And The package will arrive next Wednesday implies the sentences The package will arrive next week and The package will arrive on a Wednesday, and the latter two sentences are not ordered one way or the other by implication.

This metaphor of many directions suggests a space of more than one dimension; and, although the structure of a collection of propositions differs not only from the 1-dimensional number line but also from the structure of ordinary 2- or 3-dimensional space, spatial metaphors and diagrams can help in thinking about its structure. These metaphors and can be associated with the term *logical space* that was introduced by the philosopher Ludwig Wittgenstein (1889-1951).

We will actually use two different sorts of spatial metaphor. One metaphor is the one used in 1.2.2 to depict propositions. In it, possible worlds are the points of logical space, and propositions determine regions in the space by drawing a boundary between the possibilities they rule out and the ones they leave open. But we use a different metaphor when we speak of increasing strength in many different directions. According to this second metaphor, propositions are points in space rather than regions, and possible worlds function in it behind the scenes as something like the dimensions of the space. If we were to apply this idea in any very realistic way, the space would have too many dimensions to be visualized, but in artificially simple cases this sort of space can be depicted by a figure in ordinary 2- or 3-dimensional space.

Let's begin to look further at these ideas by considering an very simple example of the first sort of logical space. Suppose there were only 4 possible worlds. A proposition will either rule out or leave open each of these possibilities. Figure 1.2.5-1 is intended to illustrate two such propositions.

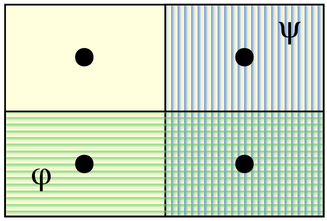

Fig. 1.2.5-1. The possibilities (the hatched bottom and right halves) that are ruled out by two propositions.

Each of these propositions rules out two of the four possibilities (in the hatched areas) and leaves open two others. The proposition expressed by the sentence φ rules out the two possibilities at the bottom of the diagram and the one expressed by ψ rules out the ones at the right. As a result both rule out the possible world in the lower right of the diagram and neither rules out the one in the upper left.

Of course, these are not the only propositions that can be expressed given this range of possibilities. A proposition has two options for each possible world: it may rule it out or leave it open. With 4 possible worlds this means that there are $2\times2\times2\times2 = 16$ propositions in all, and 6 of these rule out exactly two possible worlds.

We can illustrate all 16 of these propositions by using a logical space of the second sort. Figure 1.2.5-2 depicts (in two dimensions) a 3-dimensional figure that is one possible representation of a 4-dimensional cube. It is labeled to suggest what sorts of sentences might express these propositions.

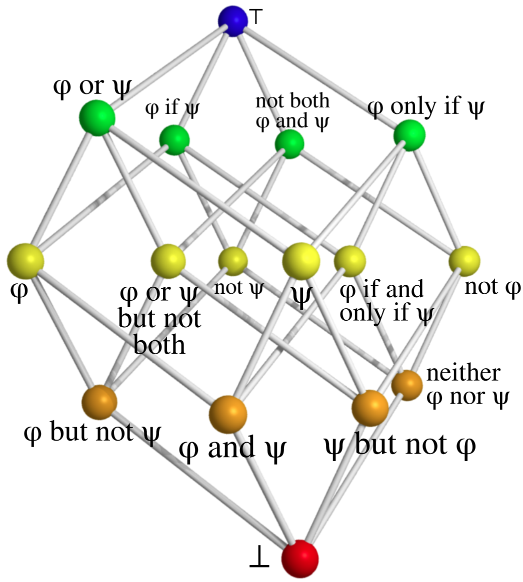

Fig. 1.2.5-2. The sixteen propositions when there are 4 possible worlds.

You can imagine that the propositions φ (which appears at the left) and ψ (near the center) are the two propositions depicted in Figure 1.2.5-1.

The levels in the structure correspond to grades of strength, with Absurdity at the bottom ruling out all possible worlds and Tautology at the top ruling out none. A line connects propositions that differ only with respect to one possible world. The proposition lower in the diagram rules out this world and the one above it leaves the world open, so the lower proposition implies the one above it. Each of the four propositions immediately above Absurdity then leaves open just one possible world. Lines connecting propositions that differ with respect to a given world are parallel (in the 3-dimensional figure, not in its 2-dimensional projection); and, in this sense, the worlds can be thought of as the dimensions on which the content of propositions can vary.

The relation between the two sorts of diagram can be seen by replacing each proposition in Figure 1.2.5-2 by its representation using a diagram of the sort illustrated in Figure 1.2.5-1. Putting the two sorts of illustration together in this way gives us the following picture of the same 16 propositions.

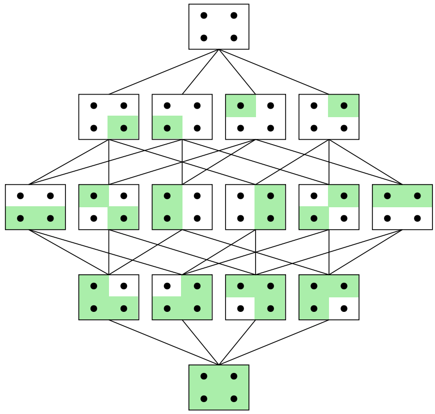

Fig. 1.2.5-3. The propositions generated by 4 possible worlds depicted as regions in one logical space (the repeated rectangle) and as points in another (the overall diagram).

The spacing of the nodes differs between Figures 1.2.5-2 and 1.2.5-3 but the left-to-right order at each level is the same, and the regions associated with φ and ψ are the same as those depicted in Figure 1.2.5-1. Since a sentence that rules out more possibilities makes a stronger claim, the size of the region occupied by the possibilities it rules out can be thought to correspond to the strength of the claim it makes. Notice that the regions ruled out here increase towards the bottom of the diagram and that they are the same in size for all nodes on the same level.

The whole structure of Figure 1.2.5-2 can be seen as a complex diamond formed of four diamonds whose corresponding vertices are linked. A simple diamond is the structure of the $2 \times 2 = 4$ propositions we would have with only 2 possible worlds. The structure in Figure 1.2.5-2 doubles the number of possible worlds and squares the number of propositions. If we were to double the number of possible worlds again to 8, we would square the number of propositions to get 256. The structure they would form could be obtained by replacing each node in the structure of Figure 1.2.5-2 by a small structure of the same form and replacing each line by a bundle of 16 lines connecting the corresponding nodes.

To get a sense of the structure of the set of propositions for a realistically large set of possible worlds, imagine carrying out this process over and over again. The result will always have an upper and lower limit (⊤ and ⊥) and many different nodes on each of its intermediate levels. As the number of possible worlds increases, the distribution of possible worlds among the various degrees of strength (which is 1, 4, 6, 4, 1 in Figure 1.2.5-2) will more and more closely approximate a bell curve. But the bell shape of the curve will also narrow significantly, and bulk of the propositions will be found in intermediate degrees of strength. In short, as the space of propositions gets closer to a realistic degree of complexity, it departs further and further from a single line with $\top$ at the top and $\bot$ at the bottom.

Glen Helman 03 Aug 2010

1.2.6. Contrasting content

We arrived at the relation of implication by considering entailment by a single premise. If we do the same with exclusion, we arrive at another relation between sentences. If φ excludes ψ, then the set {φ, ψ} formed of the two is inconsistent. When sentences φ and ψ are related in this way, it is equally true that ψ excludes φ. This reversability of this relation is reflected in the usual terminology for it: when there is no possible world in which φ and ψ are together true, φ and ψ are said to be *mutually exclusive*. There is no standard notation for the relation, and we will shortly have a way of expressing it in terms of entailment; but, when it is convenient to have special notation, we will write φ ▵ ψ to say that φ and ψ are mutually exclusive. This use of the *up-pointing triangle* is intended simply to reflect the shape of signs for some related ideas. One of these related ideas is Absurdity. In particular, notice that sentences φ and ψ are mutually exclusive if and only if they form an inconsistent set—that is, if they together entail ⊥.

Mutually exclusive sentences provide one example of the differences in propositions that made for the horizontal spread of the logical space of Figure 1.2.5-2. Indeed, one of the examples cited there, the sentences The package will arrive next Wednesday morning and The package will arrive next Wednesday afternoon was a pair of mutually exclusive sentences. Mutually exclusive sentences differ to the extent that there is no overlap in the possibilities they leave open. From one point of view, that is a pretty considerable difference; but, as the example illustrates, such sentences can still have a lot in common. And, in general, sentences that rule out many possibilities may express propositions that divide the space of possibilities in very similar ways even though they have no overlap in the ones they leave open.

Mutually exclusive sentences are opposed to one another, and they can be thought of as opposites. But there are different sorts of opposites. Some, like The glass is full and The glass is empty are extremes that may both fail in intermediate cases. Others, like The glass is full and The glass is not full cover all the ground between them and do not leave room for a third alternative. Opposites of the latter sort might be described as *exactly* opposite.

The difference between these sorts of opposition is tied to another way in which sentences can differ. Sentences φ and ψ are *jointly exhaustive* when there is no possible world in which both are false, when there is no possible world that both rule out. If we put together the possibilities left open by such sentences, the result will include all possibilities because any possibility ruled out by one must be left open by the other; and, in this sense, these sentences jointly exhaust all possibilities. Such sentences certainly differ in meaning —since there is no overlap in the possibilities they rule out, they can be said to have no common content—but they are not opposites in the sense of being incompatible. They might be thought of instead as *complementary* since, in regard to possibilities left open, each picks up where the other leaves off. We will use a *down-pointing triangle* ▿ as our notation for this relation, as in the case of ▵ because of the similarity in shape between ▿ and some ideas related to joint exhaustiveness. (Tautology is one of these ideas but we will not consider the relation between it and joint exhaustiveness until 1.4.)

When sentences are not only mutually exclusive but also jointly exhaustive, they are opposed in the second way described above: since they cannot both be false, one or the other is bound to hold and there is no room for a third alternative and they are exactly opposite. We will say that two sentences for which this is so are *contradictory*. Contradictory sentences—like The glass is full and The glass is not full—are bound to have opposite truth values. We will write φ ⧖ ψ to say that φ and ψ are contradictory (using the symbol *hourglass*). (You might think of the symbol as indicating that things get turned upside down when moving from one sentence to the other.)

Although our use of the term contradictory is the standard one in discussions of deductive logic, in ordinary speech this term is often applied to sentences that are only mutually exclusive. In particular, when a claim is said to be "self-contradictory," what is meant is that part of what it says excludes something else it says. Such a sentence will not contradict itself in the sense in which we will use the term because that would require that it be both true and false in each possible world, and that cannot happen if there are any possible worlds at all (an assumption we can feel safe in making).

Just as the propositions expressed by logically strong sentences need not be far different even when they are mutually exclusive, the propositions expressed by logically weak sentences need not be far different even when they are jointly exhaustive. It is contradictory sentences that provide the true extreme examples of difference. When logical space in Figure 1.2.5-2 is thought of in three dimensions, the contradictory sentences appear in diametically opposite positions. Notice that mutually exclusive sentences cannot both appear above the middle level (for such sentences leave open more than half the possibilities), and jointly exhaustive sentences cannot appear both below the middle. Contradictory sentences fall under both restrictions. A pair of contradictory sentences might both appear on the middle level, but it is also possible for one to be of more than average logical strength if the other is relatively weak. The extreme case of this is provided by ⊥ and ⊤, which are contradictory and constitute the only example of a contradictory pair the first of whose members implies the second.

The four basic deductive relations between two sentences that we have considered are shown in the following table:

Relation	*pattern ruled out*	
φ implies ψ (φ ⊨ ψ)	φ is **T**	ψ is **F**
φ is implied by ψ (ψ ⊨ φ)	φ is **F**	ψ is **T**
φ and ψ are mutually exclusive (φ △ ψ)	φ is **T**	ψ is **T**
φ and ψ are jointly exhaustive (φ ▽ ψ)	φ is **F**	ψ is **F**

These are the only relations that can be defined by ruling out a specific pattern of truth values for two sentences because there are only four such patterns. Ruling out more than one pattern does not give us any relations beyond those already discussed. If we rule out the first two patterns, we are saying that the sentences are equivalent, and if we rule out the last two patterns, we are saying that they are contradictory. If we were to rule out any other pair of patterns, we would simply rule out a truth value for one of the sentences in all possible worlds, so we would be saying of this sentence that it was tautologous or that it was absurd. And that meas we would be describing a property of a single sentence rather than a relation between sentences. And ruling out three patterns would leave just one pattern and would specify the truth values of both sentences, saying of each them that it was tautologous or absurd. So, in one sense, the six relations for which we have terminology are the only ones possible.

Relations between the propositions expressed by a pair of sentences can be depicted by relations of areas in logical space. The regions ruled out are shown shaded in the left column in Figure 1.2.6-1, and the regions left open are shown hatched in the right column.

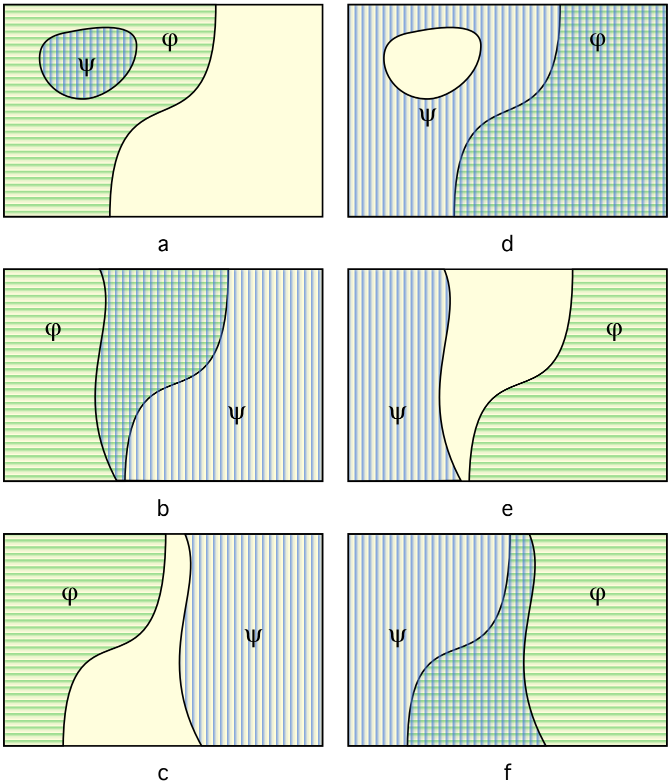

Fig. 1.2.6-1. Three relations between sentences ϕ and ψ. (a, d) ϕ implies ψ. (b, e) ϕ and ψ are mutually exclusive. (c, f) ϕ and ψ are jointly exhaustive. On the left, regions ruled out by sentences are hatched—horizontally in green for ϕ and vertically in blue for ψ. The regions left open by ϕ and ψ are hatched similarly on the right.

When $\varphi \vDash \psi$ (see a and d above), the implied sentence ψ does not rule out any possibility not already ruled out by the implying sentence φ, so the region ruled out by φ must include the region ruled out by ψ (and the region left open by φ must therefore be included in the region left open by ψ). If φ and ψ are mutually exclusive (see b and e above), there can be no overlap in the regions they leave open so the regions ruled out by the two must together cover the full range of possibilities. Here φ rules out all worlds at the left of the rectangle and ψ rules out all worlds at the right, with both ruling out a swath of worlds in the middle. It is the same thing to say that there is no overlap in the worlds they leave open, a situation depicted on the right (in e). Finally, when φ and ψ are jointly exhaustive, the situation is reversed (see c and f above): the regions left open by the two must together cover all possibilities, so the regions they rule out cannot overlap. In the diagram a swath of worlds through the middle is left open by both (see f).

When none of these relations hold between a pair of sentences φ and ψ—that is, when each of four patterns of truth values for the two appears in some possible world—we will say that φ and ψ are *logically independent*. Not only are logically independent sentences unordered by implication, they are not tied by any deductive relation. And this sort of thing holds for most pairs of sentences. Although sentences on different topics almost always provide examples, logically independent sentences do not need to differ in subject matter. For example, the sentences The package will arrive next week and The package will arrive on a Wednesday (a pair of sentences mentioned in 1.2.4) are logically independent since it is possible for the package to arrive next week but not on Wednesday (so the first doesn't imply the second), for it to arrive on a Wednesday but not next week (so the first isn't implied by the second), for it to arrive next Wednesday (so they aren't mutually exclusive), and for it to arrive neither next week nor on a Wednesday (so they aren't jointly exhaustive).

1.2.s. Summary

1 The relation of entailment concerns the possibilities of truth and falsity for premises and conclusions; that is, it concerns the truth values of these sentences in various possible worlds. The possibilities in question are logical possibilities, which may be understood as the situations whose description is permitted by the semantic rules of the language.

2 The deductive relations a sentence stands in depend on its truth values in various possible worlds. That is, they depend on its truth conditions. These truth conditions are encapsulated in the proposition it expresses, which can be thought of as a way of dividing all possibilities into those it rules out and those it leaves open. This means that a proposition can be depicted as a division of space into two regions.

3 Entailment by a single premise, or implication, is a relation between sentences that orders them by their content. More precisely, $\varphi \vDash \psi$ when φ says everything that is said by ψ. When sentences imply each other, they say the some thing, and we say they are equivalent, a relation for which we use the sign $\simeq$. When $\varphi \vDash \psi$ but these sentences are not equivalent, φ says more than ψ and we will often say that φ makes a stronger claim and ψ a weaker one.

4 At one extreme are tautologies, which rule out no possibilities and thus have no content. All tautologies are equivalent and we will distinguish one, Tautology, for which we use the notation $\top$. At the other extreme are sentences that rule out all possibilities. Such sentences are absurd and all are equivalent to the single representative Absurdity, for which we use the notation $\bot$. An argument with an absurd conclusion is valid when and only when its premises form an inconsistent set, and this will enable us to study inconsistency by way of entailment.

5 Although certain groups of sentences can be ordered linearly between $\bot$ and $\top$ as a series of claims with steadily increasing content, the full range of propositions expressed by sentences are better thought of as inhabiting a much more complex logical space. This space might be a space of possibilities with propositions appearing as ways of dividing the space into regions, or it might be a space that has as its points propositions themselves. Logical space in this second sense has a bottom in the proposition expressed by $\bot$ and a top provided by $\top$. When there are a significant number of possible worlds, there will be many more propositions with intermediate content than there are strong propositions near $\bot$ or weak ones near $\top$.

6 Sentences can also be compared by describing differences in what they say. Sentences that cannot both be true are *mutually exclusive* (a relation for which we use the sign $\vartriangle$). The claims made by such sentences are opposite but opposite in a way that permits a third alternative. Sentences which are complementary in the sense that each must be true if the other is false are *jointly exhaustive* (for which our notation is $\triangledown$). When these two relations both hold, sentences are *contradictory* (a relation for which we use the sign $\bowtie$). Contradictory sentences always have opposite truth values and thus make claims that are opposite in a way that permits no third alternative. Sentences that are neither mutually exclusive nor jointly exhaustive and neither or which implies the other are *logically independent*.

Glen Helman 03 Aug 2010

1.2.x. Exercise questions

1. Each of the following claims that a deductive relation holds between a pair of sentences. In each case, judge whether the claim is true and, if not, describe a sort of possibility that shows it is not true. Briefly explain your answers. For example, we can say that The package will arrive sometime does not entail The package will arrive next week because the possibility that it will arrive before or after next week is ruled out by the conclusion but not by the premise. In answering, it is safe to understand the sentences below all in the most straightforward way; you will miss the point of the exercise if you try to look for subtle or obscure possibilities.

- **a.** The package will arrive next Tueday $\vDash$ The package will arrive next week
- **b.** The package will arrive next week $\vDash$ The package will arrive next Tuesday
- **c.** The package will arrive next Tueday $\vartriangle$ The package will arrive next week
- **d.** The package will arrive next Tuesday $\vartriangle$ The package will arrive next Wednesday
- **e.** The package will arrive before next Tueday $\triangledown$ The package will arrive after next Tuesday
- **f.** The package will arrive next Tuesday or before $\triangledown$ The package will not arrive before next Wednesday
- **g.** The package will arrive after next Tuesday $\simeq$ The package will arrive next Wednesday or later
- **h.** The bridge will open at the end of May $\simeq$ The bridge will open before June
- **i.** The package will arrive before next Wednesday ⧖ The package will arrive after next Wednesday
- **j.** The bridge will open before June ⧖ The bridge will open in June or later or never at all

2. Some of the following claims about deductive relations hold for any sentence φ, some for no sentence φ, and others hold only if φ is a tautology or only if it is absurd. In each case, say which is so and explain your answer.

a.	$\varphi \vDash \varphi$	**b.**	$\varphi \vDash \top$	**c.**	$\varphi \vDash \bot$
		d.	$\top \vDash \varphi$	**e.**	$\bot \vDash \varphi$
f.	$\varphi \triangledown \varphi$	**g.**	$\varphi \triangledown \top$	**h.**	$\varphi \triangledown \bot$

i. $\varphi \vartriangle \varphi$ **j.** $\varphi \vartriangle \top$ **k.** $\varphi \vartriangle \bot$

l. $\varphi \simeq \varphi$ **m.** $\varphi \simeq \top$ **n.** $\varphi \simeq \bot$

o. $\varphi ⧖ \varphi$ **p.** $\varphi ⧖ \top$ **q.** $\varphi ⧖ \bot$

3. The headings at the left of the table give information about the relation of φ and ψ and those at the top give information about the relation of ψ and χ. Fill in cells of the table by indicating what, if anything, you can conclude in each case about the relation of φ and χ. For example, if $\varphi \vDash \psi$ and $\psi \vDash \chi$, we cannot have φ true and χ false, so $\varphi \vDash \chi$ (this is the transitivity of implication). However, no other patterns for φ and χ are ruled out, so "$\varphi \vDash \chi$" is the most we can say on the basis of the given information, and it can be entered in the upper left cell.

	$\psi \vDash \chi$	$\chi \vDash \psi$	$\psi \simeq \chi$	$\psi \vartriangle \chi$	$\psi \triangledown \chi$	$\psi ⧖ \chi$
$\varphi \vDash \psi$						
$\psi \vDash \varphi$						
$\varphi \simeq \psi$						
$\varphi \vartriangle \psi$						
$\varphi \triangledown \psi$						
$\varphi ⧖ \psi$						

1.2.xa. Exercise answers

1. a. The package will arrive next Tueday entails The package will arrive next week because the package arriving next Tuesday is one of ways for it to be true that it arrives next week

   b. The package will arrive next week does not entail The package will arrive next Tuesday because the premise would still be true if it arrived another day next week

   c. The package will arrive next Tuesday and The package will arrive next week are not mutually exclusive because both will be true if it does arrive next Tuesday

   d. The package will arrive next Tuesday and The package will arrive next Wednesday are mutually exclusive since the package cannot arrive both days

   e. The package will arrive before next Tueday and The package will arrive after next Tuesday are not jointly exhaustive since both will be false if it arrives on next Tuesday

   f. The package will arrive next Tuesday or before and The package will not arrive before next Wednesday are jointly exhaustive because, if the second is false—i.e., if it does arrive before next Wednesday—then the first must be true

   g. The package will arrive after next Tuesday is equivalent to The package will arrive next Wednesday or later because arriving next Wednesday or later than that are the two ways in which a package could arrive after next Tuesday

   h. The bridge will open at the end of May is not equivalent to The bridge will open before June since it is not now the end of May so the bridge could open before June by opening even earlier than the end of May

   i. The package will arrive before next Wednesday and The package will arrive after next Wednesday are not contradictory because both will be false if it arrives on next Wednesday

   j. The bridge will open before June and The bridge will open in June or later or never at all are contradictory because opening before June, opening in June, opening later than June, and not opening at all exhaust all possibilities and are mutually incompatible

2. a. $\varphi \vDash \varphi$ holds always because φ cannot fail to be true if it is true

b. $\varphi \vDash \top$ holds always because $\top$ cannot fail to be true no matter what φ is like

c. $\varphi \vDash \bot$ holds only when φ is absurd because, if there is any possibility of φ being true, there is a possibility of $\bot$ being false when φ is true

d. $\top \vDash \varphi$ holds only when φ is a tautology because if there is any possibility of φ being false, there is a possibility of it being false when $\top$ is true

e. $\bot \vDash \varphi$ holds always because there is no possibility of $\bot$ being true so no possibility of φ being false when $\bot$ is true

f. $\varphi \triangledown \varphi$ holds only when φ is a tautology because if there is any possibility of φ being false, it does not, together with itself exhaust all possibilities

g. $\varphi \triangledown \top$ holds always becuase $\top$ covers all possibilities by itself, so it certainly exhausts them when taken together with φ

h. $\varphi \triangledown \bot$ holds only when φ is a tautology becuase, since $\bot$ leaves open no possibilities, it contributes nothing to exhausting them all and φ must do that all by itself

i. $\varphi \vartriangle \varphi$ holds only when φ is absurd because, unless φ rules out all possibilities, there will be a possibility of it being true along with itself

j. $\varphi \vartriangle \top$ holds only when φ is absurd because, since $\top$ is bound to be true, any possibility of φ being true will be a possibility of both being true

k. $\varphi \vartriangle \bot$ holds always because, since $\bot$ cannot be true, it cannot be true together with any sentence (even itself)

l. $\varphi \simeq \varphi$ holds always since a sentence must have the same truth value as itself

m. $\varphi \simeq \top$ holds only when φ is a tautology because, if φ is bound to have the same truth value as a tautology, it must be one

n. $\varphi \simeq \bot$ holds only when φ is absurd because, if φ is bound to have the same truth value as an absurd sentence, it must be one

o. $\varphi \bowtie \varphi$ never holds because no sentence can be both true and false at the same time

p. $\varphi \bowtie \top$ holds only when φ is absurd because φ is bound to be false if its value is opposite that of a sentence that is bound to be true

q. $\varphi \bowtie \bot$ holds only when φ is a tautology because φ is bound to be true if its value is opposite that of a sentence that is bound to be false

3. The appearance of “—” in a cell in the table below indicates that nothing can be concluded in general about the relation between φ and χ.

	$\psi \vDash \chi$	$\chi \vDash \psi$	$\psi \simeq \chi$	$\psi \vartriangle \chi$	$\psi \triangledown \chi$	$\psi \mathbin{\text{⧖}} \chi$
$\varphi \vDash \psi$	$\varphi \vDash \chi$	—†	$\varphi \vDash \chi$	$\varphi \vartriangle \chi$	—†	$\varphi \vartriangle \chi$
$\psi \vDash \varphi$	—*	$\chi \vDash \varphi$	$\chi \vDash \varphi$	—*	$\varphi \triangledown \chi$	$\varphi \triangledown \chi$
$\varphi \simeq \psi$	$\varphi \vDash \chi$	$\chi \vDash \varphi$	$\varphi \simeq \chi$	$\varphi \vartriangle \chi$	$\varphi \triangledown \chi$	$\varphi \mathbin{\text{⧖}} \chi$
$\varphi \vartriangle \psi$	—*	$\varphi \vartriangle \chi$	$\varphi \vartriangle \chi$	—*	$\varphi \vDash \chi$	$\varphi \vDash \chi$
$\varphi \triangledown \psi$	$\varphi \triangledown \chi$	—†	$\varphi \triangledown \chi$	$\chi \vDash \varphi$	—†	$\chi \vDash \varphi$
$\varphi \mathbin{\text{⧖}} \psi$	$\varphi \triangledown \chi$	$\varphi \vartriangle \chi$	$\varphi \mathbin{\text{⧖}} \chi$	$\chi \vDash \varphi$	$\varphi \vDash \chi$	$\varphi \simeq \chi$

In cells marked with †, the fact that no relations hold in general can be seen by noting that, if ψ is a tautology, the given relations between it and φ and χ will hold no matter what sentences φ and χ are, so it is possible for φ and χ to be logically independent. And, in the cells marked with *, something similar holds in a case where ψ is absurd: the given relations between ψ and each of φ and χ will hold no matter what φ and χ are. There are various considerations which can be used to show that what is said in other cases is the most that can be said, but it is probably easiest just to confirm for yourself that no further truth values for φ and χ are ruled out by the given information about the relation of each to ψ.

Glen Helman 03 Aug 2010

1.3. Beyond saying: pragmatics

1.3.0. Overview

Our study of logic will be limited to deductive logic; and, even within those bounds, we will consider only the logical forms that are part of first-order logic. These limits imply some others that deserve consideration in their own right: although our study of deductive logic can be seen as the study of meaning, we will not study all aspects of meaning.

1.3.1. A model of language
One simple picture of language sees it as a device for conveying information by way of the proposition expressed by sentences.

1.3.2. Some complications
This simple picture of language is too simple in many respects, but four are especially important for our purposes. Each corresponds to a further way of conveying information.

1.3.3. Speech acts
Questions and commands do not appear to convey propositions, and even declarative sentences may play roles other than assertion.

1.3.4. Implicature
Communication often exploits the assumption that what a speaker says is not only true but satisfies certain other requirements.

1.3.5. Indexicality
When a sentence conveys a proposition, the proposition that is conveyed will usually depend on the context in which the sentence is used, and sentences are sometimes designed to convey information about his context.

1.3.6. Vagueness
The range of application of many terms will have fuzzy boundaries even in a given context, and sentences that apply them to things falling in this gray area may have no determinate truth value.

1.3.7. Presupposition
Another way of conveying information rests on the preconditions for a sentence to have a truth value at all.

1.3.1. A model of language

The idea of truth conditions or of a proposition suggests a simple picture of the way language works. According to this picture, each sentence has truth conditions that are determined by the semantic rules of the language. These truth conditions settle the truth value of the sentence in each possible world, something that is encapsulated in a proposition. The proposition expressed by a sentence is its meaning. The meaning of an expression smaller than a sentence is to be found in the contribution this expression makes to the propositions expressed by sentences containing it. From this point of view, the function of language is to convey propositions.

Just as the information content of a sentence is to be found by considering the range of possible worlds it rules out, the information that a person possesses is to be found by considering the possible worlds that he or she is able to rule out. The more you can rule out, the more information you have; and the kind of information you have is determined by the particular worlds you can rule out. This means that the sum total of your knowledge can be thought of as a proposition.

Anyone's aim in acquiring information could be described as an attempt to distinguish the actual state of the world among the various alternative possibilities—in short, to locate the actual world within the space of all possible worlds. The proposition representing your knowledge goes some distance towards in ruling out some possibilities. But it will leave many open, and the actual world could be any of those open possibilities. If someone conveys a proposition to you and you accept it, you are able to rule out a whole region of logical space, a region that can be added to the region ruled out by your existing knowledge. And, in general, this will reduce your uncertainty about the location in logical space of the actual world.

You can generate information to give to others by delimiting a region within the total area you know to be ruled out. Ideally, perhaps, you would simply convey the whole of what you know; but language limits your ability to do this since only a limited range of propositions are expressed by reasonably short sentences. To convey information, you select a sentence that is entailed by what you know and assert it, thereby conveying the proposition this sentence expresses.

This is process is illustrated in the following artificial example of sharing information.

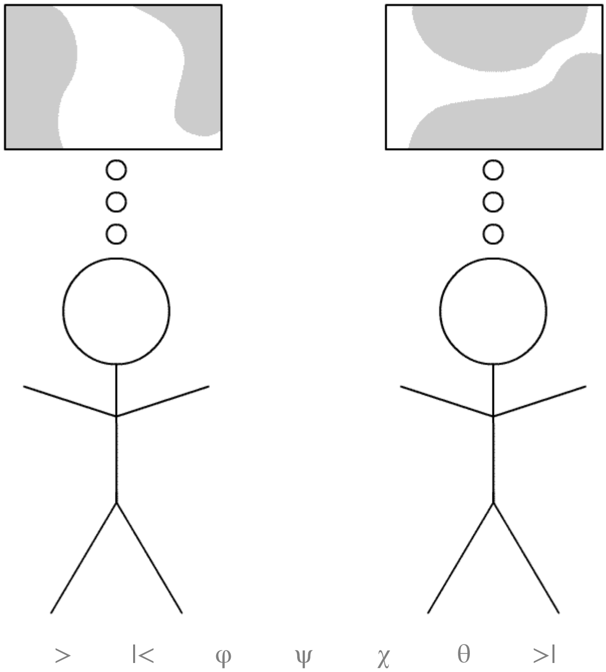

Fig. 1.3.1-1. An animation of a conversation in which information is shared. The button > will play the full conversation while the buttons φ, ψ, χ, and θ will each play one of its four stages. The buttons |< and >| move to the initial and final state, respectively.

Initially, the person on the left is able to rule out regions at the left and right of logical space as possibilities for the actual world while the person on the right is able to rule out regions at the top and bottom. The animation then shows a conversation in which each party in turn notices the truth of the one the sentences φ, ψ, χ, and θ and asserts it. The other person accepts this assertion as true and adds its content to the region ruled out by his or her beliefs. At the end of the conversation, the two people share the ability to rule out a region around the boundary of logical space though they still differ in the shape of the region left open in the middle.

In this conversation, each party is depicted as accepting what the other says as true and adding it to his or her own beliefs. The person accepting the assertion could be said to modify his or her beliefs in a way that makes it something he or she might assert. This is an example of a process that the philosopher David Lewis labeled *accommodation*. In this case of accommodation, one's beliefs are altered to accommodate an assertion someone else has made.

Of course, we do not always accept what others say—i.e., we do not always alter our beliefs to accommodate their assertions—for we may doubt that they are sincere or that they know what they are talking about. But this cannot be the ordinary case. Words can acquire and maintain a conventional meaning only if people usually mean what they say. And the act of asserting a sentence could not have the significance it does unless people were usually willing to accept assertions as well-founded. A critical attitude is important; but, at least practically, it must be the exception. Even when we are critical and ask for the grounds of someone's assertion, our request can be met only if we are at some point willing to accept assertions providing grounds as well-founded; and, when we are willing to do so, this will rarely be because there is no room for further doubt. In short, while we do not always accommodate what others say, accommodation is central to the aspects of language this model captures. We will also see that other forms of accommodation are essential to a number of the aspects of language that are not captured by the model of communication we have been considering.

There is one simplification in the picture above that is not an essential feature of the model depicted but is worth mentioning because it concerns an important use of entailment. Entailment appears in the picture in one way by setting bounds on the range of sentences that you can sincerely assert: if what you assert is to be something you believe, it must be entailed by your beliefs. But entailment also plays a role in your acceptance of what is asserted to you. Even when you do not doubt what has been asserted, you often add only some of its content to your beliefs. While, ideally, you might like to add the full content of what you hear to your beliefs, your ability to store information is limited, and what you do store is determined by your interests. And, if what you store is to be really part of what was asserted it must be implied by that assertion. That is, a fuller picture of the way a proposition is conveyed is the following:

Γ	$\vDash$	φ	$\vDash$	ψ
speaker’s beliefs		asserted sentence		proposition accepted

The first entailment turnstile marks one aspect of the process of determining what to assert (“invention” in the terminology of traditional rhetoric) while the second marks one aspect of the process of interpretation.

1.3.2. Some complications

Probably no one ever believed that the simplified model of language we have been considering was entirely accurate. But it, or something like it, was until recent decades the working model most logicians used for thinking about the function of language. Around the middle of the 20th century, philosophers became interested in a number of features of language that suggest this picture is inadequate; and these features have been incorporated into a number of richer models of language. The norms of deductive logic that we will study do not rest on the richer structure of these new models, so we will not consider them in detail. But some of the further features of language that they attempt to capture are intertwined with those we will study, so we need to take some time now to disentangle ourselves from a few of these features once and for all and to lay the groundwork for disentangling ourselves from others at later points in the course.

The complicating phenomena that we need to consider have come to be studied under the rubric of *pragmatics*. This term was originally introduced (by Charles Morris) as an alternative to *semantics* in order to distinguish issues concerning the relation between language and its users from the issues concerning the relation between language and what is spoken of. The use of the term pragmatics is no longer closely tied to this definition, and I know of no definition that really captures the way it is now used. Probably the best way to understand current usage is to consider some commonly agreed examples of pragmatic phenomena. The following ones are the most important for our purposes.

1) *Sentences are not always used to express propositions*. When a sentence is used to express a proposition, the question of its truth value is a significant one. But not all sentences have truth values or raise questions of truth value. And even when a sentence does have a truth value, its truth value may not be its most important feature. There are many ways of using sentences, many *speech acts*, besides assertion, and the way a sentence is used is one aspect of its meaning. The term *force* is often used for this aspect of meaning.

2) *The information we derive from the use of sentences is not limited to what follows from accommodating them as true*. Assertions can be expected to have properties other than truth, and there can be forms of accommodation associated with these other properties. In particular, the assumption that an assertion has a given property can be the basis for deriving information from the assertion. This produces the phenomena of *implicature*, in which a sentence suggests more than it says. Even when everything a sentence literally says is

true, an additional false suggestion can make it misleading.

These two complications suggest that propositions are not quite as central to the use of language as the simple model suggests: sentences do not serve merely to convey the propositions they express. Several further complications concern the relation between language itself and propositions: saying simply that sentences express propositions is at best a rough approximation to their meaning.

3) *The proposition expressed by a sentence (and thus its truth value) may vary with the context in which it is used*. For example, there is no way to judge the truth value of a sentence like I put that here yesterday when it is taken out of context. This dependence on context is due to various phenomena known collectively as *indexicality* or *deixis*. Both terms are etymologically related to terms for pointing, and the functions of words this and that are paradigm examples. The term *character* has been used for the way the proposition expressed depends on the context.

4) *Even with regard to a given context, a sentence may not have a definite truth value*. The meaning of *vague* terms like small and hot will vary with the context; and even in a given context there will be no sharp delineation of the cases where they apply truly. We can continue to speak of the character of a sentence containing such terms but only if we allow the proposition expressed to be depend on factors that are not fully determined by actual contexts of use.

5) *Sentences may have truth values in some possible worlds and not in others*. There can be preconditions for a sentence to have a truth value at all. Anything implied by these preconditions counts as a *(semantic) presupposition* of the sentence, and it constitutes another way in which information can be derived from it.

The force, implicatures, character, and presuppositions of a sentence are parts of its meaning in the fullest sense of the term. We will consider each at least briefly to distinguish it from the narrower sense of meaning that will be our focus. It is easy to disentangle our topic from some of these phenomena but others require more detailed consideration, and some forms of entanglement are more likely to trip us up than others. As a result we will consider some of these sorts of meaning only to dismiss them quickly, and we will set others aside without completing disentangling ourselves from them. Implicature is the only one of these aspects of meaning that we will need to pay much attention to in later parts of the course.

1.3.3. Speech acts

Although we have been speaking of sentences as if they all had truth values, there are some sentences that not only do not have truth values but cannot have them. It would be crazy to respond to a question like What time is it? by saying True enough or You're wrong! And these responses would be equally out of place in the case of an imperative sentence like Please shut the door.

Questions and imperatives are clear cases of sentences where truth values are irrelevant. But truth values may be beside the point in the case of some declarative sentences, too. Saying True enough or You're wrong! would be out of place in response to a sentence like I promise to be here tomorrow or I apologize for what I said, but the reasons they would be out of place are different here than in the case of questions and commands. The verbs promise and apologize can be used to describe certain sorts of actions that can be performed in using language; that is, they express *speech acts*. And, when they are used in the first person present tense (as in the sentences above) by the right person under the right circumstances, they can be used to perform the sort of actions they describe. That is, by saying I apologize for what I said, I can do something that can be described truly by the sentence He apologized for what he said; that is, given the right circumstances, I apologize simply by saying I do. Verbs that may be used in this way to perform the actions they describe were labeled *performative* by J. L. Austin, the philosopher who did the most to call attention to the variety of speech acts. When I use a performative verb correctly, what I say is true; but the fact that it is true is not very interesting because my saying it is what made it true.

Austin estimated that the performative verbs in English number "on the order of the third power of 10." If this estimate is accurate, there are thousands of kinds of speech act besides assertion and thousands of varieties of force beyond the sort of force we will focus on. Of course, much of this vocabulary marks only subtle differences of force between speech acts, but the fact that we have vocabulary for making such subtle distinctions indicates how important it is to us to know the specific force of an utterance. Moreover, we need not use performative verbs to perform the acts that these verbs describe. I can apologize without saying I apologize and I can make a promise without saying I promise. So we can expect that, even when we use declarative sentences, many, and perhaps most, of things we say are not simply assertions. The statement I will be there might be a simple assertion predicting the speaker's future location, but it will often (perhaps most often) be a promise.

In spite of this, we will not consider speech acts other than assertion, and our interest in assertion itself will be limited to one aspect of its force: the expression of a proposition. Although this will cut us off from much of the richness of language, it will not cut us off from much that is central to deductive reasoning. Of course, there is a sense in which conclusions can be drawn from apologies and promises, but such inferences will tend to be matched by conclusions drawn from ordinary assertions using performative verbs to describe apologies and promises (rather than make them). Moreover, many accounts of speech acts generally treat propositions as important components of their meaning, and this gives the study of assertions a central place in the study of all speech acts.

1.3.4. Implicature

As we have been using the term imply, a sentence implies anything whose content is included in the proposition it expresses. Thus we can say that the sentence My class was taught this morning implies A class was taught. The philosopher H. Paul Grice employed the term implicates to capture a different idea that is sometimes expressed by the ordinary use of the term implies.

It is not uncommon for information to be suggested by a sentence even though it is not entailed and thus is not part of what the sentence literally says. For example, my assertion of the sentence My class was taught this morning would, in most contexts, suggest that I did not teach the class myself. However, this is not part of what I said since my statement would be perfectly true if I taught the class, so My class was taught this morning implicates I did not teach my class this morning but does not imply it.

The contrasting vocabulary say and suggest was used in passing in the previous paragraph, and it is a convenient way of expressing the difference between implications and implicatures. Still, it makes a difference how the term suggest is understood. In particular, it is not intended in this use of it to convey the idea of subjective association. What a sentence implicates can be as much the product of rules of language as what it implies. The difference between the two lies in the fact that the rules leading to implicature are not (or are not only) rules assigning truth conditions.

To see what sort of rules they might be, let us consider an extension of our simple model of language use that incorporates implicature; in its outlines, it is due to Grice. To account for implicature, we extend the scope of accommodation to include not only the truth of assertions but also certain other features assertions ought to have. The maxim *Speak the truth!* is no doubt the key rule governing assertions, but other maxims, such as *Be informative!* and *Be relevant!*, also play a role. Someone who assumed I was obeying all maxims of this sort when I said, “My class was taught this morning,” might reason as follows:

> Although Helman’s assertion My class was taught this morning would have been perfectly true if he had taught his class, it would have been a strange thing to say in that case because the proposition expressed by I taught my class this morning would have contained more relevant information. So I can best accommodate his use of language if I assume he did not teach the class.

Let us say that an assertion is *appropriate* when it is in accord with all maxims governing language use and otherwise say that it is *inappropriate*. An assertion could be inappropriate even though true, so we go further when we assume it is appropriate. At that is something we usually do; that is, we usually accommodate our beliefs about the world to the assumption that the assertions others make are not only true but appropriate for the context in which they are made.

These ideas can be used to state contrasting definitions for implication and implicature. First let's restate our definition of implication in a way that will make the comparison easier:

> φ implies ψ (in a given context) if and only if φ cannot be true (in that context) when ψ is false (in that context).

To define implicature, we follow the same pattern using the concept of appropriateness instead of truth.

> φ implicates ψ (in a given context) if and only if φ cannot be appropriate (in that context) when ψ is false (in that context).

That is, while implications are conditions necessary for truth, implicatures are conditions necessary for appropriateness. (Notice that the term implicature is used here both for the things a sentence implicates and for the relation between a sentence and what it implicates. Our use of the term implication follows the same pattern.)

One aspect of the relation between implication and implicature depends on whether we understand truth itself to be one of the requirements of appropriateness. It is convenient to understand appropriateness to include truth because anything that is implied is then also implicated and implicature is a broader relation than implication. However, there is no consensus about using the terms in this way, and many would use *implicature* more narrowly to cover only those conditions necessary for appropriateness over and above those necessary for truth.

Both definitions above refer to the context in which sentences are used. We have ignored this so far in the case of implication though the phenomenon of indexicality means that such a reference is often required. In any case, it is crucial for appropriateness: while the contextual dependence of truth values is tied to specific vocabulary, appropriateness in the wider sense is always dependent on the specific context in which a sentence is used. In the example used above, if it was well known that I had made a bet that I could avoid using the word I for the next 24 hours, no one would be misled by my saying My class was taught this morning when I had in fact taught it myself.

Even though appropriateness as a whole depends on the context, there are specific conditions attached to particular words that can lead to implicatures in every context. Consider, for example, this bit of dialogue:

> Q: Was the movie any good?
> A: Yes. Even John was laughing.

The assertion Even John was laughing has a number of implicatures that depend on the conversational setting (e.g., that John was at the movie and, perhaps, that it was a comedy), but it also has one that derives from presence of the word even. This implicature is easier to recognize than to state, but it comes to something like the claim that John doesn't laugh frequently.

Implicature is a form of non-deductive inference that we will not study in its own right, but we will not be able to ignore it because it is often difficult to distinguish from implication. This is especially true for implicatures that attach to particular words because they have the same sort of uniformity across contexts that holds for the sorts of implications we will study.

One test that can be used to distinguish implicatures from implications is to ask a *yes-no question*. When asked Was even X laughing? about someone *X* who had laughed at the movie but who was known to laugh frequently, we would not answer with a simple "No" but rather say something like, "Yes, but he'll laugh at anything." Such yes-but answers indicate that the sentence we were asked about is true but inappropriate. Other qualified affirmative answers can play a similar role, and we will refer to them also as *yes-but answers* even when they do not use the term but. To simply answer "Yes" in cases where a sentence is true but has a false implicature could mislead our audience into thinking that the sentence is entirely appropriate and thus that the implicature is true. Indeed, a true sentence with a false implicature could be described as true but *misleading*. Yes-but answers acknowledge the truth of such a sentence while correcting its misleading suggestions. (There are further tests that can be used to distinguish implicatures and implications, and we will consider some others in 4.1.2.)

1.3.5. Indexicality

We will give less direct attention indexicality than to implicature, but it would be hard to ignore the phenomenon. Although indexicality is most obvious in sentences with indexical words like I, that, here, and yesterday, there are other features of a sentence, most notably its tense, that can make the proposition it expresses vary with context in which it is asserted. The sentence It's sunny is as bound to the time of assertion as is It's sunny now. And, while not every sentence contains indexical terms, it is only very special sentences that are not indexical in virtue of tense.

If the propositions expressed by sentences vary with the context, it seems that the logical properties and relations of these sentences (which we trace to the propositions they express) may vary as well. Let's look at one example. The proposition expressed by the sentence I am here will depend on the speaker, the speaker's location, and the time of utterance. And this sentence may express the same proposition as the sentence You are there when the latter is used by a second speaker in an appropriately related context. There are also many contexts in which these sentences might be asserted where they would not express the same proposition. But sentences are supposed to be logically equivalent when they express the same proposition, so it seems these sentences would be equivalent when used in some contexts and not equivalent when used in others. And the same issue arises for deductive properties as well as relations; a sentence that is a tautology when used in one context might not be a tautology when used in a different context.

More broadly it may seem that we really should not speak of sentences as having deductive properties and standing in deductive relations. If a sentence expresses no fixed proposition independent of the context in which it is asserted, we can really only talk about the deductive properties and relations of sentences-in-context, of sentences each taken together with a context of use. The term *statement* has sometimes been used to speak of a particular use of a sentence. If we use this terminology, we can say that certain statements made using the sentences I am here and You are there are equivalent and that it statements rather than sentences have deductive properties and stand in deductive relations. Something like this approach would be required if we really were to study the phenomenon of indexicality. However, the logical forms on which we will focus do not include indexical elements, so it will be possible for us to ignore this aspect of meaning.

Even when indexical elements are present, we can set aside explicit reference to contexts of use when speaking only of logical properties and relations that do not vary from context to context. For such deductive properties and relations will hold of sentences in virtue of the specific ways the propositions they express vary with the context of use—i.e., in virtue of the "characters" of these sentences. For example, we can say that sentences are equivalent if their characters lead them to express the same proposition in any context of use, and we can say that a sentence is a tautology if its character leads it to express a tautologous proposition in every context of use. Again, although the propositions expressed by The package will arrive next Wednesday and The package will arrive next week will very depending on the time of utterance, the proposition expressed by the first sentence will always entail the one expressed by the second sentence. We will limit consideration to logical properties and relations of sentences that are independent of the context of use in this way. So, even though I am here and You are there may be used to make statements that are equivalent, we will not count these sentences as equivalent because it is not the case that, in each context, the propositions expressed by these sentences are the same. (Indeed, it is not easy to think of any single context with respect to which the two would express the same proposition since a single context would require that both be spoken by the same person.)

In fact, we can use this approach without explicitly considering the characters of sentences at all. In fact, this was done in the example in 1.2.3 that included the sentences The package will arrive next Wednesday and The package will arrive next week. There we simply took it for granted that sentences were being compared with respect to some one context, and we spoke freely of the propositions they expressed in that context without bothering to note that they expressed different propositions in other contexts. This procedure is legitimate if we not to assume anything special about the context of use. And it will be easy not to make special assumptions about the context of use because the deductive properties and relations we are interested in do not depend on this context. There is an analogy here to a typical use of variables in algebra. When numerical laws are used to manipulate algebraic formulas, it is assumed that variables appearing in those formulas have been assigned numerical values. But there is often no need to consider what those values are since the laws being used apply to all numbers.

Of course, there are things we will miss by ignoring character and context. The effects of shifting context in the course of a conversation are among the things we cannot deal with. The assertion I am here followed by the confirmation Oh, so that's where you are is a simple example of this. Another phenomenon we will miss is the exploitation of some sort of dependence on context to convery information about the context. If I assert Today is Tuesday, the proposition expressed may be no more informative than is Tuesday is Tuesday since the first sentence, if true, merely tells us about Tuesday that it is Tuesday. But my assertion can still be helpful because someone who tries to accommodate it will need to take it to have been asserted on Tuesday, and will thus know what day it is. In short, even if the proposition expressed by Today is Tuesday in a given context is a tautology and conveys no information, the assumption that this sentence expresses a tautology (rather than an absurdity) in that context yields information about the context. And this way of deriving information can support a form of non-deductive inference.

On the other hand, our approach need blind us to all logical properties and relations that derive from indexical terms. We have seen this already in the case of next Wednesday and next week, but the role of the indexical terms can be less trivial than this. For example, the terms today and tomorrow are related in such a way that Tomorrow is the day after today is true in any context, so we can recognize it as a tautology. And we can also recognize that Today is Tuesday implies Tomorrow is Wednesday.

It would be too much to say, however, that our limited perspective will not blind us to any logical properties or relations that hold for all contexts of use. For there are relations between the meanings of indexical terms that hold in any context, but only with respect to the actual world of that context; and our approach will miss logical relations that derive from these aspects of meaning. For example, whoever is the speaker in a context will actually be speaking at the time of utterance, so the premise Today is Tuesday would justify the conclusion I am speaking on Tuesday. But this conclusion is not entailed by the premise—even given the contextually assigned meanings of the terms—since nothing about the day of the week of a given date logically necessitates someone speaking. To get a feel for the issue, it may help to look at a related example: although I am here now is true in the actual world of any context, it is not a tautology. That is, the proposition expressed by I am here now in a context of utterance is bound to be true in the actual world of that context, but this proposition will also be false in other possible worlds. And the fact that it is false in other worlds can be crucial for the meaning of sentences—such as I am here now but I almost didn't make it—that speak of unactualized possibilities.

1.3.6. Vagueness

One way of understanding vague terms is to suppose that their significance varies with the context of use but is not completely determined by it. The meaning of a word like small depends on the line to be drawn between what is and what is not small. This line is settled to some degree by features of the context of its use—whether the word appears in a discussion of molecules or of galaxies, for example—and some contexts will pin it down more precisely than others. But there is usually, and perhaps always, some indeterminacy remaining, and the class of things that count as small in a given context will have fuzzy edges.

Although the context dependence of vague terms means that vagueness is somewhat analogous to indexicality, the fact that sentences containing vague terms may not have definite truth values even when the context is specified means that we cannot handle such sentences in quite the same way as we do sentences exhibiting ordinary forms of indexicality. We can understand entailments involving indexical terms—such as

Today is Tuesday ⊨ Tomorrow is Wednesday

—to hold because the propositions expressed by the two sentences are related in a certain way in every context of use. But we cannot understand the entailment

Crawfordsville is small ⊨ Crawfordsville is not large

to hold for the same reason because the sentences involved may not express definite propositions in any context of use.

Still, there is a way of extending our approach to indexicality to provide an approach to vagueness. In both cases we can understand deductive properties and relations to hold for sentences because of the propositions that *would be* expressed by the sentences if certain factors were specified. In the case of the first example above, the relevant factor, the time of utterance, is specified by any actual context of use. In the second example, the relevant factors are precise *delineations* of the classes of things that the terms small and large are true of. These delineations are not fully determined by an actual context of use, but we can still say that the propositions expressed by the sentences in the second example would represent a case of entailment no matter how these delineations were specified. So, just as we will always take for granted an unspecified context of use, we will take for granted but leave unspecified precise delineations of all vague terms. And that means that we will speak of sentences as if no terms are vague.

Of course, ignoring vagueness means that we will ignore yet another important feature of language. The specific logical properties and relations we will study do not derive from vagueness, so ignoring vagueness will not limit our ability to study them. But, as with implicature and indexicality, we will miss certain ways of deriving information from things that are said. The accommodation of vague language can be analogous to accommodation of indexicality and can be an important way of conveying information. While This is hot will often be intended to provide information about whatever this refers to, it can serve instead to calibrate judgments of hotness. That is, when the audience already knows the temperature of the thing pointed to, This is hot can help someone to specify the significance of hot in a given context since accommodating this assertion requires that the thing pointed to falls within (and, indeed, some distance within) the range of hot things on any delineation of that range that is allowed by the context.

The fact that we derive information in this way provides one way of explaining a traditional logical puzzle known as the *sorites paradox* (or "paradox of the heap," after a particular ancient example trading on the vagueness of the term heap). The argument

This is hot and that is only a little cooler / That is hot

is not deductively valid because the things refered to by this and that could well fall on opposite sides of a delineation. But it seems like a reasonable argument; and, if we suppose that we accommodate vague language by considering only delineations on which what has been said is not just barely true, the conclusion will be true on any delineation that accommodates the premise. The paradox comes by imagining a series of things, with each successive thing asserted to be only a little cooler than the one before with the last clearly not hot. Each step in the series could be justified by an argument like the one above, but the final result seems unacceptable.

This result would not be surprising if we understand the displayed argument to be the result of accomodation. Suppose first that we attempted to collect all the steps in the series into a single argument. *

A is hot
B is only a little cooler than A
C is only a little cooler than B
⋮
Z is only a little cooler than Y
――――――――――――――――――――
Z is hot

This would not be reasonable because accommodating the first premise need not place the temperature assigned to A far enough from allowable delineations to support the truth of the conclusion.

On the other hand suppose we were faced with a series of arguments

α is hot
β is only a little cooler than α
――――――――――――――――――――
β is hot

one for each successive pair of terms in the series. If we really were willing to accommodate the premise at each stage, we would end up accepting the final conclusion; but the allowed delineations of hot would have shifted also at each stage and the final conclusion would end up acceptable.

Of course, someone who really refused to accept the final conclusion would probably refuse to accommodate the premise of one of the arguments along the way and would begin to be wary of them before that point. That is, these component arguments each stretch our willingness to accommodate a bit further, and it can only be stretched so far. The paradoxical inference can seem to be supported if we forget this, and think of the corresponding way of extracting information from an assertion as if it was like deductive inference in allowing us to link together inferences that are good individually. That is, the sorites paradox shows us that the non-deductive relation associated with this way of deriving information from the use of vague terms is not transitive.

There is terminological curiosity here. An argument like the one above running from A to Z—i.e., a multiple-premise argument that is associated with a series of two-premise arguments—is traditionally referred to as a *sorites argument*. But a sorites argument need have no connection with a sorites paradox. Although the term sorites is derived in both cases from the Greek term for a heap, its application to a sorites argument reflects the piling up of premises rather than any appearance in it of a vague term such as heap. A sorites argument constructed for the sorites paradox in its original form would be an argument about heaps that had a heap of premises.

1.3.7. Presupposition

When the yes answer to a yes-no question would be tantamount to making a true but misleading assertion, it is appropriate to answer yes only if we add a qualification. But it is still possible to give an affirmative answer while no qualification would make the answer no appropriate. Another of the complications of the simple picture of language appears in connection with yes-no questions for which neither answer seems legitimate.

For example, consider the question

Is John's car green?

asked about someone who does not have a car at all. In such a case, we would be at a loss to answer the question directly. This is usually explained by saying that the question *presupposes* that John has a car and has no appropriate direct answer when this presupposition does not hold. And we can say something similar about the following declarative sentences, which correspond to affirmative and negative answers to the question, respectively:

John's car is green
John's car is not green

That is, just like the question, we can take each of these assertions to presuppose John has a car.

We could capture these limits on appropriateness by regarding presupposition as a sort of implicature. That is, we might say that John having a car constitutes a necessary condition for the appropriateness of either of the assertions above. But many have held that in contexts where John has no car, it is not only the case that neither sentence is appropriate but the case that neither is true. Since one would be true if the other was false, this means that neither claim would have a truth value. If this point of view is correct, what is missing in these assertions when John has no car is not some quality like informativeness or relevance that we expect in addition to truth but instead something that is a precondition for either truth or falsity. Something that is a presupposition in this strong sense is said to be a *semantic* presupposition. If John having a car is a semantic presupposition of the two sentences above, it is easy to see why they seem equally inappropriate when John has no car: each would have no truth value so the two would be in the same position as regards truth and falsity.

Semantic presupposition is unlike the phenomena we have considered so far in that it requires fundamental changes to the simple model of language and not merely additions to it. The simple model is built around the assumption that a sentence has a truth value in every possible world, and dropping that assumption would force radical changes. And because there is no consensus, even among logicians who accept the idea of semantic presuppositions, about the exact form such changes should take, we will not attempt to incorporate failures of truth value in our model of language.

In part, we will treat semantic presupposition as we do the variety of speech acts: by not considering the examples where it may be held to occur. But we cannot avoid all the difficult cases in this way. The classic examples of semantic presupposition are sentences containing phrases employing the definite article the to refer to something by way of a description of it. Such phrases, which logicians classify as *definite descriptions*, cause problems because their success in referring depends on the existence of objects satisfying the descriptions they offer. For example, both the sentence The building between Center Hall and Sparks Center is occupied and the sentence The building between Center Hall and Sparks Center is unoccupied seem inappropriate when no such building exists because then the definite description the building between Center Hall and Sparks Center has nothing to refer to. And definite descriptions that refer contingently are so common that we cannot simply avoid all sentences containing them. The use of possessives that we saw in the example of John's car are also common, and they represent a closely related sort of case because John's car might be paraphrased by the definite description the car John has.

The approach we will take to these sorts of semantic presupposition does share two features with our approaches to other complicating phenomena. First, just as we do not attempt to capture relations of implicature in our study of logic, we will not attempt to capture relations of presupposition as such. However, the line between implication and presupposition is controversial, and relations between sentences like The building between Center Hall and Sparks Center is occupied and There is a building between Center Hall and Sparks Center fall in the disputed area. In 8.4.2 we will consider an account of definite descriptions according to which the first of these sentences implies the second.

Although we will not attempt to capture semantic relations of presupposition as such, we will need to apply our general account of logical properties and relations to sentences that may have such presuppositions. And we can do this only if we do not recognize the failures of truth value that result when semantic presuppositions are false, so we will assume that every sentence has a truth value under all possibilities. But, since we will eventually analyze sentences into units smaller than sentences, an assumption about the meanings of sentences is not enough.

We will assume in addition that any term which ought to refer does have a *reference value*. We allow this to be either an actual object or an *empty* or *nil* reference value. The latter option is designed for the case of *undefined* terms like the building between Center Hall and Sparks Center that do not refer to actual objects. We will need to distinguish these two sorts of reference value only when we consider definite descriptions in the last chapter, so, for the most part, we will merely assume the every term has been somehow given a reference value and every sentence a truth value. The references and truth values we assume for this reason can be regarded as stipulations added to the conventional meanings of these expressions, and we will consider only logical properties and relations that hold no matter how such stipulations are made. Such assignments of supplementary semantic values are usually called *supervaluations*. Both the name and this way of handling failure of presuppositions are due to Bas van Fraassen, and the assignment of precise delineations to vague terms that was discussed in the last subsection is a further application of this idea by David Lewis. As will be case in our handling of vagueness, our assumptions of references and truth values in cases of semantic presupposition will generally stay in the background. However, we will look at the assumptions we make a little more closely in 6.1.3 when we have begun to analyze sentences into expressions that are not sentences.

1.3.s. Summary

1 The idea that the norms of deductive reasoning reflect a system of relations among propositions fits into a simplified picture of the function of language. According to this picture, a person's beliefs amount to a proposition that rules out a certain range of possibilities for the actual history of the universe. The desire to know more is in part the desire to narrow the range of possibilities that are left open. When language is used cooperatively, we share our abilities to rule out possibilities by using assertions to convey propositions. The sentences we can sincerely assert are the ones that are entailed by the sum total of our beliefs, and we accommodate someone else's assertion by adjusting our beliefs so that what they asserted is now entailed by our beliefs.

2 This picture is oversimplified and something must be said about several respects in which the actual operation of language is more complex. Each is associated with an aspect of meaning:

(i) the force of a sentence that marks it as an assertion or one of the many other speech acts,
(ii) implicatures, which convey information that a sentence does not imply,
(iii) semantic presuppositions, requirements for the sentence to have a truth value,
(iv) the character of a sentence, which reflects the way the proposition it expresses varies with the context of use due to the phenomenon of indexicality, and
(v) a greater or lesser degree of vagueness.

While an account of how sentences express propositions is the province of semantics, these complicating phenomenon are usually said to be the subject matter of pragmatics.

3 Although assertion is the only speech act we will study, not even all declarative sentences have this force. J. L. Ausin estimated that assert was only one of thousands of performative verbs that can be used to both perform and describe speech acts. Although many of these speech acts do not serve to convey propositions, their force can often be described with reference to propositions.

4 We will consider only what is implied by a sentence as part of its truth conditions and not further information that may be implicated as conditions for appropriate assertion beyond the requirements for truth. A false implicature will make a sentence misleading but may leave it true. One

indication of this sort of case is a yes-but answer to the yes-no question corresponding to the sentence.

5 Indexicality means that the propositions expressed by sentences—and thus their deductive properties and relations—can depend on the contexts in which they are used. It would be possible to compare sentences only when each was associated with a specified (but perhaps different) context—such sentences-in-context are sometimes called statements. However, we will compare sentences only within a single context of use and consider only properties and relations of sentences that hold no matter what that context is. As with implicature and presupposition, accommodating sentences to the rules governing indexical phenomena provides a way of extracting information that goes beyond entailment.

6 Vagueness poses problems analogous to those posed by indexicality and presupposition. As with indexicality, we will assume a context of use; and, as with presupposition, we will assume supplementary specifications of truth value (in this case precise delineations of the boundaries of vague terms). Deductively valid conclusions will not rely on information about these factors, but accommodation to vague assertions can support non-deductive inference to extract further information. One way of explaining the sorites paradox is to suppose that it rests on a failure to distinguish this sort of inference from deductive inference.

7 Since a semantic presupposition is something that must hold in order for a sentence to have a truth value at all, sentences with non-tautologous presuppositions can fail to have truth values. The pervasiveness of definite descriptions—which can fail to refer to anything if the facts are not right—makes it hard to simply ignore sentences with non-trivial presuppositions. Instead, we will treat all terms as if they refer, simply stipulating reference values and truth values in other cases (eventually distinguishing an empty reference value) but considering only relations between sentences that hold for all such stipulations (the method of supervaluations).

1.3.x. Exercise questions

1. For each of the following sentences, give a sentence it implies and a sentence it implicates (but does not imply) in the context described:

a. My plate is clean, as reported by a small boy who has been told to finish his vegetables by a parent saying, “Clean your plate.”

b. There is a cooler in the trunk, said in reply to someone’s expressed wish to have a beer.

c. I saw the director’s last movie, said in reply to someone who asked whether the speaker has seen a certain new movie.

2. Many philosophers would argue that the sentence I’m Adam, when true, expresses the same proposition as Adam is Adam; that is, if it is true at all, it is true in every logically possible world. The phenomenon of indexicality or deixis can help to explain how I’m Adam could be informative even if these philosophers are correct and it expresses a tautology when it is true. To see how this might work, ask yourself what information can be derived about a context of utterance by accommodating the use in this context of the sentence I’m Adam.

3. J. L. Austin, the philosopher who made people aware of the variety and importance of speech acts, suggested a way of identifying them. Look for verbs that can fit in the context I hereby … (e.g., I hereby assert that … or I hereby apologize). That is, look for, verbs that (in grammarians’ jargon) can be used in “first person indicative active sentences in the simple present tense” along with the adverb hereby. These are the “performative verbs” mentioned in 1.3.3. Austin suggested that there are such verbs for most speech acts. Find half a dozen as varied in character as possible.

1.3.xa. Exercise answers

1. The following are perhaps the most likely answers though they are not the only correct ones:

a. implies: No vegetables are on the boy's plate
implicates: The boy has finished his vegetables

b. implies: The trunk is not empty
implicates: There is beer in the cooler

c. implies: The speaker has seen a movie by the director in question.
implicates: The speaker has not seen the new movie [with further implicatures depending on the tone of voice]

2. The truth value of I'm Adam depends on features of the context in which it is uttered—specifically, on the identity of the speaker. So, it is not true in some contexts of utterance. And that means that, if we assume it is used correctly, it can tell us something about the context—who the speaker is. We derive this information not simply by assuming that the actual world is a world in which the sentence is true but by assuming, more specifically, that the sentence has been uttered in a context that makes it express a true proposition. And even if it tells us nothing about the actual world to know that the person Adam is himself, it does tell us something about the context to know that the person Adam is the speaker.

3. If Austin was right, thousands of answers are possible. I will simply note a five-fold classification of speech acts along with examples of performative verbs for each sort of act. (This classification is due to the philosopher John Searle but based on Austin's ideas.) (1) *representatives* (e.g., assert and conclude) commit the speaker to the truth of something. (2) *directives* (e.g., order and ask) are attempts to get the speaker's audience to do something. (3) *commissives* (e.g., promise and threaten) commit the speaker to some future action. (4) *expressives* (e.g., apologize and congratulate) express a psychological state. (5) *declarations* (e.g., sentence and promote) effect a change in an institution.

1.4. General principles of deductive reasoning

1.4.0. Overview

All the deductive properties and relations of sentences can be seen as special cases of a single relation. We will look at this relation and also see how to study the full range of deductive logic by way of entailment and a couple of auxiliary ideas.

1.4.1. A closer look at entailment
Entailment will be at the heart of our study and we will begin by looking in some detail at a couple ways of formulating its definition.

1.4.2. Division
It will be useful to have a special term for the kind of pattern of truth values that entailment rules out.

1.4.3. Conditional exhaustiveness
Although entailment does not encompass all the concepts of deductive logic, there is a similarly defined relation that does.

1.4.4. A general framework
All the deductive properties and relations we will consider can be expressed in terms of conditional exhaustiveness and expressed in a way that corresponds directly to definitions of them.

1.4.5. Reduction to entailment
Although conditional exhaustiveness provides a way of thinking about deductive properties and relations, entailment is way that they are most naturally established, and we need to consider how this can be done.

1.4.6. Laws for entailment
The ideas behind the reflexivity and transitivity of implication provide the core of the general principles that hold for the more general relations of conditional exhaustiveness and entailment.

1.4.7. Duality
The specific principles concerning $\top$ and $\bot$ display a kind of symmetry that we will also find in principles for other logical forms.

1.4.1. A closer look at entailment

Entailment was introduced in 1.1.6 somewhat informally as a relation between premises and a conclusion that merely extracts information from them and thus brings no risk of new error. Another way of putting the latter point is that a relation of entailment provides a conditional guarantee of the truth of the conclusion: it must be true if the premises are all true.

The discussion of entailment in 1.2.1 developed the resources necessary to give a more formal general definition. In fact it is useful to have in mind two equivalent ways of stating one.

$\Gamma \vDash \varphi$	if and only if	there is no logically possible world in which φ is false while all members of Γ are true
	if and only if	φ is true in every logically possible world in which all members of Γ are true

These are not two different concepts of entailment, for the two statements to the right of if and only if say the same thing. Still, they provide different perspectives on the concept. The second—which we will speak of as the *positive form* of the definition—is closely tied to the idea of a conditional guarantee of truth and to the reason why entailment is valuable. The first form—the *negative form*—makes the content of the concept especially clear, and this form of definition will generally be the more useful when we try to prove things concerning entailment. The other deductive properties and relations we have discussed or will go on to discuss can be given analogous pairs of definitions, a negative form ruling out certain patterns of truth values and another form stating a more positive generalization.

The equivalence of the two forms of the definition reflects a feature of all generalizations. When a generalization is false, it is because of a *counterexample*, something that is the sort of thing about which we generalize but that does not have the property we have said that all such things have. A counterexample to the claim that all birds fly is a bird that does not fly. In the positive definition of entailment, the generalization is about all possible worlds in which the premises are all true and such worlds are said to all have the property that the conclusion is true in them. A counterexample to such a generalization is then a world in which the premises are all true but the conclusion is not. The negative form of the definition then affirms the same generalization but by saying that no counterexample exists. As in the case of the generalization use to define entailment, one good way to clarify a generalization is always to ask what sort of counterexample is being ruled out.

It is important to notice how little a claim of entailment says about the actual truth values of the premises and conclusion of an argument. We can distinguish four patterns of truth values that the premises and conclusion could exhibit. Of these, a claim that an argument is valid rules out only the one appearing at the far right of Figure 1.4.1-1.

	Patterns admitted			ruled out
Premises	all **T**	not all **T**	not all **T**	all **T**
Conclusion	**T**	**T**	**F**	**F**

Fig. 1.4.1-1. Patterns of truth values admitted and ruled out by entailment.

So, knowing that an argument is valid tells us about actual truth values only that we do not find the conclusion actually false when the premises are all actually true. The other three patterns all appear in the actual truth values of some valid arguments (though not all are possible for certain valid arguments because other deductive properties and relations of the sentences involved may rule them out).

To see examples of this, consider an argument of the simple sort we will focus on in the next chapter:

It's hot and sunny
It's humid but windy

It's hot and humid

This argument is clearly valid since its conclusion merely combines two items of information each of which is extracted from one of the premises. Depending on the state of the weather, the premises may be both true, both false, or one true and the other false; and, in any case where they are not both true the conclusion can be either true or false. In particular, if it's hot and humid but neither sunny nor windy, the conclusion will be true even though both premises are false. This should not be surprising: a false sentence can still contain some true information, so information extracted from a pair of sentences that are not both true might be either true or false.

Of course, seeing one of these permitted patterns does not tell us that the argument is valid; no information that is limited to actual truth values can do that because validity concerns all possible worlds, not just the actual one. In particular, having true premises and a true conclusion does not make an argument valid. For example, the following argument is not valid:

Indianapolis is the capital of Indiana

Springfield is the capital of Illinois

For, although the single premise and the conclusion are both true, there is a logical possibility of the capital of Illinois being different while that of Indiana is as it actually is, so there is a possible world that provides a counterexample to the claim that the argument is valid.

Glen Helman 03 Aug 2010

1.4.2. Division

The pattern of truth values for premises and conclusion that is ruled out by entailment (i.e., true premises with a false conclusion) will recur often enough that it will be convenient to have special vocabulary for it. Let us say that a set Γ is *divided* from a set Δ whenever all members of Γ are true and all members of Δ are false. Whatever gives the sentences in Γ and Δ such values will be said to divide these sets. The source of the truth values will differ from context to context though, for the time being, it will be a possible world. When there is something of the appropriate sort that divides a set Γ from a set Δ, we will say that Γ and Δ are *divisible*; otherwise we will say they are *indivisible*.

Notice that these ideas are asymmetric. When one set is divided from another it is the members of the first set that true and the members of the second that are false. You might think of sets being divided vertically, with the first set above the second. In this spatial metaphor, truth is thought of as higher than falsehood; and, although this is only a metaphor, it is a broadly useful one and is consistent with the appearance of Absurdity at the bottom of Figure 1.2.5-2 and Tautology at the top. The asymmetry of division is especially important to remember in the case of the terms divisible and indivisible since this way of expressing the idea could suggest a symmetric relation between the results of a division.

As with talk of sets of sentences as premises, it is really only the list of members of a set that we care about here, and we speak of sets only because the order of the list and the occurrence of repetitions in it do not matter. In particular, we will not distinguish between a sentence and a set that has only it as a member. So we can restate the negative definition of entailment as follows:

$$\Gamma \vDash \varphi \text{ if and only if there is no possible world that divides } \Gamma \text{ from } \varphi.$$

We will also say that an argument is divided when its premises are divided from its conclusion, so we can say that an argument is valid when no possible world divides it. So to say that a possible world divides an argument is to say that the world is a counterexample to the argument's valdity. The divisibility or indivisibility of an argument thus amounts to the existence or non-existence of such a counterexample.

It can help when thinking about cases of division where one or both of the sets Γ and Δ is empty to restate the requirement all members of Γ are true as no member of Γ is false and restate the requirement for Δ analogously. That is, the most generally useful form of definition of division is this:

Γ is divided from Δ	if and only if	no member of Γ is false and no member of Δ is true

Notice that the requirement this places on a set is automatically satisfied when that set is the empty set ∅. That means that we can say:

Γ is divided from ∅ if and only if no member of Γ is false

∅ is divided from Δ if and only if no member of Δ is is true

Either way, we can see in particular that the empty is bound to be divided from itself. This consequence is no more than a curiosity, but it serves to emphasize that we are using the term divides in a rather special sense.

Glen Helman 03 Aug 2010

1.4.3. Conditional exhaustiveness

We can use the idea of division to generalize entailment to a relation between sets. And it is useful to do this because the more general relation encompasses all the deductive properties and relations of sentences. Although we have focused on entailment and will continue to do so, it doesn't suffice by itself to capture all the ideas of deductive logic. In particular, we need the idea of the absurdity $\bot$ to describe inconsistency in terms of entailment, and we have not yet seen how to say, in terms of entailment, when sentences are jointly exhaustive. But the more general relation can serve to define both of these ideas.

This new relation associated with joint exhaustiveness in much the way entailment is associated with tautologousness. Actually, it is associated in this way with a more general idea of exhautiveness that concerns any number of sentences, not merely two. Just as a pair of sentences are jointly exhaustive when we can be sure that, no matter what, at least one of the two is true, we will say that a set Δ of any size is *exhaustive* when we can be sure that at least one of its members is true. We will speak of these members as *alternatives*, so a set of alternatives is exhaustive when we can be sure that always at least one of these alternatives is true.

For example, the alternatives The glass is full, The glass is empty, and The glass is partly full form a set that is exhuastive in this sense. You might notice that it happens that any two of these alternatives are mutually exclusive, but that is an accident of this example. Replacing the first two alternatives with The glass is at least 90% full and The glass is no more than 10% full would not damage exhaustiveness since the new alternatives are true in even more possibilities, and neither of them excludes the claim that the glass is partly full. For another, more artificial, example, consider The book is not red, The book is not green, and The book is not blue. It is possible for all three of these alternatives to be true, so certainly no two of them are mutually exclusive; and if one is false the other two are true, so we are bound to have at least two of them true and the three are certainly an exhaustive set of alternatives.

We will use the notation $\vDash \Delta$ for this general idea of exhaustiveness and define it more formally (in a negative and positive form, respectively) as follows:

$\vDash \Delta$	if and only if	there is no possible world in which all members of Δ are false
	if and only if	in each possible world, at least one member of Δ is true

The notation for exhaustiveness provides notation for tautologousness; for, if φ is the sole member of Δ, a guarantee that at least one alternative from Δ is true is a guarantee that φ is true. So we can write $\vDash \varphi$ to say that φ is a tautology —i.e., that $\varphi \simeq \top$. The extended use of the entailment turnstile also provides us with a new notation for the idea of joint exhaustiveness: $\varphi \triangledown \psi$ if and only if $\vDash \varphi, \psi$.

Now let us return to the project of generalizing entailment. While tautologousness is an unconditional guarantee of truth, entailment guarantees the truth of its conclusion only given the truth of a set of assumptions. Entailment is thus a guarantee of truth for a single sentence only given the conditions set out in the assumptions, and we can think about an analogous conditional guarantee that a set is exhaustive. Saying that Δ is exhaustive unconditionally tells us that ranges of possibilities left open by its alternatives taken together cover all possibilities whatsoever. We can say that a set Δ is *exhaustive given* a set Γ when the ranges of possibilities left open by the alternatives in Δ taken together cover all possibilities in which every assumption in Γ is true. When this is so we have a guarantee that in any possible world in which all assumptions in Γ are true at least one alternative in Δ is true. For example, while the two alternatives The glass is full and The glass is empty are not jointly exhaustive, they are exhaustive given the assumption The glass is not partly full since it rules out all possibilities where they are both false.

Our notation for conditional exhaustiveness will again use the entailment turnstile, writing $\Gamma \vDash \Delta$ with the set of assumptions on the left and the set of alternatives on the right. It will help in reading this notation to have vocabulary that makes Γ the subject, so we will say that Γ *renders* Δ *exhaustive* when Δ is exhaustive given Γ. The negative and positive forms of the definition of this idea are as follows:

$\Gamma \vDash \Delta$	if and only if	there is no possible world in which all members of Δ are false while all members of Γ are true
	if and only if	in each possible world in which all members of Γ are true, at least one member of Δ is true

And, as promised, this idea can be stated very directly in terms of division:

$\Gamma \vDash \Delta$ if and only if there is no possible world that divides Γ from Δ. Entailment is the special case where the set Δ consists of a single sentence, for to say that φ is entailed by Γ comes to the same thing as saying that φ is rendered exhaustive by Γ. Either way we are claiming that there is no possible world that divides Γ from φ.

In cases of conditional exhaustiveness that are not cases of entailment, what is rendered exhaustive is either a set of several alternatives or the empty set. In these cases, it does not make sense to speak of a conclusion, for when the set on the right has several members, these sentences need not be valid conclusions from the set that renders them exhaustive. Indeed, a jointly exhaustive pair of alternatives will be rendered exhaustive by any set, but often neither member of the pair will be entailed by that set. This is particularly clear in the case of sentences like The glass is full and The glass is not full that are both jointly exhaustive and mutually exlcusive—i.e., that are contradictory. Although the set consisting of such pair is rendered exhaustive by any set, only an inconsistent set could entail both of these alternatives. So the term conclusion will be reserved for cases where there is a single alternative.

Glen Helman 03 Aug 2010

1.4.4. A general framework

It was noted in the last section that conditional exhaustiveness does not merely generalize entailment and unconditional exhaustiveness but encompasses all deductive properties and relations. It is not surprising that does so if these properties and relations are understood to all consist in guarantees that certain parterns of truth values appear in no possible world. For any claim there is no world where certain sentences Γ are true and other sentences Δ are false is a claim that $\Gamma \vDash \Delta$. Of course, a given deductive property or relation may rule out a number of different patterns—i.e., rule out a number of different ways of distributing truth values among the sentences it applies to—but this just means that a deductive property or relation may consist of a number of different claims of conditional exhaustiveness. In the case of the properties and relations we will consider, only equivalence and contradictoriness involve more than one claim of conditional exhaustiveness.

The table below summarizes the deductive properties and relations that involve only one claim of conditional exhaustiveness along with the vocabulary we have used for various special cases. The ideas discussed in the last subsection appear in the three columns at the right. Moving down one of these columns, we move from an unconditional guarantee of truth somewhere in a set of alternatives to a conditional guarantee that is hedged with one or more assumptions. Moving left to right in a one of the rows, we move from a guarantee of truth that focuses on a single alternative, a definite conclusion, to one that applies to a set of two or more alternatives.

assumptions ↓ / *alternatives* →	*none*	*one*	*two*	*any no.*
none		$\vDash \psi$ *tautologous*	$\vDash \psi, \psi'$ (or $\psi \triangledown \psi'$) *jointly exhaustive*	$\vDash \Delta$ *exhaustive*
one	$\varphi \vDash$ *absurd*	$\varphi \vDash \psi$ *implies*	$\varphi \vDash \psi, \psi'$	$\varphi \vDash \Delta$
two	$\varphi, \varphi' \vDash$ (or $\varphi \vartriangle \varphi'$) *mutually exclusive*	$\varphi, \varphi' \vDash \psi$	$\varphi, \varphi' \vDash \psi, \psi'$	$\varphi, \varphi' \vDash \Delta$
any no.	$\Gamma \vDash$ *inconsistent*	$\Gamma \vDash \psi$ *entails*	$\Gamma \vDash \psi, \psi'$	$\Gamma \vDash \Delta$ *renders exhaustive*

The column to the left of these three covers the cases where the set of alternatives is empty. There can be no unconditional guarantee of this sort, so there is no entry in the first row. The entry would not be a property or relation but instead the false statement $\varnothing \vDash \varnothing$ (which asserts an unconditional guarantee that some member of the empty set is true).

Since there are no alternatives in question, the ideas in the first column are really properties of sets of assumptions (just as those in the first row are properties of sets of alternatives). Absurdity, mutual exlucsiveness, and inconsistency are negative properties, each of which guarantees that a certain group of assumptions cannot all be true. They do this indirectly by making these assumptions conditions of a guarantee of something that is bound to be false—i.e., that the empty set of alternatives exhausts all possibilities. That is, they use the same device as a sentence like If that's a good book, then I'm the King of France which denies something by stating it as a sufficient condition for an absurd claim.

So, in each of these columns, movement down from one row to the next is a matter of making a guarantee of truth conditional on further assumptions. It is possible to think of movement to the right within each row in a somewhat analogous way: adding alternatives modifies a guarantee by adding exceptions. To claim that It is raining and It isn't raining are jointly exhaustive is not to guarantee the truth of either sentence, but such a claim does assert the existence of a guarantee that for each sentence that it is true apart from cases where the other is true. Similarly, a claim of entailment is a guarantee that the premises of an argument are not all true unless the conclusion is, so it can be seen to differ from a claim of inconsistency by adding an exception.

Terms like except and unless carry implicatures that can interfere with understanding this idea. It is important to understand them as you would in a guarantee. A guarantee that a product will function for three years unless it has been abused merely makes the guarantee conditional on the absence of abuse. It does not "guarantee" in addition that the product will not function if it has been abused although the statement The product will function unless it has been abused might suggest this under other circumstances.

The ideas of division and conditional exhaustiveness also provide ways of extending to any set the idea of logical independence introduced in 1.2.6 in the case of a pair of sentences. First, let us look at this general idea of logical independence directly. We will say that a set Γ of sentences is *logically independent* when every way of assigning a truth value to each member of Γ is exhibited in at least one possible world. This is the same as saying that for every part of the set (counting both the empty set and the whole set Γ as parts of Γ) it is possible to divide that part from the rest of the set. When the set has two members, this is the same as the earlier idea. When the set $\{\varphi\}$ containing a single sentence φ is logically independent in this sense, the sentence φ is said to be *logically contingent* because there is at least one possible world in which it is true and at least one where it is false, so its truth or falsity is not settled by logic.

Conditional exhaustiveness provides an alternative way of describing this idea. When the sentences in a set are *not* independent, not every way of dividing them into a set of true sentences and a set of false sentences is logically possible. And when some way of dividing them is not possible, the set contains at least one pair of non-overlapping subsets Γ and Δ such that $\Gamma \vDash \Delta$. So the members of a set are logically independent when the relation of conditional exhaustiveness never holds between non-overlapping subsets. (It always holds between sets that overlap because there is no way of dividing such sets.)

When a set is logically independent, each member is contingent and any two of its members are logically independent, but the contingency of members and the independence of pairs does not by itself imply that the set as a whole is logically independent. For example, assume that the sentences X is fast, X is strong, X has skill, and X has stamina form an independent set. Then the sentences

X is fast and strong	X has skill and stamina	X is fast and has stamina

are each contingent, and any two of them can be seen to be independent. However, the first two taken together entail the third, so these three more complex sentences do not form an independent set.

Glen Helman 03 Aug 2010

1.4.5. Reduction to entailment

Conditional exhaustiveness relaxes the restriction to a single conclusion found in entailment to include cases where there are several alternatives or none at all. To express the ideas captured by conditional exhaustiveness in terms of entailment, we need to add ways of capturing each of these added cases.

When a claim of conditional exhaustiveness offers no alternatives, it asserts the inconsistency of the assumptions; and that comes to the same thing as entailing the specific absurdity $\bot$. That is, we can state the following:

INCONSISTENCY *VIA* ABSURDITY. $\Gamma \vDash$ (i.e., $\Gamma \vDash \varnothing$) if and only if $\Gamma \vDash \bot$.

This law holds because rendering exhaustive the empty set and entailing $\bot$ both offer conditional guarantees of a truth that cannot exist, so each has the effect of ruling out the possibility of meeting the conditions of the guarantee.

To express the idea of rendering exhaustive multiple alternatives using entailment we need help from the concept of contradictoriness. Contradictoriness comes in here because having an exception in a guarantee comes to the same thing as having its contradictory as a condition. For example, the guarantee The product will funciton for three years unless it is abused is equivalent to The product will function for three years if it hasn't been abused, and the guarantee The product will function for three years if it is serviced regularly is equivalent to The product will function for three years unless it is not serviced regularly. To make this intuitive point more formally, note first that when sentences are contradictory, they always have opposite truth values. So making one true comes to the same thing as making the other false, and contradictory sentences play opposite roles when sets are being divided. More specifically, if φ and $\overline{\varphi}$ are contradictory sentences, then

Γ is divided from (Δ together with φ)
if and only if
(Γ together with $\overline{\varphi}$) is divided from Δ

because each of these divisions requires that φ be made false and $\overline{\varphi}$ be made true. Since a claim of conditional exhaustiveness asserts that a division is not possible, having a sentence as an alternative comes to the same thing as having a sentence contradictory to it as an assumption; that is,

if φ and $\overline{\varphi}$ are contradictory, then $\Gamma \vDash \varphi, \Delta$ if and only if $\Gamma, \overline{\varphi} \vDash \Delta$

If we apply this idea repeatedly (perhaps infinitely many times), we can move any set of alternatives to the left of the turnstile, and that is the basis of the following law:

> ALTERNATIVES *VIA* ASSUMPTIONS. Let $\overline{\Delta}$ be the result of replacing each member of Δ by a sentence contradictory to it. Then $\Gamma \vDash \Delta, \Sigma$ if and only if $\Gamma, \overline{\Delta} \vDash \Sigma$.

In short, we can remove alternatives if we put sentences contradictory to them among the assumptions.

The laws we have seen give us two approaches to restating claims of conditional exhaustiveness as entailments. A claim with no alternatives—i.e., a claim of inconsistency—can be turned into an entailment by adding $\bot$ as the conclusion. And we may replace any alternatives by assumptions contradictory to them to reduce multiple alternatives to a single conclusion. The two may be combined by replacing all alternatives by contradictory assumptions and then adding $\bot$ as conclusion. The following table uses these two approaches to state all the deductive properties we have considered in terms of the general ideas of entailment and contradictoriness and of the specific absurdity $\bot$:

Concept	*in terms of entailment and other ideas*
φ is a tautology	$\vDash \varphi$
Γ entails φ	$\Gamma \vDash \varphi$
φ is absurd—i.e., $\varphi \vDash$	$\varphi \vDash \bot$
φ and ψ are mutually exclusive—i.e., $\varphi \vartriangle \psi$ (or $\varphi, \psi \vDash$)	$\varphi, \psi \vDash \bot$
Γ excludes φ—i.e., $\Gamma, \varphi \vDash$	$\Gamma, \varphi \vDash \bot$
Γ is inconsistent—i.e., $\Gamma \vDash$	$\Gamma \vDash \bot$
φ and ψ are jointly exhaustive —i.e., $\varphi \triangledown \psi$ (or $\vDash \varphi, \psi$)	$\overline{\varphi} \vDash \psi$ (or $\overline{\psi} \vDash \varphi$, or $\overline{\varphi}, \overline{\psi} \vDash \bot$)
Γ is exhaustive—i.e., $\vDash \Gamma$	$\overline{\Gamma} \vDash \bot$
φ and ψ are equivalent—i.e., $\varphi \simeq \psi$	both $\varphi \vDash \psi$ and $\psi \vDash \varphi$
φ and ψ are contradictory—i.e., $\varphi \between \psi$ (or both $\varphi, \psi \vDash$ and $\vDash \varphi, \psi$)	both $\varphi, \psi \vDash \bot$ and $\overline{\varphi} \vDash \psi$ (or $\overline{\psi} \vDash \varphi$, or $\overline{\varphi}, \overline{\psi} \vDash \bot$)

Here $\overline{\varphi}$ is any sentence contradictory to φ, and $\overline{\Gamma}$ is the result of replacing each member of Γ by a sentence contradictory to it

There are alternative ways of stating each of these ideas in terms of entailment. Any time $\bot$ appears as the conclusion and there is at least one assumption, $\bot$ could be replaced as the conclusion by a sentence contradictory to an assumption, which is then dropped. That is, $\Gamma, \varphi \vDash \bot$ if and only if $\Gamma \vDash \overline{\varphi}$. And whenever $\bot$ is not the conclusion, it could be made the conclusion if the a sentence contradictory to the previous conclusion is added to the assumptions —i.e., $\Gamma \vDash \varphi$ if and only if $\Gamma, \overline{\varphi} \vDash \bot$. Also, we may replace an assumption and the conclusion both by putting a sentence contradictory to each on the other side of the turnstile—i.e., $\Gamma, \varphi \vDash \psi$ if and only if $\Gamma, \overline{\psi} \vDash \overline{\varphi}$.

It may seem pointless to define the relation of contradictoriness in terms of entailment, as is done in the last row of this table, since we need to use the idea of contradictoriness in order to do this. But the definition does mean that, once we know a single sentence contradictory to a given sentence, we can say what other sentences are contradictory to it using only the ideas of entailment and absurdity.

1.4.6. Laws for entailment

Most of the laws of deductive reasoning we will study will be generalizations about specific logical forms that will be introduced chapter by chapter, but some very general laws can be stated at this point. We have already seen some of these. We have just seen the laws tying inconsistency to Absurdity alternatives to assumptions. And the principles of reflexivity and transitivity for implication discussed in 1.2.3 can be generalized to provide basic laws for entailment and conditional exhaustiveness. We will look first at the case of entailment.

Two basic laws suffice to capture the basic properties of entailment considered in its own right:

LAW FOR PREMISES. *Any set of assumptions entails each of its members.* That is, $\Gamma, \varphi \vDash \varphi$ (for any sentence φ and any set Γ).

CHAIN LAW. *A set of assumptions entails anything entailed by things it entails.* That is, if $\Gamma \vDash \varphi$ for each assumption φ in Δ and $\Delta \vDash \psi$, then $\Gamma \vDash \psi$ (for any sentence ψ and any sets Γ and Δ).

Taken together, these laws tell us that the relation which holds between sets Γ and Δ when Γ entails all members of Δ is both reflexive and transitive. For the law for premises tells us that any set entail every member of itself. And, if Γ entails every member of Δ and Δ entails every member of the Σ, then Γ also entails every member of Σ by the chain law. Although this reflexive and transitive relation is, like conditional exhaustiveness, a relation between sets of sentences, they are different relations, and we will see later that conditional exhaustiveness is neither reflexive nor transitive.

These two principles have as a consequence two further principles the addition and subtraction of assumptions that will play an important role in our study of entailment:

MONOTONICITY. *Adding assumptions never undermines entailment.* That is, if $\Gamma \vDash \varphi$, then $\Gamma, \Delta \vDash \varphi$ (for any sets Γ and Δ and any sentence φ).

LAW FOR LEMMAS. *Any assumption that is entailed by other assumptions may be dropped without undermining entailment.* That is, if $\Gamma, \varphi \vDash \psi$ and $\Gamma \vDash \varphi$, then $\Gamma \vDash \psi$ (for any sentence φ and set Γ).

Each of these principles is based on both the law for premises and the chain law. In the case of the first, the law for premises tells us that Γ together with Δ entails every member of Γ alone, so $\Gamma, \Delta \vDash \varphi$ if $\Gamma \vDash \varphi$ by the chain law. The assumption of the second that $\Gamma \vDash \varphi$ combines with the law for premises to tell us that Γ entails every member of the result of adding the further assumption

Φ, and the chain law then tells us that Γ entails anything Ψ entailed by this enlarged set of assumptions.

The term lemma can be used for a conclusion that is drawn not because it is of interest in its own right but because it helps us to draw further conclusions. The second law tells us that if we add to our premises Γ a lemma φ that we can conclude from them, anything ψ we can conclude using the enlarged set of premises can be concluded from the original set Γ.

The idea behind the law of monotonicity is that adding assumptions can only make it harder to find a possible world that divides the assumptions from the conclusion, so, if no possible world will divide Γ from φ, we can be sure that no world will divide from φ the larger set of assumptions we get by adding some further assumptions Δ. The term monotonic is applied to trends that never change direction. More specifically, it is applied to a quantity that does not both increase and decrease in response to changes in another quantity. In this case, it reflects the fact that adding assumptions will never lead to a decrease in the sets of alternatives rendered exhaustive by them and adding alternatives will never lead to a decrease in the sets of assumptions rendering them exhaustive.

It is a distinguishing characteristic of deductive reasoning that a principle of monotonicity holds. For, when reasoning is not risk free, additional data can show that a initially well-supported conclusion is false and do so without undermining the original premises on which the conclusion was based. If such further data were added to the original premises, the result would no longer support the conclusion. This means that risky inference is, in general, *non-monotonic* in the sense that additions to the premises can reduce the set of conclusions that are justified. This is true of inductive generalization and of inference to the best explanation of available data, but the term non-monotonic is most often applied to another sort of non-deductive inference, an inference in which features of typical or normal cases are applied when there is no evidence to the contrary. One standard example is the argument from the premise Tweety is a bird to the conclusion Tweety flies. This conclusion is reasonable when the premise exhausts our knowledge of Tweety; but the inference is not free of risk, and the conclusion would no longer be reasonable if we were to add the premise that Tweety is a penguin.

The law for premises and the chain can be shown to give a complete account of the general laws of entailment in the sense that any relation between sets of sentences and sentences that obeys them is an entailment relation for some set of possibile worlds and assignment of truth values to sentences in each world. But this is not to say that they provide a complete general account of deductive properties and relations, because our definitions of the may of these in terms of entailment also used the ideas of contradiction and the absurdity $\bot$. The laws providing for inconsistency *via* absurdity and alternatives *via* assumptions govern these ideas but they were stated for conditional exhaustiveness rather than entailment. Although laws for inconsistency and contradictoriness might be stated in terms of entailment, doing so now would pointlessly anticipate later topics, so we will let the two laws we began with suffice.

Let us look briefly at conditional exhaustiveness. As noted earlier, it is neither reflexive nor transitive. Although $\Gamma \vDash \Gamma$ whenever Γ has at least one member, we have already seen that $\varnothing \nvDash \varnothing$. And if conditional exhaustiveness were transitive every sentence φ would imply every other sentence ψ since $\varphi \vDash \varphi, \psi$ and $\varphi, \psi \vDash \psi$. In spite of this, we can state laws for relative exhaustiveness that are somewhat analogous to the basic laws for entailment. First two basic laws:

REPETITION. *A set of assumptions renders exhaustive any set of alternatives that it overlaps.* That is, $\Gamma, \varphi \vDash \varphi, \Delta$ (for any sentence φ and any sets Γ and Δ).

CHAIN LAW. *If a set of sentences each of which is a sufficient exception to a claim of exhaustiveness itself renders exhaustive a set of sentences each of which is a sufficient additional assumption for the claim, the claim holds without exceptions or additional assumptions.* Suppose (i) $\Gamma \vDash \varphi, \Delta$ for each φ in Σ, (ii) $\Gamma, \psi \vDash \Delta$ for each ψ in Θ, and (iii) $\Sigma \vDash \Theta$. Then $\Gamma \vDash \Delta$ (for any sentences φ and ψ and any sets Γ, Δ, Σ, and Θ).

Although the first of these is similar to the law for premises, it is given a different name because this law is as much about alternatives as about assumptions. The metaphor of a chain does not apply very directly to the second law, but this law does play a role for conditional exhaustiveness that is analogous to the chain law for entailment. Its verbal statement is more complex than the other laws, and it may not be clear how to fit it with what follows. The idea is that condition (i) tells us that the claim of conditional exhaustiveness of Δ given Γ holds when we add to Δ any member φ of Σ as a further alternative (i.e., as an exception to the claim). Condition (iii) guarantees the exhaustiveness of Θ given Σ, and condition (ii) tells us that the exhaustiveness of Δ holds given Γ together with any member ψ of Θ as a further assumption. The law then holds because, if each member of Γ is true, then by (i) we must have at least one member of Δ true unless each member of Σ is true; and, if the latter is the case, by (iii) we must have at least one member of Θ true and, by (ii), this is enough to insure that at least one member of Δ is true as we wished.

As with entailment, we will consider two laws that follow from this basic pair.

MONOTONICITY. *Adding assumptions or alternatives never undermines conditional exhaustiveness.* That is, if $\Gamma \vDash \Delta$, then $\Gamma, \Sigma \vDash \Delta, \Theta$ (for any sets Γ, Δ, Σ, and Θ);

CUT. *An alternative may be dropped if adding it as an assumption is enough to render the remaining alternatives exhaustive.* That is, if $\Gamma, \varphi \vDash \Delta$ and $\Gamma \vDash \varphi, \Delta$, then $\Gamma \vDash \Delta$ (for any sentence φ and any sets Γ and Δ).

The second is relatively close in form to the law for lemmas but it given a given a different name, as was the repetition law, because assumptions and alternatives play parallel roles in it. The significance of the term cut lies simply in its effect of dropping the sentence φ. The idea behind that is that, given the truth of all members of Γ, at least one of the alternatives Δ to be true in a case where φ is true because $\Gamma, \varphi \vDash \Delta$ and in a case where φ false because $\Gamma \vDash \varphi, \Delta$ and φ cannot be the alternative that is true.

One of the reasons for considering conditional exhaustiveness is that a law providing alternatives *via* assumptions follows from the basic laws. This law takes the following form:

ALTERNATIVES *VIA* ASSUMPTIONS. If both $\varphi, \psi \vDash$ and $\vDash \varphi, \psi$ (i.e., φ and ψ are contradictory), then $\Gamma \vDash \varphi, \Delta$ if and only if $\Gamma, \psi \vDash \Delta$.

To see why this follows, suppose that φ and ψ are contradictory and $\Gamma \vDash \varphi, \Delta$. We can apply the chain law with $\Gamma, \psi \vDash \Delta$ as the claim we wish to establish and $\varphi, \psi \vDash$ (i.e., $\varphi, \psi \vDash \varnothing$) as the claim cited in condition (iii). Because the Θ mentioned in the law is the empty set $\varnothing$ in this case, there is nothing to show for (ii) since there is no member of Θ for which it might fail; and (i) says merely that $\Gamma, \psi \vDash \varphi, \Delta$, which holds by monotonicty (since we have assumed that $\Gamma \vDash \varphi, \Delta$), and $\Gamma, \psi \vDash \psi, \Delta$, which holds by repetition. We can use $\vDash \varphi, \psi$ in a similar way to show $\Gamma \vDash \varphi, \Delta$ when we suppose $\Gamma, \psi \vDash \Delta$.

We cannot expect to get the law providing for inconsistency *via* Absurdity without some principle stating the logical properties of $\bot$ (something we will consider in the next subsection), but we can say that $\Gamma \vDash$ (i.e., Γ is inconsistent) if $\Gamma \vDash \varphi$ and $\varphi \vDash$ (i.e., φ is absurd). (The argument applies the chain law in a way similar too, but simpler than, the one we just saw.) In the other direction, knowing that Γ is inconsistent does not enable us to conclude that it entails some inconsistent sentence because we don't yet have a law telling us that there are any inconsistent sentences. But we can say that if Γ is inconsistent, it entails any inconsistent sentence there is because an inconsistent set entails any sentence whatsoever: we know that if $\Gamma \vDash$ (i.e., $\Gamma \vDash \varnothing$) then $\Gamma \vDash \varphi$, for any sentence φ, by monotonicity. This gives us the following law pointing the way to, if not providing, inconsistency *via* absurdity:

INCONSISTENCY *VIA* ABSURDITY. If $\varphi \vDash$, then $\Gamma \vDash$ if and only if $\Gamma \vDash \varphi$.

1.4.7. Duality

In the context of conditional exhaustiveness all that need be said about the logical properties of Tautology $\top$ and Absurdity $\bot$ is that Tautology is a tautology (i.e., $\vDash \top$) and that Absurdity is absurd (i.e., $\bot \vDash$). The first of these makes sense for entailment and, together with the basic laws of entailment provides with the sort of laws we will go on to consider shortly. However, it is the latter laws that we will focus on since they state the role of $\top$ in entailment. And, in the case of $\bot$, saying merely that it is absurd tells us nothing from the point of view of entailment since that is to say only that $\bot \vDash \bot$.

Tautology $\top$ is entailed by any set of premises (the empty set included) because it cannot go beyond the information contained in any set of sentences; and, for the same reason, the presence of $\top$ among the premises of an argument contributes nothing to the argument's validity. These two ideas can be expressed more formally in the following laws.

LAW FOR $\top$ AS A CONCLUSION. $\Gamma \vDash \top$ (for any set Γ).

LAW FOR $\top$ AS A PREMISE. $\Gamma, \top \vDash \varphi$ if and only if $\Gamma \vDash \varphi$ (for any set Γ and sentence φ).

Although they are stated for $\top$, these laws will hold for all tautologies since they hold simply in virtue of the proposition expressed by $\top$.

These laws are different in character from the ones consider in the last subsection because they concern the logical properties of a specific sort of sentence rather than the general principles governing logical relations. They are also a first sample of a common pattern in the laws of deductive reasoning that we will consider. Entailment is so central to deductive reasoning that an account of the role of a kind of sentence in entailment as a conclusion and as a premise will usually tell us all we need to know about it.

A simple law describes the role of absurdities as premises. We state it for the specific absurdity $\bot$.

LAW FOR $\bot$ AS A PREMISE. $\Gamma, \bot \vDash \varphi$ (for any set Γ and sentence φ).

An argument with an absurdity among its premises is valid by default. Since its premises cannot all be true, there is no risk of *new* error no matter what the conclusion is. There is no law restating the significance of having $\bot$ as a conclusion because that is simplest way we have of using entailment to say that a set of assumptions is inconsistent.

Although entailment will be our focus, it is enlightening to consider analogues for conditional exhaustiveness of the laws just stated. In particular, we can state a law for $\bot$ as an alternative in the context of conditional exhaustiveness, and all the properties of ⊤ and ⊥ take a particularly symmetric form when stated in terms of that relation.

	as a premise	*as an alternative*
Tautology	if $\Gamma, \top \vDash \Delta$, then $\Gamma \vDash \Delta$	$\Gamma \vDash \top, \Delta$
Absurdity	$\Gamma, \bot \vDash \Delta$	if $\Gamma \vDash \bot, \Delta$, then $\Gamma \vDash \Delta$

That is, while ⊤ contributes nothing as a premise and may be dropped, it is enough for a claim of conditional exhaustiveness to hold that it be an alternative (no matter how small the set Γ of premises or the set Δ of other alternatives). And while it is enough to have ⊥ as a premise (no matter how small the set of alternatives is), it contributes nothing as an alternative and may be dropped.

Notice that the converses of the principles at the upper left and lower right hold by monotonicity because they are just the addition of a premise in one case and an alternative in the other. If we take the if and only if principle that results from adding the converse to the lower right and consider a case where Δ is empty, we get

$$\Gamma \vDash \bot \text{ if and only if } \Gamma \vDash$$

This is the principle for conditional exhaustiveness that lies behind the law providing inconsistency *via* Absurdity of 1.4.5 (and it follows from the law promising inconsistency *via* absurdity that was stated at the end of the last section once we have stated that $\bot \vDash$). The moral is that our use of ⊥ as a conclusion to define inconsistency in terms of entailment really involves the same idea as the principle for ⊥ as an alternative that may be stated for conditional exhaustiveness.

The symmetry exhibited by the set of principles in the table above might be traced to the symmetry of conditional exhaustiveness: since ⊤ and ⊥ are contradictory, having one as an assumption comes to the same thing as having the other as an alternative according to the law of 1.4.5 providing alternatives *via* assumptions. However, there is a more general idea behind this symmetry that will apply also to cases where sentences are not contradictory.

The essential difference between the lower left and upper right in the table above lies in interchanging ⊥ and ⊤ and, at the same time, interchanging premises and alternatives. And the same is true of the upper left and lower right. That is, if we apply this transition to the lower left, we get

$$\Delta \vDash \Gamma, \top$$

and that differs from the upper right only in the order of the alternatives and the exchange of Δ for Γ. And neither of these differences is important.

Alternatives function only as a set, so the order in which they are listed does not matter. And, since each of Γ and Δ could be any set, exchanging them does not alter the content of the principle. Either way, we say that it is enough to have $\top$ as an alternative no matter what premises and what further alternatives we have. The possibility of the sort of transformation used to get from the lower left to the upper right can be expressed by saying that $\top$ and $\bot$ on the one hand and premise (or assumption) and alternative on the other constitute pairs of *dual* terms. We will run into other pairs of terms later that fit into the same sort of duality.

Glen Helman 03 Aug 2010

1.4.s. Summary

1 Entailment may be defined in two equivalent ways, negatively as the relation that holds when the conclusion is false in no possible world in which all the premises are true or positively as the relation which holds when the conclusion is true in all such worlds. The negative form has the advantage of focusing attention on the sort of possible world that serves as a counterexample to a claim of entailment. The positive form characterizes a relation of entailment as a conditional guarantee of the truth of the conclusion, a guarantee conditional on the truth of the premises.

2 The requirements for a world to serve as a counterexample to entailment suggest the general idea of dividing a pair of sets by making all members of the first true and all members of the second false. A world will be said to divide an argument when it divides the premises from the conclusion.

3 The idea of division enables us to define a relation of conditional exhaustiveness between sets: one set renders another exhaustive when there is no possible world that divides the two sets. We will extend the notation for entailment to express this relation between sets Γ and Δ as $\Gamma \vDash \Delta$. Entailment is the special case of this where Δ has only one member. When Δ has more than one member, its members will be referred to as alternatives because a relation of conditional exhaustive provides a conditional guarantee only that at least one member of the second set it true.

4 Since a set of alternatives can have more than one member or be empty, conditional exhaustiveness encompasses all the deductive properties and relations we have considered (as well as an extension of the idea of joint exhaustiveness to any set of sentences). The way a property or relation is expressed using conditional exhaustiveness is tied directly to the negative form of the definition of the property or relation. When no relation of conditional exhaustiveness holds no matter how a set is divided into two parts, all patterns of truth values for its members are possible and the set is logically independent. A single sentence that forms a logically independent set is logically contingent.

5 Definitions in terms of conditional exhaustiveness can be converted into definitions in terms of entailment by replacing empty sets of alternatives with $\bot$ and reducing the size of multiple sets of alternatives by replacing members by adding assumptions that are contradictory to them (using the basic law for conditional exhaustiveness).

6 Conditional exhaustiveness and entailment satisfy a principle of

monotonicity. The term monotonic reflects the fact that conditional exhaustiveness or entailment will never stop holding because of additions to the set of assumptions or set of alternatives. This principle is significant in distinguishing entailment from other forms of good inference, whose riskiness means that they are non-monotonic because adding information telling us that the risk does not pay off will undermine their quality. Both conditional exhaustiveness and entailment also satisfy analogues to the principles of reflexivity and transitivity for implication. In the case of reflexivity, these laws are repetition for conditional exhaustiveness and the law for premises for entailment. For transitivity, they are cut for conditional exhaustiveness and the law for lemmas for entailment. The latter licenses the use of lemmas, valid conclusions that are of interest only as premises in further arguments. A more general law, called the chain law, together with a law for premises, yields all laws of entailment, and these two principles amount to principles of reflexivity and transitivity for a relation of set entailment that holds when one set entails each member of another.

7 The laws describing the behavior of $\top$ and $\bot$ in the context of conditional exhaustiveness exhibit a kind of symmetry that we will see in other laws later. The sentences $\top$ and $\bot$ are dual as are the terms premise and alternative (or the left and right of an turnstile) in the sense that replacing each such term in a law by the one dual to it will produce another law.

1.4.x. Exercise questions

1. Any claim that a deductive relation holds can be stated as one or more claims that one set of sentences cannot be divided from another. (i) Restate each of the following claims in that way, and (ii) explicitly describe the sort of possibility that would divide the sets in question and is thus ruled out by claiming that the deductive relation holds. Nonsense words have been used to help you think to think how a possibility would be described without worrying whether that possibility could really occur.

For example, the claim that the sentences The widget plonked and The widget plinked are equivalent can be restated by saying that (i) the set consisting of the first sentence cannot be divided from the set consisting of the second sentence and vice versa. That is, (ii) it rules out any possibility in which the widget plonked but did not plink and any possibility in which the widget plinked but did not plonk.

- **a.** The gizmo is a widget and The gizmo is a gadget are mutually exclusive
- **b.** The gizmo is a widget and The gizmo is a gadget are jointly exhaustive
- **c.** The widget plinked is a tautology
- **d.** The widget plonked is absurd
- **e.** The widget was a gadget renders exhaustive the alternatives The widget plinked and The widget plonked
- **f.** The widget was a gizmo, The widget plinked, and The widget plonked are inconsistent

2. The basic law for conditional exhaustiveness can be used not only to replace alternatives by assumptions but also to replace assumptions by alternatives. For example, the claim that The widget is blue entails The widget is colored can be restated to say (i) The widget is blue and The widget is not colored are inconsistent, (ii) The widget is not blue and The widget is colored form an exhuastive set, or (iii) The widget is not colored entails The widget is not blue.

In the following, you will be asked to restate some statements of deductive relations by replacing alternatives with assumptions or assumptions with alternatives. You may add or remove ordinary negation to state the contradictories of sentences.

a. Restate the following as a claim of entailment: The gadget is red and The gadget is green are mutually exclusive

b. Restate the following as a claim of entailment: Someone is in the auditorium and There are empty seats in the auditorium are jointly exhaustive

c. Restate the following as a claim of absurdity: A widget is a widget is a tautology

d. Restate the following as a claim of tautologousness: A widget is a gadget is absurd

e. Restate the following as a claim of inconsistency: The widget is a gadget or gizmo and The widget is not a gadget entail The widget is a gizmo

f. Restate the following so that each assumption is replaced by an alternative and each alternative by an assumption: The widget has advanced and The widget has plonked render exhaustive the alternatives The widget has finished the task and The widget has broken

1.4.xa. Exercise answers

1. **a.** (i) The set consisting of The gizmo is a widget and The gizmo is a gadget cannot be divided from the empty set; that is, (ii) there is no possibility of the gizmo being both a widget and a gadget.

b. (i) The empty set cannot be divided from the set consisting of The gizmo is a widget and The gizmo is a gadget; that is, (ii) there is no possibility of the gizmo being neither a widget nor a gadget

c. (i) The empty set cannot be divided from the set consisting of only The widget plinked; that is, (ii) there is no possibility that the widget did not plink

d. (i) The set consisting of only The widget plonked cannot be divided from the empty set; that is, (ii) there is no possibility that the widget plonked

e. (i) The set consisting of only The widget was a gadget cannot be divided from the set consisting of The widget plinked and The widget plonked; that is, (ii) there is no possbility that the widget was a gadget while not either plinking or plonking.

f. (i) The set consisting of The widget was a gizmo, The widget plinked, and The widget plonked cannot be divided from the empty set; that is, (ii) there is no possibility that the widget was a gizmo and both plinked and plonked

2. **a.** The gadget is red entails The gadget is not green (*or:* The gadget is green entails The gadget is not red)

b. The auditorium is empty entails There are empty seats in the auditorium (*or:* There are no empty seats in the auditorium entails The auditorium is not empty)

c. A widget is a not widget is absurd

d. A widget is a not gadget is a tautology

e. The widget is a gadget or gizmo, The widget is not a gadget, and The widget is not a gizmo are inconsistent

f. The widget has not finished the task and The widget has not broken render exhaustive The widget has not advanced and The widget has not plonked