1.4. Comparing content: logical properties and relations

1.4.0. Overview

The properties and relations of sentence and propositions that are subject matter of deductive logic can be arranged in three groups.

Glen Helman 25 Aug 2005

1.4.1. A closer look at entailment

We will begin with a formal definition of the idea of entailment as a conditional guarantee of truth. When a conclusion φ is ***entailed*** by a set Γ of premises, we have a guarantee that φ is true provided that the members of Γ are all true. This is a strong guarantee for it holds, under the stated conditions, in all possible worlds. We can state this definition more formally in two equivalent ways.

$\Gamma \Rightarrow \varphi$	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

It is worth emphasizing that 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, there is no redundancy because each of the two emphasizes different aspects of the concept. The second—which we will speak of as the ***positive form***—is closely tied to the motivation for the concept, to the reason why the concept 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 will consider 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 pattern of truth values ruled out by entailment turns out to be one of the more cumbersome ones to state; and, since we will refer to it often, it will be useful to have special vocabulary for it. We will say that a possible world ***divides*** a set Γ from a set Δ when every member of Γ is true in the world and every member of Δ is false. This use of the word *divide* will prove helpful in a number of ways, but there is one respect in which it may be misleading. Other uses of *divide* point to a symmetric relation, when a is divided from b, b is divided from a. But the specific truth values that must be assigned to sets Γ and Δ for a world to divide Γ from Δ make this relation between the two sets fundamentally asymmetric. To counteract the suggestion of symmetry you might think of Γ being divided from Δ by being "set above" it, thinking of truth as being "higher" than falsity. As with the premises of an argument, we will have no need, when considering this concept of division, to distinguish between a sentence and a set with that sentence as its only member, so we may regard one or the other terms of the relation of division as a sentence. Using this idea, we can state the negative form of the definition of entailment as follows:

$\Gamma \Rightarrow \varphi$	if and only if	there is no logically possible world that divides Γ from φ

We will say that a possible world divides an argument when it divides its premises from its conclusion, so we can say that an argument is valid when no possible world divides it.

The kind of possible world ruled out by the negative form of the definition must, of course, also have some relation to the positive form. The positive form is generalization concerning all possible worlds of a certain sort. When a generalization is false, it is because of ***counterexample***, something of sort about which we generalize that does not have the property we have said 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. Thus a possible world that divides an argument is a counterexample to the claim that its premises entail its conclusion.

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 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 real content of a claim of entailment lies not in what it tells us about the actual world but in the fact that it makes a claim about all possible worlds. The other three patterns all appear in the actual truth values of some valid arguments (though not all are possible for certain arguments).

To see examples of this, consider the case of an argument whose conclusion is among its premises—for example,

Indianapolis is the capital of Indiana
Springfield is the capital of Illinois
———
Indianapolis is the capital of Indiana

Such an argument is trivial but, because of this, it is obviously valid. Its conclusion certainly does no more than extract information from the premises; and, because it is one of the premises, there is certainly no possible world in which it is false while the premises are all true. Now the example above has true premises and a true conclusion, the first of the patterns in Figure 1.4.1-1. The other two patterns of truth values allowed for valid arguments can be produced by changing *Illinois* and *Indiana*, respectively, to *Ohio*.

Indianapolis is the capital of Indiana
Springfield is the capital of Ohio
———
Indianapolis is the capital of Indiana

Indianapolis is the capital of Ohio
Springfield is the capital of Illinois
———
Indianapolis is the capital of Ohio

That these two patterns of truth values should be possible is clear also from the idea of extracting information. Information can be extracted from a set of sentences even though they are not all true, and the information extracted in such a case 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; 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 divides the premise from the conclusion.

Glen Helman 25 Aug 2005

1.4.2. Laws for entailment

Most of our concern with entailment will not be with particular examples, but instead with general laws. Most of these will be generalizations about specific logical forms, but some very general ones can be stated now (and a few of these appeared already in the exercise 1.1.x.1).

We will begin with single-premised entailment—i.e., with implication. Implication is ***reflexive*** in the sense that any sentence φ implies itself, and it is ***transitive*** in the sense that, if a sentence χ is implied by a sentence ψ that is in turn implied by a sentence φ, then χ is also implied directly by φ. That is,

$\varphi \Rightarrow \varphi$; and

if $\varphi \Rightarrow \psi$ and $\psi \Rightarrow \chi$, then $\varphi \Rightarrow \chi$

for any sentences φ, ψ, and χ. Notice that the second of these can equally well be described as saying that a sentence χ may be validly concluded from anything φ that implies a premise ψ from which χ may be validly concluded. In short, it tells us that we will not destroy validity if we replace the conclusion of a single-premised valid argument by something it implies, and we may replace the premise by anything that implies it. More graphically,

if φ / ψ is valid and $\psi \Rightarrow \chi$, then φ / χ is valid; and

if ψ / χ is valid and $\varphi \Rightarrow \psi$ and , then φ / χ is valid.

Laws somewhat analogous to reflexivity and transitivity apply to arguments with any sets of premises. What we will call the ***law for premises*** says that a sentence is entailed by any set of premises containing it. That is,

$\Gamma, \varphi \Rightarrow \varphi$

for any set Γ of sentences and any sentence φ. The analogue of the second law for single-premised arguments says that a set of premises that entails every premise of a valid argument also entails its conclusion: for any sets Γ and Δ and any sentence ψ,

if $\Gamma \Rightarrow \varphi$ for each premise φ in Δ and $\Delta \Rightarrow \psi$, then $\Gamma \Rightarrow \psi$

We will refer to this as the ***chain law*** since it enables us to link valid arguments together to get new valid arguments. These are not directly principles of reflexivity and transitivity since those ideas only make sense for relations between the same sorts of things; but a relation between sets of sentences that holds when Γ entails every member of Δ is reflexive and transitive.

We will consider two further general laws of entailment that follow from the law for premises and the chain law but are each valuable for special purposes. The first tells us that we can add premises without destroying the validity of an argument: for any sets Γ and Δ and any sentence φ

if $\Gamma \Rightarrow \varphi$, then $\Gamma, \Delta \Rightarrow \varphi$

This law should not be surprising because, in general, the more premises we have, the easier it is to validly conclude a given sentence. If we think of entailment as associating a collection of valid conclusions with any set of sentences, this law tells us that as the set of premises increases the set of valid conclusions will never decrease. Mathematicians apply the term *monotonic* to situations like this, so we will speak of this law as the principle of ***monotonicity*** for entailment.

Although monotonicity will play only an auxiliary role in our discussion of deductive reasoning, it is a distinguishing characteristic of deductive reasoning that such a principle holds. For, when reasoning is not risk free, additional data can show that a initially well-supported conclusion is false—and it can do this without undermining the original on which we based our conclusion. If such further information were added to our premises, we would not expect the conclusion to still be well supported. Indeed, the risk in good but risky inference can be thought of as a risk that further information will undermine the quality of the inference, so risky inference (or, more precisely, the way the quality of such inference is assessed) is, in general, ***non-monotonic***. 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 inferences that are based on features of typical or normal cases. 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 other side of the coin is that dropping premises can never help in deductive reasoning and may well destroy validity. But, while we cannot in general safely drop premises, we can drop a premise when it is entailed by others that we retain:

if $\Gamma, \varphi \Rightarrow \psi$ and $\Gamma \Rightarrow \varphi$, then $\Gamma \Rightarrow \psi$

for any set Γ and any sentences φ and ψ. 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. This law tells us that anything we can conclude using an intermediate conclusion φ is a valid conclusion from the original premises Γ, so it justifies the use of lemmas, and we will refer to it as the ***law for lemmas***.

In summary, what these laws tell us about entailment is that (i) we can validly conclude any premise (law for premises), (ii) we can validly conclude anything entailed by valid conclusions from our premises (chain law), (iii) we can add premises without destroying validity (monotonicity), and (iv) we may safely drop from our premises lemmas that are entailed by the remaining premises (law for lemmas).

1.4.3. Equivalence and tautologousness

Recall that the relation of equivalence applies to sentences that have the same informational content—for example,

Neither the shoulders nor the median are finished
The shoulders and the median are both unfinished

In each possible world, such sentences must have the same truth value as each other, which is the same as saying that neither can be false when the other is true, that each entails the other.

We could define equivalence as mutual entailment; but it will be useful to define it directly using definitions similar to those we have given for entailment. The key idea is that logical equivalence amounts to the necessary identity of truth values. Formally, we can describe the conditions under which a pair of sentences φ and ψ are ***(logically) equivalent*** as follows:

φ ⇔ ψ	if and only if	there is no possible world in which φ and ψ have different truth values
	if and only if	φ and ψ have the same truth value as each other in every possible world

Notice that the second form does not say that the truth values of these sentences do not vary from possibility to possibility, only that, if they vary, they vary in the same way.

The connection between equivalence and entailment can then be stated as the law:

φ ⇔ ψ if and only if both φ ⇒ ψ and ψ ⇒ φ

for any sentences ψ and φ. You may think of this as the ***basic law for equivalence*** because the properties of equivalence can be derived from those of entailment by using it.

The key properties are stated in the following group of laws, which hold for any sentences φ, ψ, and χ and any set Γ of sentences:

φ ⇔ φ (***reflexivity***)
φ ⇔ ψ if and only if ψ ⇔ φ (***symmetry***)
if φ ⇔ ψ and ψ ⇔ χ, then φ ⇔ χ (***transitivity***)
if Γ ⇒ φ and either φ ⇔ ψ or ψ ⇔ φ, then Γ ⇒ ψ (***conclusion replacement***)
if Γ, φ ⇒ χ and either φ ⇔ ψ or ψ ⇔ φ, then Γ, ψ ⇒ χ (***premise replacement***)

We saw in 1.4.2 that laws of reflexivity and transitivity hold for implication. But implication is not symmetric; it is one consequence of equivalence amounting to *mutual* entailment or implication. The last two laws tell us that equivalence sentences play the same role as conclusions and as premises; each of two equivalent sentences may be replaced by the other as either a conclusion or a premise without destroying validity.

Given the symmetry of equivalence the alternative *or* ψ ⇔ ϕ in the replacement laws is redundant. It is stated to emphasize that the direction of the replacement does not matter—unlike the following laws for entailment alone (which hold for any set Γ and any sentences φ and ψ):

if Γ ⇒ φ and φ ⇒ ψ, then Γ ⇒ ψ (***conclusion covariance***)

if Γ, φ ⇒ χ and ψ ⇒ φ, then Γ, ψ ⇒ χ (***premise contravariance***)

These are more general versions of a couple of ways of stating the transitivity of implication that were noted in 1.4.2. They tell us that we can replace a conclusion by something it implies and replace a premise by something that it is implied by. The terms *covariance* and *contravariance* refer to the fact that in one case the direction of replacement is same as the direction of the implication and in the other case has the opposite direction. Equivalence in either direction between a pair of sentences licenses replacement in both directions because equivalent sentences both entail and are entailed by each other. We will see later that equivalence licenses further sorts of replacement—not only of whole sentences but of their components—and this is to be expected because equivalent sentences are identical with regard to the aspects of meaning that are of concern to deductive logic.

The two forms of the definition of a ***tautology*** are as follows:

φ is a tautology	if and only if	there is no possible world in which φ is false
	if and only if	φ is true in every possible world

That is, because a tautology says nothing, it cannot be false and we have an unconditional guarantee of its truth.

Recall that, because the truth value of a tautology is fixed for every possible world, any two tautologies are equivalent. That means that laws for the specific tautology $\top$ apply also to other tautologies. The following two laws concerning $\top$ hold for every set Γ of sentences and every sentence φ:

$\Rightarrow \top$ ($\top$ ***as a conclusion***)

if $\Gamma, \top \Rightarrow \varphi$, then $\Gamma \Rightarrow \varphi$ ($\top$ ***as a premise***)

The law for $\top$ as a conclusion says that $\top$ is a valid conclusion from the empty set of premises (which we represent by leaving the left side of $\Rightarrow$ empty). It follows from the monotonicity of entailment that $\top$ is a valid conclusion from any set of premises. The law for $\top$ as a premise reflects another consequence of the fact that a tautology conveys no information: since $\top$ contributes nothing as a premise, it can be dropped from any list of premises without destroying the validity of an argument.

Glen Helman 25 Aug 2005

1.4.4. Absurdity and inconsistency

Just as a tautology is a sentence for which we have an unconditional guarantee of truth, we can define an ***absurd*** sentence as one for which we have an unconditional guarantee of falsity:

φ is absurd	if and only if	there is no possible world in which φ is true
	if and only if	φ is false in every possible world

Like tautologies absurdities are all equivalent and all have the same properties as the representative absurdity $\bot$. These properties are the opposite of those of $\top$. In particular, anything can be concluded from $\bot$ (and thus from any set of premises containing it). That is,

$\bot \Rightarrow \varphi$ ($\bot$ ***as a premise***)

for any sentence φ.

We have no law for restating conditions under which a set of sentences has $\bot$ as a valid conclusion; the property of entailing $\bot$ will be a fundamental deductive concept, an important addition to the range of ideas introduced in 1.2.2. We can, however, define this concept in terms of truth value and possible worlds. Because of the nature of entailment, when a set Γ entails $\bot$, we have a conditional guarantee of the truth of $\bot$; and, since $\bot$ cannot be true, this must be a guarantee whose conditions cannot be met. That is, a set entails $\bot$ just in case its members cannot all be true. We will say that such a set is ***inconsistent***, an idea that may be defined more formally as follows:

Γ is inconsistent	if and only if	there is no possible world in which all members of Γ are true
	if and only if	in each possible world, at least one member of Γ is false

Notice that there is no requirement here that any member of Γ be false in all possible worlds, that Γ contain an absurd sentence. There must always be an error of fact somewhere in Γ but its location may change from possible world to possible world.

Notice that an absurd sentence like $\bot$ is one that forms an inconsistent set by itself. This means that absurdity is really a special case of inconsistency. Another of moving from absurdity to a more general concept is to think of a conditional guarantee of falsity (in the way that entailment is a conditional guarantee of truth). We will speak of such a concept as ***exclusion*** or ***relative inconsistency***. It is the idea of a sentence being ***excluded by*** a set or being ***inconsistent with*** it:

φ is excluded by (or is inconsistent with) Γ	if and only if	there is no possible world in which φ and the members of Γ are all true
	if and only if	φ is false in every possible world in which all members of Γ are true

The two ways we have used to express this idea reflect the connections with other concepts that are exhibited by the two forms of its definition. On the one hand, a sentence is inconsistent with a set when adding it to the set would produce an inconsistent set. This is the idea behind the negative form of the definition. Notice that any member of an inconsistent set is inconsistent with the set formed of all other members; each member is equally liable to being singled out in this way as a scapegoat for whole set's inconsistency. The situation is symmetric in another way, too: it makes as much sense to say that Γ is inconsistent with φ as to say that φ is inconsistent with Γ since our focus is on the inconsistency of the set formed from the two.

On the other hand, the positive form of the definition describes a conditional guarantee of falsity. This is a negative analogue to entailment and the verb *exclude* provides a corresponding grammatical analogue to *entail*. When applying the earlier unqualified concept of inconsistency to pairs of sentences, we will often speak of members of the pair as ***mutually exclusive*** because, when $\{\varphi, \psi\}$ is an inconsistent set, each of φ and ψ excludes the other.

The property of inconsistency and the relation of exclusion or relative inconsistency are tied by entailment by the following basic laws:

Γ is inconsistent if and only if $\Gamma \Rightarrow \bot$ (***basic law for inconsistency***)

Γ excludes (or is inconsistent with) φ if and only if $\Gamma, \varphi \Rightarrow \bot$ (***basic law for exclusion***)

A number of further principles follow directly from the laws stated for entailment by using the laws above. In the case of simple inconsistency, the four features of entailment summarized at the end of 1.4.2 have as direct consequences the following: (i) any set with an absurd member is inconsistent, (ii) any set that entails all members of an inconsistent set is inconsistent, (iii) any sentence may be added to an inconsistent set without destroying its inconsistency, and (iv) a sentence may be dropped from an inconsistent set without destroying inconsistency provide it is entailed by the remaining members. The situation is a little more complex in the case of exclusion, but the four laws for entailment together may used to establish, among other things, the following principles for exclusion: (i) an absurd sentence is excluded by any set, (ii) an inconsistent set will exclude every sentence, (iii) a set excludes anything excluded by a set whose members it entails, (iv) a sentence excluded by a set will also be excluded by any larger set, (v) a sentence may be dropped from an excluding set without destroying the exclusion provided it is entailed by the remaining members, and (vi) a set that both entails and excludes the same sentence is inconsistent.

Entailment and exclusion are opposites in a way analogous to tautologousness and absurdity. And, although they are both deductive concepts, they combine to set bounds for other forms of reasoning. If we are confident about the accuracy of a set of data, any sentence that is entailed by the data is a reasonable conclusion but any that is excluded by the data would be unreasonable. The norms of non-deductive forms of reasoning—whether this be inductive generalization, inference to the best explanation, or non-monotonic reasoning about typical examples—draw a line between the extremes provided by entailment and exclusion. They identify reasonable conclusions that can be added to the ones entailed by the data while avoiding any that are excluded by it. In a picture,

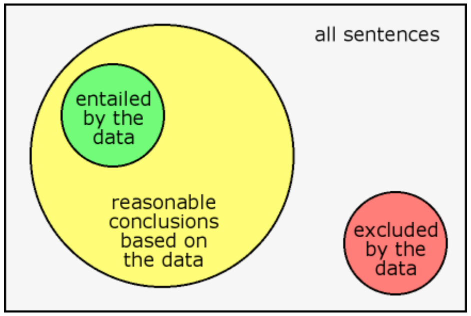

Fig. 1.4.4-1. Reasonable (but not necessarily deductive) conclusions from a body of data in relation to the sentences entailed by or excluded by the data.

Of course, the bounds provided by entailment and exclusion are firm only when we are confident in our data, and calling data into question is itself an important form of reasoning. But here, too, entailment and exclusion are relevant since they indicate ranges of claims that would never or would always lead us to call data into question.

1.4.5. Exhaustiveness

Exclusion has entailment as its natural opposite but inconsistency has a natural opposite of another sort. We will say that a set is ***exhaustive*** or that its members are ***jointly exhaustive*** under the following conditions:

Γ is exhaustive	if and only if	there is no possible world in which all members of Γ are false
	if and only if	in each possible world, at least one member of Γ is true

The term *exhaustive* reflects the fact that an exhaustive set exhausts all possibilities in the sense that any possible world is left open by at least one member of the set. That is, if we collect the possible worlds left open by each of the members of an exhaustive set and combine all these collections, we will find all possible worlds included.

Exhaustiveness is an unconditional guarantee but it applies to sentences, not individually, but as a group. That is the reason for using the qualification *jointly* when we speak about the members of the set rather than the set itself. And we often have reason to speak of the members because the most important application of this idea for our purposes is to sets with two members. In that case, we say pair of sentences φ and ψ are jointly exhaustive when the set formed of the two is exhaustive—that is, when we are guaranteed that at least one of the two is true.

The most important application of the joint exhaustiveness of pairs is in the analysis of the relation of ***contradictoriness***, which provides a kind of opposite to equivalence:

φ and ψ are contradictory	if and only if	there is no possible world in which φ and ψ have the same truth value
	if and only if	in each possible world, φ and ψ have opposite truth values

To say that a pair of sentences come with this sort of guarantee of opposite truth values is to say that we have a guarantee that they are not both true and a guarantee that they are not both false. That is, for any sentences φ and ψ,

φ and ψ are contradictory if and only if φ and ψ are both mutually exclusive and jointly exhaustive (***basic law for contradictoriness***)

Although in ordinary discourse, the term *contradictory* is often applied to sentences that are merely mutually exclusive, in logical contexts it tends to be applied only to sentences that are also jointly exhaustive. Contradictoriness will play a central role in our account of the logical properties of negation and it is crucial for this that it have ties to both inconsistency and exhaustiveness.

The final deductive concept we will consider is a very general relation that is both a conditional guarantee related to exhaustiveness and a generalization of both entailment and inconsistency. ***Relative exhaustiveness*** is a relation between sets of sentences; when it holds, we will say that one set ***renders*** the other set ***exhaustive***. Our notation for this idea will extend the use of the entailment arrow to allow a set or a list of sentences to appear on the right. Relative exhaustiveness is defined as follows:

$\Gamma \Rightarrow \Delta$	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

When a set Δ is exhaustive relative to a set Γ (that is, when $\Gamma \Rightarrow \Delta$) the collection containing of any possible world left open by any member of Δ includes all worlds in which every member of Γ is true. Notice that this is quite different from saying that Γ entails each member of Δ (a relation between sets mentioned in 1.4.2) for that would imply a conditional guarantee that all members of Δ are true while relative exhaustiveness provides instead a guarantee that at least one member of Δ is true. For this reason, we will refer to multiple sentences on the right of $\Rightarrow$ as ***alternatives*** rather than conclusions. In these terms, the definition of relative exhaustiveness tells us that a set of premises renders a set of alternatives exhaustive if and only if, in each possible world in which all the premises are true, at least one of the alternatives is true. Let us extend the idea of division from 1.4.1 to pairs of sets, saying that a possible world ***divides*** Γ from Δ when each member of Γ is true in that world while each member of Δ is false. Then we can say that $\Gamma \Rightarrow \Delta$ when there is no possible world that divides Γ from Δ. So a world that divides Γ from Δ is a counterexample to exhaustiveness of Δ relative to Γ.

There are three basic principles for relative exhaustiveness, which are rough analogues of the laws for implication and entailment. For any sentence φ and any sets Γ, Δ, Σ, and Θ of sentences:

$\varphi \Rightarrow \varphi$

if $\Gamma \Rightarrow \Delta$, then $\Gamma, \Sigma \Rightarrow \Delta, \Theta$

if $\Gamma \Rightarrow \varphi, \Delta$ and $\Gamma, \varphi \Rightarrow \Delta$, then $\Gamma \Rightarrow \Delta$

First corresponds to the reflexivity of implication and the law for premises, and the second corresponds to the law of monotonicity. The third—usually called the ***cut law***—is related to both the chain law and the law for lemmas (which are closely related to each other).

The second of these principles reflects the fact that being able to divide sets means being able to assign certain values to all their members. As a result, if it is impossible to do this for given sets Γ and Δ, it will remain impossible if members are added to either of them. The first principle reflects the fact that, because a sentence cannot be both true and false, it cannot be divided from itself. The cut law reflects the other side of the coin, the fact every sentence is either true or false. The easiest way to see the connection involves a kind of roundabout argument that is one of the reasons that the negative forms of definitions are useful. Suppose that $\Gamma \Rightarrow \Delta$ fails—i.e., that some possible world divides Γ from Δ. This world must also assign some truth value to any given sentence φ. If it makes φ false while dividing Γ from Δ, the claim $\Gamma \Rightarrow \varphi, \Delta$ will fail; and, if it makes φ true, the claim $\Gamma, \varphi \Rightarrow \Delta$ will fail. Thus, if $\Gamma \Rightarrow \Delta$ fails, then so will either $\Gamma \Rightarrow \varphi, \Delta$ or $\Gamma, \varphi \Rightarrow \Delta$. But that means that if $\Gamma \Rightarrow \varphi, \Delta$ are $\Gamma, \varphi \Rightarrow \Delta$ hold, the claim $\Gamma \Rightarrow \Delta$ must hold, too.

1.4.6. A general framework

The value of relative exhaustiveness does not lie in capturing some ordinary vocabulary for discussing deductive reasoning but instead in its ability to encompass other ideas that do. For example, entailment is the special case of relative exhaustiveness where a single alternative is rendered exhaustive. And the notation is the same. Given our notational conventions, $\Gamma \Rightarrow \varphi$ means both that φ is entailed by Γ and that the set whose only member is φ is rendered exhaustive by Γ. Inconsistency can be expressed in terms of relative exhaustiveness by way of entailment, but it can be expressed in a more direct way, too. Since there is no way that the empty set could have at least one member true in any possible world, to say that it is exhaustive relative to a set Γ is to say that the members of Γ cannot all be true. So to write $\Gamma \Rightarrow$ (with nothing to the right of the arrow) is to say that Γ is inconsistent.

The following table uses connections with relative exhaustiveness to provide notation for each of the concepts we have seen in this chapter:

Concept	*Notation*
Γ entails φ	$\Gamma \Rightarrow \varphi$
φ is a tautology	$\Rightarrow \varphi$
φ and ψ are equivalent	both $\varphi \Rightarrow \psi$ and $\psi \Rightarrow \varphi$ (abbreviated to $\varphi \Leftrightarrow \psi$)
Γ excludes φ	$\Gamma, \varphi \Rightarrow$
Γ is inconsistent	$\Gamma \Rightarrow$
φ and ψ are mutually exclusive	$\varphi, \psi \Rightarrow$
φ is absurd	$\varphi \Rightarrow$
Γ is exhaustive	$\Rightarrow \Gamma$
φ and ψ are jointly exhaustive	$\Rightarrow \varphi, \psi$
φ and ψ are contradictory	both $\varphi, \psi \Rightarrow$ and $\Rightarrow \varphi, \psi$ (abbreviated to $\Leftrightarrow \varphi, \psi$)

The double arrow notation used for equivalence and contradictoriness may suggest the more general idea of sets that render each other exhaustive—i.e., sets Γ and Δ such that $\Gamma \Leftrightarrow \Delta$. But this idea is not of much interest apart from the two special cases. In particular, sets related in this way need not have equivalent roles in deductive reasoning (as can be seen by noting that a pair of contradictory sentences form a set that renders exhaustive and is rendered exhaustive by the empty set).

The properties and relations listed above amount to guarantees that certain patterns of truth values are logically impossible. And the particular patterns are shown by the notation used, for in each case $\Gamma \Rightarrow \Delta$ says that no possible world divides Γ from Δ. It is also useful to have vocabulary for speaking about cases where *no* pattern is ruled out. We will say that a sentence is ***logically contingent*** if both truth values appear for it among possible worlds—if it is both possibly true and possibly false. This is to say that it is neither tautologous nor absurd. We will say that a pair of sentences are ***logically independent*** if each of the four patterns of truth values for the two sentences is exhibited in some possible world. This is the same as saying that the two sentences are neither mutually exclusive nor jointly exhaustive and that neither implies the other. This also implies that each of the two sentences is logically contingent.

Finally, we can extend the idea of logical independence to a set of sentences by saying that the members of a set of sentences are independent when any way of assigning a truth value to each of them is exhibited in at least one possible world. 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 that is to say that the set contains two disjoint (i.e., non-overlapping) subsets one of which renders the other exhaustive. So the members of a set Γ are logically independent when the relation of relative exhaustiveness never holds between non-overlapping subsets of Γ. (When a pair of sets do share a member, each renders the other exhaustive no matter what the sets are like otherwise.) When a set is logically independent, each member is contingent and any two of its members are logically independent, but contingency of members and independence of pairs does not by itself imply that the set as a whole is logically independent. For example, assuming the sentences *X is fast*, *X is strong*, *X has skill*, and *X has stamina* form an independent set, 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 are independent.

However, the first two taken together entail the third, so these three more complex sentences do not form an independent set.

The idea of relative exhaustiveness thus provides the resources for both surveying other deductive properties and relations and for speaking about cases where none of them hold. But the kind of symmetry built into the idea (and exhibited in the laws for it seen at the end of 1.4.5) also provides a way of describing connections among various deductive principles.

This description will also employ the idea of contradictoriness. If a pair of sentences are contradictory, then each will be true in a possible world if and only if the other is false. Since in assessing relative exhaustiveness, we consider, for each possible world, the truth of premises and the falsity of alternatives, having one of a pair of contradictory sentences as a premise comes to the same thing as having the other as an alternative. So we can remove a sentence from one side of the arrow if we add a contradictory sentence on the other side. This is stated more formally in the following, which will serve as our ***basic law for relative exhaustiveness***

if $\Leftrightarrow \varphi, \varphi'$, then $\Gamma, \varphi \Rightarrow \Delta$ if and only if $\Gamma \Rightarrow \varphi', \Delta$

We can use this principle (read from right to left) to replace alternatives by contradictory premises. If we begin with a finite set, we can eventually transform a claim of exhaustiveness into a claim of entailment. Our chief application of this will be one its simplest cases: if $\Leftrightarrow \varphi, \varphi'$, then $\Gamma \Rightarrow \varphi, \psi$ if and only if $\Gamma, \varphi' \Rightarrow \psi$. That is, a pair of sentences φ and ψ are rendered exhaustive by a set Γ of premises if and only one of the pair, ψ, is entailed by Γ together with a sentence φ′ that is contradictory to the other member of the pair. Since the order of a list of alternatives does not matter (so saying that $\Gamma \Rightarrow \varphi, \psi$ is the same as saying that $\Gamma \Rightarrow \varphi, \psi$), this law tells us that we can drop either of a pair of alternatives if we add to the premises a sentence contradictory to the alternative we drop.

The properties of ⊤ and ⊥ take a particularly symmetric form when stated in terms of relative exhaustiveness.

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

That is, while $\top$ contributes nothing as a premise and may be dropped, it is sufficient by itself as the only alternative (no matter how small our set of premises). And while $\bot$ is sufficient by itself as a premise (no matter how small the set of alternatives is), it contributes nothing as an alternative and may be dropped. The symmetry here might be traced to that of relative exhaustiveness: since $\top$ and $\bot$ are contradictory, the principles on each diagonal are connected by the basic law for relative exhaustiveness.

However, there is a more general idea behind this symmetry. To take the simplest case above, we might state the lower left and upper right as follows:

$$\bot \Rightarrow$$
$$\top \Leftarrow$$

(where an arrow running right to left is understood to have its alternatives on its left and its premises on its right). That is, the difference lies in interchanging Absurdity and Tautology and reversing the direction of the arrow—or, what comes to the same thing, interchanging premises and alternatives. If we apply the same transition to the principle at the upper left we get

$$\text{if } \Gamma, \bot \Leftarrow \Delta, \text{ then } \Gamma \Leftarrow \Delta$$

or, rewriting so the arrows run left to right (without change of premises and alternatives),

$$\text{if } \Delta \Rightarrow \bot, \Gamma, \text{ then } \Delta \Rightarrow \Gamma$$

The latter differs from the principle for Absurdity as an alternative on the lower right above only in the interchange of Γ and Δ; and, since each could be any set, their interchange does not change the content of the principle. The possibility of this sort of transformation can be expressed by saying that $\top$ and $\bot$ on the one hand and *premise* and *alternative* on the other constitute pairs of ***dual*** terms. We will run into other pairs of dual terms later.

1.4.s. Summary

Entailment may be defined in two equivalent ways, either as the relation that holds when the conclusion is false in no possible world in which all the premises are true or as the relation which holds when the conclusion is true in all such worlds. The first approach can be stated more briefly by saying that an argument is valid when no world divides the premises from conclusion; a world that does divide premises from conclusion is a counterexample to the claim of entailment or validity.

The idea of entailment can also be understood by way of certain laws governing it. For example, if we limit ourselves to single-premised arguments—i.e., to implication—the relation is reflexive and transitive. The law for premises and the chain law are analogous principles that apply to entailment more generally. Entailment also obeys a principle of monotonicity asserting that a premises may always be added without undermining entailment (something does not hold for many forms of non-deductive inference) and a law for lemmas that tells us that a premise may dropped when it is entailed by other premises.

Other properties and relations besides entailment can be given pairs of negative and positive definitions. This is true for the ideas of logical equivalence and tautologousness introduced in 1.2.2. Sentences are equivalent when they entail each other, and this basic law implies that equivalence is symmetric as well as reflexive and transitive. Moreover, equivalent statements may replace one another either as premises or conclusions of an argument without affecting its validity (unlike the case of entailment which obeys only the weaker laws of conclusion covariance and premise contravariance). The laws governing tautologies are most easily stated by focusing on the particular case of Tautology $\top$. For example, $\top$ is always a valid conclusion, but it never contributes anything as a premise and may be freely added to or dropped from the premises without changing an argument's validity.

The definitions of absurdity are in a way opposite those of tautologousness and having Absurdity $\bot$ as a premise, like having a $\top$ as a conclusion, makes an argument valid. When an argument with $\bot$ as its conclusion is valid, its premises form an inconsistent set. Inconsistency is the fundamental negative concept of deductive logic and the relative concept of being excluded by or inconsistent with a set is a kind of negative opposite to entailment. As a relation between pairs of sentences relative inconsistency is symmetric and such sentences are said to be mutually exclusive. Although inconsistency is a fundamental deductive property, it is one we will establish by using laws that describe it in terms of entailment.

The negative concepts of inconsistency and exclusiveness are opposed in one way to entailment and in another way to exhaustiveness. Contradictory sentences are ones that are bound to differ in truth value; such sentences can be characterized as both mutually exclusive and jointly exhaustive. Exhaustiveness can be conditional and this is a relation between sets that generalizes entailment to allow a set of alternatives rather than a single conclusion. This relation fails when a possible world divides its premises from its alternatives by making the former all true and the latter all false. Relative exhaustiveness obeys cut law which are analogous to, but more symmetric than, the principles governing entailment.

Relative exhaustiveness has an important role in unifying the concepts of deductive logic. All the ones we have seen can be described as special cases of it. We can also use it to describe the absence of deductive properties and relations, whether this is the logical contingency of individual sentences or the logical independence of pairs or larger sets. Laws governing relative exhaustiveness in its own right tend to be symmetric in form. Relative exhaustiveness can be connected with entailment by law employing the idea of contradictoriness. This law exhibits a kind of symmetry that is found also in the laws for $\top$ and $\bot$ stated in terms of relative exhaustiveness. Their symmetry can also be seen as one instance of a relation of duality that we will encounter in other cases as well.

1.4.x. Exercise questions

1. Restate each of the following claims about logical properties and relations, putting into symbolic notation those stated in English and into English those stated in symbolic notation:

- **a.** $\varphi, \psi \Rightarrow \chi$
- **b.** φ is entailed by ψ
- **c.** $\varphi \Leftrightarrow \varphi$
- **d.** $\psi \Rightarrow$
- **e.** φ is inconsistent with Γ
- **f.** φ is entailed by the members of Γ together with ψ

2. The following steps lead you to construct a proof of the law for lemmas

if $\Gamma,\varphi \Rightarrow \psi$ and $\Gamma \Rightarrow \varphi$, then $\Gamma \Rightarrow \psi$

Begin by supposing that $\Gamma, \varphi \Rightarrow \psi$ and $\Gamma \Rightarrow \varphi$ are both true. We want to show that, under this supposition, $\Gamma \Rightarrow \psi$ is also true. To do that, we consider any possible world w in which all members of Γ are true and try to show that ψ is true in w.

- **a.** Our supposition that $\Gamma, \varphi \Rightarrow \psi$ and $\Gamma \Rightarrow \varphi$ are both true combined with what we know about w enables us to conclude that φ is true. Why?
- **b.** Adding the information that φ is true in Γ to what we already knew, we can conclude that ψ is true. Why?

So, knowing about w only that all members of Γ were true, we are able to conclude that ψ is true. And that shows us that ψ is true in every world in which all members of Γ are true, which means that $\Gamma \Rightarrow \psi$.

Another approach to proving the law is to show that $\Gamma \Rightarrow \psi$ fails only if at least one of $\Gamma, \varphi \Rightarrow \psi$ and $\Gamma \Rightarrow \varphi$ fails. The following three steps show this:

- **c.** Suppose that w is a counterexample to $\Gamma \Rightarrow \psi$. What truth values do ψ and the members of Γ have in w?
- **d.** What truth values are needed to have a counterexample to $\Gamma \Rightarrow \varphi$? To have a counterexample to $\Gamma, \varphi \Rightarrow \psi$?

e. The world w from **c** will be a counterexample to either $\Gamma, \varphi \Rightarrow \psi$ or $\Gamma \Rightarrow \varphi$. Why?

1.4.xa. Exercise answers

1. **a.** φ and ψ together entail χ

b. $\psi \Rightarrow \varphi$

c. φ is equivalent to itself

d. ψ is absurd
or: ψ taken by itself forms an inconsistent set

e. $\Gamma, \varphi \Rightarrow$
or: $\Gamma, \varphi \Rightarrow \bot$
(Strictly speaking, $\Gamma, \varphi \Rightarrow \bot$ expresses entailment rather than inconsistency, but it is true if and only if φ is inconsistent with Γ.)

f. $\Gamma, \psi \Rightarrow \varphi$

2. **a.** We have supposed that $\Gamma \Rightarrow \varphi$. That is, we have supposed that φ is **T** in any possible world in which all members of Γ are **T**. But w is a world in which all members of Γ are **T**, so φ, too, must be **T** in w.

b. We now know that φ and all members of Γ are **T** in w. But we supposed that $\Gamma, \varphi \Rightarrow \psi$ and we now know that all the premises of this entailment are **T** in w, so ψ also must be **T** also.

c. For w to be a counterexample to $\Gamma \Rightarrow \psi$, it must make give ψ the value **F** and give all the members of Γ the value **T**.

d. A counterexample to $\Gamma \Rightarrow \varphi$ must give φ the value **F** and give all the members of Γ the value **T**. A counterexample to $\Gamma, \varphi \Rightarrow \psi$ must give ψ the value **F** while giving φ and all the members of Γ the value **T**.

e. We know that w gives ψ the value **F** and gives all the members of Γ the value **T**. But it also must make φ either **T** or **F**. If it does the former, it is a counterexample to $\Gamma, \varphi \Rightarrow \psi$; and if it does the latter, it is a counterexample to $\Gamma \Rightarrow \varphi$.