

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.

Glen Helman 03 Aug 2010

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 its 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**.




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.

Glen Helman 03 Aug 2010

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.

Glen Helman 03 Aug 2010

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 $\phi \models \psi$ 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 $\phi \models \psi$ and $\psi \models \phi$ —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 \simeq (*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 $\phi \simeq \psi$, 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 $\phi \models \psi$ 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 $\phi \models \psi$ but $\psi \not\models \phi$ —the possibilities left open by ψ will include all those left open by ϕ (because $\phi \models \psi$) 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 $\phi \models \psi$ and $\psi \models \chi$, then $\phi \models \chi$. 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 $\phi \simeq \psi$ then $\psi \simeq \phi$. 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.

Glen Helman 03 Aug 2010

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 \top (*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 \top will hold for all tautologies, and we will often simply speak of \top 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 \perp (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 \top 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 \perp can stand at the bottom of any ordering by logical strength.

Any sentence implying \perp is thus equivalent to it and is itself absurd. More generally, the idea of entailing \perp provides way characterizing inconsistency. That is, we can have $\Gamma \vDash \perp$ 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 \perp . 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 \perp . 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.

Glen Helman 03 Aug 2010

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 $2 \times 2 \times 2 \times 2 = 16$ 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.

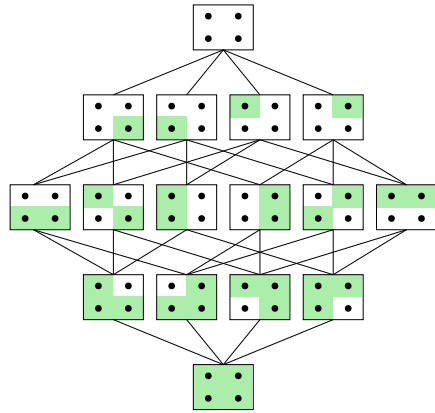


Fig. 1.2.5-3. The propositions generated by 4 possible worlds depicted as regions in one logical space (the repeated rectangle) and as points in another (the overall diagram).

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 (\top and \perp) 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 \perp 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 $\{\phi, \psi\}$ 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 reversibility 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 $\phi \Delta \psi$ 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 \perp .

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 Δ 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 $\phi \text{X} \psi$ 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 diametrically 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 \perp and \top , 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:

<i>Relation</i>	<i>pattern ruled out</i>
ϕ implies ψ ($\phi \vDash \psi$)	ϕ is T ψ is F
ϕ is implied by ψ ($\psi \vDash \phi$)	ϕ is F ψ is T
ϕ and ψ are mutually exclusive ($\phi \Delta \psi$)	ϕ is T ψ is T
ϕ and ψ are jointly exhaustive ($\phi \nabla \psi$)	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 means 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.

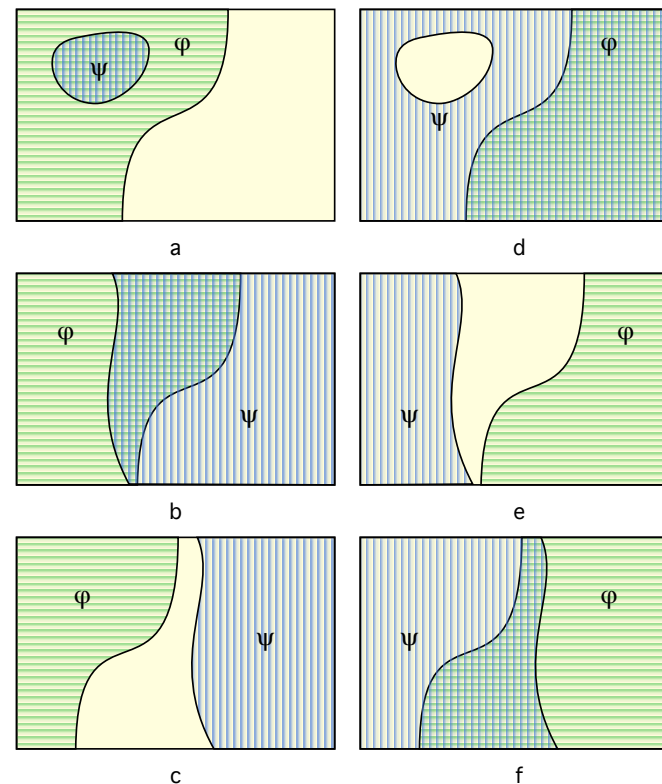


Fig. 1.2.6-1. Three relations between sentences ϕ and ψ . (a, d) ϕ implies ψ . (b, e) ϕ and ψ are mutually exclusive. (c, f) ϕ and ψ are jointly exhaustive. On the left, regions ruled out by sentences are hatched—horizontally in green for ϕ and vertically in blue for ψ . The regions left open by ϕ and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

When $\phi \vDash \psi$ (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).

Glen Helman 03 Aug 2010

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, $\phi \models \psi$ when ϕ says everything that is said by ψ . When sentences imply each other, they say the same thing, and we say they are equivalent, a relation for which we use the sign \simeq . When $\phi \models \psi$ 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 \perp . 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 \perp 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 \perp 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 \perp 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 Δ). 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 ∇). When these two relations both hold, sentences are contradictory (a relation for which we use the sign \times). 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 Tuesday** \models **The package will arrive next week**
 - b. **The package will arrive next week** \models **The package will arrive next Tuesday**
 - c. **The package will arrive next Tuesday** Δ **The package will arrive next week**
 - d. **The package will arrive next Tuesday** Δ **The package will arrive next Wednesday**
 - e. **The package will arrive before next Tuesday** ∇ **The package will arrive after next Tuesday**
 - f. **The package will arrive next Tuesday or before** ∇ **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** \times **The package will arrive after next Wednesday**
 - j. **The bridge will open before June** \times **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. $\phi \models \phi$	b. $\phi \models \top$	c. $\phi \models \perp$
	d. $\top \models \phi$	e. $\perp \models \phi$
f. $\phi \nabla \phi$	g. $\phi \nabla \top$	h. $\phi \nabla \perp$

- i. $\phi \Delta \phi$ j. $\phi \Delta \top$ k. $\phi \Delta \perp$
 l. $\phi \simeq \phi$ m. $\phi \simeq \top$ n. $\phi \simeq \perp$
 o. $\phi \Sigma \phi$ p. $\phi \Sigma \top$ q. $\phi \Sigma \perp$

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 $\phi \vDash \psi$ and $\psi \vDash \chi$, we cannot have ϕ true and χ false, so $\phi \vDash \chi$ (this is the transitivity of implication). However, no other patterns for ϕ and χ are ruled out, so “ $\phi \vDash \chi$ ” is the most we can say on the basis of the given information, and it can be entered in the upper left cell.

	$\psi \vDash \chi$	$\chi \vDash \psi$	$\psi \simeq \chi$	$\psi \Delta \chi$	$\psi \nabla \chi$	$\psi \Sigma \chi$
$\phi \vDash \psi$						
$\psi \vDash \phi$						
$\phi \simeq \psi$						
$\phi \Delta \psi$						
$\phi \nabla \psi$						
$\phi \Sigma \psi$						

Glen Helman 03 Aug 2010

1.2.xa. Exercise answers

1.
 - a. *The package will arrive next Tuesday* 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 Tuesday* 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. $\phi \vDash \phi$ holds always because ϕ cannot fail to be true if it is true

- b.** $\varphi \models \top$ holds always because \top cannot fail to be true no matter what φ is like
- c.** $\varphi \models \perp$ holds only when φ is absurd because, if there is any possibility of φ being true, there is a possibility of \perp being false when φ is true
- d.** $\top \models \varphi$ 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.** $\perp \models \varphi$ holds always because there is no possibility of \perp being true so no possibility of φ being false when \perp is true
- f.** $\varphi \nabla \varphi$ 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.** $\varphi \nabla \top$ holds always because \top covers all possibilities by itself, so it certainly exhausts them when taken together with φ
- h.** $\varphi \nabla \perp$ holds only when φ is a tautology because, since \perp leaves open no possibilities, it contributes nothing to exhausting them all and φ must do that all by itself
- i.** $\varphi \Delta \varphi$ 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.** $\varphi \Delta \top$ 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.** $\varphi \Delta \perp$ holds always because, since \perp 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.** $\varphi \simeq \perp$ 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.** $\varphi \boxtimes \varphi$ never holds because no sentence can be both true and false at the same time
- p.** $\varphi \boxtimes \top$ 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.** $\varphi \boxtimes \perp$ 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 .

	$\psi \models \chi$	$\chi \models \psi$	$\psi \simeq \chi$	$\psi \Delta \chi$	$\psi \nabla \chi$	$\psi \boxtimes \chi$
$\varphi \models \psi$	$\varphi \models \chi$	—†	$\varphi \models \chi$	$\varphi \Delta \chi$	—†	$\varphi \Delta \chi$
$\psi \models \varphi$	—*	$\chi \models \varphi$	$\chi \models \varphi$	—*	$\varphi \nabla \chi$	$\varphi \nabla \chi$
$\varphi \simeq \psi$	$\varphi \models \chi$	$\chi \models \varphi$	$\varphi \simeq \chi$	$\varphi \Delta \chi$	$\varphi \nabla \chi$	$\varphi \boxtimes \chi$
$\varphi \Delta \psi$	—*	$\varphi \Delta \chi$	$\varphi \Delta \chi$	—*	$\varphi \models \chi$	$\varphi \models \chi$
$\varphi \nabla \psi$	$\varphi \nabla \chi$	—†	$\varphi \nabla \chi$	$\chi \models \varphi$	—†	$\chi \models \varphi$
$\varphi \boxtimes \psi$	$\varphi \nabla \chi$	$\varphi \Delta \chi$	$\varphi \boxtimes \chi$	$\chi \models \varphi$	$\varphi \models \chi$	$\varphi \simeq \chi$

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