1.2. What is said: possibilities of truth and falsity

1.2.0. Overview

In 1.1.3, we noted two closely related properties of a deductive inference: it is a transition from premises to conclusion that is free of any new risk of error, and the information provided by its conclusion is already present in its premises. The second of these properties points to the general idea of the information content of a sentence. This provides a useful perspective from which to consider many aspects of deductive reasoning, and we will explore it further in this section.

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 characterize the semantic content of a sentence, to give an account of what it says.

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

1.2.4. A model of language
One simple picture of language sees it as a device for communicating semantic content, for sharing propositions.

Glen Helman 01 Aug 2004

1.2.1. Truth values and possible worlds

When an inference is deductive—when its premises and conclusion constitute a valid argument—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. This means that the ideas of truth and falsity have a central place in our discussion of 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 the pattern of truth values for its premises and conclusion shown in Figure 1.2.1-1 is impossible. That is (using some of the other terminology we have available), a conclusion is entailed by a set of assumptions when it is impossible for the truth value of the conclusion to be **F** when each of the assumptions has the truth value **T**.

premise	**T**
premise	**T**
⋮	⋮
premise	**T**
conclusion	**F**

Fig. 1.2.1-1. The pattern of truth values that is impossible when an argument is valid.

Since to speak of a risk of error is to speak of a possibility of error, it is also useful to have some vocabulary for speaking of possibility and impossibility. The sort of impossibility in question when we speak of entailment is a very strong sort that may be referred to as ***logical 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. Knowledge that someone was in Boston at 10:00 A.M. EST on January 1, 1980 renders the conclusion that she was not in New York at 10:01 A.M. EST that same day unassailable from the point of view of common sense. But this is not a deductively valid conclusion, for presence in those places at those times is logically possible and it would remain so if the difference in time was made so small as to imply travel faster than the speed of light.

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.

The term *world* could be misleading since it might suggest something that a physicist would describe as a system that can have many different states; in our usage, however, a possible world is less like a system than a particular one of its states or a particular history tracing these states through time; and these are not states of one among many systems but of a global system including everything there is. That is, the phrase *possible world* could be restated more accurately, though less conveniently, using the phrase *possible history of the universe*.

To say that something is logically *im*possible is to say that it is true in *no* logically possible world. The weaker claim of physical impossibility says only that a claim is found to be true nowhere in the narrower range of physical possibilities or physically possible worlds. But what sets the limits of the particular range of possible worlds considered in deductive logic? If we move beyond the bounds of physical laws, what is left to stop us from saying anything? I claimed earlier that travel faster than the speed of light was logically possible. Is the same true for presence in two places at the same time or travel backward in time? We will not need specific answers to such questions, but we should have some general sense of the basis for answering questions like them.

Unfortunately, there is no uncontroversial view of this. However, the issues at dispute between the alternatives are largely independent of the content of deductive logic, so we will consider only one respectable account without giving equal attention to its rivals. The association between logic and language mentioned in 1.1 provides this account. It traces the norms of deductive inference to rules governing the meanings of sentences, and these rules might be held to set the limits of the range of logical possibilities. So one way to answer the question of whether a claim of impossibility is a claim of *logical* impossibility is ask whether the description of what is said to be impossible is really disallowed by the semantic rules of the language in which it is stated. To decide whether being in two places at the same time is a logical impossibility is to decide whether a sentence like *I was in two places at the same time* is even a meaningful description. To say that being in two places at the same time was logically impossible would not be to describe such a situation and go on to say that it is impossible but instead to claim that no such situation could even be coherently described.

Glen Helman 01 Aug 2004

1.2.2. Truth conditions and propositions

If we use the terminology of the last section, we can say that the conclusion of an inference brings with it no new risk of error when there is no possible world in which it has the truth value **F** while the premises all have the value **T**. This means that, when judging whether a relation of entailment holds between premises Γ and a conclusion ϕ, it is enough know the truth value of ϕ and the truth values of the members of Γ in each possible world. We will assume that a sentence always has one or the other of the truth values **T** and **F**, so we will know the truth value of a sentence in each possible world if we know which possible worlds are the ones in which it is true. We will refer to this aspect of the meaning of a sentence as the ***proposition*** it expresses.

A proposition is thus an abstract entity that encapsulates the aspects of a sentence's meaning which are relevant to questions of entailment (and deductive reasoning generally). The relations of entailment that hold between sentences derive from relations among the propositions they express. This can help in thinking about the analogies between mathematics and logic: deductive logic is concerned with propositions in the way algebra and the theory of functions are concerned with numbers. Since the relation of entailment depends only on the propositions expressed by sentences, we will use the term *entailment* both for a relation between sentences and for the corresponding relation between the propositions expressed by them. And this relation among propositions provides the fundamental order for the range of propositions in the way that the relation ≤ provides the fundamental order of the numbers. We will look further at this analogy between propositions and numbers in 1.2.3, but first we will look in more detail at the idea of specifying a truth value in each of the various logically possible worlds.

The proposition expressed by a sentence determines its truth value in each of the full range of logical possibilities and, in this sense, specifies its ***truth conditions***. A sentence leaves open a certain range of possibilities, those in which it is true, and rules out others, those in which it is false. Comparing the ranges of possibilities left open and ruled out by sentences—i.e., comparing their truth conditions—provides a way of comparing the their contents, the strengths of the claims they make. Consider the successively more specific statements:

The package will arrive sometime
The package will arrive next week
The package will arrive next Wednesday

All leave open many possibilities since they say nothing about, for example, the time of day the package will arrive, the form of transportation used, or its condition on arrival (to say nothing of the myriad of other bits of information in a full possible history of the universe); but each statement rules out some possibilities left open by the statement above it. And in general, the less information contained by some statement, the more possibilities it leaves open; the more it says, the more possibilities it rules out.

Let us look at entailment from this perspective. Begin with our recent characterization of risk-free inference: a sentence ϕ is entailed by a set Γ if there is no possible world in which ϕ is false while all members of Γ are true. This is to say that ϕ does not rule out any possibility that is left open by all members of Γ. Now, a possibility is left open by all members of Γ if there is no member that rules it out, so we can say that ϕ is entailed by Γ if it does not rule out any possibility that is not already ruled out by at least one member of Γ. If we think of possibilities ruled out as an indication of content, this amounts to the claim that ϕ is entailed by Γ if its content does not extend beyond the content already found among the members of Γ. And that brings us back to another characterization of deductive inference—namely, as the extraction of information, for we are saying that ϕ is entailed by Γ if its informational content is already to be found among the members of Γ taken together.

Of course, two sentences may provide exactly the same information, differing only 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 rule out exactly the same possibilities for Sam's residence and thus express the same proposition. 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).

Logically equivalent sentences divide logical space between possibilities ruled out and possibilities left open in exactly the same way. Since they have the same content they are valid conclusions from the same premises and can be freely substituted for each other as premises of an argument without changing its validity. In particular, logically equivalence sentences will entail each other. Also, when sentences are not logically equivalent, there must be some possible world where one is true and the other false. This cannot happen if a pair of sentences entail each other, so any pair of sentences that entail each other must be logically equivalent. Our notation for logical equivalence—the sign ⇔

(***left right double arrow***)—reflects this characterization of equivalence as mutual entailment.

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. That is, any two tautologies are logically equivalent. It is convenient to distinguish a particular tautology and mark it by special notation, which is specified simply as a sentence that says nothing. We will call this sentence ***Tautology*** and use the sign ⊤ (***down tack***) as our notation for it. Notice that ⊤ entails nothing except other tautologies (since any other sentence would add information) and is entailed by any set of premises (because it will not add information to any set of sentences).

At the other extreme of truth conditions is a sentence that rules out all possibilities. The fact that this is the opposite of 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. 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 like this 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. Also as with tautologies, we will introduce a particular example of an absurdity, named ***Absurdity***, and we use the special notation ⊥ (the perpendicular sign, or ***up tack***) for it. Taken as a single premise, ⊥ entails every sentence. And that is another way of indicating the problem with absurd sentences: they try to go off in every direction at once. A set of premises that entails ⊥ need not contain an absurd sentence but, taken together, the sentences of such a set rule out all possibilities, and the set is therefore equally problematic. Although such a set is itself useless as a source of information, the idea of set of sentences whose members cannot all be true together will prove to be important, and ⊥ will serve us by helping us to characterize this sort of set in terms of entailment.

If we add Tautology and Absurdity to our earlier English examples, we get the following series of successively stronger claims:

⊤
The package will arrive sometime
The package will arrive next week
The package will arrive next Wednesday
⊥

It's not hard to find propositions that lie between those expressed by each successive pair here. In particular, a proposition that leaves open only one of the possible worlds in which the package arrives next Wednesday would lie between *The package will arrive next Wednesday* and ⊥. As before, each sentence entails itself and any above it. It is worth noting that, in this example, apart from the relations involving ⊤ and ⊥, the relations of entailment will not follow from the laws we will consider in later chapters. For those laws will concern the logical words *and*, *not*, etc., while these entailments depend on connections among the time designations *next Wednesday*, *next week*, and *sometime*.

1.2.3. 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 and perhaps the most notable difference between them is that, while numbers can be ordered in a linear way, the collection of propositions has a more complex structure. Although our examples in 1.2.2 formed a single chain from ⊤ to ⊥, it should be clear that we could have gone in many different directions to add content to any of the propositions that was not already absurd.

This metaphor of many directions suggests a space of more than one dimension. Although the structure of a collection of propositions 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 to elucidate its structure. These metaphors and can be associated with the term ***logical space*** introduced by the philosopher Ludwig Wittgenstein (1889-1951). We will actually use two different sorts of spatial metaphor. In one metaphor, 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. In the other metaphor, possible worlds are the dimensions of the space and propositions are its points. In cases where there are finitely many possible worlds and propositions, the points of such a space can be thought of as the vertices of a figure that might be depicted in ordinary 2- or 3-dimensional space.

To see an 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 1.2.3-1 illustrates two such propositions; each rules out two of the four possibilities (in the shaded and hatched areas) and leaves open two others: the sentence ϕ rules out the two possibilities at the bottom of the diagram and 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.

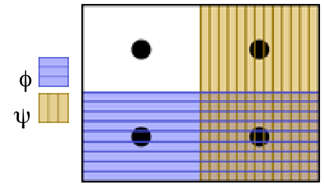

Fig. 1.2.3-1. The possibilities (shaded and hatched) ruled out by two propositions.

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 × 2 × 2 × 2 = 16 propositions in all. Of these, six rule out two of the four possible worlds.

We can illustrate all 16 of these propositions by using a logical space of the second sort. Figure 1.2.3-2 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. You can imagine that the propositions ϕ (which appears at the left) and ψ (near the center) are the two propositions depicted in Figure 1.2.3-1.

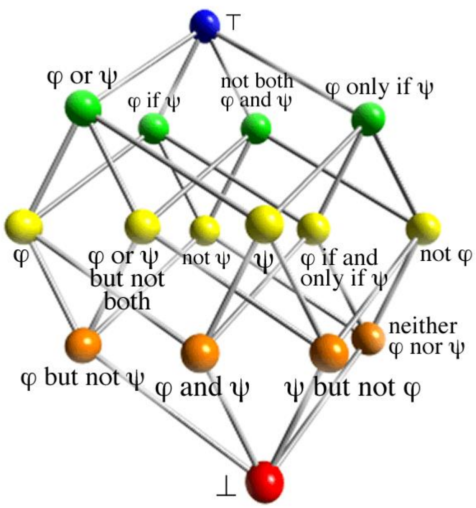

Fig. 1.2.3-2. The sixteen distinct propositions when there are 4 possible worlds.

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. Each of the four propositions immediately above Absurdity then leaves open just one possible world. Lines connecting worlds that differ with respect to a given proposition are parallel (in the 3-dimensional figure, not in its 2-dimensional projection). This is the sense in which possible worlds are the dimensions of this logical space.

The relation between the two sorts of diagram can be seen by replacing each proposition in Figure 1.2.3-2 by its representation using a diagram of the sort illustrated in Figure 1.2.3-1. Putting the two sorts of illustration together in this way gives us the following picture of the same 16 propositions.

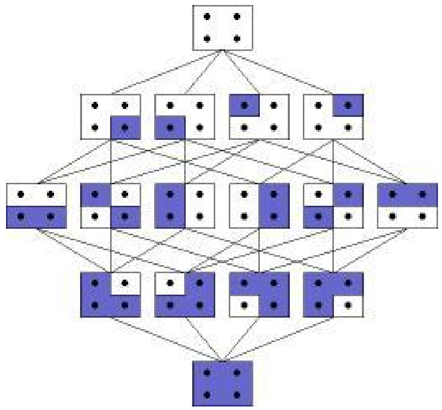

Fig. 1.2.3-3. The propositions for 4 possible worlds displayed as both regions and points.

The spacing of the nodes differs between Figures 1.2.3-2 and 1.2.3-3 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 1.2.3-1.

The whole structure of Figure 1.2.3-2 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 1.2.3-2 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 1.2.3-2 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 $\bot$) 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 1.2.3-2) will more and more closely approximate a bell curve. But the bell shape of the curve will also narrow significantly. When there are only 10 possible worlds, the two propositions at the extremes still make up 0.2% of all propositions. But once there are 100 possible worlds, 99.9% of the propositions are found in the middle of third of the degrees of strength and with 1000 possible worlds 99.9% are found in the middle 1/10 (i.e., among those that rule out between 450 and 550 possible worlds). Eventually, a graph of the distribution would appear as a single spike in middle. And the height of this spike would be so great relative to number of different grades of strength that, were the scales of the two axes the same, the whole graph would be this single line. In short, the space of propositions is so far from a being a line from $\top$ to $\bot$ that the distance from $\top$ to $\bot$ vanishes in comparison to the spread of the space near its mid-point.

Of course, with any substantial number of possible worlds, we are dealing with vast numbers of propositions (for example, already about 10^{301} when there are only 1000 possible worlds), and tiny percentages of vast numbers can still be enormous. So there is no shortage of propositions whose strength lies towards the extremes. For example, again with 1000 possible worlds, there are better than 8 trillion propositions that leave open 5 or fewer of these 1000 possibilities.

If we suppose that there are infinitely many possible worlds, these counts and percentages cease to make sense. But it is possible to distinguish different infinite quantities and the propositions that leave open or rule out a finite number of possibilities will be less numerous in this sense than those both leave open and rule out an infinite number. For example, if we suppose that there are as many possible worlds as there are integers (of which there are $\aleph_0$, ***aleph-null***—the first of a series of infinite quantities labeled with the ***alef symbol*** $\aleph$), we might arrange the grades of strength as follows:

number of worlds ruled out	left open	number of propositions
0	$\aleph_0$	1
1	$\aleph_0$	$\aleph_0$
2	$\aleph_0$	$\aleph_0$
⋮	⋮	
$\aleph_0$	$\aleph_0$	$2^{\aleph_0}$
⋮	⋮	
$\aleph_0$	2	$\aleph_0$
$\aleph_0$	1	$\aleph_0$
$\aleph_0$	0	1

Each grade other than the extremes and the central one contains $\aleph_0$ propositions. If we group all grades other than the central one together, we still have $\aleph_0$ (one of the oddities about the arithmetic of infinite quantities). The population of the central grade is the larger infinite quantity $2^{\aleph_0}$, which is the size of the set of all real numbers and is also the size of the set of all propositions formed from a collection of $\aleph_0$ possible worlds. So in this space, the single central grade of strength has as many propositions as the whole space, and both the size of the next largest grades and the total number of grades is a smaller infinite quantity that vanishes in comparison.

1.2.4. A model of language

The idea of content we have been exploring suggests a simple picture of the nature of language and the way it is used. According to this picture, each sentence has truth conditions that are determined by the semantic rules of the language. These truth conditions settle the truth value of the sentence in each possible world and thus determine the proposition it expresses. The proposition expressed by a sentence is its meaning. The meanings of other sorts of expressions—words, phrases, clauses—is to found by identifying the contributions they make to the propositions expressed by sentences containing them.

From this point of view, the function of language is to convey propositions. Just as the information content of a sentence is to be found by considering the range of possible worlds it rules out, the information that a person possesses is to be found by considering the possible worlds that he or she is able to rule out. The more I can rule out, the more information I have; and the kind of information I have is determined by the particular worlds I can rule out. This means that the sum total of my knowledge can be thought of as a proposition.

Our aim in acquiring information could be described as an attempt to distinguish the actual state of the world among the various alternative possibilities—in short, to locate the actual world within the space of all possible worlds. The proposition representing our knowledge goes some distance towards in ruling out some possibilities. But we can expect to still leave many open, and the actual world could be any of them. A new proposition helps us go further by ruling out a whole region of logical space. It can be added to the proposition representing someone's existing knowledge to rule out further possible worlds and narrow the range within which the actual world might lie. Consequently, conveying a proposition to someone can help him or her determine the precise location of the actual world.

When we acquire information, we are able to add the content of a new proposition to the content of the proposition expressing what we already know. And we can generate information to give to others by delimiting a region within the total area we know to be ruled out. Ideally, perhaps, we would simply convey the whole of what we know; but language and, more generally, the various costs of transmitting information limit our ability to do this. Instead we select a proposition from among those entailed by what we know, balancing the costs of transmission against the value a proposition might have to someone else.

If I assert a sentence, I commit myself to its truth and thus to the actual world being one of the possibilities it leaves open; equivalently, I commit myself to the actual world not being one of the worlds it rules out. So someone may garner information from my assertion by accepting it as true and using the line it draws between the possibilities it leaves open and those it rules out to further pin down the location of the actual world.

This is illustrated in the following artificial example. Initially, the person on the left is able to rule out regions at the left and right of logical space as possibilities for the actual world while the person on the right is able to rule out regions at the top and bottom.

Fig. 1.2.4-1. An animation of a conversation in which information is shared. The button > will play the full conversation while the buttons φ, ψ, χ, and θ will each play one of its four stages. The buttons |< and >| move to the initial and final state, respectively.

The animation then shows a conversation in which each party in turn notices the truth of the one the sentences ϕ, ψ, χ, and θ and asserts it. The other person accepts this assertion as true and adds its content to the region ruled out by his or her beliefs. At the end the conversation, the two people share the ability to rule out a region around the boundary of logical space though their beliefs still differ in the shape of the region left open in the middle.

As was noted above, one constraint on this sort of communication is the fact that not every proposition entailed by what we believe is expressible by a sentence, not even in principle (there are too many propositions) let alone in practice. This is suggested in Figure 1.2.4-1 by the fact that only a very limited range of ways of dividing logical space are used by the four sentences used to convey information; each sentence illustrated divides logical space simply by a vertical or a horizontal line.

This constraint figures into the cooperative character of the conversation. For example, it is only after learning the truth of ϕ that the second person is in a position to express a proposition dividing logical space in this way. And the sentence, ψ, that he or she then asserts puts the first person in a position to find an appropriate proposition (expressed by χ) in a region of logical space where there was none before. One of the chief functions of deductive inference—and one expressed by the traditional concept of a syllogism mentioned in 1.1.2 — is to “put 2 and 2 together,” to combine disparate sources of information and extract information that, while it does not go beyond what the sources provide when combined, does go beyond what any of them provides individually.

Entailment figures in the picture we have been considering in one way by setting bounds on the range of sentences that convey information we can sincerely share: we can sincerely assert only sentences entailed by what we believe. But entailment can be seen to play a second role also. We assert things that we think will be of interest to our audience. But the full content of what we assert may not be of interest to everyone we assert it to. Consequently, someone listening to us may extract only some of the information we provide in order to add it to his or her beliefs. While, ideally, we might like to add the full content of what we hear to our beliefs, our ability to store information is limited, and what we do store is determined by our interests.

Glen Helman 01 Aug 2004

1.2.s. Summary

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 those situations whose description is permitted by the semantic rules of the language.

Information about the truth values of a sentence in all possible worlds is information about its truth conditions, and these truth conditions determine the proposition it expresses. Sentences that have the same truth conditions, that express the same proposition, are logically equivalent (and idea for which we use the sign $\Leftrightarrow$). From the point of view of deductive logic, equivalent sentences have the same properties and stand in the same relations to other sentences. Entailment is one relation among sentences that depends only on the propositions they express. A conclusion ϕ is entailed by a set Γ of premises when ϕ rules out only worlds that are ruled out by at least one member of Γ. This is a way of saying that the content of ϕ does not exceed that of the members of Γ taken together, so entailment is a comparison of sentences in terms of their informational content. At one extreme are tautologies, which rule out no possibilities and thus have no content. All tautologies are equivalent and we will distinguish one, Tautology, for which we use the notation $\top$. At the other extreme are sentences that rule out all possibilities. Such sentences are absurd and all are equivalent to the single representative Absurdity, for which we use the notation $\bot$.

Although certain groups of sentences can be ordered linearly between $\bot$ and $\top$ as a series of claims with steadily increasing content, the full range of propositions expressed by sentences are better thought of as inhabiting a much more complex logical space. This can be thought of, on the one hand, as a space of possibilities, with an individual proposition constituting a division of the space into two regions, the possibilities it rules out and the possibilities it leaves open. Another sort of space has as its points not possible worlds but propositions, with different possible worlds representing different dimensions with respect to which the location of propositions can differ. Logical space in this sense has a bottom in the proposition expressed by $\bot$ and a top provided by $\top$. So long as there are alternative possibilities (that is, more than just one possible world), there will be more propositions with intermediate content than there are degrees of content intermediate between $\bot$ and $\top$.

This picture of deductive reasoning fits into a simplified picture of the function of language. Our beliefs, the information we think we have, amount to a proposition that rules out a certain range of possibilities for the history of the universe. In general, we would like to narrow down the range of possibilities left open even further. When language is used cooperatively (something that must be the standard case), we share the ability to rule out possibilities by asserting sentences that rule out some of the possibilities our beliefs lead us to exclude. The sentences we can sincerely assert are the ones that express propositions that are entailed by the proposition expressing the sum total of our beliefs.

Glen Helman 01 Aug 2004

1.2.x. Exercise questions

1. Suppose you know that a certain argument is valid but do not know whether its premises and conclusion are true or false. If you are given one of the further items of information **a-c** about the premises of the argument, what if anything can you say about the truth value of its conclusion?

- **a.** The premises are all true.
- **b.** The premises are all false.
- **c.** Some premises are true and some are false.

2. Suppose that ϕ, ψ / χ is an argument that you know to be valid. If you find that the conclusion χ is false, what if anything can you say about the truth values of the premises ϕ and ψ?

3. For each of the following items of information, tell what you can conclude from it about the equivalence of sentences ϕ and ψ.

- **a.** ϕ and ψ are both true
- **b.** ϕ and ψ are both false
- **c.** ϕ is true and ψ is false
- **d.** There is a sentence χ such that χ and ϕ together entail ψ, and χ and ψ together entail ϕ (i.e., $\chi, \phi \Rightarrow \psi$ and $\chi, \psi \Rightarrow \phi$)

4. For each of the following pieces of information, tell what if anything you can conclude about the possibilities left open and the possibilities ruled out by the sentence ϕ:

- **a.** ϕ is equivalent to a tautology ψ
- **b.** ϕ entails $\top$
- **c.** a tautology ψ entails ϕ
- **d.** ϕ is equivalent to $\bot$
- **e.** ϕ entails an absurdity ψ
- **f.** $\bot$ entails ϕ

1.2.xa. Exercise answers

1.

- **a.** The conclusion must be true; an argument with true premises and a false conclusion would not be valid.
- **b.** Nothing can be said; some valid arguments whose premises are all false have true conclusions and others have false conclusions. (Here is an example of a valid argument with false premises and a true conclusion: *Indianapolis is the capital of Illinois*, *the capital of Illinois is east of the Wabash* / *Indianapolis is east of the Wabash*.)
- **c.** Nothing can be said; some valid arguments with both true and false sentences among their premises have true conclusions and others have false conclusions.

2. You can say that the premises are not both true, that at least one of the two is false.

3.

- **a.** Nothing; a pair of sentences with the same truth value may be equivalent but only if each has the same value as the other not only in the actual world but in all other possible worlds as well.
- **b.** Nothing; a pair of sentences with the same truth value may be equivalent but only if each has the same value as the other in ***all*** possible worlds.
- **c.** The sentences are not equivalent; equivalent sentences can never have different truth values.
- **d.** Nothing; such sentences need not even have the same truth value. In any possible world in which χ is true, we know that ψ is true if ϕ is and that ϕ is true if ψ is. But ϕ and ψ might have different truth values in worlds where χ is false.

4.

- **a.** ϕ expresses the same proposition as ψ so, like ψ, it leaves open all possible worlds and rules out none.
- **b.** Nothing can be concluded about ϕ; because $\top$ cannot be false, it has no content and every sentence entails it.
- **c.** Since ϕ is entailed by ψ, it must be true in every possible world in which ψ is true; therefore, ϕ must, like ψ, leave open all possibilities (that is, it is also a tautology)
- **d.** ϕ expresses the same proposition as $\bot$ so, like $\bot$, it rules out all possible worlds and leaves open none.
- **e.** Since ϕ is entails ψ, it must rule out every possibility that ψ rules out; but ψ rules out all possibilities so ϕ must as well (that is, it is also absurd).
- **f.** Nothing can be concluded about ϕ; because $\bot$ rules out all possibilities, no sentence can have any further content and thus no sentence fails to be entailed by $\bot$.