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.