

From: Bertrand Russell, *The Problems of Philosophy* (1912)

...

13 Knowledge, Error, and Probable Opinion

13.1 The question as to what we mean by truth and falsehood, which we considered in the preceding chapter, is of much less interest than the question as to how we can know what is true and what is false. This question will occupy us in the present chapter. There can be no doubt that *some* of our beliefs are erroneous; thus we are led to inquire what certainty we can ever have that such and such a belief is not erroneous. In other words, can we ever *know* anything at all, or do we merely sometimes by good luck believe what is true? Before we can attack this question, we must, however, first decide what we mean by 'knowing', and this question is not so easy as might be supposed.

13.2 At first sight we might imagine that knowledge could be defined as 'true belief'. When what we believe is true, it might be supposed that we had achieved a knowledge of what we believe. But this would not accord with the way in which the word is commonly used. To take a very trivial instance: If a man believes that the late Prime Minister's last name began with a B, he believes what is true, since the late Prime Minister was Sir Henry Campbell Bannerman. But if he believes that Mr. Balfour was the late Prime Minister, he will still believe that the late Prime Minister's last name began with a B, yet this belief, though true, would not be thought to constitute knowledge. If a newspaper, by an intelligent anticipation, announces the result of a battle before any telegram giving the result has been received, it may by good fortune announce what afterwards turns out to be the right result, and it may produce belief in some of its less experienced readers. But in spite of the truth of their belief, they cannot be said to have knowledge. Thus it is clear that a true belief is not knowledge when it is deduced from a false belief.

13.3 In like manner, a true belief cannot be called knowledge when it is deduced by a fallacious process of reasoning, even if the premisses from which it is deduced are true. If I know that all Greeks are men and that Socrates was a man, and I infer that Socrates was a Greek, I cannot be said to *know* that Socrates was a Greek, because, although my premisses and my conclusion are true, the conclusion does not follow from the premisses.

13.4 But are we to say that nothing is knowledge except what is validly deduced from true premisses? Obviously we cannot say this. Such a definition is at once too wide and too narrow. In the first place, it is

too wide, because it is not enough that our premisses should be *true*, they must also be *known*. The man who believes that Mr. Balfour was the late Prime Minister may proceed to draw valid deductions from the true premiss that the late Prime Minister's name began with a B, but he cannot be said to *know* the conclusions reached by these deductions. Thus we shall have to amend our definition by saying that knowledge is what is validly deduced from *known* premisses. This, however, is a circular definition: it assumes that we already know what is meant by 'known premisses'. It can, therefore, at best define one sort of knowledge, the sort we call derivative, as opposed to intuitive knowledge. We may say: '*Derivative* knowledge is what is validly deduced from premisses known intuitively'. In this statement there is no formal defect, but it leaves the definition of *intuitive* knowledge still to seek.

13.5 Leaving on one side, for the moment, the question of intuitive knowledge, let us consider the above suggested definition of derivative knowledge. The chief objection to it is that it unduly limits knowledge. It constantly happens that people entertain a true belief, which has grown up in them because of some piece of intuitive knowledge from which it is capable of being validly inferred, but from which it has not, as a matter of fact, been inferred by any logical process.

13.6 Take, for example, the beliefs produced by reading. If the newspapers announce the death of the King, we are fairly well justified in believing that the King is dead, since this is the sort of announcement which would not be made if it were false. And we are quite amply justified in believing that the newspaper asserts that the King is dead. But here the intuitive knowledge upon which our belief is based is knowledge of the existence of sense-data derived from looking at the print which gives the news. This knowledge scarcely rises into consciousness, except in a person who cannot read easily. A child may be aware of the shapes of the letters, and pass gradually and painfully to a realization of their meaning. But anybody accustomed to reading passes at once to what the letters mean, and is not aware, except on reflection, that he has derived this knowledge from the sense-data called seeing the printed letters. Thus although a valid inference from the letters to their meaning is possible, and *could* be performed by the reader, it is not in fact performed, since he does not in fact perform any operation which can be called logical inference. Yet it would be absurd to say that the reader does not *know* that the newspaper announces the King's death.

13.7 We must, therefore, admit as derivative knowledge whatever is the result of intuitive knowledge even if by mere association, provided there is a valid logical connexion, and the person in question could become aware of this connexion by reflection. There are in fact many ways, besides

logical inference, by which we pass from one belief to another: the passage from the print to its meaning illustrates these ways. These ways may be called 'psychological inference'. We shall, then, admit such psychological inference as a means of obtaining derivative knowledge, provided there is a discoverable logical inference which runs parallel to the psychological inference. This renders our definition of derivative knowledge less precise than we could wish, since the word 'discoverable' is vague: it does not tell us how much reflection may be needed in order to make the discovery. But in fact 'knowledge' is not a precise conception: it merges into 'probable opinion', as we shall see more fully in the course of the present chapter. A very precise definition, therefore, should not be sought, since any such definition must be more or less misleading.

13.8 The chief difficulty in regard to knowledge, however, does not arise over derivative knowledge, but over intuitive knowledge. So long as we are dealing with derivative knowledge, we have the test of intuitive knowledge to fall back upon. But in regard to intuitive beliefs, it is by no means easy to discover any criterion by which to distinguish some as true and others as erroneous. In this question it is scarcely possible to reach any very precise result: all our knowledge of truths is infected with some degree of doubt, and a theory which ignored this fact would be plainly wrong. Something may be done, however, to mitigate the difficulties of the question.

13.9 Our theory of truth, to begin with, supplies the possibility of distinguishing certain truths as *self-evident* in a sense which ensures infallibility. When a belief is true, we said, there is a corresponding fact, in which the several objects of the belief form a single complex. The belief is said to constitute *knowledge* of this fact, provided it fulfils those further somewhat vague conditions which we have been considering in the present chapter. But in regard to any fact, besides the knowledge constituted by belief, we may also have the kind of knowledge constituted by *perception* (taking this word in its widest possible sense). For example, if you know the hour of the sunset, you can at that hour know the fact that the sun is setting: this is knowledge of the fact by way of knowledge of *truths*; but you can also, if the weather is fine, look to the west and actually see the setting sun: you then know the same fact by the way of knowledge of *things*.

13.10 Thus in regard to any complex fact, there are, theoretically, two ways in which it may be known: (1) by means of a judgement, in which its several parts are judged to be related as they are in fact related; (2) by means of *acquaintance* with the complex fact itself, which may (in a large sense) be called perception, though it is by no means confined to objects of the senses. Now it will be observed that the second way of knowing a

complex fact, the way of acquaintance, is only possible when there really is such a fact, while the first way, like all judgement, is liable to error. The second way gives us the complex whole, and is therefore only possible when its parts do actually have that relation which makes them combine to form such a complex. The first way, on the contrary, gives us the parts and the relation severally, and demands only the reality of the parts and the relation: the relation may not relate those parts in that way, and yet the judgement may occur.

It will be remembered that at the end of Chapter XI we suggested that there might be two kinds of self-evidence, one giving an absolute guarantee of truth, the other only a partial guarantee. These two kinds can now be distinguished. 13.11

We may say that a truth is self-evident, in the first and most absolute sense, when we have acquaintance with the fact which corresponds to the truth. When Othello believes that Desdemona loves Cassio, the corresponding fact, if his belief were true, would be 'Desdemona's love for Cassio'. This would be a fact with which no one could have acquaintance except Desdemona; hence in the sense of self-evidence that we are considering, the truth that Desdemona loves Cassio (if it were a truth) could only be self-evident to Desdemona. All mental facts, and all facts concerning sense-data, have this same privacy: there is only one person to whom they can be self-evident in our present sense, since there is only one person who can be acquainted with the mental things or the sense-data concerned. Thus no fact about any particular existing thing can be self-evident to more than one person. On the other hand, facts about universals do not have this privacy. Many minds may be acquainted with the same universals; hence a relation between universals may be known by acquaintance to many different people. In all cases where we know by acquaintance a complex fact consisting of certain terms in a certain relation, we say that the truth that these terms are so related has the first or absolute kind of self-evidence, and in these cases the judgement that the terms are so related *must* be true. Thus this sort of self-evidence is an absolute guarantee of truth. 13.12

But although this sort of self-evidence is an absolute guarantee of truth, it does not enable us to be *absolutely* certain, in the case of any given judgement, that the judgement in question is true. Suppose we first perceive the sun shining, which is a complex fact, and thence proceed to make the judgement 'the sun is shining'. In passing from the perception to the judgement, it is necessary to analyse the given complex fact: we have to separate out 'the sun' and 'shining' as constituents of the fact. In this process it is possible to commit an error; hence even where a *fact* has the first or absolute kind of self-evidence, a judgement believed to 13.13

correspond to the fact is not absolutely infallible, because it may not really correspond to the fact. But if it does correspond (in the sense explained in the preceding chapter), then it *must* be true.

13.14 The second sort of self-evidence will be that which belongs to judgements in the first instance, and is not derived from direct perception of a fact as a single complex whole. This second kind of self-evidence will have degrees, from the very highest degree down to a bare inclination in favour of the belief. Take, for example, the case of a horse trotting away from us along a hard road. At first our certainty that we hear the hoofs is complete; gradually, if we listen intently, there comes a moment when we think perhaps it was imagination or the blind upstairs or our own heartbeats; at last we become doubtful whether there was any noise at all; then we *think* we no longer hear anything, and at last we *know* we no longer hear anything. In this process, there is a continual gradation of self-evidence, from the highest degree to the least, not in the sense-data themselves, but in the judgements based on them.

13.15 Or again: Suppose we are comparing two shades of colour, one blue and one green. We can be quite sure they are different shades of colour; but if the green colour is gradually altered to be more and more like the blue, becoming first a blue-green, then a greeny-blue, then blue, there will come a moment when we are doubtful whether we can see any difference, and then a moment when we know that we cannot see any difference. The same thing happens in tuning a musical instrument, or in any other case where there is a continuous gradation. Thus self-evidence of this sort is a matter of degree; and it seems plain that the higher degrees are more to be trusted than the lower degrees.

13.16 In derivative knowledge our ultimate premisses must have some degree of self-evidence, and so must their connexion with the conclusions deduced from them. Take for example a piece of reasoning in geometry. It is not enough that the axioms from which we start should be self-evident: it is necessary also that, at each step in the reasoning, the connexion of premiss and conclusion should be self-evident. In difficult reasoning, this connexion has often only a very small degree of self-evidence; hence errors of reasoning are not improbable where the difficulty is great.

13.17 From what has been said it is evident that, both as regards intuitive knowledge and as regards derivative knowledge, if we assume that intuitive knowledge is trustworthy in proportion to the degree of its self-evidence, there will be a gradation in trustworthiness, from the existence of noteworthy sense-data and the simpler truths of logic and arithmetic, which may be taken as quite certain, down to judgements which seem only just more probable than their opposites. What we firmly believe, if it is true, is called *knowledge*, provided it is either intuitive or inferred

(logically or psychologically) from intuitive knowledge from which it follows logically. What we firmly believe, if it is not true, is called *error*. What we firmly believe, if it is neither knowledge nor error, and also what we believe hesitatingly, because it is, or is derived from, something which has not the highest degree of self-evidence, may be called *probable opinion*. Thus the greater part of what would commonly pass as knowledge is more or less probable opinion.

In regard to probable opinion, we can derive great assistance from *coherence*, which we rejected as the *definition* of truth, but may often use as a *criterion*. A body of individually probable opinions, if they are mutually coherent, become more probable than any one of them would be individually. It is in this way that many scientific hypotheses acquire their probability. They fit into a coherent system of probable opinions, and thus become more probable than they would be in isolation. The same thing applies to general philosophical hypotheses. Often in a single case such hypotheses may seem highly doubtful, while yet, when we consider the order and coherence which they introduce into a mass of probable opinion, they become pretty nearly certain. This applies, in particular, to such matters as the distinction between dreams and waking life. If our dreams, night after night, were as coherent one with another as our days, we should hardly know whether to believe the dreams or the waking life. As it is, the test of coherence condemns the dreams and confirms the waking life. But this test, though it increases probability where it is successful, never gives absolute certainty, unless there is certainty already at some point in the coherent system. Thus the mere organization of probable opinion will never, by itself, transform it into indubitable knowledge.

13.18

...