Problem-Solving and the Problem of Induction*

by Donald Gillies, University College London

Contents

1. Introduction

Popper devotes the first chapter of his 1972 book: *Objective Knowledge* to an extended treatment of the problem of induction. He begins this chapter, and indeed the book as a whole, with his famous claim (1972, p. 1):

"I think that I have solved a major philosophical problem: the problem of induction."

Later in the book Popper proposes a general schema of problem-solving.[1] In fact there are discussions of this in Ch.3, p. 119; Ch. 4, pp. 164-5; Ch.6, p. 243; and Ch. 8, p. 297. The formulations of the schema differ slightly in these different discussions. I will use a formulation from Ch. 6, which seems to me the best. The idea of this paper is to apply this general schema of problem-solving to Popper's treatment of the problem of induction, and to see what results.

2. Popper's general schema of problem-solving

Popper writes (1972, p. 243):

"Using 'P' for problem, 'TS' for tentative solutions, 'EE' for error-elimination, we can describe the fundamental evolutionary sequence of events as follows:

$$P \rightarrow TS \rightarrow EE \rightarrow P.$$

But this sequence is not a cycle: the second problem is, in general, different from the first: it is the result of the new situation which has arisen, in part, because of the tentative solutions which have been tried out, and the error-elimination which controls them. In order to indicate this, the above schema should be rewritten:

$$P_1 \rightarrow TS \rightarrow EE \rightarrow P_2.\text{”}$$

Earlier Popper describes EE (error-elimination) as (1972, p. 164) “a severe critical examination of our conjecture.” He also describes the change from P_1 to P_2 as a ‘problem-shift’, observing (1972, p. 165):

“This leads to the … relation called ‘*problem shift*’ by I. Lakatos, who distinguishes between progressive and degenerating problem shifts.”

Later I will consider whether, in the case of the problem of induction, my suggested change from P_1 to P_2 is a progressive or a degenerating problem-shift.

3. The initial problem (P_1) and Popper’s tentative solution (TS)

I will now begin applying Popper’s general schema of problem-solving to Popper’s treatment of the problem of induction. Obviously we have to start by identifying P_1. This is clearly what Popper calls (1972, p. 2): “*the traditional philosophical problem of induction*”. Popper actually gives two formulations of this traditional problem. For simplicity, I will consider only the second, and this gives us:

P_1 = What is the justification for inductive inferences?

So what then is Popper’s solution of P_1? The gist of it is contained in the following passage (1972, p. 2):

“The second formulation assumes that there are inductive inferences, and *rules* for drawing inductive inferences, and this, again, is an assumption which should not be made uncritically, and one which I also regard as mistaken.”

So Popper thinks that there no inductive inferences. He goes on to describe the idea that there are such things as (1972, pp. 6-7): “a kind of optical illusion”. Now if there really are no such things as inductive inferences, then we do not have to justify them, and this solves P_1. Popper holds that science progresses through conjectures and refutations, and this is a process which does not involve any inductive inferences – only deductive ones. Perhaps Popper’s most emphatic denial of induction comes in the following passage which I will refer to as ‘the 1963 induction is a myth quotation’.[2] It runs as follows (1963, p. 53):

“Induction, i.e. inference based on many observations, is a myth. It is neither a psychological fact, nor a fact of ordinary life, nor one of scientific procedure.”

It should be noted that Popper here speaks of induction in connection with psychology, ordinary life, and scientific procedure. In the present paper, however, I will confine myself to considering induction in the context of scientific procedure.

Returning to our main theme, we can sum up Popper's tentative solution (TS) to P_1 as follows:

TS = There are no inductive inferences and so there is no need to justify them. Science progresses by conjectures and refutations, and this procedure does not use inductive inferences.

4. The EE phase: (i) computer induction

I now pass to the EE phase of the general schema of problem-solving. This consists of a severe critical examination of Popper's tentative solution (TS) of P_1. Here I will no longer continue quoting from Popper, but rather give the criticisms of other philosophers including myself. The first criticism I wish to present of Popper's TS came from discovering that researchers in artificial intelligence (AI) have developed a form of computer induction. I have given a detailed account of this criticism in my book (1996) *Artificial Intelligence and Scientific Method*, Chs. 1-3, pp. 1-71. Here I will summarise briefly the results.

There is a branch of AI known as 'machine learning' whose aim is to generate hypotheses automatically from data, in other words to carry out mechanical induction. In my 1996 book, some examples are given of successful machine learning programs (particularly Quinlan's ID3 and Muggleton's GOLEM), and it is argued that these show the existence of inductive rules of inference (or IRIs). It is worth noting the form which these IRIs take. Let e be the data and h the hypothesis inferred from the data. We might think that an IRI would take the form: 'From e, infer h'. However it turns out that it always has the form: 'From e & K infer h', where K is background knowledge (cf. my 1996, p. 18).

Developments in machine learning since 1996 have only reinforced the claim that inductive rules of inference exist. Hence it can be argued (cf. my 1996, p. 56) that Popper's 1963 induction is a myth quotation can no longer be regarded as correct. In fact programs such as Quinlan's ID3 or Muggleton's GOLEM (and more recently developed machine learning programs) do make inductive inferences based on many observations and have become a part of scientific procedure.

This criticism of Popper's TS to P_1 must be tempered by the following observations (cf. my 1996, p. 66). Popper's 1963 induction is a myth quotation, as applied to scientific procedure, contained a good deal of truth *at the time when it was published*. The first machine learning programs to be used successfully in science were Buchanan and Feigenbaum's Meta-DENDRAL, and Michalski's INDUCE. These appeared in the late 1970s and early 1980s, i.e. more than fifteen years after Popper's induction is a myth quotation. Moreover, I can find hardly any genuine and significant uses of Baconian or mechanical induction in science before the machine learning programs just mentioned. Thus my conclusion is that Popper's 1963 induction is a myth quotation has become incorrect because science itself is changing. This change is of course brought about by the introduction of computers and is some ways analogous to the changes brought about at any earlier phase of science's development by the introduction of instruments of observation such as telescopes and microscopes (cf. my 1996, p. 69). In effect, the current development of computers and AI is likely to change science in such a way that Baconian or mechanical induction becomes a standard part of scientific procedure.

The view that computer induction exists is now generally held, but it has been challenged in an interesting paper by Tamburini (2006).[3] Tamburini only considers one of the two main examples of computer induction which I give in my 1996. This is ID3. He remarks quite correctly (2006, p. 273): "Popper's anti-inductivism was questioned on the basis of ID3 performances (Gillies 1996); …" ID3 is a system which infers decision trees from data. Like any machine learning system, it assumes some background knowledge (K) which, in this case, is that the domain of objects under consideration is appropriately described by a specified set of attributes. Given K, decision trees are automatically generated from data e using built-in algorithms. As the correctness of the decision tree (D say) by no means follows deductively from e & K, it seems almost inescapable that we have here an inductive rule of inference (IRI) of the form: From e & K, infer D. As I show in my 1996 (pp. 36-38), this IRI can be considered as generated by the iteration of more basic IRIs. It should also be observed that the decision trees generated, which were not previously known, turn out in many cases to contain substantial knowledge and to be very successful in practice. An example, given in my 1996, p. 46, comes from the work of Bratko. He used a developed form of ID3, known as ASSISTANT, to induce decision trees for medical diagnosis. These decision trees performed better than human specialists in all cases in which an objective statistical comparison was possible.

ID3 seems to be such a clear example of mechanical induction that it is difficult to see how Tamburini can hold the following opinion (2006, p. 268): " … I maintain here that AI investigations on learning systems do not compel one to relinquish Popper's radical scepticism towards induction." The core of Tamburini's defence of this position, in the case of ID3, seems to lie in the following passage (2006, p. 276):

"If the presuppositions of the first kind (ID3 biases) can be suitably stated in declarative form, then a concept learning algorithm such as ID3 can be redescribed as a theorem prover."

The line of thought here seems to be the following. The algorithms in ID3 involve some presuppositions (or, as one might say, heuristics). Suppose we were able to state these explicitly as say P. ID3 involves an inductive rule of inference of the form: From e & K, infer D. However, if we added P to e & K, we might be able to turn this inductive rule into a deductive rule of the form: From e & K & P, infer D, where now D follows from e & K & P by deductive logic. In this way an inductive machine learning system would become like an automated theorem prover which involves only deductive rules of inference.

The first point to note here is that Tamburini's claim is only hypothetical (2006, p. 276): "If the presuppositions of the first kind (ID3 biases) can be suitably stated in declarative form …" Of course it would be very difficult indeed, if not downright impossible, to suitably state these presuppositions in declarative form, and Tamburini doesn't attempt to do so. Thus his suggested reduction of ID3 induction to deduction is purely hypothetical and most unlikely ever to be accomplished.

Let us, however, suppose, as our second point, that this reduction could really be carried out. There is no doubt that ID3, redescribed as a theorem prover, would be much more complicated that the original ID3 presented as an inductive learning system. Why should we introduce all this unnecessary complication which would never be adopted in practice? The question before us is whether to allow the introduction of inductive rules of inference (IRIs), or to allow only deductive rules of inference (DRIs). If we adopt the former position, we get simple computer induction systems which are successful in practice. If we adopt the second position, we are forced to try to transform these systems into equivalent theorem provers which involve only DRIs. This is a difficult, probably hopeless, task which adds complexity with no practical gain. The case therefore for allowing the introduction of IRIs is overwhelming.

I will come back briefly to computer induction later on, but, as I have already presented this particular criticism of Popper's tentative solution (TS) to P_1 in some detail in an earlier publication, I will devote most of the rest of this paper to another criticism.

5. The EE phase: (ii) is corroboration in some sense inductive?

In the course of his discussion of the problem of induction, Popper introduces the notion of degree of corroboration which he characterises as follows (1972, p. 18):

> "By the degree of corroboration of a theory I mean a concise report evaluating the state (at a certain time *t*) of the critical discussion of a theory, with respect to the way it solves its problems; its degree of testability; the severity of tests it has undergone; and the way it has stood up to these tests."

The degree of corroboration of a hypothesis h given evidence e is written C(h, e), or perhaps better C(h, e&K) where K stands for the background knowledge. Popper's term 'corroboration' was introduced to contrast his theory of corroboration with Carnap's theory of confirmation (cf. Carnap, 1950). Indeed the two theories differ in important ways. For example, Carnap is a Bayesian which means that he thinks that C(h,e) satisfies the axioms of probability, whereas Popper is a non-Bayesian and denies that C(h,e) satisfies the axioms of probability. In symbols, Popper's claim here is that $C(h,e) \neq P(h|e)$. Despite these differences I prefer to use the terms 'corroboration' and 'confirmation' as synonyms, and to abbreviate them both by C. The difference between

Popper’s theory and Carnap’s is expressed by saying that they characterise the C-function differently.

Returning now to Popper’s 1972 discussion of induction, one might ask whether he really needed to introduce corroboration at all. Suppose there are, in a particular area of investigation, n theories between which scientists have to decide. Could they not devise tests which refute n-1 of these theories, leaving only one unrefuted which will then become the preferred theory? No notion of corroboration is needed to carry out such a procedure. However, Popper does point out a possible difficulty here (1972, p. 15):

“On the other hand, *among the theories actually proposed* there may be more than one which is not refuted at a time *t*, so that we may not know which of these we ought to prefer.”

Suppose, however, we have a measure of corroboration, we can then prefer the best corroborated theory among those which are unrefuted.

Corroboration is also involved in what Popper calls the *pragmatic problem of induction*. Popper gives two formulations of this problem. I will focus on the second of these which he states as follows (1972, p. 21):

“Pr_2 Which theory should we prefer for practical action, from a rational point of view?”

He goes on to say (1972, p. 22):

“My answer to Pr_2 is: … we should *prefer* as basis for action the best -tested theory.”

Given Popper’s characterisation of corroboration quoted earlier, we can roughly identify “the best-tested theory” with “the best-corroborated theory”. This leads to the following pragmatic principle, which, it should be stressed, is a modification of what Popper writes:

(1) Use, as the basis for action, the best corroborated theory

This principle (1) does not quite correspond to what happens in practice as we can see by considering the following example. Suppose a pharmaceutical firm has developed a new drug X to treat some illness. Before X is put on the market, it is important to make sure that it does not have any harmful side effects. Let us therefore formulate the following hypothesis:

H_X: X, when taken in the appropriate dosage, does not have any harmful side effects

Now before X can be put on the market H_X must, by law, be subjected to a series of severe tests – first with animals, and then in the form of clinical tests on humans. Only if

H_X passes all these tests can it be marketed.[4] To put it another way, X can only be put on the market if H_X has a sufficiently high degree of corroboration.

This leads to the following pragmatic principle (2):

(2) Use, as the basis for action, theories which have a sufficiently high degree of corroboration

What is meant by ‘sufficiently high degree of corroboration’ is specified in the case of drugs by the government regulations on what tests a new drug must pass before it can be put on the market. In general it would be understood contextually as part of the practice of the area in question.

Now we come to the problem. It seems that pragmatic principles such as (2) are indeed accepted as guides to action. But in accepting such a principle, are we not implicitly giving an inductive significance to corroboration? Suppose a theory has a high degree of corroboration. This means that it has explained correctly the results of past observations, and perhaps also given the correct predictions in a number of tests. Let us say in these circumstances that the theory has so far performed well. However, if we adopt the theory as the basis for actions, are we not assuming that it will continue to perform well in the future? In other words accepting a pragmatic principle such as (2) seems to be implicitly adopting an inductive assumption.

The criticism of Popper involving computer induction could not have been formulated before the late 1970s and early 1980s, because it was only then that successful systems for computer induction were created. However, the criticism involving corroboration and the pragmatic problem of induction is much older. In his 1994 (pp. 20-23), Miller gives a list of no less that 11 philosophers who have made criticisms along these lines, and goes on (1994, pp. 38-45) to try to answer all these objections. One formulation of such a criticism is to be found in Salmon (1968). Salmon considers whether we are acting rationally if we prefer a prediction based on a well-corroborated scientific theory to a prediction based on some theory which has low or even negative corroboration. He writes (1968, p. 97):

“Either corroboration has an inductive aspect or there is no logic of prediction. If there is no logic of prediction, it is hard to see how any choice would be ‘rational’.”

Salmon thinks that we can make rational choices here, so that the conclusion of his argument is that corroboration has an inductive aspect. Salmon further elaborates this criticism of Popper in his 1981.

O’Hear also gives an elegantly formulated criticism of this kind in his 1980, where he writes (pp. 40-41):

“ … it is unclear how Popper is in a position to tell us that it is more rational to act on a well-corroborated theory than to adopt any other policy when it comes to action. … High corroboration shows only that a theory has done well up to now. Hume's point is that our world might suddenly change to being one where chance might be a good method or where previously falsified theories might be the best to act on or where we might be better off having no method at all. I cannot see how Popper is justified in claiming that these methods are, in the light of his acceptance of Hume's point, worse methods for basing practical actions on."

But what does Popper himself say about the relation between corroboration and induction? There is in fact one passage in which he seems to come close to giving an inductive significance to his measure of corroboration (1959, New Appendix *ix, p. 418):

"It might well be asked at the end of all this whether I have not, inadvertently, changed my creed. For it may seem that there is nothing to prevent us from calling *C(h, e)* 'the inductive probability of *h*, given *e*' or – if this is felt to be misleading, in view of the fact that *C* does not obey the laws of the probability calculus – 'the degree of the rationality of our belief in *h*, given *e*'. A benevolent inductivist critic might even congratulate me on having solved, with my *C* function, the age-old problem of induction *in a positive sense* – on having finally established, with my *C* function, the validity of inductive reasoning."

However, it should be noted that the view given here is that of 'a benevolent inductivist critic'. It is not Popper's own as the following passage from *Objective Knowledge* clearly shows (1972, p. 18):

"Corroboration (or degree of corroboration) is thus an evaluating *report of past performance*. … Being a report of past performance only, … *it says nothing whatever about future performance* …"

But if degree of corroboration really said nothing whatever about future performance, why should we use it to guide our actions? In using it in this way, we are surely implicitly assuming that degree of corroboration does say something about future performance. In other words we are giving degree of corroboration an inductive significance.

That concludes the EE phase of the general schema of problem-solving, and I will next consider what new problem arises from all this.

6. The new problem (P_2): choosing a C-function

The new problem which arises is, so I claim, that of choosing a C-function. I will first explain what I mean by this, and then explain why it arises from the EE part of the preceding discussion.

Choosing a C-function sounds like giving the full specification of a mathematical function which for any values of h and of e&K gives a real number C(h, e&K). Perhaps

Carnap dreamed of constructing such a fully specified mathematical function, but it is not a very realistic aspiration as far as current practice is concerned. In some AI cases, a C-function is precisely specified in the mathematical sense, and this function is coded into the machine learning program. However, the language used to specify the function, the nature of the background knowledge K, and the precise details of the function would all depend on the specific application, and would be different in a different application (even if the C-functions used in different applications have some features in common). In ordinary human science, the specification of the function is also highly context-dependent, but here it is qualitative as well. This is clearly shown in the drug case described above, in which what is meant by ‘a sufficiently high degree of corroboration’ is specified by listing the tests which must be performed and passed to achieve this grade. This listing of the necessary tests is one way of choosing a C-function for a problem.

Then again there are debates concerning general features of the C-function. For example, as already mentioned, the Bayesians hold that C(h, e&K) should satisfy the standard axioms of probability while some non-Bayesians such as Popper deny this. Of course this suggests that the C-function might be Bayesian in some contexts and non-Bayesian in others (for a suggestion along these lines, see my 1998, section 4, pp. 155-6).

So, to sum up, what we are here referring to as ‘choosing a C-function’ is actually quite a complex and context-dependent process. It may involve, in a particular AI context, choosing a specific mathematical function. However, in more general contexts, it may be no more than a specification of certain general features of the C-function, and of the circumstances in which the C-function attains some key value.

Let us next analyse how the problem of choosing a C-function arises out of the EE discussion given earlier. I argued that Popper’s own treatment of the problem of induction involves introducing corroboration, and so gives rise to the problem of choosing a C-function. My first criticism of Popper’s approach was that his claim that induction is a myth is wrong because inductive rules of inference are used in successful AI machine learning programs. Now we can connect this criticism with the subsequent discussion of corroboration, because, once a C-function is chosen, we can use it to justify an inductive rule of inference (see my 1996, p. 105). An inductive rule of inference takes the form: ‘From e & K, infer h’, where K is the background knowledge, e is the evidence, and h is a hypothesis which explains the evidence. The justification of such a rule given a C-function is simple. The rule is justified, if C(h, e&K) is sufficiently high.

So the suggestion is to change our original problem (P_1): ‘What is the justification of inductive inferences?’ into the new problem (P_2) of choosing a C-function. But is this problem-shift progressive or degenerating? Naturally I would like to argue that it is progressive, and will now explain my reasons for thinking that it is.

I argued earlier that there are some exceptions to Popper’s claim that rules of inductive inference do not exist. However, these exceptions are relatively rare. They occur for example in the machine learning programs of AI. For the vast bulk of human science both in the past and present, rules of inductive inference do not exist. For such science, Popper's model of conjectures which are freely invented and then tested out seems to me more accurate than any model based on inductive inferences. Admittedly, there is talk nowadays in the context of science carried out by humans of 'inference to the best explanation' or 'abductive inference', but such so-called inferences are not at all inferences based on precisely formulated rules like the deductive rules of inference. Those who talk of 'inference to the best explanation'or 'abductive inference', for example, never formulate any precise rules according to which these so-called inferences takes place. In reality the 'inferences' which they describe in their examples involve conjectures thought up by human ingenuity and creativity, and by no means inferred in any mechanical fashion, or according to any precisely specified rules. Now the advantage of the new problem (P_2) of choosing a C-function is that it solves the original problem (P_1), as we have seen, but it also deals with the case of hypotheses generated not by any inductive inference but by a process of conjecture and testing. Such conjecture-generated hypotheses are justified if, as a result of testing, they become well corroborated. So, if we have agreed on the choice of a C-function, we can provide justification whether a hypothesis is generated by some inductive inference, or whether it is obtained by conjecture and testing. Thus the new problem is more general than the old.

Another advantage of the new problem over the old one is that the traditional problem of induction (P_1) suggested a series of approaches which proved to be very unsatisfactory (for some details about these, see my 1993, pp. 8-11 & 34). Formulating the problem of induction in the form: 'What is the justification for inductive inferences?' suggested to many thinkers, particularly those of the Cambridge school such as Keynes and Russell, that inductive inferences needed to be justified in terms of general principles such as *the uniformity of nature* or *the principle of induction*. However, two problems emerged with this approach. First of all it proved almost impossible to formulate these alleged principles in any clear fashion. For example, Russell's formulation of the principle of induction contains an error which vitiates it (see my 1993, p. 10). Secondly it seemed to be impossible to give any convincing justification of such principles. One great merit of the shift from P_1 to P_2 is that it enables us to dispense with such obscure and unsatisfactory principles.

These then are the merits of the shift from P_1 to P_2 , and they seem to me to justify the claim that this shift is a progressive one. However, it should be stressed that Popper's general schema for problem-solving is a never-ending process. As Popper himself says (1972, p. 164):

"P_2 is the problem situation as it emerges from our first critical attempt to solve our problems. It leads up to our second attempt (*and so on*)."

So we should now look more closely at the new problem P_2 , and see if we can propose a preliminary tentative solution. This I will do in the next (and final) section.

7. A tentative solution (TS) to the new problem: an extension of Neurath's principle

So far I have formulated the new problem P_2 rather loosely, as being that of choosing a C-function. We can, however, split this into two questions. The first question is: 'how do we set about choosing a particular C-function in a specific situation?' The best way of approaching this question is to examine the practice of good scientists, and see if we can formulate general principles which underlie this practice. This could be described as *codifying practice* in the choice of C-functions. It is normally studied in philosophy of science under the heading 'confirmation theory'. But now suppose we have chosen a particular C-function. Then a second question arises, namely:

What is the justification for particular choices of C-function?

Because this question is analogous to the traditional philosophical problem of induction (P_1) with which we started, I will from now on take this second question as our P_2. So the problem shift is from

P_1 = What is the justification for inductive inferences?

to

P_2 = What is the justification for particular choices of C-function?

Our tentative solution (TS) to P_2 is based on some ideas of Neurath's. These are expressed by Neurath in the following famous passage (1932/3, p. 201):

"We are like sailors who must rebuild their ship on the open sea, never able to dismantle it in dry-dock and to reconstruct it there out of the best materials."

Here Neurath gives his view as an analogy. However, in a previous work (1993, p. 138), I have tried to formulate his position in a more explicit fashion as what could be called *Neurath's principle*. This is a conjunction of two parts, (A) and (B) which may be stated as follows:

(A) In order to test any scientific statement, we have to assume for the time being some other scientific statements. (This corresponds in the analogy to the fact that we can remove a plank of the ship only if we leave some others in place, since otherwise the ship would sink.)

(B) There is, however, no scientific statement which cannot be subjected to testing, and perhaps abandoned as a result of tests. (This corresponds in the analogy to the fact that any plank of the ship can be removed and checked to see if it is rotten.)

The first application we can make of Neurath's principle is to the body of scientific theories. Here we cannot question, and demand justification for, all our scientific theories at the same time. To test out one scientific theory, we have, for the time being, to assume others – in particular the theories used to interpret the relevant observations and experimental results. Similarly, I now argue, we cannot question, and demand justification for, all our choices of C-function at the same time. What we can do, and what has actually been done in the course of scientific and technological development, is to test out, and perhaps reject or modify, *particular* choices of C-function, while assuming, for the time being, other such choices. There is a circle here, just as there is in the case of testing out scientific theories, but it is no more vicious in the one case than in the other. Although we cannot criticize our choices of C-function all together, there is no particular such choice which cannot be criticized, tested out, and evaluated.

My suggestion then is to extend Neurath's principle from scientific theories to C-functions. We test out our choices of C-functions by experience just as we test out our scientific theories by experience. The only rule in both cases is that we cannot question all our assumptions at the same time. In order to question one thing, other things must, for the time being at least, be assumed. I will conclude by illustrating this with an example.

Let us return to our consideration of testing new drugs to make sure that they have no harmful side effects before they are put on the market. Earlier we formulated this problem by considering the following hypothesis:

H_X: X, when taken in the appropriate dosage, does not have any harmful side effects

where X is a new drug developed by a pharmaceutical firm to treat some illness. In this case the C-function for H_X is chosen informally as follows. A series of tests $t_1, \ldots, t_n$ is specified. Some of these will be on animals and some will be clinical tests on humans. Suppose our evidence (e_n say) is that all these tests have been carried out, and H_X has passed them all. Then the value of $C(H_X, e_n)$ is judged to be sufficiently high to allow the drug to be put on the market.

Suppose this choice of C-function is well-established, and has been used successfully for a number of years. Then a new drug T is devised.[5] The standard tests are performed on T and it passes them all successfully. So T is put on the market. However, tragedy ensues. T is actually very successful at curing the illnesses for which it is prescribed, but a quite unexpected side-effect occurs. When it is prescribed to pregnant women, they give birth to babies with very severe defects. This disaster leads to a modification of the choice of C-function which has been used hitherto by the pharmaceutical industry. A new test (t_{n+1} say) is introduced which consists in giving the drug to experimental animals which are pregnant, and then checking whether the resultant offspring have any defects. The drug only passes this test successfully if no birth defects are discovered.[6] Let e_n be as before, and e_{n+1} be the evidence that the tests $t_1, \ldots, t_n, t_{n+1}$ have been carried out, and H_X has passed them all. Then the value of $C(H_X, e_{n+1})$ is judged to be sufficiently high to allow the drug to be put on the market. However, the value of $C(H_X, e_n)$ is no longer judged to be sufficiently high to allow the drug to be put on the market. The choice of C-function in this particular context has been changed.

Suppose further that the new choice of C-function works well, and there are no further disasters occasioned by its use. We can then conclude – of course implicitly assuming other choices of C-function which have not been changed – that the new choice of C-function in this particular context is an improvement and was justified. This shows how Neurath's principle allows us to justify changes in our choice of C-function. It shows indeed that choices of C-function can be steadily improved along with the rest of science.

Notes

* This paper incorporates modifications to earlier drafts suggested by David Corfield, David Miller, and David Teira. I am very grateful to them for their comments.

1. On the origins of this schema of Popper's, see ter Hark (2004), particularly pp. 128 & 175.

2. I heard Popper himself utter the fateful words: "Induction is a myth" when I attended his lectures as a graduate student in the academic year 1966-7. As far as I can remember, Popper continued: "… and those who use the term 'induction', do not know what they are talking about."

3. I would like to thank David Miller for drawing this paper to my attention.

4. I have simplified somewhat here since, in practice, drugs are allowed to have harmful side effects in some classes of patients, provided these classes can be specified clearly in advance so that it is known that patients in one of these classes should not be prescribed the drug. For example, drugs for some heart conditions may have no harmful side effects for the normal patient, but might have harmful side effects for patients suffering from diabetes. I have ignored these complications since they do not affect the points about corroboration being made here.

5. As the letter T indicates, this hypothetical example is based on the real case of thalidomide. However, my example is a considerable simplification of what actually happened in that case. Some details about the actual thalidomide case are to be found in Timmermans and Leiter (2000). One interesting thing to which they draw attention is that thalidomide has been partially rehabilitated as a drug. It is in fact very effective as a treatment for some very severe conditions such as a tissue inflammatory syndrome which occurs in leprosy and AIDS wasting syndrome. It is now prescribed for these conditions, while taking precautions to prevent it ever being used by pregnant women.

6. Timmermans and Leiter say (2000, p. 45): "After the thalidomide disaster, studies in which pregnant rabbits were given thalidomide produced phocomelia birth defects."

References

Carnap, Rudolf (1950) *Logical Foundations of Probability*, 2nd Edn., University of Chicago Press, 1963.

Gillies, Donald (1993) *Philosophy of Science in the Twentieth Century: Four Central Themes*, Blackwell.

Gillies, Donald (1996) *Artificial Intelligence and Scientific Method*, Oxford University Press.

Gillies, Donald (1998) Confirmation Theory. In D.M.Gabbay and P. Smets (eds.) *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Volume 1, Kluwer, pp. 135-167.

Miller, David (1994) *Critical Rationalism. A Restatement and Defence*, Open Court.

Neurath, Otto (1932/3) Protocol Sentences. Reprinted in English translation in A.J.Ayer (ed.), *Logical Positivism*, Free Press, 1959, pp. 199-208.

O'Hear, Anthony (1980) *Karl Popper*, Routledge, 2002.

Popper, Karl (1959) *New Appendices to the Logic of Scientific Discovery*, in 6th revised impression of the 1959 English translation, Hutchinson, 1972, pp. 307-464.

Popper, Karl (1963) *Conjectures and Refutations*, Routledge and Kegan Paul.

Popper, Karl (1972) *Objective Knowledge. An Evolutionary Approach*, Oxford at the Clarendon Press.

Salmon, Wesley C. (1968) The Justification of Inductive Rules of Inference (+ Discussion). In Imre Lakatos (ed.) *The Problem of Inductive Logic*, North-Holland, pp. 24-97.

Salmon, Wesley C. (1981) Rational Prediction, *British Journal for the Philosophy of Science*, **32**, pp. 115-125.

Tamburini, Guglielmo (2006) Artificial Intelligence and Popper's Solution to the Problem of Induction. In Ian Jarvie, Karl Milford, and David Miller (eds.) *Karl Popper: A Centenary Assessment*, Volume II Metaphysics and Epistemology, Ashgate, pp. 265-284.

Ter Hark, Michel (2004) *Popper, Otto Selz and the Rise of Evolutionary Epistemology*, Cambridge University Press, Paperback Edition 2007.

Timmermans, Stefan and Leiter, Valerie (2000) The Redemption of Thalidomide:: Standardizing the Risk of Birth Defects, *Social Studies of Science*, **30**, pp. 41-71.