PREDICTIVE POLICIES

II—Aaron Sloman

WHAT MAKES SOME POLICIES BETTER THAN OTHERS?

1. First I must apologise for ever agreeing to present a paper on induction. Only after it was too late to withdraw did I learn how incompetent I was to deal with the subject, in view of the vast and varied literature and my unfamiliarity with all but a very small proportion of it. In particular, a philosopher who has not mastered the formal and mathematical developments in the fields of subjective probabilities, statistical decision theory, games theory and inductive logic should tread very warily indeed, for he risks repeating mistakes which have already been corrected in the literature or making suggestions which are out of date. It is impossible now for me to heed this warning: perhaps other philosophers will.

2. Mr. McGowan's paper seems to have two main aims, first, to say what an inductive inference policy is and how it differs from alternative non-deductive policies, and secondly, to show that the inductive policy is better, or more rational, than the alternatives. I shall criticise his characterisation of induction, his arguments to show its superiority, and some of his undiscussed assumptions. Finally, I shall take the risks mentioned in the previous paragraph by discussing the nature of attempts to justify induction and suggesting some lines of further enquiry, based on an analysis of the logic of ' better '. I start with some comments on Mr. McGowan's preliminary discussion, before turning to his recursive characterisation of an inductive inference policy.

3. In paragraph 2 (c) Mr. McGowan allows that " systems of inductive logic may be constructed within which we may assign a (high) ' probability ' to an hypothesis on the basis of an evidence statement ". Later in the paragraph he suggests that it is necessary to establish that " when the sentences and the technical terms of the calculus are given an empirical meaning, the world, and in particular the future, works in accordance with its probability assignments. " But this requirement is scarcely intelligible if statements of the form ' given evidence e the probability of h is x ' are analytic in such a calculus (as in Carnap's systems), for then there is no question of the world or the future conforming or not conforming. If the statement were given a non-analytic interpretation in terms of proportions among the *possible* worlds, there

would still be no questions of the actual world, or the future, conforming or not, any more than there is a question of *one* man conforming or failing to conform to a statement about the average height of all men. Moreover, if the interpretation is supposed to be in terms of ' long-run ' frequencies among *actual* events, the difference between conformity and non-conformity appears to lack any empirical significance for a finite person since whatever he observes will be compatible with any frequency between 0 and 1. If the interpretation of probability statements is in terms of confidence levels or rational degree of belief, the only thing which can conform or fail to conform is a person making predictions, not the future which he predicts. I conclude that Mr. McGowan is wrong to hint that ' reinforcement ' of a probability calculus with some presupposition of uniformity or limited independent variety could establish conformity between the world and the calculus, since the question of conformity does not arise.

4. I shall be making several comments on Mr. McGowan's formulations of the problem of induction. The first concerns his restriction of the discussion to what he calls *projective* arguments in his first paragraph, namely, arguments in which " the conclusion is either a generalisation of which the premisses constitute but a proper subset of its instances, or a further instance not included in the premisses ". My first complaint is that many conclusions which " go beyond " the available evidence bear very much more complex and indirect relations to the evidence, for example the conclusion that the charge on an electron is approximately 4.77×10^{-10} e.s.u., or the conclusion that neutrinos exist. The danger in ignoring arguments which are more complex than the ones he calls projective is that it is quite possible that the simplicity of projective arguments is deceptive. For it may be that when their conclusions are accepted it is not because they are conclusions of projective arguments, but because they are derivable from available evidence by a sequence of steps essentially involving the more complex and indirect relations mentioned above. Perhaps this point will become clearer later on. (A further minor complaint about the definition of ' projective ' is that it is not clear that it includes arguments whose conclusions are statements of probabilities, such as apparently occur in loop III of his diagram. This is partly because he does not tell us which of the many possible senses of ' probability ' he is using.)

5. In his section 3 Mr. McGowan draws up a table of ' forms '

of inference. My first complaint is that his assumption that the second category, pseudo-deductive inferences, is distinct from the third, pseudo-inductive inferences, and the fourth, inductive inferences, is very doubtful. Secondly, he goes on to say that " there is fairly common agreement as to which actual forms of argument fall into classes (1) to (4) ", which just seems to be wrong, in view of the frequent disagreements about what conclusions should be drawn from available evidence. Further, even if there are commonly used intuitive principles for making such classifications, it would seem that the well known paradoxes of confirmation and Nelson Goodman's paradoxes (see his *Fact Fiction and Forecast*) show these principles to be inconsistent or at least incomplete.

Suppose laymen would all ascribe some sort of validity to the example given by Mr. McGowan:

> " All lemons which have ever been tasted have been sour.
> Therefore, all lemons taste sour."

Couldn't the reasonableness easily be disputed, for example, if it is known that some lemons (as yet untasted) grow in circumstances in which some chemical essential for sourness is not available, or if it is known that there is a species of lemon (also as yet untasted) which differs from all tasted lemons in several very marked respects? Even without such definite counter-evidence, it could be argued, along lines hinted at in my previous paragraph, that accepting the conclusion would be rash unless one had some reason to believe that there is something about the constitution of lemons (or all sour-tasting fruit), rather than, say, the circumstances in which they had previously been tasted, which *explains* why they taste sour. In section 4 Mr. McGowan claims that we adopt the inductive method with perfect ease and unconcern, and that we (who?) are convinced that induction is reasonable. I have been suggesting that neither is true of reflective people.

6. However, even it is true that there is much that reflective and unreflective people all agree about, the agreement may be in *beliefs* concerning specific common sense generalisations, well known scientific laws, and particular predictions based on them. It is by no means clear that such agreement implies that there is agreement concerning *methods* of arriving at such beliefs, or even that there *are* any methods. If some generally accepted statement were written down as the conclusion of an explicit inductive argument whose premisses described a lot of the available evidence, it is not at all obvious that most people would be found to agree

that the premisses in some sense make the conclusions more probable than its negation, even if they all agreed that the conclusion was true. Moreover, even the assertion that most people *believe the conclusions* of such arguments may misleadingly suggest conscious consideration of the question and rejection of the alternative, in cases where belief concerning the conclusion amounts to nothing more than a disposition to act *as if* it were true. Perhaps there is a sense in which the conclusion ' The floor will not give way when I take the next step ' is inductively related to a particular person's experience, and perhaps he acts on the proposition: but how many people adopt such beliefs as a result of conscious deliberation? Now it may well be that a solution to the problem of induction takes the form of a demonstration that in some circumstances it is rational to *act on* a certain hypothesis irrespective of whether the available evidence gives good reason for *believing* it. (More on this below.) But even if there is general agreement on which propositions people should act on, we cannot assume without investigation that this is because of general agreement on some method of inference.

7. So far I have commented mainly on Mr. McGowan's preliminary discussions, and it is time now to turn to his attempts to give a more precise characterisation of inductive reasoning. He starts by giving some necessary conditions for an inference's being inductive. The second of these is:

> " (b) The conclusion/prediction must be determined by the results of previous observations. . . . "

This seems to be a more condensed version of two requirements which, in section 5, he says are necessary conditions for rational action whether based on inductive policies or not, namely:

> (i) Requirement of non-arbitrariness: if the circumstances had been different the act would have been different.
>
> (ii) Requirement of consistency: if the circumstances were the same again the act would be the same.

Clearly (i) is too strong: not all differences need lead to different acts. We should at most say that if the circumstances had been *sufficiently* different the act would have been different. What is wanted is that an inference rule should be sensitive to the results of previous observations in that the conclusions should not be the same irrespective of the available evidence: but a one-one correlations between sets of evidence statements and conclusions should not be required. But can a *general* answer be given to the question: how much difference is sufficient? It seems unlikely. Further, it seems that (ii) is also too strong, for it rules out the possibility that in some cases it may be best to adopt a rule according to which if the evidence is that a certain proportion of A's, but not all, are B's the prediction should be determined by a probability device geared to the observed proportion: *e.g.,* if two-thirds of all A's have been B's, then spin a numbered wheel and predict that the next A will be not-B if a number divisible by 3 turns up, otherwise predict that it will be a B. Games theory shows how actions based on partly random policies may maximise expected utilities: why rule out in advance the possibility of such a rational predictive policy? It seems that in (b) ' partly determined ' should replace ' determined '.

8. Perhaps my most important criticism is that Mr. McGowan's fourth necessary conditions for inductive inference, (d) the Principle of Persistence, is also formulated too strongly. As stated, it postulates " that whatever (sic) characteristics, correlations, sequences and concatenations of properties the world has displayed in its observed segments will persist to the greatest degree in its hitherto unobserved segments ", and is thus equivalent to a Principle of Uniformity of Nature. This is much too strong, for it implies that nothing ever changes. Since many regularities observed in the past have subsequently ceased to exist, the available evidence should lead us to reject the principle on its own terms. Moreover, this principle can easily lead us to make contradictory assertions on the basis of observational evidence. Thus, suppose all A's which have been fully examined are also B's and C's, and that an A which has not been fully examined is known not to be a C. Then we can say that all fully examined A's have been D's where being a D is defined as being (C & B) or (not-C & not-B). In these circumstances the principle of persistence requires us to predict that the A which has not been fully examined is a B and a D, which is incompatible with its being known to be a C. More generally, we should have to assert all three of the following inconsistent propositions on the basis of available evidence:

(i) Some A's are not C's.

(ii) All A's are B's.

(iii) All A's are D's.

(Goodman's paradoxes generate similar problems for the Principle of Persistence, but at the cost of using much stranger predicates.)

9. Clearly we need a Weak Principle of Persistence, or a Weak Principle of Uniformity, in which ' whatever ' is replaced by ' some ': ' Some characteristics, *etc.*, will persist in the future '. But as soon as this replaces the strong principle it can no longer provide the traditional " justification " for particular inductive inferences, and it appears that the concept of an inductive policy needs to be supplemented by some indication of the criteria by which to select from among the observed regularities those which are to be projected. So the philosophical problem of induction acquires two facets: first the problem of justifying the Weak Principle of Persistence, and secondly the problem of justifying the chosen method of selecting which regularities to project. We shall see in the next paragraph that the Weak Principle is almost too weak to be worth justifying. But can there be an *a priori* justification of a method for deciding *which* observed characteristics of the universe will persist? (If such a method could be formulated it would presumably turn much of scientific research into a mechanical process.) Of course we do have methods of *elimination* on the basis of new evidence: reject generalisations which turn out to have counter-instances. And it is possible to argue, following Popper, that searching for counter-instances of hitherto undisproved generalisations is part of the scientific methodology which enables us to increase our scientific knowledge most quickly. But *before* the new evidence has been found this gives us no method of selection, and when it has been found the elimination of certain generalisations as false is purely deductive: there is no special kind of inductive inference at work. Besides, a vast number of possibilities always remains after each elimination. Could there be any rule for deciding on the basis of *already available* evidence which regularities should be projected? Perhaps instead of a rule specifying precisely which regularities should and which should not be projected we could have a rule specifying sufficient conditions for projecting some regularities, *e.g.*, ' predict the persistence of some regularity if in the past it has enabled more correct predictions to be made, in a wider variety of circumstances, than any other hitherto observed regularity incompatible with it '. Later on, I shall discuss the possibility of incorporating such considerations in Mr. McGowan's system of predictive policies. But it should be noted at once that the history of science shows that even regularities projected in these circumstances (*e.g.*, Newton's dynamics) have almost always turned out false

eventually. Is it possible nevertheless to find some justification for this or some other rule?

10. Before discussing the general problem of justification let us look at Mr. McGowan's attempt to show that induction is better than counter-induction. He argues (a) that the adoption of inductive policies for making singular predictions commits one to accepting general hypotheses (though at the end of section 11 he weakens this and says induction allows us to retain undisproved generalisations: but this much could be said of many alternatives to induction); (b) that a counter-inductive policy commits one to rejecting all general hypotheses; (c) mixed policies, allowing counter-induction to replace induction on occasions, contain the same flaws as counter-induction itself: since the flaws are " dominant rather than recessive ". The rejection of all generalisations is a flaw, he argues, because we are quite unjustified in building such a suppositions into our method: " there may be no true universal generalisations ", he says, " but this we must try to learn ". I have several objections to this. First, the argument to establish (a) does not work. Secondly, (c) is asserted without argument and seems to be simply false. Finally, the importance of the assumption that there are some true universal generalisations is over-rated, and it is by no means clear that we could ever ' learn ' that the assumption is true, or that it is false. Consider the following two possible basic assumptions of predictive policies:

(i) There are no true universal generalisations.

(ii) There are some true universal generalisations. (*Cf.* paragraph 9.)

Since the latter is consistent with the assumption that any particular universal generalisation, or even *every* universal generalisation expressible in any given language, is false, it cannot make a difference to one's policy for making singular predictions which of these two assumptions one makes, unless some further specification is added to (ii), which would take us back to the problem raised in the previous paragraph. As it stands (ii) is too weak to generate more than vague hopes, (which is why I previously said it is hardly worth justifying), and (i) does no more than contradict vague hopes: both are compatible with our going on indefinitely making singular predictions of an inductive sort, and both are compatible with the assumption that all hitherto observed regularities will not persist for long. (Compare paragraph 12, below.) Let us now return to steps (a) and (c) in the argument.

11. Mr. McGowan's argument to establish that the use of P_i (the inductive policy) for making singular predictions commits one to accepting general hypotheses seems to rely on some such principle as the following:

(A) If an unobserved object can be inferred to be a B according to P_i then P_i may subsequently be applied as if the object had been *observed* to be a B.

In section 11 he attempts to establish that one cannot be " serious in accepting P_i " without also accepting (A), and it is quite possible that there is something in the meaning of " serious " which justifies this claim. However, the argument does not show that accepting P_i without doing so " seriously " is somehow less rational than accepting it " seriously ", and still less does it show that P_i *entails* (A), which, I suspect, is what he really wants to show. Can anything be said to justify (A)? Possibly an argument could be constructed using principles like the following:

(B) If it is rational to believe p on evidence e, and p logically entails q, then it is rational to believe q on evidence e.

Or, more generally,

(C) If evidence e makes it rational to believe p, and evidence p & e makes it rational to believe q, then evidence e makes it rational to believe q.

However, it is assumptions like these which generate the well-known paradoxes of confirmation, so until satisfactory resolutions of those paradoxes have been found, the assumptions must be suspect. In any case, it would appear to be reasonable to modify (C) to imply that the degree of support given to q by e may be less than that given to p. In this way very remote predictions could turn out to be given very little support, and hence Mr. McGowan's step to the general hypothesis covering all the predictions, however remote, must also lead to a conclusion which has very little support. I conclude that there are too many loopholes in this step in the argument to vindicate induction. (The three principles (A), (B) and (C) really require far more discussion than I have time for.)

12. Let us now look in more detail at the criticism of counter-induction, and policies with counter-inductive elements. In section 5 paragraph (4) Mr. McGowan defines counter-induction as a policy which replaces the Principle of Persistence with the

Principle of Desistence, to the effect that all observed regularities come to an end, and sooner rather than later. Contrast this principle with the following Weak Principles of Desistence, which, it appears from his final paragraph, Mr. McGowan would claim are fatally infected with the same flaws as the strong principle:

(D_1) All observed regularities come to an end eventually (some later rather than sooner).

(D_2) Some, but not all, observed regularities come to an end, and sooner rather than later.

(D_2) Most, but not all, observed regularities come to an end and sooner rather than later.

Although the strong principle of desistence always leads to predictions which are the negations of inductive predictions, these weak principles need not. In fact, as remarked two paragraphs ago, the weakest one, (D_1) leaves it open to us to go on making inductive predictions indefinitely, since it does not tell us which regularities will end when. Moreover, (D_2) and (D_3) allow us to go on making inductive predictions in at least some cases, and *allow* us to hope that some of them are instances of true universal generalisations. But this is all that the Weak Principle of Persistence does (see paragraphs 9 and 10 above), and indeed the latter is even *implied* by the ' but not all ' in (D_2) and (D_3). Since the strong version of the Principle of Persistence has to be rejected in favour of the weak version for the reasons mentioned in paragraph 8, and since the weak principle is implied by the weak versions of the Principle of Desistence, which appear to have the advantage of fitting observed facts (namely, some temporarily acceptable generalisations have been observed to be refuted), there seems to be nothing left of Mr. McGowan's attempt to show that induction is superior to all alternatives. Indeed, policies based on one of the weak principles of desistence have the great advantage over Mr. McGowan's inductive policy that they don't lead to the contradiction mentioned in paragraph 8, since they allow us to reject as false some inductive conclusions. This allows adjustments to be made to avoid inconsistencies. (I am not claiming to be able to justify any particular method of selecting regularities to project, merely that inconsistencies are unavoidable if no selection is made.)

13. Can we modify Mr. McGowan's recursive rules for inductive inference in section 7, and the flow chart representing their operation, so as to take account of all these criticisms? It seems that the most important purpose served by the system of

rules and the chart is to ensure that all possibilities are accounted for, so that whatever happens there is a determinate next step, or in the words of section 10, " there is a perfectly good further move for every outcome ". This property is supposed to be preserved when ' yes ' and ' no ' are swapped in each starred pair, showing that a strongly counter-inductive policy need not break down as sometimes alleged. However, in order not to have to give up this claim on account of incompatible predictions resulting from application of the policy, Mr. McGowan says, in section 10, that the policy does not *force* one to accept both of two incompatible predictions. This is by no means clear from the formulations of his rules, given in section 7, for they contain the imperative ' predict ', whereas he seems to want them to say ' it is permitted but not obligatory to predict ', or some such thing. But if several contradictory predictions may all be permitted at a certain stage, it is not clear in what sense there is a " perfectly good further move for every outcome ". Perhaps one solution would be to add a further rule:

(5) If the application of the other rules leads to an inconsistent set of predictions, replace the set by the single prediction consisting of the disjunction of all its members.

This is equivalent, of course, to replacing the strong principle of persistence (all regularities persist) with a weaker one (some observed regularities persist), in the case of the inductive policy.

14. However, if we are left with nothing stronger than a policy requiring us to predict that some regularity or other will persist, it is pretty certain to be of little practical use. Can we narrow down the disjunction of predictions by adding principles enabling us to select from the vast range of undisproved generalisations? Perhaps something can be achieved by adding the Principle of Total Evidence to Mr. McGowan's set of rules, a principle which he failed to incorporate despite mentioning it in section 3. The fact that all observed A's have been observed to be B's need not be all the relevant evidence available in connexion with the problem whether the next A will be a B, as pointed out in connexion with the lemon example in paragraph 5, above. So it would seem to be advisable to incorporate the following cautionary questions concerning other possible relevant evidence before the step which reads: ' Predict $n+1$st A is a B ' (see loop I in Mr. McGowan's flow chart).

(i) Is there a C such that the first n A's are C's and the $n+1$st A is not a C?

(ii) Is there a property X possessed by all observed A's and the $n+1$st A, and a property Y possessed by the $n+1$st A such that all observed X's which are not Y's have been B's while all observed X's which are Y's have been non-B's?

(iii) Is there a generalisation which has previously led to many correct predictions concerning A's and other things and which entails that in certain conditions no A is a B?

No doubt similar cautionary questions could be inserted prior to the predictive steps in the part of the procedure concerned with probabilistic predictions. The important thing which all these questions have in common is that if the answer is ' yes ' then there is some reason, according to whatever principles underlie inductive predictions, for doubting that the next A will be a B, whereas if the answer is 'no' then no reason has been given for not going ahead with the prediction.

15. However, even if the foregoing questions were all answered in the negative it seems that more and more complex questions in similar vein could be formulated and inserted before ' Predict $n+1$st A is a B ', especially when inductions concerning functional relationships are considered, since mathematical notations provide us with a systematic way of generating infinitely many functional relationships which obviously cannot all be investigated. It is not clear how there could be any rational principle for deciding when to stop asking such cautionary questions and to go ahead with the prediction. Caution should be limited, perhaps, but how limited? Perhaps considerations of simplicity are relevant here. Perhaps in particular contexts the limits are determined in a purely *conventional* way: it is simply " not the done thing " to go on asking sceptical questions beyond a certain point. (The conventions would change with scientific progress, and one aspect of scientific genius would consist in the ability to ask the right unconventional questions.) Whatever the answer to this may be, there are two ways in which the insertion of these questions about available evidence change the character of Mr. McGowan's inductive policy. First, finding the important answers to the cautionary questions will generally require the discovery of new concepts (taking the rôle of C, X, Y, *etc.*) and there is not much hope of a mechanical rule for doing this or for ensuring that a particular attempt to answer one of the questions has exhausted all possibilities. This extends the sort of indeterminacy already present in his third rule of section 7, which poses the question " Is there a C such that the first n A's are C's and the $n+1$st A is not a

C? " in cases where the $n+1$st *A* has turned out not to be a *B*. Secondly, when the answers to the cautionary questions are positive, and there are reasons against predicting that the next *A* will be a *B*, we are left with no way of assessing whether to give more weight to the reasons *for* the prediction or the reasons *against*. It is now no longer true (if it ever was) that " there is a perfectly good move for every outcome ".

16. But should the policy even be definite in cases where the answers to the cautionary questions formulated two paragraphs ago are all in the negative, *i.e.*, when there are no positive reasons against making the prediction? Should we not include such further cautionary principles as the following?

> P_a: Don't project a regularity unless it is part of a more general regularity for which there is independent evidence, or unless it explains several independently observable less general regularities.
>
> P_b: Hedge all predictions with the qualification: ' Provided that the circumstances are sufficiently similar to those hitherto investigated '.

If it is possible to learn from past experience at all, then it seems that the refutation of many apparently well supported regularities in the past should teach us that some kind of cautionary principles should be added to any inductive policy. But even if we can find principles which, if acted on in the past, would have minimised our errors (and can we?) it is not clear how we could show that they would have similar effects in the future: after all, the consistent application of cautionary principles would require us to be cautious even about the prediction that just these cautionary principles are required for the avoidance of error in the future!

17. This last point, together with the previously mentioned difficulties in ensuring that all available evidence has been taken account of, leave me very sceptical about the possibility of constructing a set of rules which (a) incorporates everything that we can be said to have learnt from experience about the best way to learn from experience, and (b) can be shown, without circularity, to be better than all alternatives, even for future attempts to learn from experience. But perhaps a deeper reason for being sceptical lies in the unclarity of this use of ' better '. ' Better in relation to what? ' is the question which immediately arises and is usually left unanswered. One thing cannot simply be better than another: there must always be a basis of comparison relative to

which it is better. We indicate a basis (or at least part of a basis) when we answer questions like: ' better for what? ', ' better at what ', ' better in what respect? ' (I hope to publish another paper elaborating on this in the near future.) What then is the basis relative to which inductive policies are supposed to be better than others? None is mentioned by Mr. McGowan when he says:

> " The problem of induction is to show that the choice of inductive policies, in preference to any possible alternative predictive method, is a rational choice. "

What is better in relation to one basis of comparison may be worse in relation to another: in relation to the aim of being surprised as often as possible, it may well be that some obviously non-inductive method of making predictions is better, and to that extent more rational, than inductive methods.

18. Consider another example: a pack of cards is turned up one at a time and you have to guess for each one whether it will be red or black. You are given a prize for every correct guess of ' red ' and you have to have to pay a forfeit for every incorrect guess of ' black ', but nothing happens in other cases. Clearly the best (and most rational) policy in relation to the aim of winning as many prizes as possible, or the aim of minimising the number of forfeits, or both, is to guess ' red ' every time, even if you know that most of the cards will be black. But in relation to the aim of guessing correctly as often as possible, some other policy may well be better (and more rational). Now imagine the pack of cards to be replaced by a set of undisproved empirical generalisations and guessing ' red ' or ' black ' replaced by guessing ' true ' or ' false '. If the only " pay-off " is disappointment when one selected as true turns out false, then in relation to the aim of minimising disappointment the best policy is to guess ' false ' for *all*.

19. Thus, the question ' Which policy is best? ' or ' Which policy is most rational? ' cannot be answered until a basis of comparison has been specified. Is there any one basis relative to which inductive and alternative methods of inference can illuminatingly be compared? Perhaps the situation will become clearer if we look at deductively valid inferences, and see how they fare in relation to *the aim of arriving only at true conclusions,* which I shall call the basis B. Investigation of inferences of various sorts shows (trivially) that some satisfy the following condition while some do not:

C_0: the conclusion of the inference is true.

Then clearly those which satisfy the condition are better in relation to the basis B than those which do not, no matter how the conclusions are related to the premisses, *i.e.*, whether deductively or not. But in general we have to assess inferences in circumstances in which it is not known whether the condition C_0 is satisfied. We find that some inferences satisfy the following condition:

> C_1: the premisses and conclusion are related in such a way that it is impossible for the former to be true while the latter is false.

Some inferences can be shown to satisfy this condition even when it is not known whether they satisfy C_0. *If we restrict ourselves to inferences having true premisses,* those which satisfy C_1 can be guaranteed to be at least as good as any others in relation to the basis B, since they will also satisfy C_0. So for selecting among the class of inferences with true premisses an inference policy which yields only inferences satisfying C_1 is better relative to the basis B than a policy which does not, since the latter can lead to false conclusions. Thus, for selecting among the inferences with true premisses, deductive policies are better in relation to B than those which are not deductive, including inductive policies. (So much for the claim that induction and deduction cannot be compared on account of having their own standards of excellence!) Among the possible non-deductive inference policies can some be shown to be better than others, relative to the basis B (the aim of arriving only at true conclusions)?

20. The clarification of this question would require an excursion into the logic of ' better ', for which there is not space in this paper. But some of the main points can be hinted at briefly by means of an example. Suppose one wants a jacket which fits perfectly: this determines a basis of comparison. Among jackets which do not fit perfectly some are better than others relative to this basis, according to which approximate most closely to the perfect fit. But if jacket X and jacket Y differ in that the sleeves of X fit more closely than the sleeves of Y while the waist of Y fits more closely than the waist of X, we may only be able to say that in relation to the basis one is better in some respects while the other is better in other respects, and neither is better on the whole. Can any parallel to these situations be found in the attempt to compare various types of non-deductive inference policies?

21. I can only outline some suggestions for further investigation of this question: a full analysis would be too lengthy. First,

since we are considering inferences which do not satisfy condition C_1 we have to find alternative conditions which an inference can be shown to satisfy independently of discovering the truth-value of its conclusion, and then we have to examine their connexion with the basis B in order to assess their relative merits. Now if C is such a condition several questions can be asked about it and its relation to B:

(i) Is the assertion that an inference satisfies condition C justifiable *a priori, i.e.*, without empirical investigation?

(ii) Does the assertion that an inference satisfies C involve a *prediction?*

(iii) Is the assertion of a certain relation between satisfying C and the basis B justifiable *a priori?* (*e.g.*, can it be shown *a priori* that satisfying C is better than not satisfying C in relation to the basis B, other things being equal?)

(iv) Does the assertion of the relation between C and B involve any *prediction?*

The first and third questions are mentioned in order to contrast them with the second and fourth. Thus, a necessary condition for an inference to satisfy the condition C may be that it contains all known facts as premisses, in which case, unlike satisfying condition C_1 there is an empirical element. This need not matter for present purposes. However, if the answer to question (ii) or (iv) is affirmative, then in the context of attempts to justify induction circularity is unavoidable.

22. The foregoing questions could be applied to each of the conditions listed below, which, in one way or another, have been supposed by philosophers or logicians to be of importance in connexion with inferences. Most of these conditions would ideally have to be supplemented by a clause specifying that all (relevant) known facts are included in the premisses: I shall leave this out for the sake of brevity.

C_2: the inference is of a type T such that the policy of accepting premisses while rejecting conclusions of inference of that type must lead to contradictions if consistently followed;

C_3: as above, but . . . must lead to breakdown in communication;

C_4: the inference contains premisses and conclusion so related that human reason is incapable of doubting the conclusion when the premisses are accepted;

C_5: the inference is of a type T such that the policy of consistently denying conclusions while accepting premisses in inferences of that type presupposes that all generalisations are false, and so frustrates the aim of science (this is the condition stressed by Mr. McGowan);
C_6: rejecting the conclusions while accepting the premisses in inferences of this sort makes all practical decision impossible;
C_7: the conclusion is derived from the premisses in accordance with generally accepted standards of reasoning;
C_8: the conclusion is derived from the premisses in accordance with principles which, in the past, have led to a higher proportion of true than false conclusions;
C_9: the conclusion is derived deductively from the premisses by conjoining as further premisses generalisations which, in the past, have led only to correct predictions, in a wide variety of circumstances (*i.e.*, " highly corroborated " generalisations);
C_{10}: the conclusion is derived from the premisses in accordance with a rule which, if consistently followed, ensures that we discover our mistakes as quickly as possible;
C_{11}: the inference is of a type T such that the inference from all (presently) known facts to the conclusion " most inferences of type T with true premisses have true conclusions " is itself of type T;
C_{12}: the probability (in some specified sense) of the conclusion relative to the premisess is greater than $\frac{1}{2}$;
C_{13}: the conclusion is derived from the premisses according to a rule R such that in the long run R derives more true than false conclusions from true premisses;
C_{14}: as above but . . . the *probability* that in the long run R derives more true than false conclusions from true premisses, is greater than $\frac{1}{2}$, relative to all known facts;
C_{15}: the conclusion is so related to the premisses that if N believes the premisses N will feel more confident that the conclusion is true than that it is false (where N is a specified individual, or group of persons, *e.g.*, a scientist generally thought to be competent);
C_{16}: acting on the conclusion when the premisses are true will maximise expected utility;
C_{17}: acting on the conclusion when the premisses are true will maximise actual utility;
C_{18}: as above but . . . will maximise minimum possible gain (utility);

C_{19}: as above but . . . will minimise maximum possible loss;
C_{20}: the conclusion is derived from the premisses according to a rule R such that in the long run acting on conclusions derived according to R from true premisses will lead to greater gains than if any other prediction policy is used.

This is by no means a complete list.

23. Most of these conditions bear no direct relationship to the basis B (the aim of arriving only at true conclusions). And those that are related very closely to it cannot be ascribed to any particular inference without making a prediction (see paragraph 21, above). Several of the conditions make use of technical terms (' utility ', ' probability ', ' acting on ', ' in the long run ') to which it is very difficult to attach any clear and precise meaning, and in some of the possible interpretations of these terms conditions using them are irrelevant to our basis of comparison B. The condition stressed by Mr. McGowan, namely C_5, does appear to be related to the aim of arriving at true conclusions, but I have already argued above that he has not shown that it can be used to establish the superiority of inductive inferences over all alternatives.

24. But should we stick to the basis B? Some of the conditions irrelevant to B may be very important relative to some other basis of comparison of a pragmatic nature (think of the aim of winning the maximum number of prizes in the card game of paragraph 18). But even amongst those conditions which are relevant to alternative bases, some involve a predictive element, so that it cannot be conclusively settled at a time when a decision has to be taken, whether a particular inference satisfies the condition or not. Further, where it can be settled, satisfaction of the condition need not depend solely on the relation between the premisses and the conclusion, as in the case of deductive inferences and Carnap's system of inductive logic, but may depend also on the " pay-off matrix ", that is on the " prizes " and " punishments " for making correct and incorrect guesses in various possible states of affairs. Thus, even relative to some pragmatic basis of comparison the question whether a particular inference policy is better (more rational) than alternatives may depend on the answers to questions like ' What have you got to gain or lose? '—questions not normally considered by philosophers in discussing induction. (For an excellent introduction to such topics see *Games and Decisions* by R. D. Luce and H. Raiffa.)

25. To sum up: Mr. McGowan has assumed that there is a

clear distinction between inductive inferences and others, that we all know how to make the distinction, that we all agree that the inductive ones are somehow better or more reasonable than the alternatives, and I have criticised all of these assumptions. Further he has formulated the philosophical problem of induction as the problem of showing *why* the inductive ones are better, and he has attempted to show that inductive policies as represented in his flow chart are better than non-inductive and counter-inductive ones. I have criticised some of the details of his argument and put forward the counter-claim that policies based on the weaker principles of desistence are better at avoiding contradictions and conform to past experience more closely than policies based on his strong principle of persistence. Accordingly, some modifications of his predictive rules have been suggested. Perhaps most importantly of all, I have argued that the assertion that one policy or inference is better or more rational than another is an incomplete assertion until a basis of comparison has been specified, since different policies may be better or more rational in relation to different bases, and I have indicated some possible approaches for further investigation of this point. A final line of investigation which should be mentioned is the problem of deciding which of two bases of comparison is better relative to some higher-order basis of comparison, a problem which may turn out to be very important in connexion with justifications of predictive policies. It seems that I have asked more questions than I have answered. Perhaps formulating them will help someone more familiar with the field than I am to find interesting answers.