IMMODEST INDUCTIVE METHODS*

DAVID LEWIS[1]

Princeton University

Inductive methods can be used to estimate the accuracies of inductive methods. Call a method *immodest* if it estimates that it is at least as accurate as any of its rivals. It would be unreasonable to adopt any but an immodest method. Under certain assumptions, exactly one of Carnap's lambda-methods is immodest. This may seem to solve the problem of choosing among the lambda-methods; but sometimes the immodest lambda-method is $\lambda = 0$, which it would not be reasonable to adopt. We should therefore reconsider the assumptions that led to this conclusion: for instance, the measure of accuracy.

Suppose you are looking for an inductive method to trust. By an *inductive method*, I mean a systematic way of letting the available evidence govern your degree of belief in hypotheses. We can represent a method by a function **C** from pairs of propositions to real numbers in the unit interval. You *trust* the method if your degree of belief in any hypothesis h, conditionally on evidence e, is $\mathbf{C}(h \mid e)$.

One thing you can do given an inductive method **C** is to estimate the values of numerical magnitudes. By a (*numerical*) *magnitude*, I mean a function from all possible worlds to numbers. The speed of a given racehorse (in a given race), for instance, is that function whose value at any possible world w is the speed of that horse in that race in the world w (or some arbitrarily chosen value if that horse does not exist, or does not run in that race, in the world w). For any magnitude m there is a set V_m of its possible values. For each value v in V_m, there is a proposition p_v which holds at all and only those possible worlds at which m has the value v; we regard p_v as the proposition that m has the value v. The **C**-*mean estimate*, on evidence e, of a magnitude m may now be defined thus:

$$\text{(D1)}\quad \mathbf{E}_{\mathbf{C}}(m \mid e) =_{\text{df}} \sum_{v} v \cdot \mathbf{C}(p_v \mid e)$$

where v ranges over V_m. You might, for instance, wish to use your inductive method **C** to bet on the horse whose **C**-mean estimated speed, on the evidence available to you, is highest.

You should hope to give your trust to an inductive method **C** that will give you accurate estimates; you want there to be no more difference than you can help between the actual values and your **C**-mean estimates, on the available evidence, of

* Received May, 1969.

[1] I am grateful to Robert L. Goble for many valuable comments; to the University of California for a Faculty Fellowship supporting the work reported in this paper; to the U.C.L.A. Campus Computing Network for an allocation of computer time used to prepare the numerical examples shown in Figures 1 and 2; to Diane Wells for drawing the figures; and to an anonymous referee for *Philosophy of Science* who suggested certain improvements in exposition.

the various magnitudes in which you are interested. It is plausible that the undesirability of errors might rise more than linearly with the size of the errors. Thus you might wish to measure accuracy by considering, say, the squared error of the inductive method in various estimating tasks.

But you cannot just pick the most accurate method—not unless you already know the actual values of the magnitudes you wish to estimate, in which case you do not need to estimate them. The best you can do is pick the inductive method with the highest *estimated* accuracy, just as you might bet on the horse with the highest estimated speed.

The trouble is that you need an inductive method to estimate anything, even to estimate the accuracy of various inductive methods. And your selection of a method with the highest estimated accuracy will come out differently depending on which method you use to make the estimate. It is as if *Consumer Reports*, *Consumer Bulletin*, etc., each published rankings of the consumers' magazines, as they do of other products. You would have to know which one to read in order to find out which one to read.

Let us say that an inductive method **C** *recommends* an inductive method **C**′ if the **C**-mean estimate of the accuracy of **C**′ is not exceeded by the **C**-mean estimate of the accuracy of any rival method. An inductive method might or might not recommend itself. If it does, let us call it *immodest*. When asked which method has the best estimated accuracy, the immodest method answers: "I have."

We may restate these definitions more precisely, making explicit reference to the class of competing inductive methods, to the way in which accuracy is measured, and to the total evidence available for use in estimating accuracies of methods.

(D2) Method **C** *recommends* method **C**′ in the class M of methods, under the accuracy-measure A, on the evidence e, iff $\mathbf{E_C}(A(\mathbf{C}') \mid e) \geq \mathbf{E_C}(A(\mathbf{C}'') \mid e)$ for any method **C**″ in the class M.

(D3) Method **C** is *immodest* in M, under A, on e, iff $\mathbf{E_C}(A(\mathbf{C}) \mid e) \geq \mathbf{E_C}(A(\mathbf{C}') \mid e)$ for any method **C**′ in M.

Notice that it depends on the evidence which methods are immodest. It may happen according to (D3), and as we shall see it does happen under the assumptions we are about to make, that a method which is immodest on evidence e_1 is not immodest on different evidence e_2.

Does the immodesty of an inductive method give you any good reason to trust it? Certainly not every immodest inductive method deserves your trust. Consider Barker's gypsy ([1], p. 17) who, when asked whether her method of crystal-gazing is reliable, answers:

> Oh, yes indeed, you may be sure that my crystal gazing yields reliable answers to all your questions. I know that it does, for I conducted an empirical inquiry into the matter; seeking an answer to the question whether my crystal gazing is a reliable way of answering questions, I looked into my crystal ball, and the answer that I saw there was "Yes."

Immodesty is too easy to come by.

But reverse the question: would *non*-immodesty give you any good reason *not* to trust an inductive method? Indeed it would. Suppose you did trust some non-immodest method. By definition, it estimates some competing method to be more accurate than itself. So if you really did trust your original method, you should take its advice and transfer your trust to one of the competing methods it recommends. It is as if *Consumer Bulletin* were to advise you that *Consumer Reports* was a best buy whereas *Consumer Bulletin* itself was not acceptable; you could not possibly trust *Consumer Bulletin* completely thereafter.

The answer to our first question, whether immodesty is a good reason to trust an inductive method, ought to be: it depends on the competition. Any immodest method deserves your trust more than any non-immodest method, but the immodest methods must compete among themselves on other grounds. Immodesty is a necessary but not sufficient condition of adequacy for inductive methods.

The requirement of immodesty will not help you much in choosing an inductive method unless few of the otherwise adequate methods are immodest. We might expect all methods to be immodest; in that case, it will get you nowhere to require immodesty as a condition of adequacy. How many methods are immodest?

We cannot answer this question yet; we must first specify the class M of methods you wish to choose from, the accuracy-measure A you have adopted, and your total available evidence e.

Assume you intend to apply your chosen inductive method only to a limited range of inductive tasks. Your interests are confined to propositions about a certain universe of N things, and about a certain family of k kinds (mutually exclusive and jointly exhaustive) to which those things may belong. You will be content with an inductive method defined on any pair of such propositions. Assume further that you have already adopted several conditions of adequacy which rule out all but the λ-methods: the inductive methods discussed by Carnap in [2]. Your remaining problem is to pick *one* of the λ-methods.

The λ-methods are of interest partly because they comprise all and only the methods that satisfy certain rather plausible conditions of adequacy ([2], sections 4, 8), but even more because they are simple and well understood. The quickest motivation for them is as follows. First, let us stipulate that any adequate method $\mathbf{C}$ must conform to the basic axioms of conditional probability. Second, if e is a complete description of a sample and h is the proposition that a certain thing not in the sample belongs to a certain kind i, what should $\mathbf{C}(h \mid e)$ be? It should be close to the relative frequency of kind i in the sample if the sample is large, but close to $1/k$ if the sample is small. Therefore, let us stipulate that $\mathbf{C}(h \mid e)$ is the relative frequency of kind i in an *augmented* sample consisting of the actual sample plus a fictitious sample of λ things, λ/k of each kind, where λ is any positive real, or 0 or ∞. This has the desired result since when the actual sample is large compared to λ the actual part of the augmented sample predominates, whereas when the actual sample is small compared to λ the fictitious part of the augmented sample predominates. Carnap shows ([2], sections 5, 10) that these two stipulations suffice to determine a unique inductive method $\mathbf{C}_\lambda$; the parameter λ measures the method's caution in learning from experience.

Suppose next that you have adopted one of the accuracy-measures employed by Carnap in [2], sections 21–24: an accuracy-measure based on mean square errors of estimates of relative frequencies. Choose arbitrarily some small number t; we shall see that the value of this parameter does not matter. Let i range over the k kinds of things, and let r_i be the magnitude whose value $r_i(w)$ at any possible world w is the relative frequency in the universe, at world w, of things of kind i. Let T_{tw} be the set of all pairs $\langle i, j\rangle$ where i is any one of the kinds and j is any proposition true in w that completely describes a sample containing t things. We regard each such pair in T_{tw} as representing an estimating task: the task of estimating the relative frequency r_i of kind i on evidence j. Note that the set T_{tw} is finite. The error of inductive method **C** on the task represented by $\langle i, j\rangle$ at the world w is the difference between the **C**-mean estimate $\mathbf{E}_{\mathbf{C}}(r_i \mid j)$ of r_i on evidence j and the true value $r_i(w)$ of r_i at w. Carnap suggests that we measure the inaccuracy of **C** at the world w by the mean square error of **C** at w on all such tasks: we take the mean over all pairs $\langle i, j\rangle$ in the set T_{tw}. Carnap shows ([2], section 21) that the mean square error (with parameter t) of a λ-method $\mathbf{C}_\lambda$ at w is given approximately by (1).

$$(1) \qquad \frac{t - \lambda^2/k + (\lambda^2 - t)\sum_i r_i(w)^2}{k(t + \lambda)^2}.$$

The approximation consists in taking estimated relative frequencies in the rest of the universe excluding the t things described by j, rather than estimated relative frequencies in the entire universe; thus it is a good approximation when t is sufficiently small compared to N. Carnap does, in practice, use the approximate mean square error as given by (1); so let us follow his practice, defining the family of accuracy-measures $\mathbf{A}_t$ for the λ-methods as follows.

$$(\text{D4}) \quad \mathbf{A}_t(\mathbf{C}_\lambda)(w) =_{\text{df}} -\left(\frac{t - \lambda^2/k + (\lambda^2 - t)\sum_i r_i(w)^2}{k(t + \lambda)^2}\right).$$

Thus $\mathbf{A}_t(\mathbf{C}_\lambda)$ is that magnitude whose value at any world w is the approximate negative mean square error of $\mathbf{C}_\lambda$ at w on estimating tasks represented by the pairs in T_{tw}.

Suppose finally that as you set out to pick an inductive method your total available evidence e is a complete description of a certain sample containing s things such that, for each kind i, s_i things in the sample belong to kind i.

Having specified the class of inductive methods you wish to choose from, the accuracy-measure you wish to maximize, and the total evidence at your disposal, we are ready to reconsider the question: how many methods are immodest? The answer is: exactly one.

To show this, we begin by noting that $\mathbf{C}_\lambda$ recommends $\mathbf{C}_{\lambda'}$ on the evidence e if and only if λ' is chosen to maximize the $\mathbf{C}_\lambda$-mean estimate of $\mathbf{A}_t(\mathbf{C}_{\lambda'})$. Since $\mathbf{A}_t(\mathbf{C}_{\lambda'})$ is linear with $\sum_i r_i^2$—we thus denote the magnitude whose value at any possible world w is $\sum_i r_i(w)^2$—we find that the $\mathbf{C}_\lambda$-mean estimate on e of $\mathbf{A}_t(\mathbf{C}_{\lambda'})$ is given by (2).

$$(2) \qquad -\left(\frac{t - \lambda'^2/k + (\lambda'^2 - t)\mathbf{E}_{\mathbf{C}\lambda}(\sum_i r_i^2 \mid e)}{k(t + \lambda')^2}\right).$$

Setting the derivative of (2) with respect to λ' equal to 0 and solving, we obtain (3). We can easily verify that (3) gives a maximum value of (2).

$$\lambda' = \frac{1 - \mathbf{E}_{\mathbf{C}\lambda}(\sum_i r_i^2 \mid e)}{\mathbf{E}_{\mathbf{C}\lambda}(\sum_i r_i^2 \mid e) - 1/k}. \tag{3}$$

Observe that the parameter t has now vanished. That is why we were free to choose t arbitrarily. Equation (3) gives a necessary and sufficient condition for $\mathbf{C}_\lambda$ to recommend $\mathbf{C}_{\lambda'}$ under *any* one of the accuracy-measures $\mathbf{A}_t$.

Next, let h range over all possible *statistical distributions*: that is, propositions giving the relative frequencies in the universe of each of the k kinds. Since the statistical distributions are mutually exclusive and jointly exhaustive propositions, and since every statistical distribution h implies that each r_i has a definite value r_{ih}, we can easily obtain (4).

$$\sum_{ih} r_{ih}^2 \mathbf{C}_\lambda(h \mid e) = \mathbf{E}_{\mathbf{C}\lambda}(\sum_i r_i^2 \mid e). \tag{4}$$

Using certain properties of the λ-methods ([2], section 11) we obtain (5).

$$\sum_h r_{ih} \mathbf{C}_\lambda(h \mid e) = \left(\frac{s_i + \lambda/k}{s + \lambda}\right)\left(\frac{N - s}{N}\right) + \frac{s_i}{N}. \tag{5}$$

Next, let $\mathbf{C}_\Delta$ be the λ-method which would be most accurate, under any $\mathbf{A}_t$, if the relative frequencies of the kinds in the universe were the same as the relative frequencies of the kinds in the sample described by the evidence e—that is, at a possible world w where, for each kind i, $r_i(w) = s_i/s$. The results of [2], section 22, yield (6).

$$\sum_i s_i^2 = \frac{s^2\Delta + s^2 k}{k\Delta + k}. \tag{6}$$

Next, let d be the proposition that a certain two things—an arbitrarily chosen two not in the sample described by e—both belong to the same kind. This proposition is of no relevance to our topic in itself; but it happens that by considering it we can obtain some useful equations. Since the statistical distributions are mutually exclusive and jointly exhaustive, we obtain (7).

$$\mathbf{C}_\lambda(d \mid e) = \sum_h \mathbf{C}_\lambda(d \cdot h \mid e) = \sum_h \mathbf{C}_\lambda(d \mid h \cdot e)\mathbf{C}_\lambda(h \mid e). \tag{7}$$

Again using properties of the λ-methods ([2], section 11) we obtain (8).

$$\mathbf{C}_\lambda(d \mid e) = \sum_i \left(\frac{s_i + \lambda/k}{s + \lambda}\right)\left(\frac{s_i + 1 + \lambda/k}{s + 1 + \lambda}\right). \tag{8}$$

We obtain equation (9) by considering corresponding terms on the left and right. Whenever h is a statistical distribution inconsistent with e, $\mathbf{C}_\lambda(h \mid e) = 0$ and the term vanishes on both sides; but whenever h is a statistical distribution consistent with e, it can easily be shown that the left-hand factors of the terms on the left and right are equal.

(9) $$\sum_h \mathbf{C}_\lambda(d \mid h\cdot e)\mathbf{C}_\lambda(h \mid e) = \sum_h \left[\sum_i \left[\left(\frac{Nr_{ih} - s_i}{N - s}\right)\left(\frac{Nr_{ih} - 1 - s_i}{N - 1 - s}\right)\right]\mathbf{C}_\lambda(h \mid e).\right.$$

Given the system of equations (3)–(9), it is now merely a matter of laborious algebra to solve for λ' in terms of N, s, and Δ.

The plan is as follows. First substitute the right hand side of (8) for the left hand side of (7) and the right hand side of (9) for the right hand side of (7). After simplifying the resulting equation with the aid of (4), (5), and (6), it becomes possible to solve for $\mathbf{E}_{\mathbf{C}\lambda}(\sum_i r_i^2 \mid e)$ in terms of N, s, Δ, and k. Substituting this solution into (3) and simplifying further, we eventually obtain (10) as a necessary and sufficient condition for $\mathbf{C}_\lambda$ to recommend $\mathbf{C}_{\lambda'}$.

(10) $$\lambda' = \frac{C_1\lambda^2 + C_2\lambda + C_3}{C_4\lambda^2 + C_5\lambda + C_6}$$

where
$$\begin{aligned} C_1 &= N^2 + \Delta N^2 - N - \Delta N + s + \Delta s - s^2 \\ C_2 &= 2\Delta N^2 s + 2N^2 s - 2Ns^2 + \Delta s^2 \\ C_3 &= \Delta N^2 s^2 + \Delta N s^2 \\ C_4 &= N + \Delta N - s - \Delta s + s^2 \\ C_5 &= N^2 + \Delta N^2 + 2Ns^2 - \Delta s^2 \\ C_6 &= N^2 s^2 + N^2 s + \Delta N^2 s - \Delta N s^2. \end{aligned}$$

Notice that the parameter k has vanished and that the only relevant properties of the evidence e turn out to be those given by the two numbers s and Δ. The behavior of the recommendation relation between λ' and λ, specified by (10), is illustrated in Figure 1.

Setting $\lambda' = \lambda$ in (10), we obtain the cubic equation (11) as a necessary and sufficient condition for $\mathbf{C}_\lambda$ to be immodest.

(11) $$0 = C_4\lambda^3 + (C_5 - C_1)\lambda^2 + (C_6 - C_2)\lambda - C_3.$$

If the sample described by e is empty, $s = 0$; if the sample is uniform, containing things of only one kind, $\Delta = 0$. In either case, $\lambda = 0$ is the only non-negative real solution of (11). If s and Δ are both positive, on the other hand, C_4 and C_3 are positive; hence (11) has a positive solution, and $\lambda = 0$ is not a solution. If s and Δ are positive, moreover, $C_5 - C_1$ cannot be negative unless $C_6 - C_2$ is also negative, so by Descartes' rule of signs (11) can have at most one positive solution. Thus in every case (11) has a unique non-negative solution. The behavior of such solutions of (11) is illustrated in Figure 2.

This completes the proof of the result stated earlier; exactly one λ-method is immodest. More precisely:

> *Theorem*: Let M be the class of λ-methods; let A be one of Carnap's mean square error accuracy-measures $\mathbf{A}_t$; and let e be a complete description of a sample. Then exactly one λ-method is immodest in M, under A, on e. The immodest λ-method is $\mathbf{C}_0$ iff the sample described by e is either empty or uniform.

This theorem may seem welcome. Our new condition of adequacy, immodesty, and

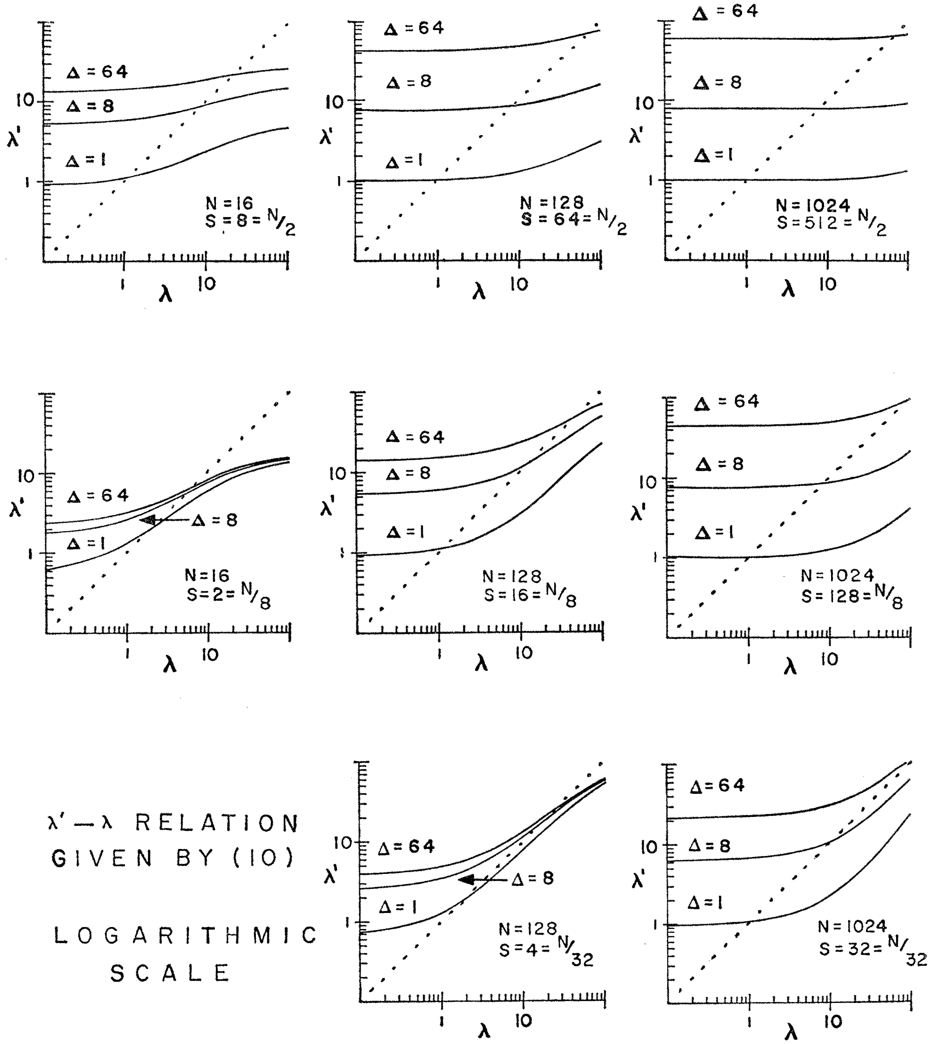

FIGURE I

the conditions of adequacy that restrict you to the λ-methods are enough to solve completely your problem of choosing an inductive method. You need only choose the one remaining adequate method. The conditions of adequacy thereby determine the degrees to which you should believe all propositions of interest to you, given total evidence consisting of a complete description of a sample.

But you should not accept this seeming solution to your inductive problem. In view of the second part of the theorem, it is not satisfactory. In some cases the immodest λ-method is $\mathbf{C}_0$; and $\mathbf{C}_0$, as Carnap argues ([2], section 14), is an extremely

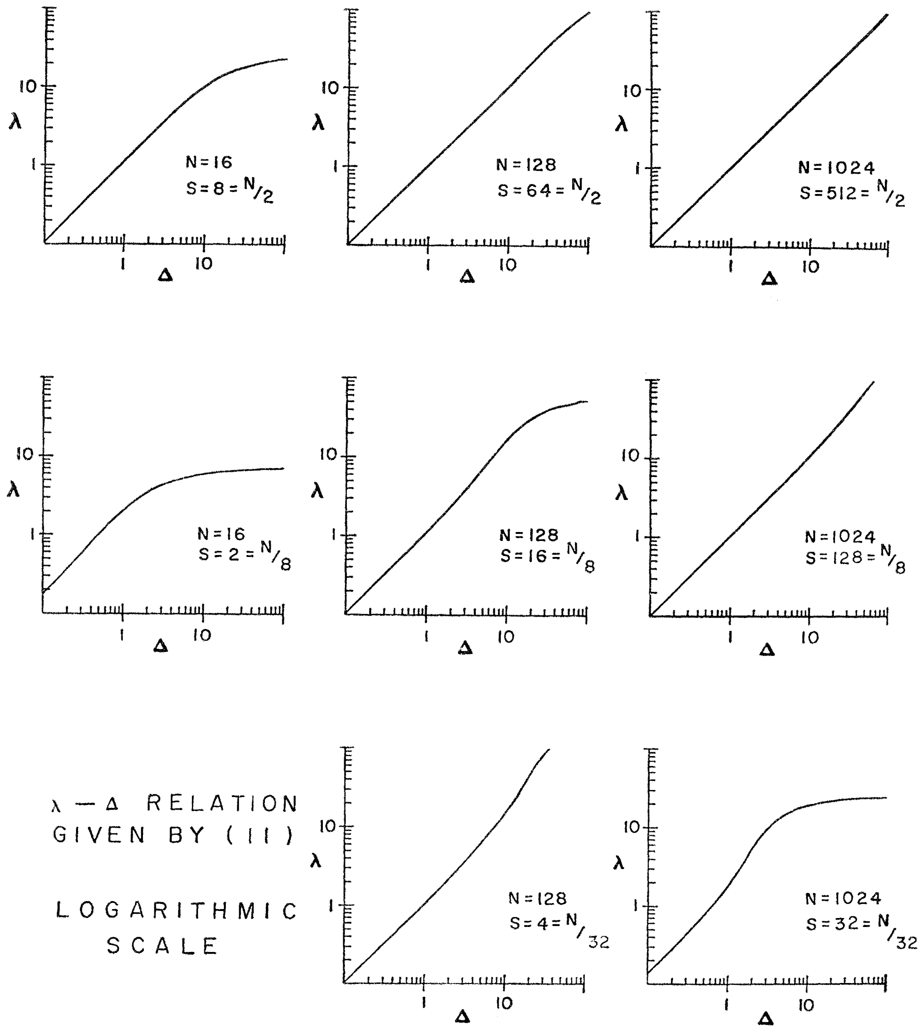

FIGURE 2

unreasonable method. It calls for jumping to conclusions, with absolute certainty, on little or no evidence.

Suppose, for instance, that your total evidence *e* is the proposition that a certain thing belongs to a certain kind *i*. Thus *e* describes a uniform sample of size 1. Then $\Delta = 0$, so the immodest λ-method on *e* is $\mathbf{C}_0$. Let *h* be the proposition that a certain other thing is also of kind *i*; $\mathbf{C}_0(h \mid e) = 1$. In terms of our original motivation for the λ-methods, the fictitious part of the augmented sample vanishes, so we take the relative frequency of kind *i* in the actual sample. The small size of the sample makes no difference. Knowing only that one thing is of kind i, you are supposed to believe to degree 1—to be absolutely certain—that the other thing is too. Surely that would be unreasonable.

Or suppose you have no evidence; e is merely the necessary proposition. Then $s = 0$, so the immodest λ-method on e is again $\mathbf{C}_0$. Let d be the proposition that a certain two things are both of the same kind; $\mathbf{C}_0(d \mid e) = 1$. You are supposed to be certain *a priori* of the contingent proposition d. That too would be unreasonable.

If $\mathbf{C}_0$ is inadequate, and if only λ-methods are adequate, and if only immodest methods are adequate, and if your total evidence happens to be a description of an empty or uniform sample, then you will be left with *no* adequate inductive method. You will have no reasonable way to assign degrees of belief to propositions on the basis of your evidence.

What can we do about this conclusion? We might accept it; but it amounts to a severe inductive scepticism. Ordinary scepticism is content to claim that there is no good reason to adopt any particular inductive method, but this scepticism is worse: it claims that there are good reasons *not* to adopt any given inductive method. So much the worse for any philosophical argument that leads to such a conclusion!

It will not help to use an exact expression for mean square error in place of the approximation (1). I used that approximation for simplicity and to stay close to [2]. But an exact expression is known (R. Carnap, personal communication); and when that is used in place of (1) in defining accuracy, it turns out to lead to the same unwelcome conclusion: $\mathbf{C}_0$ is uniquely immodest on evidence consisting of an empty or uniform sample.

We can hardly overcome our objections to choosing $\mathbf{C}_0$. If trusting $\mathbf{C}_0$ in the cases I described would not be a clear case of inductive unreason, what would be?

I do not think we should escape by rejecting immodesty as a condition of adequacy. Consider what that would mean. If you wish to maximize accuracy in choosing a method, and you have knowingly given your trust to any but an immodest method, how can you justify staying with the method you have chosen? If you really trust your method, and you really want to maximize accuracy, you should take your method's advice and maximize accuracy by switching to some other method that your original method recommends. If that method also is not immodest, and you trust it, and you still want to maximize accuracy, you should switch again; and so on, unless you happen to hit upon an immodest method. Immodesty is a condition of adequacy because it is a necessary condition for stable trust.

We might escape by looking beyond the λ-methods, hoping that in some larger class of inductive methods we will always find an immodest method better than $\mathbf{C}_0$. Carnap gives conditions of adequacy that rule out all but the λ-methods; but, as he recognizes, some of these conditions are only *prima facie* plausible. Moreover, there are certain well-known objections to the λ-methods, independently of the problem of the unique immodesty of $\mathbf{C}_0$.

Alternatively, we might escape by rejecting Carnap's mean square error accuracy-measures; I prefer this way out. The reasons for demanding immodesty under whatever accuracy-measure you want to maximize seem to me strong, but it is not at all obvious that you should want to maximize accuracy as measured by mean square error of estimates of relative frequencies of kinds. These measures are suggested by well-established practices in statistics, for instance least-squares curve-fitting. We have studied them because Carnap used them in [2], but Carnap did not argue for them there. If rejecting them is an easy way out of the problem of the unique immodesty of $\mathbf{C}_0$, that seems a rather good reason for rejecting them.

One plausible change in the accuracy-measure comes to mind at once. Perhaps in taking the mean square error of estimates of relative frequencies of kinds on the basis of samples of size t, we should take the mean not over *all* such samples but only over those which include the sample described by our total evidence e. (This would mean choosing $t \geq s$.) Why care about error in cases we already know cannot arise? This change might be appropriate on other grounds, but it will not solve our difficulty: $\mathbf{C}_0$ is still uniquely immodest when $s = 0$.

To summarize: I have argued that immodesty—in the class of otherwise adequate methods, under an appropriate accuracy-measure, on the total evidence—is a necessary condition of inductive adequacy. Whether it is a condition that will help much in choosing a method depends on how selective it is. When it is applied to the λ-methods, using Carnap's accuracy measures, it is extremely selective. But it is *too* selective, since sometimes there is no adequate method left. I take this not as an objection to the condition of immodesty, but rather as a reason to expand the class of eligible inductive methods, to find a different accuracy-measure, or both. Having done one or both, we will face a new version of the question: how many, and which, inductive methods are immodest?

REFERENCES

[1] Barker, S. F., *Induction and Hypothesis*, Cornell University Press, Ithaca, New York, 1957.
[2] Carnap, R., *The Continuum of Inductive Methods*, University of Chicago Press, Chicago, 1952.