

A Quantitative Occam's Razor¹

Rafael Sorkin

Institute for Advanced Study, Princeton, New Jersey 08540 and Center for Theoretical Physics, Department of Physics and Astronomy, University of Maryland, College Park, Maryland 20742

Received May 31, 1982

Interpreting entropy as a prior probability suggests a universal but "purely empirical" measure of "goodness of fit." This allows statistical techniques to be used in situations where the correct theory—and not just its parameters—is still unknown. As developed illustratively for least-squares nonlinear regression, the measure proves to be a transformation of the R^2 statistic. Unlike the latter, however, it diminishes rapidly as the number of fitting parameters increases.

1. INTRODUCTION

Statistics, as commonly practiced, suffers from well-known conceptual difficulties. The textbook procedure is to assume provisionally a "null hypothesis," H_0 , and then reject it if it leads to too small a probability for the actual outcome or "data," D . But D , being only one of very many possible outcomes, is never very likely. To make up for this one lumps D with other unlikely outcomes, but the manner of lumping—equivalently the choice of statistic—is subject to whim.² Besides, it does not really make sense to reject H_0 unless some other tenable hypothesis H_1 explains D better.³ The Bayesian method overcomes these problems, but introduces a new source of subjectivity with the need to assign "prior probabilities." The

¹Supported in part by the NSF under Grant No. PHY78-24275.

²One could remove this ambiguity by finding and using a "universal statistic". This seems to be the proposal of Kullback (1959) with the statistic being the (log of the) "likelihood ratio."

³Of course this sketch is a caricature, but the dangers referred to are ones which thoughtful statisticians have been able more to warn against than to eliminate systematically. (Kruskal and Tanur, 1978).

lack of a systematic way to do this becomes especially disturbing when one is dealing with an infinite-dimensional parameter space (function space).

Probably none of these difficulties is too serious when it is a question of estimating a parameter in a theory of known form, and certainly they are all unimportant in the large N limit, where normality always prevails and all applicable prescriptions agree. But what about purely “phenomenological” applications where there are relatively few observations and even the *functional form* of the true distribution is unknown? In this type of situation one might doubt whether statistics makes sense at all, but unfortunately such a “phenomenological fit” to the data is often all that one really has, especially in fields such as sociology and economics where usually one knows neither the form of the underlying functional relation (if any) nor the distribution of the random deviations from that relation (the “errors”). (See, e.g., Hanushek and Jackson, 1977.) In the physical sciences one usually does know the theory in advance, but even here “phenomenological” questions of the sort “Are galaxies randomly distributed in the sky?” can arise. And then how is one to decide whether a seeming regularity—say, a clumping or a presence of filaments—is really present?

As I just said, “Bayesian statistics” allows one to *pose* such questions, but it can answer them objectively only to the extent that the notions of a “theory” and of its “prior probability” can be freed of the subjective interpretation they usually carry. This paper will attempt such a liberation in two steps. First we will reinterpret “probability of the theory, T ” to mean “probability that the state of the universe is such that T holds”; and then we will try to use the formulas $S = k \lg N$ and $S/k = -I$ (I being information) to estimate this probability. In other words we take seriously the fact that even such things as societies or economies are ultimately physical systems and therefore try to apply *universally* the basic formulas of statistical mechanics and information theory.

To see what this might mean in practice we will focus on the particular problem of (nonlinear) regression or “curve fitting” and will derive a quantitative criterion of “goodness of fit” which will be given explicitly for the class of least-squares fits.

2. ENTROPY AND PRIOR PROBABILITIES

In “curve fitting” one is seeking the functional dependence—if any—of some variable y on a second variable x . Usually one writes the presumed dependence in the form

$$y = f(x) + u \tag{1}$$

where u is the "error," but for simplicity and generality we can deal directly with the collection of probabilities $P(y; x)$.

Definition. The data D are the collection of "observed" pairs.

(For example, we might be seeking the dependence of height on age. Then x_i could be the age in months of the i th individual and y_i his or her height to the nearest centimeter.) We will also *assume* that the observations are independent, and can then state the following:

Definition. A (phenomenological) theory T is an assignment to each possible pair (y, x) of a number $P(y; x)$ representing the *hypothetical probability* of y given x .

Here x , the variable regarded as "independent," must range over some finite set, and for each of its values, y must range over some finite set appropriate to x .

Now let $-I(T)$ be the log of the (unnormalized) "prior probability" of theory T , and let $-I(D|T)$ be the log probability of D according to T (cf. Penrose, 1970):

$$I(D|T) = \sum_{i=1}^N \lg P(y_i; x_i)^{-1} \quad (2)$$

Then the probability that *both* T is true *and* D occurs is $p(T) \equiv e^{-I(D,T)}$, where

$$I(D, T) = I(D|T) + I(T) \quad (3)$$

As is well known this p results as the "posterior probability" of T when the standard rules of probability are applied, and then the theory which maximized $p(T)$ would be the "best bet."

But why should the rules of probability apply to theories? Well the y_i and the x_i must be observables of some physical system (a person, a collection of galaxies, an economy, etc.) and it is not necessarily unwarranted to ask for the probability for a *general* system in a given class to possess (commuting) observables y, x and to be in a state such that the $P(y; x)$ have the values specified by a given theory T .

Having made this leap let us continue a bit further. What we have called the "system" is really just a state of a more general system. (For example, an economy is a collection of atoms in a particular quantum state.⁴) Treating the system as a "black box" characterized by the probabili-

⁴It is not clear, however, that the notion of observable can be assimilated to that of state this way.

ties $P(y; x)$, we can ask what is the number (suitably defined) of states which yield a given set of P 's, that is to say, a given T . If N is the number of states in this collection—let us call it the collection of “realizations” of T —then $\log N = S(T)/k = -I(T)$ is the corresponding *entropy*.

Conversely by estimating S we would discover the needed “prior” probabilities e^S . Equivalently we could estimate the *information* needed to build (at least) one of the possible realizations, R , of T from the given constituents. If $I(R)$ is the information needed for a particular realization (equal by definition to the minimum entropy *created* in actually building R), then as in thermodynamics, $\max_R \{-I(R)\}$ is often a reliable estimate of $-I(T)$. Our rule then will be to estimate $I(T)$ from the *simplest* concrete realization of T .

This is still quite vague in practice, but we can make it much more definite by restricting ourselves to simple realizations of a special sort, namely, to Turing machines which can compute the function $(y, x) \rightarrow P(y; x)$.⁵ (It is clear that an actual system realizing T could be constructed from such a Turing machine with little extra effort.⁶)

Of course with sufficient *general* knowledge about the type of system under consideration one could conceivably show in a particular case that such Turing machine estimates of $I(T)$ were inappropriate. But such knowledge would be precisely the theory which, by definition, is lacking in “phenomenological” applications. When a theory is not lacking, one should look directly thereto for the prior probability of T . Also notice that many general systems could actually be set up to mimic Turing machines, and the idea that a computationlike process exists as a *quotient* of the system is not nearly as bizarre as the idea that it exists within it as a *subsystem*.

We have now arrived at a prescription which (although it is still not fully defined) might have more immediate plausibility for some people than the chain of steps leading to it. Let us formulate it directly.

Criterion. That theory T best fits given data D which minimizes $I(D, T) = I(D|T) + I(T)$, where $I(T)$ is the complexity (information content) of the simplest⁷ Turing machine which can compute the probabilities $P(y; x)$ defining T .

⁵By assuming such a machine exists we are of course restricting the possible theories, but only in a way harmless on the phenomenological level.

⁶In this connection note that quantum systems can furnish truly random numbers.

⁷This will always be well defined (relative to a definition of Turing machine complexity) because the number of machines with complexity less than a given value is finite. In fact, one can go further and ask whether the prior probability $e^{-I(T)}$ is normalizable over the class of all Turing machines. If it is, or at least if $e^{-I(D, T)}$ is normalizable for fixed D , then our earlier interpretation, following equation (3), of $p(T)$ as a “posterior probability” takes on a precise meaning. In connection with our criterion notice also that when D is data for a large number

Thus stated, the criterion appears as a kind of Occam's razor, balancing *naive* goodness of fit [as measured by the log likelihood, $-I(D|T)$] against complexity of the theory achieving the fit [as measured by $I(T)$].

Unfortunately the correct definition of Turing machine complexity for our purposes is far from evident despite the efforts of many workers (Fine, 1973; Yasuhara, 1971). Nevertheless we will see in the next section that attention to parameter-storage requirements can enable one to say something about $I(T)$ even in the absence of a general theory of Turing machine complexity.

A development strictly analogous to that which follows would thus go through for *any* theory whose information content can be estimated in terms of its need for parameter storage. In particular the treatable theories are in no way restricted to those hypothesizing normal errors.⁸ However, for analytical convenience in estimating $I(D, T)$ we *will* henceforth restrict ourselves to this class of theories, i.e., to fits of the least-squares type.

3. APPLICATION TO (NONLINEAR) LEAST-SQUARES REGRESSION

Although an information-theoretical viewpoint really presupposes bounded discrete data (otherwise $I = \infty$) it will be harmless, and analytically convenient, to treat y as continuous and ranging from $-\infty$ to ∞ . The character of x will not matter, but for definiteness we can imagine it as a column vector with rational entries. Also, we will define for any class of theories, C ,

$$I(D[C]) = \min_{T \in C} \{I(D|T) + I(T)\}$$

as the value of $I(D, T)$ attained by that theory of type C which affords the best "phenomenological fit" to D .

3.1. The Class of Theories C_0 . Consider first the class C_0 of theories asserting that y does not depend on x at all but is normally distributed with

N of actually independent repetitions of the same experiment, then as $N \rightarrow \infty$ the T which predicts the true probabilities will eventually do better than every other. More precisely let $D = (D_1, D_2, \dots, D_N)$, let the true probability of D be $\prod p(D_i)$, and consider only theories T of the form $P(D) = \prod_{i=1}^N q(D_i)$. For large N the (essentially) constant term $I(T)$ becomes negligible in comparison to $I(D|T)$. Hence the maximum likelihood estimate for q becomes best, and this is known to approach the true probability function, p .

⁸In fact, our formal criterion is not really restricted to curve fitting at all nor to the assumption of repeated independent observations. As long as T is an assignment $Y \rightarrow P(Y)$ [or $Y \rightarrow P(Y; X)$] of probabilities to overall outcomes, a calculation analogous to that given below could be performed.

mean μ , and standard deviation σ , i.e., which postulate

$$P(y; x) = (2\pi\sigma^2)^{-1/2} \exp\left[-(y - \mu)^2 / 2\sigma^2\right] \tag{4}$$

for some μ and σ . Actually (4) cannot be a perfect equality. Rather, since we are identifying theories with (equivalence classes of) Turing machines and no Turing machine calculates with infinite precision, the theories of type C_0 are really those which *approximate* a normal distribution to some degree of accuracy. Let us estimate $I(D|C_0)$ for this class.

Taken literally the distribution (4) leads to an $I(D|T)$ of the form

$$I(D|\sigma, \mu) = \frac{N}{2} \lg(2\pi\sigma^2) + \sum_{i=1}^N \frac{(y_i - \mu)^2}{2\sigma^2} \tag{5}$$

N being the number of data points, or “observations.” For given D this is a minimum when

$$\mu = \mu_0 = \bar{y} \equiv N^{-1} \sum y_i$$

and

$$\sigma = \sigma_0 = \left(\overline{y^2} - \bar{y}^2\right)^{1/2}$$

Expanding (5) to second order about these values yields

$$I(D|\sigma_0 + \Delta\sigma, \mu_0 + \Delta\mu) \approx \frac{N}{2} \lg(2\pi e \sigma_0^2) + N \left(\frac{\Delta\sigma}{\sigma_0}\right)^2 + \frac{N}{2} \left(\frac{\Delta\mu}{\sigma_0}\right)^2 \tag{6}$$

(where $e = 2.71828\dots$), which shows how imprecision in σ and μ affects the information content of D with respect to T .

Now let TM be a Turing machine whose corresponding theory T [= the set of values (P, y, x) computed by TM] is in the class C_0 . We can imagine TM as a computer with program and hardware and regard the length of the program residing in core as an estimate of (more properly a lower bound for) the information content of TM. (See Fine, 1973.) Then the length of the shortest program for T will be an estimate of $I(T)$. Since $P(y; x)$ depends on the parameters σ, μ , our machine TM must in effect have access to them, and the simplest way to achieve this will probably (except for very special values such as $\mu = 0$) be just to store σ and μ directly as part of the program itself. This need to store its parameters implies a lower bound for the information content of TM, and therefore for $I(T)$. In

the sequel we will simply replace $I(T)$ by this lower bound, thereby acquiring an approximation which neglects the information corresponding, in our computer image, to the hardware and to the part of the program carrying out the actual computation.

The storage required to hold σ and μ to precisions $\delta\sigma$, $\delta\mu$ is approximately

$$\lg \delta\mu^{-1} + \lg \delta\sigma^{-1} = I(\delta\sigma, \delta\mu) \quad (7)$$

(For convenience we evade here and below the question—related to the actual discreteness and boundedness of y —of the units in which y , σ , and μ are stored.) If, further, we approximate in (6) $\Delta\mu^2$ by its mean, $\delta\mu^2/12$, with respect to a uniform distribution in the interval $[-\delta\mu/2, \delta\mu/2]$, do the same for $\Delta\sigma^2$, and add the result to (7), we find for $I(D|T) + I(T)$ the approximation⁹

$$\frac{N}{2} \lg(2\pi e\sigma_0^2) + \frac{N}{12} \left(\frac{\delta\sigma}{\sigma_0}\right)^2 - \lg \delta\sigma + \frac{N}{24} \left(\frac{\delta\mu}{\sigma_0}\right)^2 - \lg \delta\mu \quad (8)$$

Minimizing this with respect to $\delta\mu$ and $\delta\sigma$ furnishes finally the approximation for $I(D[C_0])$:

$$I(D[C_0]) \approx \frac{N-2}{2} \lg(2\pi e\sigma_0^2) + \lg N + 1.70 \quad (9a)$$

$$\sigma_0^2 = \overline{y^2} - \bar{y}^2 \quad (9b)$$

which corresponds to storing σ and μ with precisions

$$\delta\sigma = (6/N)^{1/2} \sigma_0, \quad \delta\mu = (12/N)^{1/2} \sigma_0$$

It seems remarkable that these values of $\delta\sigma$ and $\delta\mu$, which here represent the optimal precision for *storing* σ and μ , equal the fluctuations one would expect to see (for large N) in the sample mean and variance!

3.2. The Class C_1 . Let us generalize to theories which continue to assume independent normal errors but which allow for a functional dependence of y on x . In particular consider the class C_1 of theories corresponding

⁹In line with the neglect of the part of the program concerned with actual computation as opposed to parameter storage, we neglect also the increase in $I(D|T)$ due to computational round-off error.

to a least-squares fit of the data to a fixed set of functional forms

$$y = f(x; \theta_\alpha) \tag{10}$$

where $\theta_0, \theta_1, \dots, \theta_K$ are parameters on which f depends. [For example, $f(x; \theta)$ might be a polynomial in the components of x and the θ_α its coefficients.] In other words, a particular $T \in C_1$ says (to some approximation) that

$$P(y; x) = (2\pi\sigma^2)^{-1/2} \exp\left\{ \frac{-[y - f(x; \theta)]^2}{2\sigma^2} \right\} \tag{11}$$

Notice that T involves a total of $K + 2$ parameters: σ and the $K + 1$ parameters θ_α .

As before we have

$$I(D, T) = I(D|\sigma^* + \Delta\sigma, \theta_\alpha^* + \Delta\theta_\alpha) + I(\delta\sigma, \delta\theta_\alpha) \tag{12}$$

where σ^* and θ_α^* are the optimum values of σ and θ_α (the values we attempt to store) and $\Delta\sigma$ and $\Delta\theta_\alpha$ are the deviations therefrom when σ and θ_α are stored with precisions $\delta\sigma$ and $\delta\theta_\alpha$.

Without having the exact functional form of f one cannot find $I(D[C_1])$ precisely. Nonetheless one can without specializing f still estimate very crudely the minimum of (12) with respect to $\sigma^*, \theta_\alpha^*, \delta\sigma$, and $\delta\theta$. This is done in the Appendix. To express the result most conveniently, let us define

$$V = \overline{y^2} - \bar{y}^2 \\ = \text{sample variance of } y \tag{13}$$

and denote by ' s^2 ' the minimum mean-square residual, corresponding to the choice $\theta = \theta^*$. Then the minimum of (12) occurs for

$$\sigma^* \approx [(N - 2)/(N - K - 2)]^{1/2} s \tag{14}$$

and has the value

$$I(D[C_1]) \approx \frac{N - K - 2}{2} \lg \frac{2\pi e N s^2}{N - K - 2} + \frac{K}{2} \lg(2\pi e V) \\ + \frac{K}{2} \lg \frac{N - 2}{12} + \lg N + 1.7 \tag{15}$$

Notice incidentally that σ^* has been automatically revised upward from s by

an amount which, for large K and N , reproduces perfectly the usual adjustment for " K degrees of freedom"!

3.3. Comparing Different Functional Forms. Given any pair of theories, T' and T'' , the difference

$$H(T'', T') = I(D, T') - I(D, T'') \quad (16)$$

estimates the log of the ratio of the posterior probabilities of T'' and T' , and can be said to measure how much more "informative" theory T'' is than T' with respect to the data D . Within the class of "least-squares fits," i.e., of theories of some class C_1 , it is convenient to refer everything to the simple class C_0 and define

$$H(C_1) = I(D[C_0]) - I(D[C_1]) \quad (17)$$

or, equivalently, $H(C_1) = H(T_1^*, T_0^*)$, where T_1^* and T_0^* are the best theories in their respective classes. To evaluate H we need only observe that with the present notation (suited to C_1) the parameter σ_0^2 in equations (9) should be called V . With this substitution the difference of (9a) and (15) becomes

$$H \approx \frac{N - K - 2}{2} \lg \frac{V/(N-2)}{s^2/(N-K-2)} - \frac{K}{2} \lg \frac{N-2}{12} \quad (18)$$

Before discussing this expression let us recall that it can be a good approximation to H only when we can ignore the terms in $I(T)$ which we identified with the computer hardware and with the parts of the programs concerned with actually performing calculations. If we imagine that both programs use the same machine then the hardware term drops out of the difference, H . But the terms due to the programs cannot really be ignored; in fact they are needed in general to disfavor extremely elaborate forms for f in (10). In practice, however, f is usually very simple—in fact it is often linear¹⁰—so that our approximation may not be too bad. Nevertheless (18) does unfairly favor C_1 over C_0 to some extent.

¹⁰For linear $f [f(x) = \sum b_\alpha x^\alpha + \mu]$ the very crude estimates of the Appendix can be improved because $I(D|\sigma, \theta) = I(D|\sigma, \mu, b_\alpha)$ is then a calculable explicit function of its arguments. With respect to a particular strategy for storing the parameters, $I(\delta\sigma, \delta\theta)$ can also be found explicitly and the effect of the boundedness and discreteness of y can be taken into account. For a seemingly reasonable choice of storage-strategy such an analysis leads to an H differing from (18) by an expression AK where A depends on the particular data D and is of order unity unless the matrix $x^\alpha x^\beta - \bar{x}^\alpha \bar{x}^\beta$ is badly ill conditioned or R^2 is very small (very poor fit).

Written in terms of the standard definition, $R^2 = 1 - s^2/V$, equation (18) reads

$$H \approx \frac{N - K - 2}{2} \lg \frac{(N - K - 2)/(N - 2)}{1 - R^2} - \frac{K}{2} \lg \frac{N - 2}{12} \quad (19)$$

That R^2 has an information-theoretic meaning has long been recognized ([Kullback, 1959], Chap. 10, Sections 3 and 4). In fact from (5), (11), and the definition (2), it follows immediately that

$$\min_{\sigma, \mu} I(D|\sigma, \mu) = \frac{N}{2} \lg(2\pi eV)$$

$$\min_{\sigma, \theta} I(D|\sigma, \theta) = \frac{N}{2} \lg(2\pi e s^2)$$

Hence

$$\min_{\sigma, \mu} I(D|\sigma, \mu) - \min_{\sigma, \theta} I(D|\sigma, \theta) = \frac{N}{2} \lg \frac{V}{s^2} = \frac{N}{2} \lg(1 - R^2)^{-1} \quad (20)$$

to which we would have been led in place of H had we set out to minimize $I(D|T)$ alone, rather than its sum with $I(T)$.¹¹

For purely phenomenological purposes, however, a theory with only two parameters is not equivalent to one with many parameters, even if the latter reduces to the former for certain values of the extra parameters. Other things being equal the former, *simpler* theory is preferable, and (19) takes account of this by diminishing (20) in a manner that grows more and more severe as K approaches N . In particular a fit with more parameters than data points is automatically excluded because of the factors $(N - K - 2)$ in (19).

4. DISCUSSION

Taken together with (2) and (3), any workable definition of $I(T)$ provides a completely general measure, H , in terms of which one can

¹¹Notice that (20) is just the log of the maximum-likelihood ratio for C_1 versus C_0 . It follows immediately from this that as $N \rightarrow \infty$ with K, V, s^2 fixed, the leading terms of (18) coincide with the expression, $\lg(\text{max. likelihood ratio}) - K/2 \lg N$, derived in [Schwarz (1978)] as the prior-independent, *asymptotic* form of a Bayesian posterior. To discover the basis for this agreement one might begin by asking what must be added to the assumptions of [Schwarz (1978)] to obtain the particular *finite-sample* expression, H .

compare *any* two phenomenological theories intended to describe given data D .

A phenomenological theory in the sense of this paper admits realization by Turing machines and, insofar as one can ignore ambiguity arising from the possibility of different representations of input (x, y) and output $P(y; x)$, a given T is realized by a unique class of machines. With respect to such a realization—and with respect to a definition of the information content of a Turing machine—the evaluation of $I(T)$ becomes in effect a problem in recursive function theory.

The work of the last section attempts to approach the least-squares method of curve fitting from this standpoint by regarding such a fit as a phenomenological theory characterized by the parameters σ , θ_α , $\delta\sigma$, $\delta\theta_\alpha$ of equations (10)–(12). In this important special case, equation (19) should be a reasonable first approximation to H , at least insofar as most of $I(T)$ can be regarded as subsisting in the parameters θ rather than in the capacity¹² to compute f . The difference in H between, for example, a logarithmic and a quadratic fit would then express in absolute units the difference in the amount of information “extracted” from the data by the one fit compared to the other. In particular H itself compares a given fit to one denying any functional dependence of y on x .

When (19) is large for a particular functional form C_1 , we would like to say that the observed variation of y with x is “meaningful.” But how do we know that some other, non- C_1 theory based on a completely different functional relation between y and x (perhaps in conjunction with a highly non-normal and x -dependent distribution) might not offer a still smaller $I(D, T)$ than the least-squares fit in question? From the point of view advocated here, the only fully satisfactory course would be to examine every Turing machine small enough to be relevant to the data and call significant only those features shared by all T which came sufficiently close to realizing the absolute minimum of $I(D, T)$. Until we have a simple way to survey so many possibilities, though, it looks like (19) could serve as a useful guide to whether an observed variation of y with x should be taken seriously.

As a first step beyond this one could try to develop tests which, when applied to particular data D , would yield a lower bound to $I(D[C])$ for a reasonably large class C of choices of a functional form $f(x; \theta)$ and of a

¹²One can probably estimate this extra complexity for simple functional forms (in particular for the linear form) and correct H accordingly. Alternately one could render it negligible by lumping together a sufficiently large number of phenomenological studies all of which used the same functional form for f . This would make sense only when all the studies really were in some way part of a larger social project, such as determining the toxicity of a large number of industrial chemicals.

probability distribution for the error term, $y - f(x; \theta)$. Of course statisticians have already constructed many such alternative C 's (generalized least-squares, logistic...) and for each one it should be easy to obtain the estimate of $I(D[C])$ analogous to the estimates, (9) and (15). The more of these theories one tried with given data, the more confident one could feel with that particular one (or ones) which attained the least value of $I(D|T) + I(T)$.

APPENDIX: ESTIMATION OF $I(D|C_1)$

We will assume that the parameter θ_0 is an overall constant, μ , in f and reserve the symbol " θ " for θ_α , $\alpha \geq 1$, so that f has the form

$$f(x; \mu, \theta) = \mu + g(x; \theta_1 \cdots \theta_\kappa) \quad (\text{A1})$$

For arbitrary σ , μ , θ , $I(D|\sigma, \mu, \theta)$ is from (11) and the definition (2) of $I(D|T)$

$$I(D|\sigma, \mu, \theta) = \frac{N}{2} \lg(2\pi\sigma^2) + N \frac{\overline{u^2}}{2\sigma^2} \quad (\text{A2})$$

where

$$u = y - \mu - g(x; \theta) \quad (\text{A3})$$

and the overhead bar denotes the sample mean, as always.

In terms of $z := y - g(x; \theta)$, (A2) takes the form

$$\frac{N}{2} \lg(2\pi\sigma^2) + \frac{\overline{(z - \mu)^2}}{2\sigma^2}$$

For fixed θ , this is a function of σ and μ ; and we can approximate it by expanding about its minimum. As with equation (5) the minimum occurs for

$$\mu = \bar{z}, \quad \sigma = v^{1/2} \quad (\text{A4})$$

where

$$v := \overline{z^2} - \bar{z}^2$$

and the expansion analogous to (6) is

$$I(D|\sigma, \mu, \theta) \approx \frac{N}{2} \lg(2\pi ev) + N \frac{\Delta\sigma^2}{v} + N \frac{\Delta\mu^2}{2v} \quad (\text{A5})$$

where $\Delta\sigma = \sigma - v^{1/2}$, $\Delta\mu = \mu - \bar{z}$. If σ and μ in (A4) are stored with precisions $\delta\sigma$ and $\delta\mu$ then, as before, we can estimate $\Delta\sigma^2$ and $\Delta\mu^2$ in (A5) by $\delta\sigma^2/12$ and $\delta\mu^2/12$, respectively. Adding the result to

$$I(\sigma, \mu, \theta) = \lg \delta\sigma^{-1} + \lg \delta\mu^{-1} + I(\theta)$$

yields

$$I(D, T) \approx \frac{N}{2} \lg(2\pi ev) + \frac{N\delta\sigma^2}{12v} - \lg \delta\sigma + \frac{N\delta\mu^2}{24v} - \lg \delta\mu + I(\theta) \quad (A6)$$

Here we have assumed that σ and μ are stored independently of each other and of θ and have written “ $I(\theta)$ ” for the storage required by the θ_α . Minimizing (A6) with respect to $\delta\mu$ and $\delta\sigma$ gives

$$I(D, T) \approx \frac{N-2}{2} \lg(2\pi ev) + \lg \frac{N\pi e^2}{\sqrt{18}} + I(\theta) \quad (A7)$$

with the minimum occurring at $\delta\sigma = (6v/N)^{1/2}$, $\delta\mu = (12v/N)^{1/2}$, as in (9).

We now imagine that each of the θ_α is stored with the same relative precision, ϵ , i.e., that $\Delta\theta_\alpha = \eta_\alpha \theta_\alpha$ with $-\epsilon/2 < \eta_\alpha < \epsilon/2$. (Allowing ϵ to depend on α would not change anything.) Then

$$I(\theta) = -K \lg \epsilon$$

whose sum with (A7) is to be minimized.

To do so we must know how v , the sample variance of $y - g(x; \theta)$, depends on $\Delta\theta$. In order to have a general expression for this let us expand v about the value $\theta = \theta^*$ which minimizes it, and furthermore estimate—very roughly—that the deviations $\Delta\theta_\alpha$ give rise to corresponding relative deviations of $g(x; \theta)$ from $g(x; \theta^*)$. In other words we assume that, owing to the imprecision in θ_α , each $g(x; \theta)$ suffers an error of about

$$\pm \eta_\alpha V^{1/2} \quad (A8)$$

(Here $V^{1/2} = \overline{y^2} - \bar{y}^2$ [defined in equation (13)] is being taken as a typical scale for y , such scales for the θ 's (with respect to which it will be most efficient to store them) being derived from V and from typical scales for the components of x .) For each of the N data points (y, x) there are K such errors and assuming these NK errors to combine roughly independently [and independently of $y - g(x; \theta^*)$] we can expect

$$v \approx s^2 + K\epsilon^2 V/12 \quad (A9)$$

where $s^2 = v|_{\theta=\theta^*}$ is the minimum of v .

Under these assumptions and estimates (A7) becomes

$$I(D, T) \approx \frac{N-2}{2} \lg[2\pi e(s^2 + K\epsilon^2 V/12)] - K \lg \epsilon + \lg \frac{N\pi e^2}{(18)^{1/2}}$$

Minimizing this with respect to ϵ (and neglecting a small negative term which never exceeds -1) produces equation (15) of the text; and substituting the minimizing value, $\epsilon^2 = 12s^2/(N - K - 2)V$, into (A9) yields (14).

Remark 1. It is possible to carry out the minimization leading to (9) and the analogous minimization leading to (14) and (15) without expanding in $\Delta\sigma$; i.e., one minimizes $\langle I(D|T) \rangle + I(T)$, where the brackets denote exact expectations with respect to $\Delta\sigma$ and $\Delta\mu$ in the ranges $[-\delta\sigma/2, \delta\sigma/2]$, $[-\delta\mu/2, \delta\mu/2]$. This alters (9), (14), and (15) slightly but does not affect H , which in *this* sense can therefore be regarded as exact, even for small N .

Remark 2. Strictly speaking it would be better to replace V by $V - s^2$ in (A8) since only this part of the variance of y can be associated to variations in f . This would increase H in (19) by $K \lg R^{-1}$.

REFERENCES

- Fine, T. L. (1973). *Theories of Probabilities*. Academic Press, New York.
- Hanushek, E. A., and Jackson, J. E. (1977). *Statistical Methods for Social Scientists*. Academic Press, New York.
- Kolmogorov, A. (1968). Logical basis for information theory and probability theory, *IEEE Trans. Info. Theory*, **14**, 663.
- Kruskal, W., and Tanur, J. M. (eds.) (1978). *International Encyclopedia of Statistics*. The Free Press, New York. See the articles on hypothesis testing and significance tests.
- Kullback, S. (1959). *Information Theory and Statistics*. Wiley, New York.
- Penrose, O. (1970). *Foundation of Statistical Mechanics*. Pergamon Press, Oxford, p. 227.
- Schwarz, G. (1978). Estimating the dimension of a model, *Ann. Stat.*, **6**, 461-464.
- Yasuhara, A. (1971). *Recursive Function Theory and Logic*, Academic Press, New York.