

PREDICTIVE LIKELIHOOD

by

David Hinkley*

Technical Report No. 299
(Revision of TR No. 254)

*Preparation supported by National Science Foundation Grant
MCS77-00959.

PREDICTIVE LIKELIHOOD

by

David Hinkley
University of Minnesota

Summary

The likelihood function is the common basis of all parametric inference. However, with the exception of an ad hoc definition by R. A. Fisher, there has been no such unifying basis for prediction of future events, given past observations. This article proposes a definition of predictive likelihood which can help to remove some non-uniqueness problems in sampling-theory predictive inference, and which can produce a simple prediction analog of the Bayesian parametric result, posterior \propto prior \times likelihood, in many situations.

AMS 1970 subject classifications:

Key words and phrases. Prediction, confidence regions, exponential families, likelihood, Bayesian methods, statistical inference

¹The original work was supported by NSF grant MPS75-08778 and preparation of this account was supported by NSF grant MCS77-00959.

1. Introduction

In 1920 Karl Pearson [8] posed "The Fundamental Problem of Practical Statistics" as follows:

'An "event" has occurred p times out of $p + q = n$ trials, where we have no a priori knowledge of the frequency of the event in the total population of occurrences. What is the probability of its occurring r times in a further $r + s = m$ trials?'

Pearson's purpose was to reexamine the general applicability of Bayes's earlier solution, and the resulting controversy, described by Edwards [2], is of some interest. However, the main question seems to have been largely ignored in the intervening years, while parametric inference has dominated statistical thought.

In parametric inference, a fundamental concept is that of mathematical likelihood, on which both frequentist and Bayesian methods rest. A corresponding concept for prediction has been lacking. The present paper describes a definition of predictive likelihood developed independently by the author and Lauritzen [6], and shows how the definition relates to parametric likelihood and to Bayesian posterior predictive distributions. Most of the results are different from, and developed independently of, those in Lauritzen [6]. This paper is based on the author's Special Invited Lecture at the 1975 I.M.S. meeting in Atlanta, originally written up in [5].

Outline. The basic problem is discussed in Section 2. Section 3 describes the main definition of predictive likelihood, establishes natural consistency properties and demonstrates a correspondence with Bayes posterior predictive densities. The discussion is deliberately non-mathematical and

simple illustrative examples are given. Section 4 briefly deals with an extension of predictive likelihood that takes advantage of the existence of invariant ancillary statistics. Section 5 briefly review possible inferential uses of predictive likelihood, e.g. in defining unique frequentist prediction intervals.

Notation. Throughout the paper upper case letters denote random variables, lower case letters their realized values. The probability density, or mass, function (p.d.f.) of any random variable X at the value x is denoted by $f(x)$; conditional p.d.f.'s are denoted $f(x|\theta)$, $f(x|y)$, etc. Special notation is used for data and predictand random variables: Y is the observable random variables, and Z is the predictand random variables (to be predicted after observing Y). S and T correspond to functions (e.g. sufficient reductions) of Y and Z respectively, and R is always the minimal sufficient reduction of all variables (Y, Z) . Where Y and Z correspond to collections of individual X 's, m and n respectively denote the sample sizes.

2. The Problem

Suppose that X_1, \dots, X_{m+n} are random variables whose distribution is indexed by the unknown parameter θ . Our problem consists of being able to observe $Y = (X_1, \dots, X_m)$ and wanting then to make a predictive statement about $Z = (X_{m+1}, \dots, X_{m+n})$ based on the observation $Y = y$. The predictive statement is to be in the form of one or more confidence intervals, possibly centered on a point estimate.

From a Bayesian point of view our problem is straightforward. Once the prior distribution of θ is determined, and after $Y = y$ is observed, the posterior predictive distribution of Z given $Y = y$ may be calculated.

To be specific, if θ has prior p.d.f. $p(\theta)$, then

$$f(\theta|y) \propto p(\theta) \text{lik}(\theta|y) \quad (2.1)$$

and the posterior predictive p.d.f. of Z is

$$f(z|y) = \int f(\theta|y)f(z|y,\theta)d\theta = \frac{\int f(z|y,\theta)p(\theta)\text{lik}(\theta|y)d\theta}{\int p(\theta)\text{lik}(\theta|y)d\theta} \quad (2.2)$$

The standard form for a Bayesian 100% prediction confidence region would be the Highest Posterior Density region

$$R_{\beta}(y) = \{z: f(z|y) \geq k_{\beta}\} \quad (2.3)$$

where k_{β} is a function of y chosen so that

$$\int_{R_{\beta}(y)} f(z|y)dz = \beta .$$

The non-Bayesian does not find our problem straightforward (standard frequentist methods are outlined in Section 5). This is in some contrast to the situation where inference about θ is required, when the notion of likelihood provides a possible parallel between Bayesian and frequentist methods. Thus we have Highest Posterior Density Bayesian confidence regions for θ of the form

$$\{\theta: p(\theta)\text{lik}(\theta|y) \geq a\} ,$$

and likelihood-based frequentist confidence regions for θ of the form

$$\{\theta: \text{lik}(\theta|y) \geq b\} .$$

Both methods correspond to likelihood ordering of the θ -values when $p(\theta)$ is constant, although of course the confidence regions may still differ. There is no apparent parallel of (2.3) for the frequentist, because there is no apparent analog of $\text{lik}(\theta|y)$ for the predictand Z . Thus for the frequentist there is no statistical instrument for uniquely defining a prediction confi-

dence region, whereas likelihood provides such an instrument when dealing with θ .

By analogy with (2.1), a prediction likelihood function $\text{lik}^*(z|y)$ would satisfy

$$f(z|y) = a(y,z)\text{lik}^*(z|y), \quad (2.4)$$

where $a(\cdot, \cdot)$ is determined by the marginal prior distributions of Y and Z . The remainder of the paper is devoted to defining such a function $\text{lik}^*(z|y)$ and to describing some of its properties. Thus the discussion focusses on a basis for predictive inference, rather than on methods for such inference.

3. The Lauritzen - Hinkley Predictive Likelihood and Its Properties

R. A. Fisher [4, p. 134] suggested an ad hoc definition of predictive likelihood which used the notion of the degree to which values of y and z support the true hypothesis of a common θ -value for the two sets of variables. We use the same basic notion, but tie it more closely to parametric likelihood. Essentially we define the predictive likelihood of the value z to be the relative frequency of the observation $Y = y$ given the value of the minimal sufficient reduction of $(Y, Z) = (X_1, \dots, X_{m+n})$. According to this definition, what we have observed (y) becomes more likely as the predictive likelihood of the unknown (z) increases. As we shall see, the relationship to parametric likelihood is very close. Similar concepts have been used directly in the construction of frequentist confidence intervals by Faulkenberry [3] and Vit [10].

We start with a slight simplification of the problem introduced in Section 2. We suppose that $Y = (X_1, \dots, X_m)$ and $Z = (X_{m+1}, \dots, X_{m+n})$ are independent, but that individual components X_i need not be otherwise independent nor need they be identically distributed. The minimal sufficient reductions of Y , Z and (Y, Z) will be denoted by S , T and R respectively; R is clearly the

minimal sufficient reduction of (S,T) . To emphasize that the value of r is determined by particular values s and t we sometimes write $r(s,t)$.

Because of sufficiency, prediction of Z given $Y = y$ is statistically equivalent to prediction of T given $S = s$; the conditional distribution of Z given T is completely known. We then make

Definition 1. Let Y and Z be independent, with distributions indexed by the common unknown parameter θ . Let $R = r(S,T)$, S and T be the minimal sufficient reductions of (Y,Z) , Y and Z respectively. Then if t is uniquely defined by r and s , the predictive likelihood of $T = t$ given $S = s$ is

$$\text{lik}^*(t|s) = f(s|r) = f(s|r(s,t)), \quad (3.1)$$

and the predictive likelihood of $Z = z$ is

$$f(z|t) \text{lik}^*(t|s).$$

The predictive likelihood is independent of θ by sufficiency of R .

Notice that this definition is symmetric with respect to S and T , so that $\text{lik}^*(t|s) = f(t|r(s,t))$. With few exceptions (such as Example 2 below) a non-trivial result is obtained from (3.1) only in the case of exponential family random variables.

As a simple illustration of Definition 1, we use Pearson's Bernoulli problem.

Example 1. Let X_1, \dots, X_{m+n} be independent Bernoulli variables with $\text{pr}(X_j = 1|\theta) = \theta$. Then if $Y = (X_1, \dots, X_m)$ and $Z = (X_{m+1}, \dots, X_{m+n})$, we have

$$R = \sum_{j=1}^{m+n} X_j = S + T, \quad S = \sum_{j=1}^m X_j, \quad T = \sum_{j=m+1}^{m+n} X_j.$$

A simple evaluation of (3.1) shows that the predictive likelihood is hypergeometric,

$$\text{lik}^*(t|s) = \binom{m}{s} \binom{n}{t} / \binom{m+n}{s+t}.$$

For the special case $n = 1$, this result gives "relative likelihoods" corresponding to the Laplace Law of Succession, i.e.

$$\frac{\text{lik}^*(X_{m+1} = 1 | s)}{\text{lik}^*(X_{m+1} = 0 | s)} = \frac{s+1}{m-s+1} ;$$

this reflects a relationship between predictive likelihood and posterior predictive density established in Lemma 2 below.

There are two desirable consistency properties of predictive likelihood that are easily verified in Example 1 and that hold for general exponential families. In words, these properties are (i) predictive likelihood converges to the true density of T as the number m of observations increases: (ii) predictive likelihood converges to the parameter likelihood $\text{lik}(\theta | s)$ as the number n of variables to be predicted increases. More formally, we have

Lemma 1. Let $Y = (X_1, \dots, X_m)$ and $Z = (X_{m+1}, \dots, X_{m+n})$ be such that the X_j are independent and identically distributed with exponential family density

$$f(x | \theta) = \exp\{-\theta x + c(\theta) + d(x)\}.$$

Let $S = \sum_1^m X_j$, $T = \sum_{m+1}^{m+n} X_j$ and $R = S + T$. Denote the maximum likelihood estimators of θ based on S alone and T alone by $\hat{\theta}_S$ and $\hat{\theta}_T$ respectively.

Then, with $\text{lik}^*(t | s)$ defined by (3.1),

(i) as $m \rightarrow \infty$ with n fixed,

$$\text{lik}^*(t | s) = f(t | \hat{\theta}_S) + O_p(m^{-1}) \quad (3.2)$$

and

$$\text{lik}^*(t | s) = f(t | \theta) + O_p(m^{-\frac{1}{2}});$$

(ii) as $n \rightarrow \infty$ with m fixed,

$$\text{lik}^*(T | s) = f(x | \hat{\theta}_T) + O_p(n^{-1}) \quad (3.3)$$

and

$$\text{lik}^*(T | s) = f(s | \theta) + O_p(n^{-\frac{1}{2}}) .$$

Proof is outlined in the Appendix, where some extensions are suggested.

In (3.3) the parameter θ summarizes the infinite future (X_{m+1}, \dots) and

the parameter likelihood should be expected as a limiting result. It is conjectured that similar consistency results hold for the full generality of Definition 1. We might note that none of the above consistency properties hold for the definition given by Fisher [4, p. 134].

Up to this point we have required that the components of Y and Z be independent, but Definition 1 can be extended to cover the case of dependent sequences of random variables. The essential points of the definition are that R be minimal sufficient for (Y,Z) ; that S be a sufficient reduction of Y ; and that R be determined by S and T where T is a function of Z . The extended definition takes account of the facts that S may need to be larger than the minimal sufficient reduction of Y , and that the minimal sufficient reduction of Z must be determined by T and S (i.e. not necessarily by T alone).

Definition 2. Let R be minimal sufficient for (Y,Z) . Let S be sufficient for Y and let T be a function of (Z,S) such that (i) R is determined by (S,T) and (ii) the minimal sufficient reduction of Z is determined by (S,T) . In addition, we require the function $r(s,t)$ to have a unique inverse $t(r,s)$ for each value of s . Then the predictive likelihood of $T = t$ given $S = s$ is

$$\text{lik}^*(t|s) = f(s|r) = f(s|r(s,t)) \quad (3.4)$$

and the predictive likelihood of $Z = z$ is

$$f(z|s,t) \text{lik}^*(t|s) \quad (3.5)$$

The standard partition algorithm for determining R ([1, p. 24]) may be used to define minimal S and T .

The following two examples illustrate the differences between Definitions 1 and 2.

Example 2. Let X_1, \dots, X_m, X_{m+1} be i.i.d. with uniform density on the range $(0, \theta)$, so that X_{m+1} is to be predicted from X_1, \dots, X_m . We define $X_{(j,j)} = \max(X_1, \dots, X_j)$. Then the minimal sufficient statistic R is $X_{(m+1, m+1)}$, which is determined by

$$S = X_{(m,m)} \text{ and } T = \begin{cases} 0 & (X_{m+1} \leq s) \\ X_{m+1} & (X_{m+1} > s) \end{cases} .$$

These statistics satisfy the conditions of Definition 2, and a direct calculation of (3.4) gives

$$\text{lik}^*(t|s) = \begin{cases} m/(m+1) & (t \leq 0) \\ ms^m / \{(m+1)t\} & (t > 0) \end{cases} .$$

It follows that the full predictive likelihood (3.5) for $Z = X_{m+1}$ is

$$\text{lik}^*(z|s) = \begin{cases} \frac{m}{(m+1)s} & (z \leq s) \\ \frac{ms^m}{(m+1)z} & (z > s) \end{cases} .$$

Thus the likelihood is uniform over the observed range $(0, s)$.

If we had tried to apply Definition 1 with $T = X_{m+1}$, we could not have distinguished different values of x_{m+1} below s . That is, t would not be determined by s and r when $r = s$.

Example 3. Let $\{X_j: j = 0, \pm 1, \dots\}$ be a stationary first-order Markov binary sequence with

$$\text{pr}(X_{j+1} = b | X_j = a) = \theta_{ab} \quad (a, b = 0, 1) .$$

Suppose that $Y = (X_1, \dots, X_m)$ is to be observed and that $Z = (X_{m+1}, \dots, X_{m+n})$ is then to be predicted. For any time-connected sequence $(X_c, X_{c+1}, \dots, X_d)$ we define the matrix of transition frequencies $Q(c, d)$ by

$$Q_{ab}(c, d) = \# \text{ of } a \rightarrow b \text{ transitions in } (X_c, \dots, X_d) .$$

The minimal sufficient reduction of $(Y, Z) = (X_1, \dots, X_{m+n})$ is

$$R = (X_1, Q(1, m+n)) .$$

The minimal sufficient reduction of Y must be augmented by X_m , so that

$$S = (X_1, Q(1,m), X_m) ,$$

in order for R to be obtainable from S and Z . The necessary function of Z is

$$T = (X_{m+1}, Q(m+1,m+n)) ,$$

which is minimal sufficient for Z . Note that S is minimal totally sufficient in the sense of Lauritzen [7].

Calculation of (3.4), although complicated, follows from results of Whittle [9]. In particular, for the special case $n = 1$ we obtain

$$\text{lik}^*(x_{m+1}|s) = \frac{q_{x_m, x_{m+1}}^{(1,m) + 1}}{\sum_{j=0}^1 q_{x_m, j}^{(1,m) + 1}} ,$$

which corresponds to the result of Example 1 applied to row x_m of $Q(1,m)$.

The augmentation of the minimal sufficient reduction of Y is clearly necessary for all Markov processes.

In Section 2 we alluded to the fact that there was no apparent analog of the factorization (2.1) for the Bayes posterior predictive density (2.2) of Z given $Y = y$. It is very easy to prove that such a factorization is possible using Definition 2.

Lemma 2. The Bayes posterior predictive density of T given $S = s$ factorizes as

$$f(t|s) = \text{lik}^*(t|s) \frac{f(r(s,t))}{f(s)} , \quad (3.6)$$

where the last factor is the ratio of prior marginal densities of R and S .

Proof For S and T , the Bayes posterior predictive density (2.2) is

$$f(t|s) = \frac{\int f(s,t|\theta) p(\theta) d\theta}{\int f(s|\theta) p(\theta) d\theta} .$$

By the sufficiency of R , the right hand side is equal to

$$f(s,t|r) \frac{\int f(r|\theta) p(\theta) d\theta}{\int f(s|\theta) p(\theta) d\theta}$$

with $r = r(s,t)$. But, by Definition 2, $f(s,t|r(s,t)) = f(s|r(s,t)) = \text{lik}^*(t|s)$, which proves the result.

In principle, then, the only prior probabilities needed for prediction are those connected with samples of size m and $m+n$. This requires less than specification of the full prior density $p(\theta)$ in many cases; in exponential families the specification of $f(r)$ for all $m+n$ is equivalent to specification of $p(\theta)$.

It should be noted that the posterior predictive density is statistically equivalent to predictive likelihood only if the marginal prior density of R is constant for all admissible values of $r = r(s,t)$, rather than if the prior density $p(\theta)$ is constant. In particular, the equivalence will depend on the sampling model. For example, if in Example 1 $p(\theta) = 1$, then R has a uniform distribution. But with same $p(\theta)$, the sufficient statistic under inverse sampling does not have a uniform distribution. This reflects the fairly obvious fact that the predictive likelihood alone does not satisfy the (strong) likelihood principle ([1, p. 39]) as a basis for inference.

4. Conditional Predictive Likelihood

The predictive likelihood of Definition 2 gives non-trivial results only when sufficiency provides a genuine reduction of (Y,Z) . Thus, for example, the definition is essentially vacuous for most location-parameter models. However, in several common statistical models where sufficiency provides inadequate reduction, the minimal sufficient statistic can be expressed in terms of an ancillary statistic (with distribution independent of θ) and a conditionally sufficient statistic of low dimension. In such

cases a conditional predictive likelihood may be defined in terms of probability distributions conditional on the values of the ancillary structure, which parallels the usual conditional approach to parametric inference (see [4] and [1, p. 38]). Here we briefly discuss the conditional version of Definition 1.

Suppose then that Y and Z are independent, with R , S and T as in Definition 1. Let R , S and T be expressed respectively as (R^+, A) , (S^+, B) and (T^+, C) where A , B and C are ancillary. Because C is ancillary with known distribution, the prediction problem is now equivalent to prediction of T^+ given C and S . The essential change from Definition 1 is that the probability calculation for S given $R = r$ must now be made conditional also on the observed value of the ancillary B . We therefore have Definition 3. Under the above conditions

$$\text{lik}^*(t|s) = f(c) \text{lik}^*(t^+|s, c) = f(c) f(s^+|r, b). \quad (4.1)$$

To avoid possible difficulties with non-uniqueness of ancillary statistics, we should limit the definition to cases where A , B and C are maximal invariants of R , S and T with respect to the group structure of the probability model. With this limitation it is immediate that $\text{lik}^*(t^+|s, c)$ is determined solely by the maximal invariant function of (s^+, t^+) together with b and c . In the case of location parameter models for i.i.d. variables, the predictive likelihood (4.1) is equal to the conditional density of the maximal invariant function of (S^+, T^+) given $B = b$ and $C = c$, so that the conditional predictive likelihood is then equal to the fiducial density of T .

Definition 3 does not in general permit a factorization of the Bayes posterior predictive density as in Lemma 2, unless B is empty.

5. Relation of Predictive Likelihood to Frequentist Prediction Confidence Regions

For the problem formulated at the beginning of Section 2, an unbiased $1 - \beta$ prediction confidence region for Z is a set $P_\beta(Y)$ satisfying

$$\text{pr}\{Z \in P_\beta(Y) | \theta\} = 1 - \beta \quad (5.1)$$

uniformly in θ . There are two classical methods for determining $P_\beta(Y)$, to each of which the likelihood definitions of preceding sections relate.

The first method, attributed to Neyman, uses test critical regions as follows. We suppose that Y and Z have p.d.f.'s $f(y|\theta_1)$ and $f(z|\theta_2)$ respectively, and consider the hypothesis $H_0: \theta_1 = \theta_2$. If for a specified alternative H_A an unbiased (similar) test critical region Q_β of level β can be found, then

$$\text{pr}\{Y, Z \in Q_\beta; H_0\} \equiv \beta. \quad (5.2)$$

One region $P_\beta(y)$ satisfying (5.1) is, then, the projection of the complement Q_β^c onto the subspace $Y = y$. For the particular alternative H_A , the "natural" choice of Q_β is the uniformly or locally most powerful critical region, if such exists, and in many cases this leads to smallest confidence regions $P_\beta(y)$. However, the resulting confidence region is characterized by the alternative hypothesis H_A used in constructing Q_β . Thus the method as described does not uniquely define one system of regions $P_\beta(y)$, because H_A is not uniquely defined. This difficulty is removed by ordering the values of Z according to values of the predictive likelihood $\text{lik}^*(t|s)$, i.e. by requiring that

$$\inf_{z \in P_\beta(y)} \text{lik}^*(t|s) \geq \sup_{z \notin P_\beta(y)} \text{lik}^*(t|s). \quad (5.3)$$

This corresponds to the notion of likelihood-based confidence regions for θ (Cox and Hinkley [1, p. 218]).

Example 4. Suppose that X_1, \dots, X_m, X_{m+1} are i.i.d. $N(\mu, \tau)$ with both μ and τ completely unknown, and suppose that X_{m+1} is to be predicted after observing X_1, \dots, X_m . Here the possible artificial hypotheses H_0 and H_A are numerous. Definition 1 shows that $\text{lik}^*(x_{m+1} | \sum_{j=1}^m x_j, \sum_{j=1}^m x_j^2)$ is

proportional to a student-t density for

$$X_{m+1} - \frac{\sum_{j=1}^m x_j}{m} \quad ,$$

which forces the usual symmetric student-t interval for X_{m+1} by appeal to (5.3).

The second classical method for solving (5.1) is the pivotal method. Here a pivotal or invariant function $h(s, t)$ is used to determine a region w_β such that

$$\text{pr}\{h(S, T) \in w_\beta | \theta\} = 1 - \beta \quad (5.4)$$

for all θ , and then the prediction confidence region is

$$P_\beta(y) = \{z: h(s, t) \in w_\beta\} \quad .$$

The predictive likelihood may again be used to uniquely define $P_\beta(y)$ via (5.3). Notice that in problems with group structure, to which the pivotal method is restricted, Definition 3 provides an immediate solution: $h(S, T)$ is the maximal invariant function of (S^+, T^+) , whose conditional distribution determines $P_\beta(y)$. Example 4 is a trivial illustration of this.

Appendix: Proof of Lemma 1

The following is an outline of the proof for Lemma 1, part (i). The essential requirements are that a Central Limit Theorem expansion hold for S and R with standardizing constant $m^{\frac{1}{2}}$, and that the maximum likelihood estimator $\hat{\theta}_S$ given $S = s$ satisfy

$$s = E(S|\hat{\theta}_S) . \tag{A.1}$$

The latter is a standard property of linear exponential families. For the case stated in Lemma 1, if the X_j are i.i.d. with p.d.f.

$$f(x|\theta) = \exp\{-\theta x + c(\theta) + d(x)\} ,$$

then for $S = \sum_{j=1}^m X_j$ we have

$$s + mc'(\hat{\theta}_S) = 0 , \quad E(S|\theta) = -mc'(\theta) .$$

To obtain the first result in Lemma 1, we write

$$\text{lik}^*(t|s) = f(s|s+t) = \frac{f(s|\theta)f(t|\theta)}{f(s+t|\theta)} \quad \text{for all } \theta .$$

Then, choosing $\theta = \hat{\theta}_S$, we have

$$\text{lik}^*(t|s) = f(t|\hat{\theta}_S) \frac{f(s|\hat{\theta}_S)}{f(s+t|\hat{\theta}_S)} . \tag{A.2}$$

It remains to show that the ratio in (A.2) is $1 + O_p(m^{-1})$. To see this we apply a Central Limit Theorem expansion, noting that (A.1) holds and that consequently

$$s + t = E(R|\hat{\theta}_S) + O(1) . \tag{A.3}$$

For example, in the i.i.d. case with the assumption of fourth moments for X , the usual Edgeworth expansion may be applied.

The second result in part (i) of the Lemma is a consequence of an expansion of the first result using $\hat{\theta}_S - \theta = O_p(m^{-\frac{1}{2}})$. Note that in order

to show weak convergence to $f(t|\theta)$ we need only assume that a Central Limit Theorem applies to S and R so that by (A.1) and (A.3) the ratio term in (A.2) is $1 + o_p(1)$. Thus simple consistency of $\text{lik}^*(t|s)$ will hold quite generally if (A.1) is satisfied.

Part (ii) of the Lemma follows from part (i) by reversing the roles of S and T .

Acknowledgements. I thank several readers of earlier versions of this paper for their help in clarifying the presentation. Particular thanks go to Steffen Lauritzen.

REFERENCES

- [1] Cox, D. R. and Hinkley, D. V. (1974). Theoretical Statistics. London: Chapman-Hall.
- [2] Edwards, A. W. F. (1974). A problem in the doctrine of chances. Proc. Conference on Foundational Questions in Statistical Inference, eds. Barndorff-Nielsen, O. et al. University of Aarhus.
- [3] Faulkenberry, G. D. (1973). A method of obtaining prediction intervals. J. Amer. Statist. Assoc. 68, 433-5.
- [4] Fisher, R. A. (1973). Statistical Methods and Scientific Inference. (3rd ed.) New York: Hafner.
- [5] Hinkley, D. V. (1976). On predictive inference. University of Minnesota School of Statistics Technical Report No. 254.
- [6] Lauritzen, S. L. (1974). Sufficiency, prediction and extreme models. Scand. J. Statist. 1, 128-134.
- [7] Lauritzen, S. L. (1975). General exponential models for discrete observations. Scand. J. Statist. 2, 23-33.
- [8] Pearson, K. (1920). The fundamental problem of practical statistics. Biometrika 13, 1-16.
- [9] Whittle, P. (1955). Some distribution and moment formulae for the Markov chain. J. Roy. Statist. Soc. B 17, 235-242.
- [10] Vit, P. (1973). Interval prediction for a Poisson process. Biometrika 60, 667-8.