On the Consistency of the Fiducial Method

By D. A. S. Fraser

University of Toronto

[Received July 1961. Revised January 1962]

Summary

Lindley (1958) and Sprott (1960) have proposed several criteria for the satisfactory application of the fiducial method of inference to statistical problems. The criteria are concerned with consistency when there are various ways of applying the method. In sections 1 and 2 these criteria are discussed briefly, and the range of problems that yield fulfilment of the criteria is shown to be considerably larger than that indicated by Lindley and Sprott. A frequency interpretation for a fiducial distribution has been proposed and discussed in two recent papers (Fraser, 1961a, b); this interpretation is available for problems that are invariant in a natural sense under translation and transformation groups. If the purpose of the fiducial method of inference is to produce posterior probability statements about the parameter that have a frequency interpretation and are consistent according to criteria such as Lindley's and Sprott's, then the range of validity can be substantially extended by making a small modification to the method; this is discussed in sections 3 to 6.

1. Introduction

The fiducial method of inference was introduced by R. A. Fisher in 1930. It has been discussed and developed in many of his papers since then. The purpose of the method is to obtain from data valid probability statements concerning the parameter value. The conflict with confidence theory indicates that this purpose is not achieved in general using the specialized interpretation of a probability as a relative frequency in a repetitive situation and that a broader and perhaps more scientific interpretation as a proportion in a conceptual reference set is needed (Fisher, 1956, p. 109). An attempt in this direction for a certain kind of model uses transformations to generate the reference set (Fraser, 1961a, b).

There seem to be two aspects to analysing whether the fiducial method achieves its purpose. One aspect is to investigate whether the probability statements can be interpreted and correlated with the physical situation in a satisfactory manner. Another is to investigate whether the probability statements yield consistent results when manipulated by rules conventional to probability theory. In section 2 of this paper consistency for real parameters is shown to be limited to translation-parameter models. Interestingly, it is for such models that the frequency interpretation or correlation with the physical situation is available (Fraser, 1961a, b). Thus the two aspects are touched upon by investigating the second.

The investigation of consistency will be in terms of criteria proposed by Lindley (1958) and Sprott (1960). Lindley considers two samples coming from the same distribution and bases his criterion on a statement by Fisher (1956) that fiducial probability "is entirely identical with the classical probability of the early writers, such as Bayes". Lindley's criterion can be formalized as follows:

Criterion L. The fiducial distribution from a first sample used as a prior distribution for a Bayesian analysis of a second sample should yield the fiducial distribution as obtained directly from the combined sample.

In his investigation of the range of problems that yield fulfilment to this criterion, Lindley limits himself to Koopman-Darmois models, which have a sufficient statistic for any size sample. In so limiting his investigation, Lindley neglects the large class of problems that have instead a conditionally sufficient statistic given an ancillary statistic. For the Koopman-Darmois models, Lindley shows that consistency is limited to variables that in an appropriate mode of expression have a normal distribution with a location parameter or have a gamma distribution with a scale parameter.

Sprott (1960) considers some conditions to be imposed on distributions *a posteriori*, and then uses two specializations of these for an investigation of the fiducial method. The first of these is a weaker version of Criterion L.

Criterion S_1. The fiducial distribution from one sample used as a prior distribution for a Bayesian analysis of another sample should yield a result independent of the order of combination.

Sprott also limits his investigation to Koopman-Darmois models, and for these models finds that consistency according to criterion S_1 is limited to the same normal and gamma variables obtained by Lindley.

Sprott also considers a criterion that is applicable with two distribution forms having the same parameter.

Criterion S_2. For two distribution forms each involving the parameter θ, the fiducial distribution as obtained from a sample from one form and used as a prior distribution for a Bayesian analysis of a sample from the other form should yield a result independent of the order of combination.

Again for Koopman-Darmois models consistency is limited to the normal and gamma variables.

As part of his investigation of the fiducial method, Lindley (1958) determines the conditions under which the fiducial distribution for a real parameter is identical to the Bayes posterior distribution as obtained from some prior distribution. Lindley's argument can be put in a simple form and is needed in the next section for the investigation of consistency. The argument develops from the equation

$$\left|\frac{\partial F}{\partial \theta}\right| = n(\theta)\left|\frac{\partial F}{\partial x}\right| r(x). \tag{1}$$

The left side is the fiducial density function as obtained from a variable x having distribution function $F(x|\,\theta)$; the right side is the product of a prior density function $n(\theta)$, the likelihood function for θ obtained from x, and a normalizing constant $r(x)$. The problem is to determine the kind of function $F(x|\,\theta)$ that will satisfy this equation. Rearrangement yields

$$n(\theta)\,r(x) = \left|\frac{\partial F/\partial \theta}{\partial F/\partial x}\right| = \left|\frac{\partial x}{\partial \theta}\right| \tag{2}$$

and, with F fixed in value,

$$n(\theta)\,d\theta = \frac{dx}{r(x)}.$$

Integration then produces

$$\Theta(\theta) = X(x) + K,$$

where $\Theta(\theta)$, $X(x)$ are the indefinite integrals and K is the constant of integration; the value of K depends on the value of F. Thus the quantity $X(x) - \Theta(\theta)$ has a fixed distribution and the parameter is essentially a location parameter. The condition is thus that $F(x \mid \theta)$ be of translation-parameter form.

2. The Range of Consistency

For their investigation of consistency Lindley and Sprott restrict the underlying distribution to being of standardized Koopman-Darmois form

$$f(x \mid \theta) = r(\theta)\, s(x)\, e^{-x\theta}. \tag{3}$$

In this form the variable and the parameter have been coded or put in a natural mode of expression with the simple product $x\theta$ as argument of the negative exponential function; the more general uncoded form introduces nothing essentially different. By simple manipulation equation (3) can be exhibited as

$$f(x \mid \theta) = r'(\theta)\, s'(x)\, N(x \mid \theta), \tag{4}$$

where $N(x \mid \theta)$ designates a normal density function with θ as a location parameter. If $f(x \mid \theta)$ attaches probability only to the positive axis, then equation (3) can also be written

$$f(x \mid \theta) = r''(\theta)\, s''(x)\, \Gamma(x \mid \theta), \tag{5}$$

where $\Gamma(x \mid \theta)$ designates a gamma density function with θ as a scale parameter, or equivalently with $\ln \theta$ as a location parameter. Both $N(x \mid \theta)$ and $\Gamma(x \mid \theta)$ are translation-parameter models.

A simple sampling interpretation can be given to density functions of the form (4) and (5). Consider a density function

$$f(x \mid \theta) = R(\theta)\, S(x)\, T(x \mid \theta), \tag{6}$$

where $T(x \mid \theta)$ is also a density function. Observations from this distribution can be obtained by the following procedure. Take observations from the distribution with density function $T(x \mid \theta)$; subject each observation to a random selection operation so that an observation is retained with retention probability $KS(x)$ and is rejected with rejection probability $1 - KS(x)$. If $S(x)$ is bounded, K can be chosen so that $KS(x) \leqslant 1$; if $S(x)$ is unbounded, K can be chosen so that $KS(x) < 1$ except on a region having arbitrarily small probability according to $T(x \mid \theta)$. The observations that survive the selection operation will have a probability density function $f(x \mid \theta)$ given by (6). We can thus say that $f(x \mid \theta)$ is obtained by selection from $T(x \mid \theta)$ or that $f(x \mid \theta)$ is a selected form of $T(x \mid \theta)$.

The two alternative forms (4) and (5) for the Koopman-Darmois model can be described as selected normal or selected gamma models respectively; as such they are two special kinds of selected translation-parameter models. It is natural then to enquire whether the Koopman-Darmois model can be exhibited in the form

$$f(x \mid \theta) = R(\theta)\, S(x)\, T(x \mid \theta), \tag{7}$$

where $T(x \mid \theta)$ is a translation-parameter model other than the normal and the gamma. Equations (7) and (3) yield

$$T(x \mid \theta) = \frac{r(\theta)}{R(\theta)} \frac{s(x)}{S(x)}\, e^{-x\theta}. \tag{8}$$

The analysis of Lindley (1958, section 3) shows that the only translation-parameter solutions are the normal $N(x|\theta)$ and the gamma $\Gamma(x|\theta)$. Thus, the restriction to a Koopman-Darmois model is essentially a restriction to models obtained by selection from normal and gamma density functions.

Lindley and Sprott find for Koopman-Darmois models that consistency occurs only for the normal and gamma translation-parameter models, and hence that the effect of the consistency criteria is to restrict to the cases in which there is effectively no selection. Consistency, however, is obtained generally for translation-parameter models (Fraser, 1961b). The remainder of this section shows that consistency is limited to translation-parameter models and that the consistency criteria in general suppress selection on such models.

Consider now the delimitation of the range of consistency according to criteria L, S_1 and S_2. Let x be a real variable with distribution function $F(x|\theta)$ and let y be an independent variable with distribution function $G(y|\theta)$; F and G may be identical or different. The fiducial density for θ obtained from the first variable is

$$f_f(\theta|x) = |\partial F(x|\theta)/\partial\theta|.$$

A Bayesian analysis on the second variable with this as prior density yields the posterior relative frequency function

$$\left|\frac{\partial F(x|\theta)}{\partial\theta}\right| \left|\frac{\partial G(y|\theta)}{\partial y}\right| \tag{9}$$

for θ. The analogous posterior relative frequency function obtained by starting with y is

$$\left|\frac{\partial G(y|\theta)}{\partial\theta}\right| \left|\frac{\partial F(x|\theta)}{\partial x}\right|. \tag{10}$$

According to criterion L or criterion S_1 (F and G identical), or according to criterion S_2 (F, G different), these posterior distributions must be the same. Thus consistency implies that expressions (9) and (10) as functions of θ must be proportional; hence

$$\frac{\left|\dfrac{\partial F(x|\theta)}{\partial x}\right|}{\left|\dfrac{\partial F(x|\theta)}{\partial\theta}\right|} = \frac{\left|\dfrac{\partial G(y|\theta)}{\partial y}\right|}{\left|\dfrac{\partial G(y|\theta)}{\partial\theta}\right|} h(x,y), \tag{11}$$

where $h(x,y)$ is the constant of proportionality. If y is given any fixed value then equation (11) takes the form (2) for F; if x is given any fixed value, then equation (11) takes the form (2) for G. Thus F and G must be of translation-parameter form.

3. A Modification to the Fiducial Method

For a real parameter the fiducial method is consistent precisely for translation-parameter models; and for such models it has a frequency interpretation. If the purpose of the method is to obtain posterior distributions with such properties then the range of validity can be substantially extended by a small modification to the method.

Consider first a real variable x with a translation-parameter distribution located by θ. Let $p(g)$ be the probability density function of the pivotal variable $g = x - \theta$; then

$$f(x \mid \theta) = p(x-\theta),$$
$$f_f(\theta \mid x) = p(x-\theta).$$

The fiducial distribution $f_f(\theta \mid x)$ has a frequency interpretation by means of transformations (Fraser, 1961a, b) and in terms of a general reference set (Fisher, 1956, p. 109; Fraser, 1961a, b). Bayesian manipulations can be performed in a consistent manner for any combination of such distributions provided that different distribution forms use the same translation groups on the parameter space.

Consider now a selected translation-parameter specification,

$$f^*(x \mid \theta) = r(\theta)\, s(x)\, p(x-\theta), \tag{12}$$

and suppose the translation group is in some way natural to the variable x. A three-step probability model can be used to give structure to this specification.

First, for a given θ let A be an event that can occur with probability $Jr(\theta)$ and not occur with probability $1 - Jr(\theta)$; the constant J is chosen so that $Jr(\theta) \leqslant 1$ for some suitably wide range of θ values.

Second, for a given θ let x be an observation from the translation-invariant distribution with density $p(x-\theta)$.

Third, for a given x let R, designating retention of that x, be an event that can occur with probability $Ks(x)$ and not occur, rejection, with probability $1 - Ks(x)$.

If A occurs in the first step, choose an x according to the second step; then subject the x to selection according to the third step. The result of this three-step procedure is an x if both A and R occur, and is nothing otherwise. The first probability mechanism is dependent on θ only; the second is translation invariant and has an "error" deviation $g = x - \theta$ that has a fixed distribution; and the third is dependent on x only.

The probability of obtaining an x by this procedure is

$$\begin{aligned} P(AR \mid \theta) &= \int_{-\infty}^{\infty} Jr(\theta)\, p(x-\theta)\, Ks(x)\, dx \\ &= JK \int_{-\infty}^{\infty} JK f^*(x \mid \theta)\, dx = JK \end{aligned}$$

and does not depend on θ. The conditional distribution of x given the occurrence of A and R is

$$\begin{aligned} P(x \mid A, R, \theta) &= \frac{Jr(\theta)\, p(x-\theta)\, Ks(x)\, dx}{JK} \\ &= f^*(x \mid \theta)\, dx. \end{aligned} \tag{13}$$

If the procedure is performed repeatedly until an observation is obtained then the distribution of that observation is

$$P(x \mid \theta) = f^*(x \mid \theta)\, dx.$$

The repeated performance of the procedure thus generates observations on the selected translation specification (12).

If the translation group is accepted as being meaningful for the variable x, and if this probability model is accepted as a satisfactory interpretation or description of the specification (12), then an inversion can be formed to obtain a posterior distribution for θ.

The distribution relating x to θ in the second step is translation-invariant. For a given x, the distribution of θ under $A \cup \tilde{A}$ is available from the preliminary discussion at the beginning of this section and has probability element

$$p(x-\theta)\, d\theta.$$

The joint probability of this θ and the event A is

$$Jr(\theta)\, p(x-\theta)\, d\theta. \tag{14}$$

The conditional distribution given A and, of course, x is obtained from the normalized version of (14). Thus, the posterior distribution of θ given x has the following relative frequency function

$$r(\theta)\, p(x-\theta)$$

and is equivalent to the likelihood function. The use of the three-step probability model and the results on translation-invariant models has produced a distribution for θ given x that has a frequency-interpretation. In the next section the consistency of this procedure for obtaining posterior distributions will be investigated.

This real-variable example suggests that the following modification be made to the fiducial method. Consider any specification that describes selection applied to a transformation-parameter model. The natural pivotal quantity for the transformation-parameter model can be used to obtain a preliminary fiducial distribution for the underlying model; the posterior distribution for the selected model is obtained by weighting or modulating the preliminary fiducial density function by the θ-function in the original density function.

To describe the modification in symbols, suppose that x, θ, g take values in a group and that the pivotal variable, $g = \theta^{-1}x$, has probability density function $p(g)$ with respect to left Haar measure μ on the group. The probability element for x given θ is

$$p(\theta^{-1}x)\, d\mu(x) \tag{15}$$

and the frequency element for θ given x is

$$p(\theta^{-1}x)\, \Delta(\theta)\, d\mu(\theta), \tag{16}$$

where $\Delta(\theta)$ is the modular function relating right Haar measure to left Haar measure.

Consider now a variable x with the probability density element

$$f(\theta \mid x)\, d\mu(x) = r(\theta)\, s(x)\, p(\theta^{-1}x)\, d\mu(x). \tag{17}$$

This can be viewed as the result of a selection function $s(x)$ operating on the transformation-parameter model just described.

The posterior distribution for θ before selection can be obtained from the pivotal quantity $g = \theta^{-1}x$ and is given by (16). The posterior frequency element for θ given, in addition, the analogue of the event A in the example is

$$f_F(\theta \mid x)\, dx = c(x)\, r(\theta)\, p(\theta^{-1}x)\, \Delta(\theta)\, d\mu(\theta), \tag{18}$$

where $c(x)$ is the normalizing constant. This posterior distribution has a frequency interpretation as in the example and will be referred to as the modified fiducial distribution for θ given x for the model (17). Note that the modified fiducial density function is just the likelihood function $r(\theta)\,p(\theta^{-1}x)$ modulated by the modular function $\Delta(\theta)$, and is thus equivalent to the likelihood function in the amount of information concerning x that is retained in the function form. This equivalence is not true in general for ordinary fiducial distributions and has been viewed as an unpleasant feature of such distributions; see an example of this in the next section.

4. Lindley's Example

Lindley (1958) discusses a simple example that illustrates the breakdown of the consistency condition L. Let x have the following probability element

$$f(x\,|\,\theta)\,dx = \frac{\theta^2}{\theta+1}(x+1)\,e^{-x\theta}\,dx$$

$$= \frac{\theta}{\theta+1}(x+1)\,x\theta\,e^{-x\theta}\,d\ln x$$

for $x>0, \theta>0$. The second expression uses the left Haar measure element $d\ln x$ for scale changes for the variable x. The ordinary fiducial element for θ given x is recorded by Lindley as

$$f_f(\theta\,|\,x)\,d\theta = \frac{\theta}{(\theta+1)^2}\{(2+x)+\theta(1+x)\}\,x\theta\,e^{-x\theta}\,d\ln\theta.$$

Consider now two observations x, y from this distribution. The marginal element for the sufficient statistic $z = x+y$ is

$$h(z\,|\,\theta)\,dz = \frac{\theta^4}{(\theta+1)^2}(z+z^2+\tfrac{1}{6}z^3)\,e^{-z\theta}\,dz$$

$$= \frac{\theta^2}{(\theta+1)^2}(1+z+\tfrac{1}{6}z^2)\,(z\theta)^2\,e^{-z\theta}\,d\ln z.$$

The ordinary fiducial element for θ given z is recorded by Lindley:

$$h_f(\theta\,|\,z)\,d\theta = \frac{\theta^2}{(\theta+1)^3}\{(2+\tfrac{4}{3}z+\tfrac{1}{6}z^2)+\theta(1+z+\tfrac{1}{6}z^2)\}\,(z\theta)^2\,e^{-z\theta}\,d\ln\theta.$$

Lindley compares this overall fiducial distribution with the following obtained by using the fiducial distribution from x for a Bayesian analysis on y; this latter has frequency element

$$f_f(\theta\,|\,x)f(y\,|\,\theta)\,d\theta = \frac{\theta^2}{(\theta+1)^3}(y+1)\{(2+x)+\theta(1+x)\}\,(x\theta)\,(y\theta)\,e^{-(x+y)\theta}\,d\ln\theta.$$

These fiducial distributions are different and neither has a frequency function that can be expressed as a function of the likelihood function.

The modified fiducial frequency function for θ given z is simply obtained from formula (18) and is

$$h_F(\theta\,|\,z)\,d\ln\theta = \frac{\theta^2}{(\theta+1)^2}\{1+z+\tfrac{1}{6}z^2\}\,(z\theta)^2\,e^{-z\theta}\,d\ln\theta,$$

and is in fact equivalent to the likelihood function since the modular function of the group is identically one.

The same frequency element is obtained if the modified fiducial distribution from x or y is used for a Bayesian analysis on the other variable, since this amounts to putting two likelihood functions together. The modified fiducial distribution given by $h_F(\theta|z)$ has a frequency interpretation and is consistent according to the Lindley-Sprott criteria.

5. Consistency of the Modified Method

Consider n independent variables $x_1, \ldots, x_n$ and suppose that each variable has a selected transformation-parameter model of the kind discussed in section 3 and that the same parameter θ is involved in each model. In a selected transformation-parameter model there is a transformation group operating on the parameter space. In this section it will be assumed that each of the models has the same transformation-group on the parameter space and that a product such as $\theta^{-1} x_i$ indicates multiplication according to this group. The case involving differing groups will be considered in the next section.

Let the probability element for the variable x_i be

$$r_i(\theta)\, s_i(x_i)\, p_i(\theta^{-1} x_i)\, d\mu(x_i). \tag{19}$$

The assumption just made does not mean that the transformation groups as they apply to some given modes of expression of the different x_i need be the same, but it does mean that the individual variables after the coding operation that represents x_i as a group element must have the same transformation group. For example, Sprott's (1960, p. 317, case 3) transformation groups on the variables T_1 and T_2 are different, but become the same when the variables are coded in the natural form for the transformation-parameter model.

The joint probability element for $(x_1, \ldots, x_n)$ is

$$\prod_{i=1}^{n} r_i(\theta) \prod_{i=1}^{n} s_i(x_i) \prod_{i=1}^{n} p_i(\theta^{-1} x_i) \prod_{i=1}^{n} d\mu(x_i). \tag{20}$$

The modified fiducial distribution for θ obtained from x_j has frequency element

$$r_j(\theta)\, s_j(x_j)\, p_j(\theta^{-1} x_j)\, \Delta(\theta)\, d\mu(\theta). \tag{21}$$

Using this as a prior distribution for a Bayesian argument on $(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)$ produces the posterior frequency element for θ

$$\prod_{i=1}^{n} r_i(\theta) \prod_{i=1}^{n} s_i(x_i) \prod_{i=1}^{n} p_i(\theta^{-1} x_i)\, \Delta(\theta)\, d\mu(\theta); \tag{22}$$

this is obtained by multiplying (21) by the likelihood function corresponding to the additional variables. Note that this posterior distribution (22) does not depend on j and hence satisfies Sprott's condition S_2.

If a sufficient statistic exists for the underlying transformation models with element

$$\prod_{i=1}^{n} p_i(\theta^{-1} x_i) \prod_{i=1}^{n} d\mu(x_i), \tag{23}$$

then there is a mode of expression for it, say $t(x_1, \ldots, x_n)$ that has transformation properties; this follows from the work of Fraser (1961b, section 11). Let the marginal probability element for t be

$$P(\theta^{-1} t)\, d\mu(t),$$

as derived from the transformation model (23). The same reference then shows that the joint probability element (20) for $(x_1, \ldots, x_n)$ can be written as

$$\prod_{i=1}^{n} r_i(\theta) \prod_{i=1}^{n} s_i(x_1) P(\theta^{-1} t)\, w(x_1, \ldots, x_n) \prod_{i=1}^{n} d\mu(x_i), \tag{24}$$

where $w(x_1, \ldots, x_n)$ is an invariant function. The modified fiducial distribution for θ obtained from t has frequency element

$$\prod_{i=1}^{n} r_i(\theta) P(\theta^{-1} t)\, \Delta(\theta)\, d\mu(\theta), \tag{25}$$

which is seen to be equivalent to (22) by noting the relationship between (20) and (24). Lindley's condition L is therefore fulfilled for this model.

A sufficient statistic need not exist for the parameter θ. The underlying model has, however, an ancillary statistic and a conditionally sufficient statistic (Fraser, 1961b, section 11). The application of selection to this underlying model may be such that this ancillary statistic is not an ancillary statistic for the selected model. In such a case the fiducial approach by means of the ancillary method seems blocked: for the selected model the ancillary statistic is "informative" for θ. If, however, the selection is such that the ancillary statistic remains ancillary, then the conditional probability element of, say, x_1, given the ancillary statistic, has the form

$$\prod_{i=1}^{n} r_i(\theta) \prod_{i=1}^{n} p_i(\theta^{-1} x_1 . x_1^{-1} x_i)\, k(x_1;\, x_1^{-1} x_2, \ldots, x_1^{-1} x_n)\, d\mu(x_1);$$

in this, the ancillary statistic is represented by $(x_1^{-1} x_2, \ldots, x_1^{-1} x_n)$, and the normalizing constant

$$\prod_{i=1}^{n} r_i(\theta)$$

is the same as in the joint probability element, a consequence of the ancillary statistic remaining ancillary for the selected model. The modified fiducial method applied to this conditional distribution again yields (22); condition L is again fulfilled.

6. With Differing Transformation Groups on the Parameter Space

Consider n independent variables $x_1, \ldots, x_n$ and suppose that each of the first k variables has a selected transformation-parameter model. Let the probability element for x_i with $i \leqslant k$ be

$$r_i(\theta_i)\, s_i(x_i)\, p_i(\theta_1^{-1} x_i)\, d\mu_i(x_i),$$

where μ_i is the left Haar measure for the transformation group for the ith variable, and let the probability element for x_i with $i > k$ be

$$f_i(x_i \mid \theta)\, d\mu_i(x_i),$$

where μ_i is a measure on the space of the variable x_i. Suppose also that there is a single parameter θ that determines each of these distributions. The subscript is added to the θ for the first k distributions to indicate that differing group multiplications may be present for combination with the differing variables $x_1, \ldots, x_n$.

A posterior distribution for θ can be obtained by using the modified fiducial distribution for θ as obtained from, say, x_j for a Bayesian analysis on the remaining variables $(j \leqslant k)$; it has the frequency element

$$\prod_{i=1}^{k} r_i(\theta_i) \prod_{i=1}^{k} s_i(x_i) \prod_{i=1}^{k} p_i(\theta_i^{-1} x_i) \prod_{i=k+1}^{k} f_i(x_i \mid \theta)\, \Delta_j(\theta_j)\, d\mu_j(\theta_j).$$

Thus in general a different posterior distribution will be obtained for each choice of $j = 1, \ldots, k$, and Sprott's condition S_1 will not be fulfilled $(k > 1)$. If $k = 1$ and none of the distributions $f_i(x_i \mid \theta)$ is of selected transformation form, then seemingly there is a unique posterior distribution for θ. If $k > 1$, there is a variety of posterior distributions each with a certain kind of frequency interpretation, and a claim to validity in reference to a particular initiating observation. It seems reasonable then to contemplate the convex hull of such distributions. This would yield a frequency distribution for the percentage points of the posterior distribution of θ and perhaps relate to the levels of uncertainty considered by Fisher (1957).

ACKNOWLEDGEMENTS

This material was developed during a research appointment at Bell Telephone Laboratories, Murray Hill, N.J. I gratefully express my appreciation to Bell Telephone Laboratories for their support of this research, and to Professor John Tukey who made available to me extensive notes on fiducial inference. The material was then revised and extended at Stanford University under a National Science Foundation Grant.

REFERENCES

FISHER, R. A. (1930), "Inverse probability", *Proc. Camb. phil. Soc.*, **26**, 528–535.
—— (1956), *Statistical Methods and Scientific Inference*. Edinburgh: Oliver and Boyd.
—— (1957), "The underworld of probability", *Sankhyā*, **18**, 201–210.
FRASER, D. A. S. (1961a), "On fiducial inference", *Ann. math. Statist.*, **32**, 661–676.
—— (1961b), "The fiducial method and invariance", *Biometrika*, **48**, 261–280.
LINDLEY, D. V. (1958), "Fiducial distributions and Bayes' theorem", *J. R. statist. Soc.* B, **20**, 102–107.
SPROTT, D. A. (1960), "Necessary restrictions for distributions *a posteriori*", *J. R. statist. Soc.* B, **22**, 312–318.