Applications of Basu's Theorem

Dennis D. Boos and Jacqueline M. Hughes-Oliver *

January 1997

Abstract

This brief note suggests ways to motivate Basu's Theorem, reinforce some basic ideas of classical inference, and show applications to various statistical problems.

KEY WORDS: Ancillary statistic; Complete family; Pivotal quantity; Sufficient statistic.

Institute of Statistics Mimeo Series No. 2290

Department of Statistics, North Carolina State University

*Dennis D. Boos is Professor and Jacqueline M. Hughes-Oliver is Assistant Professor, Department of Statistics, North Carolina State University, Raleigh, NC 27695–8203. Email addresses are: boos@stat.ncsu.edu, hughesol@stat.ncsu.edu

1. INTRODUCTION

Basu's Theorem is one of the most elegant results of classical statistics. Succinctly put, the Theorem says: if $\boldsymbol{T}$ is a complete sufficient statistic for a family of probability measures $\mathcal{P}$, and $\boldsymbol{V}$ is an ancillary statistic, then $\boldsymbol{T}$ and $\boldsymbol{V}$ are independent. (See, for example, Lehmann, 1983, p. 46, Casella and Berger, 1990, p. 262, or the original source, Basu, 1955, Theorem 2.) We often think of $\mathcal{P}$ as a parametric family $\mathcal{P} = \{P_{\boldsymbol{\theta}}, \boldsymbol{\theta} \in \Omega\}$, where $\boldsymbol{\theta}$ has dimension q, but it is important to note that the theorem holds for the more general families $\mathcal{P}$ that arise from nonparametric representations.

In Section 2 we review some basic definitions and ideas from classical inference that motivate Basu's Theorem. In Section 3 we give some useful technical results which may be obtained from Basu's Theorem. These results include applications to parametric and nonparametric families and to single and multiple samples.

2. STATISTICAL INFERENCE

The goal of statistical inference is to make statements about the underlying distribution that generated a given observed sample. Suppose the n-dimensional random vector $\boldsymbol{X} = (X_1, \ldots, X_n)$ represents the sample. We assume that $\boldsymbol{X}$ has a distribution from some family $\mathcal{P}$, but we do not know the specific member, call it P_0, that generated $\boldsymbol{X}$. The basic goal is to determine P_0 based on $\boldsymbol{X}$. In the case of a parametric family, we need to determine the value of $\boldsymbol{\theta}$, say $\boldsymbol{\theta}_0$, that corresponds to the distribution of $\boldsymbol{X}$.

To understand the role that Basu's Theorem plays in inference, we first need to define some basic concepts.

2.1 Definitions

A *sufficient statistic* $\boldsymbol{T}$ is defined by the property that the conditional distribution of the data $\boldsymbol{X}$ *given* $\boldsymbol{T}$ is the same for each member of $\mathcal{P}$. When densities exist, the density of $\boldsymbol{X}$ factors into two parts: the conditional density of $\boldsymbol{X}$ given $\boldsymbol{T}$, which is the same for every member of $\mathcal{P}$ and thus provides no information about P_0; and the density of $\boldsymbol{T}$, which contains all the information in the data which points to P_0 as the distribution of $\boldsymbol{X}$. It is thus reasonable that statistical inference be based solely on the sufficient statistic $\boldsymbol{T}$ which is usually of much smaller dimension that the data vector $\boldsymbol{X}$.

Reduction to a sufficient statistic, however, is neither unique nor guaranteed to yield maximum reduction. When it exists, a *minimal sufficient statistic* achieves the greatest reduction in data without losing information, that is, $\boldsymbol{T}$ is minimal sufficient if it is sufficient and can be computed from (is a function of) *every* sufficient statistic.

An *ancillary statistic* $\boldsymbol{V}$ is one whose distribution is the same for all members of $\mathcal{P}$. Therefore, $\boldsymbol{V}$ contains no information about the distribution of $\boldsymbol{X}$.

A statistic $\boldsymbol{T}$ is said to be *complete* with respect to the family $\mathcal{P}$ if there are no functions ϕ such that $E\phi(\boldsymbol{T}) = 0$ for all members of $\mathcal{P}$, except $\phi(t) \equiv 0$. The relationship between completeness and sufficiency may now be discussed.

2.2 Basu's Theorem in Inference

The essential difference between a sufficient (or even minimal sufficient) statistic and a complete sufficient statistic is that the sufficient statistic $\boldsymbol{T}$ may contain extraneous information which is not relevant for determining P_0. That is, there may exist a function of $\boldsymbol{T}$ which is ancillary. On the other hand, Basu's theorem tells us that if $\boldsymbol{T}$ is complete in addition to being sufficient, then no ancillary statistics can be computed from $\boldsymbol{T}$. Therefore $\boldsymbol{T}$ has been successful in "squeezing out" any extraneous information.

Note that minimality of $\boldsymbol{T}$ does not guarantee this "squeezing out" but completeness does. Complete sufficient statistics are, however, always minimal sufficient (Lehmann 1983, p. 46). A nice discussion of the ideas in this subsection is found in Lehmann (1981).

Consider, for example, the Laplace location family

$$\mathcal{P} = \left\{ f_\theta(x) = f_0(x-\theta), f_0(x) = \frac{1}{2}exp(-|x|), -\infty < x < \infty, -\infty < \theta < \infty \right\}$$

The order statistic $\boldsymbol{T} = (X_{(1)}, \ldots, X_{(n)})$ is sufficient for $\mathcal{P}$ but $\boldsymbol{V} = (X_{(n)} - X_{(1)}, \ldots, X_{(n)} - X_{(n-1)})$ is ancillary even though it may be obtained from $\boldsymbol{T}$.

Is $\boldsymbol{T}$ complete? Arguing from the contrapositive to Basu's theorem, we can say that since $\boldsymbol{T}$ and $\boldsymbol{V}$ are not independent (since $\boldsymbol{V}$ is a function of $\boldsymbol{T}$), then $\boldsymbol{T}$ is not a complete sufficient statistic for $\mathcal{P}$. A more typical method of showing lack of completeness comes from noting that $E[X_{(n)} - X_{(1)}] = c_1$, where c_1 just depends on $f_0(x) = (1/2)exp(-|x|)$ and can be computed. Thus $\boldsymbol{T}$ is not complete because $E\phi(\boldsymbol{T}) = 0$ for all members of $\mathcal{P}$, where $\phi(\boldsymbol{T}) = X_{(n)} - X_{(1)} - c_1$.

The Fisherian tradition suggests that when a sufficient statistic is not complete, then we should condition on ancillary statistics for the purposes of inference. This approach runs into problems because there are many situations where several ancillary statistics exists but there is no "maximal ancillarity" (Basu 1964). Fortunately, when a complete sufficient statistic exists, Basu's Theorem assures us that we need not worry about conditioning on ancillary statistics since they are all independent of the complete sufficient statistic.

2.3 Examples of Complete Sufficient and Ancillary Statistics

Before moving to applications, it may be helpful to remember that complete sufficient statistics exist for regular full rank exponential families (Lehmann 1983, p. 46). These include for example the Poisson, gamma, beta, binomial, many normal distribution models (univariate, multivariate, regression, ANOVA, etc.), truncated versions (Lehmann, 1983, prob. 5.31, p. 68), and censored versions (Bhattacharyya, Johnson, and Mehrotra 1977).

The order statistic $(X_{(1)}, \ldots, X_{(n)})$ is sufficient for regular families and is complete sufficient if $\mathcal{P}$ is the set of all continuous distributions or the set of all absolutely continuous distributions (Lehmann 1986, p. 143-144). This is very useful for applications to nonparametric problems.

In location models with densities of the form $f(x - \theta)$, random variables such as $\overline{X} - \theta$ are called pivotal quantities. In scale models with densities of the form $f(x/\theta)/\theta$, $\overline{X}/\theta$ is a pivotal quantity. Pivotal quantities are similar to ancillary statistics in that their distributions are the same for each member of the family, but pivotals are not statistics because their computation requires unknown quantities. In fact, a common method of verifying ancillarity of $\boldsymbol{V}$ is to re-express it as a function of a pivotal quantity $\boldsymbol{W}$.

It is easy to show that all location-invariant statistics are ancillary if $\mathcal{P}$ is a location family; all scale-invariant statistics are ancillary if $\mathcal{P}$ is a scale family; and all location-and-scale-invariant statistics are ancillary if $\mathcal{P}$ is a location-and-scale family.

3. TECHNICAL APPLICATIONS

Basu's theorem may be applied to a variety of problems. A selective presentation follows.

3.1 Independence of $\overline{X}$ and S^2.

Suppose $\boldsymbol{X}$ represents a random sample from a $N(\mu, \sigma^2)$, where σ^2 is known. Standard exponential family results yield that the sample mean $\overline{X}$ is complete sufficient for this normal location family. Moreover, the residual vector $\boldsymbol{V} = (X_1 - \overline{X}, \ldots, X_n - \overline{X})$ is seen to be ancillary because it is location-invariant and so may be written as $((X_1 - \mu) - (\overline{X} - \mu), \ldots, (X_n - \mu) - (\overline{X} - \mu))$, which is a function of the pivotal vector $\boldsymbol{W} = (X_1 - \mu, \ldots, X_n - \mu)$ whose distribution clearly does not depend on μ. Thus, $\boldsymbol{V}$ has the same distribution for each μ.

Basu's Theorem now tells us that since $\overline{X}$ is complete sufficient and S^2 is a function of the ancillary statistic $\boldsymbol{V}$, then $\overline{X}$ and S^2 are independent. Although we assumed that σ^2 is known, this result holds true for *any* normal distribution since our knowledge of σ^2 has nothing to do with the joint distribution of $\overline{X}$ and S^2. In fact, assuming σ^2 unknown only unnecessarily complicates the task of finding the joint distribution of $\overline{X}$ and S^2.

In general a key to applying Basu's Theorem is deciding which parameters to assume known and which to leave unknown.

3.2 Monte Carlo Swindles.

A *Monte Carlo swindle* is a simulation technique which allows a small number of replications to produce statistical accuracy at the level one would expect from a much larger number of replications. In some swindles, Basu's Theorem can be used to derive the basic independence needed to define the method.

Suppose, for example, that we want to estimate the variance of the sample median M of a random sample $\boldsymbol{X}$ of size n from a $N(\mu, \sigma^2)$ distribution. The straightforward approach is to generate N samples of size n from a $N(\mu, \sigma^2)$ distribution, compute the median for each sample resulting in $M_1, \ldots, M_N$, and then compute the sample variance of those medians as the estimate of $\text{var}(M)$.

The *Monte Carlo swindle* used here is to estimate $\text{var}(M - \overline{X})$ (instead of $\text{var}(M)$) by the N Monte Carlo samples and add $\text{var}(\overline{X}) = \sigma^2/n$ to get a more precise estimate of $\text{var}(M)$.

Why does the swindle work? First note that as in the previous subsection, $\overline{X}$ and $\boldsymbol{V} = (X_1 - \overline{X}, \ldots, X_n - \overline{X})$ are independent by Basu's Theorem because $\overline{X}$ is complete sufficient and $\boldsymbol{V}$ is ancillary. Consequently, $\overline{X}$ is also independent of the median of $\boldsymbol{V}$ which is $\text{median}(X_1 - \overline{X}, \ldots, X_n - \overline{X}) = M - \overline{X}$. Hence

$$\text{var}(M) = \text{var}(M - \overline{X} + \overline{X}) = \text{var}(M - \overline{X}) + \text{var}(\overline{X}).$$

Thus, from the simulation we would find the sample variance of $M_1 - \overline{X}_1, \ldots, M_N - \overline{X}_N$ and add it to σ^2/n to get our estimate of $\text{var}(M)$.

To see that the swindle estimate is more precise than the straightforward estimate, note that the variance of these estimates are approximately $2[\mathrm{var}(M - \overline{X})]^2/(N-1)$ and $2[\mathrm{var}(M)]^2/(N-1)$ respectively, due to the asymptotic normality of $M - \overline{X}$ and $\overline{M}$. Thus the swindle estimate is seen to have smaller variance than the first approach since $\mathrm{var}(M - \overline{X})$ $< \mathrm{var}(M)$ due to the high correlation between M and $\overline{X}$.

More intuitively we note that the Monte Carlo swindle approach is more efficient because the $\mathrm{var}(\overline{X}) = \sigma^2/n$ contribution to $\mathrm{var}(M)$ is much larger than the contribution from $\mathrm{var}(M - \overline{X})$ which is being estimated by simulation. The error in estimation by simulation is therefore limited to a very small part of $\mathrm{var}(M)$.

3.3 Moments of Ratios.

It is easy to derive the following simple result: if X/Y and Y are independent and appropriate moments exist, then

$$E\left(\frac{X}{Y}\right)^k = \frac{E(X^k)}{E(Y^k)}. \tag{1}$$

We now use Basu's Theorem and (1) in several applications.

3.3.1 Score Tests

In a number of testing problems, the score statistic has the form $(\sum X_i)^2/\sum X_i^2$ and is typically compared to a chisquared distribution with one degree of freedom. In a particular example from Liang (1985), Zhang and Boos (1997) discovered by simulation that this asymptotic testing procedure is conservative in small samples, that is, the χ_1^2 critical values are too large on average. Initially this result seemed counter-intuitive since the form of the statistic appears closer to that of an t^2 statistic than to a Z^2 statistic.

To investigate further, we may simplify the problem by assuming that $\boldsymbol{X}$ is a random sample from a $N(\mu, \sigma^2)$ distribution with σ^2 unknown and that the goal is to test the null hypothesis $\mu = 0$ using the statistic

$$t_s^2 = \frac{(\sum X_i)^2}{\sum X_i^2} = \frac{n\overline{X}^2}{\frac{1}{n}\sum X_i^2}.$$

Basu's Theorem can help us calculate the null moments of t_s^2.

If $\boldsymbol{X}$ is a random sample from the scale family $N(0, \sigma^2)$, then $T = \sum X_i^2$ is complete sufficient and $V = (\sum X_i)^2 / \sum X_i^2 = t_s^2$ is ancillary because it is scale-invariant and so may be written as the ratio of the two pivotal quantities $(\sum X_i/\sigma)^2$ and $\sum(X_i/\sigma)^2$. Basu's theorem therefore yields independence of T and V. In addition, result (1) tells us that

$$E(V^k) = \frac{E[(\sum X_i)^{2k}]}{E[(\sum X_i^2)^k]},$$

so that the mean and variance of t_s^2 are 1 and $2 - 6/(n+2)$. Compared to the mean and variance of the χ_1^2 distribution which are 1 and 2, respectively, we can see that comparing t_s^2 to the percentiles of a χ_1^2 distribution will give conservative test levels. In a class setting, it might be worth mentioning than an exact solution to this simplified testing problem is to transform monotonically to the square of the usual one sample t statistic and use $F(1, n-1)$ distribution critical values.

3.3.2 Minimum Variance Unbiased Estimation

Suppose again that $\boldsymbol{X}$ is a random sample from the scale family $N(0, \sigma^2)$. The search for minimum variance unbiased estimators is limited to functions of $T = \sum X_i^2$, since T is a complete sufficient statistic. One way to get unbiased functions of T is to use conditional expectations.

Consider, for example, $X_{(n)}$ the largest order statistic. If $aX_{(n)} + b$ is unbiased for the quantity to be estimated, then a minimum variance unbiased estimator is based on the the conditional expectation of $X_{(n)}$ given T. Basu's theorem helps us find this conditional expectation. Since T is complete sufficient and $V = X_{(n)}/\sqrt{T}$ is ancillary (it is scale-invariant), T and V are also independent. Thus

$$E(X_{(n)}|T) = E\left[X_{(n)}\frac{\sqrt{T}}{\sqrt{T}} \mid T\right] = E(V\sqrt{T}|T) = \sqrt{T}E(V).$$

Moreover, using result (1) we have that $E(V) = E(X_{(n)})/E(\sqrt{T})$.

3.4 Behrens-Fisher Problem.

Suppose that $X_1, \ldots, X_m$ is a random sample from $N(\mu_1, \sigma_1^2)$ and $Y_1, \ldots, Y_n$ is a random sample from $N(\mu_2, \sigma_2^2)$. When no assumption is made that $\sigma_1^2 = \sigma_2^2$, the testing situation $H_0 : \mu_1 = \mu_2$ is called the Behrens-Fisher problem. Numerous approaches have been proposed including Welch's solution (see Wang, 1971): compare

$$t_w = \frac{\overline{X} - \overline{Y}}{\sqrt{\frac{S_1^2}{m} + \frac{S_2^2}{n}}}$$

to a t distribution with estimated degrees of freedom f. Wang (1971) showed how to calculate the exact probability of a Type I error for Welch's solution:

$$P(t_w \geq t_{\alpha,f}) = P(t(\rho) \geq h(\alpha, m, n, S_1^2/S_2^2)),$$

where h is a somewhat messy function,

$$t(\rho) = \frac{\overline{X} - \overline{Y}}{\sqrt{\left(\frac{1}{m} + \frac{\rho}{n}\right) \frac{(m-1)S_1^2 + (n-1)S_2^2/\rho}{(m+n-2)}}},$$

and $\rho = \sigma_2^2/\sigma_1^2$. The key to evaluating this probability is to note that $t(\rho)$ has a t distribution with $m + n - 2$ degrees and that $t(\rho)$ is independent of S_1^2/S_2^2.

But why are $t(\rho)$ and S_1^2/S_2^2 independent? This fact was merely stated in Wang (1971) without argument. If ρ is assumed known, then S_1^2/S_2^2 is ancillary, and $t(\rho)$ is a function of the complete sufficient statistic $\boldsymbol{T} = (\overline{X}, \overline{Y}, (m-1)S_1^2 + (n-1)S_2^2/\rho)$. Basu's theorem then says that $t(\rho)$ and S_1^2/S_2^2 are independent. As with the example in Section

3.1, the assumption that ρ is known has no consequence on the joint distribution of $t(\rho)$ and S_1^2/S_2^2, so the result holds true even in the case that ρ is unknown.

As mentioned previously the trick in many applications is deciding what to assume is known in order to get both a useful complete sufficient statistic and an ancillary statistic. Consider, for example, assuming that ρ is unknown. Then the complete sufficient statistic is $\boldsymbol{T} = (\overline{X}, \overline{Y}, S_1^2, S_2^2)$, and while S_1^2/S_2^2 is a function of $\boldsymbol{T}$, $t(\rho)$ is not ancillary (even though it is pivotal under H_0).

3.5 Rank Statistics

For random samples from continuous distributions, the rank vector and the order statistic are independent. A simple proof of this fact is to note that the rank vector is ancillary and thus independent of the complete sufficient order statistic. This result applies to any continuous distribution. The key here is to take the set of probability measures large enough in order for the order statistic to be complete.

A direct proof of the independence of the rank vector and the order statistic may be found in Randles and Wolfe (1979, p. 273). Application of this independence result in the proof of asymptotic normality of linear rank statistics may be found on p. 285 of that text. Application of the result to adaptive rank tests may be found in Hogg, Fisher, and Randles (1975, p. 658).

REFERENCES

Basu, D. 1955, "On Statistics Independent of a Complete Sufficient Statistic," *Sankhya*, 15, 377–380.

Basu, D. 1964, "Recovery of Ancillary Information," *Sankhya*, Ser. A, 26, 3–16.

Bhattacharyya, G. K., Johnson, R. A., and Mehrotra, K. G. 1977, "On the Completeness of Minimal Sufficient Statistics with Censored Observations," *The Annals of Statistics*, 5, 547–553.

Casella, G., and Berger, L. R. (1990), *Statistical Inference*, Wadsworth & Brooks/Cole.

Hogg, R. V., Fisher, D. M., and Randles, R. H. (1975), "A Two-Sample Adaptive Distribution-Free Test," *Journal of The American Statistical Association*, 70, 656–661.

Lehmann, E. L. (1981), "An Interpretation of Completeness and Basu's Theorem," *Journal of The American Statistical Association*, 76, 335–340.

Lehmann, E. L. (1983), *Theory of Point Estimation*, Wiley.

Lehmann, E. L. (1986), *Testing Statistical Hypotheses*, Wiley.

Liang, K. Y. (1985), "Odds Ratio Inference with Dependent Data," *Biometrika* 72,678-682.

Randles, R. H. and Wolfe, D. A. (1979), *Introduction to the Theory of Nonparametric Statistics*, Wiley.

Wang, Y. Y. (1971), "Probabilities of the Type I Errors of the Welch Tests for the Behrens-Fisher Problem," *Journal of the American Statistical Association*, 66, 605–608.

Zhang, J., and Boos, D. D. (1997), "Mantel-Haenszel Test Statistics for Correlated Binary Data," to appear in *Biometrics.*