ENVIRONMETRICS
Environmetrics 2001; **12**: 437–449 (DOI: 10.1002/env.470)

Environmental power analysis – a new perspective

David R. Fox*,†

CSIRO, Private Bag, PO Wembley, Western Australia 6018, Australia

SUMMARY

Power analysis and sample-size determination are related tools that have recently gained popularity in the environmental sciences. Their indiscriminate application, however, can lead to wildly misleading results. This is particularly true in environmental monitoring and assessment, where the quality and nature of data is such that the implicit assumptions underpinning power and sample-size calculations are difficult to justify. When the assumptions are reasonably met these statistical techniques provide researchers with an important capability for the allocation of scarce and expensive resources to detect putative impact or change. Conventional analyses are predicated on a general linear model and normal distribution theory with statistical tests of environmental impact couched in terms of changes in a population mean. While these are 'optimal' statistical tests (uniformly most powerful), they nevertheless pose considerable practical difficulties for the researcher. Compounding this difficulty is the subsequent analysis of the data and the impost of a decision framework that commences with an assumption of 'no effect'. This assumption is only discarded when the sample data indicate demonstrable evidence to the contrary. The alternative ('green') view is that any anthropogenic activity has an impact on the environment and therefore a more realistic initial position is to assume that the environment is already impacted. In this article we examine these issues and provide a re-formulation of conventional mean-based hypotheses in terms of population percentiles. Prior information or belief concerning the probability of exceeding a criterion is incorporated into the power analysis using a Bayesian approach. Finally, a new statistic is introduced which attempts to balance the overall power regardless of the decision framework adopted. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS: sample-size determination; statistical inference; environmental assessment; percentiles

1. INTRODUCTION

The use of power analysis (PA) and sample-size determination (SSD) to plan environmental studies has become almost de facto practice among environmental researchers (Green, 1989, 1994; Peterman, 1990; Fairweather, 1991). The intention is clear: we wish to determine the amount of resource that is required to confidently assert that an impact or change in environmental condition has occurred when indeed this is the case. At the same time, we wish to protect ourselves from incorrectly reaching this conclusion. The probabilities associated with these outcomes are respectively referred to as the *power* and *size* (or level of significance) of the statistical test. The techniques are well known to statisticians, although interestingly this group is perhaps less enthusiastic about the utility of PA and SSD than those who routinely apply them. The reasons for this may in part be due to the fact that power and sample size calculations are predicated on some reasonably strong statistical assumptions–assumptions that

* Correspondence to: David Fox, CSIRO, Private Bag, PO Wembley, Western Australia 6018, Australia
†E-mail: david.fox@epo.csiro.au

Received 5 July 2000
Accepted 26 September 2000

are often difficult to justify in practice. Mother nature, it seems, is an unwilling participant in designed experimentation and studies to assess her health. Data obtained from environmental studies are notorious for their paucity, non-normality, heterogeneous error structure, over-dispersion, and spatial–temporal dependency.

The presence of one or more of these anomalies will affect both the power and size of the statistical test. Data paucity will reduce statistical power. The non-normality issue is not as important as many researchers would believe, and much of the effort in eliciting a suitable transformation of the data is unwarranted. Serious departures from normality will have an unpredictable effect on any reported *p*-value (Horton, 1978). The presence of heterogeneous error structure, over-dispersion, and spatial-temporal dependency will generally invalidate the results of PA, SSD, and ANOVA analyses.

Notwithstanding these difficulties, power and sample-size analyses implicitly assume a statistical model for the data. Assessment of the adequacy of this model is often overlooked or ignored once the sample data have become available. Clearly, if the model and the assumptions on which the power analysis has been based are incorrect, then any determination of sampling effort is likely to be seriously flawed. Although it is possible in certain instances to derive power and sample size 'rules' or formulae for non-normal cases (eg. inference for Poisson means), these techniques are less well known and consequently not as frequently employed. The plethora of computer software for PA and SSD (Thomas and Krebbs, 1997) for inference concerning means of normal distributions has only served to cement the foundations of these techniques in environmental assessment.

More recently, the applicability of the statistical significance-testing paradigm for environmental assessment has been called into question (McBride *et al.*, 1993; Johnson, 1999; Suter, 1996). These concerns have also been echoed by Nelder (1999), who is highly critical of the myopic focus that has resulted from a dependency on binary hypotheses, meaningless multiple comparison procedures and an unhealthy obsession with *p*-values. Given this rising tide of discontent with significance testing, we might feel inclined to abandon sample-size and power analyses altogether. However, while reform is clearly indicated, we should not throw the statistical baby out with the bathwater. If performed with due diligence, a mild degree of scepticism and appropriate attention to assumptions, power analysis, and sample-size determinations give us some quantitative measure of the likely performance of statistical tests and help focus our minds on the allocation of scarce resources. Rather than reject these methods out of hand and in the absence of credible alternative strategies, this article attempts to reconcile some of the philosophical and operational difficulties by developing a compromise framework for significance testing.

The remainder of this article is structured as follows. In section 2 we discuss some of the conceptual difficulties with the present hypothesis testing paradigm as it relates to environmental investigations and establish the argument for a compromise or hybrid approach. A new statistic called the 'environmental power' (*EP*) is introduced and its theoretical development presented in Section 3 with an example illustrating its use given in Section 4. Section 5 establishes the nexus between the new *EP* statistic and its conventional counterpart for sample-size determination and use in the planning of environmental studies. Some additional properties of the *EP* statistic are discussed in Section 6 and some indications for further work provided in Section 7.

2. A DECISION FRAMEWORK FOR ASSESSING ENVIRONMENTAL CONDITION

Consider the problem of determining whether or not some human activity has resulted in an environmental impact (however defined). A fairly typical, if not standard, approach commences with the formulation of a hypothesis to be tested (this hypothesis translates the notion of 'impact' into a statement about a parameter value) followed by the collection of sample data and concludes with a decision about the true state of nature. The use of conventional statistical inference has forced environmental scientists to adopt the frequentist's dichotomous decision-making process where one is forced to conclude that an impact has occurred or it hasn't. The wisdom of this approach in an environmental context is questionable. It could be argued that a more sensible strategy is afforded by Bayesian methods, for example, which seek to refine prior belief in the light of new evidence and to express the outcome as a probability distribution. An alternative view is that the binary decision-making process should be replaced by improved methods to represent and model the spatial and temporal extent of the quantity of interest. A product of this approach is typically some sort of map that makes no attempt to classify the impact, but simply describes it. Those who view the map must then ascribe to it their own subjective assessment of the 'significance' of what they see. Furthermore, the notion of an 'environmental control' in BACI (before-after-control-impact) type assessments (Underwood, 1990, 1991, 1994) may be regarded as an oxymoron. The difficulty is that by necessity an environmental control has to be sufficiently far removed from a potentially impacted site that one can no longer guarantee that the integrity of the underlying spatial–temporal processes have been preserved. Additionally, while it may be possible to identify a site that has not been disturbed by the impact under investigation, it is highly unlikely that it has not been disturbed by some other human activity. Smith *et al.* (1993) discuss other issues in the context of BACI designs.

It is generally recognized that statistical significance and ecological significance are not exchangeable concepts, although a proclamation of the latter is usually based on detection of the former. Figure 1 shows some scenarios to illustrate the relationships between the two concepts.

Case 1: The observed effect is not statistically significant, and the results are inconsistent with an effect large enough to be ecologically significant.

Case 2: The observed effect is not statistically significant, although the results are consistent with effects large enough to be ecologically significant.

Case 3: The observed effect is statistically significant but the results are not consistent with effects which are large enough to be ecologically significant.

Case 4: The observed effect is statistically significant and the effect may or may not be large enough to be ecologically significant.

Case 5: The observed effect is statistically significant and we can be confident that the effect is large enough to be ecologically significant.

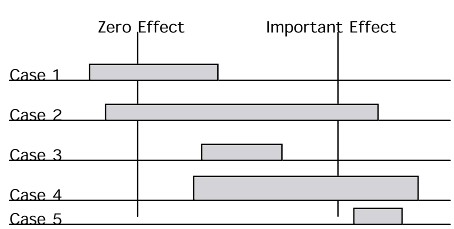

Figure 1. Confidence intervals for five different environmental scenarios.

Notwithstanding the previously identified difficulties with environmental power analysis, there are clear benefits in adopting a statistical approach to the planning of environmental studies. As noted by Green (1994), Fairweather (1991) and others, the adoption of a statistical framework for inference forces the researcher to provide a precise and explicit statement of null and alternative hypotheses. Power analysis provides an additional 'reality check' by ensuring that the planned allocation of effort is neither inadequate nor profligate. This notion of identifying a sampling and analysis regime that is 'fit for purpose' is one that deserves more attention in the environmental sciences. Yet there are problems with the application of Neyman–Pearson-type decision-making processes. The ANOVA framework devised by Fisher and Yates some 60 years ago at Rothamstead Experimental station in the U.K. was a triumph of statistical ingenuity. The methodology is perhaps the single most recognized, used, and abused of all statistical tools. When applied correctly and the attendant assumptions have been satisfied, ANOVA methods work superbly. Under these conditions ANOVA is an optimal (uniformly most powerful) statistical procedure. This single technique has, over the years, demonstrated its utility in applications as diverse as psychology, medicine, epidemiology, physics, nutrition, pharmacology, education, business, and marketing. It is not surprising therefore to see analysis of variance adopted as an almost de facto standard for environmental monitoring and assessment. While mean-based inference is appropriate in many instances, it is less likely to be relevant to environmental studies. There are parallels between environmental protection and engineering design and construction. Designers of dams, bridges, and other hydrologic structures are not so much interested in average conditions as they are with extremes (eg. the *maximum* load or the 1-in-100 year flood). Environmental scientists have a similar interest in extrema–high rainfall events, low river flows, high contaminant concentrations, low abundance, and/or diversity of biological populations. Thus a more appropriate formulation of statistical hypotheses would be couched in terms of the *percentiles* of a distribution rather than its mean.

The environmental scientist is perhaps also uncomfortable with other aspects of conventional statistical hypothesis testing procedures, in particular, the specification of a null hypothesis that invariably assumes no impact or change has occurred. This hypothesis is only rejected when the sample evidence is incontrovertible in its support for the alternative hypothesis. While this formulation of the problem and testing strategy appeal to our notion of 'innocent until proven guilty' it is somewhat at variance with the spirit of 'environmental protection' and the precautionary principle. Indeed, the more environmentally 'green' among us would argue that any anthropogenic activity is not without an environmental effect and therefore we should adopt as our initial premise a statement that the impact has been deleterious. The onus of environmental sampling is thus to demonstrate beyond reasonable doubt the veracity of the complimentary hypothesis–that the impact is *not* deleterious. A mere re-jigging of the size and power of the statistical test does not extinguish concerns over the choice of an appropriate null hypothesis. The roles of power and level of significance cannot simply be reversed to accommodate a re-labelling of the null and alternative hypotheses. Thus we can identify two distinct paradigms or statistical frameworks for environmental hypothesis testing (Table 1).

Table 1. Two frameworks for environmental hypothesis testing

'Green' decision framework (Ω_1)	'Brown' decision framework (Ω_2)
H_0: impact *is* deleterious	H_0: no deleterious impact
H_1: impact is *not* deleterious	H_1: deleterious impact

Table 2. Identification of various error types under different decision frameworks

Decision framework	Action	True state of nature: Impact not deleterious	True state of nature: Impact is deleterious
Ω_1	Accept H_0	β_1	
	Reject H_0	$(1-\beta_1)$	α_1
Ω_2	Accept H_0		β_2
	Reject H_0	α_2	$(1-\beta_2)$

Some basic relationships between level of significance under the 'green' and 'brown' formulations of the statistical hypotheses are illustrated in Table 2.

The αs and βs in Table 2 are probabilities of committing the indicated error. An environmentally responsible monitoring program would seek to minimize either α_1 or β_2 depending on the analysis framework adopted. However, we acknowledge that, a priori, either decision framework may be adopted depending on the experimenter's preferences and bias. Equally, it is often the case where decision framework Ω_2 is adopted by the notional 'polluter' only to have the outcomes assessed by an agency whose philosophical stance on environmental protection is best reflected by decision framework Ω_1. Thus, in an attempt to reconcile these differences we might suggest that tests be designed to minimize the sum $\alpha_1 + \beta_2$.

In the next section we provide details of a framework for environmental power analysis that is couched in terms of population percentiles and which endeavours to accommodate the dual nature of the decision-making framework. A useful by-product of this approach is that, unlike mean-based inference, estimates of variance parameters are not required. Although this is not a new idea (see, for example, Guenther, 1977), its application as a tool for environmental assessment has not been previously developed.

3. ENVIRONMENTAL POWER ANALYSIS DEFINED

Let Θ be the true (population) proportion of the population that is numerically less than some threshold or criterion τ. For some variable, X, of interest, define the pth percentile, ξ_p as that value of X for which. $P(X \leq \xi_p) = p$. Furthermore, it will be assumed that environmental compliance is achieved when $\xi_p \leq \tau$. Our null and alternative hypotheses under each of the decision frameworks Ω_1 and Ω_2 have been identified in the previous section.

We first develop key results under decision framework Ω_2. An extension to Ω_1 is then straightforward.

The null and alternative hypotheses under Ω_2 can be re-written in terms of ξ_p:

$$\begin{aligned} H_0 &: \xi_p = \tau \\ H_1 &: \xi_p > \tau \end{aligned} \tag{1}$$

We treat θ as a hyper-parameter having some pdf $g_\Theta(\theta)$ (e.g. a beta distribution).

It is further assumed that the conditional distribution $X \mid \theta$ is $N(\mu_\theta, \sigma^2)$.

The hypotheses of Equation (1) can be equivalently expressed in terms of the population mean, μ:

$$\begin{aligned} &\mathrm{H}_0 : \mu = \mu_\theta \qquad \theta = p \\ &\mathrm{H}_1 : \mu > \mu_\theta \end{aligned} \tag{2}$$

where

$$\mu_\theta = \tau - z_\theta \sigma \tag{3}$$

and z_θ is the quantile from a unit normal distribution for probability θ.

Now an α-level test of the hypothesis in Equation (2) rejects H_0 if

$$\bar{X} > \tau + \frac{z_\alpha \sigma}{\sqrt{n}} - z_p \sigma \tag{4}$$

where $\bar{X}$ is the mean of a sample of size n from the population of X values.

The *conditional* power of this test is

$$P\left[\bar{X} > \tau + \frac{z_\alpha \sigma}{\sqrt{n}} - z_p \sigma \mid \mu = \mu_0\right] \tag{5}$$

Equation (5) can be alternatively written as

$$P[Z > z_\alpha - \sqrt{n}(z_p - z_\theta)] = (\Psi_2 \mid \Theta = \theta) \tag{6}$$

The *unconditional* power of the test (Ψ_2) is given by the expectation $\mathrm{E}_\theta[\Psi_2 \mid \theta]$. Thus,

$$\Psi_2 = \int_{-\infty}^{\infty} (\Psi_2 \mid \theta) g_\Theta(\theta) \mathrm{d}\theta \tag{7}$$

Under decision framework Ω_1 the unconditional power Ψ_1 is obtained using Equation (7) with appropriate substitution of subscripts and where

$$(\Psi_1 \mid \Theta = \theta) = P[Z > z_\alpha + \sqrt{n}(z_p - z_\theta)] \tag{8}$$

Observe that neither Equation (6) nor Equation (8) require the specification of σ^2, the population variance. This is a very attractive feature and removes a source of uncertainty in the power calculations.

A new statistic called the *environmental power* (*EP*) is introduced in an attempt to acknowledge the different approaches to power analysis as represented by frameworks Ω_1 and Ω_2.
The *EP* statistic is defined as follows:

$$\kappa(\theta) = 1 - \frac{\Psi_1(\theta)\Psi_2(\theta)}{\alpha^2} \tag{9}$$

where α is the (common) level of significance used for both test regimes. To understand how κ works consider the following scenarios.

Case 1: $\theta = p$

In this case the true proportion of compliances is as specified and $\kappa = 0$, implying that there is no basis for arguing either way as to the deleterious nature of the impact.

Case 2: $\theta > p$

In this case the true proportion of compliances is *greater* than that specified and $\Psi_1(\theta)$ will be relatively large and $\Psi_2(\theta)$ small. In the extreme as $\theta \to 1, \Psi_1(\theta) \to 1$ and $\Psi_2(\theta) \to 1$, and thus $\kappa(\theta) \to 1$. Under these circumstances $\kappa(\theta)$ will approach unity more rapidly than $\Psi_1(\theta)$, which is appropriate. Under the decision framework Ω_2, the null hypothesis of no deleterious impact is accepted with probability $1 - \alpha$ for $\theta > p$. Given that $1 - \alpha$ is typically quite large (eg. ≥ 0.95), the implicit readiness of declaring 'no deleterious impact' may be troublesome, particularly when θ is only marginally greater than p. The probability of making this declaration as indicated by the environmental power will be somewhat less than $1 - \alpha$ whenever θ is only marginally greater than p (the value of θ for which $\kappa(\theta)$ equals $1 - \alpha$ is found as the solution to $\Psi_1(\theta)\Psi_2(\theta) = \alpha^3$).

Case 3: $\theta < p$

In this case the true proportion of compliances is *less* than that which is specified and $\Psi_1(\theta)$ will be small and $\Psi_2(\theta)$ larger. In the extreme as $\theta \to 0, \Psi_1(\theta) \to 0$ and $\Psi_2(\theta) \to 1$ and thus $\kappa(\theta) \to 1$. Again, $\kappa(\theta)$ possesses more desirable characteristics than either $\Psi_1(\theta)$ or $\Psi_2(\theta)$ alone.
An example illustrating these principles is now provided.

4. EXAMPLE

Suppose that for a particular water quality criterion, the 90th percentile is specified to be τ. We wish to assess environmental compliance by an examination of specimen concentrations at 15 randomly selected locations, using an $\alpha = 0.05$ test. Thus in this case we have $n = 15$, $\alpha = 0.05$, and $p = 0.9$.

Now, $\theta = P[X \leq \tau]$ is the true proportion of the population that does not exceed the guideline value. We consider the following five scenarios for θ:

$\theta = 0.9$ (compliance in a strict sense, but not unequivocal)
$\theta = 0.925$ (marginally compliant)
$\theta = 0.98$ (demonstrably compliant)
$\theta = 0.88$ (marginally non-compliant)
$\theta = 0.6$ (demonstrably non-compliant).

Table 3 indicates the probabilities of making the correct judgement under Ω_1 and Ω_2 separately and also gives the *EP* for each case.

For the above example it can be established that $\kappa(0.9614) = 0.95$. Observe that $\theta = 0.9614$ lies between our subjective classifications of marginally compliant and demonstrably compliant. Under decision framework Ω_1 there is a considerably smaller probability (0.593) of declaring compliance than would be the case under Ω_2 or as measured by the *EP*.

Table 3. Probabilities of correct decision-making for various combinations of θ and decision frameworks

	θ				
	0.6	0.88	0.9	0.925	0.99
Ω_1	0.95	0.95	0.95	0.151	0.911
Ω_2	0.990	0.109	0.95	0.95	0.95
EP	1.00	0.1364	0.0	0.275	1.00

5. SAMPLE SIZE DETERMINATIONS USING *EPA*

It is clear from the example and the ensuing discussion that the *EP* as determined by the statistic $\kappa(\theta)$ has properties which are more intuitively appealing than those associated with either Ω_1 or Ω_2. On this basis we claim it makes more sense to plan sampling effort around the attainment of a 'reasonable' *EP* (e.g. > 0.8) rather than focusing on the conventional definitions of statistical power (Cohen, 1988). The immediate difficulty, however, is the dependency of the *EP* on the true (but unknown) θ. As with Equations (7) and (8) the unconditional *EP* is obtained as the integral

$$\int_{-\infty}^{\infty} \kappa(\theta) g_\Theta(\theta)\mathrm{d}\theta \tag{10}$$

An obvious candidate distribution for Θ is the beta, although we have found a suitably constrained Burr Type XII distribution (Burr, 1942) to be particularly useful. The form of the pdf is given by Equation (11):

$$g_\Theta(\theta; k, c) = \frac{kc(1-\theta)^{c-1}}{\left[1 + \left(\frac{1-\theta}{\theta}\right)^c\right]^{k+1} \theta^{c+1}} \quad 0 < \theta < 1 \tag{11}$$

The range of possible shapes available using the family of curves using Equation (11) is illustrated in Figure 2.

Continuing with the previous example, suppose our prior belief in θ is adequately expressed by the pdf $g(\theta; k = 5, c = 1.1)$ that has approximately 65 per cent of its probability mass below 0.9 and 35 per cent above. This prior probability distribution is consistent with a subjective belief that tends to favour non-compliance, although not overwhelmingly so. A plot of the pdf is shown in Figure 3.

Then the *unconditional EP* for a sample of size $n = 15$ is 0.725. To achieve an unconditional *EP* of 0.8 the sampling effort would need to be increased to $n = 29$. By way of comparison, the unconditional power associated with the conventional hypothesis formulation (Ω_2) is 0.461. Note that there is an upper limit on the power which could be achieved under the conventional ('brown') decision framework and no amount of sampling can increase this figure. This is a direct consequence of the Bayesian prior that we impose on the parameter Θ. This is readily explained as follows.

Under decision framework Ω_2, the *conditional* power of the test is given by Equation (5). As $n \to \infty$ we have $z_\alpha\sigma/\sqrt{n} \to 0$ and $\bar{X} \to \mu$. Thus, in the limit, Equation (5) becomes $P[\mu_\Theta > \tau - z_p\sigma]$.

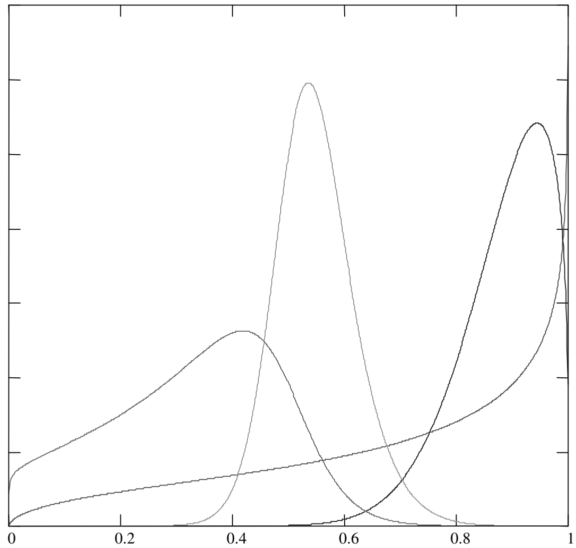

Figure 2. Examples of Burr's distributions generated using Equation (11).

But $\mu_\Theta = \tau - z_\theta\sigma$ and so this probability becomes $P[z_\Theta < z_p]$which is equivalent to $P[\Theta < p]$ (note that Θ is a random variable having prior pdf $g_\Theta(\theta)$ and hence z_Θ is also a random variable). The probability $P[\Theta < p]$ thus represents that maximum power attainable. In the example above this is 0.65.

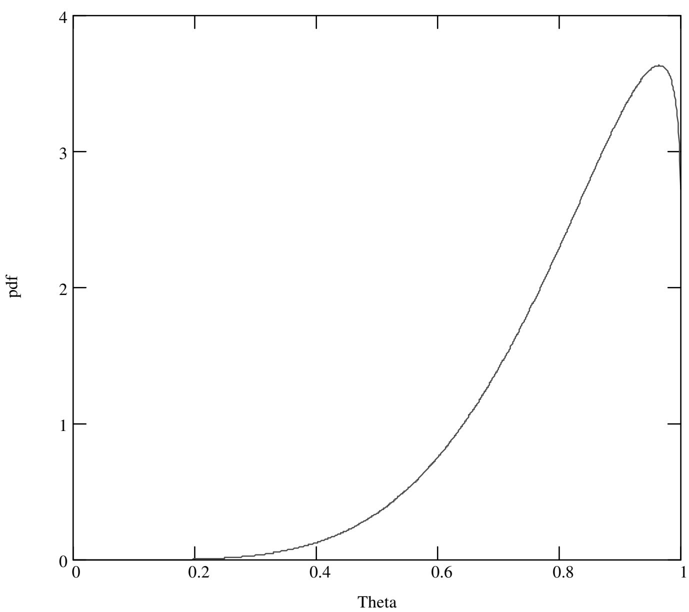

Figure 3. Burr's distribution with $k = 5$ and $c = 1.1$.

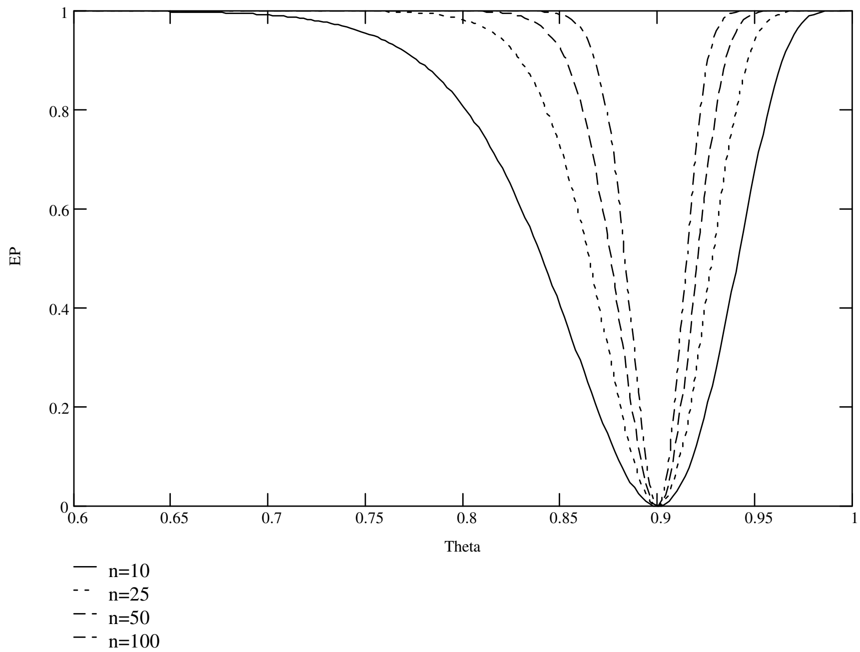

Figure 4. Relationship between *EP* and sample size for the case $p = 0.9$.

6. SOME PROPERTIES OF THE *EP* STATISTIC

Figure 4 shows conditional *EP* curves for samples of size 10, 25, 50, and 100 for the case $p = 0.9$. It is evident that the plots are asymmetrical although the asymmetry about $\theta = 0.9$ diminishes with increasing sample size. From Figure 4 we see that for a sample of size $n = 10$ the *EP* exceeds a nominal 0.8 for $\theta < 0.8$ or $\theta > 0.96$. In other words, even small samples have a high *EP* when we are in compliance (i.e. > 0.9). On the other hand, such a small sample can only detect demonstrably out of compliance situations (i.e. those for which $\theta < 0.8$). If we assumed that it was equally likely that the true value of θ lay between 0.85 and 0.95, say, then the unconditional *EP* would be 0.202 for $n = 10, 0.378$ for $n = 25, 0.530$ for $n = 50$, and 0.663 for $n = 100$. These figures are obtained by effectively averaging the appropriate curve in Figure 4 between θ in the range 0.85 to 0.95.

In the example of the previous section we chose a prior pdf for θ which reflected our belief that 'non-compliance' was more likely than 'compliance'. Thus, the 'averaging' process just described gives more weight to the out-of-compliance part of the *EP* curve. This situation is depicted in Figure 5.

7. CONCLUSIONS

In this article we have developed a compromise approach to power analysis and sample size determination for environmental assessment. The method acknowledges the legitimacy of different formulations of the null and alternative hypotheses for environmental monitoring and introduces a new statistic called the *environmental power* (or *EP*) as a means of reconciling potentially conflicting results that can arise from different approaches. Furthermore, the use of Bayesian priors to

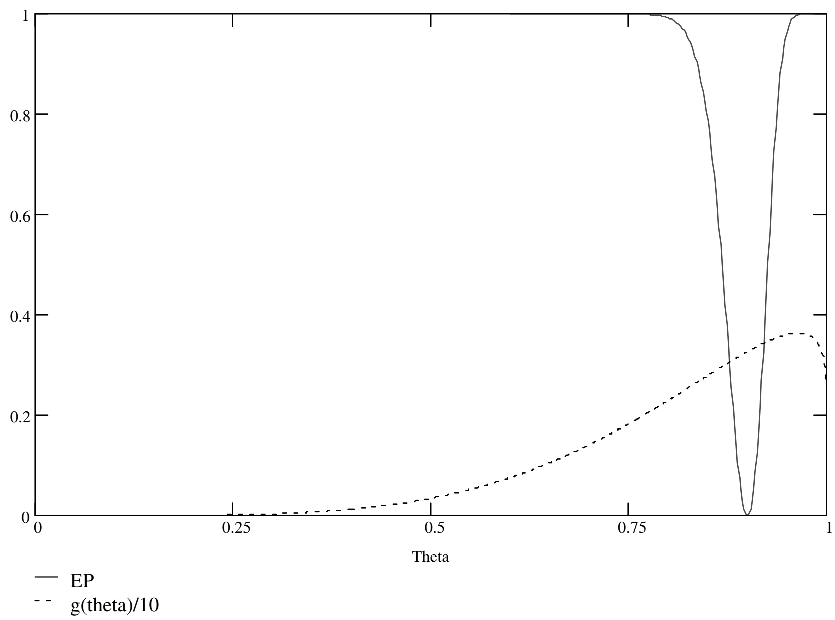

Figure 5. *EP* curve (solid line) for $n = 30$ and $p = 0.9$ with assumed prior distribution (dashed line) for θ overlaid.

accommodate imprecise information about the true state of nature coupled with an emphasis on population percentiles makes for a powerful and flexible planning tool.

It is recognized that these benefits come at the price of increased complexity and they demand a higher level of technical sophistication than do standard techniques. However, the procedures described are particularly amenable to computer implementation, and this would be a necessary next step if our *environmental power analysis* were to gain widespread popularity among natural resource managers.

One remaining difficulty concerns the appropriate choice of distribution for the Bayesian prior. We have introduced a modified Burr distribution as a flexible family of pdfs in the development of our methodology. Although the choice of parameter values for c and k will not be immediately obvious, a number of approaches can be contemplated. For example, by specifying the *mode* and *mean* of the distribution we can solve the following two equations for c and k:
Mode:

$$\frac{\mathrm{d}}{\mathrm{d}y}\frac{kc(1-y)^{c-1}}{\left[1+\left[\frac{(1-y)}{y}\right]^{c}\right]^{k+1}y^{c+1}}=0$$

Mean:

$$\int_0^1 y\frac{kc(1-y)^{c-1}}{\left[1+\left[\frac{(1-y)}{y}\right]^{c}\right]^{k+1}y^{c+1}}dy=\mu$$

Table 4. Parameter estimates for modified Burr distribution for specified mean and mode.

Mean	Mode	c	k
0.8	0.9	1.257	5.123
0.7	0.9	1.109	2.462
0.6	0.9	1.041	1.512
0.7	0.8	1.285	2.701

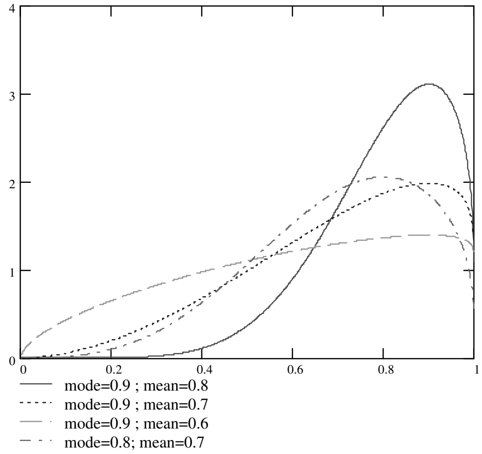

Figure 6. Modified Burr distributions for parameter values given in Table 4.

This has been done for a limited number of cases and the results displayed in Table 4 and the corresponding pdfs. illustrated in Figure 6.

ACKNOWLEDGEMENT

The author gratefully acknowledges the helpful suggestions made by Dr. Ray Correll (CSIRO Mathematical and Information Sciences).

REFERENCES

Burr IW. 1942. Cumulative frequency functions. *Annals of Mathematical Statistics* **13**, 215–232.

Cohen J. 1988. Statistical Power Analysis for the Behavioural Sciences (2nd edn). Lawrence Erlbaum Associates: Hillsdale NJ.

Fairweather PG. 1991. Statistical power and design requirements for environmental monitoring *Australian Journal of Marine and Freshwater Research* **42**: 555–567.

Green RH. 1989. Power analysis and practical strategies for environmental monitoring. *Environmental Research* **50**: 195–205.

Green RH. 1994. Aspects of power analysis in environmental monitoring. In *Statistics in Ecology and Environmental Monitoring*, eds. Fletcher DJ and Manly B. University of Otago Press: Dunedin, New Zealand.

Guenther WC. 1977. Power and sample size determinations when the alternative hypotheses are given in quantiles. *The American Statistician* **31**: 117–118.

Horton RL. 1978. *The General Linear Model Data – Analysis in the Social and Behavioural Sciences*. McGraw-Hill.

Johnson DH. 1999. The insignificance of statistical significance testing. *Journal of Wildlife Management* **63**(3): 763–772.

McBride GB, Loftis JC, and Adkins NC. 1993. What do significance tests really tell us about the environment? *Environmental Management* **17**: 423–432.

Nelder JA, 1999. Statistics for the millennium – from statistics to statistical science. *The Statistician* **48**(2): 257–269.

Peterman RM. 1990. Statistical power analysis can improve fisheries research and management. *Canadian Journal of Fisheries and Aquatic Sciences* **47**(1): 2–15.

Smith EP, Orvos DR, and Cairns J. 1993. Impact assessment using the before-after-control-impact (BACI) model: concerns and comments. *Canadian Journal of Fisheries and Aquatic Sciences* **50**: 627–637.

Suter GW. 1996. Abuse of hypothesis testing statistics in ecological risk assessment. *Human and Ecological Risk Assessment* **2**: 331–347.

Thomas L. and Krebbs CJ. 1997. A review of statistical power analysis software. *Bulletin of the Ecological Society of America* **78**(2): 128–139.

Underwood AJ. 1990. Experiments in ecology and management: their logics, functions, and interpretations. *Human and Ecological Risk Assessment* **2**: 331–347.

Underwood AJ. 1991. Beyond BACI: experimental designs for detecting human environmental impacts on temporal variation in natural populations. *Australian Journal of Marine and Freshwater Research* **42**: 569–587.

Underwood AJ. 1994. On beyond BACI: sampling designs that might reliably detect environmental differences. *Ecological Applications* **4**: 3–15.

