

Raindrop Plots: A New Way to Display Collections of Likelihoods and Distributions

Nicholas J. BARROWMAN and Ransom A. MYERS

In a variety of settings, it is desirable to display a collection of likelihoods over a common interval. One approach is simply to superimpose the likelihood curves. However, where there are more than a handful of curves, such displays are extremely difficult to decipher. An alternative is simply to display a point estimate with a confidence interval, corresponding to each likelihood. However, these may be inadequate when the likelihood is not approximately normal, as can occur with small sample sizes or nonlinear models. A second dimension is needed to gauge the relative plausibility of different parameter values. We introduce the raindrop plot, a shaded figure over the range of parameter values having log-likelihood greater than some cutoff, with height varying proportional to the difference between the log-likelihood and the cutoff. In the case of a normal likelihood, this produces a reflected parabola so that deviations from normality can be easily detected. An analogue of the raindrop plot can also be used to display estimated random effect distributions, posterior distributions, and predictive distributions.

KEY WORDS: Likelihood interval; Log odds ratio; Meta-analysis; Nonlinear model; Nonnormality; Variable-height plot.

1. INTRODUCTION

Likelihood plays a central role in the frequentist, evidentialist, and Bayesian statistical paradigms. In a variety of settings, we are faced with a collection of likelihoods. For example, in meta-analysis, we are interested in comparing and possibly combining information about a parameter of interest from a collection of studies. It may be possible to represent this information by a collection of single-parameter likelihoods. However, the display of these likelihoods can be challenging, as can be seen for the 41 studies of the treatment of peptic ulcers analyzed by Efron (1996) (Figure 1). Although the figure does give a general impression of the parameter values consistent with the studies, the likelihood for any given study is very difficult to discern and the overall effect is a tangled mess. An alternative is to display all 41 likelihoods separately. However this requires much more space

and makes it difficult to compare one likelihood to another. It is also noteworthy that because likelihoods are only defined up to a multiplicative constant, the choice of scaling when several likelihoods are graphed is arbitrary and different scalings may be perceived differently.

Rather than displaying collections of likelihoods, it is conventional simply to display point estimates with corresponding bars representing confidence intervals. Throughout this article, we refer to this as the “traditional display.” The simplest and most common confidence intervals are based on normal theory and are symmetric about the point estimate on the appropriate scale. A number of the likelihoods in Figure 1 indicate departures from normality: several of them are markedly asymmetric or monotonic. In general, normal approximations may be poor with small sample sizes, nonnormal distributions, and nonlinear models. Likelihood intervals (which are typically asymmetric) can be used in place of normal theory intervals; however, these may be inadequate when a point estimate does not exist, when there is no upper or lower limit, or when the likelihood has multiple local modes.

This article proposes a new display for collections of likelihoods, the *raindrop plot*, that overcomes these deficiencies. We then introduce a related display for displaying distributions. Our goal is to build on recent advances in the display of raw data (e.g., Lee and Tu 1997 and Hintze and Nelson 1998) to provide displays of parameter uncertainty and predictive distributions.

To illustrate the construction of the raindrop plot, we begin with the very simple case introduced in Figure 1, log odds ratios for 2×2 tables. A nonlinear regression model furnishes a second,




Figure 1. Conditional likelihoods for the log odds ratio comparing recurrent bleeding following different types of surgery in 41 studies of the treatment of peptic ulcers. The data were reported by Efron (1996) based on the meta-analysis by Sacks et al. (1990). Morris (1996) noted several discrepancies between Efron's dataset and that of Sacks et al.; we use Efron's dataset for comparability. Following Efron, we have scaled each likelihood to integrate to 1 over the interval shown.

Nicholas J. Barrowman is Chief Biostatistician, Chalmers Research Group, Children's Hospital of Eastern Ontario Research Institute, 401 Smyth Road, Ottawa, Ontario, Canada K1H 8L1 (E-mail: nbarrowman@cheo.on.ca). Ransom A. Myers is Killam Memorial Chair in Ocean Studies, Department of Biology, Dalhousie University, Halifax, Nova Scotia, Canada B3H 4J1 (E-mail: ransom.myers@dal.ca). This work was supported by the Killam Foundation, the Natural Sciences and Engineering Research Council of Canada, and the Sloan Foundation Future of Marine Animal Population (FMAP) project. We thank Robert Platt, David Moher, Ba' Pham, David Hamilton, and Chris Field for helpful suggestions.

more complex example, in which the raindrop plot is contrasted with the traditional display.

2. THE RAINDROP PLOT

Consider a 2×2 table (a, b, c, d) of data from a randomized controlled trial, where a and b are the numbers of failures and successes in the treatment group and c and d are the numbers of failures and successes in the control group. It is often desired to estimate the log odds ratio,

$$\theta = \log \left(\frac{P(\text{Failure}|\text{Treatment})P(\text{Success}|\text{Control})}{P(\text{Success}|\text{Treatment})P(\text{Failure}|\text{Control})} \right).$$

A good estimator of θ in terms of bias and mean squared error is

$$\hat{\theta} = \log \left(\frac{(a + \frac{1}{2})(d + \frac{1}{2})}{(b + \frac{1}{2})(c + \frac{1}{2})} \right) \quad (1)$$

with standard error

$$\text{SE}(\hat{\theta}) = \left(\frac{1}{a + \frac{1}{2}} + \frac{1}{b + \frac{1}{2}} + \frac{1}{c + \frac{1}{2}} + \frac{1}{d + \frac{1}{2}} \right)^{\frac{1}{2}}. \quad (2)$$

(Agresti 1990). $\hat{\theta}$ is asymptotically normal, but for small sample sizes its sampling distribution is highly skewed. Treating the margins of the 2×2 table as fixed, the conditional likelihood for θ is given by

$$L(\theta) = \binom{a+b}{a} \binom{c+d}{c} e^{\theta a} / S(\theta),$$

where $S(\theta)$ is the sum of the numerator above over the allowable choices of a subject to the marginal constraints of the table. Let θ^{MLE} be the MLE of θ , and denote the log of the conditional likelihood (or indeed, the log of any likelihood of a parameter θ) by $\ell(\theta)$. An approximate $100 \cdot (1 - \gamma)\%$ confidence interval for θ based on the asymptotic distribution of the likelihood ratio (Cox and Hinkley 1974) is given by

$$\{\theta : 2[\ell(\theta^{\text{MLE}}) - \ell(\theta)] \leq \chi_1^2(1 - \gamma)\},$$




Figure 2. Log-conditional likelihood for the log odds ratio, based on data $a = 1$, $b = 10$, $c = 15$, $d = 10$, showing how the raindrop shape is obtained. The log-conditional likelihood (solid curve) has been graphed with its maximum (indicated by the dotted vertical line) equal to 0. A drop in the log-likelihood of approximately 1.92 (indicated by the horizontal line) is significant at the 95% level. The corresponding 95% confidence interval for the log odds ratio is $(-5.65, -0.85)$. By reflecting the part of the curve above -1.92 about the horizontal line, we obtain a symmetric region (shaded in the figure). The height of the region at a particular value of the log odds ratio relative to the maximum height gauges the relative plausibility of that value. The vertical dashed lines on the left and right hand sides are the lower and upper 95% confidence limits based on the normal approximation and Equations (1) and (2).

where $\chi_1^2(p)$ is the $100 \cdot p$ th percentile of a chi-square distribution with 1 degree of freedom. Because the likelihood is only defined up to a multiplicative constant, for convenience we set $\ell(\theta^{\text{MLE}}) = 0$. Because $\chi_1^2(0.95)/2 \approx 1.92$, a 95% confidence interval is equivalent to

$$\{\theta : \ell(\theta) \geq -1.92\}. \quad (3)$$

In the evidential paradigm (Royall 1997), a $1/r$ support interval for θ is defined as

$$\{\theta : \ell(\theta) \geq \log(1/r)\}.$$

The benchmark value $r = 8$ gives a support interval slightly wider than the approximate 95% confidence interval (3).

As an example, consider data $a = 1$, $b = 10$, $c = 15$, $d = 10$. The estimated log odds ratio given by (1) is -2.34 , fairly close to the MLE, -2.65 . However, the likelihood is asymmetric about its maximum: the 95% confidence interval, computed as $\hat{\theta} \pm 1.96\text{SE}(\hat{\theta})$, is $(-4.22, -0.45)$, whereas the likelihood-ratio based interval (3) is substantially different: $(-5.65, -0.85)$.

This traditional display is sometimes misinterpreted by users as implying that all values within a confidence interval are equally plausible. Because the likelihood ratio provides a gauge of the relative plausibility of different values of θ and allows for asymmetry, a graphical display of the likelihood ratio would be desirable. An extra dimension is needed to represent the likelihood ratio. The principles of graph construction and the paradigm of graphical perception developed by Cleveland (1985) provide a guide to designing an appropriate display. To display likelihood ratios, we introduce a new kind of display, the *raindrop plot*, so called because the visual effect is reminiscent of raindrops streaking across a car window. To produce a “raindrop” shape, the log-likelihood (with its maximum chosen to be 0) is graphed over the range where it is greater than -1.92 . It is then reflected about the horizontal line at -1.92 and the resulting region is shaded, producing vertically symmetric raindrop shapes (Figure 2).

We believe that the reflection and shading simplify the perceptual task for the viewer who wishes to compare several likelihoods and detect asymmetries and other differences in shape. Similar approaches have been used by Lee and Tu (1997) and Hintze and Nelson (1998).

If x is the realization of a normal random variable X with known variance σ^2 , the log-likelihood for the mean θ of X has a particularly simple form: $\ell(\theta) \approx -(x - \theta)^2 / (2\sigma^2)$. Its graph—a parabola—has constant curvature, so that the “quadratic raindrop” has a particularly simple shape. Subtle deviations from normality can thus be detected by the eye; for example, a raindrop with any straight edges cannot be normal. Modern graphical data analysis environments, such as S, make producing such figures very easy. S-Plus/R code to produce a quadratic raindrop is given in the Appendix. It should be noted that deviation from normality can sometimes be extreme. For example, likelihoods can be multimodal and it is even possible that raindrops may consist of disconnected pieces, representing confidence sets rather than intervals.

A modification of the scheme for producing raindrop shapes can be used to display distributions. In place of the log-likelihood, we use the log probability density, with a cutoff

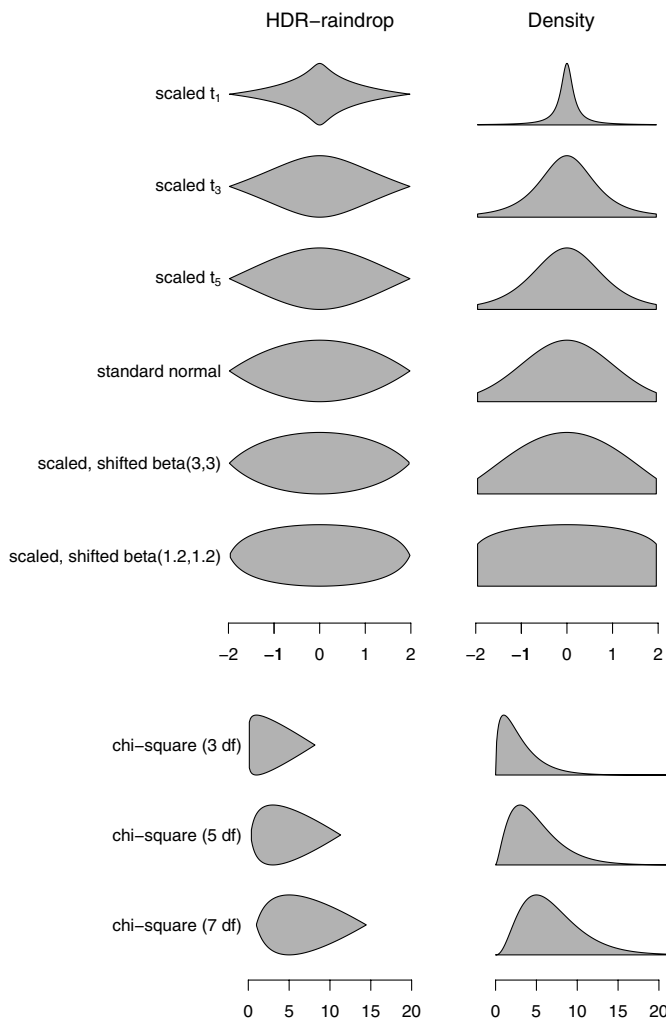


Figure 3. Reference 95% HDR raindrops for selected t distributions, the standard normal distribution, and selected β distributions (left panel) along with their probability density functions over the same interval (right panel). The β distributions have been shifted so they are centered at zero. The t and β distributions have been scaled so that their 2.5th and 97.5th percentiles match those of the standard normal distribution.

corresponding to a probability of 0.95. In other words, we use a raindrop based on the log density over the highest density region, or HDR (Hyndman 1996). We call this an HDR-raindrop, and it provides more information than the interval or region defined by the simple highest density region. For reference, Figure 3 shows HDR-raindrops for several distributions. HDR-raindrops provide information not only on where distributions have the bulk of their mass but also on other properties such as skewness and kurtosis. As with normal likelihoods, the HDR-raindrop for a normal distribution is a reflected parabola. Leptokurtic distributions give HDR-raindrops whose sides are “pinched” (as for the t_1) or angular (as for the t_3), while platykurtic distributions give HDR-raindrops with flat tops (as for the beta(1.2,1.2)). Note that the conventional measure of kurtosis, the standardized fourth central moment, is not defined for t -distributions with less than five degrees of freedom. [For a review of the concept and measurement of kurtosis, see Balanda and MacGillivray (1988).] Although an HDR-raindrop does not show the extreme tails of a distribution, the shapes of the sides of an HDR-raindrop are suggestive of the tail behavior. The concave sides of the t_1 HDR-

raindrop suggest that the tails of the distribution are very heavy, while the convex sides of a beta(1.2,1.2) HDR-raindrop suggest that the tails of the distributions are very light. The skewness of chi-squared distributions is clearly shown in HDR-raindrops. The angular sides of the HDR-raindrop for the chi-square with three degrees of freedom suggests a heavy tail compared to the chi-square with seven degrees of freedom.

Raindrops can be used to produce displays reminiscent of the forest plots often featured in meta-analyses (Lewis and Clarke 2001), with raindrops for individual sampling units followed by a summary raindrop. To visually distinguish the different types of raindrops we suggest the following. For likelihood-based intervals, plot 95% raindrops using gray shading, possibly superimposed on 99% raindrops with no shading. Plot summary rain-




Figure 4. Raindrop plot of log odds ratios and meta-analytic summaries for studies of ulcer treatments from the analysis of Efron (1996). For each study, 95% raindrops (shaded) are superimposed on 99% raindrops (unshaded). At the bottom, two meta-analytic summaries from Efron (1996) are shown: a normal 95% raindrop for the mean (based on a point estimate and jackknife standard error reported by Efron) and a 95% HDR (highest density region) raindrop for the “MLE prior” (shaded darker to emphasize that its interpretation is different). Note that the summary raindrops have been drawn taller than the individual study raindrops, both for visual emphasis and to aid in detecting deviation from normality. Note that Efron’s analysis excluded study number 40 (which was judged to be an overly influential outlier) and study number 41 (which has a flat likelihood).

drops using double height for emphasis and to facilitate detection of nonnormality. For HDR-raindrops, plot using double height and a different color or darker shading, to emphasize that their interpretation is different.

3. ULCER DATA

Raindrops for the 41 studies of ulcer treatment introduced earlier are shown in Figure 4.

Note that nine of the studies have zeros in at least one cell. The resulting likelihoods do not have peaks: except for study number 41 (which has a zero marginal total and hence a completely flat likelihood), the likelihoods are “ramped”: the likelihood increases monotonically, and there is no finite maximum likelihood estimate. Such cases are sometimes excluded from analyses, yet they clearly contain information. Commonly, point estimates and confidence intervals based on Equations (1) and (2) are used in these cases. Likelihood intervals are an alternative and raindrops automatically show these along with information about the relative plausibility of parameter values within the interval.

Efron (1996) illustrated his empirical Bayes methods for combining likelihoods using the ulcer studies. His approach involves a hierarchical model featuring a data-based approach to the choice of prior distribution (known as the random effects distribution in mixed model terminology). Efron’s “MLE prior” is an approximation to the Bayes predictive distribution for the log odds ratio of an as-yet unobserved study. Two summary raindrops based on the MLE prior are shown at the bottom of Figure 4. The first is a quadratic 95% raindrop based on the mean of the MLE prior and a jackknife standard error reported by Efron (1996). The second is an HDR-raindrop for the MLE prior itself, showing a subtle deviation from normality.

The summary raindrops at the bottom of Figure 4 are reminiscent of the diamond-shaped summary symbols often used in meta-analytic plots. However, unlike the diamonds, the second dimension of the raindrops is informative.

4. COHO SALMON DATA

To contrast the traditional display with the raindrop plot, we now consider a more complex example. In recent years there have been alarming declines in coho salmon (*Oncorhynchus kisutch*) populations on the west coast of North America. A population dynamics approach may help to explain why this is happening. Adult coho spawn in streams and rivers. About 1.5 years later, their offspring—known as smolts at this life stage—migrate to sea. Another 1.5 years later, the survivors return to spawn. Let S represent the quantity of spawners, and let R represent the quantity of young fish (or “recruits”) produced by those spawners. The dependence of R on S is of central importance in understanding the population biology of the species and for fisheries management. Coho salmon exhibit a characteristic spawner recruitment relationship for which Barrowman and Myers (2000) proposed a piecewise linear “hockey stick” model. However, the sharp bend in the hockey stick may not be biologically realistic; additionally, estimation of such models may lead to numerical difficulties and poor asymptotic behavior. Barrowman and Myers (2000) therefore proposed two smoother generalizations of the hockey stick. Here we consider

the “logistic hockey stick” with positive parameters α , θ , and μ :

$$R = \alpha\theta\mu(1 + e^{-1/\theta}) \left(\frac{S}{\theta\mu} - \log \left(\frac{1 + e^{(S-\mu)/(\theta\mu)}}{1 + e^{-1/\theta}} \right) \right). \quad (4)$$

Here α is the slope at the origin, μ is the inflection point, and θ is a smoothness parameter (assumed fixed and known). The parameter α is of particular interest because it can be used to estimate extinction rates, and obtain predictions concerning the recovery of over-exploited populations. Barrowman et al. (in press) analyzed data on 14 coho populations in the Pacific north-west region, using several models including the logistic hockey stick (Figure 5). So that the α parameter would be on a common scale across different populations, the values of spawners and recruitment were standardized by a measure of habitat size: S was measured as the number of spawning females per kilometer of river, and R was measured as the number of female smolts per kilometer of river.

Because α and μ are both positive parameters, it is convenient to perform computations in terms of $\log \alpha$ and $\log \mu$. To make inferences about the distribution of $\log \alpha$, we computed a profile likelihood for each population

$$L_\theta(\log \alpha) = \sup_{\log \mu} L_\theta(\log \alpha, \log \mu),$$

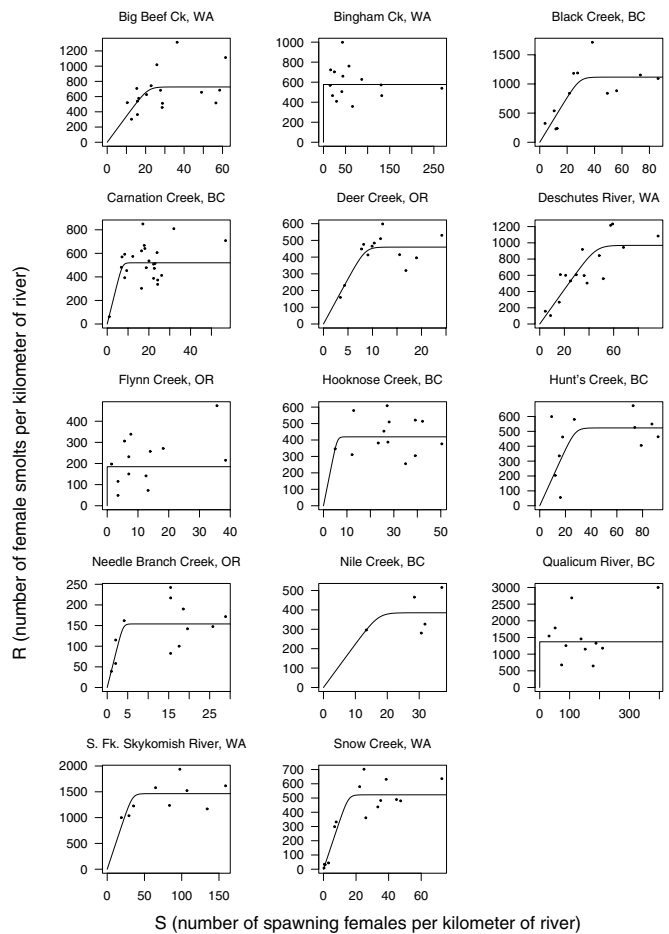


Figure 5. Coho salmon data with superimposed median recruitment curves from individual maximum likelihood fits of the logistic hockey stick model with smoothness parameter $\theta = 0.1$, assuming a lognormal recruitment distribution.

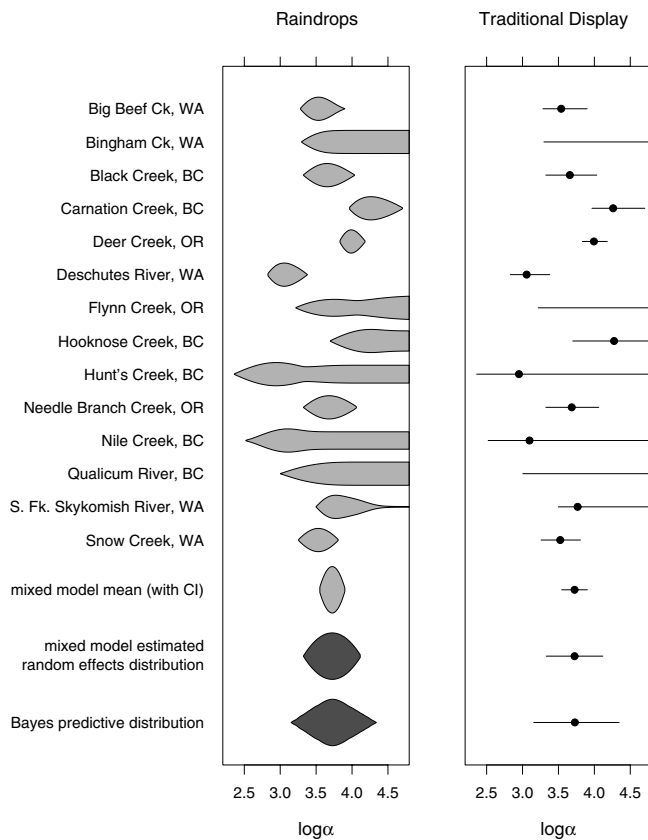


Figure 6. Estimates of the $\log \alpha$ parameter of the logistic hockey stick model with smoothness parameter $\theta = 0.1$ for 14 coho salmon populations showing 95% raindrops (left-hand panel) contrasted with traditional display (right-hand panel). The traditional display shows point estimates and approximate 95% confidence intervals derived from the profile likelihood. (The usual confidence intervals reported by nonlinear regression programs based on normal approximations are entirely inadequate and are not shown here.) At the bottom of the figure several summary estimates are shown, based on the nonlinear mixed effects model of Barrowman et al. (in press). The “mixed model mean” and the “mixed model estimated random effects distribution” were both estimated using the method of Lindstrom and Bates (1990). The “Bayes predictive distribution” was obtained using BUGS (Gilks et al. 1994; Gilks et al. 1996). The corresponding raindrop was based on a kernel density estimator applied to a sample of size 100,000 from the predictive distribution; the interval shown in the right-hand panel was obtained from the 2.5th and 97.5th percentiles of this sample.

where $L_{\theta}(\log \alpha, \log \mu)$ is the joint likelihood for $\log \alpha$ and $\log \mu$ given a fixed value of θ . Because the joint likelihood can be irregularly shaped, the maximization was carried out using a grid search algorithm (Lerman 1980). Raindrops were then formed based on the profile likelihood (Figure 6). Bates and Watts (1988) suggested several related profiling techniques for parameters of nonlinear regression models, which are useful for individual datasets.

To obtain summary estimates across the coho salmon populations, Barrowman et al. (in press) proposed a nonlinear mixed effects model assuming that $\log \alpha$ and $\log \mu$ are normally distributed population-specific random effects, and that conditional on the values of $\log \alpha$ and $\log \mu$ for each population, the values of $\log R$ are iid normal. Models were fitted using the method of Lindstrom and Bates (1990) and also using a Bayesian approach with vague priors for the means and variances of the random ef-

fects, implemented using Markov chain Monte Carlo methods with BUGS software (Gilks et al. 1994; Gilks et al. 1996).

It is useful to contrast the information provided by the raindrops (left-hand panel of Figure 6) with that of the traditional display (right-hand panel). For some populations, such as Needle Branch Creek, the log profile likelihood is very nearly quadratic, and the traditional display is sufficient. For other populations, such as Bingham Creek, $\log \alpha$ does not have a finite maximum likelihood estimate and the traditional display provides a poor summary. Several populations show local maxima: Flynn Creek, for example, shows a local maximum near $\log \alpha = 3.7$ but a slightly more plausible maximum at $\log \alpha = \infty$. Asymmetric confidence intervals are evident in the traditional display, but in some cases the raindrops provide much more information. For example, in the traditional display, the confidence intervals for Nile Creek and South Fork Skykomish River show the same type of asymmetry but the raindrops indicate that the likelihoods are very different. The likelihood for Nile Creek exhibits a local maximum near $\log \alpha = 3$ and then with increasing $\log \alpha$ gently declines to a plateau of slightly less likely values. The likelihood for South Fork Skykomish River exhibits a maximum near $\log \alpha = 3.7$ and then declines to a much greater extent with increasing $\log \alpha$. We conclude that large values of α are more plausible in the case of Nile Creek than in the case of South Fork Skykomish River.

Summary estimates from frequentist and Bayesian mixed model analyses are displayed at the bottom of the figure. Because commonly used software for fitting nonlinear mixed models implicitly treats the estimated mean as being normally distributed, the “mixed model mean (with CI)” is represented by a quadratic raindrop based on the point estimate and standard error obtained using the method of Lindstrom and Bates (1990). The “mixed model estimated random effects distribution” is normal by assumption and so it is also represented by a quadratic raindrop. Apart from emphasizing that not all values are equally likely, these two raindrops provide no more information than the traditional display. However, they do provide a reference to compare with the preceding individual population raindrops. Considered in the light of the summary estimates, the local maxima evident in the raindrops for Flynn Creek, Hooknose Creek, Hunt’s Creek, and Nile Creek are consistent with the mixed model estimate of a mean $\log \alpha$ of approximately 3.7. Finally, the HDR raindrop for the “Bayes predictive distribution” clearly shows nonnormality: the diamond shape of the raindrop indicates greater kurtosis than a normal distribution. The raindrop resembles that of a t distribution with three degrees of freedom (Figure 3). This information is not revealed by the traditional display.

5. DISCUSSION

Effectively displaying collection of likelihoods or distributions is challenging. Commonly a collection of superimposed curves is shown (Efron 1996; van Houwelingen and Zwinderman 1993). However, judgments involving superimposed curves are difficult (Cleveland 1985, p. 271) and if there are more than, say five studies involved, the result is hard to decipher. Plotting the curves in separate panels makes it possible to discern individual curves but requires much more space. A more compact solution is the traditional display of point estimates with confidence intervals, which is sufficient if the log-likelihood or

log density is close to being quadratic. But with small sample sizes, nonlinear models, and Bayesian procedures, marked departures from normality can occur, and the traditional display may be uninformative or even misleading. In such cases, a second dimension is required: raindrop plots use this dimension to effectively communicate the relative plausibility of different parameter values. The result is a compact and informative way of displaying the groups of likelihoods or distributions.

The raindrop plot is based on reflecting the log-likelihood; however, alternative approaches are possible using, for example, the likelihood or the p value function (Miettinen 1985). One advantage of using the log-likelihood is that the log scale permits a tighter display vertically, so that many raindrops can be displayed on a single page. Also, the raindrop's simple shape under normality—a reflected parabola—facilitates detection of nonnormality. Like Q-Q plots (Wilk and Gnanadesikan 1968), raindrops can be used to detect departures from normality.

A related display, the HDR-raindrop, can be used to show the features of estimated random effects distributions, posterior distributions, and predictive distributions. As complex Bayesian analyses have become feasible with methods such as Markov chain Monte Carlo, the need to display collections of posterior and predictive distributions is growing. HDR-raindrops are related to the violin plots of Hintze and Nelson (1998), although the violin plot was designed for investigating data samples rather than distributions per se and uses the density instead of the log density.

The visual impression conveyed by a raindrop plot is determined in part by the area enclosed by each raindrop. It should be noted that this area does not correspond to a likelihood, nor does the area enclosed by an HDR-raindrop correspond to a probability. As with confidence intervals, the width of a raindrop indicates the range of parameter values consistent with the observed data under a given model. The variation in the height of the raindrop shows the relative plausibility of parameter values within this range. The area of a raindrop thus relates to the uncertainty associated with an estimate. For HDR-raindrops, the width indicates the range over which the bulk of the distribution has its support. The variation in the height of the HDR-raindrop provides information on other properties of the distribution, such as skewness and kurtosis (Figure 3). The area of an HDR-raindrop thus relates to the shape and scale of a distribution.

We have found that the raindrop plot is particularly useful for Bayesian analysis because posterior distributions are rarely normal. In the simple case of estimating the mean μ given a sample from a normal distribution with unknown variance σ^2 , if a noninformative prior $p(\mu, \sigma^2) \propto 1/\sigma^2$ is used, the posterior distribution is a shifted, scaled t distribution (Box and Tiao 1973). For more complex estimates, other nonnormal posteriors can result. Boxplot-type displays are often used to show the median, quartiles, and upper percentiles of posterior distributions. We prefer HDR-raindrops since they provide more information about the shape of the distribution.

The raindrop plot provides a compromise between displaying the full set of likelihoods or distributions at one extreme and simply displaying point estimates with error bars at the other extreme. We believe that raindrop plots will be a useful addition to the data analyst's toolbox.

APPENDIX

The following code runs in S-Plus and R and draws a quadratic 95% raindrop:

```
theta <- seq(-3, 3, length=500)
loglik <- -0.5*theta^2
cutoff <- -1.92
loglik <- loglik - max(loglik)
select <- loglik > cutoff
THETA <- theta[select]
LOGLIK <- loglik[select]
thickness <- 1-LOGLIK/cutoff
plot(0, xlim=range(theta), ylim=c(-1, 1),
     type="n", xlab="theta", ylab="")
polygon(c(THETA, rev(THETA)), c(-thickness,
                                rev(thickness)), col=2)
```

[Received June 2003. Revised July 2003.]

REFERENCES

- Agresti, A. (1990), *Categorical Data Analysis*, New York: Wiley.
- Balanda, K. P., and MacGillivray, H. L. (1988), "Kurtosis: A Critical Review," *The American Statistician*, 42, 111–119.
- Barrowman, N. J., and Myers, R. A. (2000), "Still More Spawner-Recruitment Curves: The Hockey Stick and its Generalizations," *Canadian Journal of Fisheries and Aquatic Sciences*, 57, 665–676.
- Barrowman, N. J., Myers, R. A., Hilborn, R., Kehler, D. G., and Field, C. A. (in press), "The Variability Among Populations of Coho Salmon in the Maximum Reproductive Rate and Depensation," *Ecological Applications*.
- Bates, D. M., and Watts, D. G. (1988), *Nonlinear Regression Analysis and its Applications*, New York: Wiley.
- Box, G. E. P., and Tiao, G. C. (1973), *Bayesian Inference in Statistical Analysis*, New York: Wiley.
- Cleveland, W. S. (1985), *The Elements of Graphing Data*, Monterey, CA: Wadsworth.
- Cox, D. R., and Hinkley, D. V. (1974), *Theoretical Statistics*, London: Chapman & Hall.
- Efron, B. (1996), "Empirical Bayes Methods for Combining Likelihoods," *Journal of the American Statistical Association*, 91, 538–565.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (eds.) (1996), *Markov Chain Monte Carlo in Practice*, London: Chapman & Hall.
- Gilks, W. R., Thomas, A., and Spiegelhalter, D. J. (1994), "A Language and Program for Complex Bayesian Modelling," *The Statistician*, 43, 169–178.
- Hintze, J. L., and Nelson, R. D. (1998), "Violin Plots: A Box Plot-Density Trace Synergism," *The American Statistician*, 52, 181–184.
- Hyndman, R. J. (1996), "Computing and Graphing Highest Density Regions," *The American Statistician*, 50, 120–126.
- Lee, J. J., and Tu, Z. N. (1997), "A Versatile One-Dimensional Distribution Plot: the BLiP Plot," *The American Statistician*, 51, 353–358.
- Lerman, P. M. (1980), "Fitting Segmented Regression Models by Grid Search," *Applied Statistics*, 29, 77–84.
- Lewis, S., and Clarke, M. (2001), "Forest Plots: Trying to See the Wood and the Trees," *British Medical Journal*, 322, 1479–1480.
- Lindstrom, M. J., and Bates, D. M. (1990), "Nonlinear Mixed Effects Models for Repeated Measures Data," *Biometrics*, 46, 673–687.
- Miettinen, O. S. (1985), *Theoretical Epidemiology*, New York: Wiley.
- Morris, C. N. (1996), Comment on "Empirical Bayes Methods for Combining Likelihoods," *Journal of the American Statistical Association*, 91, 555–558.
- Royall, R. M. (1997), *Statistical Evidence: A Likelihood Paradigm*, London: Chapman and Hall.
- Sacks, H. S., Chalmers, T. C., Blum, A. L., Berrier, J., and Pagano, D. (1990), "Endoscopic Hemostasis, An Effective Therapy for Bleeding Peptic Ulcers," *Journal of the American Medical Association*, 264, 494–499.
- van Houwelingen, H. C., and Zwinderman, K. H. (1993), "A Bivariate Approach to Meta-analysis," *Statistics in Medicine*, 12, 2273–2284.
- Wilk, M. B., and Gnanadesikan, R. (1968), "Probability Plotting Methods for the Analysis of Data," *Biometrika*, 55, 1–17.