(2) provides an approximate formula that can be used to estimate the standard error of $\hat{\beta}$. As an example, we consider a series of 1632 average monthly temperatures over the Northern Hemisphere (land and sea) used by the IPCC (Intergovernmental Panel on Climate Change) for its global warming analyses. Figure 3a presents the data, smoothed by applying a 49-month moving average and centered around the 1950–1979 mean value for each month. It shows a pattern of steady behavior until about 1920, followed by a sharp rise between 1920 and 1940, then a gradual decrease until about 1975, followed by the sharp rise that has triggered the present alarm about global warming. Over the whole series, there is a clear rise in temperature, but whether it is due to the greenhouse effect is a matter of intense debate among climate scientists.

A linear trend was fitted to this data series (unsmoothed) and resulted in a estimated trend of 0.40°C per century, a figure consistent with several other estimates of global warming over the last century and a half. The first 120 periodogram ordinates of the residuals are plotted on log-log scales in Figure 3b. The pattern is quite similar to the two series quoted by Beran, and again seems to show evidence of long-range dependence. This is confirmed by the estimates $H = 0.90$ with standard error 0.05, based on $n_0 = 1$, $n_1 = 120$; also $\hat{b} = 0.0033$. When these figures are inserted into (2) (adjusted for the unit of trend) the standard error of the estimated trend is around 0.1, which is again consistent with earlier estimates of standard error including those quoted by Bloomfield (1992). My main doubt about this conclusion is whether the series can really be assumed stationary, given the obvious inconsistencies in methods of measurement over the last century and a half, but this would take us into other aspects beyond the scope of the present discussion.

I believe the message of all three examples is that the concept of long-range dependence must be taken seriously. At the same time, exactly how these examples are to be interpreted could be a matter of considerable debate. Jan Beran is to be congratulated on his very clear and comprehensive review, and I hope it will act as a springboard for much further research in this area.

ACKNOWLEDGMENTS

Example 3 is based on data provided by Dr. P. Jones of the Climate Research Unit, University of East Anglia. I thank Jan Beran for providing me with the Nile and NBS data, and Peter Bloomfield for copies of his references. This work was supported in part by NSF Contract DMS-91-15750.

Rejoinder

Jan Beran

I would like to thank the discussants for their stimulating comments and valuable suggestions. Their comments emphasize once more that long memory is an important issue to anybody who uses statistical inference, since it occurs rather frequently in real data and strongly influences the validity (and power) of standard tests and confidence intervals. Particularly interesting are the data examples analyzed by Smith (global warming – climatological data), Haslett and Raftery (wind speed – meteorological data) and Dempster and Hwang (employment series – economic data), since these are examples that concern everyone (and not just a selected group of scientists). Parzen summarizes the main message of the paper very clearly by saying that in data analysis, we always have to decide whether the data (either the original measurements or residuals, e.g., after subtracting a regression function) are white noise, a short-memory process or a long-memory process. The same view is expressed in a more general context by Mosteller and Tukey (1977, p. 119 ff): "even in dealing with so simple a statistic as the arithmetic mean, it is often vital to use as direct an assessment of its internal uncertainty as possible. Obtaining a valid measure of uncertainty is not just a matter of looking up a formula." In other words, no formula should be applied without checking its approximate validity. Naturally, this does not only refer to "classical" formulas, such as $\mathrm{var}(\overline{X}) = \sigma^2 n^{-1}$, but also to the "new" formulas, such as $\mathrm{var}(\overline{X}) = L(n)n^{2H-2}$ $(0 < H < 1)$, given in the present review paper.

One major reason why the question of long memory is usually not dealt with in daily statistical practice is the lack of statistical software packages. Haslett and Raftery's program (and its implementation in the next release of SPLUS) is therefore a welcome contribution. As already mentioned briefly after formula (12) and