DISCUSSION OF: A STATISTICAL ANALYSIS OF MULTIPLE TEMPERATURE PROXIES: ARE RECONSTRUCTIONS OF SURFACE TEMPERATURES OVER THE LAST 1000 YEARS RELIABLE?

BY JONATHAN ROUGIER

University of Bristol

The authors are to be congratulated on the clarity of their paper, which gives discussants and readers much to sink their teeth into. My comments are somewhat critical, but this should in no way devalue this paper as an important contribution to the ongoing debate concerning the information about historical climates that is recoverable from proxies. Figure 14, in particular, provides much food for thought.

In Section 3.2, comparing the proxy-based reconstruction of climate to measures based on actual climate (in-sample mean and ARMA model) is not very helpful for assessing the performance of the proxy—in fact, it confirms information already presented about the nature of the climate process and the relative variability of the proxies. This distracts from the more pertinent finding in Section 3.3 that the proxy-based reconstruction seems to perform no better than various random proxies. Again, though, this result is not necessarily detrimental to the proxy. If one generates 1138 random sequences of length 149 with roughly the right time-series properties, one should not be surprised to find that a 1139th sequence is near the span of a small subset, and it is a testament to the Lasso procedure that it seems to be doing a good job at picking this subset out. Hold-outs at the end of the calibration period would provide a more powerful test; for hold-outs in the middle, one can be fairly confident that if the Lasso finds a match at both ends, then the middle will fit reasonably well. In Section 3.5, the finding that large numbers of pseudo-proxies are selected can be explained in the same way. Moreover, the Lasso procedure will have a bias against selecting actual proxies, if they are correlated with each other. Overall, I do not think that Section 3 presents evidence against the proxies.

I am bemused by Section 5. First, let us be very clear that this is not a "fully Bayesian" analysis. What we have here is a normalised likelihood function over β and σ masquerading as a posterior distribution, in order to implement a sampling procedure over the model parameters. This seems a perfectly reasonable ad-hockery [although a Normal Inverse Gamma conjugate analysis would be more conventional; see O'Hagan and Forster (2004), Chapter 11], but to call it "fully

---

Received September 2010.