

FORMAL STATISTICAL MODELS FOR ESTIMATING RADIOCARBON CALIBRATION CURVES

C E Buck¹ • P G Blackwell

Department of Probability & Statistics, University of Sheffield, United Kingdom.

ABSTRACT. We report on the development and implementation of a model-based statistical method for the estimation of radiocarbon calibration curves using diverse data. The method takes account of uncertainty on both the ¹⁴C and calendar scales, coherently integrating data, the calendar age estimates of which arise from different dating methods. It also allows for correlation between observations, if they have particular sources of uncertainty in common. We adopt an approach based on a random walk model, tailoring it to take account of possible calendar age offsets between different data sources by adding a random effect component. The latter allows us to use the same modeling framework for constructing the new calibration curve IntCal04, the comparison curve NotCal04, the Southern Hemisphere curve SHCal04, and the marine calibration curve Marine04.

1. INTRODUCTION

When the IntCal98 group put together the data for the 1998 internationally agreed radiocarbon calibration curve, they saw their primary task as collecting high-quality data. The task of constructing a calibration curve from these data was then undertaken using relatively simple data averaging methods based on those of Ward and Wilson (1978). Since 1998, however, the statistical community has shown interest in estimation of the underlying ¹⁴C calibration curve from raw data, and it has become clear that the methods adopted in 1998 ignore several important features of the data, and hence of the problem that we are seeking to solve. As a result, in this paper we outline an approach to curve construction that we have devised for IntCal04. The work we report here has benefited greatly from a considerable body of literature on Gaussian process modeling within the Bayesian statistics research community, but in particular from the work of Gómez Portugal Aguilar et al. (2002), Christen and Nicholls (2000), and from detailed discussions with Andrew Millard in the Department of Archaeology at the University of Durham to whom we are greatly indebted.

In seeking methods for curve construction for IntCal04, we were concerned with several specific aspects of the problem that had not been handled satisfactorily in the past. In particular, we wanted to be able to take account of the fact that:

1. Most samples in the IntCal database do not derive their carbon from a single year of metabolism;
2. Many samples in the IntCal database have calendar age estimates that derive from methods other than dendrochronology (e.g. U/Th dating) and are, hence, not precisely known (but have measures of uncertainty associated with them);
3. Not all samples in the IntCal database have independent calendar age estimates—some, for example, derive their calendar age estimates from a wiggle match, and hence their calendar age estimate is related to all the others in the same sequence;
4. The ¹⁴C calibration curve is intended to be our best estimate of the relationship over time between ¹⁴C age and calendar age and, since this relationship is a continuously varying one, observations from it are correlated.

In the past, all four of these factors (and several other more subtle ones discussed below) have been ignored both in the construction of internationally agreed calibration curves and in the methodological literature cited above. Knox and McFadgen (1997) do address the issue of the averaging of

¹Corresponding author. Email: c.e.buck@sheffield.ac.uk.

blocks of tree rings, as part of an approach to calibration curve construction based on Fourier analysis, but their method relies on constant block size and equally spaced observations.

Here, we address all of these issues, by formally modeling the relationship between the data and the calibration curve.

2. THE UNDERLYING MODEL

2.1 Simplest Case

In the simplest case, we have a single ^{14}C determination X (BP) with known calendar date θ (cal BP); we then typically assume that X is given by the true value of the calibration curve at date θ , written $\mu(\theta)$, plus a term ε which has a Gaussian distribution with variance σ^2 (and mean zero), giving

$$X = \mu(\theta) + \varepsilon,$$

$$\varepsilon \sim N(0, \sigma^2).$$

To learn about the curve for a particular date, θ^* say, we do not use only those observations with $\theta = \theta^*$, because we believe that for any θ near θ^* , knowing $\mu(\theta)$ would tell us something about $\mu(\theta^*)$. To formalize this, we express our prior beliefs about the relationships between different points on the curve in terms of a random walk. We represent our beliefs about changes in the curve from one year to the next by a Gaussian distribution, with mean or “drift” β and variance (per year) r^2 . So we have the following:

$$\mu(\theta+1) = \mu(\theta) + Z_{\theta+1},$$

$$Z_{\theta+1} \sim N(\beta, r^2),$$

or

$$\mu(\theta+1) \sim N[\mu(\theta) + \beta, r^2],$$

and more generally,

$$\mu(\theta+\delta) \sim N[\mu(\theta) + \delta\beta, |\delta|r^2].$$

If we wish to evaluate the curve at a date θ^* , then we have

$$X \sim N[\mu(\theta^*) + (\theta - \theta^*)\beta, |\theta - \theta^*| r^2 + \sigma^2].$$

From now on, we will imagine that we are interested in a single such point at a time; the reasoning behind this, and its advantages and disadvantages, are discussed in Section 4.3.

If we have multiple observations relevant to $\mu(\theta^*)$, as we always do in practice, then we can write them as X_1, \dots, X_N , with

$$X_j \sim N[\mu(\theta^*) + (\theta_j - \theta^*)\beta, |\theta_j - \theta^*| r^2 + \sigma_j^2] \quad (1).$$

However, we can *not* take the X_j s to be independent, since the component of our uncertainty that derives from the random walk is common to multiple observations.

To allow for this appropriately, we need to write X_j explicitly in terms of the “steps” of the random walk. If $\theta_j > \theta^*$, we have

$$\begin{aligned} X_j &= \mu(\theta_j) + \varepsilon_j \\ &= \mu(\theta^*) + Z_{\theta^*+1} + \dots + Z_{\theta_j} + \varepsilon_j \\ &= \mu(\theta^*) + \sum_{\theta=\theta^*+1}^{\theta_j} Z_{\theta} + \varepsilon_j, \end{aligned}$$

with a corresponding expression if $\theta_j < \theta^*$. It is then clear that some of the Z_{θ} s may be common to different X_j s.

This gives the covariance between two observations X_i and X_j , say, as

$$\text{Cov}(X_i, X_j) = r^2(\min\{\theta_i, \theta_j\} - \theta^*) \tag{2}$$

if $\theta_i, \theta_j > \theta^*$, since $\min\{\theta_i, \theta_j\} - \theta^*$ is the number of Z terms in common, with a similar expression if $\theta_i, \theta_j < \theta^*$ and zero covariance otherwise.

Finally, we can represent the uncertainty in all the X_j s, and the dependence between them, by relating all of them to a collection of Z_{θ} s. In practice, we use only the potentially relevant X_j s, namely those associated with calendar dates in some “window” $(\theta_{\min}, \theta_{\max})$ around θ^* . Writing

$$\mathbf{X} = \begin{bmatrix} X_1 \\ \dots \\ X_N \end{bmatrix}, \quad \boldsymbol{\varepsilon} = \begin{bmatrix} \varepsilon_1 \\ \dots \\ \varepsilon_N \end{bmatrix}, \quad \mathbf{M} = \begin{bmatrix} \mu(\theta_{\min}) \\ \dots \\ \mu(\theta^*) \\ \dots \\ \mu(\theta_{\max}) \end{bmatrix}, \quad \mathbf{Z} = \begin{bmatrix} -Z_{\theta_{\min}+1} \\ \dots \\ -Z_{\theta^*} \\ Z_{\theta^*+1} \\ \dots \\ Z_{\theta_{\max}} \end{bmatrix},$$

we have

$$\mathbf{M} = \mathbf{BZ},$$

where the matrix B is of the form

$$\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix},$$

extended to cover the range $(\theta_{\min}, \theta_{\max})$, and so

$$\mathbf{X} = \mathbf{ABZ} + \boldsymbol{\varepsilon},$$

where the matrix A is defined by $a_{j,\theta} = 1$ if $\theta = \theta_j$ and $a_{j,\theta} = 0$ otherwise. Since the Z_{θ} s are independent, this matrix representation allows us to readily calculate covariances between all the X_j s.

Our approach for this simplest case is related to the approaches of Christen and Nicholls (2000) and Gómez Portugal Aguilar et al. (2002), but does not in itself model any part of the calibration curve in sufficient detail; for that, we need to allow a more general structure for our observations.

2.2 Averaging Over Blocks

Often, a single determination X_j arises not from material that dates to a single year, but from an average over n_j successive years. In principle, this could be a weighted average (allowing say for different amounts of datable material from different tree rings), but since our main applications—Hughen et al., McCormac et al., Reimer et al., and van der Plicht et al. (this issue)—will assume equal weights in such cases, we will do the same here to avoid unduly complex notation. We can write

$$X_j \sim N(m_j, \sigma_j^2),$$

where

$$m_j = [\mu(\theta_j) + \mu(\theta_j+1) + \dots + \mu(\theta_j+n_j-1)] / n_j.$$

More generally, we can extend the matrix formulation above by redefining A so that $a_{j,\theta} = n_j^{-1}$ if $\theta_j \leq \theta \leq \theta_j + n_j - 1$ and $a_{j,\theta} = 0$ otherwise. Again, the covariances between all the X_j s can be obtained immediately in this way.

This structure enables us to estimate those parts of calibration curves that involve only negligible error on the calendar scale, essentially those using tree-ring data only. This includes the most recent parts of IntCal04 (Reimer et al., this issue), to 12.4 cal kyr BP, and the part of SHCal04 based purely on Southern Hemisphere data (McCormac et al., this issue).

2.3 Uncertainty in Calendar Dates

In practice, we sometimes do not observe each θ_j directly. We may instead observe T_j where

$$T_j \sim N(\theta_j, \tau_j^2).$$

In the absence of any prior information about θ_j , this implies

$$\theta_j \sim N(T_j, \tau_j^2).$$

Note that if we extend the possible values of θ_j from exact years to any possible value, then the random walk defined above can be thought of as a Gaussian process or diffusion, and Equations 1 and 2 still hold. Then X_j has mean $\mu(\theta^*)$ and variance

$$r^2 E[(\theta_j - \theta^* | T_j)] + \sigma_j^2,$$

which can be calculated analytically when $\theta_j \sim N(T_j, \tau_j^2)$.

The corresponding covariance in the calendar age errors for observations X_i and X_j involves

$$E[c(\theta_i - \theta^*, \theta_j - \theta^*) | T_i, T_j],$$

where

$$c(\phi, \psi) = \begin{cases} \min\{|\phi|, |\psi|\} & \text{if } \phi, \psi \text{ have the same sign,} \\ 0 & \text{otherwise.} \end{cases}$$

This term cannot be calculated analytically, but must be obtained by numerical integration. In many cases, it is appropriate to regard the calendar age errors in θ_i and θ_j as independent, e.g. if the observations relate to completely separate data sources, or if they are directly dated using U/Th dating. More generally, the uncertainties in θ_i and θ_j may be linked, e.g. if they have some layer-counting or sequence-matching error in common. Then we take

$$T_i, T_j \mid \theta_i, \theta_j$$

to be bivariate normal and assume a uniform prior, as before; numerical integration can still be used to evaluate $\text{Cov}(X_i, X_j)$. For example, if X_i and X_j are in the same “floating” sequence of observations, assigned absolute dates by wiggle matching, then any wiggle-matching error, ω , will be common to the two, and so we might have

$$T_i = \theta_i + \omega,$$

$$T_j = \theta_j + \omega,$$

$$\omega \sim N(0, \tau^2),$$

if there is no other calendar age error. If, on the other hand, T_i and T_j are determined by layer counting within the same varved sequence, and $T_i < T_j$, then any error accrued up to layer T_i will also affect T_j , and we might have

$$T_i = \theta_i + \omega_i,$$

$$T_j = \theta_j + \omega_i + \omega_j,$$

where both ω_i and ω_j have Poisson distributions that, like the ¹⁴C counting errors, can be approximated by normal distributions for present purposes.

As a final example, in addition to having calendar date estimates for individual samples, in some cases we also have prior information that observations are ordered in time. For example, because of the physical deposition process, the speleothems giving rise to data used here have annual laminae whose ordering should be respected in the statistical analysis. Enforcing a strict ordering is not easy to do within this framework (although it would be straightforward in the more general approach described in Section 5.1). Instead, in this case we specify a high positive correlation between the errors on the calendar dates, so that if, e.g., one layer is much older than estimated, the same is very likely to be true for nearby layers.

The approach described in this section enables us to allow in a coherent way for calendar errors on observations, as is needed for the older part of IntCal04 (Reimer et al., this issue), beyond 12.4 cal kyr BP.

2.4 Offsets Between Curves

Looking at Lake Suigetsu and the Bahamas speleothems (and others) which form part of the NotCal data (discussed in detail in van der Plicht et al., this issue), it is clear that there are real differences between the individual “comparison” data sets; we can think of each source having its own offset from the overall, average curve. Looking at plots of the data (given in van der Plicht et al., this

issue), it is clear that such an offset can be large and does change over time, although short-term changes tend to be small. Similarly, in constructing the Southern Hemisphere calibration curve, SHCal (McCormac et al., this issue), it is clear that an offset exists between the curves for the Northern and Southern Hemispheres, and that allowance must be made for it if we are to use Northern Hemisphere data to help learn about the Southern Hemisphere calibration curve

To take account of such offsets, we can extend our model as follows. We denote the offset for the k th source of data at date θ by $v_k(\theta)$. At any given date of interest θ^* , $v_k(\theta^*)$ can be thought of as random; we take its prior distribution to have the form

$$v_k(\theta^*) \sim N(\gamma_k, W_k^2)$$

for all θ^* . For any given k , we assume that $v_k(\theta)$ changes only slowly compared with changes in $\mu(\theta)$, as is clear from casual inspection of the data. We represent $v_k(\theta)$ as a random walk² of the same form as $\mu(\theta)$ with steps having variance u_k^2 .

An observation X_j from source k will then have mean

$$\mu(\theta^*) + (\theta_j - \theta^*)\beta + \gamma - k,$$

and variance

$$|\theta_j - \theta^*| (r^2 + u_k^2) + W_k^2 + \sigma_j^2 \quad (3)$$

(cf. Equation 1). If two observations, X_i and X_j say, are *both* from source k , the covariance between them will be

$$\text{Cov}(X_i, X_j) = W_k^2 + (r^2 + u_k^2)(\min\{\theta_i, \theta_j\} - \theta^*) \quad (4)$$

if $\theta_i, \theta_j > \theta^*$ with a similar expression if $\theta_i, \theta_j < \theta^*$; if they are from the same data set and $\theta_i < \theta^* < \theta_j$ or $\theta_j < \theta^* < \theta_i$, their covariance is simply W_k^2 .

The ways in which this framework is applied in NotCal and SHCal are outlined in Section 4.2.

3. IMPLEMENTATION

Given the model described above, it is in principle straightforward to take a collection of observations, with knowledge of the size and type of errors both in calendar age and ¹⁴C date, and produce estimates and standard errors for points on the calibration curve.

For each date θ^* at which we wish to estimate the curve, we select those observations X_1, \dots, X_N associated with calendar dates in a suitable window ($\theta_{\min}, \theta_{\max}$) around θ^* . Then X_1, \dots, X_N given $\mu(\theta^*)$ have a multivariate normal distribution with all means equal to $\mu(\theta^*)$, and with covariance given by the calculations above. We use a uniform prior for $\mu(\theta^*)$; it is then straightforward to invert the covariance matrix numerically and obtain a posterior estimate for $\mu(\theta^*)$ and an associated standard deviation. Results of this form, based on different data sets with different error structures, are given in Hughen et al., McCormac et al., Reimer et al., and van der Plicht et al. (this issue).

²Strictly speaking, since we take the typical size of these offsets to be constant over time, the process $v(\cdot)$ should have some tendency to revert to zero, and should be modeled by, e.g., a stationary autoregressive process. However, provided the rate of change of the offset is small ($u \ll W$), the random walk is a reasonable approximation to such a process.

In practice, this relies on some decisions that have to be made about details of the implementation, and these are discussed in Section 4.

4. PRACTICALITIES

4.1 Specifying the Random Walk

The implementation outlined Section 3 depends crucially on how the random walk of Section 2.1 is specified. For the purposes of IntCal, it was felt vital to have an explicit, repeatable numerical method (as opposed to a simulation-based approach), and, in practice, that means assuming particular values for the mean β and variance (per year) r^2 .

It is natural to take a value of 1 for β , since we would expect the calibration curve to change by approximately 1 (calibrated) year per (calendar) year. In fact, our experimentation shows that the results are very insensitive to the specific value used for β ; a more precise value based on current knowledge of ^{14}C decay rates makes no difference to the curve obtained.

It is less clear *a priori* what is a reasonable value for r . Christen and Nicholls (2000) estimate a value of $r = 7.9$ on the basis of single-year data from Stuiver et al. (1998) by using a combination of a full MCMC implementation (see Section 5.1) of a random walk model and by repeated analyses of 100-yr blocks of the same data. In the same way, we looked in detail at particular high-resolution subsets of the IntCal04 data. For the results in Hughen et al., McCormac et al., Reimer et al., and van der Plicht et al. (this issue), we therefore take $r = 8$.

4.2 Describing the Offsets

In a similar way to fixing r in Section 4.1, for NotCal and SHCal we also need to specify the parameters defining the random offsets in Section 2.4 and the rate at which they may change over time. Formal estimation of these quantities is difficult without the full simulation-based framework mentioned in Section 5.1; instead, we experimented with repeated analyses of the data to determine which values were consistent with the data.

For NotCal, we had little prior information about the offsets; for all non-IntCal data sources, offsets were taken to be centered at zero (“random effects” in the usual terminology of the statistical literature) but with some large variance. We took $\gamma_k = 0$, $W_k = 1000$, and $u_k = 10$ for the non-IntCal data sources (and $\gamma_k = W_k = u_k = 0$ for those sources meeting the IntCal criteria). Results are very insensitive to the exact values used.

The offset between SHCal and IntCal is better understood, and we expect its direction and approximate magnitude to be fairly consistent. We need to allow for the offset to extend the Southern Hemisphere curve back beyond 990 cal BP, the limit of the Southern Hemisphere data; our prior for that analysis comes from considering the separate IntCal and SHCal curves back to that point. The offset between them varies gradually over time, and its value at 990 cal BP (and hence our “best guess” at its value at any earlier date) is 55 yr. Based on this, and the variability of the offset between 50 cal BP and 990 cal BP, we take $\gamma_k = 55$, $W_k = 25$, and $u_k = 1$ for the Northern Hemisphere data used, and of course, $\gamma_k = W_k = u_k = 0$ for the Southern Hemisphere data themselves.

4.3 Why One Point at a Time?

As mentioned in Section 2.1, we choose to organize our calculations by treating separately each desired point on the curve. The advantage of this approach is that it enables us to limit the observations that have to be considered in any one calculation to a reasonable number, typically much

smaller than the size of the whole data set. Obviously, this involves some element of approximation, but for each point on the curve we choose a suitable “window” of data points to use, so that the effect of excluding the remaining points is negligible (see Section 4.4).

Calculating points on the calibration curve separately has the theoretical disadvantage that it does not give the posterior covariances between the estimated values. However, given the form in which the calibration curves are generally reported and used, these covariances are not likely to be used routinely even if available.

4.4 Choosing the Window

To calculate each $\mu(\theta^*)$ for the IntCal curve, we used the 100 observations X_j with values T_j (or θ_j if known) that were nearest to θ^* . Experimentation showed that this was sufficient to calculate the estimate of $\mu(\theta^*)$ accurately. The exception was a small part of the curve near $\theta^* = 15,500$, where the presence of a large number of observations with large errors necessitated using more points; 200 points proved sufficient.

For the NotCal and SHCal curves, more care was needed in selecting those data points to be included in each calculation, because of the widely varying levels of uncertainty of the observations and the offsets between data from different sources. As far as possible, for the calculation of each point we used equal numbers of observations from each source. At the same time, as far as possible within each source, equal numbers from before and after θ^* were selected. In undertaking this data selection, we regarded all IntCal data as coming from a single source when using them in NotCal. In selecting observations for calculation in SHCal, all Southern Hemisphere observations were treated as a single source; IntCal observations were treated as 2 distinct sources, one for observations before 990 cal BP (the limit of the Southern Hemisphere data) and one for observations after 990 cal BP, to ensure sufficient overlap between Northern and Southern Hemisphere data for stable estimation of the offset. In both these cases, 150 observations were used in estimating each point on the curve.

4.5 Choosing Output Values

To actually generate a calibration curve, a decision has to be made as to which specific points on the curve to estimate, i.e. the values of θ^* to use. We chose to calculate the most recent part of the curve at 5-yr intervals, which seems to be the most detailed level justified by the data and which is compatible with the 10-yr intervals used in IntCal98. Older parts of the curve are calculated at sparser points, in line with the decreasing density (and increasing calendar error) of the data.

4.6 Numerical Integration

Evaluation of covariances between observations, in the presence of calendar age error (Section 2.3), requires calculation of the expectation of

$$c(\phi, \psi) = \begin{cases} \min\{|\phi|, |\psi|\} & \text{if } \phi, \psi \text{ have the same sign,} \\ 0 & \text{otherwise} \end{cases}$$

when ϕ and ψ have a bivariate normal distribution. This is carried out by numerical integration, evaluating $c(\cdot, \cdot)$ at a grid of points centered at the mean of the distribution of ϕ and ψ . A 25×25 grid gives a sufficiently accurate answer for present purposes.

5. DISCUSSION

5.1 Alternative Implementations

It has been pointed out (A Millard, personal communication) that some of the modeling in Section 2 fits within the framework of Dynamic Linear Modeling (see e.g. West and Harrison 1997). This would have the advantage of allowing us to estimate multiple points on the curve simultaneously, and obtain the appropriate covariances; however, it does not readily extend to incorporating the errors in calendar dates necessary beyond 12 kyr.

As mentioned in Section 4.2, a more flexible approach to fitting the models of Section 2 would be to use a simulation-based Markov chain Monte Carlo (MCMC) approach. Such an approach would allow the incorporation of uncertainty about r and W (and other parameters if necessary) and would automatically give covariances between points on the curve, as well as estimates and standard deviations for any other aspect of the curve. However, for the specific purpose of producing a standard calibration curve, this approach has the disadvantages that it is harder to describe and that it is inherently harder to reproduce. Thus, for present purposes we have used the numerical approach described in Section 3.

5.2 Extending the Modeling

There are some aspects of the calibration process, taken as given in the methodology described here, that might usefully be incorporated in the process of estimating the calibration curve. One is the “wobble matching” of floating sequences, which should itself be affected by the calibration curve, so the 2 processes should ideally be carried out jointly. For present purposes, in estimating the curve we have taken the “wobble match” and associated error as given. Another similar case is the interpolation of calendar ages within U/Th-dated sequences; again, the process of interpolation and of calibration should ideally be treated as interdependent.

More ambitiously, since *any* sample with both absolute and ^{14}C dating information, however uncertain, tells us *something* about the calibration curve, there is an argument that calibration of individual samples is itself interdependent with the process of constructing the curve.

Finally, there is scope for incorporating prior information about mechanisms thought to be involved in the dynamics of atmospheric ^{14}C . For example, it would be interesting to allow a component of periodic behavior corresponding to solar cycles. The current approach does not do this, since it lies beyond the primary purpose of the standard calibration curves, although it is interesting to note that there are such cycles, with periods of 11 and 25 yr for example, detectable in our estimated curve.

ACKNOWLEDGMENTS

We are grateful to Andrew Millard and Andrés Christen for helpful comments on this work.

Caitlin Buck acknowledges, with gratitude, the support and encouragement of her colleagues during a period of ill health while working on this project. In addition, she wishes to thank the staff of the Access to Work team in Sheffield and the support worker whom they funded, Sammy Rashid, without whose patience and practical and intellectual input she could not have completed her contribution to this work.

REFERENCES

- Christen JA, Nicholls G. 2000. Random-walk radiocarbon calibration. Technical report #457. Auckland: Mathematics Department, University of Auckland.
- Gómez Portugal Aguilar D, Litton CD, O'Hagan A. 2002. Novel statistical model for a piece-wise linear radiocarbon calibration curve. *Radiocarbon* 44(1): 195–212.
- Hughen KA, Baillie MGL, Bard E, Bayliss A, Beck JW, Blackwell PG, Buck CE, Burr GS, Cutler KB, Damon PE, Edwards RL, Fairbanks RG, Friedrich M, Guilderson TP, Herring C, Kromer B, McCormac G, Manning S, Bronk Ramsey C, Reimer PJ, Reimer RW, Remmele S, Southon JR, Stuiver M, Talamo S, Taylor FW, van der Plicht J, Weyhenmeyer CE. 2004. Marine04 marine radiocarbon age calibration, 0–26 cal kyr BP. *Radiocarbon*, this issue.
- Knox FB, McFadgen BG. 1997. Least-squares fitting a smooth curve to radiocarbon calibration data. *Radiocarbon* 39(2):193–204.
- McCormac G, Hogg AG, Blackwell PG, Buck CE, Higham TFG, Reimer PJ. 2004. SHCal04 Southern Hemisphere calibration, 0–11.0 cal kyr BP. *Radiocarbon*, this issue.
- Reimer PJ, Baillie MGL, Bard E, Bayliss A, Beck JW, Blackwell PG, Buck CE, Burr GS, Cutler KB, Damon PE, Edwards RL, Fairbanks RG, Friedrich M, Guilderson TP, Herring C, Hughen KA, Kromer B, McCormac G, Manning S, Bronk Ramsey C, Reimer RW, Remmele S, Southon JR, Stuiver M, Talamo S, Taylor FW, van der Plicht J, Weyhenmeyer CE. 2004. IntCal04 terrestrial radiocarbon age calibration, 0–26 cal kyr BP. *Radiocarbon*, this issue.
- Stuiver M, Reimer PJ, Bard E, Beck JW, Burr GS, Hughen KA, Kromer B, McCormac FG, van der Plicht J, Spurk M. 1998. IntCal98 radiocarbon age calibration, 24,000–0 cal BP. *Radiocarbon* 40(3):1041–83.
- van der Plicht J, Beck JW, Bard E, Baillie MGL, Blackwell PG, Buck CE, Friedrich M, Guilderson TP, Hughen KA, Kromer B, McCormac FG, Bronk Ramsey C, Reimer PJ, Reimer RW, Remmele S, Richards DA, Southon JR, Stuiver M, Weyhenmeyer CE. 2004. NotCal04—comparison/calibration ¹⁴C records 26–50 cal kyr BP. *Radiocarbon*, this issue.
- Ward GK, Wilson SR. 1978. Procedures for comparing and combining radiocarbon age determinations: a critique. *Archaeometry* 20(1):19–31.
- West M, Harrison PJ. 1997. *Bayesian Forecasting and Dynamic Models*. 2nd edition. New York: Springer-Verlag.