MILLENIAL SCALE SEA LEVEL CURVE ESTIMATION

BU-1544 -M **November 2000**

John Staudenmayer
Greg Balco
W. Roland Gehrels
Naomi Altman
Ciprian Crainiceanu
Jing Qui

Keywords: geology, measurement error, segmented regression, maximum likelihood, EM algorithm, smooth transition.

Abstract:
Understanding past relative sea-level (RSL) changes is important to a number of social and scientific questions, including the effects of global climate change, the distribution of coastal environments, and future land-use planning. In particular, accurately characterizing millenial sea-level changes is important in evaluating their relationship to ice-sheet dynamics and long-term climate trends. In this paper, we analyze sea-level data from several Maine salt marshes previously reported in a paper from the geological literature. We address these data with a "smooth transition"/approximate segmented regression model and account for the presence of covariate measurement error. Our model builds on previous methodological work and includes a more flexible model, two alternate computational techniques for estimating the parameters in the model, and asymptotic confidence intervals for our estimates. For these specific data, our model shows that the data from one site are anomalous and therefor warrant additional investigation. This important aspect of the data was not previously identified by a non-statistical analysis. In general, our technique offers a statistically justifiable method for using geological data with measurement errors to estimate actual RSL curves that can be compared to each other, independently determined climate records, or geophysical model results.

BU-1544-M

Millennial Scale Sea Level Curve Estimation

John Staudenmayer*
Department of Biostatistics
Harvard School of Public Health

Naomi Altman
Department of Biometrics
Cornell University

Greg Balco
Department of Geological Sciences
University of Washington

Ciprian Crainiceanu
Department of Statistical Science
Cornell University

W. Roland Gehrels
Department of Geographical Sciences
University of Plymouth, UK

Jing Qiu
Department of Statistical Science
Cornell University

November 1, 2000

Abstract

*jstauden@hsph.harvard.edu. This work was supported by NIEHS training grant number ES07261 at Cornell University. We would like to thank Patrick Sullivan for useful discussions and suggestions. Collection of sea-level index points along the coast of Maine was funded by a variety of sources, including NSF, NOAA, the Geological Society of America, and the US Nuclear Regulatory Commission. This work is a contribution to IGCP Project 437 "Coastal Environmental Change During Sea-Level Highstands". Greg Balco was supported by a Hertz Foundation Graduate Fellowship.

Understanding past relative sea-level (RSL) changes is important to a number of social and scientific questions, including the effects of global climate change, the distribution of coastal environments, and future land-use planning. In particular, accurately characterizing millennial sea-level changes is important in evaluating their relationship to ice-sheet dynamics and long-term climate trends. In this paper, we analyze sea-level data from several Maine salt marshes previously reported in a paper from the geological literature. We address these data with a "smooth transition" / approximate segmented regression model and account for the presence of covariate measurement error. Our model builds on previous methodological work and includes a more flexible model, two alternate computational techniques for estimating the parameters in the model, and asymptotic confidence intervals for our estimates. For these specific data, our model shows that the data from one site are anomalous and therefore warrant additional investigation. This important aspect of the data was not previously identified by a non-statistical analysis. In general, our technique offers a statistically justifiable method for using geological data with measurement errors to estimate actual RSL curves that can be compared to each other, independently determined climate records, or geophysical model results.

Keywords: geology, measurement error, segmented regression, maximum likelihood, EM algorithm, smooth transition.

1 Introduction

From stories of lost cities and migrations across now sunken land bridges to questions of the effects of global warming and investigations of differential crustal motion, there has long been both public and scientific interest in understanding the history of changes in the height of the sea relative to the height of the land [13]. In this paper, we analyze data from four salt marshes along the coast

of Maine (USA) [9] and provide new methodology for similar studies.

We fit the sea level data with a smooth segmented regression model taking into account the measurement error in the dates (Section 2). This approach is methodologically similar to [11], and we develop a more flexible model (Section 2), present two alternate computational methods to estimate the parameters (Sections 3.1 and 3.2), and use approximate asymptotic confidence intervals for the estimates (Section 3).

To the best of our knowledge, this paper is the first statistical treatment of millennial scale relative sea-level data. It provides the geologist with a way to objectively reduce interpretive biases that occur when comparing features of estimated sea level curves. We illustrate the method using the data from [9] in Sections 3 and 4. Further, in Section 5 we briefly suggest how this statistical approach, combined with simulation, can be used to plan future experiments and estimate the required number of radiocarbon dates. One reason estimating relative sea-level curves is important is because sea-level change records over the past several millenia combined with climate data over the same time period need to be used to calibrate the models that are used to predict future sea levels [10].

1.1 Investigating Sea Level Change

Accurate tidal records have been collected for several centuries, but determining *millennial scale* changes in sea level requires dating deposited materials. Early attempts to do this relied on identification of exposed aquatic formations and dating either man made or natural terrestrial items below the current sea level [13]. The advent of radiometric dating has made it possible to construct quantitative estimates of relative sea level over time (see Figure 1). Core samples are collected

from the ocean shore and organic matter in the core is identified and dated (the x-axis estimate) [14]. Sea level at the date of sample deposition (the y-axis estimate) is inferred from the identity of fossilized *foraminifera* in the sample. We explain the procedure in more detail in Section 1.2.

The two most common methods currently used by geologists to summarize millennial scale sea level curves are interpolation and heuristic smooths of the data. "Error envelopes" created by connecting the corners of ± 1 standard deviation (in x and y directions) error boxes centered at each point are often displayed. Numerical summaries consisting of confidence intervals for the slopes of linear segments of the curves are sometimes included. These are based on simple linear regression of estimated height on estimated date. Comparisons of curve estimates (across sites or across researchers) seem to rely on either visual identification of seemingly common features or on comparisons of the confidence intervals around the estimated slopes. The measurement error in *both* the *time* and *height* scales does not appear to be addressed in the literature.

1.2 Gehrels, Belknap and Kelley (1996)

[9] creates four site specific scatterplots of estimated sea level (relative to present) over time in the late Holocene period. The data are displayed in Figure 1 and Table 1. This section describes the data.

Each point in Figure 1 represents a subsample from a core sample taken from the relatively undisturbed area near the bottom of one of the salt marshes. Nearly all the data points are from separate core samples. Each subsample was then analyzed to yield estimates of two quantities: deposition date (the x-axis of the scatterplot) and sea level at deposition time (the y-axis). The analysis methods are explained in detail in [9]. We summarize them here briefly.

Table 1: Site by site data from [9]. These sites are numbered in order of their relative positions from southwest to northeast along the Maine USA coast. These numbers also are used as index parameters in our models.

Site Number	Site Name	^{14}C Age (1000's ^{14}C years)	s.d. of Age	Mean High Tide at Deposition (m)	s.d. of Mean High Tide
1	Wells	0.30	0.05	-0.50	0.23
1	Wells	0.71	0.17	-1.09	0.16
1	Wells	1.09	0.05	-1.00	0.25
1	Wells	3.27	0.07	-3.08	0.24
1	Wells	3.34	0.06	-1.85	0.19
1	Wells	3.59	0.06	-3.82	0.16
1	Wells	3.90	0.15	-2.91	0.30
1	Wells	4.24	0.07	-3.40	0.17
1	Wells	4.26	0.06	-3.93	0.17
1	Wells	4.36	0.06	-3.08	0.16
1	Wells	4.74	0.07	-4.47	0.17
2	Phippsburg	0.99	0.06	-0.56	0.15
2	Phippsburg	2.68	0.07	-1.78	0.21
2	Phippsburg	2.77	0.07	-1.60	0.22
2	Phippsburg	3.44	0.05	-2.31	0.17
2	Phippsburg	3.47	0.15	-3.55	0.17
2	Phippsburg	3.76	0.06	-3.01	0.17
2	Phippsburg	4.27	0.07	-3.32	0.14
2	Phippsburg	4.84	0.10	-4.11	0.17
2	Phippsburg	4.95	0.08	-6.35	0.15
2	Phippsburg	4.98	0.06	-6.87	0.13
2	Phippsburg	6.30	0.06	-15.3	0.33
3	Gouldsboro	0.57	0.05	-0.74	0.22
3	Gouldsboro	1.23	0.07	-0.94	0.13
3	Gouldsboro	2.01	0.06	-1.39	0.22
3	Gouldsboro	2.33	0.07	-1.16	0.22
3	Gouldsboro	2.38	0.07	-1.63	0.14
3	Gouldsboro	2.73	0.08	-2.08	0.12
3	Gouldsboro	3.58	0.08	-4.21	0.20
3	Gouldsboro	4.04	0.07	-5.14	0.16
4	Machiasport	0.49	0.07	-0.75	0.12
4	Machiasport	1.07	0.09	-1.16	0.17
4	Machiasport	2.54	0.11	-1.36	0.14
4	Machiasport	3.01	0.07	-2.56	0.14
4	Machiasport	4.08	0.08	-3.65	0.14
4	Machiasport	4.80	0.08	-4.58	0.14

Deposition time was estimated by ^{14}C dating the organic matter in the subsample and assuming that the age of the organic matter coincides with deposition time. The ^{14}C years can be transformed into calendar years with a non-linear transformation [16], but, for simplicity, we have not done so here.

Sea level height is estimated from fossilized *foraminifera* in the subsample. The geologists inferred the man high tide level at the time the *foraminifera* were alive from the modern distribution of these organisms on the surface of the salt marshes. This estimates the relative mean tide level at the time the subsample was deposited when it is combined with two other estimated quantities: the depth of the subsample relative to current high tide level, and an estimate of the tidal range when the subsample was deposited.

[9] addressed three hypotheses with these data. We focus on the second one: "Differential crustal motion along the coast of Maine has occurred during the past 5 ka and can be detected by comparing local relative sea level chronologies from which the influence of changes in tidal range has been removed." *In the following sections, we formulate and fit a parametric model to these data and address this hypothesis by means of statistical estimation of the parameters in the model.*

Although this paper focuses on analyzing data from [9], the methods are generally applicable to other data from this subfield. Two additional goals of this paper are to provide a standard methodology for comparing sea level curves and to develop tools for planning future experiments.

The model we use is empirical in the sense that it is meant to approximate the observed data rather than to reflect any knowledge about the mechanism that causes sea level change. However, the model and extensions with more knots are flexible enough to approximate sea level curves for

many underlying mechanisms.

2 A Statistical Model for Sea Level Curve Estimates

We describe our model for the relative sea level curve estimates in two parts: Section 2.1 statistically characterizes the univariate data distributions, and Section 2.2 describes the regression model relating the subsample age and estimated sea level. Section 3 describes two methods to estimate the parameters. The results of estimation and a discussion of geological interpretations are in Section 4.

2.1 Model for the Collected Data

The presence of measurement error in both the sea level and age measurements makes this dataset especially interesting. A general reference for non-linear measurement error models is [4].

Let $i = 1, \ldots, 4$ represent the sites and $j = 1, \ldots, n_i$ represent the observations from a given site. Since the age of each subsample is estimated by radiocarbon dating, the observed age $W_{i,j}$ is the true age $X_{i,j}$ plus error $U_{i,j}$ ($W_{i,j} = X_{i,j} + U_{i,j}$). The radiocarbon lab supplies an estimate of the standard deviation of each estimated age ($\hat{\sigma}_{u,i,j}$).

We model the distribution of each $U_{i,j}$ (f_u) as independent $N(0, \sigma^2_{u,i,j})$. Empirical support for this is in [16]. Since ^{14}C dating involves using spectroscopy to count a very large number of atoms of a specific isotope in a sample, a loose interpretation of the central limit theorem suggests that one would expect the resulting estimate to be approximately normally distributed. Treating the $\hat{\sigma}_{u,i,j}$ as if they are known quantities is reasonable since the ^{14}C dating labs have a large amount

of validation data with which they assess the variability of an estimate from a given sample [16].

In measurement error models, a distinction is made between *functional* models where the unobserved true values ($X_{i,j}$) are considered to be nuisance parameters, and *structural* models that assume the true values are observations from a distribution with a known form. The two assumptions often result in different estimates. Since the locations of the subsamples within the sites are selected at random by the geologists, we use a structural model. We assume that the true values are independently sampled from a Beta distribution with support from zero to K and a density:

$$f_x(x_{i,j}|K,a,b) = \frac{1}{Beta(a,b)K}\left(\frac{x_{i,j}}{K}\right)^{a-1}\left[1-\left(\frac{x_{i,j}}{K}\right)\right]^{b-1}. \quad (1)$$

We use the Beta since it both has bounded support and is flexible enough to describe uniform, skewed, concave, and convex probability densities. We estimate a and b from the data, but we found the maximum likelihood estimate of K to be unstable. As a result, we set $K = \max_{i,j} W_{i,j} + 1$. Varying K in a neighborhood of that value did not change any of the other parameter estimates substantially.

Relative to the "naive" regression estimator that ignores measurement error, an estimator accounting for measurement error generally will be less biased and more variable. It is therefore possible that the "naive" estimator has a smaller mean squared error than the measurement error corrected estimator. Nonetheless, we prefer the measurement error corrected estimator here for two reasons. First, our estimates (Section 3) do not appear to be highly variable. Second, and perhaps more importantly, since one of our goals is to suggest a methodology that can be used for other applications of radiometric dating, we prefer not to ignore an important feature of the data.

There are also errors in the y-direction. First, there is measurement error in each estimated sea level, $Y_{i,j}$, and second, there is error in the equation of our model (Section 2.2) at each (i,j). We group both these errors into $\epsilon_{i,j}$ and assume that $\{\epsilon_{i,j}\}_{1,1}^{4,n_i}$ are independent, mean zero, and normally distributed. Each $\epsilon_{i,j}$ has its own variance, $\sigma^2_{\epsilon,i,j}$, with two components: $\sigma^2_{\epsilon(1),i,j}$, the variance of the measured sea level, and $\sigma^2_{\epsilon(2)}$, the variance around the equation for the mean of the model. [9] provides estimates of $\{\sigma^2_{\epsilon(1),i,j}\}_{1,1}^{4,n_i}$. We estimate $\sigma^2_{\epsilon(2)}$ from the data. Let $\sigma^2_{\epsilon,i,j} = \sigma^2_{\epsilon(1),i,j} + \sigma^2_{\epsilon(2)}$.

We define the vectors $\mathbf{Y}, \mathbf{W}, \mathbf{X}$, and $\boldsymbol{\epsilon}$ as follows:

$$\begin{aligned}
\mathbf{Y} &= [Y_{1,1}, \ldots, Y_{1,n_1}, Y_{2,1}, \ldots, Y_{2,n_2}, Y_{3,1}, \ldots, Y_{3,n_3}, Y_{4,1}, \ldots, Y_{4,n_4}]', \\
\mathbf{W} &= [W_{1,1}, \ldots, W_{1,n_1}, W_{2,1}, \ldots, W_{2,n_2}, W_{3,1}, \ldots, W_{3,n_3}, W_{4,1}, \ldots, W_{4,n_4}]', \\
\mathbf{X} &= [X_{1,1}, \ldots, X_{1,n_1}, X_{2,1}, \ldots, X_{2,n_2}, X_{3,1}, \ldots, X_{3,n_3}, X_{4,1}, \ldots, X_{4,n_4}]', \text{ and} \\
\boldsymbol{\epsilon} &= [\epsilon_{1,1}, \ldots, \epsilon_{1,n_1}, \epsilon_{2,1}, \ldots, \epsilon_{2,n_2}, \epsilon_{3,1}, \ldots, \epsilon_{3,n_3}, \epsilon_{4,1}, \ldots, \epsilon_{4,n_4}]'.
\end{aligned}$$

Let $n = n_1 + n_2 + n_3 + n_4$, the total number of data points. Figure 1 and Table 1 show the observed data and their estimated errors.

2.2 Model for Mean Sea Level

For each of the four sites, we model the expected relative sea level given ^{14}C age as a smooth approximation to a segmented linear regression (or linear spline) with one breakpoint (or knot). We arrive at this model by first fitting quadratic P-spline regressions [3] to each scatterplot (ignoring the measurement error). This flexible fit indicates that a linear spline with one knot is a good approximation to the shape of the curve. A benefit of this model is the ease with which the parameters can be interpreted. Of course, data from other studies might require more knots or

other modifications.

The model that we use is similar to the smooth transition models developed by [1]. Let $\Phi(X)$ be the normal cumulative distribution function and $(X)_+$ be a function that returns X if $X > 0$ and 0 otherwise. The parameters of interest are, for sites $i = 1, \ldots, 4$, $\beta_{0,i}$, the intercept, $\beta_{1,i}$, slope prior to the knot, $\beta_{2,i}$, the additional slope following the knot, and τ_i, the knot location.

The first regression function for the model we consider has the general form:

$$\begin{aligned} E[Y_{i,j}|X_{i,j}] &= \beta_{0,i} + \beta_{1,i}X_{i,j} + \beta_{2,i}(X_{i,j} - \tau_j)\Phi\left(\frac{X_{i,j} - \tau_i}{\sigma}\right) \\ &:= \mu_{saturated}(X_{i,j}). \end{aligned} \tag{2}$$

This function is a differentiable approximation to the segmented regression function:

$$E[Y_{i,j}|X_{i,j}] = \beta_{0,i} + \beta_{1,i}X_{i,j} + \beta_{2,i}(X_{i,j} - \tau_i)_+. \tag{3}$$

We prefer the differentiable approximation since it allows us to estimate the parameters in the model with efficient quasi-Newton methods (Section 3.1) and maintains the interpretability of the segmented regression model. We select σ to be 0.05 since we find that value to be both small enough to provide a good approximation to the segmented linear function and large enough to keep the condition number of the approximate inverse Hessian in the quasi-Newton algorithm small enough for numerical stability. With more data, we could choose σ by maximum likelihood.

Table 2 shows that a simpler model describes our data nearly as well as the saturated model. This simpler model uses a β_0 and β_1 that are common across sites, allows β_2 to vary from site to site, and gives site 2 (Phippsburg) its own knot. Let $k = 1$ for $i = 1, 3$, or 4 and 2 otherwise. The

Table 2: This table illustrates the pattern of how the fit degrades as the model is simplified.

Site Specific Parameters	Number of Parameters in Regression Model	2Log(Likelihood)	P-value for Improved Fit
$\beta_{0,i}, \beta_{1,i}, \beta_{2,i}, \tau_i$ ("saturated")	16	-152.52	0.32
$\beta_{2,i}$ and two τs ("simple")	8	-161.17	0.001
None	4	-189.06	

simple model is:

$$\begin{aligned} E[Y_{i,j}|X_{i,j}] &= \beta_0 + \beta_1 X_{i,j} + \beta_{2,i}(X_{i,j} - \tau_k)\Phi\left(\frac{X_{i,j} - \tau_k}{\sigma}\right) \\ &:= \mu_{simple}(X_{i,j}). \end{aligned} \tag{4}$$

We fit and discuss both models in the next sections.

As a vector, denote the parameters to be estimated as $\boldsymbol{\theta}$. For the saturated model,

$$\boldsymbol{\theta}_{saturated} = [\beta_{0,1}, \ldots, \beta_{0,4}, \beta_{1,1}, \ldots, \beta_{1,4}, \beta_{2,1}, \ldots, \beta_{2,4}, \tau_1, \ldots, \tau_4, \sigma_{\epsilon(2)}, a, b]'.$$

For the simpler model,

$$\boldsymbol{\theta}_{simple} = [\beta_0, \beta_1, \beta_{2,1}, \ldots, \beta_{2,4}, \tau_1, \tau_2, \sigma_{\epsilon(2)}, a, b]'.$$

The parameters of geological interest are the intercepts, slopes and knots.

3 Fitting the model

We estimate the parameters in our model using maximum likelihood (ML) and compare two computational methods: direct maximization using numerical integration and the Expectation Maximization (EM) algorithm [5]. We find that both methods yield nearly identical estimates, and discuss their strengths and weaknesses in Section 3.3. As mentioned earlier, the direct method is similar to the approach taken by [11]. Our EM implementation treats the unobserved $\{X_{i,j}\}$ as latent variables [15], [4].

Similarly to [12], we assess the variability around the estimates from the direct method using confidence intervals based on likelihood ratios. We also present an approximate confidence band around the mean response. Since we use a smooth approximation to the segmented regression model, one would expect the standard asymptotic theory about the null distribution of the likelihood ratio statistic to be valid. We looked at the profile likelihoods and they were smooth, relatively symmetric and apparently convex in a neighborhood of their maxima.

For the segmented regression problem with no measurement error, [7] and [6] showed that under certain reasonable conditions, likelihood ratio statistics often have the standard chi-squared asymptotic null distributions. We mention these results to point out that even though our smooth approximation is in some sense close to a non-standard problem (segmented regression), this particular non-standard problem is fairly well behaved.

3.1 Direct Maximization

The likelihood function for this model is:

$$L(\mathbf{Y},\mathbf{W}|\boldsymbol{\theta}) = \int_{\mathbf{x}\in R^n} \prod_{i=1}^{4}\prod_{j=1}^{n_i} f_y(y_{i,j}|\mu(x_{i,j}),\sigma^2_{\epsilon,i,j}) f_u(w_{i,j}|x_{i,j},\sigma^2_{u,i,j}) f_x(x_{i,j}|K,a,b)\mathbf{dx}. \quad (5)$$

The likelihood has the same general form whether $\mu(x_{i,j})$ is the simplified or saturated model. Since we assumed conditional independence of $(Y_{i,j}, W_{i,j}, X_{i,j})$, this integral is simply a product of n 1-dimensional integrals. The log-likelihood is:

$$l(\mathbf{Y},\mathbf{W}|\boldsymbol{\theta}) = \sum_{i=1}^{4}\sum_{j=1}^{n_i} \log\left[\int_{x_{i,j}\in R} f_y(y_{i,j}|\mu(x_{i,j}),\sigma^2_{\epsilon,i,j}) f_u(w_{i,j}|x_{i,j},\sigma^2_{u,i,j}) f_x(x_{i,j}|K,a,b) dx_{i,j}\right]. \quad (6)$$

The direct approach relies on the fact that we can use numerical methods both

1. to evaluate (approximate) each of those one dimensional integrals,

2. and to compute the maximum likelihood estimates of the parameters.

Approximations of the integrals are quite straightforward and fast since the functions involved are 1-dimensional, smooth and have essentially bounded support. We use efficient quasi-Newton methods to find $\hat{\boldsymbol{\theta}}$ and maximize the likelihood, and find that using the smoother approximation to the segmented regression results in more numerically stable estimates. Although the second derivative of the mean function at the breakpoint is large, the condition number of the inverse approximate Hessian used in the quasi-Newton method is small enough to maintain numerical stability. Tables 3 and 4 contain the estimates of the parameters of interest for the simple and saturated models respectively. The fitted curves and approximate confidence bands are displayed in Figures 2 and 3. We use MATLAB to compute these estimates. The computations required less than 10 seconds on a Pentium II 133MHz computer running Windows 98.

Table 3: Maximum likelihood estimates of the parameters of interest in the simplified model and 90% confidence intervals.

Parameter	MLE	90% Confidence Interval
β_0	-0.15	(-0.31, 0.00)
β_1	-0.68	(-0.73,-0.63)
$\beta_{2,1}$	-0.39	(-1.61,-0.17)
$\beta_{2,2}$	-5.61	(-7.61,-3.61)
$\beta_{2,3}$	-1.62	(-2.49,-1.00)
$\beta_{2,4}$	-0.54	(-0.88,-0.24)
τ_1 (sites 1,3,4)	2.67	(2.36 ,3.05)
τ_2 (site 2)	4.48	(4.27 ,4.72)

Table 4: Maximum likelihood estimates of the parameters in the saturated model and 90% confidence intervals.

Parameter	MLE	90% Confidence Interval
$\beta_{0,1}$	-0.42	(-0.65,-0.19)
$\beta_{0,2}$	0.51	(0.25, 0.76)
$\beta_{0,3}$	-0.49	(-0.75,-0.24)
$\beta_{0,4}$	-0.64	(-0.92,-0.37)
$\beta_{1,1}$	-0.52	(-0.59,-0.45)
$\beta_{1,2}$	-0.91	(-0.98,-0.84)
$\beta_{1,3}$	-0.39	(-0.51,-0.28)
$\beta_{1,4}$	-0.64	(-0.45,-0.27)
$\beta_{2,1}$	-0.80	(-1.01,-0.59)
$\beta_{2,2}$	-5.70	(-8.70,-3.43)
$\beta_{2,3}$	-1.91	(-2.55,-1.46)
$\beta_{2,4}$	-0.95	(-1.16,-0.73)
τ_1	2.86	(2.53, 3.21)
τ_2	4.58	(4.35, 4.80)
τ_3	2.45	(2.19, 2.75)
τ_4	2.41	(2.03, 2.80)

Since the mean response functions are nonlinear in the estimated parameters and the number of parameters is large, developing a confidence band for the mean response is difficult. As a result, we develop approximate confidence bands in two steps: first we compute bands that treat the knots as if they were fixed, and then we adjust to account for the variability of the knot estimates.

More specifically, when the knot is fixed, the model is linear in the estimated parameters. As a result, we can use minus the inverse of the Hessian of the observed log-likelihood (Equation 6) with respect to the βs as an estimate of covariance matrix of the estimated βs and appeal to asymptotic normality to get an asymptotic confidence band. For a $1-\alpha\%$ band, we account for the variability of the knot by using the limit of the $1-2\alpha\%$ univariate confidence interval for the breakpoint in the lower bound of the confidence band and the upper limit of the knot in the upper bound. While we would expect the claimed coverage probabilities to differ somewhat from actual coverage probabilities, the band provides a sense of the variability around the fitted model. For statistical inference, we recommend using the parameter estimates such as Tables 3 and 4 and Figure 4.

3.2 Expectation Maximization (EM) Algorithm

We also implemented the EM algorithm [5] to compute the maximum likelihood estimates treating $[\mathbf{X}]$ as a latent variable.

For this problem, the algebra involved is quite difficult, but is relatively straightforward to handle numerically. The algorithm converges to the same estimates as the direct method (Tables 3 and 4) after about thirty iterations.

3.3 Comparison of the Direct Method and the EM Algorithm

We found both the EM algorithm and the direct maximization method to be acceptable approaches for this application. The direct method is simpler to explain but is tractable only because the likelihood only involves low integrals. If the measurement errors were correlated or the $Y_{i,j}$s were not independent for instance, the direct method becomes intractable, but one could still easily implement the EM algorithm.

4 Discussion of Inferences and Geological Significance

Figure 4 shows parameter estimates for the saturated model at all sites (arrayed from SW to NE), and can be used to assess the geographic variability of RSL curves along the Maine coast. These parameter estimates suggest three important conclusions.

First, the modeled RSL curve at Phippsburg is significantly different from those at other sites: sea-level rise slowed from middle Holocene rates earlier at this site, but was more rapid than at others thereafter. This has several possible explanations. First, it is possible that there has been an unrecognized experimental error. Second, some hydrographic effect related to this site's backbarrier setting, perhaps a change in tidal inlet geometry, may have caused changes in tidal range at the site that are not directly related to changing water levels in the open Gulf of Maine. Third, it is possible that unusual crustal deformation took place along this portion of the Maine coast. In any case, further examination of this site is warranted. It is also interesting that the conclusion derived from our statistical analysis, that Phippsburg has an anomalous RSL history, differs from the conclusion of [9] that unusually rapid RSL rise took place at Gouldsboro. This

points out the need for a consistent and statistically justifiable method of comparing inherently variable radiocarbon-based RSL histories.

Second, with the exception of anomalous results at Phippsburg, estimated model parameters change slowly with geographic position: the rate of late Holocene RSL rise increases gradually to the SW, and other parameters are not distinguishable at the 95% confidence interval among Wells, Gouldsboro, and Machiasport. This suggests that these three sites are recording essentially the same RSL history, reinforces the conclusion of [9] that neotectonic deformation has not taken place in extreme eastern Maine, and is consistent with the general paradigm of a stiff continental crust generally assumed in geophysical models for RSL change.

Third, with the exception of Phippsburg, the fact that all the estimated intercepts are negative suggests that in very recent years the rate of RSL rise has increased. The same result, that the rate of RSL rise over the last century is unprecedented in the past two millenia, was obtained by [9] from comparison of salt-marsh and tide gauge data and by [8] from more detailed analyses of very recent salt-marsh stratigraphy. While it is tempting to constrain the regressions to go through the origin, this would not be appropriate since our model applies to the time scale range of the observed data which does not include the present. Constraining the regressions to go through zero (without adding another knot) would be similar to adding a leverage point to the regression.

5 Conclusion and Using the Model to Help Plan Future Experiments

In this paper we developed and fitted a flexible parametric model to describe the data generated from basal peat based relative sea-level estimates. The method provides parameter parameter estimates and an objective way of comparing curves from site to site. It can also be used to plan future experiments and determine the required sample size at each site as follows: Given

1. a hypothesized relative sea level curve,
2. roughly feasible ranges for the distributional parameters described in Section 2,
3. and a proposed sample size,

one easily can simulate the data one might expect to collect and fit a model to that simulated data. Monte Carlo methods can then be used to estimate the probability of detecting a specific hypothesized difference of a certain size between two sites. If those probabilities are deemed too low, this procedure can be repeated with a larger proposed sample size. The procedure is stopped when the probabilities are deemed to be high enough.

While the one knot linear spline model is both flexible enough to fit these data well and parsimonious enough to provide reasonable inferential power, if we had more data, it would be interesting to modify [2]'s recently developed Bayesian method of estimating penalized regression splines in the presence of measurement error to suit this problem.

References

[1] D. W. Bacon and D. G. Watts. Estimating the transition between two intersecting straight lines. *Biometrika*, 58:525–534, 1971.

[2] S. Berry, R. Carroll, and D. Ruppert. Bayesian smoothing and regression splines for measurement error problems. *Submitted to Journal of the American Statistical Association*, 2000.

[3] B. Brumback, D. Ruppert, and M. P. Wand. Comment on "Variable selection and function estimation in additive nonparametric regression using a data-based prior" by Shively, Kohn, and Wood. *Journal of the American Statistical Association*, 94:794–797, 1999.

[4] R. Carroll, D. Ruppert, and L. Stefanski. *Measurement Error in Nonlinear Models.* Chapman and Hall, 1995.

[5] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B, Methodological*, 39:1–38, 1977.

[6] P. I. Feder. The log likelihood ratio in segmented regression. *Annals of Statistics*, 3:84–97, 1975.

[7] P. I. Feder. On asymptotic distribution theory in segmented regression problems–identified case. *Annals of Statistics*, 3:49–83, 1975.

[8] W. R. Gehrels. Middle and Late Holocene Sea-Level Changes in Eastern Maine Reconstructed from Foraminiferal Saltmarsh Stratigraphy and AMS ^{14}C Dates on Basal Peat. *Quaternary Research*, 52:350–359, 1999.

[9] W. R. Gehrels, D. F. Belknap, and J. T. Kelley. Integrated high-precision analyses of Holocene relative sea-level changes: Lessons from the coast of Maine. *Geological Society of America Bulletin*, September:1073–1088, 1996.

[10] I. Goodwin, N. Harvey, O. van de Plassche, R. Oglesby, and F. Oldfield. Research targets 2000 year of climate-induced sea level fluctuations. *EOS Transactions, American Geophysical Union*, 81:311–312, 2000.

[11] H. Küchenhoff and R. Carroll. Segmented regression with errors in predictors: Semi-parametric and parametric methods. *Statistics in Medicine*, 16:169–188, 1997.

[12] P. M. Lerman. Fitting segmented regression models by grid search. *Applied Statistics*, 29:77–84, 1980.

[13] P. A. Pirazzoli. *Sea-Level Changes : The Last 20000 Years (Coastal Morphology and Research).* John Wiley and Sons, 1996.

[14] A. C. Redfield and M. Rubin. The age of salt marsh peat and its relations to recent changes of sea level at Barnstable, Massachusetts. *Proceedings of the National Academy of Sciences*, 48:1728–1735, 1962.

[15] D. Schafer. Covariate measurement error in generalized linear models. *Biometrika*, 74:385–391, 1987.

[16] M. Stuiver, P. J. Reimer, E. Bard, J. W. Beck, G. S. Burr, K. A. Hughen, B. Kromer, G. McCormac, J Van Der Plicht, and M. Spurk. INTCAL98 Radiocarbon Age Calibration, 24,000-0 cal BP. *Radiocarbon*, 40:1041–1083, 1998.

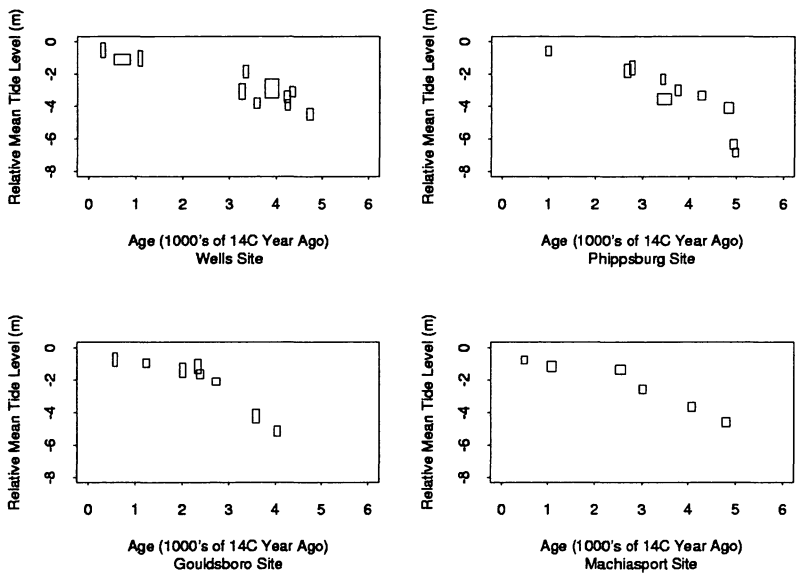

Figure 1: These four plots contain scatterplots of the collected data. Relative sea level is computed as (previous mean sea level) - (current mean sea level). The sea level relative to the land has been increasing over time. This figure does not contain the last observation from the Phippsburg site (see Table 1), which has been deleted because it was outside the data range of the other sites. Note however that this data point agrees with the fitted model quite well. The widths and heights of the rectangles are equal to two times the estimated standard deviations of the age and mean tide estimates respectively.

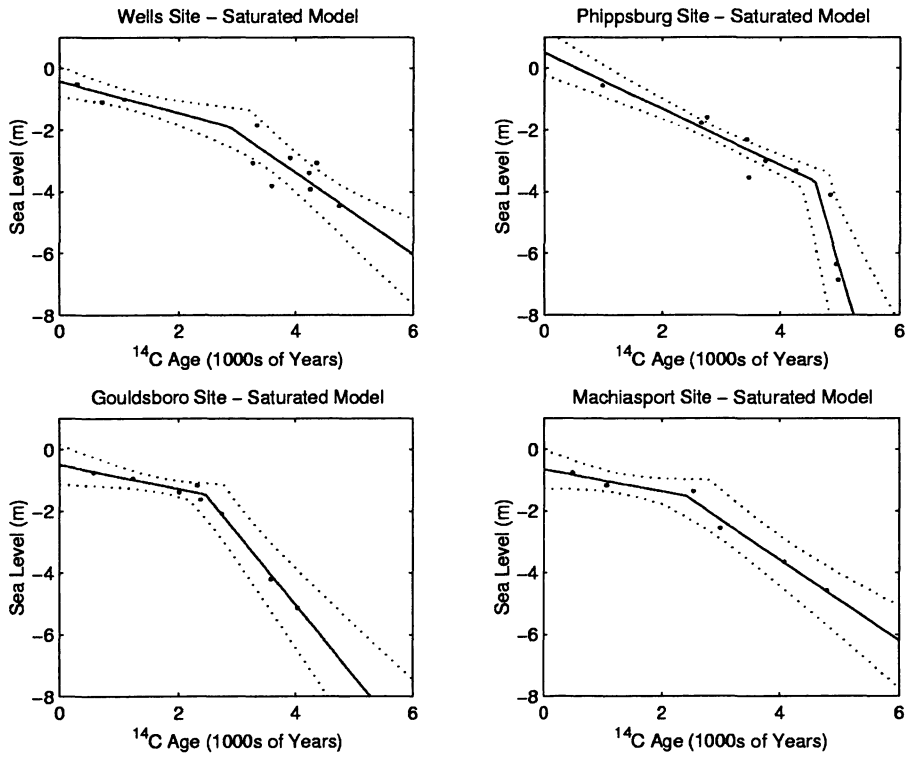

Figure 2: These four plots illustrate an approximate 95% confidence band around the estimated saturated model mean response.

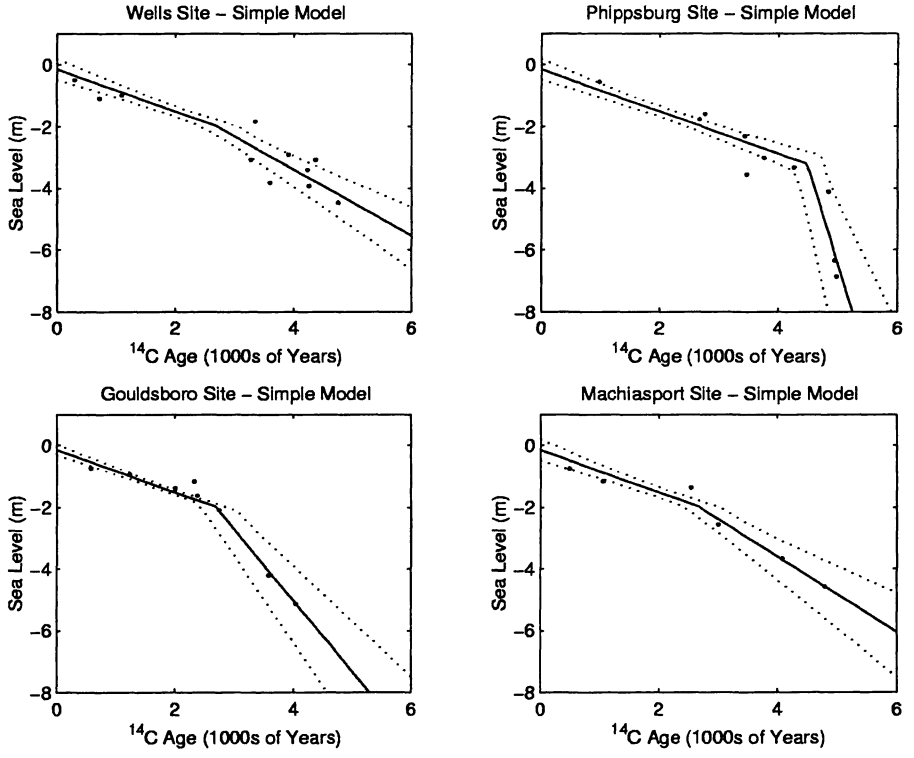

Figure 3: These four plots illustrate an approximate 95% confidence band around the estimated simple model mean response.

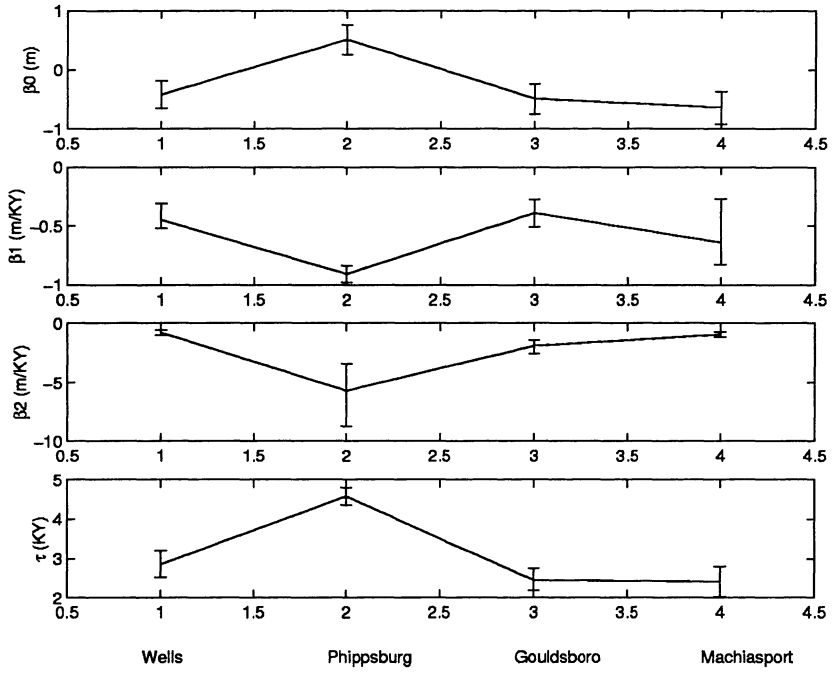

Figure 4: These plots summarize geographical variation of the parameter estimates and 90% confidence intervals from the saturated model. Left to right, the sites are arrayed from SW to NE. Note that β_0 is the intercept; β_1 is the slope before the knot; β_2 is the additional slope after the knot; and τ is the knot location.