An Assessment of Climate Change in the Ocean, Mixed-Layer Depth, and Subjective Likelihood*

Ana Grohovac Rappold
US EPA

Michael Lavine
Duke University

Susan Lozier
Duke University

August 25, 2005

1 Introduction

This paper describes a Bayesian statistical analysis of long term changes in the depth of the ocean's mixed layer. The data are thermal profiles recorded by ships. For this data there is no good sampling model and therefore no obvious likelihood function. Our approach is to elicit posterior distributions for small data sets directly from the expert. Then we infer the likelihood function and use it on large data sets.

The typical Bayesian analysis posits data from a parametric family of sampling distributions, as in Equation 1.

$$\vec{y} \sim p(\vec{y} \,|\, \theta) \tag{1}$$

After $\vec{y}$ has been observed it is treated as fixed, and the likelihood function is defined to be $\ell(\theta) \equiv p(\vec{y} \,|\, \theta)$, a function of θ. The interpretation is that likelihood ratios $\ell(\theta_1)/\ell(\theta_2)$ quantify $\vec{y}$'s evidence for θ_1 as opposed to θ_2.

*supported by NSF grant #ATM–0221939

For our data set there is no believable sampling model $p(\vec{y}\,|\,\theta)$, so we cannot assign $\ell(\theta) \equiv p(\vec{y}\,|\,\theta)$. We take a different approach, wherein lies the statistical novelty of this paper. For several values of i, about a dozen, we show the expert $\vec{y}_i$ and directly elicit her posterior distribution $p(\theta\,|\,\vec{y}_i)$. Elicitation is done under conditions where the expert has an approximately uniform prior for θ. After elicitation we know the prior and posterior and can therefore infer the likelihood function.

After examing the dozen or so elicited posteriors and conferring with the expert, we constructed an algorithm that accepts a $\vec{y}$ as input and yields a likelihood function $\ell(\theta)$ as output. After contructing the algorithm we checked that it gave sensible results on several hundred more $\vec{y}$'s. We then applied the algorithm to our full collection of data $\{\vec{y}_i\}_{i=1}^{T}$ which, when combined with our real prior, yields our posterior. We call ℓ a likelihood function because it approximately summarizes the expert's weight of evidence and, when multiplied by the uniform prior, yields her posterior.

The data arise in a study of the ocean's climate. The situation is more complicated than described in this introduction, as the data are a time series and our model must account for an annual cycle. Section 2 describes the scientific background while Section 3 describes the data. Section 4 describes our subjective likelihood, how it was elicited, and how it is modelled. It contains whatever statistical novelty is in this paper. Section 6 describes our prior, accounting for the annual cycle, year-to-year variation, heteroscedasticity, and a possible secular trend. The posterior is found by MCMC; it is described in Section 7. Section 8 is a discussion and Section 9 is an appendix describing an alternative model that we ultimately rejected.

2 Application to an Ocean Mixed Layer

Recent evidence that the world's oceans have warmed over the past fifty years (Levitus et al. [2000]) and that the attendant increase in the ocean's heat content is an order of magnitude larger than the increase in the atmospheric and cryospheric heat content (Levitus et al. [2001]) have made it abundantly clear that a determination of how our global climate is changing in response to long-term natural and/or anthropogenic forcing depends on the effectiveness of the ocean as a heat reservoir. However, the effectiveness of the ocean as a reservoir is curtailed by increasing thermal stratification which limits the extent to which surface signals can be transmitted to depth. Interest,

therefore, has focused on the upper ocean.

To a first approximation, oceanographers regard the ocean as having two layers: a mixed layer from the surface down to as much as several hundred meters, and a stratified layer beneath. The mixed layer is that part of the surface ocean that displays uniformity in properties such as temperature, salinity, and density. The mixed layer forms because the upper waters of the ocean are mixed via waves and wind and also through thermal convection where the surface waters overturn upon losing buoyancy to the atmosphere. Such overturning creates a mixed layer. The depth M of the mixed layer evolves through an annual cycle and depends on geographic location. Because M depends crucially on the heat content of the upper ocean, long-term changes in heat content may result in long-term changes in M. Essentially we address in this application, the question as to whether or not there has been a secular trend in M in the North Atlantic subtropical gyre. Such an evaluation will increase our understanding of potential physical and biological consequences of global warming.

3 Data

Hydrographic data, such as temperature, salinity and pressure, is collected from ships, sent to the National Oceanographic Data Center (NODC), where it is quality controlled and then made publicly available (`www.nodc.noaa.gov`). This paper reports on an analysis of NODC historical hydrographic data recorded over a small spatial region near Bermuda. A sufficiently small region was chosen for our study so that we can safely ignore spatial variability. An analysis of data from a wide region of the North Atlantic will appear elsewhere. A map of the data is in Figure 1; the data's temporal distribution is shown in Figure 2.

The i'th data point $\vec{y}_i$ is recorded at time t_i and consists of temperatures $\vec{y}_i = (y_{i1}, \ldots, y_{in_i})$ measured at depths $\vec{d}_i = (d_{i1}, \ldots, d_{in_i})$. (Other properties are also measured, but we do not report on them here.) The first depth is always sufficiently close to the surface so that (1) we can assume y_{i1} is approximately equal to the surface temperature and (2) we can take $d_{i1} = 0$ without serious error. Temperature as a function of depth is called a temperature *profile*. Figure 3 shows one such profile.

Generally speaking, the upper layer of the ocean, because it is vertically mixed, should have a uniform temperature, while the stratified layer should

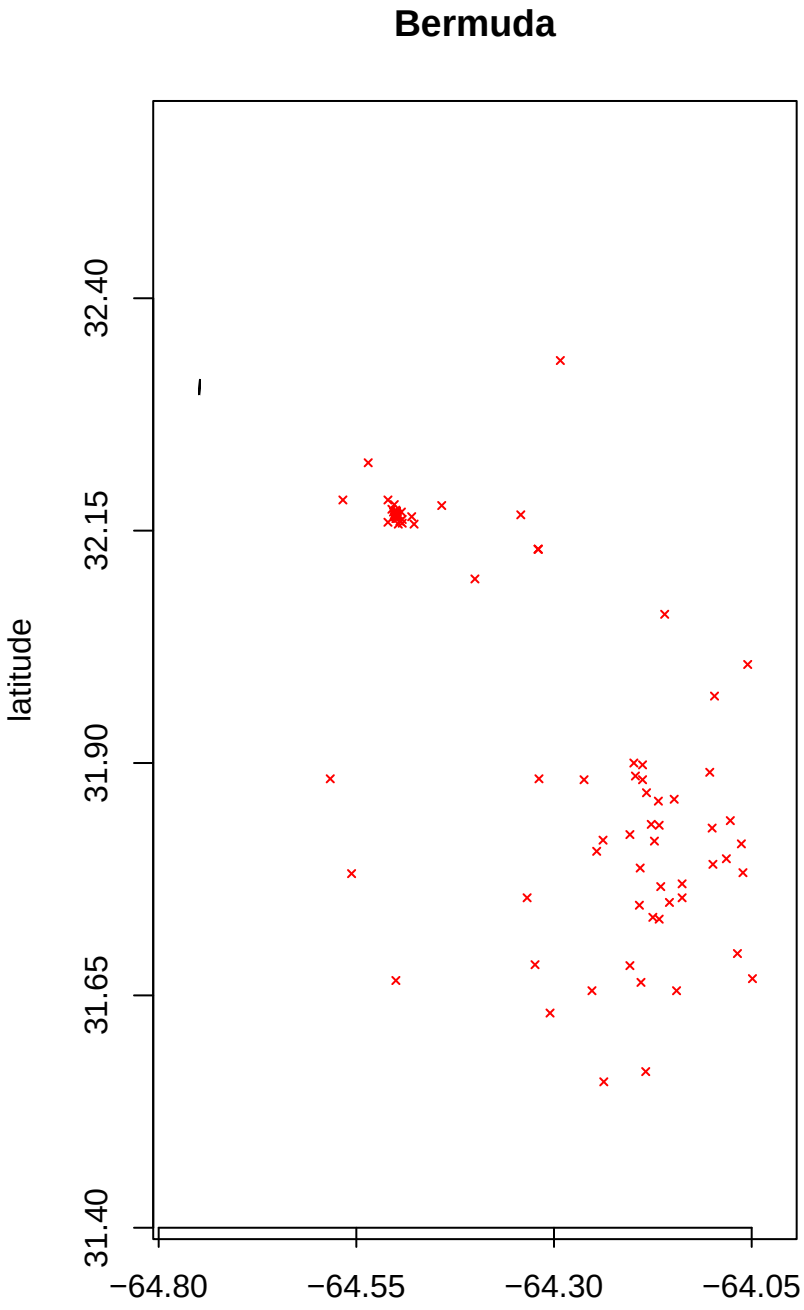

Figure 1: data locations near Bermuda

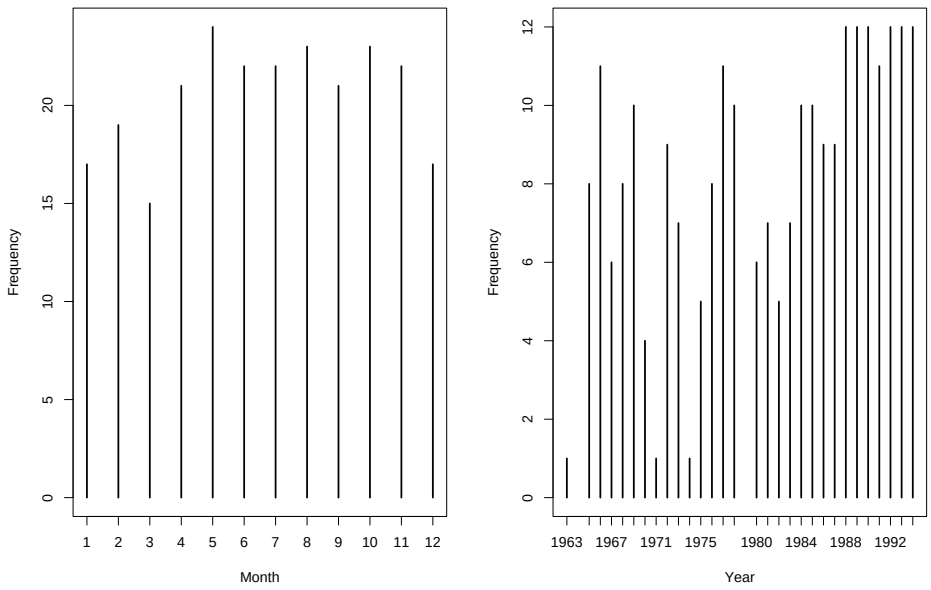

Figure 2: Sampling times near Bermuda

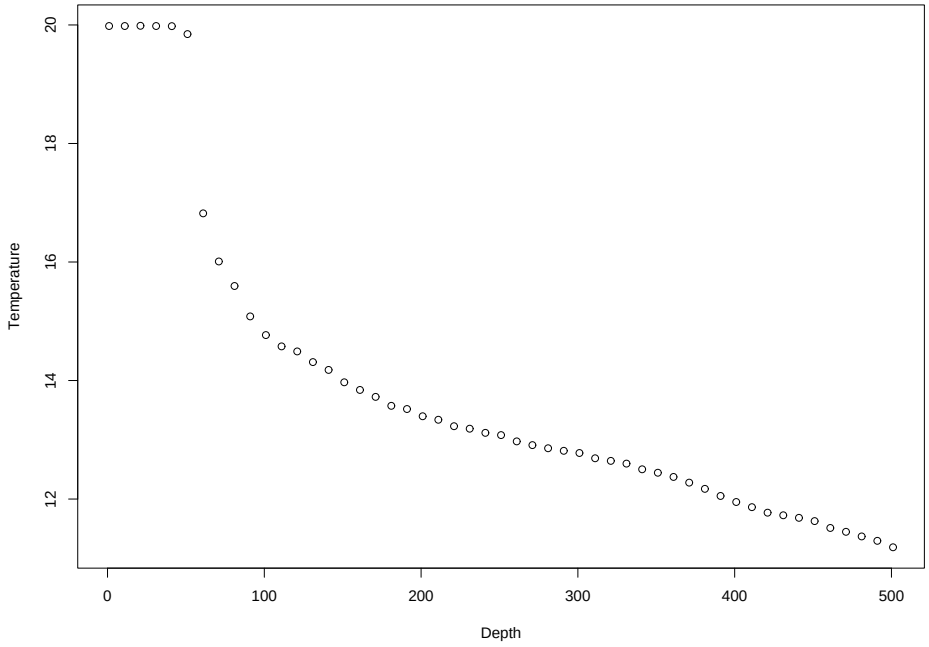

Figure 3: A nice profile: temperature as a function of depth measured at intervals of 10m. M appears to be between about 40 and 60 meters.

have a monotonically decreasing temperature. Figure 3 exemplifies these characteristics. Locating the bottom of the mixed layer is fairly straightforward in such profiles. In Figure 3 we can say with confidence that M is somewhere near 50 meters. Not all profiles, however, present such a clear delineation of M. For examples, see Figure 4. Each of the profiles in Figure 4 has several depths which might plausibly be the bottom of the mixed layer. A key question is: What do profiles such as those in Figure 4 tell us about M?

If we had a physical model that predicted the spatial and temporal distribution of temperature with depth (i.e., a temperature profile), we could fit the model to the data, and the likelihood function would quantify the information for the parameter M. However, one of the remaining fundamental problems in oceanography is a complete theoretical description of the thermocline, which expresses the horizontal and vertical change in the temperature field. In the absence of a complete description, simplifications have been offered. One such simplification of the nonlinear physics that governs the thermocline yields a prediction that the temperature should decrease exponentially from M to the ocean floor (Mellor [1996]), where it is about 2 °C throughout the world. If such a simplification were generally valid, then a change-point model with three parameters — surface temperature, M, and decay rate — would fit the data well. However, the data in our region of interest do not exhibit exponential decay, as illustrated by Figure 5 which shows the same profile displayed in Figure 3 along with several exponential decay curves. Overall, our understanding of the physics of the thermocline has not advanced to the point where we could offer a model to which to fit the data. In such a void, we turn instead to the methods described in Section 4.

4 Subjective Likelihood

4.1 Heuristic Explanation

In this section we consider what can be learned from a single profile $\vec{y}$ sampled at depths $\vec{d}$; the parameter of interest is M.

From either a likelihood or Bayesian perspective, what is needed is a sampling model $p(\vec{y} \mid M)$. However, in Section 3 we argued that no reliable sampling model exists. So how can we describe what is learned from $\vec{y}$? Our

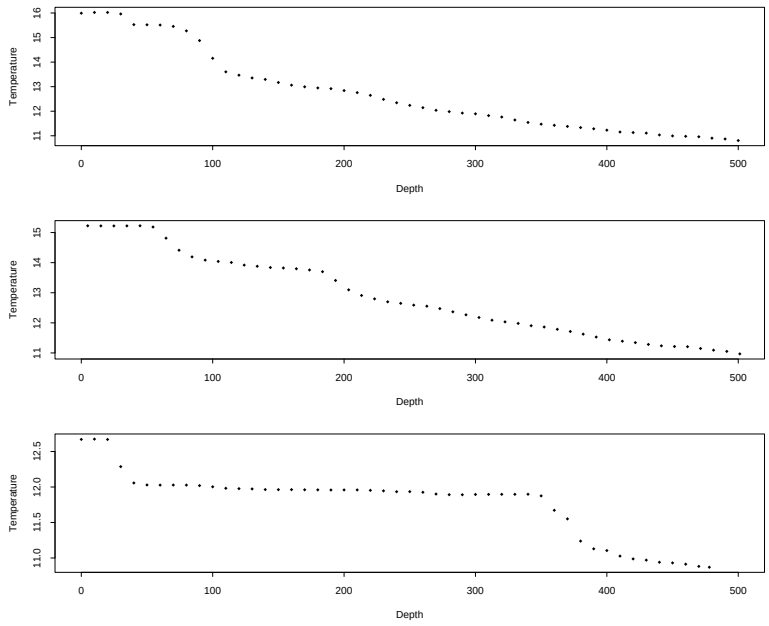

Figure 4: Three profiles. None indicate M with clarity.

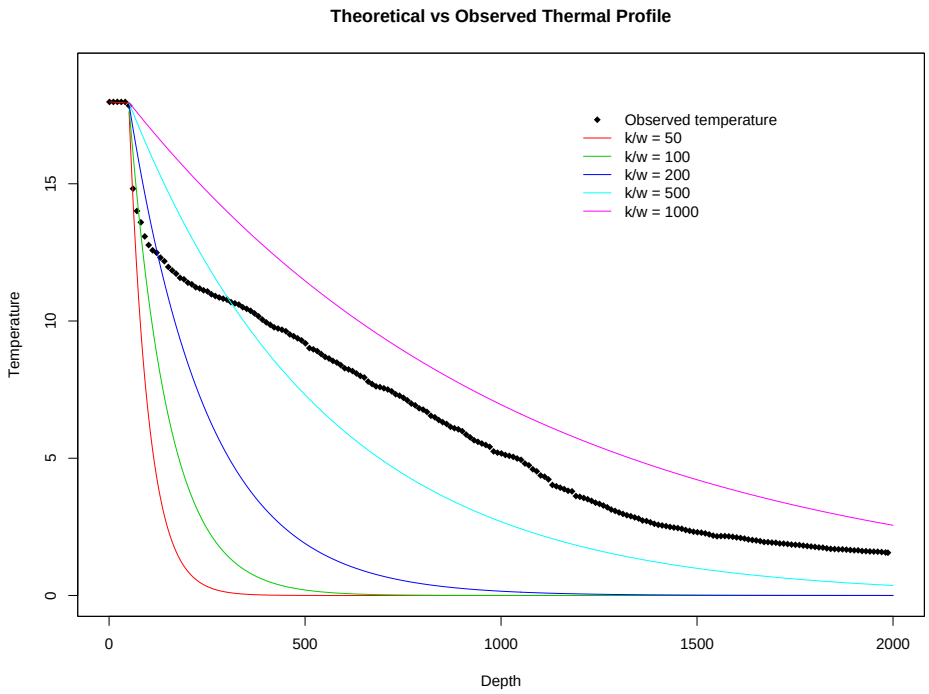

Figure 5: The same profile as in Figure 3 with several exponential decay curves. None fit the data well.

approach is to ask an oceanographer directly.

We showed the oceanographer a profile $(\vec{y}, \vec{d})$ and asked what she could tell us about M, based on that profile. Her response was to tell us $\Pr[M \in I_j]$ for $j = 1, \ldots, n$ where I_j is the interval from d_j to d_{j+1}. ($d_1 = 0$ is the ocean surface; d_{n+1} is the floor.) The process was repeated for several profiles, about a dozen. Figure 3 is an example. For this profile the oceanographer's assessments were

$$\Pr[M \in I_6] \approx 10 \Pr[M \in I_5]$$
$$\Pr[M \in \text{some other interval}] \approx \text{very small}$$

The next step was to ask the oceanographer how she makes her assessments. She told us that for an interval $I_j = [d_j, d_{j+1}]$ she considers two quantities: $y_1 - y_j$, the temperature drop from the surface to the top of I_j, and $(y_j - y_{j+1})/(d_{j+1} - d_j)$, the rate of temperature drop within I_j. Small values of $y_1 - y_j$ and large values of $(y_j - y_{j+1})/(d_{j+1} - d_j)$ imply that I_j is likely to contain M. She also said these are the only quantities that matter; other aspects of the profile carry so little information as to be ignorable. To formalize, define

$$\vec{\Delta}_1 = (\Delta_{11}, \ldots, \Delta_{1n}) = (0, y_1 - y_2, \ldots, y_1 - y_n)$$
$$\vec{\Delta}_2 = (\Delta_{21}, \ldots, \Delta_{2n}) = ((y_1 - y_2)/(d_2 - d_1), \ldots, (y_n - 2)/(d_{n+1} - d_n))$$

Then, according to the oceanographer, the probabilities $\Pr[M \in I_j]_{j=1}^n$ are some function $g(\vec{\Delta}_1, \vec{\Delta}_2)$.

The reasoning is this. If interval I_j contains M, then the temperature above I_j should be roughly uniform, and $y_1 - y_j$ should be small. So $g(\vec{\Delta}_1, \vec{\Delta}_2)$ should be a decreasing function of $\vec{\Delta}_1$. Similarly, if I_j contains M, then the rate of temperature decrease in I_j should be large. So $g(\vec{\Delta}_1, \vec{\Delta}_2)$ should be an increasing function of $\vec{\Delta}_2$.

After considering multiple profiles — some real and some artificially constructed to learn particular aspects of the oceanographer's thinking — we settled on

$$\Pr[M \in I_j \,|\, \vec{y}, \vec{d}]_1^n \propto g(\vec{\Delta}_1, \vec{\Delta}_2) \approx \exp\{-2\vec{\Delta}_1\} \times [1 - \exp\{-\vec{\Delta}_2/0.3\}] \quad (2)$$

as providing sufficiently close approximation to the oceanographer's posterior probabilities for the intervals I_j. (The left-hand side of (2) is a vector with one

component for each interval. The right-hand side is also a vector; notation such as $\exp\{-2\vec{\Delta}_1\}$ means element-wise multiplication and exponentiation.)

Figures 6 and 7 show how Equation 2 works for the profiles in Figures 3 and 4. Unnormalized posterior probabilities calculated according to Equation 2 are plotted as horizontal bars over their respective intervals. The points to note are (1) the relative heights of the bars in Figure 6 and (2) the fact that the probability distributions in Figure 7 are bimodal.

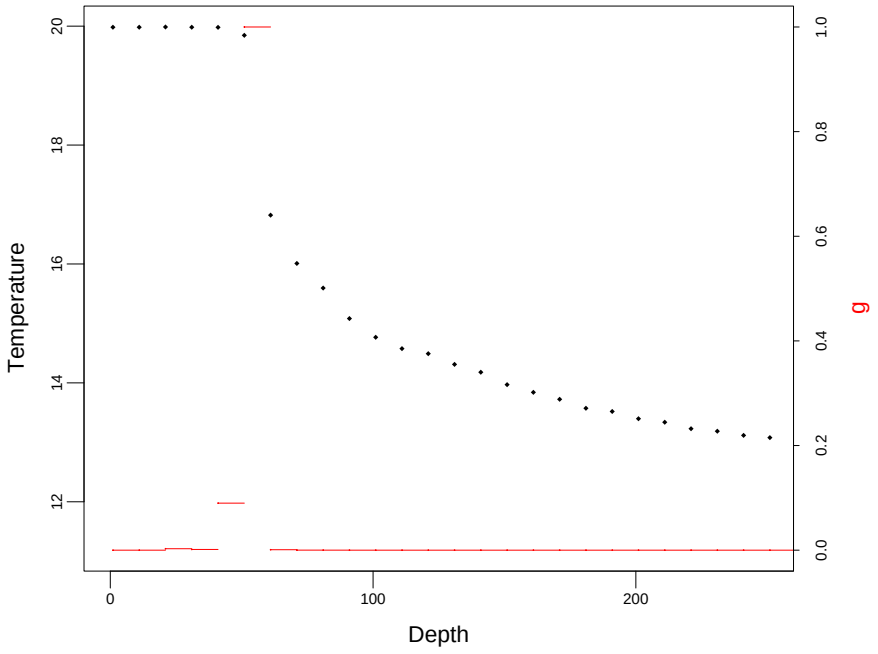

Figure 6: The profile of Figures 3. The horizontal bars show posterior probabilities of intervals calculated according to Equation 2, assuming a uniform prior. The bar over I_6 is approximately 10 times as high as the bar over I_5. The oceanographer agrees that the horizontal bars approximately match her subjective evaluation.

Two more facts are needed to complete the specification of the likelihood.

1. Because we did not tell the oceanographer the time of year or physical location of the profiles, her prior for M was approximately uniform.

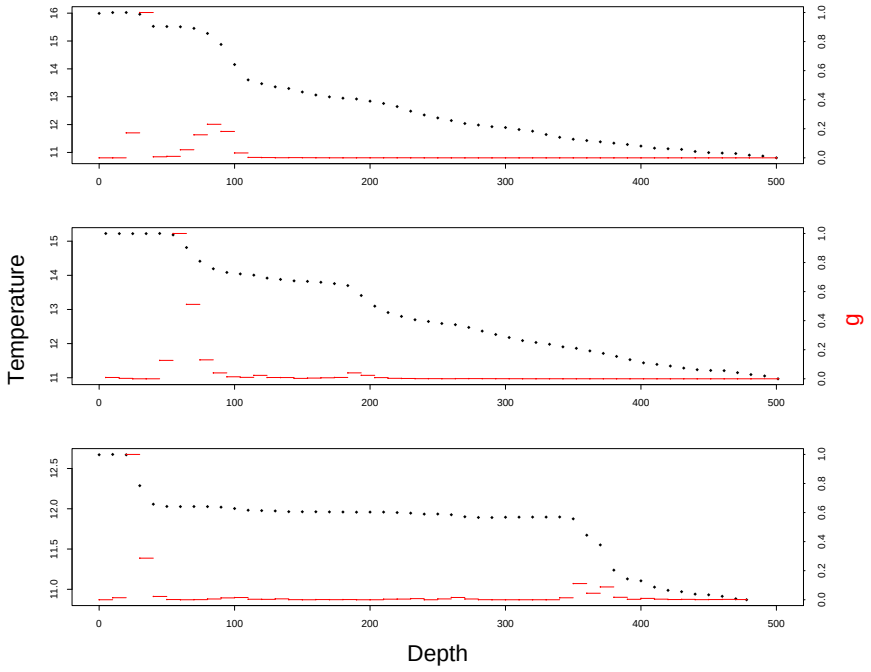

Figure 7: The three profiles of Figure 4. The horizontal bars show posterior probabilities of intervals calculated according to Equation 2, assuming a uniform prior. The oceanographer agrees that the horizontal bars approximately match her subjective evaluation.

(She agrees.)

2. Her posterior for M is approximately uniform within each interval. (She agrees.)

Equation 2 and Fact 2 specify the posterior density of M. It is piecewise constant on the intervals I_j; its value on I_j is

$$p'(m \mid \vec{y}) = \frac{\exp\{-2\Delta_{1j}\} \times [1 - \exp\{-\Delta_{2j}/0.3\}]}{(d_{j+1} - d_j) \sum_k (\exp\{-2\Delta_{1k}\} \times [1 - \exp\{-\Delta_{2k}/0.3\}])} \tag{3}$$

for $m \in I_j$. And since the prior was uniform, Equation 3 is also the likelihood function.

More formally, for any depth m let $j(m)$ be the interval containing m, i.e., $m \in I_{j(m)}$. Then we define the subjective likelihood function $\ell'(m)$ by

$$\ell'(m) \propto \frac{\exp\{-2\Delta_{1j(m)}\} \times [1 - \exp\{-\Delta_{2j(m)}/0.3\}]}{d_{j(m)+1} - d_{j(m)}} \tag{4}$$

Equation 4 gives us a rule for computing the likelihood function for any profile. In subsequent sections we apply the rule to a large collection of profiles, more than the oceanographer can assess individually.

4.2 Semiformal Justification

Bayesian analysis derives from the joint distribution $p(\vec{y}, M)$. To arrive at $p(\vec{y}, M)$ one typically specifies $p(M)$ and $p(\vec{y} \mid M)$. Our analysis also uses $p(\vec{y}, M)$ but arrives at it differently.

Since the oceanographer cannot tell us directly about $p(\vec{y}, M)$ or $p(\vec{y} \mid M)$, but can tell us about the relationship between M and $(\vec{\Delta}_1, \vec{\Delta}_2)$, we work with a transformation of variables

$$(\vec{y}, M) \longleftrightarrow (M, \vec{\Delta}_1, \vec{\Delta}_2, \vec{s}),$$

find the joint distribution of $(M, \vec{\Delta}_1, \vec{\Delta}_2, \vec{s})$, and derive that of $(\vec{y}, M)$ by back transformation. Specifically,

$$(y_1, \ldots, y_n, m) \longleftrightarrow (m, \Delta_{1j(m)}, \Delta_{2j(m)}, s_1, \ldots, s_{n-2})$$

where

$$s_k = \begin{cases} y_{d_k} - y_{d_{k-1}} & \text{for } k = 1, \ldots, j(m) - 1, \\ y_{d_{k+2}} - y_{d_{k+3}} & \text{for } k = j(m), \ldots, n - 2. \end{cases}$$

For later use, we note that for a fixed m this is a linear transformation and the absolute value of the Jacobian is $1/(d_{j(m)+1} - d_{j(m)})$. But the transformation and Jacobian vary according to $j(m)$.

The joint density of $(M, \vec{\Delta}_1, \vec{\Delta}_2, \vec{s})$ is given by

$$\begin{aligned} p(m, \Delta_{1j(m)}, \Delta_{2j(m)}, \vec{s}) &= p(m)p(\Delta_{1j(m)}, \Delta_{2j(m)} \,|\, m)p(\vec{s} \,|\, m, \Delta_{1j(m)}, \Delta_{2j(m)}) \\ &\approx p(m)p(\Delta_{1j(m)}, \Delta_{2j(m)})p(\vec{s}) \\ &\propto p(m)p(\Delta_{1j(m)}, \Delta_{2j(m)}) \\ &= p(\Delta_{1j(m)}, \Delta_{2j(m)}) \end{aligned}$$

The second line follows from the oceanographer's judgement that $M \perp (\vec{\Delta}_1, \vec{\Delta}_2)$ and $\vec{s} \perp (M, \vec{\Delta}_1, \vec{\Delta}_2)$, at least approximately. The third line follows because, while $\vec{s}$ does depend on m — i.e., $\vec{s} = h(m, \vec{y})$ for some function h — the dependence of $p(\vec{s})$ on m is judged so slight as to be ignorable. The term $p(\Delta_{1j(m)}, \Delta_{2j(m)})$ is retained because it does contain useful information about m. The fourth line follows because in this section $p(m) = 1$ in accord with the conditions of elicitation.

The oceanographer's judgment expressed in Equation 2 is equivalent to setting

$$\begin{aligned} p'(\Delta_{1j(m)}, \Delta_{2j(m)}) &\propto g(\Delta_{1j(m)}, \Delta_{2j(m)}) \\ &= \exp\{-2\Delta_{1j(m)}\} \times [1 - \exp\{-\Delta_{2j(m)}/0.3\}] \end{aligned}$$

The joint density of $(\vec{y}, M)$ comes from the transformation

$$(M, \vec{\Delta}_1, \vec{\Delta}_2, \vec{s}) \to (\vec{y}, M).$$

Accounting for the Jacobian, it is

$$p'(\vec{y}, m) \propto \frac{\exp\{-2\Delta_{1j(m)}\} \times [1 - \exp\{-\Delta_{2j(m)}/0.3\}]}{d_{j(m)+1} - d_{j(m)}} \tag{5}$$

Since the prior was uniform, this is also the likelihood function, as heuristically justified in Equation 4.

Starting from Equation 5 one could, if one wished, try to derive $p(\vec{y} \,|\, m)$, the sampling model implied by our subjective elicitation exercise. In practice, since we already have $\ell(m)$, there is little need to do so.

5 Three Modifications

The probabilities and likelihoods derived in Section 4 are not the final word. Three more considerations come into play.

1. Of all the measurements in a temperature profile, the surface measurement is the most noisy. Furthermore, the surface measurement can be influenced by daily warming; it may undergo a temperature cycle each day, and measurements made during the warm part of the cycle may not reflect fundamental properties of the water column. The key to recognizing such a situation is $t_1 > t_2 \approx t_3$.

   When this happens, $\vec{\Delta}_1 = (0, y_1 - y_2, \ldots, y_1 - y_n)$ is not a good reflection of the information in the profile. So, for profiles where $y_1 = \max y_i$ and $|y_2 - y_3| < 0.05$, we set $\vec{\Delta}_1 = (0, y_2 - y_2, \ldots, y_2 - y_n)$.

2. The idea behind $\vec{\Delta}_{2j} = (y_j - y_{j+1})/(d_{j+1} - d_j)$ is, obviously, that an interval containing M is likely to have a large temperature drop. But the oceanographer's judgment is that temperature drops bigger than about 0.3 °C are not more indicative of M than drops of about 0.3 °C. Therefore, we set $\vec{\Delta}_{2j} = \min((y_j - y_{j+1}), 0.3)/(d_{j+1} - d_j)$.

3. Equation 2 was constructed to mimic the oceanographer's assessments of high probability intervals. For low probability intervals it accurately captures the fact that they have low probability, but might misstate their probability ratios by several orders of magnitude. In addition, it is widely believed that subjective probability assessments are often too sharp; they understate the true amount of uncertainty. The result in our analysis could be oversensitivity to outliers, or intervals with low probabilities. Therefore, we modify Equation 3 and set

$$p(m \mid \vec{y}) = \max[p'(m), .01p'(\hat{m})] \tag{6}$$

6 Model and Prior

The observations $\vec{y}_1, \ldots, \vec{y}_T$, shown in Figure 1, run over a time index $t = t_1, \ldots, t_T$ that spans multiple years, as shown in Figure 2.

Figure 8 shows the result of applying Equation 2 to each profile separately. Each vertical bar is the maximum probability interval of some profile; its

extent is indicated on the ordinate. Its abcissa is the day-of-year when that profile was recorded. (The curve is a posterior mean and will be discussed later.) There are two features of note. First, there is an annual cycle with deeper mixed layers in the winter. And second, there is greater variability in the winter. Our model will accomodate both features.

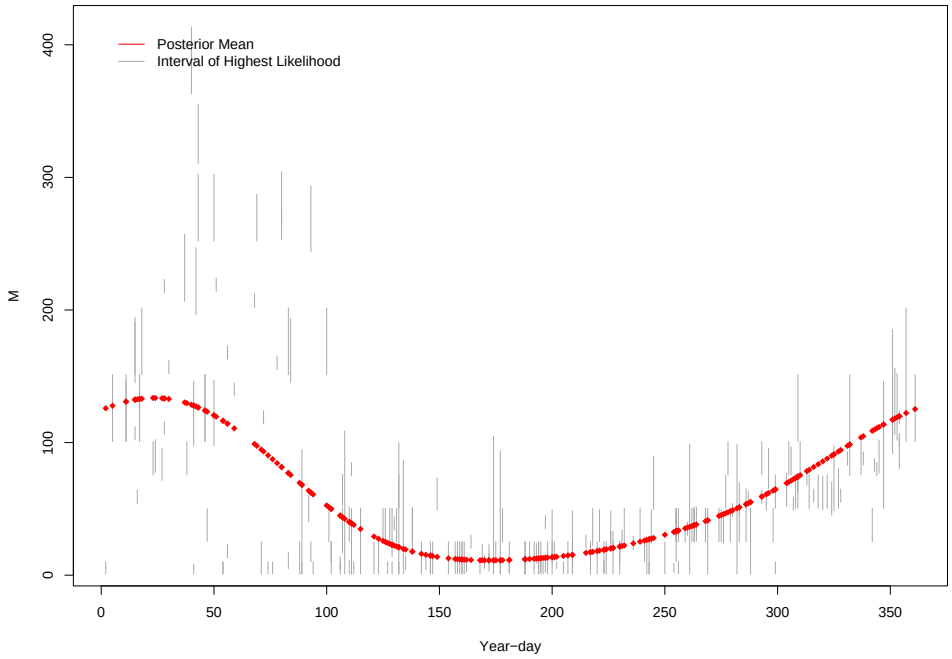

Figure 8: Vertical bars are maximum probability intervals of all profiles versus day-of-year. The curve is the posterior mean of $\mu(t)$.

Let $M(t)$ be the mixed layer depth at time t. It is apparent that $M(t)$ undergoes an annual cycle, the precise details of which are unknown and may vary from year to year.

Mean annual cycle We use μ to denote the mean annual cycle. Thus for any time t,

$$\mu(t) = \mu(t \bmod 365).$$

We model μ with a *process convolution* (See Higdon [1998] for details.), as described in the next paragraph.

Let $0 \leq v_1 < v_2 < \cdots < v_{12} \leq 365$ be an equally spaced sequence of points (mod365); let $x_1, \ldots, x_{12}$, be values "at" those points; and let k be a kernel. Our model is

$$\mu(t) = \sum_{\ell=1}^{12} x_\ell k(t - v_\ell)$$

The x_ℓ's are modelled as unknowns to be fit from the data. A priori,

$$x_1, \ldots, x_{12} \sim \text{i.i.d. } \mathbb{N}(0, \sigma_x^2)$$

The posterior means of the x_ℓ's yields the estimate $\hat{\mu}(t) = \sum \hat{x}_\ell k(t - v_\ell)$ plotted in Figure 8.

Deviations from μ The mixed layer deepens in the Fall as atmospheric temperatures decrease. Colder air means colder sea surface temperature, which in turn means that surface waters become dense and sink. The sinking causes surface waters to mix with deeper waters. The process continues through the winter, leading to deeper mixing and increased values of M. The effect is reversed in the Spring as surface waters warm. Spring heat is conveyed downward through diffusion, a much slower process than Fall's convection. Thus the annual cycle is asymmetric and the summer mixed layer is relatively shallow and stable.

Because weather varies from year to year, the actual cycle for M in a particular year may differ from μ, especially in the winter. To account for those differences we allow each winter month (November through April) to have its own random effect. Our notation for the random effect of month m in year y is $b_{y,m}$. Letting $y(t)$ and $m(t)$ denote the year and month of time t, respectively, the random effect associated with a profile taken at time t is $b_{y(t),m(t)}$.

In addition to annual variations due to weather, we are interested in a possible secular trend. We account for that in our model with a linear regression term βt. We will be interested in the posterior distribution of β.

Finally, to account for M's greater variability in the winter than in the summer, we use a piecewise constant variance: σ^2 in the summer and $3\sigma^2$ in the winter. Modelling the variance as piecewise constant is crude but, we believe, effective.

Putting everything together, the full model and prior are

$$\begin{gathered}
\sigma_x^2, \sigma_b^2, \sigma_\beta^2, \sigma^2 \sim \text{i.i.d. InvGam}(.001, .001) \quad \text{mutually independent} \\
x_1, \ldots, x_{12} \,|\, \sigma_x^2, \sigma_b^2, \sigma_\beta^2, \sigma^2 \sim \text{i.i.d. } \mathbb{N}\left(0, \sigma_x^2\right) \\
\mu(t) = \sum_{\ell=1}^{12} x_\ell k(t - v_\ell) \\
\{b_{y,m}\} \,|\, \sigma_x^2, \sigma_b^2, \sigma_\beta^2, \sigma^2, \vec{x} \sim \text{i.i.d. } \mathbb{N}\left(0, \sigma_b^2\right) \\
\beta \,|\, \sigma_x^2, \sigma_b^2, \sigma_\beta^2, \sigma^2, \vec{x}, \{b_{y,m}\} \sim \mathbb{N}(0, \sigma_\beta^2) \\
\nu(t) = \mu(t) + b_{y(t),m(t)} + \beta t \\
\tau(t) = \begin{cases} \sigma^2 & \text{if } t \text{ is in winter} \\ 3\sigma^2 & \text{if } t \text{ is in summer} \end{cases} \\
M_{t_i} \,|\, \sigma_x^2, \sigma_b^2, \sigma_\beta^2, \sigma^2, \vec{x}, \vec{b}, \beta \sim \mathbb{N}(\nu(t_i), \tau(t_i)) \quad \text{mutually independent} \\
\vec{y} \,|\, \sigma_x^2, \sigma_b^2, \sigma_\beta^2, \sigma^2, \vec{x}, \vec{b}, \beta, \vec{M} \sim \text{subjective likelihood, mutually independent}
\end{gathered} \tag{7}$$

7 Computations and Posterior

The full conditional distributions are available for all parameters except the M_{t_i}'s. Therefore, one efficient method to sample from the posterior distribution is a Gibbs sampler with Metropolis-Hasting steps for the M_{t_i}'s.

Conditionally on current values of $\nu(t_i)$, we propose a new move $M_{t_i}^* \sim \mathbb{U}[0, \text{ocean floor}]$ and accept the move with probability

$$\min\left\{1, \frac{\ell(M_{t_i}^*)\exp\{-\frac{(M_i^* - \nu(t_i))^2}{2\tau(t_i)^2}\}}{\ell(M_{t_i})\exp\{-\frac{(M_i - \nu(t_i))^2}{2\tau(t_i)^2}\}}\right\}$$

where ℓ is the subjective likelihood function defined by Equation 4.

Figure 8 shows $\hat{\mu}(t) = \sum \hat{x}_\ell k(t - v_\ell)$, the estimate of μ calculated from the posterior means $\hat{x}_\ell$. Note that it is asymmetric, as expected from our understanding of the physical process. The asymmetry is partial justification for modelling μ as a process convolution rather than a summation of sinusoids. Figure 9 is another way to view the fit of the model. For each profile it shows the posterior mean of ν_i on the abcissa and the maximum probability interval according to Equation 2 on the ordinate. Figure 9 indicates an

overall reasonably good fit and heteroscedasticity associated with the deeper mixed layers of winter. The x_ℓ's themselves are not of interest, so we don't show any pictures.

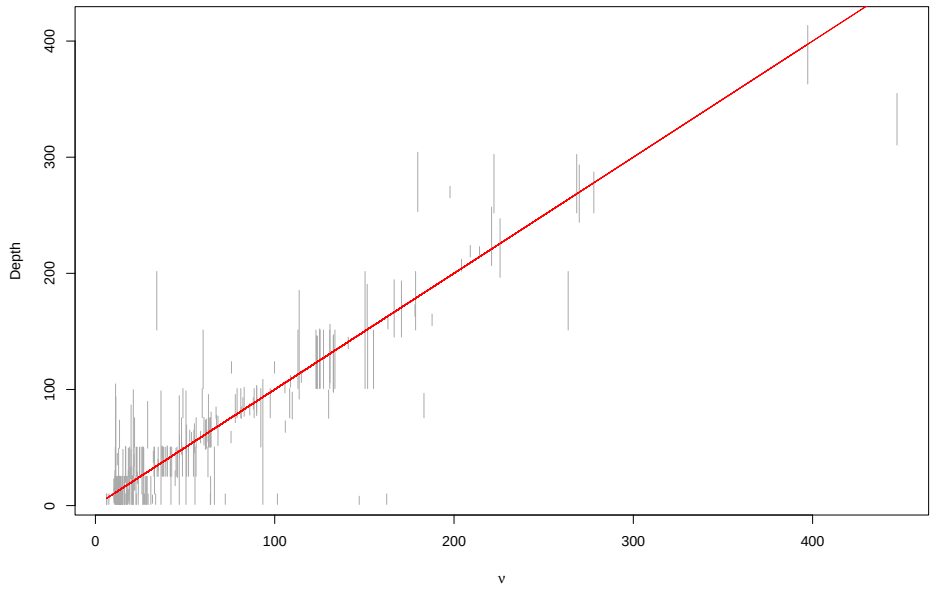

Figure 9: Maximum probability interval from Equation 2 vs. posterior mean of ν_i.

One way to look for a secular trend is to fit a model identical to Equation 7 but without β and examine the random effects for evidence of trend. Figure 10 does just that; it shows the posterior mean $b_{y,m}$'s from such a fit. For each month they are plotted as a function of year. There is no evidence of trend except possibly for December. However, the apparent trend in December is due to confounding. December measurements in recent years tended to be early in the month, hence with shallower mixed layers and lower M's, and therefore with negative $b_{y,m}$'s.

There is evidence of greater variability in the late winter since about 1990. We're not sure why that is, or even whether it's real.

Another way to look for a secular trend is to examine the posterior distribution of β. See Figure 11. Again, there is no evidence of trend.

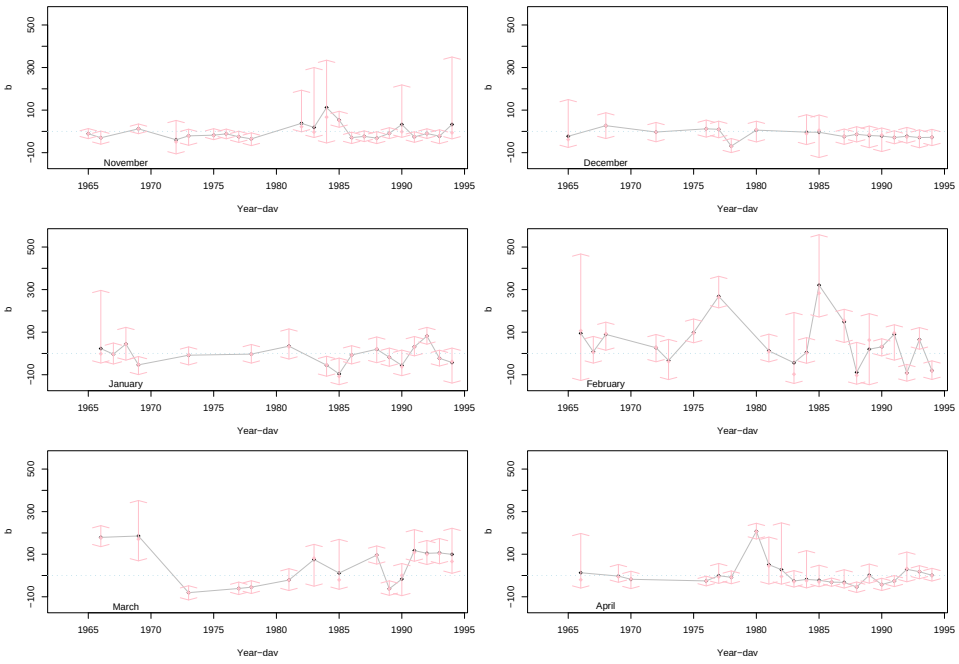

Figure 10: Random effects: $\beta_{y,m}$ vs. year

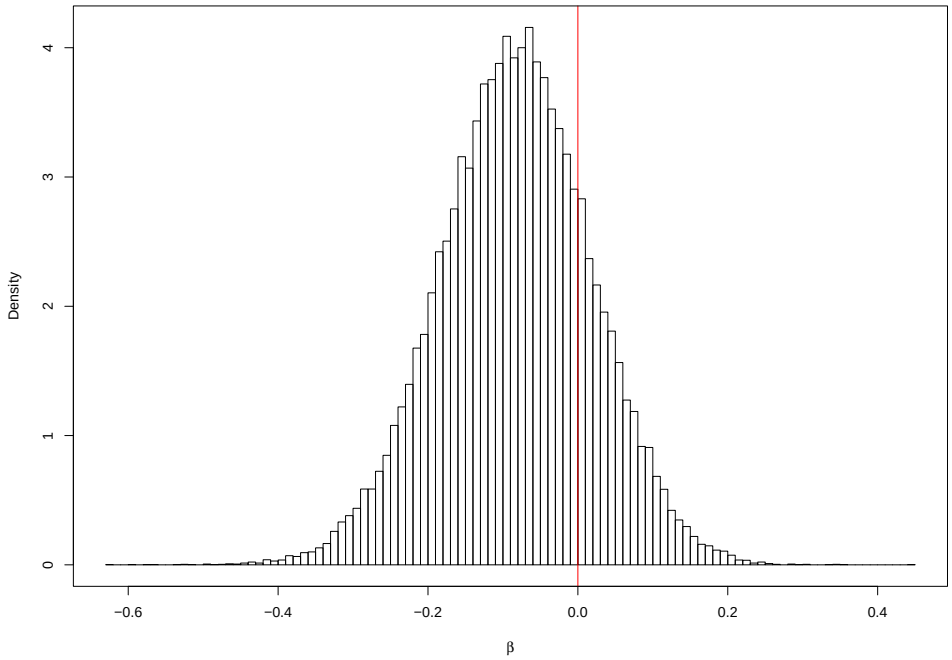

Figure 11: Posterior distribution of β: draws from the MCMC sampler.

8 Discussion

Foundations We like our subjective likelihood approach in this analysis because it incorporates expert judgment in the most direct way we can imagine. We find it an elegant and practical solution to a difficult modelling problem. Of course, it does not have the usual philosophical underpinnings of the customary Bayesian analysis. Also it raises the question: if the oceanographer can give her posterior for M_{t_i} after seeing one profile, why can't she give her posterior for β after seeing them all? We think the answer is obvious. First, studying hundreds of profiles in detail (thousands when we analyze the entire Atlantic Ocean) is too difficult, and second, carefully assimilating information from all the profiles into an opinion about β is far removed from oceanographers' experience.

We could have analyzed the data with a nonparametric change-point model in which temperature is roughly constant down to a depth M, then decreases monotonically, but otherwise nonparametrically, to the ocean floor. We tried and rejected such an approach for two reasons. First, it yields results for single profiles that disagree with the expert's opinions, and second, it is a more circuitous use of expert opinion than the subjective likelihood function. Our nonparametric change-point model is described in the Appendix.

Will subjective likelihood be accepted by the scientific community; will there be other applications; will someone discover a fatal flaw? We don't know. We hope the answers are yes, yes, and no; but time will tell.

Modelling We made some fairly crude modelling choices. In particular, there are probably better and more sophisticated ways to handle random effects, heteroscedasticity, and the secular trend. We made our choices for simplicity and in the belief that more sophistication would not change the analysis very much. We acknowledge that improvements are possible.

Our model treats the profiles as conditionally independent given $\vec{M}$. That's a questionable choice. Our expert says that the time constant of the ocean, at least for this purpose, is about 5 days or so. Profiles taken on consecutive days are probably created and influenced by almost the same physical forces and are certainly dependent. Whether they are conditionally dependent given M is another matter. In any case, there are few profiles in our data set at time intervals of less than a week, so we are willing to treat them as conditionally independent. A possible model enhancement is to use a temporal covariance function with a range on the order of about 5 days.

Another way to improve the model is to include more covariates. M is related to air temperature, so one might try including covariates such as a running mean monthly temperature, etc. But temperature covariates will be highly correlated with time of year, so the potential for improved fit might be only minor. M is also thought to be related to El-Nino/La-Nina, so the NAO (North Atlantic Oscillation) index might be a good covariate to try.

Diagnostics and Robustness Figures 8 and 9 show overall goodness of fit. Deviations from the fitted curves are a bit hard to interpret because the likelihood function is not symmetric and need not even be unimodal. For example, Figures 8 and 9 both show three profiles in which the interval of highest probability is deeper than 300m, yet the fitted value for one is around 80m and for the other two around 120m. The model appears not to fit those profiles well. What Figures 8 and 9 do not reveal is that each of those profiles has a secondary probability mode, according to Equation 2, so the apparent misfit is not so severe after all. Other profiles taken around the same dates as these three have only shallower modes; so the model favors the shallower modes in its posterior mean.

Determination of M can sometimes be complicated by daily warming of surface waters. During a hot, sunny day surface waters can warm as much as perhaps a couple of degrees down to a depth of perhaps 20m or so, then cool again during the night. Consequently, a profile taken in midafternoon might incorrectly indicate that M appears to be less than 20m. Our expert oceanographer was aware of this possibility when assessing the profiles we showed her. The coefficient of Δ_1 in Equations 2 and 5, i.e. -2, reflects her belief that temperature drops in the upper ocean are probably due to mixed layers, not daily warming. But perhaps she's overly confident on this point. As a robustness check, we recalculated our posterior using a coefficient of -0.15. It made little difference to the inference regarding trend.

Our agreement that $\ell(m)/\ell(\hat{m}) \geq .01$ ensures that the posterior will not be too heavily influenced by a small number of profiles. Still, more thorough model diagnostics are useful and will be reported elsewhere.

9 Appendix

This Appendix briefly describes a nonparametric changepoint model as an alternative to the subjective likelihood analysis. In the end we prefer subjec-

tive likelihood because (1) it directly tackles the strength of evidence problem and (2), the changepoint model doesn't accurately capture the expert's posterior opinion. In this Appendix it suffices to consider a single profile at a single location in space and time.

We consider the profile, temperature as a function of depth, or $t(d)$, to be a realization of a random function. The function t is continuous, flat on $[0, M]$ and monotonically decreasing on $[M, \text{ocean floor}]$. If temperature is properly rescaled to the unit interval then t is a cdf, and we model it with a Dirichlet process. (See Lavine and Mockus [1995] for details of modelling monotone functions with Dirichlet processes.) We take M to be random, and the shape parameter of the Dirichlet process to be flat on $[0, M]$ and exponential on $[M, \text{ocean floor}]$, to agree with first order physical calculations. We tried three different values for the total mass parameter of the Dirichlet process, $\alpha = .005, .01, .1$. Finally, we model the observations as

$$y_i \sim \mathbb{N}(t(d_i), \sigma^2) \quad \text{mutually independent given } t$$

The posterior is calculated by MCMC. Figures 12 and 13 show results for two profiles. Figure 12 is the same profile shown in Figure 3; Figure 13 is the third profile in Figure 4. The curves in Figures 12 and 13 are posterior densities of M from the Dirichlet process changepoint model.

The expert's assessment of Figure 12 is that

$$\Pr[M \in I_6] \approx 10 \Pr[M \in I_5] \quad \text{and} \quad \Pr[M \notin I_5 \cup I_6] = \text{small}$$

But the Dirichlet process changepoint model gives most of its posterior mass to I_5.

The expert's assessment of Figure 13 is bimodal; there are two likely locations for M, one around 30m and one around 350m. However, as the figure shows, the Dirichlet process changepoint posterior for M is concentrated around one or the other of the two regions, but does not have two modes.

While it might be possible to refine the changepoint model to better reflect expert opinion, we are not sure how to go about it. Since the expert's opinion is about the posterior and not about the sampling model, we thought it more straightforward and justifiable to model that opinion directly through subjective likelihood rather than indirectly through a sampling model.

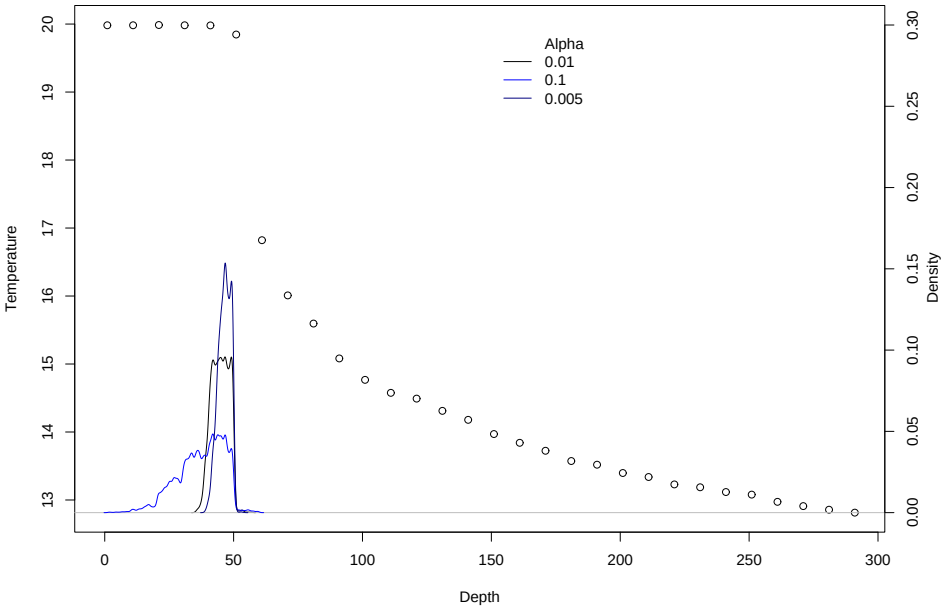

Figure 12: A profile and three posterior densities for M, according to the Dirichlet process changepoint model.

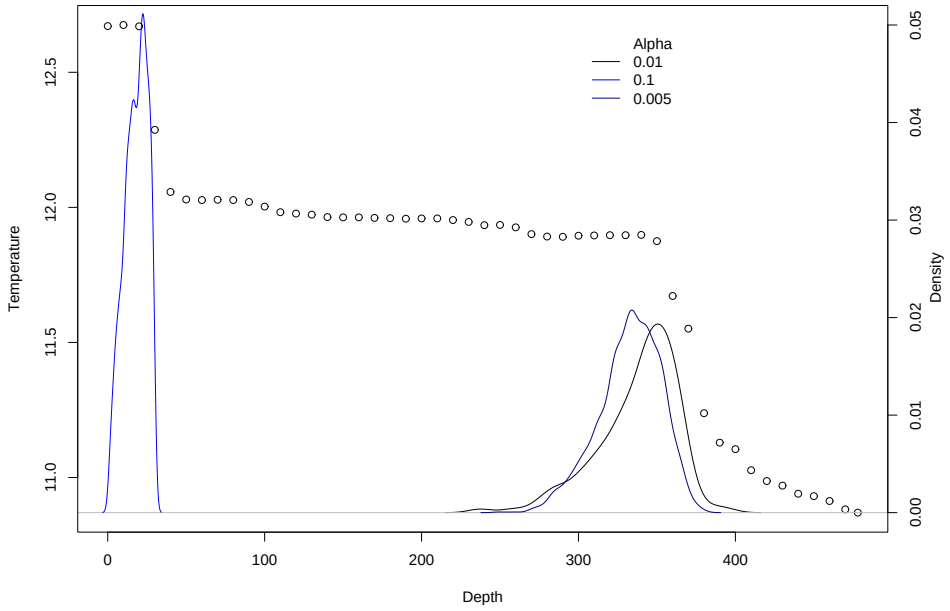

Figure 13: A profile and three posterior densities for M, according to the Dirichlet process changepoint model.

References

David Higdon. A process-copnvolution approach to modelling temperatures in the North Atlantic ocean. *Environmental and Ecological Statistics*, 5: 173–190, 1998.

Michael Lavine and Audris Mockus. A nonparametric Bayes method for isotonic regression. *jspi*, 46:235–248, 1995.

S. Levitus, J.I. Antonov, T.P. Boyer, and C. Stephens. Warming of the world ocean. *Science*, 294:840–842, 2000.

S. Levitus, J.I. Antonov, J. Wang, T.L. Delworth, K.W. Dixon, and A.J. Broccoli. Anthropogenic warming of earth's climate system. *Science*, 292: 267–270, 2001.

G.L. Mellor. *Introduction to Physical Oceanography.* Woodbury, American Institute of Physics, NY, 1996.