

Landes, Reid; Loutzenhiser, Peter; Vardeman, Stephen

Working Paper

Hierarchical Bayes statistical analyses for a calibration experiment

Technical Report, No. 2004,14

Provided in Cooperation with:

Collaborative Research Center 'Reduction of Complexity in Multivariate Data Structures' (SFB 475),
University of Dortmund

Suggested Citation: Landes, Reid; Loutzenhiser, Peter; Vardeman, Stephen (2004) : Hierarchical Bayes statistical analyses for a calibration experiment, Technical Report, No. 2004,14, Universität Dortmund, Sonderforschungsbereich 475 - Komplexitätsreduktion in Multivariaten Datenstrukturen, Dortmund

This Version is available at:

<https://hdl.handle.net/10419/49327>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Hierarchical Bayes Statistical Analyses for a Calibration Experiment

by

Reid Landes^{*}

Peter Loutzenhiser^{**}

Stephen Vardeman^{***}

Iowa State University

February 9, 2004

Abstract

We consider hierarchical Bayes analyses of an experiment conducted to enable calibration of a set of mass-produced resistance temperature devices (RTDs). These were placed in batches into a liquid bath with a precise NIST-approved thermometer, and resistances and temperatures were recorded approximately every 30 seconds. Under the assumptions that the thermometer is accurate and each RTD responds linearly to temperature change, we use hierarchical Bayes methods to estimate the parameters of the linear calibration equations. Predictions of the parameters for an untested RTD of the same type, and interval estimates of temperature based on a realized resistance reading are also available (both for the tested RTDs and for an untested one produced under the same production process conditions).

^{*} Reid Landes is a Ph.D. candidate in Statistics.

^{**} Peter Loutzenhiser is a Ph.D. student in Mechanical Engineering.

^{***} Stephen Vardeman is a Professor of Statistics and Industrial Engineering and gratefully acknowledges the financial support of the Deutsche Forschungsgemeinschaft (SFB 475, Reduction of Complexity for Multivariate Data Structures) through the University of Dortmund.

I. INTRODUCTION

In some statistical calibration problems, it is of interest to extend the calibration from a set of devices tested to other devices made under the same conditions. (We will refer to the calibration equations on the original set of devices as “direct” calibrations, and the extension of these equations to similar devices as “indirect” calibration.) Both direct and indirect calibrations are subject to error, and the precision with which the parameters of calibration equations are known is therefore of interest. (Of course, parameters of an indirectly determined calibration equation are known less precisely than those from the direct calibrations.) In addition, users of devices that have been calibrated (either directly or indirectly) often want to know the precision with which they know a measurand – the quantity being measured – based on a value read from a device in hand. Again, this precision differs depending on whether the device was directly or indirectly calibrated.

How to produce estimates and predictions of interest (particularly interval versions of these and especially those concerning an indirect calibration) via classical (or “frequentist”) statistical methods is not completely clear. (Even when the model for read values for a particular measurand is a simple linear regression model, this turns out to be a problem involving nonlinear random effects.) However, in a Bayesian framework, producing such measures of precision and such intervals is, at least in principle, straightforward. One simply consults the posterior distribution of any quantity of interest (its conditional distribution given the data from calibration studies). We present a case study (of a real multiple-device calibration study) illustrating how this can be

accomplished in practice using WinBUGS, a (currently free) Bayesian statistical software package.

An interesting byproduct of our analysis is that in cases like the one discussed, where calibration data on multiple devices is simultaneously collected for comparison to that from a single accurate (and typically more precise) “standard device” or “reference point,” it is not only possible to do the calibrations, but also to simultaneously estimate the precision of the standard device.

II. THE CALIBRATION EXPERIMENT

The case we will discuss concerns a calibration experiment for a set of interchangeable resistance temperature devices (RTDs) using a single NIST-approved thermometer. The RTDs under consideration were used in a heating, ventilation, and air conditioning (HVAC) research facility affiliated with the Iowa Energy Center. Large-scale tests requiring simultaneous temperature measurements at multiple locations in an HVAC system are performed at this facility.

[1] describes the physical methods used in temperature sensor calibration work at the Iowa Energy Center. In order to calibrate the RTDs (that we will assume came from a single set of production process conditions), subsets of them were placed in a liquid bath alongside the thermometer. Observations from each of the RTDs (resistances) and the thermometer (in degrees Fahrenheit) were taken over time. The readings from the RTDs were synchronized and taken every 30 seconds. Those of the thermometer were on a different (and not always completely regular) schedule. Times between thermometer readings were usually 30 seconds, but on occasion were as little as 28 seconds or as much

as 31 seconds. The data set we analyzed consists of series of thermometer readings (taken roughly every 30 seconds) paired with sets of RTD readings taken at time points no more than 15 seconds away from the thermometer readings.

The bath temperature was purposely manipulated by the researchers over the course of the calibration experiment. We have not used data from “transition periods” when energy was being actively put into or taken out of the bath in order to change the temperature. Rather, we have used only data from “steady state/equilibrium” portions of the experiment, where the bath temperature control system was attempting to hold a constant temperature. The approximately 30-second period between observations was long enough that we found no statistically significant auto-correlations in the (several) steady state thermometer reading series.

In the part of the data set we consider, a dozen RTDs are represented. These were tested in subsets of four in four different baths. One subset was subjected to two baths, the remaining two subsets were subjected to a single bath each. Across the baths, the minimum and maximum temperatures used in the calibration experiment were 22 and 104 °F. The within-bath temperatures had ranges of between 9 and 64 °F. The minimum number of observation vectors (thermometer temperature and four RTD resistances) from a bath was 180, and the maximum was 781. Altogether, in the data set we consider, there were a total of 2146 observation vectors.

III. STATISTICAL MODELS AND METHODS

We assume that the RTDs all respond linearly to changes in real temperature over the range used in the experiment, and that those linear relationships between real

temperature and mean resistance did not change in time. Further, we will assume that the thermometer was an accurate indicator of temperature in °F (i.e. the expected temperature reading was the true temperature). Under these assumptions we proceed to specify a probability model for our data and specify prior distributions for the model parameters.

Let

T_{ij} = the true bath temperature at time-point j in bath i

X_{ij} = the temperature read from the thermometer at time-point j in bath i

Y_{hij} = the resistance read at time-point j in bath i on RTD h

A model for thermometer readings is

$$X_{ij} = T_{ij} + \eta_{ij} \quad (1)$$

where we assume that the η_{ij} are iid (independent and identically distributed as)

$N(0, \sigma_\eta^2)$. (This is commonly known in the statistical literature as an additive

measurement error model.) Conditioned on the intercept b_{1h} and slope b_{2h} peculiar to

RTD h , we will model the resistance readings from that RTD as

$$Y_{hij} = b_{1h} + b_{2h}T_{ij} + \varepsilon_{hij} \quad (2)$$

where we assume that the ε_{hij} are iid $N(0, \sigma_\varepsilon^2)$ and independent of the η_{ij} . Note that

using equation (1), equation (2) may be rewritten as

$$Y_{hij} = b_{1h} + b_{2h}(X_{ij} - \eta_{ij}) + \varepsilon_{hij} \quad (3)$$

(This explicitly displays the role of the potential measurement error of the thermometer in determining what is observed.)

Finally, to complete the hierarchical model for the observables in terms of random effects, we assume that the intercept-slope pairs for the RTDs are themselves random draws from some distribution. Specifically suppose that

$$\begin{pmatrix} b_{1h} \\ b_{2h} \end{pmatrix} \sim \text{iid MVN}_2(\boldsymbol{\beta}, \boldsymbol{\Sigma}) \quad (4)$$

independent of the η_{ij} and ε_{hij} .

To this point, the model specified by displays (3) and (4) has parameters $\sigma_\eta^2, \sigma_\varepsilon^2, \boldsymbol{\beta}$, and $\boldsymbol{\Sigma}$. These describe respectively the imprecision of the thermometer, the variability in measured resistance of a given RTD at a fixed temperature, the mean of the intercepts and slopes for the “real” calibration equations across RTDs, and the variability in those intercept-slope pairs. This is a nonlinear random effects model. (See, for example, [2].)

We take a Bayesian approach to the analysis of data based on this model and specify a “prior” (joint) distribution for the model parameters. (A prior distribution is a “pre-data” distribution meant to describe what are usually somewhat vague pre-experimental beliefs about model parameters. It is typical in defining a joint prior to specify some marginal distributions, and then adopt a product or independence form based on these.)

To begin, we used the prior distribution

$$\sigma_\eta^2 \sim \text{InvGamma}(a_\eta, b_\eta) \quad (5)$$

This means that the reciprocal of σ_η^2 – the precision of the η distribution – has a $\text{Gamma}(a_\eta, b_\eta)$ prior. We chose parameters a and b such that the mean of the prior

distribution for σ_η^2 (say $m = b/(a-1)$) was equal to the thermometer manufacturer's reported measurement variance (namely .000891), and the variance of the prior (say $v = b^2 / ((a-1)^2 (a-2))$) was relatively large (we used the value 10). (Setting $m = b/(a-1)$ and $v = b^2 / ((a-1)^2 (a-2))$ and solving for a and b gives $a = 2 + m^2 / v$ and $b = m(1 + m^2 / v)$. We thus used $a_\eta = 2.000000008$ and $b_\eta = .0008910007$. We similarly used a prior distribution

$$\sigma_\varepsilon^2 \sim \text{InvGamma}(a_\varepsilon, b_\varepsilon) \quad (6)$$

with mean 1 and variance 1000.

Finally, to complete the specification of the prior, we modeled the mean intercept-and-slope vector as

$$\boldsymbol{\beta} \sim \text{MVN}_2(\mathbf{0}, 10^6 \mathbf{I}) \quad (7)$$

and modeled the variance-covariance matrix of the intercept-and-slope vectors as

$$\boldsymbol{\Sigma} \sim \text{InvWishart}_2(\boldsymbol{\Omega}, \lambda) \quad (8)$$

for appropriate $\boldsymbol{\Omega}$ and λ . (Here $k = 2$ is the number of rows and columns in $\boldsymbol{\Sigma}$.)

Provided $\boldsymbol{\Omega}$ is positive definite and the “degrees of freedom” $\lambda > k + 1 = 3$, this type of distribution over covariance matrices has mean $(\lambda - k - 1)^{-1} \boldsymbol{\Omega}$. For $\lambda > k + 3$ the elements of $\boldsymbol{\Sigma}$ have finite variances. As $|\boldsymbol{\Omega}^{-1}| = |\boldsymbol{\Omega}|^{-1} \rightarrow 0$, and as $\lambda \rightarrow -1$, this distribution approaches the multivariate Jeffrey's density, which is a noninformative prior distribution for $\boldsymbol{\Sigma}$ (see [3], page 88).) Using the distributions easily available in

WinBUGS we let $\Sigma^{-1} \sim \text{Wishart}_2(10\mathbf{I}, 5.1)$, which is equivalent to

$$\Sigma \sim \text{InvWishart}_2(.1\mathbf{I}, 5.1).$$

The specifications (3)-(8) then (assuming prior independence) provide a (joint) probability model for the observables and parameters. Bayes inference for parameters makes use of their “posterior distribution,” – the conditional distribution of the parameters given the observables. Note that prediction of unobservable random effects or functions of them (even for RTDs not represented in the data set) is also possible by considering a joint distribution for parameters, observables, and random effects, and then the posterior distribution of parameters and random effects (the conditional distribution given observables).

For a calibration experiment involving t interchangeable devices, prediction of the random effects

$$\begin{pmatrix} b_{1h} \\ b_{2h} \end{pmatrix} \text{ for } h = 1, 2, \dots, t, t+1 \quad (9)$$

is of interest, the first t of these representing the direct calibrations and the $(t+1)^{\text{st}}$ representing an indirect calibration. Additionally, if Y_{new} is a new observed resistance value, and ε_{new} a corresponding unobserved new draw from the $N(0, \sigma_\varepsilon^2)$ distribution, the quantity

$$T_{\text{new}} = \frac{Y_{\text{new}} - \varepsilon_{\text{new}} - b_{1h}}{b_{2h}} \quad (10)$$

is also of interest. In the present context, this is the real temperature corresponding to the measured resistance, and its prediction can be made both for RTDs represented in the original data set and for the case of $h = t+1 = 13$ corresponding to an indirect calibration.

MCMC (Markov Chain Monte Carlo) methods implemented in WinBUGS (based on successive substitution/Gibbs sampling and Metropolis-Hastings algorithms) were employed to produce vectors with empirical distribution approximating the full posterior distribution. (See [4]. The software is presently available at no cost and can be downloaded from www.mrc-bsu.cam.ac.uk/bugs . Example WinBUGS code is included as an Appendix to this paper.) Three chains were run simultaneously and monitored (using “time series” plots of parameters and random effects of interest and values of the Gelman-Rubin statistic) for burn-in to a single posterior. (Initial values of variances σ_ε^2 and σ_η^2 for the three chains were 0.001, 1, and 1000. All other initial values were pseudo-randomly generated.) Upon burn-in, the sequences of vectors of parameters and random effects visited by the chain were thinned until autocorrelations for all variables were negligible. Updating of the three chains continued until an additional 25,000 vectors from each chain were produced at this sampling rate. Marginal posterior distributions were then estimated from the resulting 75,000 vectors. Estimates of the posterior moments and quantiles for quantities of interest were obtained from the approximate marginal posteriors.

IV. RESULTS

The analyses were done in WinBUGS Version 1.3. As indicated above, three MCMC chains (each with ~2,200 variables to update at each iteration) were run simultaneously. On a machine with a 3.0 GHz Pentium IV processor and 1 Gb of RAM, a burn-in of 254,000 iterations took approximately 4 hours. After the burn-in, every 25th iterate was saved until 25,000 vectors were obtained from each chain.

As point predictions of the calibration intercepts and slopes for the RTDs tested in the experiment, we used sample means from the estimated marginal posterior distributions. These point predictions and 95% credible limits for these (obtained from the lower and upper 2.5% points of the approximate posteriors) are collected in Table 1. Further, an indirect calibration was made for an additional hypothetical RTD produced under the process conditions represented by the tested thermometers. The 13th set of regression coefficients (marked with *) in Table 1 are the predictions for the hypothetical RTD.

For β , the bivariate point prediction (approximate posterior mean) was (1181, 9.190). A 95% credible set for β_1 was (1175, 1188), and the 95% credible set for β_2 was (9.076, 9.302).

TABLE 1. Summary Statistics for RTDs from WinBUGS Analysis

Node	Mean	2.5%	Median	97.5%	Width
b _{1,1}	1192	1191	1192	1194	3
b _{2,1}	8.860	8.798	8.860	8.922	0.124
b _{1,2}	1191	1189	1191	1193	4
b _{2,2}	8.967	8.905	8.967	9.030	0.125
b _{1,3}	1189	1188	1189	1189	1
b _{2,3}	9.153	9.147	9.153	9.158	0.011
b _{1,4}	1166	1166	1166	1167	1
b _{2,4}	9.441	9.435	9.441	9.447	0.012
b _{1,5}	1199	1198	1199	1199	1
b _{2,5}	9.012	9.007	9.012	9.018	0.011
b _{1,6}	1178	1176	1178	1180	4
b _{2,6}	9.161	9.098	9.161	9.225	0.127
b _{1,7}	1158	1158	1158	1159	1
b _{2,7}	9.540	9.535	9.540	9.546	0.011
b _{1,8}	1182	1181	1182	1182	1
b _{2,8}	9.232	9.227	9.232	9.238	0.011
b _{1,9}	1186	1186	1186	1187	1
b _{2,9}	9.171	9.166	9.171	9.177	0.011
b _{1,10}	1175	1175	1175	1176	1
b _{2,10}	9.294	9.289	9.294	9.299	0.010
b _{1,11}	1174	1172	1174	1175	3
b _{2,11}	9.243	9.179	9.243	9.307	0.128
b _{1,12}	1186	1186	1186	1187	1
b _{2,12}	9.195	9.190	9.195	9.200	0.010
* b _{1,13}	1181	1159	1181	1203	44
* b _{2,13}	9.191	8.791	9.191	9.595	0.804

It is interesting that for RTDs 1, 2, 6, and 11, the credible sets for the regression coefficients are considerably wider than those for the other RTDs. This makes sense, because all four of these were in the bath that produced only 180 observations. This count is less than a quarter of that for the other RTDs. Since RTDs 8, 9, 10, and 12 were in two baths, they had 404 observations more than RTDs 3, 4, 5, and 7 (which had 781 observations). When comparing the widths of the credible sets for these two sets of

RTDs, we see very little to no change in the width of the credible sets with increasing “sample size.”

We were also able to obtain an estimate of the size of the measurement errors of the thermometer. (Notice that in our analysis we did not have to assume that the thermometer reads without error, only that it is accurate.) For this data set, the point estimate of the parameter σ_η (again a posterior mean) was 0.1548°F with a 95% credible set of $(.1498, .1601)$. The manufacturers of the thermometer probe, thermometer readout device, and bath apparatus list their respective “accuracies” as $\pm 0.018^\circ\text{F}$, $\pm 0.045^\circ\text{F}$, and $\pm 0.018^\circ\text{F}$ respectively. Further, it seems that their interpretations of these figures is that they specify uniform conditional distributions of the truth given a reading associated with the apparatus. (This interpretation seems to be consistent with the 1993 ISO *Guide to Expression of Uncertainty in Measurement* [5] discussed in [6].) If one assumes independence and compounds such errors additively, the resulting standard deviation is approximately 0.030°F . So our analysis suggests that there is substantially more uncertainty associated with the thermometer readings than naïve compounding of manufacturer precision figures would predict. Whether there is some source of variation we have failed to account for (like, for example, the lack of exact synchronization of observation between the thermometer and RTDs), age has degraded performance of the equipment below a “brand new condition” level considered by the manufacturers, or the manufacturer values are simply too optimistic is unknown.

Arriving at prediction intervals for the measurand represented in (10) is strikingly simple. From (10), the distribution of T_{new} conditioned on $Y_{\text{new}}, b_{1h}, b_{2h}$, and σ_ε^2 is plausibly taken to be Normal with mean $(Y_{\text{new}} - b_{1h})/b_{2h}$ and variance $\sigma_\varepsilon^2/b_{2h}^2$. The

MCMC can provide predicted/estimated values for b_{1h} , b_{2h} , and σ_ε^2 at each iteration.

(This is true even for $h = 13$ if the current values of β and Σ are used to generate an intercept-slope pair from the distribution (4).) Then (using the

$N\left((Y_{\text{new}} - b_{1h})/b_{2h}, \sigma_\varepsilon^2/b_{2h}^2\right)$ conditional distribution of T_{new} suggested above) a

prediction of T_{new} can be generated. The empirical distribution of the predictions of the measurand then functions as an approximation for the posterior predictive distribution of T_{new} . (Notice that this posterior provides clearly interpretable quantification of the uncertainty in temperature T_{new} associated with a new reading Y_{new} from an RTD.) The sample moments and quantiles of this empirical distribution provide the point and interval estimates for T_{new} .

By way of example, when the 12th RTD reads a resistance of 2100, the estimated temperature is 99.35 °F with 95% posterior probability that the temperature is between 99.15 and 99.55 °F. When an indirectly calibrated RTD reads a value of 2100, the estimated temperature is 99.99 °F with a 95% credible set being (97.26, 102.90).

In order to gain some understanding of the sensitivity of our analysis to the choices made in specifying priors, we experimented with a wide range of priors for Σ . We used “conjugate” (Inverse Wishart) priors with Ω ranging from .001 $\mathbf{I}$ to 1000 $\mathbf{I}$ and $\lambda = 2$ and $\lambda = 5.1$. In addition, we also experimented with a range of non-conjugate priors for Σ (ones with independent inverse gamma marginals for variances and a uniform marginal for the correlation coefficient). We found the ability to get the MCMC to burn-in (so that empirical representations of posteriors can be obtained) to be

somewhat dependent on the choice of the prior for Σ . But burn-in was possible for most of the possibilities we considered (14 of 18).

Across those choices of prior for which burn-in was achieved, inferences for intercepts, slopes, and T_{new} for RTDs actually tested changed very little. This is probably simply a reflection of a very large number of observations upon which to estimate the parameters of the direct calibrations. The data “overwhelm the prior.” The situation is somewhat different for an indirectly calibrated RTD. For the hypothetical 13th RTD, point estimates changed very little, but intervals changed substantially with choice of the prior for Σ . This makes sense. In terms of estimating parameters for Σ , the effective sample size is small (there are only 12 intercept-slope pairs represented in the data set). The relevant data (the fairly precisely known 12 pairs) did not overwhelm prior assumptions on the parameters of the distribution (4), and those show up as affecting the posterior precision of the intercept, slope, and inverse predictions (the T_{new}) for an untested RTD.

V. DISCUSSION

If we had not incorporated element (1) and subsequently element (3) into our analysis (i.e. had treated thermometer readings as exact temperatures), the model here would have been a (simpler) linear mixed effects model. But even in that simplified $\sigma_{\eta} = 0$ context, it was not completely clear to us how to handle interval estimation for a measurand from an indirect calibration using classical/frequentist statistical methods. The present Bayesian approach seemed more promising. And using it, we were able to

estimate the precision of the thermometer, a quantity that we would have otherwise had to take on faith from the manufacturer.

We are not the first to notice the possibility of taking a Bayes approach to calibration problems. See, for example, [7] and the references therein. However, the particular model we've used (specifically including imprecision in thermometer readings) and the multiple device and (particularly the) “indirect calibration” aspects of our problem seem to us to be non-standard.

WinBUGS made computations manageable, if not completely trivial. Without this software, programming the MCMC would have been formidable. But using WinBUGS (keeping in mind the caution of its developers that “MCMC sampling can be dangerous!”) we see that quite complicated calibration problems can be handled via a Bayesian approach.

VI. ACKNOWLEDGEMENTS

We would like to thank the staff at the Energy Resource Station for their efforts in calibrating the temperature sensors and providing us with the raw measurements. The people specifically involved in the endeavor included: Emir Kadic, David Perry, Xiao Hui Zhou, and Jin Wen.

VII. REFERENCES

[1] Wen, J., and Smith, T.F. (2001). “Calibration and check procedures for temperature sensors at the Iowa Energy Center Energy Resource Station,” *Technical Report ME-TFS-01-009*, Department of Mechanical Engineering, University of Iowa, Iowa City.

- [2] Pinheiro, J.C., and Bates, D. M. (2000). *Mixed-effects models in S and S-PLUS*, Springer-Verlag Inc., Berlin; New York.
- [3] Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B. (1995) *Bayesian Data Analysis*, Chapman & Hall, London.
- [4] Spiegelhalter, D.J., Thomas, A., and Best, N.G. (2000). “WinBUGS version 1.3 user manual,” MRC Biostatistics Unit, Cambridge. (www.mrc-bsu.cam.ac.uk/bugs)
- [5] ISO Technical Advisory Group, Working Group 3 (1993). *Guide to expression of uncertainty in measurement*, International Organization for Standardization, Geneva.
- [6] Gleser, L.J. (1998). “Assessing uncertainty in measurement,” *Statistical Science*, 13, pp 277-290.
- [7] Hamada, M., Pohl, H., Spiegelman, C., and Wendelberger, J. (2003). “A Bayesian approach to calibration intervals and properly calibrated tolerance intervals,” *Journal of Quality Technology*, 35, pp. 194-205.

VIII. Appendix (WinBUGS Code)

```

model
{
##### PRIORS #####

# THE PRIOR ON THE PRECISION OF THE THERMOMETER
tau.eta ~ dgamma(2.000000008, .0008910007)

# THE PRIOR ON THE PRECISION OF A RTD
tau.epi ~ dgamma(2.001, 1.001)

# THE PRIOR ON THE MEAN INTERCEPT AND SLOPE OF ALL RTDs
beta[1:2]~ dmnorm(m.beta[],p.beta[ , ])

# THE PRIOR ON THE PRECISION MATRIX OF THE RANDOM INTERCEPTS & SLOPES
P[1:2, 1:2] ~dwish(Omega[ , ], 5.1)

##### MODELS #####

# MODEL FOR THE RANDOM EFFECTS OF THE RTDs FOR WHICH THERE IS DATA
for (h in 1:12){
    b[h , 1:2] ~ dmnorm(beta[], P[ , ])
}

# MODEL FOR THE ERROR TERM OF THE THERMOMETER (eta[i,j])
for (i in 1:B){
    for (j in 1:n[i]){
        eta[i, j] ~ dnorm(0, tau.eta)
    }
}

# MODEL FOR THE OBSERVATIONS FROM THE RTDs (Y[h,i,j,k])
for (m in 1:N){
    # THE EXPECTED VALUE OF y[h,i,j,k]
    E.Y[m] <- b[RTD[m], 1] + b[RTD[m], 2]*(X[m] - eta[BATH[m],OBS[m]])

    # THE MODEL FOR Y[h,i,j,k]
    Y[m] ~ dnorm( E.Y[m], tau.epi)
}

##### PREDICTIONS #####
##### STEP 1 #####

# THE FOLLOWING STEPS ARE TO OBTAIN THE CURRENT ITERATE VALUE FOR STOCHASTIC NODES
# NEEDED FOR THE CALIBRATION OF AN UNTESTED RTD AND PREDICTION OF T.new.

for (h in 1:12){
    for (i in 1:2){
        bstar[h,i]<-b[h,i]
    }
}

for(i in 1:2) {
    beta.star[i]<-beta[i]
}

for (i in 1:2){
    for (j in 1:2){
        P.star[i,j] <- P[i,j]
    }
}

tau.epi.star<- tau.epi

```

```
##### PREDICTIONS #####
##### STEP 2 #####

# CALIBRATION OF AN UNTESTED RTD
bstar[13,1:2]~dmnorm(beta.star[], P.star[,])

# PREDICTION FOR THE TEMPERATURE T.new GIVEN y.new, b[h,], AND var.epsilon
for (h in 1:13){
  # THE EXPECTED VALUE OF T.new GIVEN y.new, b[h,], AND var.epsilon
  E.T.new[h] <- (Y.new[h] - bstar[h,1])/bstar[h,2]
  # THE PRECISION OF T.new GIVEN y.new, b[h,], AND var.epsilon
  P.T.new[h] <- (bstar[h,2]*bstar[h,2])*tau.epi.star
  # THE MODEL FOR T.new
  T.new[h] ~ dnorm(E.T.new[h], P.T.new[h])
}

##### CALCULATIONS #####
# STANDARD DEVIATION NODE
sig.eta <- sqrt(1/tau.eta)
# STANDARD DEVIATION NODE
sig.epi <- sqrt(1/tau.epi)
# INVERSION OF THE PRECISION MATRIX (P)
# TO OBTAIN THE VARIANCE-COVARIANCE MATRIX (Sigma)
for (i in 1:2){
  for (j in 1:2){
    V[i,j] <- inverse( P[,],i,j)
  }
}

### NODES TO FOLLOW: sig.epi, sig.eta, V, b, bstar[13,], beta, T.new

### INITIAL VALUES
list(tau.epi=1000, tau.eta=1000)
list(tau.epi=1, tau.eta=1)
list(tau.epi=.001, tau.eta=.001)

### DATA

list(N=8584, B=4, n=c(779,781, 406,180),
      m.beta = c(0,0), p.beta = structure(.Data = c(1.0E-6, 0, 0, 1.0E-6), .Dim = c(2, 2)) ,
      Y.new=c(1411, 1641, 1411, 1641, 1870, 1870, 2100, 1411, 1641, 1870, 2100, 2100, 2100),
      Omega = structure(.Data = c(0.1, 0, 0, 0.1), .Dim = c(2, 2)))

Y[]      X[]      BATH[] OBS[]  RTD[]
1567     40.9667      1      1      8
1570     40.9667      1      1      9
..... etc.....
```