



Durham E-Theses

A Bayesian analysis of luminescence dating

Huntriss, Alicia

How to cite:

Huntriss, Alicia (2008) *A Bayesian analysis of luminescence dating*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/2928/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

A Bayesian Analysis of Luminescence Dating

Alicia Huntriss

A Thesis presented for the degree of
Doctor of Philosophy

The copyright of this thesis rests with the author or the university to which it was submitted. No quotation from it, or information derived from it may be published without the prior written consent of the author or university, and any information derived from it should be acknowledged.



Durham
University

13 NOV 2008

Statistics and Probability Group
Department of Mathematical Sciences
University of Durham
England

June 2008



A Bayesian Analysis of Luminescence Dating

Alicia Huntriss

Submitted for the degree of Doctor of Philosophy

June 2008

Abstract

Luminescence dating is a widespread dating method used in the fields of archaeology and Quaternary science. As an experimental method it is subject to various uncertainties in the determination of parameters that are used to evaluate age. The need to express these uncertainties fully, combined with the prior archaeological knowledge commonly available, motivates the development of a Bayesian approach to the assessment of age based on luminescence data. The luminescence dating procedure is dissected into its component parts, and each is considered individually before being combined to find the posterior age distribution. We use Bayesian multi-sample calibration to find the palaeodose in the first stage of the model, consider the problem of identifying a plateau in the data, and then use this, along with the annual dose, to estimate age. The true sample age is then modelled, incorporating any prior information available, both for an individual sample and for a collection of samples with related ages.

Declaration

The work in this thesis is based on research carried out at the Statistics and Probability Group, the Department of Mathematical Sciences, England. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

Copyright © 2008 by Alicia Huntriss.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

Acknowledgements

I would like to thank my supervisors, Prof. Michael Goldstein, Prof. Ian Bailiff and Dr. Andrew Millard for their endless help, guidance and patience, and Dr. Alex Bayliss from English Heritage for her enthusiastic encouragement.

My thanks also go to Matthew, Heather, Caroline and Laura for all their support and understanding.

This research was funded by a CASE studentship with English Heritage and EPSRC.

Contents

Abstract	ii
Declaration	iii
Acknowledgements	iv
1 Introduction	1
2 Luminescence Dating	5
2.1 Luminescence Signal from quartz	6
2.1.1 Measurement of the Luminescence Signal	8
2.2 Palaeodose Evaluation	9
2.2.1 Single Aliquot Additive Approach	10
2.2.2 Single Aliquot Regeneration Protocol	10
2.2.3 Diagnostics of the Palaeodose Estimate from the SAR procedure	14
2.2.4 Single Grain Methods	15
2.2.5 Palaeodose Models	16
2.2.6 Uncertainty in Palaeodose Evaluation	17
2.3 Preheat Plateau	18
2.4 Dose Rate	20
2.4.1 Measurement of Dose Rate	21
2.4.2 Moisture Content	23
2.4.3 Dose Rate Equation	23
2.4.4 Uncertainty in Dose Rate Evaluation	24
2.5 Age Ratio	25

2.6	Assumptions Made in the Thesis	25
2.7	Motivation of the Model Strategy	26
3	Palaeodose Evaluation at a Single Preheat Temperature	32
3.1	Calibration	32
3.1.1	Bayesian Calibration	35
3.2	Palaeodose Evaluation using the Combined Aliquot Model	36
3.2.1	Likelihood	38
3.2.2	Prior Distributions	39
3.2.3	Prior Elicitation	39
3.3	Posterior Distributions	41
3.4	Stability and Convergence of the Sampler	46
3.4.1	Data	46
3.4.2	Analysis of Chains	47
3.4.3	Gelman and Rubin Method	49
3.4.4	CUSUM path plots	50
3.4.5	Geweke Method	52
3.4.6	Summary of Convergence Analysis	53
3.5	Example	53
3.5.1	Prior Specification	53
3.5.2	Data	55
3.5.3	Posterior Distribution	56
3.6	Sensitivity to Prior Parameters	57
3.6.1	γ_R as a Random Variable	62
3.6.2	Summary	63
3.7	Diagnostics	63
3.7.1	Linear Model Diagnostics	63
3.7.2	Diagnostics for the Bayesian Model	65
3.8	Palaeodose Evaluation for Sample 311-6	66
3.9	Summary Guide to evaluating palaeodose at each preheat temperature	68

4	Preheat Plateau Model	70
4.1	Motivation for modelling preheat plateau	70
4.2	Model	72
4.2.1	Location of the plateau	73
4.2.2	Data Input	74
4.2.3	Likelihood	76
4.2.4	Posterior Distribution for T_a	76
4.2.5	Prior Distributions	77
4.3	Example	78
4.3.1	Prior Specification	78
4.3.2	Posterior Distribution	80
4.3.3	Influence of the Prior Parameters	81
4.4	Sample Palaeodose Evaluation	83
4.5	Extension of the Preheat Plateau	84
4.5.1	Influence of Prior Parameters	86
4.5.2	Sample Palaeodose Evaluation	87
4.6	Summary Guide to evaluating sample palaeodose	89
5	Annual Dose	91
5.1	Dose Rate Equation	91
5.2	Example	92
5.2.1	Choice of Parameter Values	93
5.3	Influence of parameters	94
5.4	Water Content Variations	99
6	Age Evaluation	103
6.1	Age Ratio	103
6.1.1	Example	104
6.2	Sample Age	108
6.2.1	Example	110
6.2.2	Influence of Priors	111
6.3	Extended Plateau Example	115

6.4	Comparison with Current Luminescence Age Evaluation	115
7	Inference with Related Samples	116
7.1	Coeval Model	116
7.1.1	General Model for m Coeval Samples	119
7.1.2	Example	119
7.2	Similar Age Model	124
7.2.1	General m Similar Age Model	126
7.2.2	Example	127
7.2.3	Influence of Prior Specifications	129
7.3	Ordered Age Model	131
7.3.1	Example	132
8	Example	136
8.1	Palaeodose Evaluation at each Preheat Temperature	138
8.1.1	Preliminary Experiments	138
8.1.2	Prior Elicitation	139
8.1.3	Posterior Distributions	140
8.2	Preheat Plateau	146
8.2.1	Prior Elicitation	147
8.2.2	Posterior Distribution	148
8.3	Sample Palaeodose	150
8.4	Dose Rate	153
8.5	Age Ratio	154
8.6	Age	155
8.6.1	Sensitivity to Prior Specifications	155
8.7	Sample 318-1	156
8.8	Age Analysis	160
9	Conclusions	164
9.1	Future Work	167
	Bibliography	170

Appendix	182
A Notation	182
A.1 Notation for Palaeodose Evaluation	182
A.2 Notation for Preheat Plateau Model	183
A.3 Notation For Dose Rate Model	183
A.4 Notation for Age evaluation	184
B Convergence Diagnostics	185
B.1 Independence of Starting Values	185
B.2 Gelman and Rubin Method	185
C Conditional Posterior Distributions for the Combined Aliquot Model	188
C.1 Conditional Posterior Distribution for Palaeodose	189
C.2 Conditional Posterior Distribution for Aliquot Palaeodose Estimates .	189
C.3 Conditional Posterior Distributions for the Regression Coefficients . .	190
C.4 Conditional Posterior Distribution for σ^2	193
C.5 Conditional Posterior Distribution for γ_R^2	194
D Simulated Data used to Test Convergence and Stability of the Gibbs Sampler	195
E Sensitivity of the Combined Aliquot Model to Prior Judgements	198
F Influence of Prior Parameters in the Plateau Model	205
F.1 Influence of curve parameters	205
G Data	208
G.1 Fydell House, 311-6	208
G.2 Data at Lower Preheat Temperatures	210
G.3 Fydell House 311-2	212
G.4 Fydell House 311-4	214
G.5 Tattershall Castle 318-2	216
G.6 Tattershall Castle 318-1	218

H Code	221
H.1 Code for the Combined Aliquot Model	221
H.2 Code for Plateau Model	228
H.3 Code for Dose Rate	231
H.4 Code for Age Evaluation	233
H.5 Coeval Model	234
H.6 Similar Age Model	236
H.7 Ordered Age Model	239
I Analysis of 311-2, 311-4 Fydell House	242
I.1 Sample 311-2	242
I.2 Sample 311-4	244

List of Figures

2.1	The single aliquot additive dose method for palaeodose evaluation. . .	10
2.2	Example of estimating palaeodose using the SAR protocol	14
2.3	Summary diagram of current routine procedure for evaluating the sample age using OSL dating.	28
2.4	Flow chart of the steps in the Bayesian model for palaeodose evaluation.	30
2.5	Flow chart of the model steps for age evaluation.	31
3.1	Comparison of the classical estimate using regression of y on x and the inverse estimate from the regression of x on y , using a simulated data set.	34
3.2	Raw trace plots of x_R simulations, with all iterations (top) and just the first 1000 iterations. These calculations are based on a simulated data set, where 500 mGy is the true value for x_R	48
3.3	CUSUM path plots for x_R from the sampler (top) and from draws from a normal distribution with the same mean and variance as the x_R simulations (bottom), for comparison.	51
3.4	Plotted data from one of the preliminary experiments, for sample 311-6 from FydeH House.	54
3.5	Data from aliquots with a preheat temperature of 210°C, with fitted lines being used to estimate palaeodose by back-interpolation.	56
3.6	Posterior distribution of palaeodose based on three aliquots with a preheat of 210°C.	57
3.7	Posterior palaeodose distributions at a preheat of 210°C, for different prior standard deviations, σ_R , with the prior mean fixed at $\mu_R = 1000$.	58

3.8	The influence of prior standard deviation, σ_R , on posterior mean palaeodose for three different choices of prior palaeodose mean.	59
3.9	Posterior distribution for palaeodose, and palaeodose estimates from the 3 aliquots with preheat 210°C.	60
3.10	Residuals from aliquots with a preheat of 210°C, plotted against fitted values.	64
3.11	Posterior distribution for palaeodose, estimates of palaeodose from each aliquot with a preheat of 250°C and prior palaeodose distribution.	66
3.12	Posterior distribution for palaeodose at preheat (a) 220°C, (b) 230°C, (c) 240°C, (d) 250°C	67
4.1	An example of variation in palaeodose evaluation with preheat temperature given in Madsen et al 2007 [69].	71
4.2	Variation of palaeodose evaluation with preheat temperature given as an example in Wintle and Murray 2006 [120], originating from Jacobs et al 2003 [57].	72
4.3	Examples of the shape of curve that can be achieved using (4.5) to represent possible relationships between palaeodose and preheat temperature.	74
4.4	Posterior palaeodose mean and two standard deviation uncertainty bars for each group of aliquots at a certain preheat temperature. These values will be used as the data input in the model to find the start of the preheat plateau	79
4.5	Posterior distribution of plateau starting temperature T_a for sample 311-6.	80
4.6	Posterior palaeodose distribution for sample 311-6.	84
4.7	Mean and 2 standard deviation bars of posterior palaeodose distributions obtained at each preheat temperature over the extended range.	86
4.8	Posterior distribution for plateau starting temperature over the extended data set.	87
4.9	Posterior distribution for palaeodose based the aliquots of 311-6 with preheat of 200°C and above.	88

4.10	Normal mixture distribution of posterior palaeodose based on the weighted mixture of the two possible plateau locations, with the dashed line giving the normal distribution with the same mean and variance.	89
5.1	Posterior dose rate distribution for sample 311-6 from Fydell House.	95
5.2	Mean and standard deviation of the posterior dose rate against the prior standard deviation of (a) \dot{D}_β , (b) \dot{D}_γ , (c) \dot{D}_c , (d) W	96
5.3	Influence of the standard deviation of parameters (a) b , (b) g , (c) H_β , (d) H_γ , on the standard deviation of dose rate, \dot{D}	98
5.4	Influence of the mean of F on the mean and standard deviation of \dot{D}	99
5.5	Possible variation of the fractional water saturation of the dating environment since the resetting of the luminescence clock.	100
5.6	Mean dose rate with one standard deviation uncertainty bars calculated with mean of F as 0.15 for different proportions of the dating lifetime (0.7 for the remainder of the time). The dashed lines represent the mean and one standard deviation of the dose rate evaluated with mean F at 0.15 at all times.	101
5.7	Hypothesised Variation in the water saturation levels of the sampling environment.	102
6.1	Distribution of \bar{A}_E for sample 311-6. The dashed line gives the distribution of $A_E \bar{A}_E$	105
6.2	Posterior distributions for (a) Palaeodose and (b) Annual Dose for 311-6, as found in Chapters 3 and 5, overlaid with normal distributions with the same mean and variance for comparison.	106
6.3	Ratio estimate distribution, as a ratio of two independent normal distributions, along with the normal distribution with the same parameters (the distribution of the age ratio).	108
6.4	Posterior mean and standard deviation of age with (a) m_A , (b) σ_A , (c) σ_E	112

7.1	Posterior mean and standard deviation of age from the coeval model with σ_E , ρ_E and ρ_A	124
7.2	Posterior age distributions from the similar age model for 311-2 and 311-4.	128
7.3	Influence choice of σ_E has on the posterior mean and standard deviation for A_1 and A_2 under the similar age model.	130
7.4	Posterior distributions for A_1 (Sample 311-6) and A_2 (Sample 311-2) using (a) Ordered age model with $A_1 < A_2$ and (b) Similar age model.	133
8.1	Tattershall Castle	137
8.2	The initial preliminary experiment carried out by Bailiff for sample 318-2, Tattershall Castle.	138
8.3	Posterior palaeodose distributions for aliquots from sample 318-2 with preheat temperatures (a) 200°C (b) 210°C (c) 220°C (d) 240°C.	142
8.4	Raw trace plot of x_R simulations for the first 1000 iterations with a preheat of 200°C.	143
8.5	Residuals plotted against fitted values of the linear model for aliquots with a preheat of 200°C.	145
8.6	Posterior palaeodose means plotted against preheat temperature for sample 318-2, with two standard deviation uncertainty bars of the posterior distributions for palaeodose.	147
8.7	Posterior distribution of plateau starting temperature, T_a	149
8.8	Posterior palaeodose distribution of 318-2.	151
8.9	Posterior palaeodose distribution for 318-2 based on aliquots with preheat treatments $< 240^\circ\text{C}$	152
8.10	Posterior dose rate distribution for sample 318-2.	154
8.11	Posterior age ratio of sample 318-2, with the dashed line giving the density of the normal distribution with the same statistics.	154
8.12	Influence of the prior mean on the posterior mean age with a prior age standard deviation of 20 and 50 years.	156
8.13	Posterior palaeodose distribution at preheat temperature (a) 200°C, (b) 220°C, (c) 240°C.	157

8.14	Posterior palaeodose means plotted against preheat temperature for sample 318-1, with two standard deviation uncertainty bars.	158
8.15	Posterior distribution for sample palaeodose for 318-1.	159
8.16	Dose rate distribution for 318-1.	159
8.17	Age ratio of sample 318-1, with the dashed line giving the density of the normal distribution with the same statistics.	160
8.18	Posterior age distribution for (a) 318-1, (b) 318-2, (c) 318-1 and 318-2 under the coeval model,(d) 318-1 and 318-2 under the similar age model.	162
B.1	Raw trace plots from the Gibbs sampler for x_R , α , β , showing convergence independent of the starting values chosen.	186
E.1	Influence of σ_R on posterior palaeodose distribution when $\mu_R = 900$ mGy.	199
E.2	Influence of σ_R on posterior palaeodose distribution when $\mu_R = 1100$ mGy.	201
F.1	Posterior T_a distributions with prior (a) $T_a \sim N(200, 10^2)$, (b) $T_a \sim N(215, 10^2)$, (c) $T_a \sim N(230, 10^2)$	206
F.2	Posterior T_a distributions with prior (a) $T_a \sim N(200, 30^2)$, (b) $T_a \sim N(230, 30^2)$	206
F.3	Posterior T_a distributions with prior (a) $T_a \sim N(200, 50^2)$, (b) $T_a \sim N(215, 50^2)$, (c) $T_a \sim N(230, 50^2)$	207
I.1	Posterior Palaeodose distribution for sample 311-2 at preheats (a) 200°C, (b) 220°C, (c) 230°C, (d)240°C, (e)250°C, (f)260°C.	246
I.2	Palaeodose estimates plotted against preheat temperature for 311-2, with two standard deviation uncertainty bars.	247
I.3	Posterior Palaeodose distribution for sample 311-4 at preheats (a) 200°C, (b) 210°C, (c) 220°C, (d)230°C, (e)240°C.	248
I.4	Palaeodose estimates plotted against preheat temperature for 311-4, with two standard deviation uncertainty bars.	249

List of Tables

2.1	Generalised single aliquot regeneration protocol, taken from Murray and Wintle (2000) [81]	13
3.1	Mean and standard deviation of the posterior x_R distribution from the sampler for different burn-in periods, prior to thinning.	49
3.2	Mean and standard deviation of the posterior x_R distribution, calculated using different lengths of chain, prior to thinning.	49
3.3	Mean and standard deviation of the posterior x_R distribution, calculated using different levels of chain thinning. A burn-in of 1000 and chain length of 50,000 iterations was used.	50
3.4	Geweke's convergence diagnostic computed for different chain lengths.	53
3.5	Values assigned to the prior hyperparameters for the combined aliquot model, when it is applied to 311-6, FydeH House.	53
3.6	Influence of γ_R on posterior mean and standard deviation of palaeodose	61
3.7	Influence of the choice of prior parameters for the linear coefficients on posterior palaeodose distribution.	62
3.8	Posterior mean and standard deviation for palaeodose when γ_R^2 is treated as a unknown, a priori.	62
3.9	The residual standard error from the linear fit of the data from each of the three aliquots with a preheat of 210°C.	65
3.10	Posterior mean and standard deviation of palaeodose estimates at each preheat, from sample 311-6.	68

4.1	Mean and standard deviations of the posterior palaeodose at each preheat temperature for sample 311-6; the data input in the plateau model.	78
4.2	Prior distributions used for Sample 311-6, Fydell House, in the plateau model.	80
4.3	Posterior probability that the preheat plateau starts at a temperature higher than 210°C, for varying prior specifications of T_a	81
4.4	Posterior probability that the preheat plateau starts at a temperature higher than 210°C, for varying prior specifications of x_{R*}	83
4.5	Posterior palaeodose mean and standard deviation for aliquots at each preheat temperature, for the extended data set from sample 311-6. . .	85
4.6	Influence of the prior distribution for T_a on its posterior for the extended data set of sample 311-6.	88
5.1	Measurements used to compute the dose rate for sample 311-6, Fydell House.	93
5.2	Standard attenuation parameter values [1] used in the dose rate calculation.	93
5.3	Distributions assigned to the parameters in the model for dose rate for sample 311-6 of Fydell House.	94
6.1	Input in the age model for sample 311-6, Fydell House	111
6.2	Mean and standard deviation of the posterior mean age and date of the sample, for different prior judgements.	114
7.1	Ratio estimates and their standard deviations achieved for samples 311-2, 311-4.	120
7.2	Parameters of the joint prior distributions used to estimate the joint dose rate distributions of samples 311-2 and 311-4 from Fydell House.	122
7.3	Parameter values for the coeval model.	122
7.4	Posterior age and date estimates obtained with different prior age distributions.	123
7.5	Parameter values for the similar age model for samples 311-2, 311-4. .	127

7.6	Influence of prior judgements on the posterior ages achieved with the similar age model.	129
7.7	Parameter values for the ordered age model for samples 311-6, 311-2.	132
7.8	Posterior Ages resulting from the ordered age model and similar age model.	134
7.9	Posterior Age using the ordered age model with different prior specifications.	135
8.1	Prior Parameters for the combined aliquot model to evaluate palaeodose for sample 318-2 from Tattershall Castle.	140
8.2	Posterior palaeodose distributions at each preheat temperature for sample 318-2.	141
8.3	Posterior palaeodose mean and standard deviation at preheat 200°C evaluated with different burn-in lengths.	143
8.4	Posterior palaeodose mean and standard deviation at preheat 200°C evaluated for different length of chains.	144
8.5	Posterior palaeodose mean and standard deviation at preheat 200°C evaluated for different thin levels.	144
8.6	Prior and posterior palaeodose means at each preheat temperature for sample 318-2.	146
8.7	Prior distributions assigned to the plateau model parameters for Sample 318-2.	148
8.8	Posterior probabilities for plateau starting temperature for different prior judgements	148
8.9	Distributions assigned to parameters to evaluate the dose rate for sample 318-2, based on prior judgements.	153
8.10	Prior parameter values used to calculate posterior sample age using the model outlined in Chapter 6.	155
8.11	Posterior palaeodose mean and standard deviation of palaeodose evaluated at each preheat temperature of sample 318-1.	156
8.12	Comparison of different age models for 318-1, 318-2.	161

D.1	The simulated luminescence response values for the given doses, standardised to a natural luminescence value of 5000 counts	197
E.1	Influence of σ_R on posterior palaeodose when $\mu_R = 900\text{mGy}$	199
E.2	Influence of σ_R on posterior palaeodose distribution when $\mu_R = 1000$	200
E.3	Influence of σ_R on posterior palaeodose distribution when $\mu_R = 1100$ mGy.	200
E.4	Influence of μ_R, σ_R on palaeodose when $\gamma_R = 50$	202
E.5	Influence of μ_R, σ_R on palaeodose when $\gamma_R = 25$	203
E.6	Influence of μ_R, σ_R on palaeodose when $\gamma_R = 10$	203
E.7	Influence of μ_R, σ_R on palaeodose when $\gamma_R = 5$	204
F.1	Posterior probability for the preheat plateau starting above 210°C , for different prior judgements about the curve parameters	207
I.1	Values assigned to the prior hyperparameters for the combined aliquot model, when it is applied to	242
I.2	Posterior mean and standard deviation of palaeodose at each preheat temperature, for sample 311-2.	243
I.3	Prior distributions used for Sample 311-6, Fydell House, in the plateau model.	243
I.4	Distributions assigned to the parameters in the model for dose rate for sample 311-2 of Fydell House.	244
I.5	Posterior mean and standard deviation of palaeodose at each preheat temperature, for sample 311-4.	245
I.6	Distributions assigned to the parameters in the model for dose rate for sample 311-4 of Fydell House.	245

Chapter 1

Introduction

Luminescence is the light emitted from crystalline materials following absorption of energy from ionising radiation and subsequent external stimulation by some source. For example, thermoluminescence (TL) is emitted in response to heat and optically stimulated luminescence (OSL) in response to light from the visible spectrum. For luminescence, in contrast to other light emissions, such as fluorescence, there is a time lapse between the absorption and the emission of energy [46]. TL and OSL are the main types of luminescence used for dating.

Luminescence dating is a powerful chronometric technique which can be applied to a wide range of materials containing abundant minerals. The event being dated in luminescence dating is either the most recent heating (above 200-300°C) of mineral grains (such as in pottery, burnt flint, bricks), or the most recent exposure to daylight (deposition of sediments). Currently objects with an age from around a century up to a few hundred thousand years can be dated [5], though the theoretical upper age limit is thought to be up to a million years [68]. Luminescence dating is widely used to date materials from the Quaternary, a period of significant interest as it spans great climatic change and the emergence of the first modern humans [68].

The basis of luminescence dating relies on the properties of minerals, such as quartz and feldspar grains, which enable energy to be stored as charge trapped at defect sites (traps) within their crystalline structure. Exposure to ionising radiation naturally occurring in the environment (alpha and beta particles, gamma radiation), and cosmic radiation originating from space, results in these electron traps being



filled [5]. The charge is stored cumulatively, and the amount stored is proportional to the time passed since a resetting event (unless the traps have become full i.e. reached saturation before this event occurs). The dating clock is reset when the minerals are exposed to heat or light ('bleached'), referred to as the zeroing event.

The luminescence signal results when the charge is ejected from the traps ('de-trapped') when stimulated by either heat or light. Some detrapped charge goes to sites known as recombination centres and, if these centres are radiative, the excess energy is released as photons (the luminescence signal). The rate at which the trapped charge accumulates is proportional to the rate of energy absorption by a grain. In turn, the intensity of the luminescence signal on stimulation is proportional to the energy absorbed [3].

A radiation dose is defined as the energy absorbed per unit mass, measured in Gray (1 Gy= 1 J kg⁻¹) [5]. The dose the sample has received since the zeroing event is known as the palaeodose. The dose received each year is known as the dose rate or the annual dose. The number of years since resetting, or the age of the sample, is estimated using the age equation

$$A_E = \frac{Palaeodose}{Annual\ Dose}. \quad (1.1)$$

In recent years significant advances have been made in both the development of the methodology [38, 120] and in instrumentation [108], though there are still some aspects of luminescence behaviour that are not fully understood. The experimental nature of this dating method, and the inherent variations of luminescence properties between samples, means that the validity of its application and the results obtained need to be carefully analysed [68]. This, along with the careful consideration of the uncertainty and expected luminescence behaviour may be handled within a Bayesian framework, which allows the numerous uncertainties in the dating process to be fully expressed. Such analysis is the subject of this thesis.

Archaeological dating is a natural application of Bayesian methodology [28], as experts often have some degree of belief about the date of an object, or its context in relation to other such objects, before the scientific dating process has been carried out. The Bayesian philosophy has been used to combine dating information

from a number of different sources and dating techniques. For example in radiocarbon (^{14}C) dating, additional information from stratigraphy, typology and dendrochronology have been used in a Bayesian framework [28, 105]. Bayesian methods have also been applied to electron spin resonance (ESR) dating [72], TL dating [73] and archeomagnetic dating [66].

Radiocarbon dating has been the biggest archaeological application area for Bayesian statistics, with particular emphasis on the construction of chronologies [30]. This has been facilitated by the the development of the radiocarbon calibration programmes OxCal [23] and BCal [29], which are widely available, where the chronological ordering of sample is used as prior information.

OxCal has also been utilised for Bayesian chronology building with luminescence dates [92], which is applicable when there is certainty in the relative chronology. The potential has also been realised through the development of a mixture model for the evaluation of palaeodose in sediment mixtures [101].

These examples of the use of Bayesian statistics in luminescence dating apply the Bayesian methodology once the ages [92] or palaeodose estimates [101] have been computed. The luminescence dating process is complex and a number of separate calculation stages can be identified. Here we look at the uncertainties and assumptions made at each stage, modelling each stage and using the prior information available to find the posterior distribution for the parameters of interest at each step, culminating in modelling the relationship between the age ratio (1.1) and the true sample age.

Luminescence dating requires a wide variety of measurements. Chapter 2 gives an overview of the important aspects of luminescence and routine methodologies used to date materials and explains the conditions and assumptions that are made in the dating situations that our Bayesian model will cover.

Assessment of the palaeodose estimate is in part a problem in Bayesian multi-sample calibration of a linear model, and the evaluation of a single palaeodose is described in Chapter 3. Combining the palaeodose estimates to evaluate the palaeodose of the sample is covered in Chapter 4. Chapter 5 examines the uncertainties in the annual dose distribution. In Chapter 6, the uncertainty in palaeodose and

annual dose is considered to give uncertainty in their ratio A_E , and then prior information on the age is utilised to produce a posterior distribution for the sample age. In Chapter 7 inference from multiple samples and the relationships between the sample ages is considered. A working example will be presented for each step of the model, and Chapter 8 details a further example to illustrate the procedure. Chapter 9 contains a concluding discussion.

Chapter 2

Luminescence Dating

The application of luminescence to dating was first proposed in 1953 [33], though it took a further 10 years before reliable dates were produced [2] using thermoluminescence (TL). TL was initially used to date pottery, and extended to other heated materials such as tephra¹ and burnt flint [89]. The observation that exposure to light reset the luminescence clock in a similar way to heat [89], along with the separate recognition that luminescence can be stimulated by light [55], led to the development of optically stimulated luminescence dating (OSL), with the main motivation being its application to dating Quaternary sediments [4].

The main minerals that are used in luminescence dating are quartz and feldspar. Quartz is mineralogically simple, whereas feldspar is often more complex, with a wide range of structures and compositions [36]. Ages obtained using feldspar are frequently underestimated [110]. This underestimation is often caused by a phenomenon known as anomalous fading [116] where electrons are depleted from their traps over a short timescale compared to the predicted trap lifetime [113]. Dating using feldspar does have advantages over using quartz crystals, including the ability to stimulate the luminescence signal using wavelengths in the infra-red region (IRSL), which has technical advantages for measuring the luminescence signal, as the stimulating wavelength is significantly different from the wavelength of the resulting luminescence emission [36] which makes it easier to eliminate the stimulation

¹Tephra is the glass material ejected into the air by a volcanic eruption

photons from the detection system. This issue, along with the relative brightness of the luminescence compared to quartz, has led to a number of techniques being developed to overcome the problem of anomalous fading [45, 65, 110]. To eliminate any problems with anomalous fading, in this thesis we will only consider OSL dating using quartz, which does not exhibit this unfavourable property.

The quartz minerals that are used for optically stimulated luminescence dating can be divided into two main groups: those which have had their luminescence clock reset by light, or 'bleached' (e.g. sediments) and those for which the resetting occurred by heating (e.g. bricks, pottery). When the material is heated, for example in the firing of bricks, then the luminescence signal is completely reset (all the luminescence traps are emptied). However, the resetting of the signal by daylight can be highly heterogeneous with some grains being more completely bleached than others [109]. A sample containing grains which have not been completely reset can lead to overestimation of the sample age [100] and a broad, asymmetric distribution of palaeodoses [86].

The completeness of the signal bleach depends on the environment in which the sediment is deposited, which in turn dictates the length of time the grains are exposed to light and the spectrum of light available for bleaching [100]. For example, fluvial deposits are well known for exhibiting properties of heterogeneous bleaching [75] as the high energy wavelengths are attenuated through water [5], and it is these wavelengths that are most efficient at bleaching [100].

2.1 Luminescence Signal from quartz

When quartz that has been irradiated with ionising radiation is exposed to light, a luminescence signal (OSL) is emitted. This signal decays with time as the trap(s) are depopulated. The decay curve does not fit a single exponential, indicating that the luminescence does not originate from a single trap and radiative recombination centre [79].

It is thought that the luminescence signal is a result of charge being released from several trap types [9], and that the decay can be adequately represented by three ex-

ponential components [104]. These components are known as the fast, medium and slow components, with reference to their relative rate of decay [9]. The components result from the trap types having different photo-ionisation cross-sections and this implies different de-trapping probabilities [67]. Up to seven components in the OSL decay have been identified [58], depending on the properties of the quartz and the experimental procedure [62]. These can be divided into ‘fast’, ‘medium’ and ‘slow’ categories [99]. An ‘ultra-fast’ component is sometimes identified [59], which decays rapidly and is thermally unstable. This component is thought to result from a trap type distinct from that responsible for the fast component [59].

The fast and medium components dominate the initial part of the decay, after which the slow component can be identified [11]. The fast component is primarily used for luminescence dating [67] as it exhibits the desirable properties that it bleaches easily [59] and the electrons in the associated traps are stable over millions of years [99]. The slow component has much greater thermal stability which indicates that this part of the signal originates from deep traps separate from those responsible for the fast and medium components [11]. During the initial decay, the proportion of signal from the slow component is usually relatively small. However, this is not true in all cases and if it is ignored (or assumed to be constant) this can induce errors in the evaluation of palaeodose [11].

A number of complex physical models have been developed to replicate the luminescence process [9, 12, 71], though some observed luminescence behaviours are yet to be fully understood [68]. These models are based on a series of trap types and recombination centres. Some traps are highly photosensitive, and are emptied rapidly on exposure to light, while others are not affected by light exposure. The majority of traps are sensitive to heat, and are emptied when raised to high temperatures. A small proportion have very short retention lifetimes, and are unable to store electrons for more than a few days at ambient temperatures, whereas some ‘deep’ traps are thought to be able to hold electrons for millions of years.

The characteristics of the various traps are important since they affect the degree to which sediments are zeroed/reset before burial. Incompletely reset grains will give rise to erroneously young ages.

2.1.1 Measurement of the Luminescence Signal

The probability of charge eviction from a trap is dependent on the sensitivity of the trap to photoeviction, and the rate at which photons from the stimulation source arrive at the trap [108]. The sensitivity of the trap is dependent on the wavelength of the stimulation light; the probability of eviction increases for shorter wavelengths [19]. Another important factor in the choice of stimulation wavelength is the emission wavelength of the luminescence. The intensity of the luminescence emission is $\sim 10^{-19}$ of the intensity of the stimulating light [108], and so it is critical that the wavelength of the stimulation source is well separated from the main emission wavelengths.

Quartz OSL emits strongly in the blue and ultra-violet spectrum [38] (320-380nm), with the peak around 365 nm [5]. The intensity of the luminescence signal is measured using a photomultiplier (PM) tube detector coupled with appropriate coloured filters. These filters block any scattered stimulation light and define the wavelength range for detection of the emitted signal [5]. A number of different stimulation sources have been used in OSL dating, including lasers, halogen lamps and light emitting diodes (LEDs) [5]. In recent times blue and green LEDs have become popular [108].

The luminescence signal measured by the PM tube includes some unwanted components [5]; when the luminescence signal from an aliquot whose charge traps have been emptied is measured, some signal is still detected. This originates from sources including PM noise, backscattering of the stimulation source (not blocked by the colour filters) and long term luminescence [4]. This is known as the background signal, and is subtracted to evaluate the intensity of the OSL signal. The background is measured as the intensity of the OSL signal after a period of stimulation (specified by the laboratory) once the signal does not decay further with time. It is desirable for the background signal to be constant for each measurement made [114].

2.2 Palaeodose Evaluation

The luminescence signal measured when the prepared sample is stimulated for the first time in the laboratory is referred to as the natural signal. The radiation dose that the grains have experienced in the natural environment resulting in this intensity of luminescence is known as the palaeodose, and to be evaluated. The methods adopted for palaeodose evaluation in routine dating have evolved with theoretical and empirical knowledge [38], and laboratory equipment [21].

The two main approaches to palaeodose evaluation are the additive method and the regenerative method [5]. These approaches can be applied to multiple-aliquot, single aliquot or single grains, where an aliquot of grains is typically 1-2mg of crystals. Each of these will be reviewed here, with a detailed examination of the single aliquot regeneration (SAR) protocol. The SAR procedure is often used in routine dating, and is the basis for the Bayesian model for palaeodose evaluation developed in Chapter 3.

Historically, multiple aliquot methods were adopted to estimate palaeodose [3,5], where typically 24-48 aliquots were used to obtain a single palaeodose value. A number of problems were encountered with this approach, which is based on the assumption that each aliquot has the same palaeodose. This assumption cannot be verified using this method. Also, the large number of measurements and sample size required to produce one palaeodose estimate means that repeat estimates are not usually feasible, so that the uncertainty of the measurement is difficult to calculate [38].

Duller [35] first used a single aliquot approach for dating feldspar, with Murray and Wintle [81] developing the widely adopted SAR procedure for quartz grains. The ability to measure palaeodose from single aliquots enables repeat measurements to be made, and thus the uncertainty in the sample palaeodose to be evaluated [40]. Single grain approaches take this to the extreme, with large numbers of palaeodose estimates being made (and often required as the estimates are sensitive to grain to grain variations and as such may have a wider distribution).

2.2.1 Single Aliquot Additive Approach

The single aliquot additive dose (SAAD) protocol [77] involves measuring the luminescence signal, after a brief exposure to light, then irradiating the aliquot with a laboratory dose and sampling the trap population again. The measurement and irradiation cycle is repeated, and used to define an additive growth curve. This curve is extrapolated back to the intercept on the dose axis to estimate palaeodose, as illustrated in Figure 2.1. This method of palaeodose evaluation is not considered in our Bayesian model as it is no longer used routinely in OSL dating.

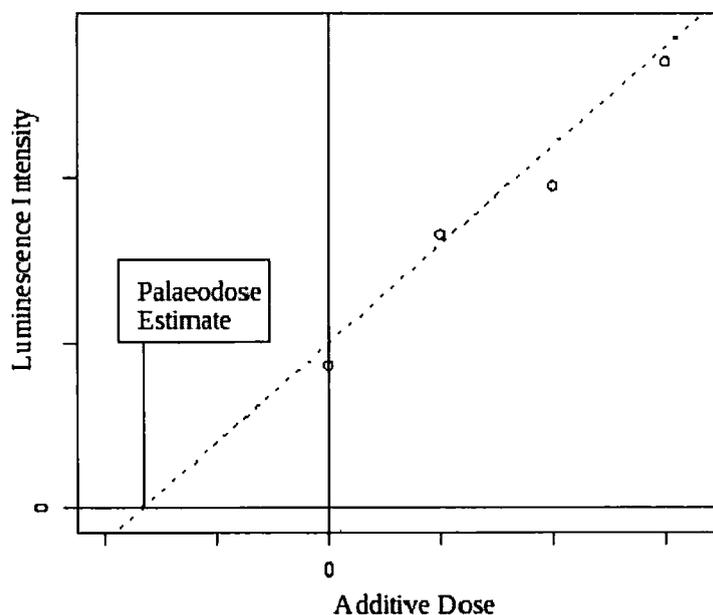


Figure 2.1: The single aliquot additive dose method for palaeodose evaluation.

(This approach is not considered in the remainder of the thesis.)

2.2.2 Single Aliquot Regeneration Protocol

The single aliquot regeneration procedure measures the natural luminescence signal resulting from the unknown palaeodose, and a series of OSL signals arising from known doses applied in the laboratory. The relationship between the laboratory doses and the luminescence signals subsequently measured is then used to estimate

the palaeodose.

The single aliquot regeneration protocol is based on three main assumptions [81]:

1. When the traps are being filled, the competition for charge is the same during the natural irradiation process as in the laboratory.
2. The luminescence sensitivity (the OSL response per unit trapped charge) is the same for the natural signal as for the signals from the regenerative doses applied in the laboratory.
3. The traps which contribute to the OSL signal, both natural and laboratory-induced, are stable over archaeological timescales.

The first assumption is difficult to check, and can only be verified by dating known-age samples where all other sources of systematic error have been accounted for.

The sensitivity of the sample is liable to change during the experimental procedure [107], and can be dependent on both thermal history and time [118]. It is the ability of the SAR procedure to monitor and compensate for these sensitivity changes which has fuelled its popularity in routine luminescence dating [38]. After each regenerative dose is applied and the luminescence signal (L) measured, a further test dose is applied which produces an OSL signal T . This test dose is the same for each of the laboratory doses used, and is an indicator of the luminescence sensitivity at that particular time [40]. The ratio of the regenerative dose and test dose signals (L/T) is used as a measure of luminescence response.

It is assumed that the sensitivity of the luminescence resulting from the regenerative dose is directly proportional to that of the signal arising from the test dose [82], but the constant of proportionality may be dependent on the dose [81]. In practice, there can be changes in sensitivity between the regenerative and test doses being applied [81], but these are taken to be independent of dose [32], and relatively small.

The third criterion given for the SAR procedure to be reliable is that the traps contributing to the OSL signal are stable over an archaeological period. A trap model has been developed by Bailey [12] (Section 2.1) where only some of the traps are stable over such time scales. Empirical results [80] have shown that signals arising from laboratory irradiated doses contain additional signal, with a relatively

short lifetime compared to the natural OSL, which is stable over archaeological and geological time periods. The unwanted unstable portion of the regenerated signal is removed by heating: this ‘preheat’ treatment occurs after the dose irradiation and prior to the signal being measured.

The preheating temperature applied is typically between 160 and 300°C, and the sample is held at that temperature typically for 10s. The appropriate choice of temperature is sample dependent, so a number of different preheat treatments are used across the aliquots, and comparisons made between the palaeodose estimates achieved.

The steps of the generalised SAR protocol from Murray and Wintle [81] are shown in Table 2.1. The first cycle measures the natural signal, so no dose is applied in the laboratory. The luminescence measurements are made at an elevated temperature to inhibit re-trapping of electrons associated with the 110°C TL peak that may give rise to a secondary phosphorescence signal [79]. The preheat temperature at stage 2 is the same for each cycle, but is changed systematically for different aliquots.

A palaeodose estimate is produced from each aliquot, by plotting the L_i/T_i ratios against regenerative dose, fitting a growth curve to the data and then using back interpolation from the natural signal L_0/T_0 onto the fitted curve to estimate palaeodose. The shape of the curve is commonly accepted to be a saturating exponential [5], to reflect the belief that at high doses the traps become full and so the luminescence signal saturates [39]. If ratio $R_i = L_i/T_i$ is observed after irradiating dose D , then

$$R(D) = I_{max} \left(1 - \exp \left\{ -\frac{D}{D_0} \right\} \right) + c \quad (2.1)$$

where D_0 is the characteristic dose (characterising the rate at which the traps become full), I_{max} is the saturation level and c is an offset [39].

OSL traps typically saturate at around 200 Gy [121]; here we will only consider relatively young samples whose traps are not approaching saturation. Therefore the growth curve can be considered linear in this region [39], and a least squares line is fitted to the L_i/T_i values. An illustration is shown in Figure 2.2.

Step	Treatment ^a	Observed ^d
1	Give dose D_i	-
2	Preheat ^b (160-300°C for 10s)	-
3	Stimulate ^c for 100s at 125°C	L_i
4	Give test dose D_t	-
5	Heat to 160°C	-
6	Stimulate for 100s at 125°C	T_i
7	Return to 1	-

Table 2.1: Generalised single aliquot regeneration protocol, taken from Murray and Wintle (2000) [81]

^aFor the natural sample, $i = 0$, and $D_0 = 0$ Gy.

^bAliquot cooled to $< 60^\circ\text{C}$ after heating.

^cThe Stimulation time is dependent on the stimulation light intensity

^d L_i and D_i are derived from the initial OSL signal (0.3 or 0.8s) minus a background estimated from the last part of the stimulation curve.

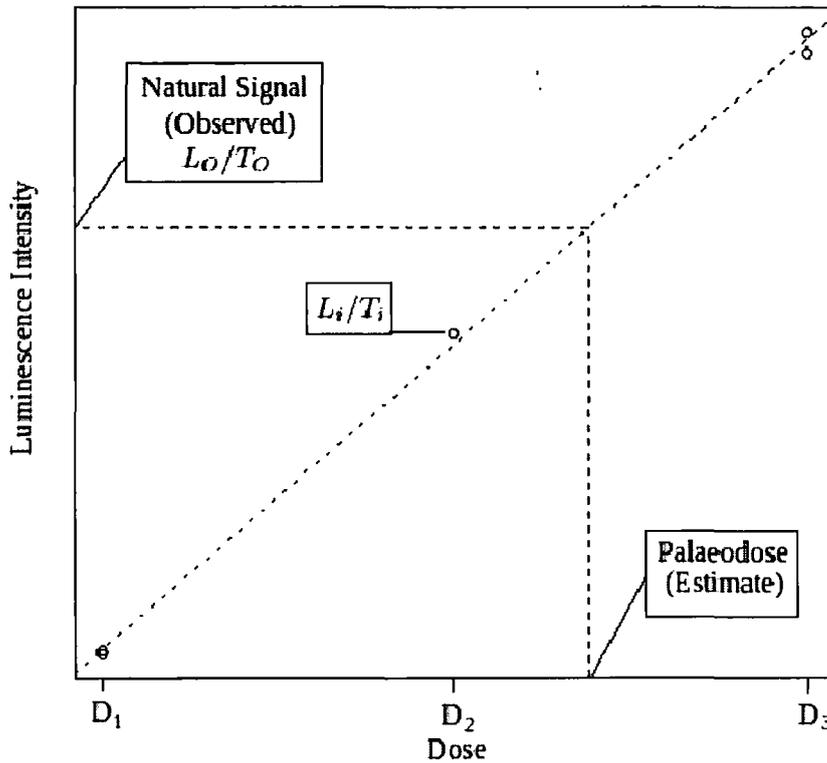


Figure 2.2: Example of estimating palaeodose using the SAR protocol

Initially, the natural signal L_0/T_0 is measured. Then, for example, at dose D_1 and D_3 two measurements of L/T were made, and at D_2 one measurement of L/T was made. The dotted line is the least squares fit to these five points, and the back extrapolation estimates the palaeodose.

2.2.3 Diagnostics of the Palaeodose Estimate from the SAR procedure

A number of tests are routinely carried out to establish the reliability of the palaeodose estimates, made and the validity of the assumptions made, by the SAR procedure.

1. **Recycling Test** During the SAR procedure, repeat measurements are often made for one or more of the regenerative doses applied. The ratio of the two corrected luminescence signals L/T with the same regenerative dose is calculated, and is known as the 'recycling ratio'. If the sensitivity changes are

correct, then this ratio should be 1 [120], though in practice a range of 0.9-1.1 is generally accepted [81].

2. **Recuperation Test.** When a regeneration dose of 0 Gy is applied, zero signal should be observed. However, prior SAR cycles of irradiation, preheating and optical stimulation may cause charge to be transferred from deeper traps, and a recuperation signal observed [6]. For the SAR procedure to produce a good estimate of palaeodose, Murray and Wintle [81] suggest that the level of recuperation should be less than 5% of the natural signal.
3. **Dose Recovery Test.** The greatest change in sensitivity is thought to occur when the sample is first heated [120]. The test signal T_0 is measured after this so may not be an appropriate measure of sensitivity of the natural signal L_0 . This can be tested by zeroing the natural signal, and applying a known dose to an unheated part of the sample [95]. The SAR procedure is then applied, with satisfactory results if this dose is ‘recovered’.
4. **Preheat Plateau Test.** A number of different preheat treatments are used across the aliquots, and the palaeodose estimates achieved are plotted against preheat temperature. A region where the palaeodose remains the same across the different preheat treatments, a ‘preheat plateau’, indicates that the preheat has suitably removed all the thermally unstable charge, and is used as a measure of self-consistency [120].

2.2.4 Single Grain Methods

A protocol for evaluating palaeodose using single grains of quartz was initially set out by Murray and Roberts in 1997 [78], after Lamothe *et al* [64] had first used a single grain approach for dating feldspar [5]. Single grain methods were motivated by palaeodose variance across aliquots, allowing a distribution of palaeodose values to be found. The inter-grain palaeodose variation is explained by [93] and summarised below.

- Grain characteristics [37]. The variability of the luminescence properties of

grains leads to some naturally ‘bright’ grains.

- Incomplete bleaching [86]. Some of the grains may not have experienced sufficient exposure to light before burial for their luminescence clock to be completely reset.
- Beta dose non-uniformity [78]. The beta dose rate may not be homogeneous across all grains during the natural irradiation process. (The beta dose rate typically contributes at least half of the total dose rate).
- Post-deposition mixing [94]. Grains from more recent deposits may contaminate older layers, e.g. by bioturbation.

Both the additive and regenerative dose methods have been applied to single grains [78], though a large number of grains are required [8] to utilise the small proportion of ‘bright’ grains with a relatively high luminescence sensitivity [37]. Single grain methods have been made possible through the development of specialist laboratory equipment [21,22], as the precise location of each of the individual grains mounted on a disc has to be determined.

2.2.5 Palaeodose Models

The development of single grain techniques (Section 2.2.4) has enabled a large number of palaeodose estimates to be made for each sample. Galbraith [47] developed a number of different models to represent the spread in the population, and so estimate the true palaeodose. The models proposed considered the logarithm of palaeodose, and are:

1. Common age model. Here the palaeodose estimates would be consistent with a common value, so that, for each of the estimated log palaeodoses $\hat{\delta}_i$, with standard error s_i ,

$$\hat{\delta}_i = \delta + \epsilon_i \quad (2.2)$$

- . with the true common log palaeodose δ , and where ϵ_i is the deviation of $\hat{\delta}_i$ from δ . This deviation is modelled as a random quantity with mean 0 and variance s_i^2 .

2. Central age model. If the log palaeodose estimates are not consistent with a common value, then the common age model can be generalised so that

$$\hat{\delta}_i = \delta_i + \epsilon_i \quad (2.3)$$

where δ_i is the true log palaeodose for grain i , and ϵ_i as above. Here the true palaeodose values for each grain are not equal but considered to be a random sample from a normal distribution.

3. Minimum age model. This is applicable in situations where the sample was incompletely bleached on deposition, and so the true log palaeodose values δ_i are considered to be a random sample from a mixed truncated normal distribution [47].

2.2.6 Uncertainty in Palaeodose Evaluation

The uncertainties in palaeodose evaluation can be categorised into systematic and random errors. The systematic errors induced usually originate from the laboratory measurement process, for example the calibration of the radioactive sources used. The sources of random uncertainty in the palaeodose estimates are discussed above in relation to grain-to-grain variation. The development of single aliquot and single grain methods of palaeodose evaluation have allowed multiple estimations to be made and allows the spread of palaeodose values to be investigated in routine dating [39].

The intensity of the luminescence signal is measured by counting the number of photons detected by the photomultiplier tube. Thus counting statistics [48] come into play in the assessment of the uncertainty in palaeodose, and are more significant when the intensity of the luminescence is low, especially when in comparison with the background signal.

The development of single grain methods has allowed a greater number of palaeodose estimates to be made for each sample, and a larger distribution of values to be observed. A number of models have been developed to calculate the most suitable palaeodose value given the distribution of aliquot estimates [13,47].

The uncertainty in each of the aliquot estimates of palaeodose can be considered to be a combination of uncertainty in the luminescence measurements, and uncertainty in converting these measurements into a palaeodose estimate. The uncertainty in the luminescence measurement depends on the intensity of the signal and that of the subtracted background count. Instrumental error is typically assumed to be around 1% [39].

The mean of the aliquot palaeodose evaluations (over a particular preheat temperature range) is routinely used to evaluate the sample palaeodose, and the uncertainty in this value is quantified using the standard deviation of these estimates (e.g. [14]).

2.3 Preheat Plateau

Preheating is an important stage of the SAR procedure (Section 2.2.2). The temperature of the aliquot is raised before the luminescence is stimulated in order to remove charge from the shallow traps which are unstable over dating timescales, and thus do not contribute to the natural luminescence signal. This allows the natural OSL signal and laboratory irradiated signal to be compared. The optimum preheat treatment empties the shallow traps while preventing a significant thermal erosion of the deep OSL traps (which contributes to the main OSL signal) [117].

The most suitable preheat treatment, in terms of both temperature and duration, has been widely debated in the past [119], with a treatment for 5 minutes at 220°C originally being suggested [103]. Currently, temperatures in the range 180-280°C are used, held for 10s. The use of a range of different preheat temperatures was recommended by Murray et al (1997) [77], to allow for sample variation in luminescence properties.

Preheating causes the luminescence properties of quartz to change [112]. Rhodes [91] observed a decrease in OSL signal intensity with preheat temperatures up to 140°C, followed by an increase in OSL. The signal has been observed to drop rapidly for high preheat temperatures, above around 280°C [76]. Most notably, the sensitivity of the sample changes, so, for example, if the quartz grains become more

sensitive, then a higher intensity of luminescence is observed for the same irradiation dose.

The SAR protocol for palaeodose evaluation (Section 2.2.2) has been developed so that any sensitivity changes are adjusted for. In order to assess the effectiveness of these corrections, the palaeodose estimates obtained are plotted against the preheat temperature applied. If there is no change in the palaeodose estimates with preheat temperature, i.e. a 'preheat plateau' is observed, then it is concluded that the sensitivity changes have been adequately compensated for [4].

A preheat plateau also indicates that the preheat treatments applied were sufficient. If the lowest preheat temperatures were not high enough to remove all the unstable charge, then the palaeodose evaluations from aliquots with these preheats would be underestimated. High temperatures can also lead to thermal erosion of the main OSL trap, and so to reducing the signal:background ratio which implies poor counting statistics, which can lead to overdispersion.

The increase in OSL signal with preheat temperature was attributed to charge transferring from non-photosensitive shallow traps to light sensitive traps and thus contributing to the OSL signal [91, 117]. However, more recently it is thought that most (if not all) of the changes in the OSL signal are due to the changes in luminescence sensitivity, and charge transfer is not significant [76].

A wide variety of trends are observed for the relationship between preheat temperature and palaeodose [120], some of which are not fully understood. However, if a preheat plateau is observed over some (or all) of the temperatures used, then these are the aliquots that are thought to be estimating the palaeodose of the sample.

The presence (or lack of) a preheat plateau in the palaeodose estimates is currently assessed by eye. The aliquots which lie on the preheat plateau are then typically averaged to estimate the palaeodose of the sample, with their spread indicating the level of uncertainty in this evaluation.

2.4 Dose Rate

The dose rate, or annual dose, is an estimate of the average natural radiation dose that the sample has received each year since the luminescence clock has been reset. This is used as the denominator in the age equation (1.1) to estimate the sample age. The methods used to evaluate the dose rate remain similar to those used in the initial determination of luminescence ages around 40 years ago [68], though there have been advances in instrumentation.

When a nucleus undergoes radioactive decay, ionising radiation is emitted and the type of radiation is dependent on the decay process. The radiation types are alpha particles, beta particles and gamma rays; cosmic radiation also contributes to the dose rate. These different radiation types exhibit different properties [5] (the ranges given below are those in typical sample medium):

- **Alpha Particles** These are heavily ionising and the radiation is highly localised; its range is $\sim 20\mu\text{m}$.
- **Beta Particles** These are lightly ionising with a range of a few mm ($\sim 3\text{mm}$).
- **Gamma Rays** These are also lightly ionising but with a range of a $\sim 30\text{cm}$.
- **Cosmic Rays** This is radiation originating from space and is lightly ionising. It has a small dependence on latitude at sea level, though at altitudes above 1km this dependency grows along with its intensity.

The radiation types all produce secondary electrons, and ionisation occurs when these electrons have been slowed down sufficiently [5]. Ionisation of atoms within host crystalline structures generate free electrons, and it is these electrons and holes which are trapped in defects in the crystal structure of the quartz or feldspar, and utilised in luminescence dating.

The main naturally occurring radioelements are uranium, potassium and thorium (lithogenic radionuclides) which are present in the sample and the surrounding environment. These often contribute to the natural OSL signal in approximately equal proportions [5], depending on the sampling environment. Rubidium isotopes also make a small contribution to the dose rate. Radioactive decay starts with the

'parent' isotope, which decays into the 'daughter' nucleus whilst emitting nuclear radiation. This decay chain continues until a stable (non-radioactive) 'daughter' nucleus is reached.

2.4.1 Measurement of Dose Rate

The evaluation of the dose rate is based on the assumption that the overall rate at which energy is absorbed is equal to the rate of energy emission, within a volume larger than the range of the radiation [5]. Uniformity of the dose rate is assumed in the conventional calculation of the luminescence age, though in some cases this it may not be a valid assumption [61]. The dose rate can either be assessed by analysis of the radionuclide composition of the material, or by measuring the rates of alpha, beta, gamma and cosmic radiation individually.

In the 'concentration approach' the content of potassium, rubidium, thorium and uranium are determined and the dose rate components are evaluated by use of conversion tables which detail the likely proportion of radioactive isotopes and their effective dose rates [1,5]. The main drawback to this method is the possibility that the thorium and uranium decay chains are not in radioactive equilibrium. This occurs when, instead of the rate of decay of the isotope being equal to the rate of formation of its daughter, there is loss of some of the daughter material. For example the daughter radon is gaseous and a portion of it can escape, especially if the material is porous [3].

Alternatively, the alpha, beta, gamma and cosmic radiation dose rates can all be measured. Thick-source alpha counting [102] can be used to evaluate the alpha radiation component in fine grain dating where the alpha radiation penetrates the full radius of grains. This applies to silt-size grains in the range 4-11 μm , whereas in coarse grain dating (100 μm) the alpha particles only reach the outer part of the grains due to their short range. Here the surface is etched using hydrofluoric acid (HF), which removes the part of the grain which is affected by the alpha radiation, and thus reduces the alpha radiation dose rate to a negligible level [3] and simplifies the dose rate calculation.

Highly sensitive phosphors can be used to measure the present day environmental

dose rates using similar dosimetry methods as radiation monitoring around nuclear plants [5]. Phosphors commonly used include aluminium oxide doped with carbon ($\text{Al}_2\text{O}_3:\text{C}$) [20]; the phosphor needs to display similar absorption properties as quartz (or feldspar) [5]. The luminescence clock of the phosphor is set to zero, typically by heating, and is then placed in a small capsule [3]. This dosimetry capsule is then left in the sampling location for, ideally, one year to account for any seasonal variation in moisture content (Section 2.4.2), though a few months is sufficient if necessary [5].

The radiation dose that the dosimetry capsule has received is then measured, either using thermoluminescence (a TLD capsule) or OSL. Depending on the radiation component rate to be assessed, the walls of the capsule are designed to absorb either alpha or alpha and beta radiation (to measure the gamma component). The cosmic radiation dose can be subtracted to isolate the gamma radiation component [3].

The beta dose rate component can either be measured using a dosimetry capsule as described above or directly using a beta-counter, for example a Geiger-Muller system [18]. Spectrometry can also be used to measure the beta dose rate.

It is optimal to measure the gamma dose rate in-situ; especially if the validity of the homogeneous dose rate assumption is in doubt [3]. This can be done by using a portable gamma spectrometer, scintillometer (which measures radioactivity levels) or with a TLD capsule [5]. A scintillometer has a short measurement time (~ 10 minutes), though it only measures the overall gamma radiation dose rate; whereas a spectrometer returns the rate from the separate potassium, thorium and uranium components.

Apart from environments with particularly low levels of radioactivity, or those at high altitudes, the cosmic radiation dose does not form a significant portion of the dose rate and so a calculated estimate of the cosmic dose rate is usually employed. It is possible to measure the present day cosmic dose rate using a portable gamma-ray spectrometer [31], though for sediments, due to the attenuation by the (increasing) overburden, this often does not represent an average for the burial period [5]. The cosmic dose rate is typically taken to be $150 \mu\text{Gy a}^{-1}$ [3], though the depth of the any sediment being dated needs to be taken into account [74]. Reconstruction of the overburden history can be attempted if this is thought necessary [5].

The dose rate due to sources located within the grains used in luminescence measurement (the internal dose rate) needs to be considered when dating feldspar and some types of quartz. In particular, potassium feldspar can be composed of up to 10-14% potassium of which around 0.05% are radioactive isotopes, which can contribute significantly to the dose rate depending on the grain size. Accurate evaluation of the internal dose rate is not straight forward [41]. Generally, it is assumed that quartz grains have no internal dose rate [5], though increasingly the validity of this assumption is being tested [111] and is now assessed for each sample in case of contamination.

2.4.2 Moisture Content

Typical dating environments will contain water in voids between the mineral grains. Water absorbs a higher proportion of the energy released by the alpha, beta and gamma radiation compared with mineral grains. Hence the dose rate reduces as the moisture content increases [5]. The moisture content of the sample in the present day environment is easily measured and appropriate adjustments can be made to the dose rate. However, the water content history is not known and an average must be estimated based on any information available about the dating environment.

2.4.3 Dose Rate Equation

In routine luminescence dating using coarse-grain quartz (so there is no alpha dose²), the dose rate \dot{D} is calculated using a standard model [3]

$$\dot{D} = \frac{b}{1 + H_{\beta}WF} \dot{D}_{\beta} + \frac{g}{1 + H_{\gamma}WF} \dot{D}_{\gamma} + \dot{D}_c \quad (2.4)$$

where

- \dot{D}_{β} , \dot{D}_{γ} , \dot{D}_c are the measured beta, gamma and cosmic radiation dose rates
- b , g are the attenuation factors for the beta and gamma radiation respectively in the sample medium: these are standard values and not sample-specific.

²As part of the laboratory preparation of the sample, chemical etching of the outer layer of the grains is carried out to remove the contribution to the dose rate from alpha particles

- H_β , H_γ are the relative attenuation factors for beta and gamma radiation in water, which also take community-wide accepted values for every dating situation.
- W is the saturation water uptake (i.e. the amount of water contained in the material at saturation), and F the time averaged fractional moisture content (on average, the water content of the material as a proportion of complete saturation).

So here $\{\dot{D}_\beta, \dot{D}_\gamma, \dot{D}_c, W\}$ are measured values, $\{b, g, H_\beta, H_\gamma\}$ take standard values and are not sample-specific, and F is estimated based on any information available about past variations in water content.

The dose rate, or annual dose is evaluated in milligray per year (mGy a^{-1}), where the Gray is the unit of absorbed dose ($1 \text{ Gy} = 1 \text{ J kg}^{-1}$).

2.4.4 Uncertainty in Dose Rate Evaluation

The sources of uncertainty in the evaluation of dose rate, beyond any experimental and measurement errors, include the following.

- **Heterogeneity**

In evaluating the dose rate, it is assumed that the sample environment has a homogeneous dose rate [5]. However this is not always true [61], particularly for beta radiation. The heterogeneous distribution of radionuclides is not necessarily a critical factor when dating with multiple grains, as palaeodose evaluated from an aliquot of grains will reflect the average dose rate if the number of grains contributing to the measured luminescence is sufficient to avoid the effects of fluctuation [61]. However where single grains are measured (Section 2.2.4) the implications of dose rate heterogeneity are more significant [84].

- **Water content history**

Unless the sample environment is known to have been arid or to have had a saturated water content throughout the dating timescale, the average moisture

content of the sample is subject to uncertainty. It is thought that in sediment dating this is a fundamental limitation on the ability to reduce error below $\pm 5\%$ of the sample age [5].

- **Standard values of attenuation factors**

Standard values are used for the attenuation factors for the radiation in water and the surrounding material [1], but nonetheless are likely to have small associated uncertainties [7]. However these differences are not expected to give rise to significant errors in the measurement of the dose rate for the majority of environments.

2.5 Age Ratio

The age is estimated by dividing the palaeodose by the dose rate, the age ratio:

$$A_E = \frac{\textit{Palaeodose}}{\textit{Annual Dose}}. \quad (2.5)$$

The values obtained by the methods described previously to estimate the palaeodose and annual dose are used in this equation. In our Bayesian analysis we go on to consider the relationship between the age ratio and the true sample age.

The uncertainty in the age is routinely assessed by summing the random and systematic errors in the palaeodose and annual dose [3]. However, there are two distinct categories of uncertainty in the sample age: the uncertainty in evaluating the age ratio, and the uncertainty in the age ratio as an estimate for the sample age.

2.6 Assumptions Made in the Thesis

Luminescence dating covers a wide range of different techniques and dating environments [38]. In this thesis we are constructing a Bayesian model for luminescence dating, and so it is necessary to start with the most basic, realistic dating environment. Therefore, the assumptions made in the remainder of the thesis are the following.

- Optically stimulated luminescence (OSL) will be used for dating.

- Quartz mineral grains free of internal radionuclide sources will be used, not feldspar, so the issues of anomalous fading and internal grain dose rate do not need to be addressed.
- The evaluation of palaeodose and annual dose are independent, though this may not always be strictly true.
- At the time of the event being dated, the material was fully zeroed. For example, heating of the quartz grains (for example firing of brick, pottery) or sediment dating where the grains were fully bleached on deposition. This means that we do not consider partially-bleached materials, where a skewed or mixture distribution for palaeodose would be expected.
- The single aliquot regeneration protocol (SAR) is used to evaluate palaeodose.
- The age of the sample and the dose rate is such that the traps are not close to being completely filled, i.e. saturation is not being approached. So, a linear relationship between dose and luminescence intensity is a good approximation, instead of the more general saturating exponential relationship.
- The dose rate is homogeneous.

2.7 Motivation of the Model Strategy

Routine analysis in luminescence dating comprises a number of separate steps to produce the final estimate of the sample age. For the situations which are being considered here (Section 2.6), for each aliquot, the palaeodose is estimated by back-interpolation from a fitted line using the SAR protocol (Section 2.2.2). These estimates are then plotted against preheat temperature as a diagnostic check; if a plateau is observed this indicates that sensitivity changes have been adequately adjusted for and the shallow traps have been emptied. The palaeodose of the sample is then evaluated by averaging the palaeodose estimates which lie on the preheat plateau.

Next, the dose rate is estimated using the dose rate equation (Section 2.4), combining experimental values with standard coefficients and estimation of the water content history. The palaeodose and dose rate evaluations are used to estimate the sample age by the age equation (1.1). This procedure is summarised in Figure 2.3.

The approach taken in this thesis to modelling luminescence dating using Bayesian methodology is to consider each of the steps to the age evaluation individually. That is, a model has been developed to evaluate the posterior palaeodose distribution at each preheat temperature. Then, the preheat plateau is modelled to find the posterior distribution for the starting temperature of the plateau. This is then used to assess which aliquots lie on the preheat plateau, and the data from such aliquots are used in the initial model to find the posterior palaeodose distribution of the sample. This process is described in Figure 2.4.

A distribution for the dose rate is found using the dose rate equation used in routine dating, along with the experimental values and expert judgements about the precision of these measurements and the water content history of the sample. The next stage in the model is to find the distribution for the age ratio, using the age equation (1.1) and the distribution for dose rate and palaeodose found previously. This is analogous to the final step in the routine age analysis (Figure 2.3).

In the Bayesian model developed here, rather than this being the culmination of the analysis, we go on to model the relationship between the age ratio and the true sample age. It is at this stage that the prior information about the sample age is utilised. The structure of the Bayesian model, once the posterior distribution for palaeodose has been found, is illustrated in Figure 2.5.

The Bayesian model developed here was designed to follow a similar structure to that of routine age analysis in luminescence dating, choosing to structure the Bayesian analysis as a series of sub-models for a number of reasons rather than accumulate the model into one large calculation.

Firstly, the motivation behind the development of a Bayesian model for luminescence dating is to enable the luminescence community to make routine use of Bayesian techniques. Thus the model needs to be accessible to those who do not have extensive knowledge of Bayesian methodology. Splitting the model into dif-

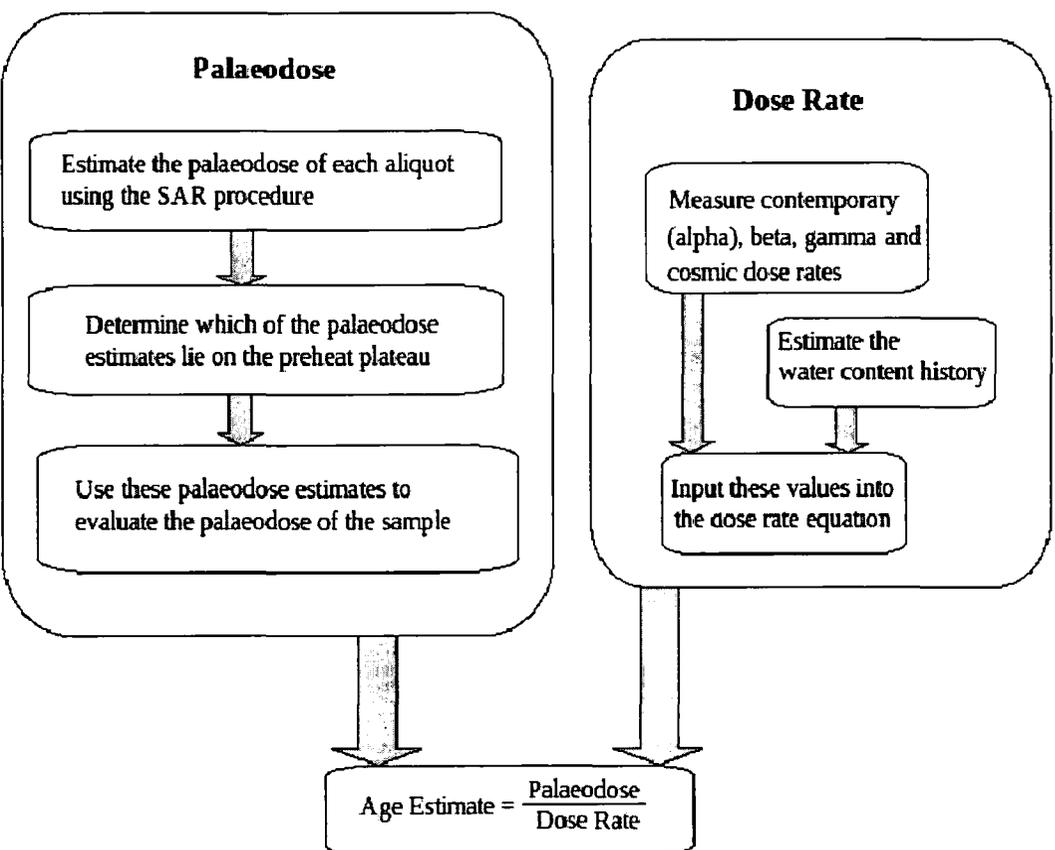


Figure 2.3: Summary diagram of current routine procedure for evaluating the sample age using OSL dating.

ferent stages enables practitioners to understand each stage of the process, as it follows a similar pattern to current practice. Also, the MCMC methods used here have proved to be robust and not caused any convergence problems, and as such are suitable for implementation by practitioners unfamiliar with MCMC techniques.

Secondly, this model strategy also allows the practitioner to view the posterior distribution for the relevant parameters at each stage. As luminescence characteristics can be extremely variable between different samples, these parameters are of significant interest to the practitioner. It also allows the practitioner to monitor the levels of uncertainty in each of the parameters as the analysis of the luminescence age progresses.

An alternative approach would have been to consider the age distribution in a single large calculation. However, this strategy presents a number of difficulties. The joint distributions of the parameters within the overall calculation are not straightforward. A lot of the details of the age evaluation would be lost within the large calculations required, and any modifications made to adapt the model to the dating environment and changes in the experimental protocol would become complex. It would be very difficult to express all of the expert prior judgements about the different ingredients of the overall age assessment within a single calculation, particularly if different aspects of the assessment are made by different individuals.

Such a solution would also somewhat oppose the aim of the model: to open the field of Bayesian statistics to luminescence dating. Overly complicated computations would only deter most luminescence practitioners from applying such Bayesian methods.

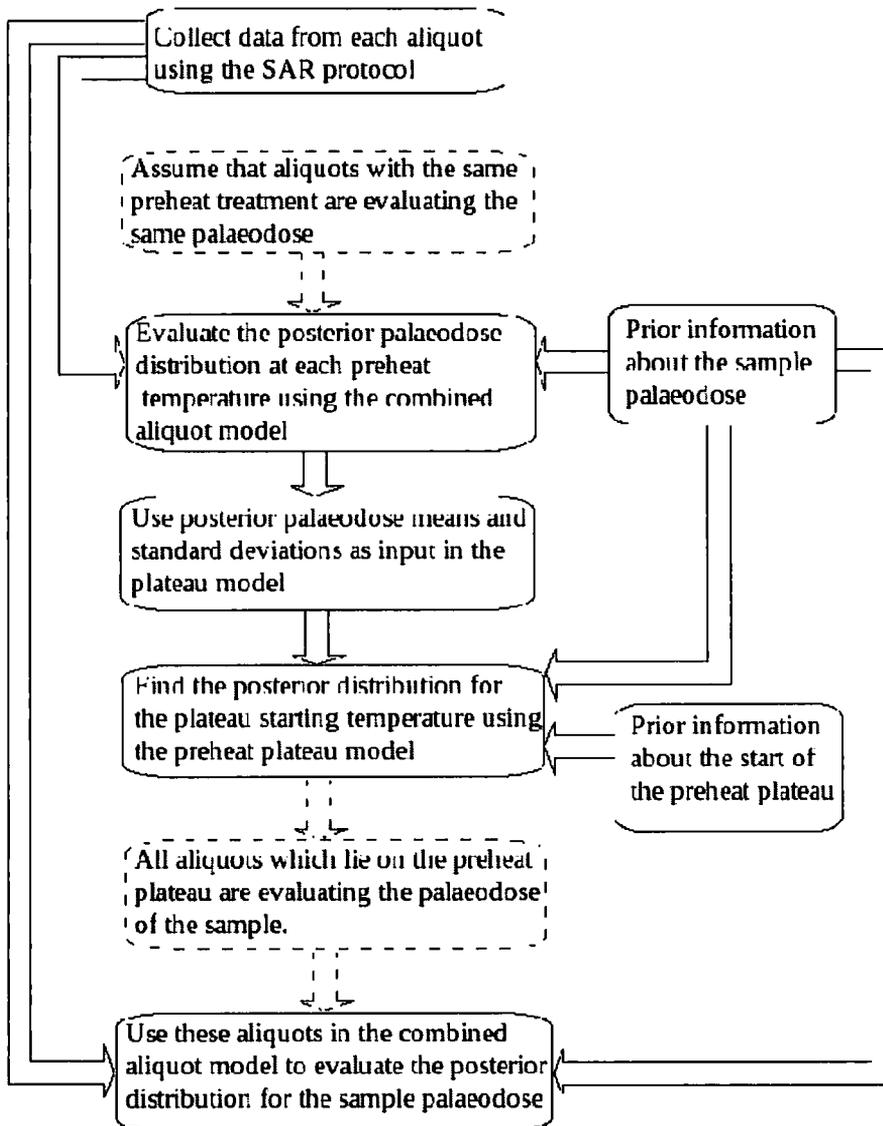


Figure 2.4: Flow chart of the steps in the Bayesian model for palaeodose evaluation.

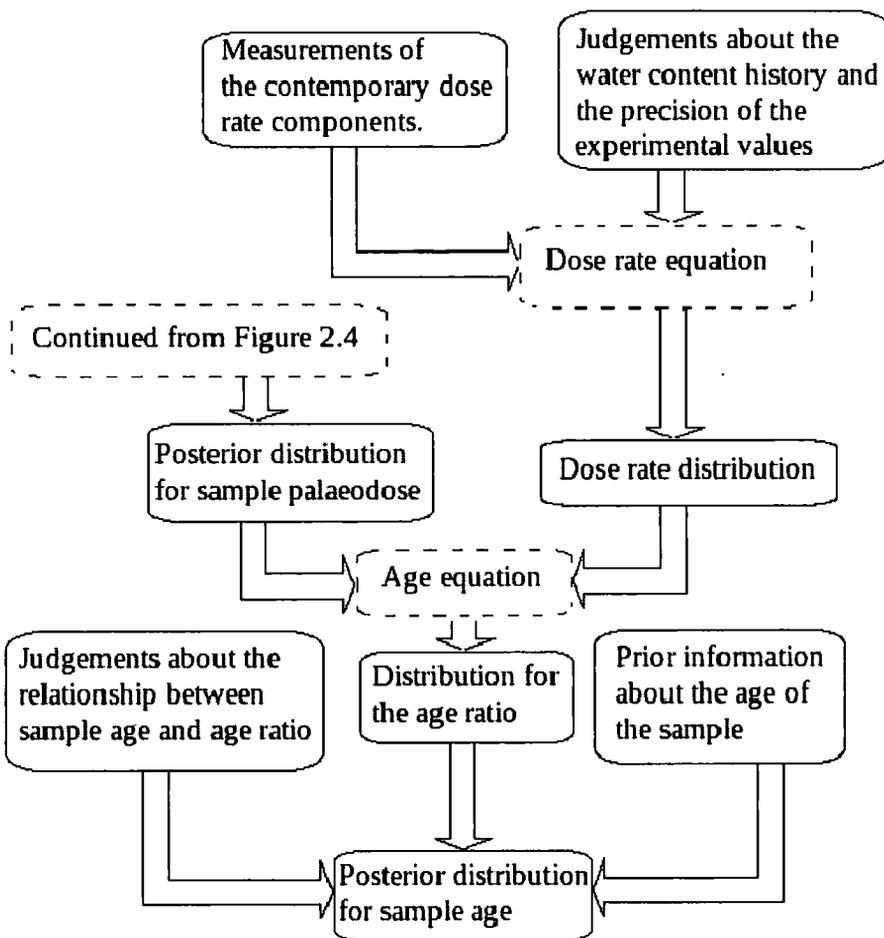


Figure 2.5: Flow chart of the model steps for age evaluation.

Chapter 3

Palaeodose Evaluation at a Single Preheat Temperature

In this chapter we look at the evaluation of palaeodose at a single preheat temperature, based on the single aliquot regeneration (SAR) protocol [81, 82]. Throughout we assume that quartz grains, which have been reset through heating, are being dated and that they have originated from a homogeneous environment. Thus the luminescence signal would be bright and well-behaved (recuperation and recycling tests perform well). We also only consider relatively young archaeological samples, so that the relationship between dose and luminescence signal is considered to be linear (Section 2.2.2).

First, the evaluation of palaeodose is considered as a calibration problem. A Bayesian model is laid out and a Gibbs sampler is detailed to estimate the posterior palaeodose distribution, and its stability and convergence are investigated. The model is tested using an example, and a sensitivity analysis carried out on the prior parameters used.

3.1 Calibration

As detailed in Chapter 2, Section 2.2.2, the SAR procedure involves the measurement of the natural luminescence signal followed by the irradiation of known regenerative laboratory doses. For each dose, luminescence response is recorded. This produces a

data set for each aliquot containing the sensitivity corrected luminescence responses, within which some of the dose values may be repeated. The luminescence response is plotted against the regenerative dose applied, and conventionally a least squares line is fitted for each aliquot. The line is used to estimate an unknown x (palaeodose) from the known response (natural luminescence), which is a calibration problem.

Several approaches have been taken to such univariate calibration problems. Suppose we have a set of data $X = \{x_1, \dots, x_n\}$ with response variable $Y = \{y_1, \dots, y_n\}$, and assume a linear relationship between x and y . Denote a further observation of the response z , where its corresponding unknown x value is ξ . There are two different estimates in the statistics literature which are commonly used for ξ :

1. **Classical estimate $\hat{\xi}$** [44]. This fits the regression model

$$y_i = \alpha + \beta x_i + \epsilon_i \quad (3.1)$$

to the data, where ϵ_i are independent and normally distributed errors. Then ξ is estimated using

$$z = \hat{\alpha} + \hat{\beta}\xi, \quad (3.2)$$

so

$$\hat{\xi} = \frac{z - \hat{\alpha}}{\hat{\beta}} \quad (3.3)$$

where $\hat{\alpha}$, $\hat{\beta}$ are the least squares estimates of α , β , with

$$\hat{\alpha} = \bar{y} - \hat{\beta}\bar{x}, \quad \hat{\beta} = \frac{S_{xy}}{S_{xx}} \quad (3.4)$$

and

$$S_{xx} = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2, \quad S_{yy} = \frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2, \quad S_{xy} = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y}). \quad (3.5)$$

2. **Inverse Estimate $\check{\xi}$** [63]. This uses the least squares estimators of the linear coefficients when x is regressed on y . The inverse estimate $\check{\xi}$ for ξ is found using

$$\check{\xi} = \hat{\alpha}_* + \hat{\beta}_* z \quad (3.6)$$

where $\hat{\alpha}_*$, $\hat{\beta}_*$ are

$$\hat{\alpha}_* = \bar{x} - \hat{\beta}_* \bar{y}, \quad \hat{\beta}_* = \frac{S_{xy}}{S_{yy}} \quad (3.7)$$

A simulated data set was used to compare the two estimates, and Figure 3.1 shows the two different regression lines fitted. The classical y on x regression has a steeper calibration line than regressing x on y as, if r is the correlation coefficient,

$$\hat{\beta}\hat{\beta}_* = \frac{S_{xy}^2}{S_{xx}S_{yy}} = r^2 \quad (3.8)$$

and when the two lines are plotted on the same set of axes the slope of the inverse regression is $\frac{1}{\beta_*}$. The exception is when the linear fit is perfect, so $r^2 = 1$ and the two lines coincide. The lines intersect at the point (\bar{x}, \bar{y}) , where $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$ and $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$. The inverse estimator $\hat{\xi}$ lies closer to the mean of the x -values, \bar{x} , than the classical estimator $\hat{\xi}$ [83]. These properties relate to the least squares method of fitting which minimises the errors in the direction of the regression, i.e., regressing y on x minimises the vertical errors about the fitted line.

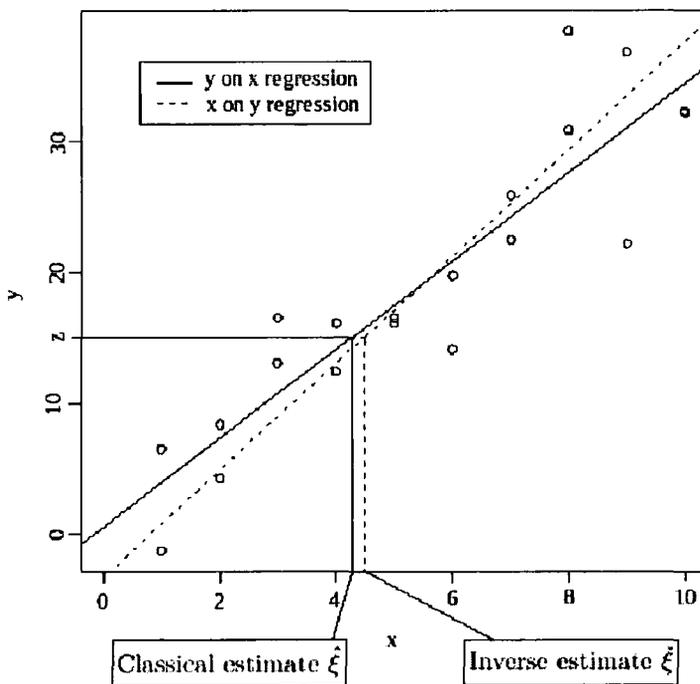


Figure 3.1: Comparison of the classical estimate using regression of y on x and the inverse estimate from the regression of x on y , using a simulated data set.

The use of the classical or inverse estimator in linear calibration problems has been heavily debated [87]. Krutchkoff (1967, [63]) promoted the use of the inverse

estimate, countering the classical approach that had been favoured since at least Einshart, 1939 [44]. Statistical arguments have been made for both estimators. Krutchkoff noted that the mean of the classical estimator $\hat{\xi}$ does not exist, and its mean square error is infinite, while the inverse estimator $\check{\xi}$ has finite mean and variance.

Brown (1993, [27]) shows that the mean and variance of the asymptotic distributions of $\hat{\xi}$ and $\check{\xi}$ are

$$E[\hat{\xi}] = \xi, \quad \text{Var}[\hat{\xi}] = \frac{\sigma^2}{\beta^2}, \quad (3.9)$$

$$E[\check{\xi}] = r^2\xi + (1 - r^2)\bar{x}, \quad \text{Var}[\check{\xi}] = r^4\frac{\sigma^2}{\beta^2}. \quad (3.10)$$

$\hat{\xi}$ is unbiased asymptotically, where $\check{\xi}$ is biased and $E[\check{\xi}]$ is the weighted average of ξ and the mean of the x -values, \bar{x} . So the bias is large for $\check{\xi}$ when the value of ξ is far from the mean.

Calibration problems can be grouped into two types; controlled or random calibration. In controlled calibration, the x -values are fixed by experimental design, where as in random calibration they are randomly selected.

The inverse estimate is more commonly used in random calibration situations, as X can be considered a random variable with (X, Y) jointly distributed [27] and so it is not as unintuitive to regress x on y . However, for controlled calibration there is no guarantee that the unknown ξ is 'like' the fixed x_1, \dots, x_n and so inferences should be restricted to be from y conditional on x , i.e. from the regression of y on x and hence the classical estimator.

The SAR procedure is a controlled calibration problem, and the classical estimate is traditionally used by practitioners of luminescence dating to estimate palaeodose from the data.

3.1.1 Bayesian Calibration

The linear calibration problem has also been tackled from a Bayesian perspective, initially by Hoadley [56]. Hoadley notes that the inverse estimator $\check{\xi}$ is a Bayes estimator if the prior distribution for ξ is student-t, with $n - 3$ degrees of freedom, scale $[(n + 1)/(n - 3)]^{\frac{1}{2}}$ and mean 0.

Hoadley goes on to propose a Bayes solution using non-informative priors, but a general form for the prior density was taken where [87]

$$P[\alpha, \beta, \sigma^2, \xi] \propto P[\alpha, \beta, \sigma^2]P[\xi]. \quad (3.11)$$

Brown [26] also makes this assumption that ξ is independent a priori of the other parameters, as well as stating that in controlled calibration,

$$P[\xi|X] = P[\xi] \quad (3.12)$$

(that is, the choice of X provides no information about ξ). However, X might provide information about ξ , if the values are chosen to lie in a region close to the expected value of ξ . Brown suggests that the prior distribution of ξ should reflect the knowledge of ξ contained in X , and so (3.12) still holds.

The Bayesian analysis of controlled calibration has been tackled by a number of other people, including Dunsmore [42], Hunter and Lamboy [54] and more recently Kacker et al [60]. However, all these approaches base the estimate of ξ on a single calibration line (though Dunsmore [42] considers each of the (x, y) points to be from independent experiments). In the application of luminescence dating, a series of calibration lines are used, one for each aliquot, and an estimate for ξ (the palaeodose) will be found from each. The distribution of these estimates around the mean value of palaeodose is then considered.

3.2 Palaeodose Evaluation using the Combined Aliquot Model

Consider the single aliquot regeneration (SAR) protocol for evaluating palaeodose (Section 2.2.2). Denote natural luminescence signal, after background correction, \mathcal{Y}_{R_j} for the $j = 1, \dots, J$ aliquots with common preheat temperature T . A series of known doses are then applied to each of the aliquots, the luminescence signal produced is measured and, for the age of samples considered here, a linear model is fitted to the sensitivity corrected response.

The J aliquots which have the same preheat temperature T will produce palaeodoses x_{Rj} , $j = 1, \dots, J$. We will suppose that these can be related in a similar manner to Galbraith's central age model [47], so that for each j

$$x_{Rj} = x_R + \delta_j^{xR} \quad (3.13)$$

where x_R denotes the the mean palaeodose value at preheat temperature T . The δ_j^{xR} are taken to be independently distributed; $\delta_j^{xR} \sim N(0, \gamma_R^2)$ for all j , γ_R to be specified by the expert. This parameter reflects the heterogeneity in the palaeodose evaluation process, from sources discussed in the previous chapter. It is also possible to adapt the approach so that γ_R^2 is treated as an unknown parameter with prior density taken as an inverse gamma distribution, and the effect of this is considered through an example in Section 3.5. The purpose of this model is to evaluate the palaeodose, x_R .

If each individual aliquot j with preheat temperature T has laboratory doses x_{ij} applied, $i = 1, \dots, n_j$ for each regeneration, let the resulting luminescence intensities after appropriate sensitivity corrections and subtraction of the background signal be denoted \mathcal{Y}_{ij} (i.e. the L_i/T_i values in Section 2.2.2).

The dose response is often variable between aliquots, i.e. the natural luminescence values for each aliquot can be quite different, even though they have come from the same environment (and thus will have been exposed to similar radiation levels). So, the luminescence sensitivity of each of the aliquots is different, so in turn will be the gradients of the linear fits. In order for the linear coefficients from each aliquot to be directly compared, the luminescence intensities are standardised by natural luminescence. That is, the natural luminescence for each aliquot is adjusted so that, say, $y_{Rj} = 10000$ counts $\forall j$. Then the standardised luminescence intensities y_{ij} are relatively adjusted,

$$y_{ij} = y_{Rj} \frac{\mathcal{Y}_{ij}}{\mathcal{Y}_{Rj}} \quad (3.14)$$

where y_{Rj} takes some suitable value and is the same for all j . We then consider the data to be $\{(y_{Rj}, y_{ij}), i = 1, \dots, n_j, j = 1, \dots, J\}$.

With the assumption that there is a linear relationship between x and y (Section

2.2.2), let

$$y_{ij} = \alpha_j + \beta_j x_{ij} + \epsilon_{ij} \quad (3.15)$$

where independent Gaussian errors are assumed, with $\epsilon_{ij} \sim N(0, \sigma^2)$, and σ^2 is unknown for $i = 1, \dots, n_j$. The use of standardised data, y_{ij} , rather than the measured values \mathcal{Y}_{ij} means that the ϵ_{ij} are not strictly independent. However, independence is assumed as in our experience the errors around the line are small.

The palaeodose estimate x_{Rj} and natural luminescence y_{Rj} also make a point on the line,

$$y_{Rj} = \alpha_j + \beta_j x_{Rj} + \epsilon_{Rj} \quad (3.16)$$

where x_{Rj} is to be found and $\epsilon_{Rj} \sim N(0, \sigma^2)$.

The assumption of normally distributed errors ϵ_{ij} is considered reasonable, based on the counting statistics involved in luminescence measurement (Section 2.2.6). The current method to evaluate x_{Rj} involves back-interpolating from the natural luminescence y_{Rj} on to the fitted line, which follows the classical estimator for a calibration problem.

The linear coefficients α_j, β_j , can be modelled in relation to the mean values α, β for preheat temperature T by

$$\alpha_j = \alpha + \delta_j^\alpha \quad \delta_j^\alpha \sim N(0, \gamma_\alpha^2) \quad (3.17)$$

$$\beta_j = \beta + \delta_j^\beta \quad \delta_j^\beta \sim N(0, \gamma_\beta^2) \quad (3.18)$$

for aliquots $j = 1, \dots, J$. The parameters $\gamma_\alpha, \gamma_\beta$ are to be specified by the expert. We specify a correlation, ρ , between α and β . This induces a correlation between α_j, β_j , but the $\delta_j^\alpha, \delta_j^\beta$ are taken to be independent, so that the covariance between α_j and β_j is the same as that between α and β .

3.2.1 Likelihood

Let D denote the data from the J aliquots with preheat temperature T . For aliquot j the data runs over $(x_{1j}, y_{1j}), \dots, (x_{n_j j}, y_{n_j j})$, with the natural luminescence values y_{Rj} for each of these J aliquots. The likelihood $L(\Theta)$ for data D can be expressed

as:

$$\begin{aligned}
 L(\Theta) = P[D|\Theta] &= \prod_{j=1}^J \prod_{i=0}^{n_j} \left[\frac{1}{\sqrt{2\pi\sigma}} \exp \left\{ -\frac{1}{2\sigma^2} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \right] \\
 &= \frac{1}{(2\pi\sigma^2)^{\sum_{j=1}^J (n_j+1)/2}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\}
 \end{aligned} \tag{3.19}$$

where $y_{0j} = y_{Rj}$ (y_{Rj} is the standardised natural luminescence and so is the same for all j), $x_{0j} = x_{Rj}$ and Θ is the set of parameters

$$\Theta = \{ \sigma^2, x_R, \alpha, \beta, x_{Rj}, \alpha_j, \beta_j, \text{ for } j = 1, \dots, J \}. \tag{3.20}$$

3.2.2 Prior Distributions

Gaussian prior distributions were employed for parameters x_R , α , β and a gamma distribution for the precision, $1/\sigma^2$, where

$$x_R \sim N(\mu_R, \sigma_R^2) \tag{3.21}$$

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} \sim N \left(\begin{pmatrix} m_\alpha \\ m_\beta \end{pmatrix}, \begin{pmatrix} \sigma_\alpha^2 & \rho\sigma_\alpha\sigma_\beta \\ \rho\sigma_\alpha\sigma_\beta & \sigma_\beta^2 \end{pmatrix} \right) \tag{3.22}$$

$$\frac{1}{\sigma^2} \sim \Gamma \left(\frac{d}{2}, \frac{a}{2} \right). \tag{3.23}$$

This choice of prior distribution is computationally convenient, while being flexible enough to allow meaningful prior information to be represented.

3.2.3 Prior Elicitation

The prior distribution over Θ reflects the judgements about the parameter values before the data are observed. Prior elicitation is an important stage in a Bayesian analysis, where an expert's knowledge is converted into prior distributions and suitable hyperparameters are specified.

In palaeodose evaluation using SAR, it is necessary for the palaeodose to lie within the range of laboratory doses applied (to ensure the linear approximation to

the dose response curve is still appropriate, i.e. they are not in the region approaching saturation). In routine dating, the practitioner thinks about the region which is likely to contain palaeodose to determine suitable choices of regenerative dose, and this can be easily translated to elicitation of the prior distributions.

In many dating situations, a rough age of the sample can be inferred using the local archaeology. An experienced practitioner will have an idea of a possible dose rate, and these two estimates will be used in the age equation to give a broad indication of palaeodose. This is used to select a possible range of regenerative doses which contains the palaeodose estimates from each aliquot. This is only a rough guide to the palaeodose value, so preliminary experiments are carried out on a small part of the sample to ensure that the laboratory doses chosen to regenerate signal in the SAR procedure are suitable. These preliminaries are also used as an indication of the luminescence characteristics for the sample, and hence the suitability of the sample for dating.

Preliminary Experiments

The preliminary experiments typically comprise one or two aliquots of the sample which are prepared and the SAR procedure used to evaluate palaeodose. The regenerative doses, initially chosen using a rough estimate of the sample age, are applied in the SAR procedure to produce an estimate for palaeodose. A single preheat temperature is usually used for all measurements. If this palaeodose estimate does not lie within the range of the regenerative doses, then the irradiated doses are adjusted and further measurements made until the palaeodose estimates produced fall in the middle of the range. So, the practitioner is aware of likely palaeodose values before the data are observed, which can be used to elicit values for μ_R , σ_R^2 , the mean and variance of the prior normal distribution for x_R .

The mean for α , m_α , is usually taken to be 0, as this indicates a judgement that no luminescence signal will be observed if no dose is applied (i.e. there is no recuperation [6]). A value of m_β is determined using m_α , μ_R and the standardised natural luminescence signal so that $y_R = m_\alpha + m_\beta\mu_R$. The prior standard deviation of α and β will be based on past dating experience. The correlation, ρ , will be

negative, and considered to be small.

The parameters γ_R , γ_α , γ_β are also to be specified by the expert. Judgements are made about how the aliquot estimates will differ from the mean values for x_R , α and β , using expertise in luminescence dating. Sensitivity analysis is carried out to investigate how influential these judgements are for the posterior distributions.

3.3 Posterior Distributions

We are interested in the posterior distribution for palaeodose with a preheat at temperature T . This probability distribution for x_R combines the information from the data with the prior judgements made, using Bayes Theorem;

$$P[x_R|D] = \frac{P[D|x_R]P[x_R]}{P[D]} \quad (3.24)$$

where D represents the data observed. This posterior distribution is difficult to calculate directly as the likelihood $P[D|x_R]$ is complicated. However, the likelihood distribution $P[D|\Theta]$, where Θ is the set of all the parameters (3.20), is known (3.19). We can therefore find the posterior distribution of x_R conditional on the remainder of the parameter set;

$$P[x_R|D, \Theta \setminus x_R] = \frac{P[D|\Theta]P[x_R|\Theta \setminus x_R]}{P[D|\Theta \setminus x_R]} \quad (3.25)$$

with $\Theta \setminus x_R$ denoting the set Θ with x_R removed.

The conditional posterior distributions for all of the parameters can be used to estimate the posterior distributions using a Gibbs Sampler [52], a Markov Chain Monte Carlo (MCMC) method. The Gibbs Sampler draws from the posterior conditional distributions of all the parameters in turn, updating the values with each draw. For example, if (for simplicity) the data comprised one aliquot ($J = 1$) then the algorithm for the Gibbs sampler would follow:

$$\begin{aligned}
\text{Draw } x_R^{(1)} & \text{ from } P \left[x_R \mid D, x_{R1}^{(0)}, \alpha^{(0)}, \alpha_1^{(0)}, \beta^{(0)}, \beta_1^{(0)}, \sigma^{2(0)} \right] \\
\text{Draw } x_{R1}^{(1)} & \text{ from } P \left[x_{R1} \mid D, x_R^{(1)}, \alpha^{(0)}, \alpha_1^{(0)}, \beta^{(0)}, \beta_1^{(0)}, \sigma^{2(0)} \right] \\
\text{Draw } \alpha^{(1)} & \text{ from } P \left[\alpha \mid D, x_R^{(1)}, x_{R1}^{(1)}, \alpha_1^{(0)}, \beta^{(0)}, \beta_1^{(0)}, \sigma^{2(0)} \right] \\
\text{Draw } \alpha_1^{(1)} & \text{ from } P \left[\alpha_1 \mid D, x_R^{(1)}, x_{R1}^{(1)}, \alpha^{(1)}, \beta^{(0)}, \beta_1^{(0)}, \sigma^{2(0)} \right] \\
\text{Draw } \beta^{(1)} & \text{ from } P \left[\beta \mid D, x_R^{(1)}, x_{R1}^{(1)}, \alpha^{(1)}, \alpha_1^{(1)}, \beta_1^{(0)}, \sigma^{2(0)} \right] \\
\text{Draw } \beta_1^{(1)} & \text{ from } P \left[\beta_1 \mid D, x_R^{(1)}, x_{R1}^{(1)}, \alpha^{(1)}, \alpha_1^{(1)}, \beta^{(1)}, \sigma^{2(0)} \right] \\
\text{Draw } \sigma^{2(1)} & \text{ from } P \left[\sigma^2 \mid D, x_R^{(1)}, x_{R1}^{(1)}, \alpha^{(1)}, \alpha_1^{(1)}, \beta^{(1)}, \beta_1^{(1)} \right]
\end{aligned}$$

which completes one iteration of the process. Here $\alpha^{(0)}$ represents the starting value for α , and $\alpha^{(1)}$ the updated value in the first cycle. So after t iterations we would have

$$\left(x_R^{(t)}, x_{R1}^{(t)}, \alpha^{(t)}, \alpha_1^{(t)}, \beta^{(t)}, \beta_1^{(t)}, \sigma^{2(t)} \right). \quad (3.26)$$

After a sufficient number of iterations the chains converge to approximate draws from the posterior distributions after appropriate thinning and burn-in period (Section 3.4).

The Gibbs sampler was chosen here as, due to the form of prior distributions assigned to the parameters (3.2.2), the conditional posterior distributions can all be explicitly found. The detailed calculation for the posterior conditional distribution of x_R follows, along with the outline of the distributions for the remainder of the parameters. Detailed derivations can be found in Appendix C.

Conditional Posterior Distribution for Palaeodose

The conditional posterior distribution for x_R is:

$$P[x_R | D, \Theta \setminus x_R] \propto P[D | \Theta] P[x_R | \Theta \setminus x_R]. \quad (3.27)$$

Here $P[D | \Theta] = P[D | \Theta \setminus x_R]$, and x_R is conditionally independent of D , Θ , given x_{Rj} , the palaeodose values from aliquots $j = 1, \dots, J$, so

$$\begin{aligned}
P[x_R|\Theta \setminus x_R] &\propto P[x_R|x_{R1}, \dots, x_{RJ}] \\
&\propto \left(\prod_{j=1}^J P[x_{Rj}|x_R] \right) P[x_R]
\end{aligned} \tag{3.28}$$

as the x_{Rj} are conditionally independent given x_R (3.13). Thus the conditional posterior distribution for x_R is:

$$\begin{aligned}
P[x_R|D, \Theta \setminus x_R] &\propto \exp \left\{ -\frac{1}{2\sigma_R^2} [x_R - \mu_R]^2 \right\} \prod_{j=1}^J \exp \left\{ -\frac{1}{2\gamma_R^2} [x_{Rj} - x_R]^2 \right\} \\
&\propto \exp \left\{ -\frac{1}{2\sigma_R^2} [x_R - \mu_R]^2 \right\} \exp \left\{ -\frac{1}{2\gamma_R^2} \sum_{j=1}^J [x_{Rj} - x_R]^2 \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{1}{\sigma_R^2} + \frac{J}{\gamma_R^2} \right] x_R^2 - 2 \left[\frac{\mu_R}{\sigma_R^2} + \frac{\sum_{j=1}^J x_{Rj}}{\gamma_R^2} \right] x_R \right) \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2} \right) \left[x_R - \frac{\frac{\sum_{j=1}^J x_{Rj}}{\gamma_R^2} + \frac{\mu_R}{\sigma_R^2}}{\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2}} \right]^2 \right\}
\end{aligned} \tag{3.29}$$

so that

$$x_R|\Theta \setminus x_R, D \sim N \left(\frac{\frac{\sum_{j=1}^J x_{Rj}}{\gamma_R^2} + \frac{\mu_R}{\sigma_R^2}}{\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2}}, \left(\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2} \right)^{-1} \right). \tag{3.30}$$

Conditional Posterior Distribution for Aliquot Palaeodose Estimates

Each of the J palaeodose values x_{Rj} , $j = 1, \dots, J$ from the J aliquots with preheat T , are only dependent on x_R and have conditional posterior distributions of the form

$$P[x_{Rj}|D, \Theta \setminus x_{Rj}] \propto P[D|\Theta] P[x_{Rj}|x_R] \tag{3.31}$$

$$\propto \exp \left\{ -\frac{1}{2} \left(\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma^2} \right) \left[x_{Rj} - \frac{\frac{\beta_j(y_{Rj} - \alpha_j)}{\sigma^2} + \frac{x_R}{\gamma^2}}{\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma^2}} \right]^2 \right\} \tag{3.32}$$

So that

$$x_{Rj}|D, \Theta \setminus x_{Rj} \sim N \left(\frac{\frac{\beta_j(y_{Rj} - \alpha_j)}{\sigma^2} + \frac{x_R}{\gamma^2}}{\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma^2}}, \left(\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma^2} \right)^{-1} \right). \tag{3.33}$$

where the mean is a weighted average of the information from the data and the mean palaeodose value x_R .

Conditional Posterior Distributions for the Regression Coefficients

The constant in the linear regression, α , is conditionally dependent on the gradient, β and the aliquot values of the intercept, α_j for $j = 1, \dots, J$. These α_j are independent conditional on α , so that

$$P[\alpha|D, \Theta \setminus \alpha] \propto P[D|\Theta]P[\alpha_1, \dots, \alpha_J|\alpha, \beta, \beta_1, \dots, \beta_J]P[\alpha|\beta] \quad (3.34)$$

$$\propto P[D|\Theta] \left(\prod_{j=1}^J P[\alpha_j|\alpha, \beta, \beta_j] \right) P[\alpha|\beta] \quad (3.35)$$

$$\propto \exp \left\{ -\frac{1}{2} \left(\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)} \right) \left[\alpha - \frac{\frac{\sum_{j=1}^J (\alpha_j - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta})}{\gamma_\alpha^2(1-\rho^2)} + \frac{(m_\alpha + \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta})}{\sigma_\alpha^2(1-\rho^2)}}{\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)}} \right]^2 \right\} \quad (3.36)$$

and thus

$$\alpha|D, \Theta \setminus \alpha \sim N \left(\frac{\frac{\sum_{j=1}^J (\alpha_j - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta})}{\gamma_\alpha^2(1-\rho^2)} + \frac{(m_\alpha + \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta})}{\sigma_\alpha^2(1-\rho^2)}}{\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)}}, \left(\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)} \right)^{-1} \right). \quad (3.37)$$

For the α_j s, the dependencies are on β_j , α and β so the conditional posterior distribution for each $j = 1, \dots, J$ is

$$P[\alpha_j|D, \Theta \setminus \alpha_j] \propto P[D|\Theta]P[\alpha_j|\alpha, \beta, \beta_j] \quad (3.38)$$

$$\propto \exp \left\{ -\frac{1}{2} \left(\frac{n_j + 1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)} \right) \right. \quad (3.39)$$

$$\left. \left[\alpha_j - \frac{\frac{\sum_{i=1}^{n_j} (y_{ij} - \beta_j x_{ij})}{\sigma^2} + \frac{\alpha + \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta}}{\gamma_\alpha^2(1-\rho^2)}}{\frac{n_j + 1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)}} \right]^2 \right\} \quad (3.40)$$

leading to

$$\alpha_j|D, \Theta \setminus \alpha_j \sim N \left(\frac{\frac{\sum_{i=1}^{n_j} (y_{ij} - \beta_j x_{ij})}{\sigma^2} + \frac{\alpha + \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta}}{\gamma_\alpha^2(1-\rho^2)}}{\frac{n_j + 1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)}}, \left(\frac{n_j + 1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)} \right)^{-1} \right). \quad (3.41)$$

Similarly, β is dependent on α and β_1, \dots, β_j ;

$$P[\beta|D, \Theta \setminus \beta] \propto P[D|\Theta]P[\beta|\alpha, \alpha_1, \dots, \alpha_J, \beta_1, \dots, \beta_J] \quad (3.42)$$

$$\propto P[D|\Theta] \prod_{j=1}^J (P[\beta_j|\beta, \alpha, \alpha_j]P[\beta|\alpha]) \quad (3.43)$$

$$\propto \exp \left\{ -\frac{1}{2} \left(\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)} \right) \left[\beta - \frac{\frac{\sum_{j=1}^J (\beta_j - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha})}{\gamma_\beta^2(1-\rho^2)} + \frac{(m_\beta + \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha})}{\sigma_\beta^2(1-\rho^2)}}{\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)}} \right]^2 \right\}, \quad (3.44)$$

so

$$\beta|D, \Theta \setminus \beta \sim N \left(\frac{\frac{\sum_{j=1}^J (\beta_j - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha})}{\gamma_\beta^2(1-\rho^2)} + \frac{(m_\beta + \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha})}{\sigma_\beta^2(1-\rho^2)}}{\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)}}, \left(\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)} \right)^{-1} \right). \quad (3.45)$$

Each β_j is only dependent on α , β and α_j , so the conditional posterior distribution is

$$P[\beta_j|D, \Theta \setminus \beta_j] \propto P[D|\Theta]P[\beta_j|\beta, \alpha_j, \alpha] \quad (3.46)$$

$$\propto \exp \left\{ -\frac{1}{2} \left(\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)} \right) \left[\beta_j - \frac{\frac{\sum_{i=1}^{n_j} (x_{ij} y_{ij} - \alpha_j x_{ij})}{\sigma^2} + \frac{\beta + \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha}}{\gamma_\beta^2(1-\rho^2)}}{\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)}} \right]^2 \right\}, \quad (3.47)$$

so that

$$\beta_j|D, \Theta \setminus \beta_j \sim N \left(\frac{\frac{\sum_{i=1}^{n_j} (x_{ij} y_{ij} - \alpha_j x_{ij})}{\sigma^2} + \frac{\beta + \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha}}{\gamma_\beta^2(1-\rho^2)}}{\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)}}, \left(\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)} \right)^{-1} \right). \quad (3.48)$$

Conditional Posterior Distribution for σ^2

Finally, the conditional posterior distribution for σ^2 is given by

$$P[\sigma^2|D, \Theta \setminus \sigma^2] \propto P[D|\Theta]P[\sigma^2], \quad (3.49)$$

so that

$$\sigma^2|D, \Theta \setminus \sigma^2 \sim \text{Inv}\Gamma \left(\frac{d + Jn}{2}, \frac{1}{2} \left(\sum_{i=0}^n \sum_{j=1}^J (y_{ij} - \alpha_j - \beta_j x_{ij})^2 + a \right) \right). \quad (3.50)$$

Conditional Posterior Distribution for γ_R^2

Here the parameter γ_R^2 (3.13) is fixed, and its value specified by the expert. The model can be adapted so that γ_R^2 is a random variable, with a prior distribution following an inverse gamma form. Another step would be added in the Gibbs Sampler, drawing from the posterior conditional distribution for γ_R^2 , and updating its value accordingly, with

$$P[\gamma_R^2|D, \Theta] \propto P[D|\Theta, \gamma_R^2] P[\gamma_R^2|x_R, x_{R1}, \dots, x_{RJ}] \quad (3.51)$$

$$\propto \left(\prod_{j=1}^J P[x_{Rj}|x_R, \gamma_R^2] \right) P[\gamma_R^2] \quad (3.52)$$

leading to

$$\gamma_R^2|D, \Theta \sim \text{Inv}\Gamma \left(\frac{J+b}{2}, \frac{1}{2} \left(\sum_{j=1}^J (x_{Rj} - x_R)^2 + c \right) \right). \quad (3.53)$$

3.4 Stability and Convergence of the Sampler

A simulated data set was used to investigate the convergence and stability of the Gibbs sampler detailed above. Using simulated data means that the palaeodose value is known, and so can be used as a comparison tool with the value achieved through the MCMC simulation. A number of different diagnostic methods have been used [43], including those proposed by Gelman and Rubin [49], Geweke [51] and graphical methods. It is useful to analyse the sampler in a variety of ways, as each method provides evidence of convergence, rather than being rigorously conclusive.

3.4.1 Data

The palaeodose x_R of the simulated data set was chosen to be 500 mGy, with $\alpha = 0$ and $\beta = 10$. The data were selected to comprise three aliquots, so $J = 3$. To simulate the aliquot palaeodose estimates (x_{Rj} , $j = 1, \dots, J$), γ_R (3.13) was set at

5, and draws taken from $x_{Rj}|x_R \sim N(x_R, \gamma_R^2)$. Similarly, the α_j , β_j were simulated, with $\gamma_\alpha = 5$, $\gamma_\beta = 5$ and $\rho = -0.3$ (3.18).

Five regenerative doses were used, 2 of which being repeated values. The luminescence responses were calculated using the linear relationship with dose (3.16) set out in the model, with the the precision, τ , where $\tau = \frac{1}{\sigma^2}$ simulated using a gamma distribution with mean 0.01, variance 0.0015. The natural luminescence values, were calculated using $y_{Rj} = \alpha_j + \beta_j x_{Rj}$. In order for the 3 aliquots to be directly comparable, the luminescence response values were standardised against the natural luminescence for each aliquot, with y_{Rj} set at 5000. The simulated data set, along with the R code used to simulate it, is shown in Appendix D.

The sampler was run using the conditional distributions detailed in Section 3.3, with 5 chains each of 50,000 iterations.

3.4.2 Analysis of Chains

The raw trace plot of the x_R simulations from one of the chains of the sampler is shown in Figure 3.2, where the actual value of the x_R is 500mGy, and x_R is the parameter of interest in this analysis. The first plot shows the full 50,000 iterations, and the second looks at the first 1000 iterations. These plots show consistent behaviour, and indicate that the sampler converges quickly, with no clear burn-in period. The true value of 500mGy for x_R is also returned. To investigate the convergence properties further, different starting values for the parameters were chosen, and even when these were far from the expected value of the parameters the chain appeared to converge. The trace plots with such starting values are shown in Appendix B.1.

Often the initial iterations are discarded, to remove the influence of the starting distributions [50]. This is known as the burn-in. Different burn-in periods were tried and the effect on the inferences made is shown in Table 3.1. The length of the burn-in period does not affect the mean and variance of x_R estimated from the iterations in this example. When no burn-in is present, the standard deviation is increased, but only a small number of iterations needed to be discarded to remove this effect.

The mean and variance of x_R was computed at various intervals along the chain,

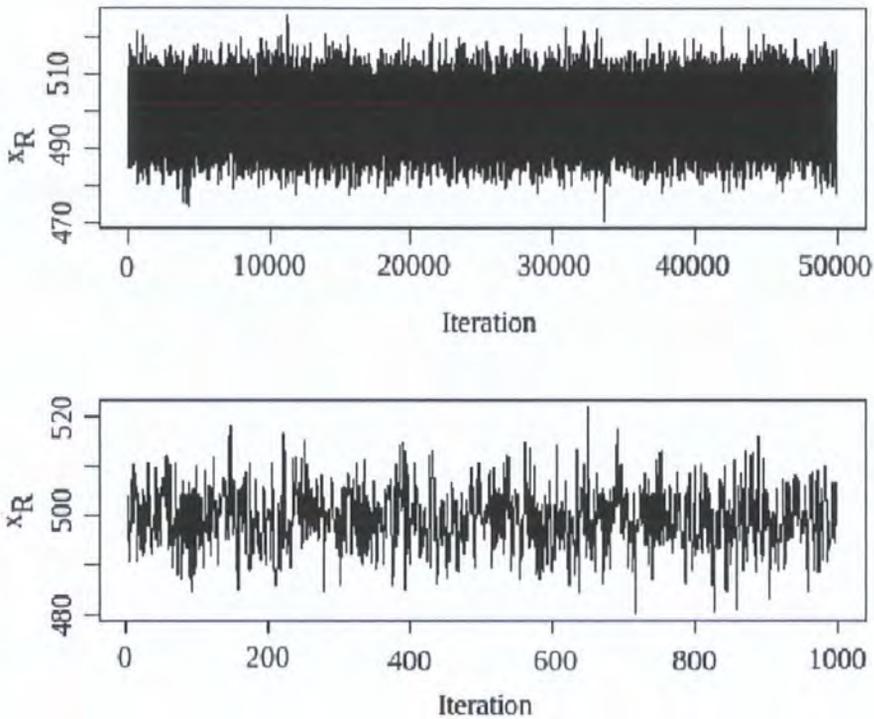


Figure 3.2: Raw trace plots of x_R simulations, with all iterations (top) and just the first 1000 iterations. These calculations are based on a simulated data set, where 500 mGy is the true value for x_R .

after a burn-in of 1000, and are shown in Table 3.2. Again the length of the chain does not affect the mean and standard deviation of the posterior palaeodose distribution, suggesting quick convergence.

Consecutive iterations in a Gibbs Sampler can be correlated [52], so to obtain approximately independent draws from the desired posterior distribution, every k^{th} value is used; that is, the chains are thinned. The level of thinning necessary to obtain approximate independence is dependent on the sampler, here Table 3.3 shows the mean and standard deviation of the posterior palaeodose distribution calculated using increasing amounts of thinning (with a burn-in period of 1000 iterations throughout). Here these summary statistics are not particularly affected by the level of thinning, as the raw trace plots shown in Figure 3.2 have good spiky characteristics. That is, each iteration does not appear to be dependent on the value of

Burn-in, n_0	Mean	SD
0	499.33	6.20
50	499.32	5.78
1000	499.32	5.78

Table 3.1: Mean and standard deviation of the posterior x_R distribution from the sampler for different burn-in periods, prior to thinning.

Iterations	Mean	SD
1000	499.32	5.76
5000	499.30	5.77
10000	499.28	5.80
30000	499.31	5.77
50000	499.32	5.78

Table 3.2: Mean and standard deviation of the posterior x_R distribution, calculated using different lengths of chain, prior to thinning.

the previous draw, and so the trace plot is seen to jump around. It was chosen to thin the chains every 5th iteration.

3.4.3 Gelman and Rubin Method

Gelman and Rubin [49] look at the convergence of m independent simulated chains with n iterations. If the sampler has converged, then the inferences made from each chain should be similar [43]. The ratio of the variance estimate to the inter-chain variance with some correction factors, R_c , is computed, the details of which are shown in Appendix B.2. $R_c \rightarrow 1$ as $n \rightarrow \infty$, and if the value is close to 1 then the sampler is considered to have reached convergence. The calculated value of R_c can be improved with further iterations, if required.

The value of R_c calculated for the simulations of x_R is 1.0013, which is close to 1, and so indicates that the sampler has reached convergence.

Thin	Mean	SD
1	499.32	5.76
2	499.31	5.74
5	499.29	5.77
10	499.31	5.73
15	499.28	5.75
20	499.30	5.74

Table 3.3: Mean and standard deviation of the posterior x_R distribution, calculated using different levels of chain thinning. A burn-in of 1000 and chain length of 50,000 iterations was used.

3.4.4 CUSUM path plots

Yu and Mykland [122] suggest using CUSUM (cumulative sum) path plots to look at the convergence and mixing of a sampler. The plots are constructed using the simulated values, say for parameter x_R , from a chain of length n (denoted $x_R^{(1)}, \dots, x_R^{(n)}$), where the first n_0 are discarded. If

$$\hat{\mu}_{x_R} = \frac{1}{n - n_0} \sum_{\tau=n_0+1}^n x_R^{(\tau)} \quad (3.54)$$

then the partial sum or CUSUM is

$$\hat{S}_t = \sum_{\tau=n_0+1}^t \left[x_R^{(\tau)} - \hat{\mu}_{x_R} \right] \quad (3.55)$$

for $t = n_0 + 1, \dots, n$, and then \hat{S}_t is plotted against t . Figure 3.3 shows the CUSUM path plot for x_R , based on one chain of the sampler with $n_0 = 1000$. For comparison, this figure also shows a CUSUM plot for $n - n_0$ draws from a normal distribution with the same mean and variance as the x_R simulations.

A smooth plot would indicate poor mixing [122], and here the CUSUM plot for the iterations of x_R from the sampler is irregular, with no sections with an increasing or decreasing trend, suggesting good mixing. It also performs well against the comparison plot computed from draws from a normal distribution, as the two paths lie within a similar range and with comparable ‘hairiness’ (not smooth), i.e.

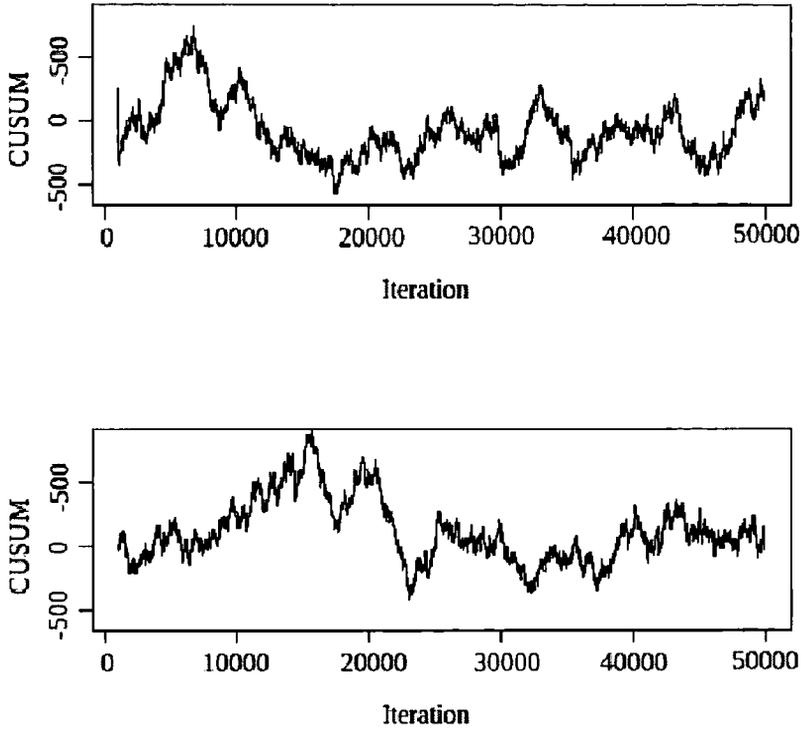


Figure 3.3: CUSUM path plots for x_R from the sampler (top) and from draws from a normal distribution with the same mean and variance as the x_R simulations (bottom), for comparison.

mixing in the sampler is comparable to that of independent draws from a normal distribution, supporting convergence of the sampler.

This method of assessing the convergence properties of the sampler is subjective, so a quantitative method based on the CUSUM plots was developed by Brooks [25]. The level of ‘hairiness’ of the CUSUM path is measured using

$$d_t = \begin{cases} 1 & \text{if } S_{t-1} > S_t \text{ and } S_t < S_{t+1} \text{ or } S_{t-1} < S_t \text{ and } S_t > S_{t+1}, \\ 0 & \text{else.} \end{cases} \quad (3.56)$$

Then

$$D_n = \frac{1}{n - n_0} \sum_{t=n_0+1}^{n-1} d_t \quad (3.57)$$

has a binomial distribution with mean $\frac{1}{2}$, variance $\frac{1}{4(n-n_0)}$ [25], and takes values between 0 and 1. For large $n - n_0$ a normal approximation can be made, and if D_n

lies outside the bounds

$$\frac{1}{2} \pm Z_{\alpha/2} \sqrt{\frac{1}{4n}} \quad (3.58)$$

100(1 - $\frac{\alpha}{2}$)% of the time then this indicates that the sampler has not converged. Brooks [25] emphasizes that these bounds are just a guide to convergence and should not be used as an exact test.

The values for D_n for x_R , with a burn-in $n_0 = 1000$ here is 0.5015, which lies inside the 95% interval, [0.4956, 0.5044], which suggests that the sampler has converged.

3.4.5 Geweke Method

Geweke [51] proposes that if a chain of the sampler has converged, then the location measures of two subsequences of the chain should be equal. Consider the two subsequences

$$\left\{ x_R^{(t)} \mid t = 1, \dots, n_A \right\} \quad \left\{ x_R^{(t)} \mid t = n^*, \dots, n \right\} \quad (3.59)$$

where $1 < n_A < n^* < n$. Then if $\left\{ x_R^{(t)} \right\}$ is stationary [51],

$$Z_n = \frac{(\bar{x}_R^A - \bar{x}_R^B)}{\sqrt{\left(\frac{1}{n_A}\right) \hat{S}^A + \left(\frac{1}{n_B}\right) \hat{S}^B}} \rightarrow^d N(0, 1) \quad (3.60)$$

where $n_B = n - n^* + 1$, $\bar{x}_R^A = \frac{1}{n_A} \sum_{t=1}^{n_A} x_R^{(t)}$, $\bar{x}_R^B = \frac{1}{n_B} \sum_{t=n^*}^n x_R^{(t)}$, and \hat{S}^A , \hat{S}^B are estimates of the variance of x_R based on the respective subsequences.

For the sampler using the simulated data, this statistic was calculated with $n_A = \frac{n}{10}$, $n^* = \frac{n}{2}$, as suggested by Geweke [51], though the choice is arbitrary [43]. This convergence diagnostic computes to 0.216 in this example, which supports chain convergence as it does not provide evidence against Z_n having a standard normal distribution.

Values for Z_n were calculated for different values of n , and the results are shown in Table 3.4. These support the initial calculations in Table 3.2 that the sampler converges quickly.

Iterations	Convergence Diagnostic
500	1.007
1000	1.012
5000	1.442
10000	1.091
50000	0.216

Table 3.4: Geweke’s convergence diagnostic computed for different chain lengths.

3.4.6 Summary of Convergence Analysis

In this simulated data example, the convergence of the sampler is clear, with each of the convergence checks carried out reaching the same conclusion, i.e. we have obtained convergence to the known palaeodose of the data. However, it is important to monitor convergence of the sampler for each example studied. Although we have only discussed the convergence of x_R here, the other parameters were also investigated in a similar manner, as convergence of one parameter does not imply convergence of the whole sample. Our experience suggests that the MCMC system is sufficiently stable not to cause problems for non-experts.

3.5 Example

Here we use the example of a heated material that was dated using OSL. one sample has been taken from the study, labelled 311-6 from Fydell House, Boston, Lincolnshire, part of a larger project on dating bricks from Medieval buildings [14].

3.5.1 Prior Specification

The hyperparameters used for the prior distributions are shown in Table 3.5.

μ_R	σ_R	γ_R	m_α	σ_α	m_β	σ_β	ρ	γ_α	γ_β	d	a
1000	100	50	0	50	10	20	-0.3	20	5	$\frac{5}{7}$	$13\frac{1}{3}$

Table 3.5: Values assigned to the prior hyperparameters for the combined aliquot model, when it is applied to 311-6, Fydell House.

One aliquot was used in the preliminary experiments for sample 311-6. The resulting luminescence responses to the laboratory doses applied are shown in Figure 3.4. Using the traditional back-interpolation from the least squares fitted line method to estimate the palaeodose, this aliquot produces a palaeodose of 934mGy. This is near the centre of the range of doses applied, indicating that the choice of laboratory dose is suitable (and so supports initial simple order of magnitude estimates of the palaeodose by the practitioner). A prior distribution for sample palaeodose was chosen to be normal with mean 1000, standard deviation 100, after discussions with Bailiff. This falls within the range of regenerative doses, but does not rely too heavily on the preliminary palaeodose estimate, as this was produced using only one aliquot.

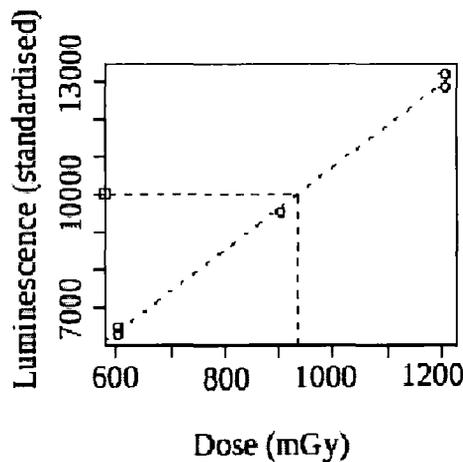


Figure 3.4: Plotted data from one of the preliminary experiments, for sample 311-6 from Fydell House.

The data have been standardised against a natural luminescence value of 10000 counts, and a linear least squares line has been fitted.

The standardised value of natural luminescence was chosen to be 10000 counts: this is an arbitrary choice. The joint prior distribution for α and β was selected to be

$$\begin{matrix} \alpha \\ \beta \end{matrix} \sim N \left(\begin{pmatrix} 0 \\ 10 \end{pmatrix}, \begin{pmatrix} 50^2 & (-0.3)(50)(20) \\ (-0.3)(50)(20) & 20^2 \end{pmatrix} \right). \quad (3.61)$$

The mean for α was chosen to be 0, as this implies that no luminescence signal will be expected if no dose has been applied. This choice, together with the choice of standardised natural luminescence, naturally leads to 10 as the prior mean of β , as then the prior mean palaeodose and natural luminescence make a point on the line with coefficients as the mean prior values of α and β .

The least squares line fitted to the preliminary data set has intercept -280, gradient 11.0. At first glance, this intercept may seem improbable considering the assumption that no signal should be seen with zero dose. However, the scale of the luminescence response needs to be considered; in comparison to the natural luminescence of 10000 counts, an intercept value of -280 is relatively close to zero. With this in mind, plus consultation from experienced practitioner Bailiff, the prior standard deviations σ_α and σ_β were selected. A smaller value was placed on σ_β because the standardisation of the data is expected to contain these values within a tighter range. The correlation between the linear coefficients is negative as, if the gradient of the fitted line were increased then, the point at which the line crosses the y-axis would be lower. This correlation is thought to be small, so a value of -0.3 was chosen.

After discussion with Bailiff, an experienced practitioner of the SAR protocol, a 5% error of the palaeodose estimates from each aliquot around the mean palaeodose was considered reasonable. So, with μ_R set at 1000mGy, γ_R was assigned a value of 50.

3.5.2 Data

The data set 311-6 comprises 20 aliquots across 5 preheat treatments in the range 210 – 250°C. The regenerative doses chosen were {603, 904, 1206, 603, 1206}, and for each aliquot the data comprises the natural luminescence, plus the luminescence response to each of the laboratory doses applied. The data are shown in Appendix G.1. Here we will initially consider evaluating the palaeodose based on the 3 aliquots with a preheat of 210°C.

The data for the 3 aliquots with a preheat temperature of 210°C are plotted in Figure 3.5, along with the least-squares fitted lines. From these lines the stan-

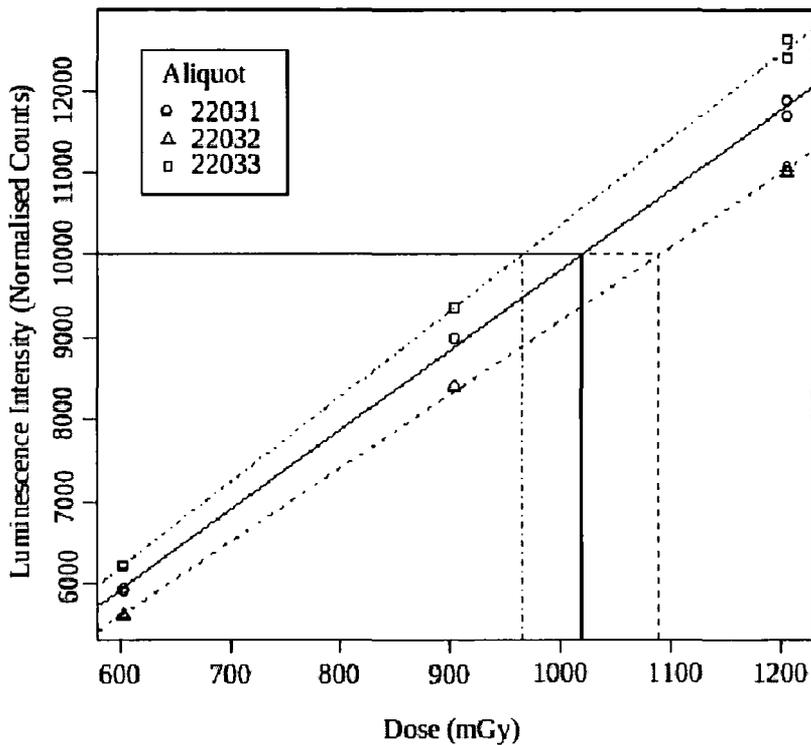


Figure 3.5: Data from aliquots with a preheat temperature of 210°C , with fitted lines being used to estimate palaeodose by back-interpolation.

The data have been standardised to a natural luminescence value of 10000 counts, so the regression lines can be directly compared.

Standardised natural luminescence is used to back-interpolate from the fitted lines to illustrate the traditional estimate for the palaeodose. Here a straight line fit to the data seems plausible, though linear diagnostics will be considered in Section 3.7.

3.5.3 Posterior Distribution

The sampler detailed in Appendix H.1, with the input as above, was run with 5 chains for 50,000 iterations. The first 1000 iterations were discarded as burn-in, and every 4th iteration taken. The convergence of the sampler was checked by looking at the trace plots of each of the parameters.

The posterior palaeodose distribution for aliquots with a preheat of 210°C is

shown in Figure 3.6, with mean 1021.5 mGy and standard deviation 28.5. This distribution is approximately normal, illustrated by the normal density that is overlaid in Figure 3.6 (dashed line), with the same mean and variance as the posterior palaeodose. In this example the posterior density for palaeodose and the normal density are almost the same.

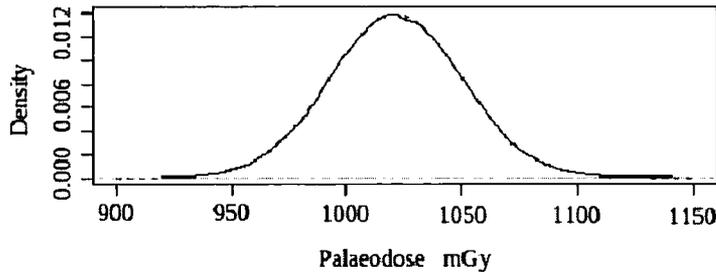


Figure 3.6: Posterior distribution of palaeodose based on three aliquots with a pre-heat of 210°C.

The dashed line represents the normal density with the same mean and standard deviation.

3.6 Sensitivity to Prior Parameters

The sensitivity of the posterior palaeodose distribution to the choice of prior parameter values is investigated here. The example of sample 311-6 from Fydell House is considered, in particular the evaluation of palaeodose at preheat 210°C. Here statistics have been quoted to a high level of precision for comparative purposes: rounding to the nearest 5mGy is accepted as appropriate in routine luminescence dating.

The hyperparameters of the prior distributions (Section 3.2.2) initially were set to the values in Table 3.5, and the reasoning for these choices is explained in Section 3.5.1.

The effect of the choice of prior mean (μ_R) and standard deviation (σ_R) of palaeodose on the posterior palaeodose distribution is now explored. With $\mu_R = 1000$, Figure 3.7 shows the influence of σ_R on the posterior distribution for palaeodose. A

small σ_R value gives greater weight to the prior distribution, and so the posterior palaeodose distribution shifts towards the prior mean and the posterior variance is reduced.

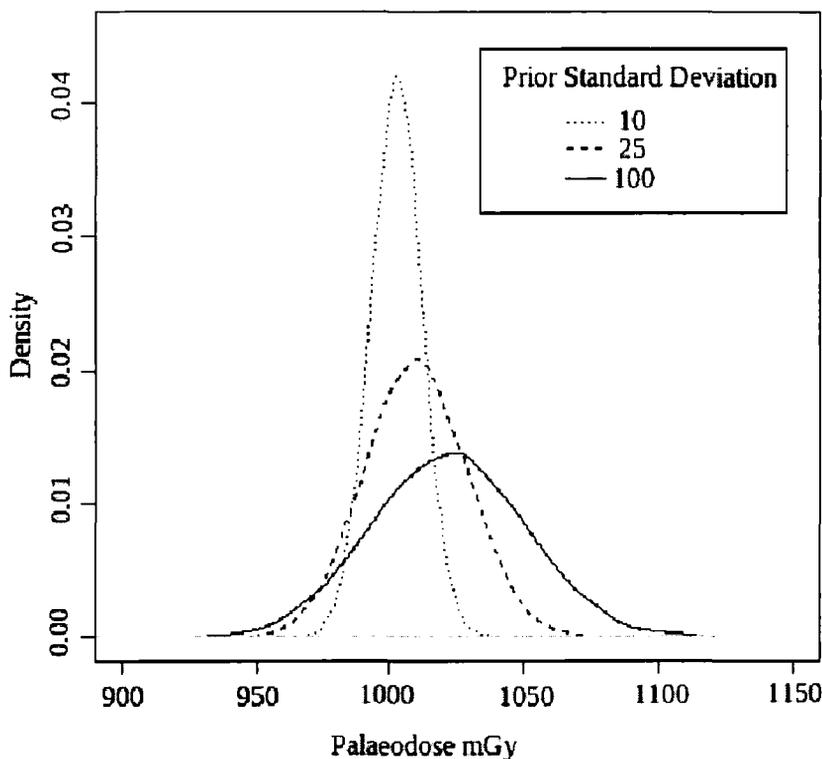


Figure 3.7: Posterior palaeodose distributions at a preheat of 210°C, for different prior standard deviations, σ_R , with the prior mean fixed at $\mu_R = 1000$.

The sensitivity of the posterior palaeodose distribution to the prior standard deviation σ_R is dependent on the choice of prior mean, μ_R . Figure 3.8 illustrates how, in this example, σ_R influences the posterior mean for palaeodose, for different μ_R . When the prior standard deviation is large, the posterior palaeodose mean is not affected by the choice of prior value for μ_R . In this example, the posterior palaeodose had mean 1020 mGy when a broader prior was used and so the information from the data dominated the analysis. When the prior beliefs are strong, this is reflected in a small value for σ_R , and so the posterior mean gravitates towards the prior.

Further details of the investigation into the influence of μ_R , σ_R on posterior palaeodose, including the statistics from the posterior distributions produced, are

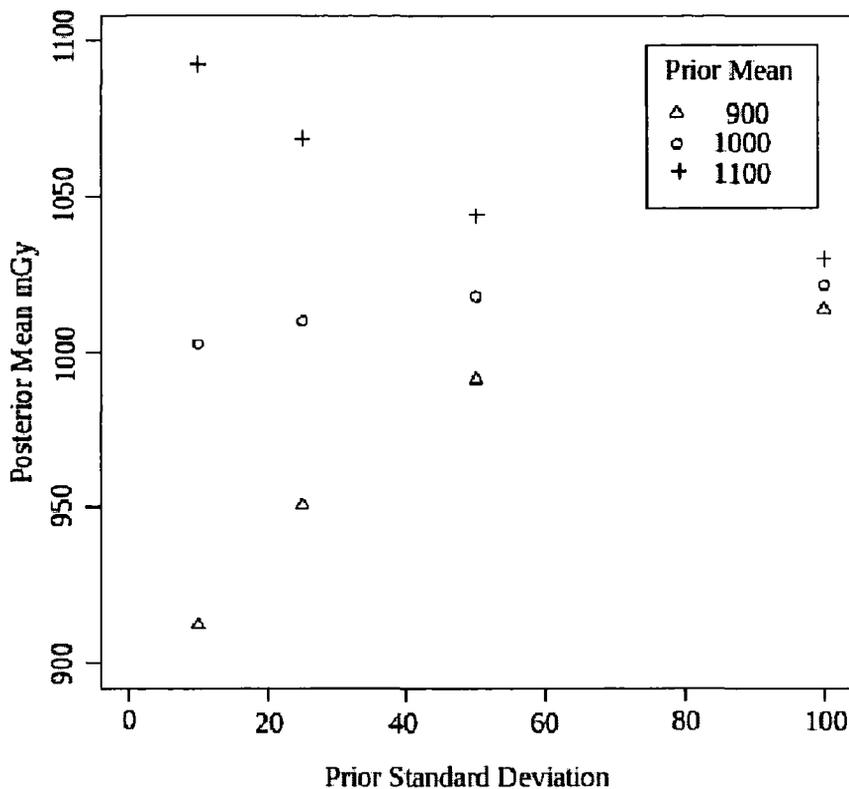


Figure 3.8: The influence of prior standard deviation, σ_R , on posterior mean palaeodose for three different choices of prior palaeodose mean.

given in Appendix E. Similar trends are observed in the posterior mean palaeodose with prior standard deviation when different prior means are used.

The sensitivity of the posterior palaeodose distribution to γ_R , the measure of spread of the palaeodose estimates from the aliquots was first considered treating γ_R as a known constant. Figure 3.9 shows how the choice of γ_R influences the posterior distributions for x_R and the aliquot estimates of palaeodose x_{R1} , x_{R2} , x_{R3} . A small γ_R value pulls the aliquot estimates towards each other and the posterior x_R distribution. As γ_R is increased, the x_{Rj} distributions spread out (as the data from the aliquots has the more dominating effect on them), and the posterior variance of x_R increases.

The posterior palaeodose distribution is formed from a combination of the aliquot estimates and the prior, and so the posterior variance will have strong influences from both γ_R and σ_R . The effect of the choice of γ_R , for different σ_R values, on the

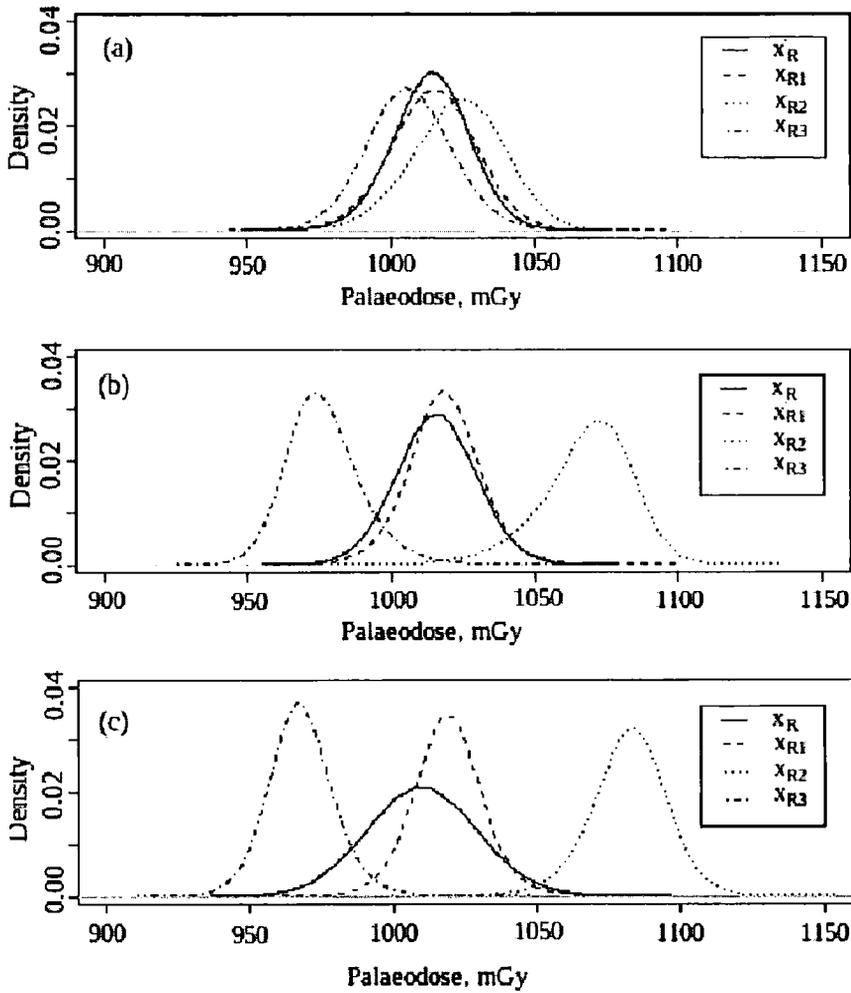


Figure 3.9: Posterior distribution for palaeodose, and palaeodose estimates from the 3 aliquots with preheat 210°C.

The values of γ_R used are (a) 10, (b) 25 and (c) 50, with prior $x_R \sim N(1000, 25^2)$.

Prior			Posterior	
μ_R	σ_R	γ_R	Mean	SD
1000	100	50	1021.4	28.2
1000	10	50	1002.4	9.4
1100	100	50	1029.9	28.7
1100	10	50	1092.3	9.5
1000	100	25	1021.8	16.7
1000	10	25	1006.0	8.6
1100	100	25	1024.9	16.5
1100	10	25	1082.8	9.8
1000	100	5	1019.0	17.5
1000	10	5	1005.6	8.5
1100	100	5	1022.3	16.7
1100	10	5	1089.4	10.0

Table 3.6: Influence of γ_R on posterior mean and standard deviation of palaeodose posterior mean and standard deviation of palaeodose is shown in Table 3.6. Results of further investigations are shown in Appendix E.

A large value for γ_R relates to the belief that the aliquot estimates x_{Rj} for palaeodose will have a large spread around x_R . So, as shown in Table 3.6, a large γ_R relates to a large posterior standard deviation if the prior standard deviation, σ_R , is also large. However, if there are strong prior beliefs, then these will dominate the posterior distribution over the data (the x_{Rj} s) when there is a large value of γ_R , and so low confidence in the data. Conversely, if γ_R is small compared to σ_R , the posterior variance reflects the spread of the estimates x_{Rj} over the aliquots.

The influence of the prior parameters for the linear coefficients on the posterior palaeodose are shown in Table 3.7. This table shows that the prior mean, variance of $(\alpha, \beta)^T$, along with the spread of the (α_j, β_j) do not have a large influence on the posterior palaeodose in this example.

Prior Parameters							Posterior Palaeodose	
m_α	m_β	σ_α	σ_β	ρ	γ_α	γ_β	Mean	Standard deviation
0	10	50	20	-0.3	5	20	1021.4	28.2
0	10	10	5	-0.3	5	20	1021.6	28.4
5	20	5	2	-0.3	5	20	1021.9	28.3
0	10	5	2	0.8	5	20	1021.9	28.8
0	10	50	20	-0.3	2	2	1021.8	28.6

Table 3.7: Influence of the choice of prior parameters for the linear coefficients on posterior palaeodose distribution.

Prior $\frac{1}{\gamma_R^2}$		Posterior x_R	
Mean	Standard deviation	Mean	standard deviation
$\frac{1}{10^2}$	10^{-3}	1019.3	42.3
$\frac{1}{10^2}$	10^{-4}	1019.4	29.1
$\frac{1}{10^2}$	10^{-5}	1021.9	18.1
$\frac{1}{25^2}$	10^{-3}	1018.8	43.2
$\frac{1}{25^2}$	10^{-5}	1022.0	25.0
$\frac{1}{50^2}$	10^{-5}	1019.7	42.1
$\frac{1}{50^2}$	10^{-8}	1021.6	28.3

Table 3.8: Posterior mean and standard deviation for palaeodose when γ_R^2 is treated as a unknown, a priori.

3.6.1 γ_R as a Random Variable

The model can be adapted to let γ_R be a random variable rather than a constant to be specified by the expert, so an extra step is added to the Gibbs sampler to update this parameter in each iteration cycle. The prior density for $\frac{1}{\gamma_R^2}$ was taken to be a gamma distribution, and the corresponding prior density is given in Equation 3.53.

Table 3.8 shows the posterior mean and standard deviation of palaeodose when this model is applied to the three aliquots with a preheat of 210°C from 311-6. These can be compared to the results calculated in Table 3.6 where γ_R^2 is a constant.

The posterior mean is not significantly affected by γ_R^2 being a random variable

instead of a constant. However, if the prior standard deviation of $\frac{1}{\gamma_R}$ is large, then the posterior variance of palaeodose is notably increased. This reflects the uncertainty about the relationship between x_R and the aliquot estimates x_{Rj} , $j = 1, \dots, J$.

The aliquot estimates x_{Rj} can be considered as the data in the posterior distribution for x_R . As, in Table 3.6, there are only 3 aliquot estimates used, the choice of prior distribution on $\frac{1}{\gamma_R}$ is influential on the posterior palaeodose. However, when there are a large number of aliquots contributing to the estimate of x_R , then this influence is not as strong, since more information from the data is available.

3.6.2 Summary

The values chosen for the prior parameters impact the posterior distribution for palaeodose in this example, most notably the posterior standard deviation. Therefore, it is important that the opinions of the expert on the parameters a priori are reflected in the specification of the prior parameters. So it would be useful if further work on prior elicitation was carried out to achieve this.

3.7 Diagnostics

3.7.1 Linear Model Diagnostics

Linear model diagnostics can be used to verify that fitting a linear model to the luminescence response to dose (3.15) is appropriate (Section 3.16). One method is to look at the residuals \hat{e}_{ij} ,

$$\hat{e}_{ij} = y_{ij} - \hat{\alpha}_j - \hat{\beta}_j x_{ij} \quad (3.62)$$

where $\hat{\alpha}_j$, $\hat{\beta}_j$ are the least squares estimates of α_j , β_j . When the residuals are plotted against fitted values, then any trend can suggest a systematic misfit [115]. The fitted values, \hat{y}_{ij} are

$$\hat{y}_{ij} = \hat{\alpha}_j + \hat{\beta}_j x_{ij}. \quad (3.63)$$

The residuals for the aliquots with a preheat of 210°C are shown in Figure 3.10. Here, for each aliquot, a funnel pattern to the residuals can be detected from the plot. This

suggests that the variance of the data around the fitted model might increase with \hat{y}_{ij} . This is often a feature of counted data [96], and here the luminescence response is measured by the counts of photons observed. The residuals plotted in Figure 3.10 are small compared with the magnitude of the luminescence response ($\sim 1\%$), with an even balance between the number of positive and negative residuals. None of the residuals have a relatively large value (compared to the residual standard errors in Table 3.9) so there are no outliers in the data. So, in this case, the residual plot does not indicate that a linear fit is inappropriate.

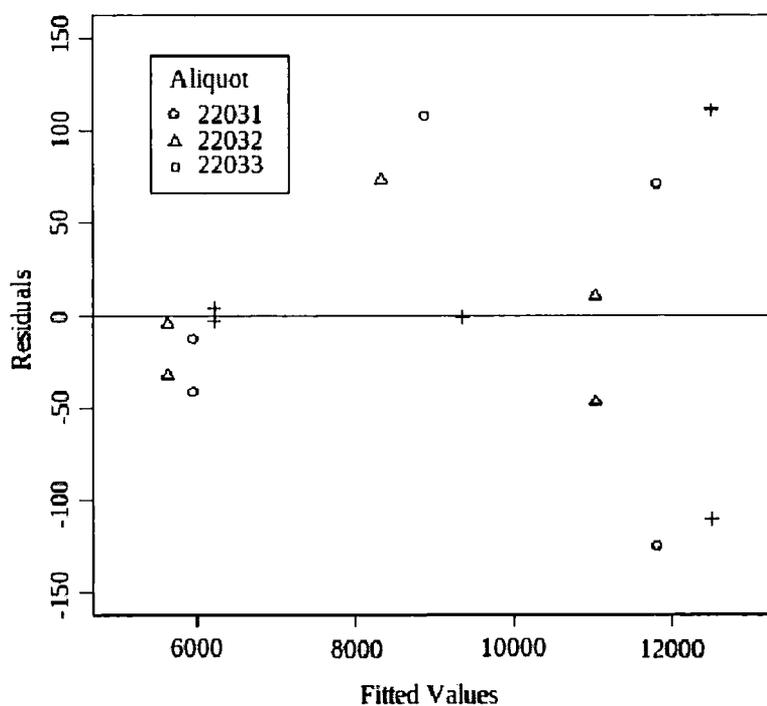


Figure 3.10: Residuals from aliquots with a preheat of 210°C , plotted against fitted values.

The correlation coefficient,

$$r = \frac{S_{xy}}{\sqrt{S_{xx}S_{yy}}} \quad (3.64)$$

can be computed to indicate the strength of the linear relationship between dose and luminescence response. If there is a perfect linear relationship between the variables, then $r = 1$. For the example considered here, the data from each of the 3 aliquots with a preheat of 210°C had a correlation coefficient of 0.999, and so indicates a

Aliquot ID	RSE
22031	106.5
22032	54.0
22033	90.4

Table 3.9: The residual standard error from the linear fit of the data from each of the three aliquots with a preheat of 210°C.

strong linear trend in the data sets. However, there are only a small amount of data points for each aliquot making it more difficult to pick up a non-linear relationship.

Another assumption made in the linear model (3.15) is that the residuals are independent and identically distributed, with $\epsilon_{ij} \sim N(0, \sigma^2)$ across the aliquots with the same preheat temperature. The value of the residual standard error for each of the aliquots with a preheat of 210°C is shown in Table 3.9. Two of these values are similar, and support the assumption that the residuals have the same standard deviation across the aliquots. The second value in the table is around half the magnitude of the other two, but with just three data points it is difficult to test the validity of the normality assumption.

When compared to the magnitude of the standardised luminescence response, the residual standard deviations are relatively small. The small sample size must also be taken into consideration, which implies a large variance on the estimates of RSE, and thus make it more difficult to detect departures from the assumptions made.

3.7.2 Diagnostics for the Bayesian Model

The difference between the prior and posterior mean for palaeodose can be compared to the posterior standard deviation as a diagnostic check of the prior specification. Here, the prior mean was chosen to be 1000 mGy, and the posterior mean for aliquots with a preheat of 210°C is 1021.5 mGy. The difference of 21.5 mGy, when compared to the posterior standard deviation of 28.5 mGy, does not suggest that there are problems with the Bayesian prior specification.

The posterior distribution for palaeodose estimated using the Gibbs sampler can

be compared to the posterior distribution of the aliquot estimates for palaeodose, x_{Rj} , $j = 1, \dots, J$ and the prior distribution. The posterior palaeodose distribution should lie within the range of these distributions, as it can be considered an average of them weighted by their variance. For example, Figure 3.11 shows these distributions for the palaeodose evaluated with a preheat of 250°C. Here the posterior palaeodose distribution straddles the distributions from the aliquot estimates, and is contained within the broad prior, so there is no reason to suggest the sampler has computed the distribution incorrectly.

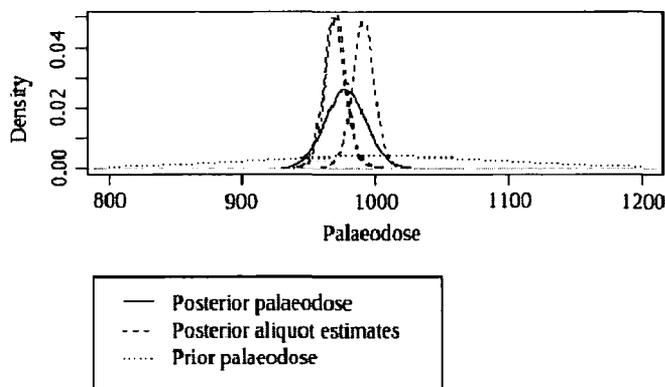


Figure 3.11: Posterior distribution for palaeodose, estimates of palaeodose from each aliquot with a preheat of 250°C and prior palaeodose distribution.

3.8 Palaeodose Evaluation for Sample 311-6

So far, only the aliquots from sample 311-6 with a preheat treatment of 210°C have been considered. However, the sample comprises 20 aliquots across 5 preheat temperatures. So, the combined aliquot model was applied to each set of aliquots, grouped by preheat temperature. The same regenerative doses were applied in each case. The same prior distributions were used at each temperature, and the posterior palaeodose mean and standard deviation achieved are shown in Table 3.10, with the distributions shown in Figure 3.12.

It is notable in Table 3.10 that the largest posterior standard deviation for palaeodose is produced using the data at a preheat of 240°C, which is based on 5 aliquots.

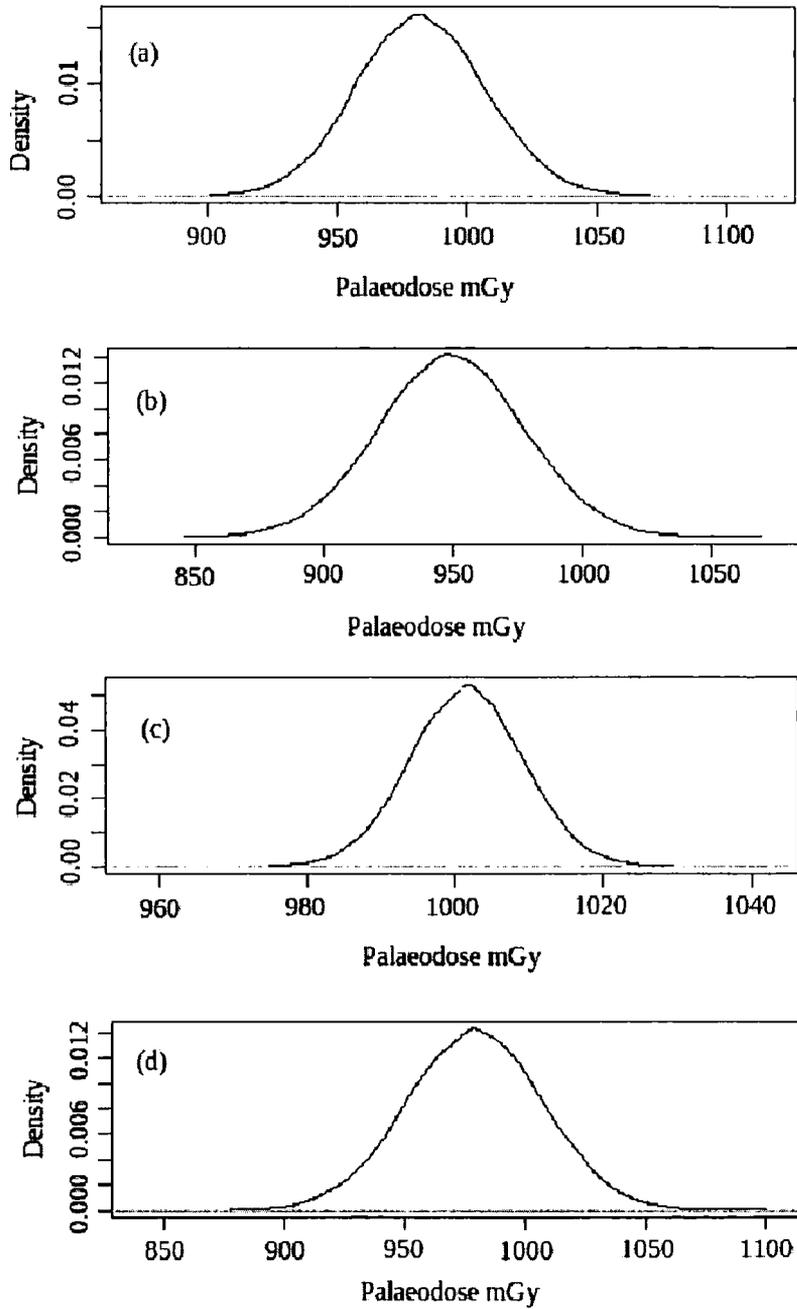


Figure 3.12: Posterior distribution for palaeodose at preheat (a) 220°C, (b) 230°C, (c) 240°C, (d) 250°C

Preheat	# aliquots	Posterior Mean	Posterior SD
210	3	1021.5	28.5
220	6	982.5	25.0
230	3	949.4	28.4
240	5	993.7	42.4
250	3	979.12	28.1

Table 3.10: Posterior mean and standard deviation of palaeodose estimates at each preheat, from sample 311-6.

This is a comparatively large number of aliquots, and so it would be expected that the posterior standard deviation would be lower at this preheat temperature. This dispersion at 240°C is due to a large spread of the posterior x_{Rj} values, rather than uncertainty around the regression lines, and illustrates the unpredictable nature of luminescence.

The method described in this chapter allows the experimenter to combine the data for a given preheat temperature, and to assess their accuracy through the posterior palaeodose standard deviation. In the next chapter we will consider the issues that arise in combining information from different preheat temperatures.

3.9 Summary Guide to evaluating palaeodose at each preheat temperature

1. Write down a range in which the palaeodose is expected to lie, based on the preliminary experiments, the choice of regenerative doses to be used, and other archaeological knowledge. Use this as the basis for the prior mean and standard deviation of palaeodose. Elicit expert opinion on the spread of the aliquot estimates of palaeodose around the palaeodose value, for γ_R .
2. Elicit values for the mean and variance of the linear coefficients, α and β , and the spread of the aliquot estimates around these values, γ_α , γ_β .
3. Standardise the (sensitivity corrected) luminescence response against natural

luminescence for each aliquot, and group the aliquots by preheat temperature.

4. For each preheat, run the Gibbs Sampler and check for convergence.
5. Perform diagnostic checks to ensure the linear fit is appropriate.
6. Look at the posterior palaeodose distribution obtained, and check that it is feasible. Note the posterior mean and variance.

Chapter 4

Preheat Plateau Model

The combined aliquot model in Chapter 3 computes an estimate for palaeodose at each preheat temperature. In standard luminescence dating, the estimates of palaeodose are plotted against preheat temperature. A region where palaeodose does not change with temperature, a preheat plateau, is considered to provide the best estimate for the sample palaeodose (Section 2.3).

In this chapter we will consider the problem of identifying a preheat plateau. Once the plateau is located, the posterior distribution for the palaeodose of the sample will be calculated. This will be illustrated using the continued example of sample 311-6 from Fydell House, Lincolnshire.

4.1 Motivation for modelling preheat plateau

The standard practice in routine luminescence dating is to identify the preheat plateau by eye. This is very subjective, and it can be easy to ‘find’ a plateau in the data when it is desired, when it could equally be argued that a different trend was present. This is particularly disconcerting considering the care taken to make the aliquot estimates as accurate and precise as possible. A robust, statistical method of plateau identification would enable the luminescence community to have a firm basis for their choice of aliquots on which to base the evaluation of sample palaeodose.

Once the posterior distribution for the location of the plateau region has been established, the set of aliquots which have undergone preheat treatments in this region

is selected, if this is clearly identified by the form of the posterior distribution. The luminescence data from these aliquots are used to evaluate the sample palaeodose distribution. Since, according to the plateau model, these aliquots all estimate the sample palaeodose, the combined aliquot model can be applied to this set to obtain the posterior sample palaeodose distribution. So, we use the methods described in Chapter 3 but apply them to the new set of aliquots which have been determined by the plateau model.

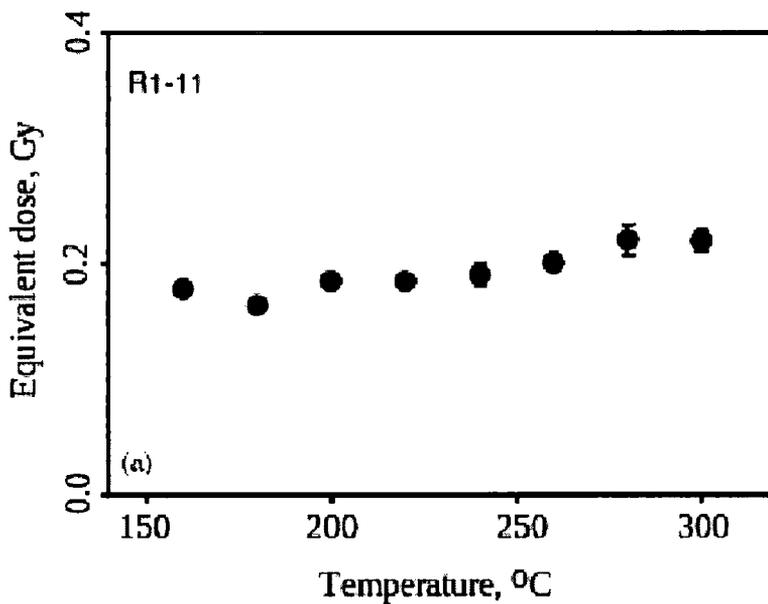


Figure 4.1: An example of variation in palaeodose evaluation with preheat temperature given in Madsen et al 2007 [69].

Various different trends for the behaviour of palaeodose with preheat temperature are observed in the literature. For example, Figure 4.1, taken from Madsen et al [69], shows a dependence of palaeodose estimation with preheat temperature. Here there is an increasing trend in the palaeodose estimates with preheat temperature, with no clear plateau region. However, a decreasing trend can also be observed, for example Figure 4.2, which Wintle and Murray used as an example of a preheat plateau, taking the data from Jacobs et al [57]. The difficulty in identifying a plateau in the data motivates the development of a model for the preheat plateau.

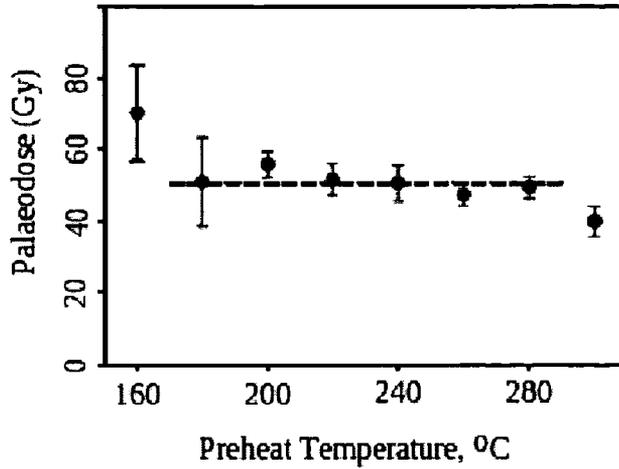


Figure 4.2: Variation of palaeodose evaluation with preheat temperature given as an example in Wintle and Murray 2006 [120], originating from Jacobs et al 2003 [57].

4.2 Model

Denote the mean palaeodose for a given preheat temperature as x_{RT} . The plateau is in the region $[T_a, T_b]$ such that

$$x_{RT} = x_{R^*} \quad \forall T \in [T_a, T_b]. \quad (4.1)$$

All aliquots which are preheated to a temperature on the plateau are evaluating the same palaeodose, x_{R^*} . This is the sample palaeodose, and used as the numerator in the age equation. The combined aliquot model (Chapter 3) can be applied to all of these aliquots to evaluate the palaeodose of the sample.

The palaeodose estimates from aliquots $j = 1, \dots, J_k$ with preheat treatments which lie in the plateau region, $\{T_k \in [T_a, T_b]\}$, are denoted x_{Rjk} where k labels the preheat temperature and j the aliquot number. These are related to the sample palaeodose, x_{R^*} , by

$$x_{Rjk} = x_{R^*} + \delta_{jk}^{xR} \quad (4.2)$$

with $\delta_{jk}^{xR} \sim N(0, \gamma_R^2)$, independent for $\{k | T_k \in [T_a, T_b], j = 1, \dots, J_k\}$. The relationship between the linear coefficients α , β and the aliquot values α_{jk} , β_{jk} from the combined aliquot model also extends to cover all aliquots which produce palaeodose

values lying on the preheat plateau;

$$\alpha_{ij} = \alpha + \delta_{jk}^{\alpha}, \quad \delta_{jk}^{\alpha} \sim N(0, \gamma_{\alpha}^2), \quad (4.3)$$

$$\beta_{ij} = \beta + \delta_{jk}^{\beta}, \quad \delta_{jk}^{\beta} \sim N(0, \gamma_{\beta}^2). \quad (4.4)$$

4.2.1 Location of the plateau

Once the location of the plateau has been identified, then the aliquots which lie on it can be used to evaluate the sample palaeodose using an extension of the combined aliquot model. However, the position of the plateau is not always obvious. The difficulty in modelling the preheat plateau stems from the uncertainty surrounding the relationship between preheat temperature T and palaeodose x_{RT} of the plateau.

Here we assume that a monotone continuous function leads to the plateau, which starts at temperature T_a at palaeodose level x_{R*} , so that for palaeodose estimate x_{RT} with a preheat T ,

$$x_{RT} = \begin{cases} x_{R*} \left(\frac{(1-\exp\{-\eta T\})\exp\{-\kappa T\}}{(1-\exp\{-\eta T_a\})\exp\{-\kappa T_a\}} \right) & T < T_a \\ x_{R*} & T \geq T_a. \end{cases} \quad (4.5)$$

We consider the four parameters x_{R*} , T_a , η , κ to be independent a priori. The two uncertain curve parameters η and κ allow a wide variety of continuous shapes of curve to be achieved before the plateau is reached. Figure 4.3 illustrates some of the different curves that can be achieved with this function, for particular values of η , κ .

Here we assume that the data do not extend to temperatures beyond the plateau, and that a ‘false plateau’ will not be observed prior to the true preheat plateau. Although such behaviour may occasionally have been observed, currently the physical reasoning behind it is not fully understood. So, as a first model for the preheat plateau, we do not consider such behaviour or the region beyond the plateau, but as theoretical knowledge is expanded the model can be further developed. As the model has been broken down into a number of separate stages, the practitioner can monitor the posterior parameter values produced and as such highlight any situations where this simplification of the problem may become an issue.

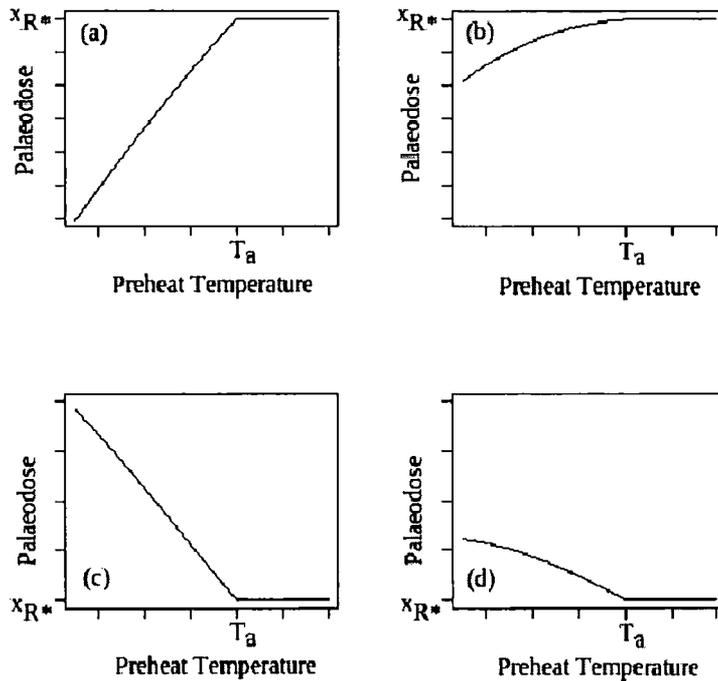


Figure 4.3: Examples of the shape of curve that can be achieved using (4.5) to represent possible relationships between palaeodose and preheat temperature.

(a) $\eta, \kappa = 0.001$, (b) $\eta, \kappa = 0.003$, (c) $\eta = 0.005, \kappa = 0.008$, (d) $\eta, \kappa = 0.008$.

4.2.2 Data Input

Current routine practice in luminescence dating is to evaluate a palaeodose estimate from each aliquot, and then plot these estimates against the preheat temperature. The plateau is then identified (or otherwise) by eye, and the palaeodose estimates which lie on the plateau are used to evaluate the palaeodose of the sample. In our Bayesian model we aim to extract the corresponding quantities which allow us to make a similar determination.

Aliquots at each preheat temperature $\{T_1, \dots, T_t\}$ have been combined to give an overall palaeodose estimate at each temperature using the combined aliquot model in Chapter 3. Let the mean of the posterior distribution for palaeodose at temperature T_i , $E[x_{RT_i}]$, be written as \bar{x}_{RT_i} , and standard deviation given by the posterior standard deviation of x_{RT_i} , denoted σ_{RT_i} for $i = 1, \dots, t$.

In the preheat plateau model, consider the data input to be the set of posterior palaeodose means at each preheat temperature, $\{\bar{x}_{RT_i}, i = 1, \dots, t\}$, found using the combined aliquot model. We assume normality of the \bar{x}_{RT_i} , with variance $\sigma_{RT_i}^2$.

As well as being analogous to current dating procedure, the reasoning behind using the information in the posterior distribution for x_{RT_i} as the data input to the plateau model is as follows. We assume here that the errors at each stage of the analysis are independent (in reality the dependencies are small). Consider observing data D to learn about x_{RT} . If $D_i \sim N(x_{RT_i}, \sigma_i^2)$ for known variance σ_i^2 , then, with vague prior information for x_{RT_i} the posterior distribution for x_{RT_i} given D_i would be approximately of the form $N(D_i, \sigma_i^2)$. So viewing the posterior means $\bar{x}_{RT_i}, i = 1, \dots, t$ as data observations with variance $\sigma_{RT_i}^2$ in the next stage of the analysis is a plausible representation of the information summarised by the posterior distribution.

This approach has been chosen as it is intuitively comparable to current procedures, is relatively easy to implement, straightforward to modify as experimental protocols change and leads to stable and manageable calculations. As seen in the example (Section 3.5.3), this normal representation is a good approximation to the form of the posterior for heated materials from a homogeneous environment.

It would be difficult to combine all of the evaluation in Chapter 3 and 4 in a single calculation because of the direct prior judgements of the magnitudes of the x_{RT_i} values. We also recognise the exploratory nature of the plateau model reflecting current theoretical understanding of the relationship between palaeodose and preheat temperature before a plateau is reached. As such it is more suited to being used mainly as a diagnostic tool to identify the presence of a preheat plateau and so in turn the aliquots which should be used to evaluate the palaeodose of the sample.

We only use the reduced form to identify the start of the preheat plateau. So, as in routine practice, the preheat plateau is used as a diagnostic tool to assess which aliquots should contribute to the assessment of the sample palaeodose evaluation. Once the plateau has been identified, we then return to the full model of Chapter 3 to evaluate the palaeodose of the sample.

4.2.3 Likelihood

The likelihood $L(x_{R^*}, T_a, \eta, \kappa)$ for data $\{\bar{x}_{RT_1}, \dots, \bar{x}_{RT_t}\}$ is

$$\begin{aligned} L(x_{R^*}, T_a, \eta, \kappa) &= P[\bar{x}_{RT_1}, \dots, \bar{x}_{RT_t} | x_{R^*}, T_a, \eta, \kappa] \\ &= \prod_{i=1}^t \left[\frac{1}{\sqrt{2\pi}\sigma_{RT_i}} \exp \left\{ -\frac{1}{2\sigma_{RT_i}^2} (\bar{x}_{RT_i} - x_{RT})^2 \right\} \right] \end{aligned} \quad (4.6)$$

with x_{RT} defined in (4.5).

4.2.4 Posterior Distribution for T_a

We now use the model as a tool to identify the plateau starting temperature, T_a . As the aliquots have preheat temperatures with increments of, for example, 10°C , the precise value of T_a is not important. The real question is between which preheat temperature interval does the plateau start, i.e. which aliquots should be used to evaluate the sample palaeodose.

Again, the Bayesian paradigm is applied to find the posterior distribution of T_a ,

$$P[T_a | \text{data}] \propto P[\text{data} | T_a] P[T_a] \quad (4.7)$$

$$\propto \int \int \int P[\text{data} | T_a, x_{R^*}, \eta, \kappa] P[T_a | x_{R^*}, \eta, \kappa] dx_{R^*} d\eta d\kappa \quad (4.8)$$

$$\propto P[T_a] \int \int \int P[\text{data} | T_a, x_{R^*}, \eta, \kappa] P[x_{R^*}, \eta, \kappa] dx_{R^*} d\eta d\kappa \quad (4.9)$$

$$\propto P[T_a] \int \int \int P[\text{data} | T_a, x_{R^*}, \eta, \kappa] P[x_{R^*}] P[\eta] P[\kappa] dx_{R^*} d\eta d\kappa. \quad (4.10)$$

The integrals (4.10) can be evaluated numerically given the prior distributions for x_{R^*} , T_a , η and κ which must be specified by the expert (Section 4.3.1).

The posterior distribution of T_a is often not that sensitive to the precise form of the curve (4.5) because most “reasonable” shapes of curve will include the data which suggest similar x_{RT} values, while eliminating x_{RT} values that are much smaller. However, more experimental investigations into the form of (4.5) would be valuable.

If the posterior distribution of T_a does not give a clear indication of the plateau starting temperature then different possible sets of aliquots should be selected in

order to evaluate the sample palaeodose, according to the posterior distribution of T_a . The distribution of T_a will give the relative probabilities of each of the posterior palaeodose distributions found using the combined aliquot model (illustrated in the example in Section 4.5).

The selection of the plateau starting temperature, and thus the aliquots chosen to evaluate the sample palaeodose, can be considered as a bias-variance trade off. If a high plateau starting temperature is used, then fewer aliquots and so less data are used to evaluate the sample palaeodose, and so the variance of the palaeodose is increased. However, a lower choice of T_a will result in more data contributing to the sample palaeodose evaluation, and thus reduce the variance but at the cost of increasing the potential bias in the evaluation.

4.2.5 Prior Distributions

The prior probability distributions for x_{R^*} , T_a , and the curve parameters η , κ are to be specified by the dating expert. The level of the plateau, x_{R^*} is the palaeodose of the sample, so it is usually appropriate to use the same prior distribution as for the palaeodose in the combined aliquot model (Section 3.5.1). The practitioner tries to choose preheat temperatures for which the palaeodose estimates will lie on the preheat plateau, and hence evaluate the sample palaeodose. So, a priori the practitioner judges all aliquots to be on the preheat plateau, and their estimates of palaeodose will correspond to the level of the plateau.

There is no purpose in routine dating for palaeodose estimates to be made in regions where the practitioner thinks the preheat temperature is too low for the plateau to have begun. Therefore the prior distribution for T_a , the temperature at which the plateau begins, is likely to lie in the lower end of the preheat temperatures applied. As with everything in luminescence dating, the behaviour of the luminescence can be unpredictable and strongly sample specific, so a reasonable degree of uncertainty is likely to be reflected in the prior distribution for T_a .

As illustrated above, the nature of the relationship between palaeodose and preheat temperature before the plateau is reached is variable amongst samples. The model has been set up in such a way that a wide variety of shapes are viable in this

region. Therefore, the prior distributions on the curve parameters η, κ should allow many reasonable curves to be adopted, so the uncertainty for these parameters will be substantial. It should also be noted that the data are likely to be relatively sparse before the plateau begins, so general priors for η, κ will prevent any available data being overpowered by somewhat arbitrary prior assumptions.

4.3 Example

The example carries on the analysis of sample 311-6 of Fydell House from Chapter 3. The data shown in Table 4.1, are the mean and standard deviations from the posterior palaeodose distributions at each preheat temperature, from the combined aliquot model. This is the same as Table 3.10, presented at the end of Chapter 3. The posterior palaeodose mean values are plotted against preheat temperature in Figure 4.4.

Preheat Temperature °C	# Aliquots	x_{RTi}	σ_{RTi}
210	3	1021.5	28.5
220	6	982.5	25.0
230	3	949.4	28.4
240	5	993.7	42.4
250	3	979.12	28.1

Table 4.1: Mean and standard deviations of the posterior palaeodose at each preheat temperature for sample 311-6; the data input in the plateau model.

4.3.1 Prior Specification

In this example, the preheat temperature range of 210-250°C was used as the practitioner, Bailiff, believed that the palaeodose evaluations made in this region were likely to form a preheat plateau (otherwise there would have been no purpose to evaluating the palaeodose at those temperatures). The prior distribution for the plateau starting temperature, T_a , was set to be normal with mean 215°C and standard deviation 30°C. A probability distribution that is unimodal and symmetrical

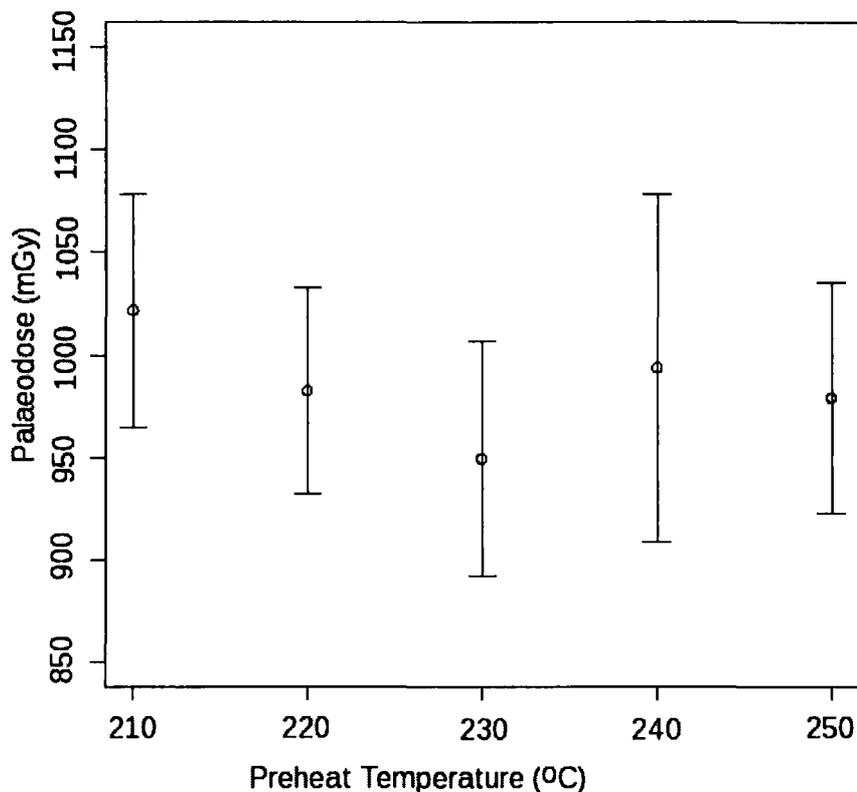


Figure 4.4: Posterior palaeodose mean and two standard deviation uncertainty bars for each group of aliquots at a certain preheat temperature. These values will be used as the data input in the model to find the start of the preheat plateau

was desired, and so the Gaussian distribution was selected. The standard deviation of 30°C reflects the uncertain nature of luminescence; its behaviour can be problem specific and so we do not have extremely strong beliefs about the starting point of the plateau.

The prior distribution of the level of the plateau, the palaeodose of the sample x_{R^*} , was also assigned a normal distribution. Its hyperparameters took the same values as the prior distribution for palaeodose in the combined aliquot model (mean 1000 mGy, standard deviation 100 mGy).

Very little is known about the curve parameters η , κ , but the prior distributions were based on likely plateau shapes (Figure 4.3). The parameters were both assigned the same prior, a normal distribution with mean 0.003 and variance 1.

The prior distributions used in this example are summarised in Table 4.2.

Parameter	T_a	x_{R^*}	η	κ
Prior	$N(215, 30^2)$	$N(1000, 100^2)$	$N(0.003, 1^2)$	$N(0.003, 1^2)$

Table 4.2: Prior distributions used for Sample 311-6, FydeH House, in the plateau model.

4.3.2 Posterior Distribution

The posterior distribution for T_a was found using (4.10) and shown in Figure 4.5. It was computed using Maple, as it is efficient in numerical integration, and the code used is shown in Appendix H.2. The integral was calculated between temperature 100-250°C, at 0.01 intervals. This was then multiplied by a suitable constant so that the total probability summed to unity (as (4.10) is based on proportionality). The region over which the function was integrated was chosen on a trial and error basis; the whole distribution needs to be covered in order to compute the posterior probabilities.

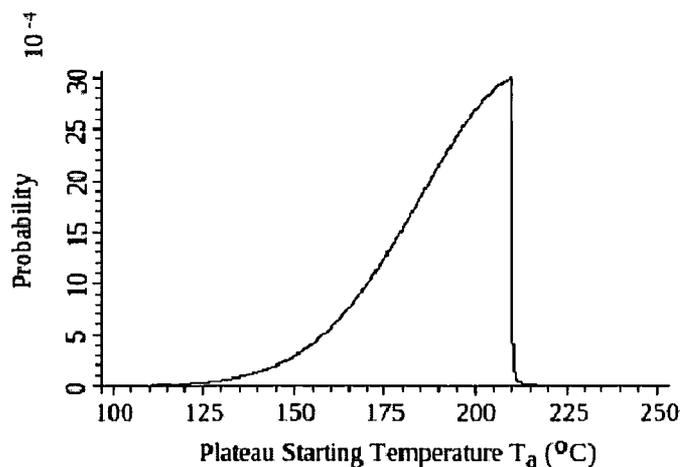


Figure 4.5: Posterior distribution of plateau starting temperature T_a for sample 311-6.

The focus of this analysis is to choose which aliquots lie on the preheat plateau. Therefore the interest lies in the probability that the plateau starts between 210 and

220°C as this would indicate whether or not to include aliquots with a preheat of 210°C in the evaluation of the sample palaeodose.

The probability that the plateau starts between 210 and 220°C is 0.007, and the model assigns no probability to the plateau starting above 220°C. This indicates that it is likely that the plateau started before 210°C and therefore the palaeodose estimates from all the aliquots of 311-6 lie on the preheat plateau.

4.3.3 Influence of the Prior Parameters

In this section, the influence that the prior distributions assigned to the parameters have on the posterior distribution of T_a is investigated. We first look at the choice of prior for T_a .

Table 4.3 shows how the choice of prior distribution for T_a influences the posterior probability that the plateau starts after 210°C. The posterior distributions achieved for T_a under these priors are shown in Appendix F. Here the focus is on which of the aliquots should be used to evaluate the sample palaeodose, i.e. do all the data points lie on a preheat plateau. The lowest preheat temperature used for aliquots from sample 311-6 was 210°C, so $P[T_a > 210|data]$ is quoted in Table 4.3 for different prior judgements made.

Prior Mean	Prior SD	$P[T_a > 210 data]$
200	10	0.0053
200	30	0.0045
200	50	0.0032
215	10	0.0253
215	30	0.0075
215	50	0.0042
230	10	0.1030
230	30	0.0107
230	50	0.0053

Table 4.3: Posterior probability that the preheat plateau starts at a temperature higher than 210°C, for varying prior specifications of T_a .

The prior distribution used in the analysis for T_a was $N(215, 30^2)$. When the prior standard deviation is decreased to 10°C , the posterior probability that the preheat plateau starts above 210°C is increased to 0.025. However, when the prior standard deviation is set to 50°C , this probability is reduced to 0.0042. The strong prior information is reflected in the posterior distribution; the model gives a higher probability to the plateau starting after 210°C (as the prior mean is set to 215°C) when the prior standard deviation is small.

When the prior mean of T_a is set at 200°C , then the corresponding posterior probabilities for the plateau starting after 210°C is reduced for each of the standard deviations used (10, 30 and 50°C). Similarly, when T_a is assigned a mean of 230°C a priori, these posterior probabilities are increased, as the prior beliefs indicate that the plateau is thought to start above 210°C .

The probability that the preheat plateau starts after the first data point at 210°C is small, for all the prior distributions assigned to T_a in Table 4.3. The data input to the model, plotted in Figure 4.4, are indicative of a plateau, and so suggests that the plateau has started at temperatures below the preheats associated with the data.

Judgements made about the level of the plateau, namely the sample palaeodose x_{R^*} , may also affect the posterior distribution for T_a . Table 4.4 gives the posterior probability that the plateau starts after 210°C when different values are assigned to the prior mean and standard deviation of x_{R^*} .

When the prior mean of x_{R^*} is set to be 1000 mGy, then lowering the prior standard deviation of x_{R^*} causes the posterior probability that $T_a > 210$ to decrease. This pattern of behaviour is also observed when the prior mean of x_{R^*} is assigned different values. When the prior mean of x_{R^*} is 900 mGy, the model assigns more probability to the plateau starting after 210°C than when the mean is 1100 mGy. As the first data point is at 1021.5 mGy (Table 4.1), the model gives more probability to this point being included in the plateau (i.e. $P[T_a > 210|data]$ is reduced) when the prior mean of the plateau level is closer to this point, at 1100mGy.

The prior distributions assigned to the curve parameters κ , η , had little influence on the posterior probability that the plateau starts after 210°C , as long as they allowed for all reasonable curve shapes (e.g. Figure 4.3). The posterior probabilities

Prior Mean	Prior SD	$P[T_a > 210 data]$
1000	200	0.0119
1000	100	0.0075
1000	50	0.0067
1000	10	0.0029
900	100	0.0083
900	10	0.0031
1100	100	0.0067
1100	10	0.0026

Table 4.4: Posterior probability that the preheat plateau starts at a temperature higher than 210°C, for varying prior specifications of x_{R*} .

achieved for different prior mean and variance values of κ , η are given in Appendix F.1.

The choice of values for the parameters a priori influences the posterior distribution of T_a . In this example, the influence is not extensive enough to alter the concluding decision on where the preheat plateau lies, i.e. which aliquots should be used to evaluate the sample palaeodose.

4.4 Sample Palaeodose Evaluation

For a given preheat plateau, we evaluate the palaeodose as in Chapter 3. The posterior distribution for the plateau starting temperature, T_a in Figure 4.5 indicates that it is likely that the preheat plateau begins at a temperature below 210°C, and so all aliquots of sample 311-6 lie on the preheat plateau. This means that all of the aliquots from the sample are estimating the same palaeodose, the palaeodose of the sample. That is,

$$x_{Rj} = x_R + \delta_j^{xR} \quad (4.11)$$

for $j = 1, \dots, N$, where N is the total number of aliquots lying on the preheat plateau (in this case the total number of aliquots in the sample). This is the same relationship used in (3.13) for aliquots at the same preheat temperature.

So, the combined aliquot model, for evaluating the posterior palaeodose distribution based on a number of aliquots (Section 3.2) was used to find the distribution for the palaeodose of the sample. The same prior specification of the parameters in the regression model was made as when the combined aliquot model was applied to aliquots at a single preheat temperature (Section 3.5.1), which are given in Table 3.5.

The Gibbs sampler detailed in Section 3.3 and Appendix H.1 was run with 5 chains over 100,000 iterations. A burn-in of 1000 was used, and the chains were thinned every 10. The resulting posterior distribution for the sample palaeodose is shown in Figure 4.6, which has mean 982.3 mGy and standard deviation 11.3.

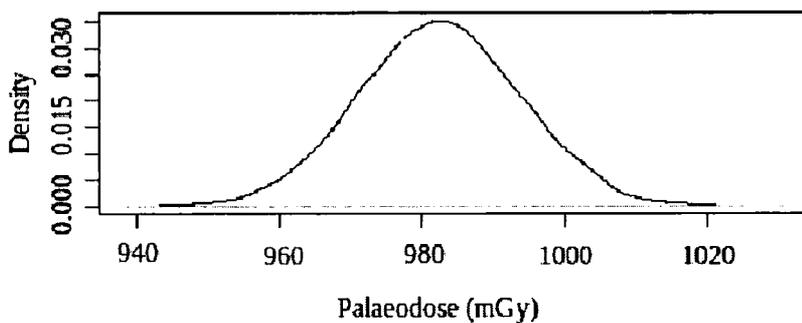


Figure 4.6: Posterior palaeodose distribution for sample 311-6.

4.5 Extension of the Preheat Plateau

It is difficult to develop the model for finding the starting point of the preheat plateau without further empirical or theoretical knowledge. Experimental data covering the low preheat temperature readings are not obtained routinely, as aliquots which are not thought to lie on the preheat plateau are not used to evaluate the sample palaeodose.

To illustrate the plateau model further, additional observations were made by Bailiff in the laboratory with sample 311-6 using lower preheat temperatures. The preheat treatments used were at 140-200°C, at 20°C intervals, and the data are shown in Appendix G.2.

A posterior distribution for palaeodose was computed at each preheat temperature, using the combined aliquot model in Section 3.2. The same prior distributions were used as those used to compute the palaeodose distributions for the original data set. Again, the sampler given in Appendix H.1 was run for 100,000 iterations, with a burn-in of 1000 and the chains were thinned every 5. The mean and standard deviation of the palaeodose distributions for each preheat temperature, including those previously computed, are shown in Table 4.5, and plotted in Figure 4.7.

Preheat (°C)	# aliquots	Posterior Mean	Posterior SD
140	2	873.5	66.9
160	2	924.1	35.9
180	2	928.57	37.0
200	2	964.5	41.1
210	3	1021.5	28.5
220	6	982.5	25.0
230	3	949.4	28.4
240	5	993.7	42.4
250	3	979.12	28.1

Table 4.5: Posterior palaeodose mean and standard deviation for aliquots at each preheat temperature, for the extended data set from sample 311-6.

It should be noted that the data from the aliquots with the lower preheat temperatures were produced at a different time to the original data set, and fewer aliquots were analysed, which could potentially be influential.

From Figure 4.7, the location of the preheat plateau is not as clear as when the original data were presented in Figure 4.4. The plateau model was applied using the same prior distributions as previously, which are given in Table 4.2. The posterior distribution found, using the code in Appendix H.2, for T_a is shown in Figure 4.8.

The model analysis indicates that the plateau starts in the ranges 180-200°C with probability 0.47 and 200-210°C with probability 0.50, with the remaining probability for T_a above 210°C. This raises the question as to whether or not to include the data with a preheat of 200°C.

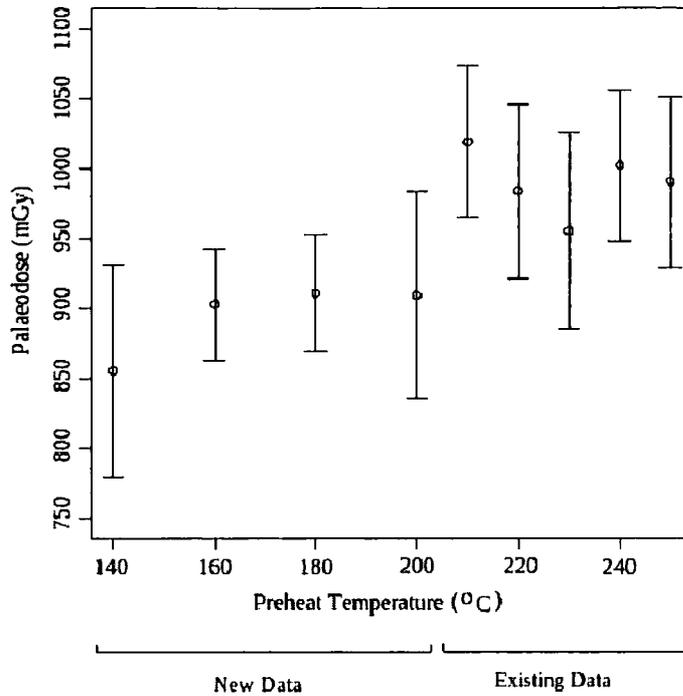


Figure 4.7: Mean and 2 standard deviation bars of posterior palaeodose distributions obtained at each preheat temperature over the extended range.

4.5.1 Influence of Prior Parameters

The influence that the prior distribution assigned to T_a has on the posterior distribution for the start of the preheat plateau is examined for the extended data set. Table 4.6 shows the posterior probabilities of the plateau starting in particular regions for different prior distributions assigned to T_a .

When the standard deviation of T_a is small, a larger probability is assigned to the plateau starting above 210°C, the region where the prior mean for T_a lies. If the prior mean for T_a is changed, then this results in the amount of posterior probability assigned to each region shifting accordingly. For all the prior distributions for T_a in Table 4.6, the majority of the posterior probability distribution falls in the region 180-210°C.

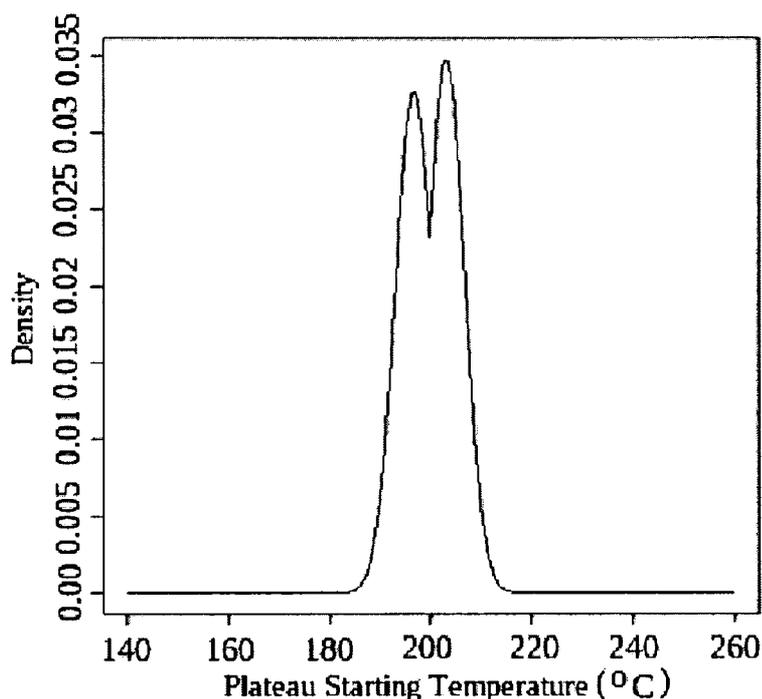


Figure 4.8: Posterior distribution for plateau starting temperature over the extended data set.

4.5.2 Sample Palaeodose Evaluation

The palaeodose has already been evaluated based on aliquots with a preheat of 210°C and above (the original data) in Section 4.4. So, the combined aliquot model was applied to all aliquots which had a preheat treatment of 200°C and above, with the same prior distributions as previously (Table 3.5). The Gibbs sampler was run for 100,000 iterations with 5 chains, a burn-in of 1000 was used and the chains were thinned to every 5th iteration. The posterior distribution for the palaeodose based on the 22 aliquots with preheats 200°C and above is shown in Figure 4.9. This distribution has mean 976.9 mGy and standard deviation 10.4.

The palaeodose distribution for the sample could be presented as a mixture of two normal distributions, with means and standard deviations from the two palaeodose evaluations given above, and scaled by the probabilities from the plateau model. That is, a mixture of a $N(976.9, 10.4^2)$ distribution with weighting 0.48 (plateau starts between 180 and 200°C) and $N(982.3, 11.3^2)$ with weight 0.52 (plateau starts

Prior	Posterior				
	T_a	$P[T_a < 180]$	$P[180 \leq T_a < 200]$	$P[200 \leq T_a < 210]$	$P[T_a \geq 210]$
$N(215, 30^2)$	0	0.47	0.50	0.03	
$N(215, 5^2)$	0	0.45	0.51	0.04	
$N(215, 50^2)$	0	0.47	0.50	0.03	
$N(195, 30^2)$	0.02	0.47	0.48	0.01	
$N(225, 30^2)$	0	0.44	0.51	0.05	

Table 4.6: Influence of the prior distribution for T_a on its posterior for the extended data set of sample 311-6.

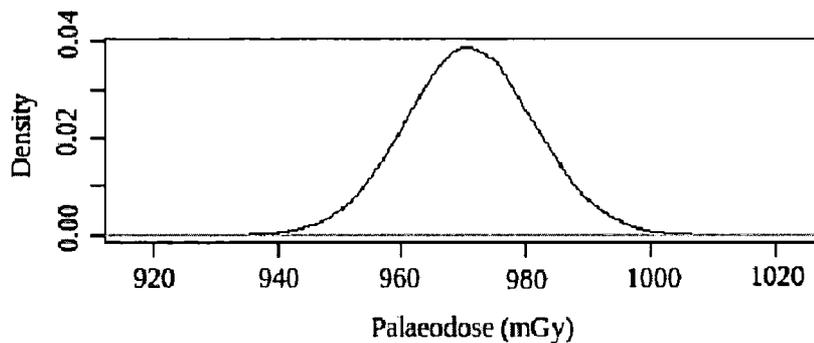


Figure 4.9: Posterior distribution for palaeodose based the aliquots of 311-6 with preheat of 200°C and above.

between 200 and 210°C). The weights were calculated by, for example, taking the posterior probability of the plateau starts between 180 – 200°C (0.47), and dividing by the probability that the plateau starts between 180 and 210°C (0.97). This mixture distribution is plotted in Figure 4.10, along with the normal distribution with the same mean and variance (mean 979, standard deviation 12.3).

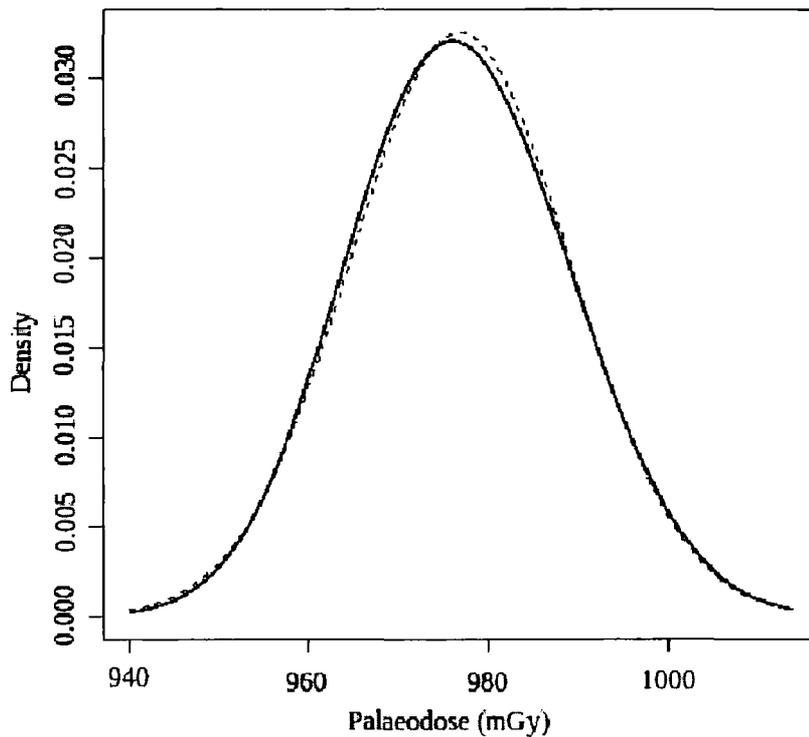


Figure 4.10: Normal mixture distribution of posterior palaeodose based on the weighted mixture of the two possible plateau locations, with the dashed line giving the normal distribution with the same mean and variance.

4.6 Summary Guide to evaluating sample palaeodose

1. Plot posterior mean and standard deviation of palaeodose against preheat temperature, to be used as the data input in the plateau model.
2. Elicit prior distributions for plateau starting temperature, and plateau level, curve parameters.
3. Compute the posterior distribution for plateau starting temperature, to find the probability that each of the data points lie on the plateau.
4. Use the combined aliquot model (Section 3.2) with the aliquots which lie on the preheat plateau to find the posterior palaeodose distribution of the sample.

5. If the plateau model results in a number of different viable locations for the preheat plateau, calculate the sample palaeodose distribution for each one and combine using a mixture of normals.
6. Sensitivity analysis of prior judgements.

Chapter 5

Annual Dose

The annual dose, or dose rate, is the estimate of the average annual radiation dose that the crystal grains have received since resetting (Section 2.4), and is the denominator of the age equation (1.1). In this chapter we consider the current methods for evaluating dose rate, adapt this into a probabilistic model and then explore the model, continuing with sample 311-6 from Fydele House, Lincolnshire, as an example.

5.1 Dose Rate Equation

The dose rate is assessed by summing the component radiation parts: alpha, beta, gamma and cosmic radiation. Here we have assumed that the grains have no internal radionuclide sources, and coarse grains have been used, where the surface has been etched in hydrofluoric acid so the alpha contribution does not need to be considered (Section 2.4).

The standard model for dose rate [3] used in luminescence dating expresses the annual dose, \dot{D} , as:

$$\dot{D} = \frac{b}{1 + H_\beta WF} \dot{D}_\beta + \frac{g}{1 + H_\gamma WF} \dot{D}_\gamma + \dot{D}_c \quad (5.1)$$

where \dot{D}_β , \dot{D}_γ are the respective beta and gamma radiation dose rates, b , g are attenuation factors and \dot{D}_c is the cosmic radiation dose rate. H_β , H_γ represent the absorption of the radiation type by water, W is the saturation water uptake and F is the fractional average water content over the burial period. The measured values

are \dot{D}_β , \dot{D}_γ , \dot{D}_c , W , with the coefficients b , g , H_β , H_γ having standard values.

The most difficult parameter to assess here is F , the time averaged fractional moisture content of the sampling environment. Here we consider dating bricks, that are taken from elevated and dry contexts. A beta distribution is assigned to F , which lies in the range $[0, 1]$, where 0 indicates complete dryness, 1 total saturation. The other parameters are considered to have independent Gaussian forms.

The parameters H_β , H_γ , b , g (the correction factors for absorption of the radiation components in water and the material), are standard values and not sample specific. These are given normal distributions, centred at the current values used across the dating community [1], with standard deviations reflecting the limit of the precision in assessing their value.

The independent distributions for \dot{D}_β , \dot{D}_γ , \dot{D}_c , W are considered to be normal, around the experimental data. The variance is dependent on how accurately the values are thought to reflect these components of the dose rate. After discussions with Bailiff, in an homogeneous material the errors in these parameters are likely to be around 5%.

The distribution for dose rate \dot{D} is found by simulating from the distributions assigned to the parameters and using (5.1), the code for this is given in Appendix H.3. The dose rate is an average measure over the lifetime of the sample. The measurements taken are contemporary, and assumptions are made by the experimenter concerning the extent to which the dose rate has varied since the luminescence clock was reset.

Hence we take \dot{D} to be the estimate for the average annual radiation dose received, based on contemporary measurements and the uncertainty associated with relating this quantity to the actual dose rate is incorporated into the model for sample age (Chapter 6).

5.2 Example

Sample 311-6 from Fydell House is again used here to illustrate the computation of the dose rate distribution. The data values for the dose rate are given in Table 5.1,

and where produced by Bailiff [14]. The standard parameter values for both the water and surrounding material are shown in Table 5.2.

\dot{D}_β	\dot{D}_γ	\dot{D}_c	W
2.21	1.30	0.2	0.033

Table 5.1: Measurements used to compute the dose rate for sample 311-6, Fydell House.

\dot{D}_β , \dot{D}_γ , \dot{D}_c are the measured contemporary rates of beta, gamma and cosmic radiation and W is the water content at saturation.

b	g	H_β	H_γ
0.92	0.93	1.14	1.25

Table 5.2: Standard attenuation parameter values [1] used in the dose rate calculation.

5.2.1 Choice of Parameter Values

The measured β , γ and cosmic dose rates, along with the saturated water content W , were assigned normal distributions, centred on the experimental values given in Table 5.1, with standard deviations 2.5% of these (i.e. 5% error represents two standard deviations). The value of uncertainty in the measurements was specified by the practitioner, Bailiff, based on past dating experience. Similarly, the attenuation and water correction coefficients were given Gaussian distributions, with mean the accepted standard value (Table 5.2) and a small standard deviation to reflect the belief's of the expert on the accuracy of these values. F , the average fractional water saturation was judged to follow a beta distribution with mean 0.15 and standard deviation 0.2. This reflects the fact that the sample originated from the a brick wall, where the saturation level will have been low and relatively stable since the brick was fired. These distributions are summarised in Table 5.3.

To compute the distribution for the dose rate of sample 311-6, 100,000 draws were made from each of the distributions in Table 5.3 using the code in Appendix

Parameter	Distribution
\dot{D}_β	$N(2.21, (0.025 \times 2.21)^2)$
\dot{D}_γ	$N(1.30, (0.025 \times 1.30)^2)$
\dot{D}_c	$N(0.2, (0.025 \times 0.2)^2)$
W	$N(0.033, (0.025 \times 0.033)^2)$
b	$N(0.92, 0.05^2)$
g	$N(0.93, 0.1^2)$
H_β	$N(1.25, 0.1^2)$
H_γ	$N(1.14, 0.1^2)$
F	$\beta(3.0375, 17, 2125)$

Table 5.3: Distributions assigned to the parameters in the model for dose rate for sample 311-6 of Fydell House.

H.3. This results in a distribution with mean 3.42, standard deviation 0.18 and is shown in Figure 5.1.

5.3 Influence of parameters

The influence that judgements about the model parameters have on the dose rate is now investigated. The beta, gamma and cosmic components of the dose rate, \dot{D}_β , \dot{D}_γ , \dot{D}_c are assigned normal distributions of the form $N(m, (mp)^2)$ where m is the measured value and p a proportion to be specified by the expert. As shown in Table 5.3, in this example the standard deviation was set to 2.5% of the mean (i.e. $p=0.025$) for each of \dot{D}_β , \dot{D}_γ , \dot{D}_c . This value was chosen as Bailiff considered the errors in these measurements to be around 5%.

Figure 5.2 shows how the posterior mean and standard deviation of dose rate changes with respect to the value placed on the prior standard deviation of \dot{D}_β , \dot{D}_γ , \dot{D}_c . The dose rate model was run for 100,000 iterations and the value of p was changed in turn to obtain the readings for the dose rate statistics.

The effect the magnitude of the standard deviation of \dot{D}_β , \dot{D}_γ , \dot{D}_c on the mean of the dose rate is minimal; in Figure 5.2 any changes in the mean are in the third

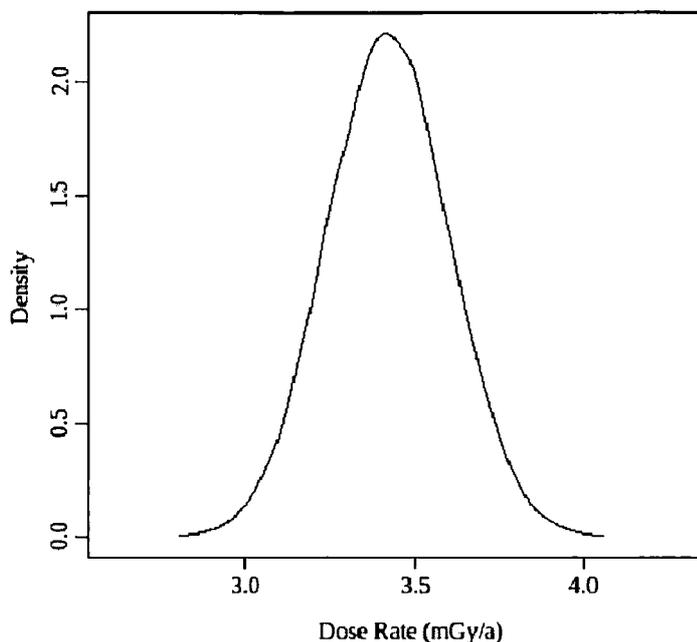


Figure 5.1: Posterior dose rate distribution for sample 311-6 from Fydell House.

decimal place for each of the three parameters. It could be possible to observe some trend in the dose rate mean, for example for the higher values of \dot{D}_β standard deviation, the mean of the dose rate looks to be falling, but these changes are small.

There is a clear trend in Figure 5.2 in the standard deviation of dose rate, when compared against the standard deviation of each of \dot{D}_β , \dot{D}_γ , \dot{D}_c . The standard deviation of the dose rate increases monotonically with the standard deviation of \dot{D}_β and \dot{D}_γ . As the dose rate equation (5.1) can be considered the sum of the beta, gamma and cosmic dose components, with correction coefficients, then this behaviour of the variance was anticipated. The standard deviation of \dot{D}_β causes greater change in the dose rate standard deviation than that of \dot{D}_γ . This is because the standard deviations of \dot{D}_β , \dot{D}_γ have been considered as a percentage of their means, and in this example the mean of \dot{D}_β is larger than that of \dot{D}_γ (2.21 compared to 1.30). The standard deviations of the dose rate components have been expressed in this way, as the mean values are experimental measurements for which the practitioner will often find it natural to specify the uncertainty in these measurements from previous dating experience.

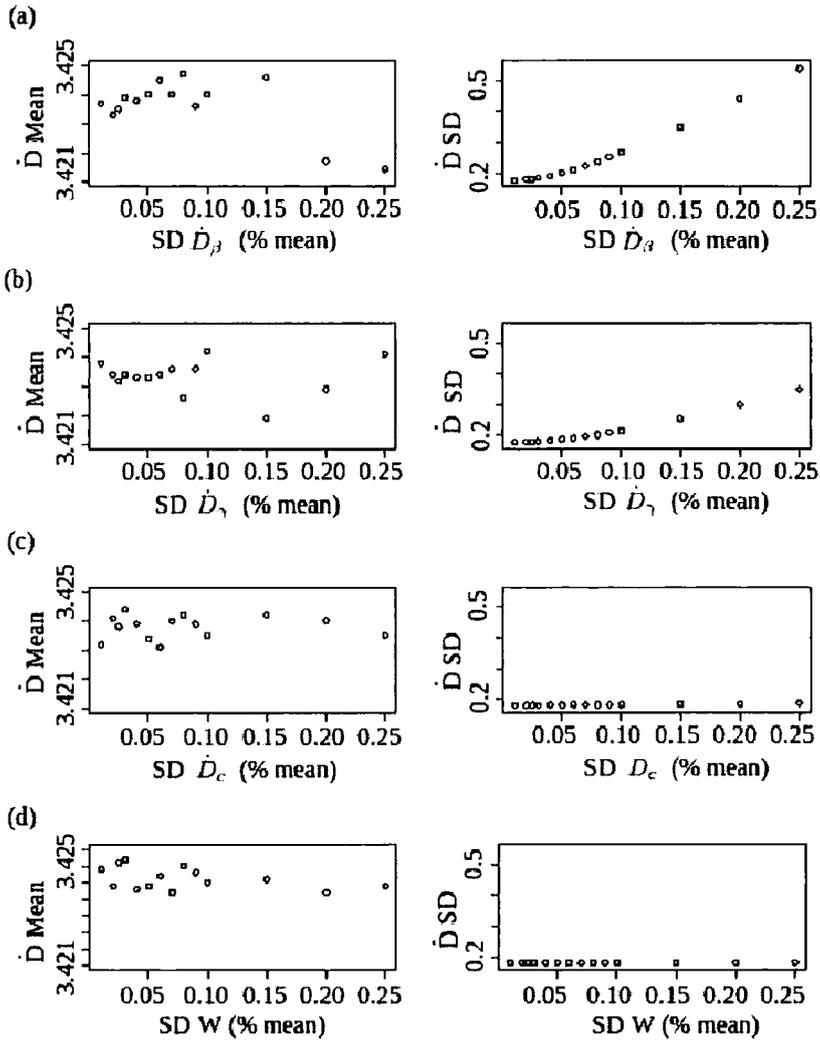


Figure 5.2: Mean and standard deviation of the posterior dose rate against the prior standard deviation of (a) \dot{D}_β , (b) \dot{D}_γ , (c) \dot{D}_c , (d) W .

The prior distributions of \dot{D}_β , \dot{D}_γ , \dot{D}_c , W , are considered normal, expressed as $N(m, (mp^2))$ where m is the measured value and p a proportion to be specified. Here we consider the influence of the choice of p on the posterior dose rate.

The standard deviation of \dot{D}_c , the cosmic ray contribution to the dose rate, has very little effect on the dose rate standard deviation, which remains at a constant level in Figure 5.2 (c). There is a very small increasing trend in the standard deviation values, which reflects the magnitude of the cosmic ray component of the annual dose compared to that of the beta and gamma radiation elements (\dot{D}_β , \dot{D}_γ).

Similarly, the standard deviation of W , the saturation water uptake of the sample, has no notable effect on the standard deviation of \dot{D} . No trend is observed in the mean of the dose rate as the standard deviation of W is increased, and on a different scale it would be possible to notice a slight upward trend in the standard deviation of \dot{D} . Again, the mean of W , its experimental value, is relatively small, so an increase in its uncertainty as a percentage of the mean is not likely to have a large effect on the overall dose rate.

The correction coefficients b , g , H_β , H_γ were next considered in this analysis. These parameters take community wide accepted values (Section 5.2.1) as the mean of their Gaussian distributions, with variance specified by the expert. Therefore, here we will only consider the influence that the judgements made about their uncertainty has on the dose rate distribution.

Figure 5.3 shows how the choice of prior standard deviation of b , g , H_β , H_γ influences the posterior dose rate standard deviation. The level of uncertainty in the attenuation correction parameters b , g has a notable influence on the standard deviation of \dot{D} , with the standard deviation of \dot{D} increasing with that of either b or g . The change in the standard deviation of \dot{D} is more marked when the standard deviation of b is increased, compared to the influence of the standard deviation of g . This is because the magnitude of the beta component of the dose rate, to which b is part of the coefficient, is larger than the gamma part (which g is associated with). So any change in b will have greater weighting in the overall dose rate, \dot{D} .

The uncertainty of the parameters H_β , H_γ are not influential on the standard deviation of the dose rate, for the range of standard deviations of H_β , H_γ investigated, and this is shown in Figure 5.3 (c), (d). The value of H_β is small, and coupled with this, H_β is just one of three parameters in the denominator, and so any influence it may have will be diminished further (similarly for H_γ).

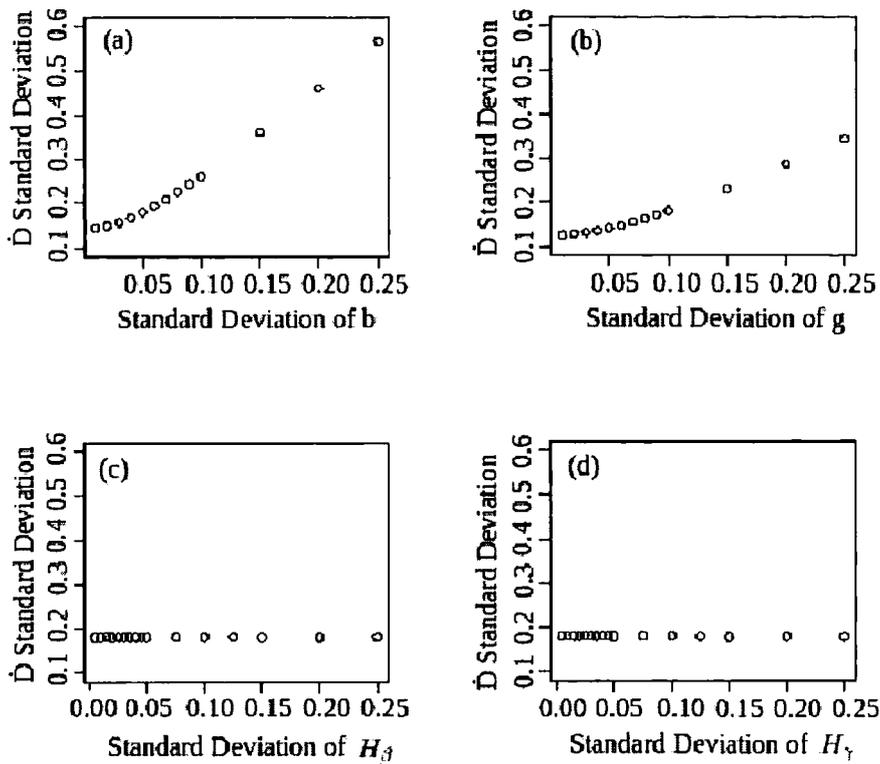


Figure 5.3: Influence of the standard deviation of parameters (a) b , (b) g , (c) H_β , (d) H_γ , on the standard deviation of dose rate, \dot{D} .

The parameter F , the average level of saturation of the sample, is the one about which least information is known. Here sample 311-6 is from a brick building, and so considered to be relatively dry (so F has a small value). However, to look at the influence of the judgements made about F , Figure 5.4 shows how the mean and standard deviation of the dose rate changes with the mean chosen for F .

There is a clear decreasing linear relationship between the mean of \dot{D} and the value given to the mean of F . A large value of F is associated with the sample containing more water (a value of 1 means the material is saturated), and as water absorbs the radiation, the quartz grains will be exposed to a lower dose of radiation, and thus the dose rate is reduced. The standard deviation of \dot{D} also shows a decreasing trend with the F mean in Figure 5.4, though there is more scatter.

Here we have considered the role each of the parameters takes in the posterior distribution for dose rate. In doing so, we have looked at extreme values of each of the

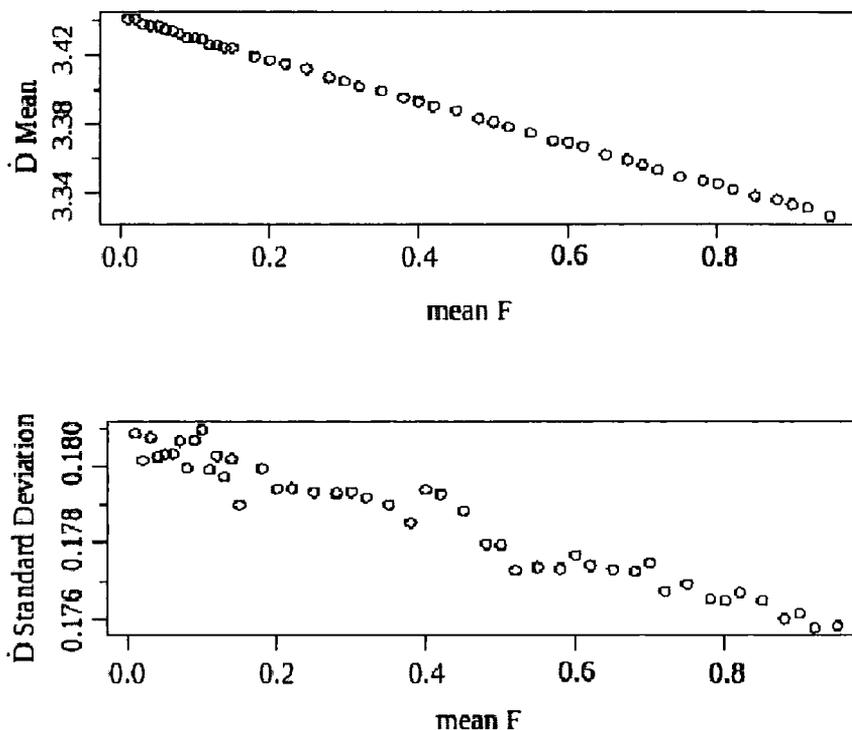


Figure 5.4: Influence of the mean of F on the mean and standard deviation of \dot{D} .

parameters, which would not realistically be specified here. The main information that is required from the expert practitioner is the error in the measured values, which here we have considered as a proportion of their mean. It is important to carry out a sensitivity analysis on this parameter to ensure that the judgements of the expert are correctly represented in the analysis.

5.4 Water Content Variations

The example which has been considered throughout is a sample taken from a brick building. One of the reasons this example was chosen is that such a sample is known to have been relatively dry throughout the dating timescale, thus eliminating a potential source of uncertainty. However, in this section we look briefly at possible variations of F , and how using a single value may induce error in the dose rate evaluation. The water saturation fraction of the dating environment will generally

have seasonal variations, so here we will look at the yearly averages.

First, consider a sample that has been in a relatively dry environment, except for a period when it has had a higher level of saturation. For example, the fractional saturation may change with time as in Figure 5.5, where $F = 0.15$ or $F = 0.7$.

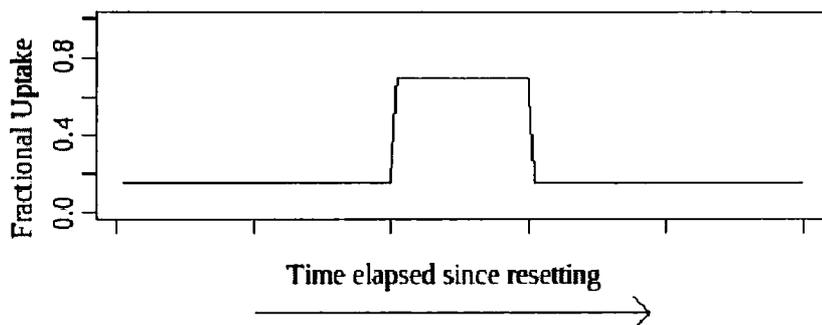


Figure 5.5: Possible variation of the fractional water saturation of the dating environment since the resetting of the luminescence clock.

Let q be the proportion of time that the fractional water saturation level was at 0.15, and calculate the overall dose rate using this information. This is then compared to the dose rate evaluation achieved if the practitioner were not aware that the fractional water uptake of the sample had been elevated (i.e. $q = 1$). For this hypothetical example, the remaining parameters were assigned the same values as those for the example above. Figure 5.6 shows how the evaluation of dose rate changes with the proportion of F computed with the mean of 0.15 (with the remainder at 0.7). The mean and one standard deviation error bars are shown, along with the mean and standard deviation of the dose rate when it is assumed that F has mean 0.15 for the whole time.

A larger value assigned to the mean of the distribution for F results in a lower dose rate (Figure 5.4), and Figure 5.6 shows that as the proportion of time spent with the mean F at 0.70 rather than 0.15, the evaluation of dose rate falls below that when the mean of F is always 0.15. Although error would be induced in the dose rate if the practitioner considered the mean of F to be 0.15, as they were unaware that the saturation level had been raised for some period, the magnitude of error is not of great concern. For example, even when the dose rate was evaluated with the

mean of F at 0.7 for half the time, the difference between the two means of \dot{D} is 2.4%.

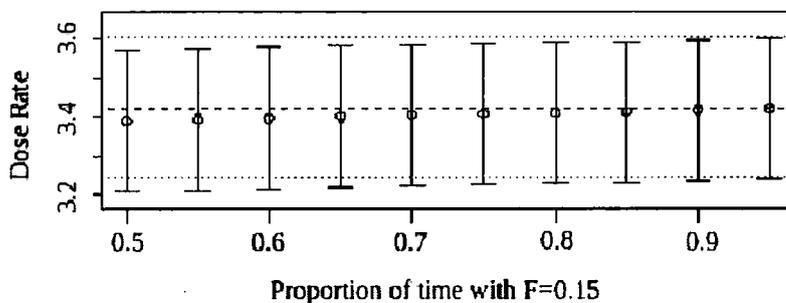


Figure 5.6: Mean dose rate with one standard deviation uncertainty bars calculated with mean of F as 0.15 for different proportions of the dating lifetime (0.7 for the remainder of the time). The dashed lines represent the mean and one standard deviation of the dose rate evaluated with mean F at 0.15 at all times.

Now consider what would happen if the water saturation fraction of the sampling environment followed a sinusoidal pattern, as illustrated in Figure 5.7. This shows the saturation varying between 0.25 and 0.75, over a period of 40 years. To compare the dose rate evaluated with F as in Figure 5.7 with a standard calculation of \dot{D} , the dose rate distribution was computed with F being assigned a mean of 0.5 (the midpoint of the cycle). Then, the mean of F was set to follow the function shown in Figure 5.7, with the remaining parameters taking the same values as above (Table 5.3). This resulted in a dose rate distribution with mean 3.380, standard deviation 0.179 which can be compared to the dose rate when mean F was constant, which has mean 3.381, standard deviation 0.178.

There is no substantial difference between the two dose rate distributions, and so there is no motivation for using an uncertain function or form for the mean of F when more basic calculation will suffice in this situation. It is extremely unlikely that a practitioner of luminescence dating would ever have such detailed information about the past water uptake levels of the sampling environment. However, it is reassuring to see that using this model for dose rate, the choice of F does not dominate the evaluation of dose rate, and thus in most dating situations it is appropriate to use



a single value for its mean.

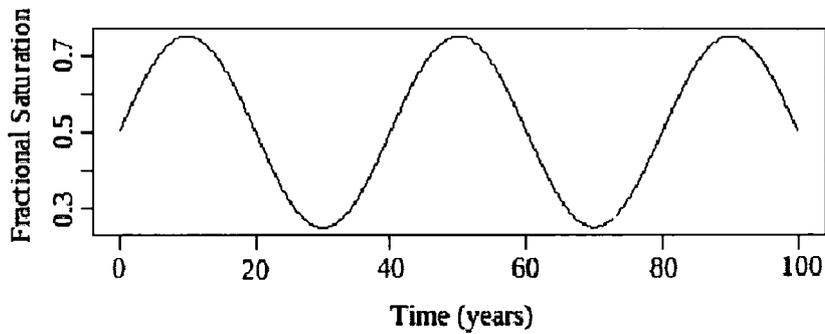


Figure 5.7: Hypothesised Variation in the water saturation levels of the sampling environment.

The uncertainty in F only has a small impact on uncertainty on \dot{D} because here the value of W is very small. When dating material with a higher water content at saturation (for example some types of sediment), the same calculation will show a higher sensitivity to variation in F , and so motivate the need to consider such uncertainty more carefully.

Chapter 6

Age Evaluation

In this chapter we consider the age ratio, the ratio of palaeodose to annual dose, which is the culmination of routine luminescence dating analysis. We then go on to consider the relationship between the age ratio and the sample age, continuing with the example of sample 311-6 from Fydell House, Lincolnshire. The date achieved is then compared with that found using a conventional analysis.

The true sample age is the number of years that have passed since the last resetting event, either by heat (bricks) or light. The sample age is estimated by the age ratio, the quotient of palaeodose and annual dose. This is evaluated using the palaeodose and annual dose distributions found previously, and this estimate of the age ratio is referred to as the ratio estimate.

6.1 Age Ratio

The age ratio A_E is given by the equation

$$A_E = \frac{\textit{Palaeodose}}{\textit{Annual Dose}} \quad (6.1)$$

and is estimated using the distributions for sample palaeodose (Section 4.4) and annual dose (Chapter 5) found previously. Draws are taken from the posterior distribution for sample palaeodose, and combined with values drawn from the annual dose distribution using (6.1) to simulate the distribution for the age ratio, A_E . This is the final step of routine dating, and the age ratio is used to evaluate the age

of the sample. A simulation approach is adopted here to allow for all forms of the posterior distributions of palaeodose and annual dose. Considering the ratio estimate as a ratio of two normals is investigated via the example in Section 6.1.1

From the posterior distribution for A_E we evaluate the posterior mean, \bar{A}_E , and the posterior variance ω_E^2 . \bar{A}_E can be considered a ratio estimate for A_E , and we observe from the examples studied that the form of the posterior distribution is approximately normal, with a unimodal and symmetric shape, or a mixture of normals if the plateau model has posterior uncertainty. Therefore, we model

$$\bar{A}_E = A_E + \delta^E \quad (6.2)$$

with $\delta^E \sim N(0, \omega_E^2)$. This use of the posterior mean and standard deviation as data input in the next stage of the model is analogous to the use of the posterior mean and standard deviation of palaeodose at each preheat temperature as data input to the plateau model (Section 4.2.2). Thus the justification for the step follows the same line of reasoning.

As there are direct prior judgements about palaeodose, the preheat plateau, and the annual dose, it would be very difficult to combine the age analysis into a single calculation. Considering the age analysis of a series of calculation steps allows the uncertainty at each stage to be considered carefully and be fully expressed. This transparent approach to the age analysis is user-friendly and easily adaptable to developments in both theoretical understanding and experimental methods in luminescence dating.

6.1.1 Example

The posterior distribution for the age ratio, A_E , was calculated for sample 311-6 of Fydell House, Lincolnshire, based on the posterior palaeodose distribution found in Section 4.4 and the annual dose distribution in Section 5.2. The distribution for \bar{A}_E is shown in Figure 6.1, which has mean 286.9 and standard deviation 15.2 years. Also shown is the normal distribution with the same mean and standard deviation, the distribution of $A_E|\bar{A}_E$. That is, \bar{A}_E is equivalent to making an observation of A_E with error variance ω^2 .

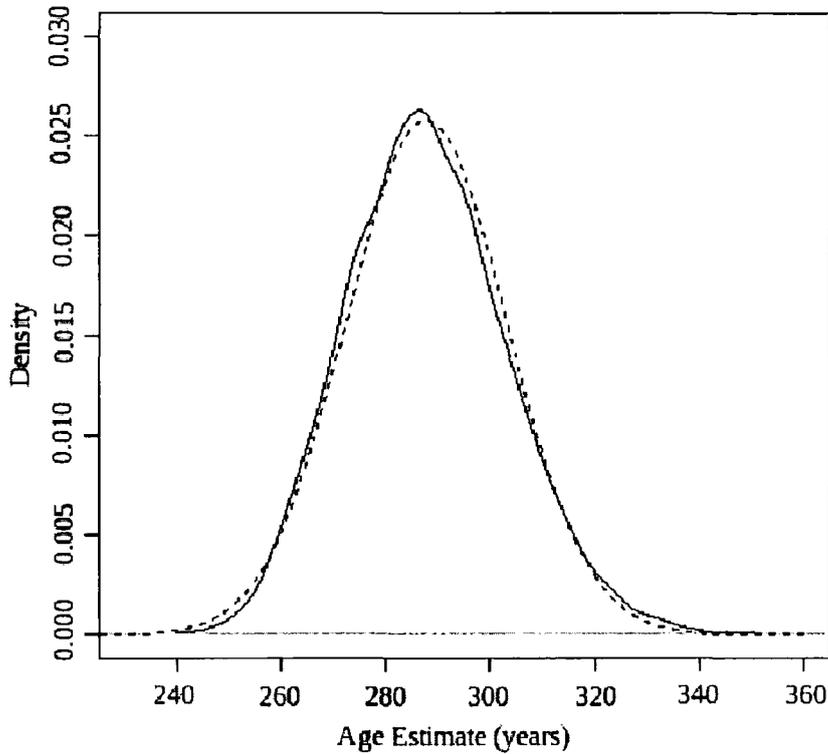


Figure 6.1: Distribution of \bar{A}_E for sample 311-6. The dashed line gives the distribution of $A_E|\bar{A}_E$.

Age ratio as a Ratio of Two Independent Normal Distributions

The posterior palaeodose distribution for sample 311-6, and the dose rate distribution found both look approximately normal. This can be observed in Figure 6.2 where normal distributions are overlaid on the palaeodose and dose rate posterior distributions, previously found in Chapters 3 and 5 respectively. We can then consider the age ratio as approximately the ratio of two normal distributions, and as assumed throughout, these two distributions are considered independent.

The distribution of the ratio of two normal distributions is a problem that has been well documented (e.g. [70]). The density of the ratio two independent normal distributions, $X \sim N(\mu_x, \sigma_x^2)$, $Y \sim N(\mu_y, \sigma_y^2)$, with $Z = X/Y$ can be expressed by [53]

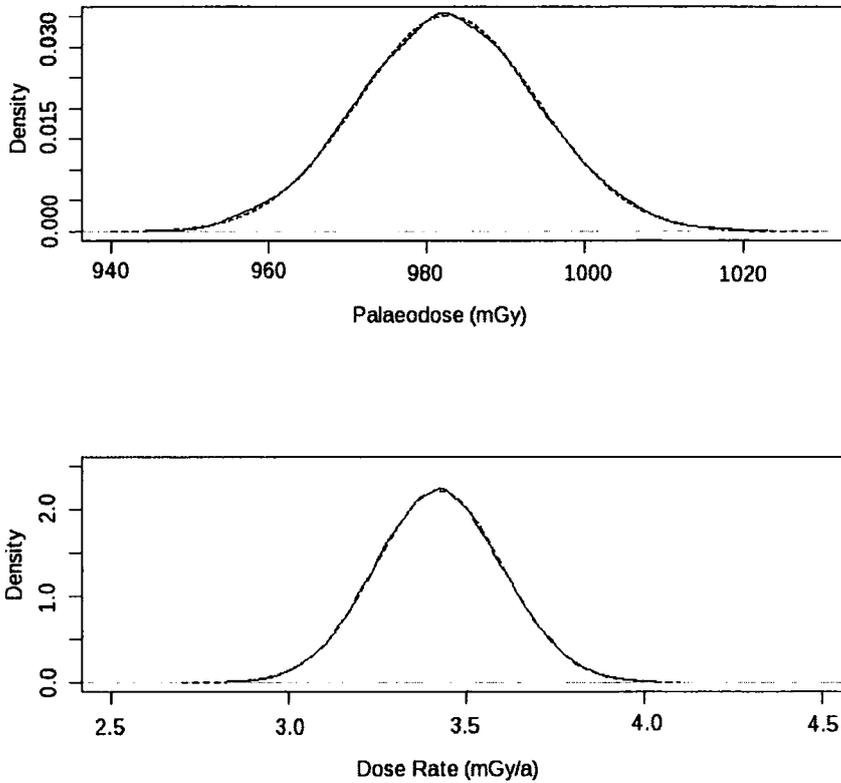


Figure 6.2: Posterior distributions for (a) Palaeodose and (b) Annual Dose for 311-6, as found in Chapters 3 and 5, overlaid with normal distributions with the same mean and variance for comparison.

$$f_Z(z) = \frac{b(z)c(z)}{a^3(z)} \frac{1}{\sqrt{2\pi}\sigma_x\sigma_y} \left[2\Phi\left(\frac{b(z)}{a(z)}\right) - 1 \right] + \frac{1}{a^2(z)\pi\sigma_x\sigma_y} \exp\left\{-\frac{1}{2}\left(\frac{\mu_x^2}{\sigma_x^2} + \frac{\mu_y^2}{\sigma_y^2}\right)\right\} \quad (6.3)$$

where

$$a(z) = \sqrt{\frac{1}{\sigma_x^2}z^2 + \frac{1}{\sigma_y^2}} \quad (6.4)$$

$$b(z) = \frac{\mu_x}{\sigma_x^2}z + \frac{\mu_y}{\sigma_y^2} \quad (6.5)$$

$$c(z) = \exp \left\{ \frac{1}{2} \frac{b^2(z)}{a^2(z)} - \frac{1}{2} \left(\frac{\mu_x^2}{\sigma_x^2} + \frac{\mu_y^2}{\sigma_y^2} \right) \right\} \quad (6.6)$$

$$\Phi(z) = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{1}{2}u^2 \right\} du. \quad (6.7)$$

Here we are considering the age ratio, the ratio of palaeodose to annual dose for sample 311-6, and we have approximated the distributions for these found in Chapter 3 and 5 respectively as normal. That is, let the posterior palaeodose distribution of sample 311-6 be $N(982.3, 11.3^2)$, and the dose rate have distribution $N(3.42, 0.18^2)$.

The distribution of the age ratio as a ratio of two independent normal distributions was found, using the density given above. This is shown in Figure 6.3. Also shown is the density of this distribution found using the ratio of simulations from each of the normal distributions. These were used to find the mean and standard deviation of the distribution, which are 288.0 and 15.6 years respectively. Overlaid is the density of the normal distribution with the same mean and variance.

This distribution can be compared to Figure 6.1, the distribution for the age ratio and its estimate, found using simulations from the posterior distribution for palaeodose and annual dose. This has mean 286.9 and standard deviation 15.2 years. The two distributions are similar, though there is a small difference in the posterior means for the age ratio. The standard deviations of the two distributions are comparable.

It is appropriate in this example to approximate the age ratio as a ratio of two independent normals. However this is not always the case, depending on the posterior palaeodose and annual dose distributions. Modelling the age ratio as a ratio of two normals could be carried through to the evaluation of the sample age with further work. However, in the model developed here we take forward only the mean and standard deviation of the age ratio to evaluate the posterior distribution sample age, and so using this method the discrepancy between the density of the ratio of normals, and the normal density in Figure 6.3 is not considered.

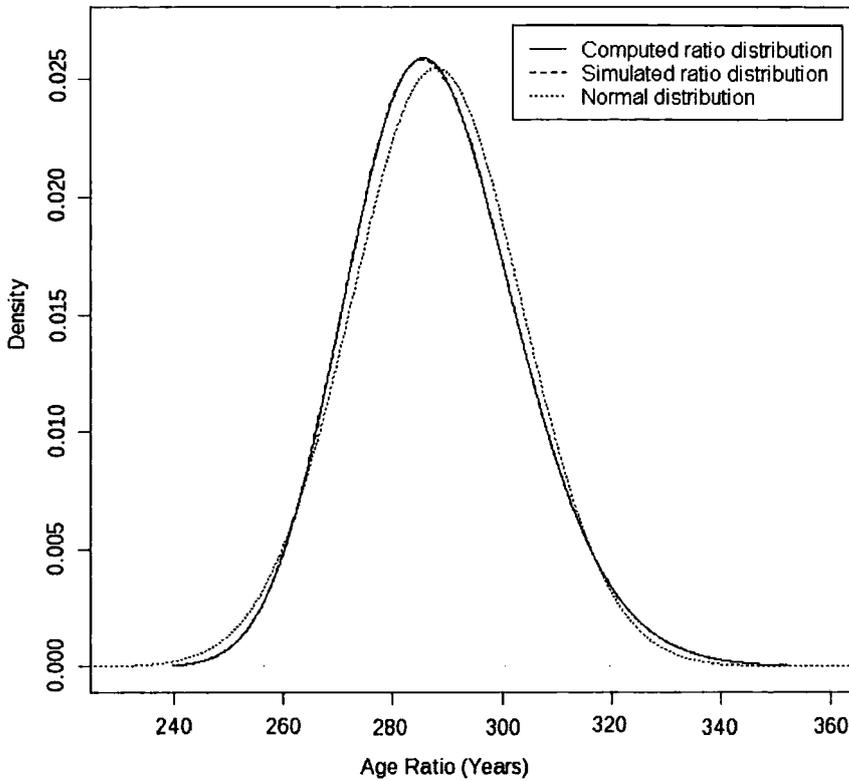


Figure 6.3: Ratio estimate distribution, as a ratio of two independent normal distributions, along with the normal distribution with the same parameters (the distribution of the age ratio).

6.2 Sample Age

The calculation of the luminescence age using routine methods culminates in the evaluation of A_E , the quotient of palaeodose to annual dose. Since this is an estimate of the true sample age A , here a model is developed considering the relationship between these two parameters.

Since even perfect evaluation of A_E would not necessarily be the true sample age, the relationship between the age ratio A_E and sample age A is modelled by

$$A_E = A + \delta^A \quad (6.8)$$

where $\delta^A \sim N(0, \sigma_E^2)$ and σ_E is specified by the expert to reflect judgements about the reliability of A_E for determining A . The lack of detailed knowledge about the

connection between A and A_E leads to this simple model being assigned. In current luminescence dating practise, this source of uncertainty is not explicitly considered. More information about the relationship between A and A_E could be found by dating known age samples, here bricks from buildings with documentary evidence of construction date have been selected.

Bringing together (6.2) and (6.8);

$$\bar{A}_E = A + \delta^A + \delta^E. \quad (6.9)$$

where δ^A , δ^E are independent. The posterior distribution for age can be found using Bayes theorem, considering \bar{A}_E with variance ω_E^2 to be the data input for the model,

$$P[A|\bar{A}_E] \propto P[\bar{A}_E|A]P[A]. \quad (6.10)$$

The justification for using the posterior mean and standard deviation from the age ratio distribution as input in the sample age model follows the same argument as outlined in Section 4.2.2.

Here, suppose a priori that $A \sim N(m_A, \sigma_A^2)$, then

$$\begin{aligned} P[A|\bar{A}_E] &\propto \exp\left\{-\frac{1}{2(\sigma_E^2 + \omega_E^2)} (\bar{A}_E - A)^2\right\} \exp\left\{-\frac{1}{2\sigma_A^2} (A - m_A)^2\right\} \\ &\propto \exp\left\{-\frac{1}{2} \left(\left[\frac{1}{\sigma_E^2 + \omega_E^2} + \frac{1}{\sigma_A^2} \right] A^2 - 2 \left[\frac{\bar{A}_E}{\sigma_E^2 + \omega_E^2} + \frac{m_A}{\sigma_A^2} \right] A \right)\right\} \\ &\propto \exp\left\{-\frac{1}{2} \left(\frac{1}{\sigma_E^2 + \omega_E^2} \right) \left[A - \frac{\frac{\bar{A}_E}{\sigma_E^2 + \omega_E^2} + \frac{m_A}{\sigma_A^2}}{\frac{1}{\omega_E^2 + \sigma_A^2} + \frac{1}{\sigma_A^2}} \right]^2\right\} \end{aligned}$$

so that

$$A|\bar{A}_E \sim N\left(\frac{\frac{\bar{A}_E}{\sigma_E^2 + \omega_E^2} + \frac{m_A}{\sigma_A^2}}{\frac{1}{\sigma_E^2 + \omega_E^2} + \frac{1}{\sigma_A^2}}, \left(\frac{1}{\sigma_E^2 + \omega_E^2} + \frac{1}{\sigma_A^2}\right)^{-1}\right) \quad (6.11)$$

is the posterior age distribution.

The prior information about the sample age is strongly context dependent; in brick dating the architectural style of the building can indicate the period in which it was built, or documentary evidence can be found. A Gaussian prior is used here due to its simplicity and tractability. In some situations other priors may be more suitable, for example a mixture of normals when the dating sample could be assigned to one of two different periods.

6.2.1 Example

The example is continued with sample 311-6 of Fydell House, Lincolnshire. The history of this building is well documented, with records indicating that it was built in the early 18th Century, and purchased by Joseph Fydell in 1726 [88]. Sample 311-6 is associated with the front façade, which is believed to have been altered with the change of ownership in 1726. The measurements were made in 2005, so the age of the sample here is thought of as years before 2005.

This sample is taken from a larger project on dating bricks from post-Medieval buildings [14], which used this ‘known-age’ sample to look at the reliability of luminescence dating methodology. Here, we also use this information for comparative purposes, rather than direct input into the prior distribution for age. Such strong prior information is not commonplace in luminescence dating; if the date of the building is known to within a few years, a luminescence date will not provide any new information. To replicate a routine dating scenario, some of this prior information will be ignored in the model for the sample age, though it will be compared to the posterior age achieved. If the prior input to the model really were as precise as the knowledge here, then this would dominate any information from the data and make the model redundant, and so this would be a poor illustrative example.

For this illustration, the age was chosen to be normal a priori, with mean 280 years before 2005 (date 1725) and standard deviation 25 years. This variance reflects common levels of uncertainty in brick dating, rather than the more detailed knowledge available for this particular building. The prior information in brick dating from a building is unlikely to be vague; architectural style will enable judgements to be made about the period from which the building originates. However, care must be taken to differentiate between the date of the building, and the date at which the brick was fired (which corresponds to the luminescence clock being reset). In particular, Medieval buildings are known for recycling bricks from older buildings [34].

The value assigned to σ_E , the error in the ratio of palaeodose to annual dose in estimating the true age of the sample (6.8) was initially taken to be 5, which is around 2% error in the sample age. This value was chosen after discussions with

Bailiff referring to his experience of brick dating.

The parameters used in the age model, both the data input and prior judgements, are summarised in Table 6.1, and the code is given in Appendix H.4 The posterior distribution achieved using these parameters is normal (6.11), with mean 284.9 and standard deviation 13 . This gives a date of 1720 ± 13 . The variance in the final age is reduced from that in the age ratio ($\omega = 15$), which is the current end point in routine dating. Here, the prior age supports the ratio estimate reducing the uncertainty in the date concluded.

Data Input		Prior Judgements		
\bar{A}_E	ω_E	m_A	σ_A	σ_E
286.9	15.2	280	25	5

Table 6.1: Input in the age model for sample 311-6, Fydell House

6.2.2 Influence of Priors

The prior judgements made about the parameters m_A , σ_A and σ_E will influence the posterior age distribution. Figure 6.4 shows how the posterior mean and standard deviation of age changes with each of these prior parameters. When the influence of any particular parameter is not being considered, its value is held at that given in Table 6.1.

In this example, the posterior mean is most influenced by the value assigned to the prior mean age, m_A . There is a linear relationship between the prior and posterior mean, which is evident from the form of the posterior distribution (6.11), when all of the other parameters are kept constant. The posterior standard deviation of age is not affected by the value assigned to m_A , as shown in Figure 6.4(a). This is clear from the equation given for posterior standard deviation in (6.11), when the prior distribution for age is normal.

The value assigned to σ_A , the prior standard deviation for age, influences both the mean and standard deviation of the posterior age distribution. When σ_A is small, the posterior mean is close to the prior mean, reflecting the strong beliefs in the prior mean chosen. However, as σ_A is increased, indicating greater uncertainty

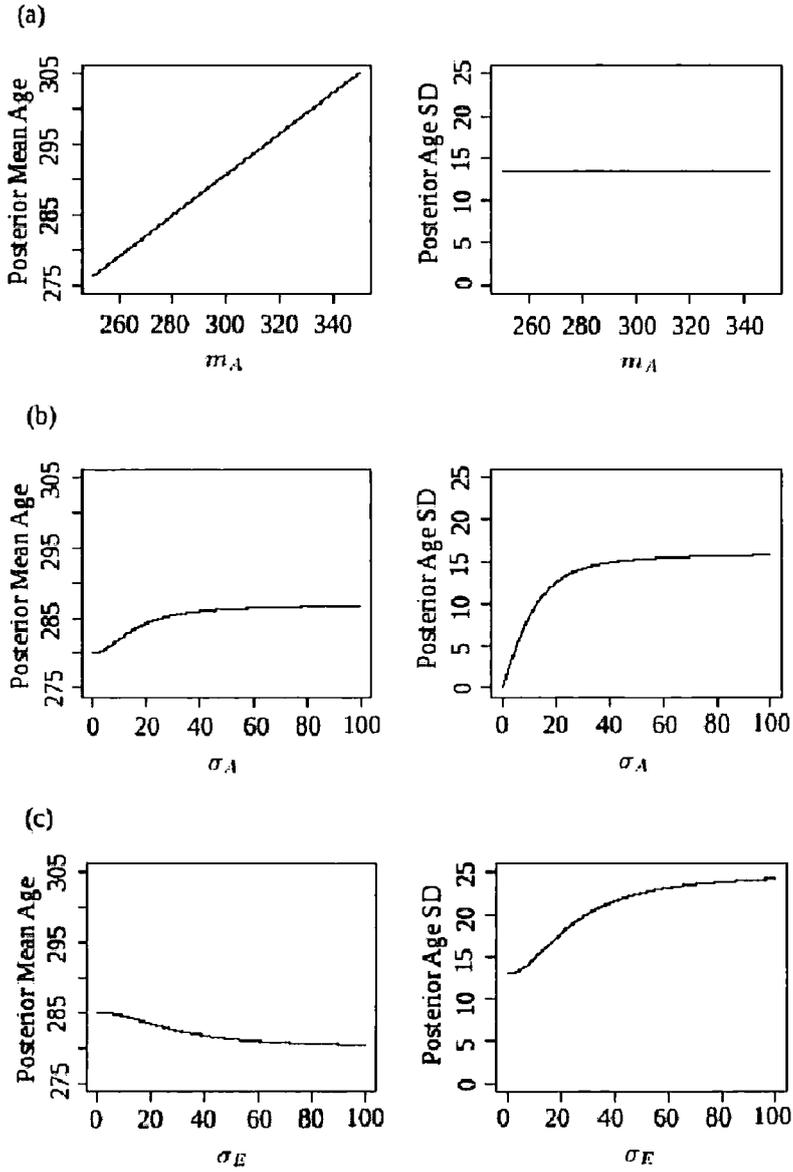


Figure 6.4: Posterior mean and standard deviation of age with (a) m_A , (b) σ_A , (c) σ_E .

a priori in the age of the sample, the posterior mean for the sample age moves towards the ratio estimate. The posterior standard deviation for age is small when there are strong prior beliefs, and increases with σ_A . This increase is sharp until the prior standard deviation reaches a similar magnitude to ω_E , the standard deviation associated with the ratio estimate of age, and in this region the prior distribution is dominant in the posterior age distribution. The posterior standard deviation for age continues to increase, but at a much slower rate as, when the prior is vague, the information from the data will dominate the posterior which is seen in (6.11).

The value assigned to σ_E represents the accuracy of the true value of the ratio of palaeodose to annual dose as an estimate for the age of the sample (6.8). When a very small value is given to σ_E , then the model assumes that the age ratio is very similar to the sample age, and thus the posterior mean is close to the ratio estimate. The mean of the posterior age gravitates towards the prior as σ_E is increased in value as, once σ_E is larger than σ_A (here set at 25), the expert has a greater belief in the prior than in the information from the data. Similarly, the posterior standard deviation for age strongly reflects the data input when σ_E is small, and moves towards the prior as the value assigned to σ_E is increased.

Here we have looked at the influence of each of the prior parameters on the posterior age distribution in turn, with the remaining parameters being set at values given in Table 6.1. However the influence of one parameter can depend on the value of another. Table 6.2 shows how the posterior mean and standard deviation (and so the date achieved) for the sample is affected by the prior parameters.

In Table 6.2 the same pattern in the standard deviation of the posterior age is seen irrespective of the value chosen for the prior mean. The posterior standard deviation increases with the values assigned to σ_E . The prior standard deviation σ_A is similarly influential, although when the prior standard deviation of age is small, ($\sigma_A = 5$) then the posterior age distribution is close to the prior, and the other prior judgements only have a minimal effect on the age.

The mean ages here, and subsequently, are quoted to a higher degree of precision than can be achieved experimentally, for comparison purposes.

This sensitivity analysis has shown how the prior information influences the

Prior Judgements			Posterior Age		Date
m_A	σ_A	σ_E	Mean	SD	(A.D.)
280	5	5	280.6	4.8	1724 ± 5
		25	280.2	4.9	1725 ± 5
		50	280.1	5.0	1725 ± 5
280	25	5	284.9	13.5	1720 ± 13
		25	282.9	19.0	1722 ± 19
		50	281.3	22.6	1724 ± 23
280	50	5	286.3	15.2	1719 ± 15
		25	285.1	25.3	1720 ± 25
		50	283.3	36.1	1722 ± 36
300	5	5	298.8	4.8	1706 ± 5
		25	299.6	4.9	1705 ± 5
		50	300.0	5.0	1705 ± 5
300	25	5	290.7	13.5	1714 ± 14
		25	294.5	19.0	1711 ± 19
		50	297.6	22.6	1707 ± 23
300	50	5	288.1	15.2	1717 ± 15
		25	290.2	25.3	1715 ± 25
		50	293.7	36.1	1711 ± 36

Table 6.2: Mean and standard deviation of the posterior mean age and date of the sample, for different prior judgements.

posterior age distribution. It is important to ensure that the prior judgements of the expert are carefully translated into prior specifications. The sensitivity analysis also shows the model behaves as expected, and as such increases confidence in the output.

6.3 Extended Plateau Example

In Chapter 4 we considered the extended data set of sample 311-6, including aliquots with lower preheat temperatures. This resulted in a mixture of normals for the posterior palaeodose distribution. The mixture contained the distribution $N(970.9, 10.4^2)$ with weight 0.48 and $N(982.3, 11.3^2)$ with weight 0.52. To find the age ratio for the sample, draws from this mixture distribution were combined with draws taken from the distribution for dose rate.

This results in a distribution with mean 286.1 and standard deviation 15.6 years. Using a prior of $A \sim N(280, 25^2)$ and $\sigma_E = 5$ this dates the sample at 1721 ± 14 .

6.4 Comparison with Current Luminescence Age Evaluation

The example of sample 311-6 from Fydell House was part of a project on late and post-medieval brick buildings [14] with strong documentary evidence of the age of the building in order to give confidence to the methodology.

The building contains a brick with the date 1726 along with ironwork similarly embossed [14] though this marks the date which Joseph Fydell purchased the house [88]. The date achieved by Bailiff [14] using conventional luminescence analysis was 1721 ± 17 , which is very similar to the 1720 ± 13 (or 1721 ± 14 using the extended data set) from the Bayesian analysis.

When these dates are compared to the ‘known’ date of 1726, the simple model used to represent the relationship between age ratio and sample age (6.8) does not seem inappropriate in this case.

Chapter 7

Inference with Related Samples

In many dating situations a number of samples may be taken from site to check for consistency and/or contemporaneity. In this chapter we look at possible relationships between samples, and the correlation structure induced by the luminescence dating methodology.

7.1 Coeval Model

The model is applied to samples which are thought to have the same age, i.e they are said to be coeval. First the model is developed for two coeval samples, and then generalised to the case of M coeval samples. An example with two samples is then presented. Dating situations where coeval samples would occur include sampling different bricks from the same building, or taking a number of samples from the same layer of sedimentation (and so they would have been bleached at the same time).

Consider two samples with the same age. An estimate for the age based on the ratio of palaeodose to annual dose will be made for each of these samples. The relationship between the measured (\bar{A}_{Ei}) and actual (A_{Ei}) value of this quotient is modelled by

$$\bar{A}_{Ei} = A_{Ei} + \delta_i^E \quad (7.1)$$

where $\delta_i^E \sim N(0, \omega_{Ei}^2)$ for sample $i = 1, 2$. As in the age model based on one sample (Section 6.1), ω_{Ei} is taken to be the variance of the posterior distribution of the

ratio of palaeodose to annual dose found for each sample. The ratio estimates of age from two samples which are the same age (e.g. taken from two bricks in the same building) are correlated, so let

$$\begin{matrix} \bar{A}_{E1} \\ \bar{A}_{E2} \end{matrix} \left| \begin{matrix} A_{E1} \\ A_{E2} \end{matrix} \right. \sim N \left(\left(\begin{matrix} A_{E1} \\ A_{E2} \end{matrix} \right), \begin{pmatrix} \omega_{E1}^2 & \rho_E \omega_{E1} \omega_{E2} \\ \rho_E \omega_{E1} \omega_{E2} & \omega_{E2}^2 \end{pmatrix} \right) \quad (7.2)$$

The correlation coefficient ρ_E will be higher when the errors in measuring the ratio of palaeodose to annual dose are systematic rather than sample specific. These errors can be subdivided into those associated with measuring the palaeodose and those with measuring the annual dose, as it is assumed that these two quantities are independent.

Systematic errors across samples will arise through measurement errors and calibration of laboratory equipment (e.g. radiation sources), assuming that the measurements are made in the same luminescence laboratory. The luminescence properties of the sample could also be prone to inducing errors in palaeodose evaluation, for example if there were poor or inconsistent dose recovery in one sample, it is likely another will have similar properties if they are taken from the same building. The dose rate measurements (Chapter 5) of \dot{D}_β , \dot{D}_γ , W will be taken for each sample. However, the values of H_β , H_γ , b , g , \dot{D}_c employ parameters or have values drawn from a common data set used by the luminescence community, and so this will induce systematic errors and thus correlation between the two dose rate evaluations, and hence the two age estimates.

The age ratios A_{Ei} represent the true value of the quotient of palaeodose and annual dose for each sample, $i = 1, 2$. These are related to the sample age, A by

$$A_{Ei} = A + \delta_i^A \quad (7.3)$$

where $\delta_i^A \sim N(0, \sigma_E^2)$ for sample $i = 1, 2$, and so

$$\begin{matrix} A_{E1} \\ A_{E2} \end{matrix} \left| A \sim N \left(\left(\begin{matrix} A \\ A \end{matrix} \right), \begin{pmatrix} \sigma_E^2 & \rho_A \sigma_E^2 \\ \rho_A \sigma_E^2 & \sigma_E^2 \end{pmatrix} \right) \right. \quad (7.4)$$

where ρ_A represents the correlation between the age estimates from the two samples. In this model, σ_E is taken to have the same value for each sample (to be specified

by the expert), though it would be possible for this parameter to take a different value for each sample if it was deemed appropriate for the context.

Pulling together (7.1) and (7.3) and assuming that δ_i^E and δ_i^A are independent,

$$\bar{A}_{Ei} = A + \delta_i^E + \delta_i^A \quad (7.5)$$

and so

$$\begin{matrix} \bar{A}_{E1} \\ \bar{A}_{E2} \end{matrix} \Bigg| A \sim N \left(\begin{pmatrix} A \\ A \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \right) \quad (7.6)$$

where, for $i = 1, 2$

$$\sigma_i = \sqrt{\omega_{Ei}^2 + \sigma_E^2} \quad (7.7)$$

$$\rho = \frac{\rho_E \omega_{E1} \omega_{E2} + \rho_A \sigma_E^2}{\sigma_1 \sigma_2}. \quad (7.8)$$

The posterior distribution for the age of the two coeval samples, A , is found using Bayes Theorem,

$$P[A|\bar{A}_{E1}, \bar{A}_{E2}] \propto P[\bar{A}_{E1}, \bar{A}_{E2}|A]P[A] \quad (7.9)$$

Again here, a Gaussian prior distribution is assigned to the sample age, $A \sim N(m_A, \sigma_A^2)$, so that

$$\begin{aligned} P[A|\bar{A}_{E1}, \bar{A}_{E2}] &\propto \exp \left\{ -\frac{1}{2(1-\rho^2)} \left[\frac{(\bar{A}_{E1} - A)^2}{\sigma_1^2} - \frac{2\rho(\bar{A}_{E1} - A)(\bar{A}_{E2} - A)}{\sigma_1\sigma_2} \right. \right. \\ &\quad \left. \left. + \frac{(\bar{A}_{E2} - A)^2}{\sigma_2^2} \right] \right\} \exp \left\{ -\frac{1}{2\sigma_A^2} (A - m_A)^2 \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{1}{(1-\rho^2)} \left(\frac{1}{\sigma_1^2} - \frac{2\rho}{\sigma_1\sigma_2} + \frac{1}{\sigma_2^2} \right) + \frac{1}{\sigma_A^2} \right] A^2 \right. \right. \\ &\quad \left. \left. - 2 \left[\frac{1}{(1-\rho^2)} \left(\frac{\bar{A}_{E1}}{\sigma_1^2} - \frac{\rho(\bar{A}_{E1} + \bar{A}_{E2})}{\sigma_1\sigma_2} + \frac{\bar{A}_{E2}}{\sigma_2^2} \right) + \frac{m_A}{\sigma_A^2} \right] A \right) \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left[\frac{1}{(1-\rho^2)} \left(\frac{1}{\sigma_1^2} - \frac{2\rho}{\sigma_1\sigma_2} + \frac{1}{\sigma_2^2} \right) + \frac{1}{\sigma_A^2} \right] \right. \\ &\quad \left. \left(A - \frac{\frac{1}{(1-\rho^2)} \left(\frac{\bar{A}_{E1}}{\sigma_1^2} - \frac{\rho(\bar{A}_{E1} + \bar{A}_{E2})}{\sigma_1\sigma_2} + \frac{\bar{A}_{E2}}{\sigma_2^2} \right) + \frac{m_A}{\sigma_A^2}}{\frac{1}{(1-\rho^2)} \left(\frac{1}{\sigma_1^2} - \frac{2\rho}{\sigma_1\sigma_2} + \frac{1}{\sigma_2^2} \right) + \frac{1}{\sigma_A^2}} \right)^2 \right\}. \quad (7.10) \end{aligned}$$

So the posterior distribution is normal,

$$A|\bar{A}_{E1}, \bar{A}_{E2} \sim N(\mu_P, \sigma_P^2) \quad (7.11)$$

where

$$\mu_P = \frac{\frac{1}{(1-\rho^2)} \left(\frac{\bar{A}_{E1}}{\sigma_1^2} - \frac{\rho(\bar{A}_{E1} + \bar{A}_{E2})}{\sigma_1\sigma_2} + \frac{\bar{A}_{E2}}{\sigma_2^2} \right) + \frac{m_A}{\sigma_A^2}}{\frac{1}{(1-\rho^2)} \left(\frac{1}{\sigma_1^2} - \frac{2\rho}{\sigma_1\sigma_2} + \frac{1}{\sigma_2^2} \right) + \frac{1}{\sigma_A^2}}$$

$$\sigma_P^2 = \left[\frac{1}{(1-\rho^2)} \left(\frac{1}{\sigma_1^2} - \frac{2\rho}{\sigma_1\sigma_2} + \frac{1}{\sigma_2^2} \right) + \frac{1}{\sigma_A^2} \right]^{-1}.$$

7.1.1 General Model for m Coeval Samples

The model can be generalised to apply to m samples with the same age. Let the m ratio estimates for age be denoted $\bar{\mathbf{A}}_{Em}$ so

$$\bar{\mathbf{A}}_{Em} = (\bar{A}_{E1}, \dots, \bar{A}_{Em})^T \quad (7.12)$$

and let

$$\bar{\mathbf{A}}_{Em}|A \sim N((A, \dots, A)^T, \Sigma). \quad (7.13)$$

When a Gaussian prior for age is used, $A \sim N(m_A, \sigma_A^2)$ then the posterior age distribution is also normal,

$$A|\bar{\mathbf{A}}_{Em} \sim N(\mu_c, \sigma_c^2) \quad (7.14)$$

where

$$\mu_c = \left(\frac{1}{\sigma_A^2} + B_m^T \Sigma^{-1} B_m \right)^{-1} \left(\frac{m_A}{\sigma_A^2} + B_m^T \Sigma^{-1} \bar{\mathbf{A}}_{Em} \right) \quad (7.15)$$

$$\sigma_c = \left(\frac{1}{\sigma_A^2} + B_m^T \Sigma^{-1} B_m \right)^{-1/2} \quad (7.16)$$

where $B_m = (1, \dots, 1)^T$, a vector of length m .

7.1.2 Example

Two samples, labelled 311-2, 311-4 were taken from two bricks that are believed to form part of the original walls of Fydell House (whereas the the example used previously, sample 311-6, is thought to be part of a later renovation). This is part of the same project as sample 311-6, where the data collection and laboratory analysis were performed by Bailiff [14].

For each sample, first the combined aliquot model was applied to aliquots at each preheat temperature (Chapter 3). The plateau model was then implemented,

and the aliquots which produce palaeodose estimates that lie on the preheat plateau were used to evaluate the sample palaeodose. The dose rate was computed using the model given in Chapter 5, and an estimate for the age ratio, \bar{A}_{Ei} was found using Section 6.1. The details of these evaluations are given in Appendix I, along with the original data for the two samples. The resulting age estimates and their standard deviations are given in Table 7.1.

Sample 311-2		Sample 311-4	
\bar{A}_{E1}	ω_{E1}	\bar{A}_{E2}	ω_{E2}
260.9	15	273.5	18

Table 7.1: Ratio estimates and their standard deviations achieved for samples 311-2, 311-4.

The majority of the correlation between the two ratio estimates \bar{A}_{E1} , \bar{A}_{E2} is considered to originate in the denominator of the age equation (1.1), the dose rate. This is because the evaluation of the dose rate relies on a number of standard parameter values that are not sample specific, and take community-wide accepted values. To assess the magnitude of the correlation between the dose rates in this case, and therefore give an indication of the value that should be assigned to ρ_E in the coeval model, the two dose rates were simulated jointly.

To simulate the two dose rates jointly, the correlation between each of the parameters in the dose rate model (Section 5.1) needs to be considered. These parameters can be divided into 3 categories:

1. **Measured Values.** These parameters are based on experimental measurements (\dot{D}_β , \dot{D}_γ , W).
2. **Community-wide values.** These parameters are not sample specific, and take the same values for each evaluation of the dose rate (H_β , H_γ , b , g) throughout the luminescence dating community.
3. **Expert Judgement** The time-averaged water uptake level of the sample takes values based on judgements made by the practitioner.

The remaining parameter in the model is \dot{D}_c , the cosmic dose rate. Usually, the cosmic dose rate is taken to be the same for similar dating situations.

As the two samples were processed in the same luminescence laboratory, and the measurements taken using the same equipment, there is likely to be a level of correlation between the measured values through systematic errors. Here, it was assumed that the correlation between each of the measured values is 0.2.

The true values of the parameters which take community wide values may not be identical for each of the samples, but are likely to be very similar. So, here the correlation between the two samples for each of these parameters was taken to be 0.999. The cosmic dose rate contribution, \dot{D}_c , will also be highly correlated, with the correlation coefficient being set to 0.9 here.

The water content history for the two samples are likely to be very similar, as they have been taken from bricks in the same buildings. So, the correlation between each of the sample's F parameter is going to be very high. It would be complex to take draws from the appropriate bivariate distribution. So here we use the same draw from a beta distribution for each sample. The purpose of this calculation is to get an idea of the correlation between the two dose rates, and any errors induced by this approximation are likely to be small.

The distributions used to simulate the two dose rates jointly are summarised in Table 7.2. After 100,000 simulations the estimated dose rates were 3.89 ± 0.20 and 3.97 ± 0.21 . These dose rates have a correlation of 0.85. This strong correlation reflects the dependence of the annual dose on the common correction parameters.

Here we assume that the palaeodose evaluations for each sample are independent. This is not strictly true, as the measurements were all taken in the same laboratory so inducing a source of systematic error. However, this is minimal and the difficulty in modelling the small amount of correlation outweighs the effect it will have on the final age. So, a correlation of 0.85 between the two dose rates, in the denominator, leads to a correlation between the ratio estimates \bar{A}_{E1} , \bar{A}_{E2} of 0.55. This is an estimate of the magnitude of ρ_E , the correlation between the ratio estimates conditional on the age ratios, A_{E1} , A_{E2} .

The parameter values used in applying the coeval model to this example are

Parameter	Mean		SD		Correlation
	311-2	311-4	311-2	311-4	
\dot{D}_β	2.80	2.91	(0.025)(2.80)	(0.025)(2.91)	0.2
\dot{D}_γ	1.22	1.20	(0.025)(1.22)	(0.025)(1.20)	0.2
\dot{D}_c	0.2	0.2	(0.025)(0.2)	(0.025)(0.2)	0.9
b	0.92	0.92	0.05	0.05	0.99
g	0.93	0.93	0.1	0.1	0.99
H_β	1.25	1.25	0.1	0.1	0.99
H_γ	1.14	1.14	0.1	0.1	0.99
W	0.033	0.033	(0.025)(0.033)	(0.025)(0.033)	0.2

Table 7.2: Parameters of the joint prior distributions used to estimate the joint dose rate distributions of samples 311-2 and 311-4 from Fydeell House.

The parameters above were all assigned bivariate normal distributions with the parameters given. The remaining parameter, F , was assigned a beta distribution, $F \sim \beta(3.0375, 17.2125)$, for each sample.

given in Table 7.3, and the R code is given in Appendix H.5.

\bar{A}_{E1}	ω_{E1}	\bar{A}_{E2}	ω_{E2}	σ_E	ρ_E	ρ_A
260.9	15	273.5	18	5	0.55	0.2

Table 7.3: Parameter values for the coeval model.

Initially, the prior distribution assigned to the sample age was normal with mean 280 years, standard deviation 25 as the samples were taken from the same building as 311-6 (Section 6.2.1). This leads to a posterior age distribution which is normal with mean 268.9 years, standard deviation 12.7 which is a date of 1736 ± 13 . Table 7.4 shows the date estimates obtained if different prior judgements are made about the age of the samples.

Figure 7.1 shows how σ_E , ρ_E and ρ_A affect the posterior mean age and standard deviation. The prior distribution for age was set to $A \sim N(280, 25^2)$, and the other parameters as indicated in Table 7.3. As σ_E is increased, both the posterior mean

Prior	Age (years)	Date (A.D.)
$A \sim N(280, 5^2)$	278 ± 5	1727 ± 5
$A \sim N(280, 25^2)$	269 ± 13	1736 ± 13
$A \sim N(280, 50^2)$	266 ± 14	1739 ± 14
$A \sim N(300, 5^2)$	296 ± 5	1709 ± 5
$A \sim N(300, 25^2)$	274 ± 13	1731 ± 13
$A \sim N(300, 50^2)$	268 ± 14	1737 ± 14

Table 7.4: Posterior age and date estimates obtained with different prior age distributions.

and standard deviation increase. A small value of σ_E implies high confidence in the age ratios as a representation of the sample age, and so the posterior statistics move towards the prior and away from the data input as σ_E rises.

The correlation between the ratio estimates \bar{A}_{E1} , \bar{A}_{E2} conditional on A_{E1} , A_{E2} is denoted ρ_E . When ρ_E is set close to one, the posterior standard deviation of age is higher than when a smaller correlation is modelled. A high level of correlation here causes the increase in posterior variance due to the difference in \bar{A}_{E1} , \bar{A}_{E2} . When ρ_E is less than 0.5, its value has little effect on the posterior age mean. However, as ρ_E is increased to one, the posterior mean falls towards the ratio estimate \bar{A}_{E1} .

The magnitude of ρ_A , the correlation between the age ratios given the sample age, has a minimal effect on the posterior age in this example. It is difficult to specify correlations, but important to establish if they are positive or negative, and low, medium or high. This sensitivity analysis has shown that, in this example, we do not need to spend a long time considering the value of ρ_A as it does not have much influence on the posterior age. The other parameters need to be more carefully considered so that the posterior distribution reflects the data and the prior judgements of the expert.

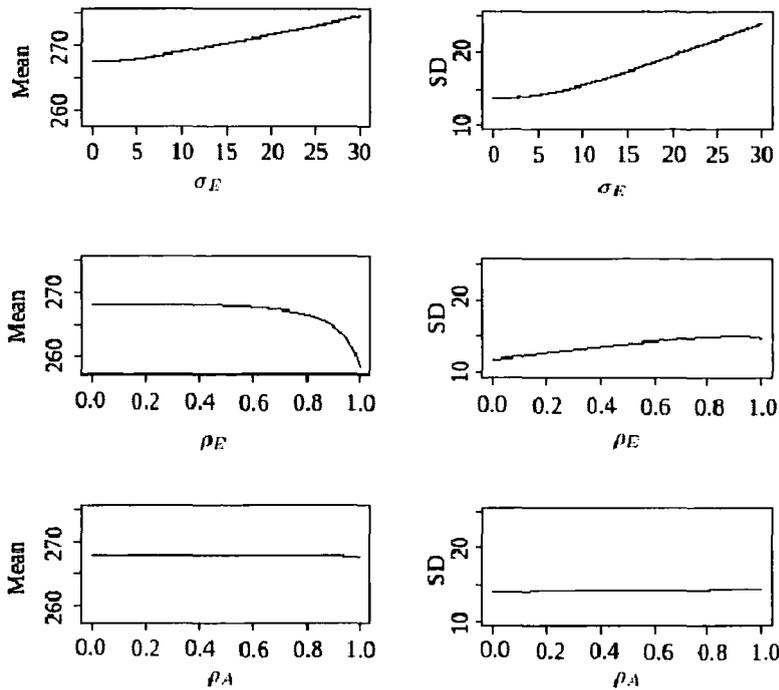


Figure 7.1: Posterior mean and standard deviation of age from the coeval model with σ_E , ρ_E and ρ_A .

7.2 Similar Age Model

Here we consider having a number of samples that have related ages, but are not thought to be coeval. The model is developed for two samples, and the relationship which is thought to exist between the sample ages is dictated by the specification of their joint prior distribution.

As above let \bar{A}_{Ei} , A_{Ei} be the measured and actual values of the ratio of palaeodose to dose rate for samples $i = 1, 2$, and denote the age of the samples A_i , $i = 1, 2$. Similarly, let

$$\bar{A}_{Ei} = A_i + \delta_i^E + \delta_i^A \quad (7.17)$$

with $\delta_i^E \sim N(0, \omega_{Ei}^2)$ and $\delta_i^A \sim N(0, \sigma_E^2)$. Again, ω_{Ei} is the standard deviation of the posterior distribution of the age estimate for each sample, and σ_E^2 represents the uncertainty in the age ratio as a representation of the sample age.

Let the prior distribution for the two sample ages be normal,

$$\begin{pmatrix} A_1 \\ A_2 \end{pmatrix} \sim N \left(\begin{pmatrix} m_{A1} \\ m_{A2} \end{pmatrix}, \begin{pmatrix} \sigma_{A1}^2 & \rho_p \sigma_{A1} \sigma_{A2} \\ \rho_p \sigma_{A1} \sigma_{A2} & \sigma_{A2}^2 \end{pmatrix} \right). \quad (7.18)$$

If the two sample were judged to be close in age, then the values placed on the two prior means would be similar and the correlation ρ_p would be large. The posterior age distribution is then computed by using a Gibbs sampler to draw from the conditional posterior distributions of $A_1|A_2, \bar{A}_{E1}, \bar{A}_{E2}$ and $A_2|A_1, \bar{A}_{E1}, \bar{A}_{E2}$ in turn, updating the values of A_1, A_2 with each iteration, so that when convergence is reached these are approximately independent draws from the posterior distributions for A_1 and A_2 .

The posterior distribution of A_1 conditional on A_2 is

$$\begin{aligned} P[A_1|A_2, \bar{A}_{E1}, \bar{A}_{E2}] &\propto P[\bar{A}_{E1}, \bar{A}_{E2}|A_1, A_2]P[A_1|A_2] \\ &\propto \exp \left\{ -\frac{1}{2(1-\rho^2)} \left(\frac{(A_{E1} - A_1)^2}{\sigma_1^2} - \frac{2\rho(A_{E1} - A_1)(A_{E2} - A_2)}{\sigma_1\sigma_2} \right. \right. \\ &\quad \left. \left. + \frac{(A_{E2} - A_2)^2}{\sigma_2^2} \right) \right\} \\ &\quad \exp \left\{ -\frac{1}{2\sigma_{A1}^2(1-\rho_p^2)} \left[A_1 - m_{A1} - \rho_p \frac{\sigma_{A1}}{\sigma_{A2}} (A_2 - m_{A2}) \right]^2 \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{1}{\sigma_1^2(1-\rho^2)} + \frac{1}{\sigma_{A1}^2(1-\rho_p^2)} \right] A_1^2 \right. \right. \\ &\quad \left. \left. - 2 \left[\frac{A_{E1}}{\sigma_1^2(1-\rho^2)} + \frac{\rho(A_2 - \bar{A}_{E2})}{\sigma_1\sigma_2(1-\rho^2)} + \frac{m_{A1} + \rho_p \frac{\sigma_{A1}}{\sigma_{A2}} (A_2 - m_{A2})}{\sigma_{A1}^2(1-\rho_p^2)} \right] A_1 \right) \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\frac{1}{\sigma_1^2(1-\rho^2)} + \frac{1}{\sigma_{A1}^2(1-\rho_p^2)} \right) \right. \\ &\quad \left. \left[A_1 - \frac{\frac{A_{E1}}{\sigma_1^2(1-\rho^2)} + \frac{\rho(A_2 - \bar{A}_{E2})}{\sigma_1\sigma_2(1-\rho^2)} + \frac{m_{A1} + \rho_p \frac{\sigma_{A1}}{\sigma_{A2}} (A_2 - m_{A2})}{\sigma_{A1}^2(1-\rho_p^2)}}{\frac{1}{\sigma_1^2(1-\rho^2)} + \frac{1}{\sigma_{A1}^2(1-\rho_p^2)}} \right]^2 \right\} \end{aligned}$$

so that

$$A_1|A_2, \bar{A}_{E1}, \bar{A}_{E2} \sim N(\mu_{P1}, \sigma_{P1}^2) \quad (7.19)$$

where

$$\mu_{P1} = \frac{\frac{A_{E1}}{\sigma_1^2(1-\rho^2)} + \frac{\rho(A_2 - \bar{A}_{E2})}{\sigma_1\sigma_2(1-\rho^2)} + \frac{m_{A1} + \rho_p \frac{\sigma_{A1}}{\sigma_{A2}} (A_2 - m_{A2})}{\sigma_{A1}^2(1-\rho_p^2)}}{\frac{1}{\sigma_1^2(1-\rho^2)} + \frac{1}{\sigma_{A1}^2(1-\rho_p^2)}} \quad (7.20)$$

$$\sigma_{P1} = \left(\frac{1}{\sigma_1^2(1-\rho^2)} + \frac{1}{\sigma_{A1}^2(1-\rho_p^2)} \right)^{-1/2}. \quad (7.21)$$

Similarly, by symmetry, the posterior distribution for A_2 given A_1 is normal,

$$A_2|A_1, \bar{A}_{E1}, \bar{A}_{E2} \sim N(\mu_{P2}, \sigma_{P2}^2) \quad (7.22)$$

where

$$\mu_{P2} = \frac{\frac{A_{E2}}{\sigma_2^2(1-\rho^2)} + \frac{\rho(A_1 - \bar{A}_{E1})}{\sigma_2\sigma_1(1-\rho^2)} + \frac{m_{A2} + \rho_p \frac{\sigma_{A2}}{\sigma_{A1}}(A_1 - m_{A1})}{\sigma_{A2}^2(1-\rho_p^2)}}{\frac{1}{\sigma_2^2(1-\rho^2)} + \frac{1}{\sigma_{A2}^2(1-\rho_p^2)}} \quad (7.23)$$

$$\sigma_{P2} = \left(\frac{1}{\sigma_2^2(1-\rho^2)} + \frac{1}{\sigma_{A2}^2(1-\rho_p^2)} \right)^{-1/2}. \quad (7.24)$$

Using a Gibbs sampler to simulate the posterior distributions gives the model scope to be adapted to, for example, age ratios which are a mixture of normals if the posterior plateau location is not certain.

7.2.1 General m Similar Age Model

The Gibbs sampler detailed above for finding the posterior ages of related samples can be generalised to a set of m samples. Consider m samples which have ages $\mathbf{A}_m = (A_1, \dots, A_m)^T$, along with the m ratio estimates $\bar{\mathbf{A}}_{Em} = (\bar{A}_{E1}, \dots, \bar{A}_{Em})$. If Σ represents the covariance matrix for the distribution of $\bar{\mathbf{A}}_{Em}|\mathbf{A}_m$, then

$$\Sigma = [\sigma_i\sigma_j\rho_{ij}] \quad i, j = 1, \dots, m \quad (7.25)$$

where

$$\sigma_i = \sqrt{\omega_{Ei}^2 + \sigma_E^2}, \quad (7.26)$$

$$\rho_{ij} = \frac{\rho_{Eij}\omega_{Ei}\omega_{Ej} + \rho_{Aij}\sigma_E^2}{\sigma_i\sigma_j} \quad i \neq j, \quad (7.27)$$

$$\rho_{ii} = 1. \quad (7.28)$$

Let the prior distribution for \mathbf{A}_m be multivariate normal, with

$$\mathbf{A}_m \sim N_m(\mu_0, \Sigma_0). \quad (7.29)$$

If \mathbf{A}_k denotes $\mathbf{A}_m \setminus A_1 = (A_2, \dots, A_m)^T$, i.e. the vector of sample ages with A_1 removed, then consider the posterior distribution for $A_1|\mathbf{A}_k$

$$P[A_1|\bar{\mathbf{A}}_{Em}, \mathbf{A}_m \setminus A_1] \propto P[\bar{\mathbf{A}}_{Em}|\mathbf{A}_m]P[A_1|\mathbf{A}_m \setminus A_1], \quad (7.30)$$

$$P[A_1|\bar{\mathbf{A}}_{Em}, \mathbf{A}_k] \propto P[\bar{\mathbf{A}}_{Em}|\mathbf{A}_m]P[A_1|\mathbf{A}_k]. \quad (7.31)$$

Then write the parameters of the prior distribution for \mathbf{A}_m as

$$\begin{aligned} \mu_0 &= \begin{pmatrix} \mu_1 \\ \mu_k \end{pmatrix}, \\ \Sigma_0 &= \begin{pmatrix} \Lambda_{11} & \Lambda_{1k} \\ \Lambda_{k1} & \Lambda_{kk} \end{pmatrix}. \end{aligned} \quad (7.32)$$

Then the conditional posterior distribution $P[A_1|\bar{\mathbf{A}}_{Em}, \mathbf{A}_k]$ is normal with mean

$$\frac{(\mu_1 + \Lambda_{1k}\Lambda_{kk}^{-1}(\mathbf{A}_k - \bar{\mathbf{A}}_k))(\Lambda_{11} - \Lambda_{1k}\Lambda_{kk}^{-1}\Lambda_{k1})^{-1} + \bar{A}_{E1}(\Sigma^{-1})_{11} + \sum_{j=2}^k (\bar{A}_{Ej} - A_j)(\Sigma^{-1})_{1j}}{((\Sigma^{-1})_{11} + (\Lambda_{11} - \Lambda_{1k}\Lambda_{kk}^{-1}\Lambda_{k1})^{-1})} \quad (7.33)$$

and variance

$$(\Sigma_{11} + (\Lambda_{11} - \Lambda_{1k}\Lambda_{kk}^{-1}\Lambda_{k1})^{-1})^{-1}. \quad (7.34)$$

Similarly the conditional posterior distributions for all other $P[A_j|\mathbf{A}_{Em}, \bar{\mathbf{A}}_m \setminus A_j]$, $j = 1, \dots, m$ can be found.

7.2.2 Example

The same example was used here as in the coeval model in Section 7.1.2, using samples 311-2 and 311-4 from Fydell House. The parameter values used in applying the similar age model are presented in Table 7.5.

\bar{A}_{E1}	ω_{E1}	\bar{A}_{E2}	ω_{E2}	σ_E	ρ_E	ρ_A
260.9	15	273.5	18	5	0.55	0.2

Table 7.5: Parameter values for the similar age model for samples 311-2, 311-4.

The prior distribution assigned to the ages here is

$$\begin{pmatrix} A_1 \\ A_2 \end{pmatrix} \sim N \left(\begin{pmatrix} 290 \\ 290 \end{pmatrix}, \begin{pmatrix} 25^2 & (0.5)(25^2) \\ (0.5)(25^2) & 25^2 \end{pmatrix} \right) \quad (7.35)$$

Here the two samples have both been taken from the original building of Fydell House, and are thought to be the same age. Thus the prior mean for A_1 and A_2

have been assigned the same value, the magnitude of which is based on documentary evidence. Similarly, the prior standard deviations of A_1 and A_2 are the same, though, as in previous analysis of this example, the standard deviation chosen reflects common levels in routine dating rather than the exceptional prior information available in this case. The influence of these judgements is looked at in Section 7.2.3 below. The R code used to find the posterior distributions is given in Appendix H.6

The posterior age distributions for 311-2 (A_1) and 311-4 (A_2) are shown in Figure 7.2. The posterior distribution for A_1 has mean 266.2, standard deviation 13.1 years while the posterior distribution for A_2 has mean 275.0, standard deviation 14.8. This leads to the dates 1739 ± 13 and 1730 ± 15 being assigned to 311-2 and 311-4 respectively.

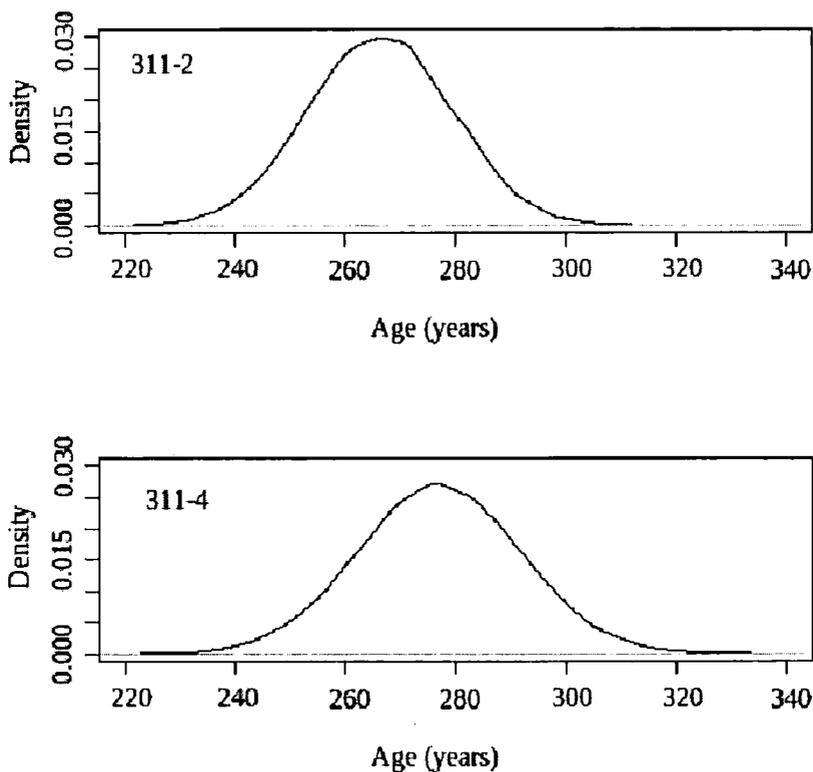


Figure 7.2: Posterior age distributions from the similar age model for 311-2 and 311-4.

These dates can be compared to the date achieved through the coeval model for the same samples with analogous prior specifications (Section 7.1.2). Modelling the

two samples as having the same age in the coeval model, a priori $A \sim N(280, 25^2)$ leads to a date of 1736 ± 13 which, as expected, falls between the two dates achieved through the similar age model.

7.2.3 Influence of Prior Specifications

The Gibbs sampler of the similar age model was run with a range of different prior specifications to look at how the prior judgements made influence the posterior age distributions and the dates evaluated for samples 311-2, 311-4. Table 7.6 gives the posterior ages and dates for the different prior distributions used.

Here we have only looked at cases where prior hyperparameters $m_{A1} = m_{A2}$ and $\sigma_{A1} = \sigma_{A2}$, as that is appropriate for this dating situation.

m_{A1}, m_{A2}	σ_{A1}, σ_{A2}	ρ_p	Age 311-2, 311-4 (years)	Date 311-2, 311-4 (A.D.)
280	5	0.5	$278 \pm 5, 279 \pm 5$	$1727 \pm 5, 1726 \pm 5$
280	5	0.8	$278 \pm 5, 279 \pm 5$	$1727 \pm 5, 1726 \pm 5$
280	25	0.2	$266 \pm 13, 277 \pm 15$	$1739 \pm 13, 1728 \pm 15$
280	25	0.5	$266 \pm 13, 275 \pm 15$	$1739 \pm 13, 1730 \pm 15$
208	25	0.8	$267 \pm 13, 273 \pm 14$	$1738 \pm 13, 1732 \pm 14$
280	50	0.5	$263 \pm 15, 274 \pm 18$	$1742 \pm 15, 1731 \pm 18$
280	50	0.8	$263 \pm 15, 273 \pm 17$	$1742 \pm 15, 1732 \pm 17$
300	5	0.5	$296 \pm 5, 298 \pm 5$	$1709 \pm 5, 1707 \pm 5$
300	25	0.5	$272 \pm 13, 282 \pm 15$	$1733 \pm 13, 1723 \pm 15$
300	50	0.5	$265 \pm 15, 277 \pm 17$	$1740 \pm 15, 1728 \pm 17$

Table 7.6: Influence of prior judgements on the posterior ages achieved with the similar age model.

When the prior standard deviation is small, then the posterior distributions are pulled towards the prior. Conversely, a large value assigned to σ_{A1}, σ_{A2} leads to the data input $\bar{A}_{E1}, \bar{A}_{E2}$ dominating the posterior distributions, and the resulting dates for 311-2, 311-4 are further apart. Similarly, the magnitude of the prior mean has a greater influence on the dates achieved when the prior standard deviation is small.

When the correlation of the two ages is high a priori this reflects the expert's judgements about the relationship between the two samples. Here, as the two prior means are set to the same value, a high ρ_p value represents the belief that the ages of the two samples are close. This is seen in Table 7.6 when $m_1, m_2 = 280$ and $\sigma_{A_1}, \sigma_{A_2} = 25$, the posterior ages are closer together when $\rho_p = 0.8$ compared to when $\rho_p = 0.2$. The effect of the value chosen for σ_E , the uncertainty of the true age around the age ratio, is shown in Figure 7.3.

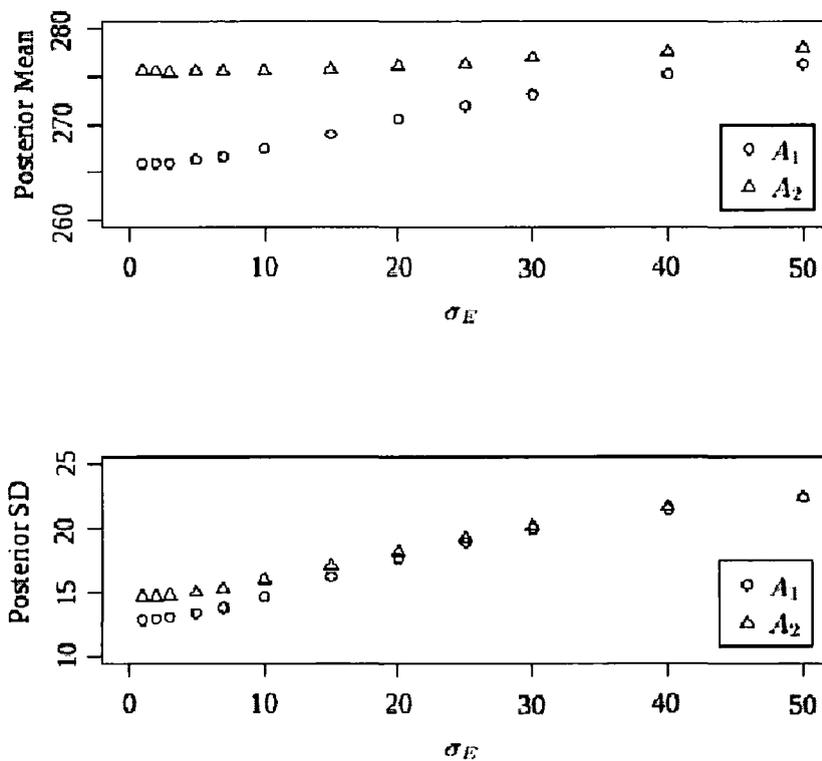


Figure 7.3: Influence choice of σ_E has on the posterior mean and standard deviation for A_1 and A_2 under the similar age model.

For both A_1 and A_2 , the posterior mean and standard deviation move towards the prior values as the magnitude of σ_E is increased. The parameter σ_E represents the confidence that the practitioner has in the age ratio as a representative for the age of the sample. Therefore, it is natural that when σ_E is large compared to the prior standard deviation $\sigma_{A_1}, \sigma_{A_2}$ then the prior distribution will gain more weight in the posterior.

In the similar age model the specification of the prior distributions represents how ‘similar’ the samples are considered to be. Therefore, it is important that this specification is done out carefully, and a sensitivity analysis carried out to ensure that this information is correctly represented.

7.3 Ordered Age Model

In many dating situations, the relative chronology of the samples is known. To model a simple example, assume that

$$A_1 < A_2. \tag{7.36}$$

There is no scope in this model to allow any uncertainty in this ordering. The application of such constraints is most suitable to sediment dating, where the chronology of the samples is dictated by their stratigraphic relationship. The samples taken from the layer closest to the surface will have been bleached by sunlight most recently (so the luminescence clock reset), and thus is the youngest. Experts can often give a precise ordering of sample ages, though in some cases post-depositional mixing or exhumation of sediments, known as pedoturbation [15] may occur. This is often caused by flora and fauna, which is known as bioturbation. The resulting vertical and lateral disturbance need to be taken into account when dating such sediments [16].

In brick dating, such issues do not occur, though unlike sediment dating it is difficult (and often not appropriate) to place a relative chronology on samples with certainty. Here we are taking these brick samples to allow comparisons to be drawn with the similar age model while illustrating the potential of the model. It would not be recommended to apply this model in the age analysis of these two samples, though we can see the potential of the model for more appropriate dating situations.

The posterior probability distribution for the sample ages are found using the order constraint along with the model above for two ages in Section 7.2. A Gibbs sampler (Section 3.3) is used, first a draw is made from the conditional posterior distribution $P[A_1|A_2, \bar{A}_{E1}, \bar{A}_{E2}]$, then using this updated value for A_1 draws are repeatedly made from $P[A_2|A_1, \bar{A}_{E1}, \bar{A}_{E2}]$ until one satisfies the condition of $A_1 <$

A_2 , and this is taken to be the updated value of A_2 . The code for the Gibbs sampler with the rejection criteria is shown in Appendix H.7, where the prior distribution for A_1 and A_2 is specified as

$$\begin{pmatrix} A_1 \\ A_2 \end{pmatrix} \sim N \left(\begin{pmatrix} m_{A1} \\ m_{A2} \end{pmatrix}, \begin{pmatrix} \sigma_{A1}^2 & \rho_p \sigma_{A1} \sigma_{A2} \\ \rho_p \sigma_{A1} \sigma_{A2} & \sigma_{A2}^2 \end{pmatrix} \right). \quad (7.37)$$

7.3.1 Example

Consider samples 311-2 and 311-6 from Fydell House, the example discussed previously. Sample 311-6 is taken from a façade from renovations when the ownership of the house was transferred to Fydell [88], whereas 311-2 is from part of the original building. Let the true age of sample 311-6 be denoted A_1 , and let A_2 correspond to the age of 311-2, so that the chronological constraint here is

$$A_1 < A_2. \quad (7.38)$$

Table 7.7 gives the values assigned to the parameters in the ordered age model. A_1 is believed to be younger than A_2 so the prior distribution reflects this: a priori let

$$\begin{pmatrix} A_1 \\ A_2 \end{pmatrix} \sim N \left(\begin{pmatrix} 280 \\ 290 \end{pmatrix}, \begin{pmatrix} 25^2 & (0.5)25^2 \\ (0.5)25^2 & 25^2 \end{pmatrix} \right) \quad (7.39)$$

311-6		311-2				
\bar{A}_{E1}	ω_{E1}	\bar{A}_{E2}	ω_{E2}	σ_E	ρ_E	ρ_A
286.9	15	260.9	15	5	0.55	0.2

Table 7.7: Parameter values for the ordered age model for samples 311-6, 311-2.

The Gibbs sampler with the chronological constraint $A_1 < A_2$ was run over 3 chains for 20,000 iterations, and the code for this is given in Appendix H.7. The resulting posterior distributions for the ages are shown in Figure 7.4 (a), and their statistics given in Table 7.8. These distributions are compared with the posterior age distributions achieved when the similar age model was applied, using the same prior specifications but without the constraint on the order, in Figure 7.4 (b) and Table 7.8.

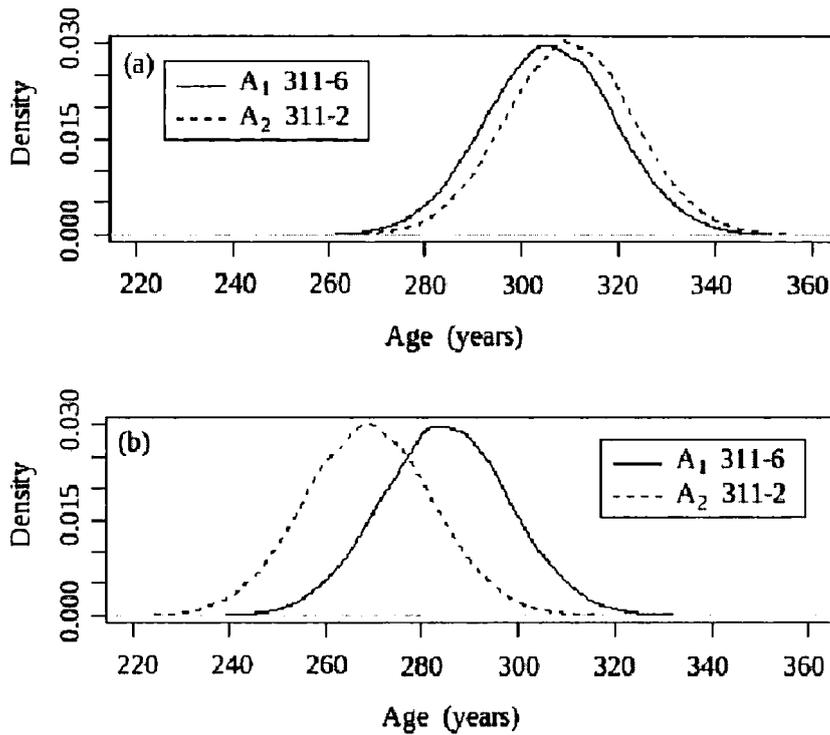


Figure 7.4: Posterior distributions for A_1 (Sample 311-6) and A_2 (Sample 311-2) using (a) Ordered age model with $A_1 < A_2$ and (b) Similar age model.

The condition applied in the ordered age model that $A_1 < A_2$ has a marked effect on the posterior age distributions in this case, particularly as the ratio estimates are ordered $\bar{A}_{E2} < \bar{A}_{E1}$. The similar age model dates both samples a lot younger than the ordered age model, though the posterior standard deviations are the same. The constraint of $A_1 < A_2$ means that each draw of A_2 is forced upwards above A_1 even though the data are implying otherwise, and it also pushes the two dates together.

The influence of the prior judgements made on the posterior age distributions is presented in Table 7.9, along with the run time of the sampler in each case.

As expected, when the prior standard deviation of age is small, the prior age distribution has a greater influence on the posterior. However, when the prior standard deviation of the ages is large, the posterior ages are not close to the data input \bar{A}_{E1} , \bar{A}_{E2} as this contradicts the additional condition that $A_1 < A_2$ in this example.

The value assigned to σ_E , the uncertainty in the age ratio as a representation of

Model	Posterior age (years)		Date (A.D.)	
	A_1	A_2	A_1	A_2
Ordered	306 ± 13	310 ± 13	1699 ± 13	1695 ± 13
Similar	285 ± 13	269 ± 13	1720 ± 13	1736 ± 13

Table 7.8: Posterior Ages resulting from the ordered age model and similar age model.

age, is also influential on the posterior distributions. When σ_E is large, this indicates low confidence in the age ratios and so the posterior ages are dominated by the prior. This also adds more uncertainty into the posterior ages. When σ_E is small, then the age ratios are given greater weighting, but again here the order $A_1 < A_2$ comes into play and so the posterior ages get older.

As well as influencing the posterior age distributions, the choice of prior parameters also has a notable effect on the run time of the sampler in this example. When the prior beliefs are dominant in the analysis, either through a small prior age standard deviation or large value of σ_E , then the sampler is very cheap to run. The sampler becomes much more expensive when there is comparatively high confidence in the ratio estimates. In the sampler, for each iteration the conditional posterior distribution for A_2 is repeatedly drawn from until a value is obtained that satisfies the condition $A_1 < A_2$. As the ratio estimates contradict this ordering, when they dominate the posterior distribution the probability that a draw from the conditional distribution meets the chronological criterion is smaller, and thus more draws need to be made before this is satisfied, and in turn this lengthens the run time of the sampler.

It can be seen that if a large number of ages were involved in such a model, then the rejection criteria for each draw would be a lot more complex, and thus lead to an inefficient sampler. This problem has been encountered in radiocarbon dating, where the use of Bayesian statistics is widespread. Such rejection algorithms are used for simple chronologies [28], though for any more substantial problems the MCMC simulation can be challenging [85].

Such chronological models are also vulnerable to the Stein effect [106]. Archae-

Prior Specifications						Posterior Age (years)		Sampler
m_{A1}	m_{A2}	σ_{A1}	σ_{A2}	ρ_p	σ_E	A_1	A_2	run time (s)
280	290	5	5	0.5	5	281 ± 5	288 ± 5	8
280	290	25	25	0.5	5	306 ± 13	310 ± 13	346
280	290	50	50	0.5	5	315 ± 15	319 ± 15	3471
280	290	25	25	0.1	5	305 ± 13	306 ± 13	2334
280	290	25	25	0.9	5	299 ± 13	303 ± 13	11
280	290	25	25	0.5	15	295 ± 16	303 ± 16	20
280	290	25	25	0.5	25	291 ± 19	303 ± 19	11
280	290	25	25	0.5	50	290 ± 22	306 ± 22	8
270	270	25	25	5	0.5	308 ± 13	312 ± 13	1248
270	270	25	25	0.5	50	294 ± 22	307 ± 22	10
260	265	25	25	0.5	5	302 ± 13	305 ± 13	687
260	265	25	25	0.5	50	276 ± 22	291 ± 21	8

Table 7.9: Posterior Age using the ordered age model with different prior specifications.

ologists are interested in the range of the dates, and thus the temporal duration of the site. However, the Stein effect can lead to over estimation of this parameter, especially when the date range is small compared to the uncertainty in each date [85]. This issue has been overcome in radiocarbon dating by using a uniform prior distribution on the span of dates which is derived from a physical model of deposition [85].

The model outlined above for ordered age samples is a simple extension of the similar age model, and as discussed above contains a number of problems. Any previous inclusion of chronological information in a Bayesian framework using luminescence dates [92] has used the radiocarbon calibration programme OxCal [23, 24]. Here radiocarbon dates are used in conjunction with the luminescence data to produce a chronology for the site being dated. However, as OxCal is designed for use with radiocarbon dates, it is only possible to input the OSL age estimates with their random errors, and any systematic errors have to be added after the analysis.

Chapter 8

Example

In this chapter we take a second example to further illustrate our general approach of Bayesian analysis for luminescence dating. Again a dating environment has been chosen that meets the assumptions made in the thesis (Section 2.6), and the example is a ‘known-age’ sample so comparisons can be made with the age produced using our analysis.

Two samples (labelled 318-1, 318-2) were taken from Tattershall Castle, Tattershall, Lincolnshire by Bailiff as part of the project on late and post-Medieval brick buildings [14], which also includes the example previously considered from Fydell House, Lincolnshire. As for many of the buildings in this project, there is significant documentary evidence for the age of Tattershall Castle, which is used as a comparative tool to evaluate the luminescence dating methodology. Here we ignore some of this unusually precise prior knowledge, applying a more common level of uncertainty to the prior distribution of age and then use the extra information for comparative purposes after the Bayesian analysis.

Samples were taken from the brick tower of Tattershall Castle, construction of which began in 1434-5 for Lord Cromwell, the treasurer of England at that time [88, 97]. A picture of Tattershall castle is shown in Figure 8.1. Accounts from 1445-6 indicate that 322,000 bricks were supplied for the tower and it was constructed under the Flemish ‘brekemaker’ Baldwin [14]. This documentary evidence, along with the architectural style of the building led Bailiff to assign a date range of 1445-1450 [14] before the data were analysed.



Figure 8.1: Tattershall Castle

This image was taken from the Geograph project collection. The copyright on this image is owned by Kate Jewell and is licensed for reuse under the Creative Commons Attribution ShareAlike 2.0 license.

The samples were taken from the interior walls, one from the ground floor of the NE tower (318-2), and one from the basement in the NW tower (318-1), 318-2 is considered here first. We will use the measurements made by the laboratory, along with their expert judgement, to carry out a Bayesian analysis on the sample age as detailed previously.

This analysis looks at each stage of luminescence dating in turn, using the following steps.

1. Evaluation of palaeodose at each preheat temperature.
2. Use these estimates to identify the start of the preheat plateau.
3. Compute the posterior distribution for sample palaeodose based on aliquots which lie on the preheat plateau.
4. Calculate the dose rate distribution.
5. The distribution for age ratio is found using the sample palaeodose and dose rate distributions.

6. The sample age distribution is based on the distribution for age ratio along with the prior specifications.

8.1 Palaeodose Evaluation at each Preheat Temperature

The first stage in the Bayesian age analysis is to evaluate the palaeodose at each preheat temperature, using the combined aliquot model detailed in Chapter 3. In the laboratory, the first step is to prepare the sample and carry out some preliminary experiments.

8.1.1 Preliminary Experiments

A number of preliminary experiments were initially carried out by Bailiff to ascertain the suitability of the sample for luminescence dating and to find appropriate regenerative doses to be used in the single aliquot regeneration (SAR) protocol (Section 2.2.2) to evaluate the palaeodose.

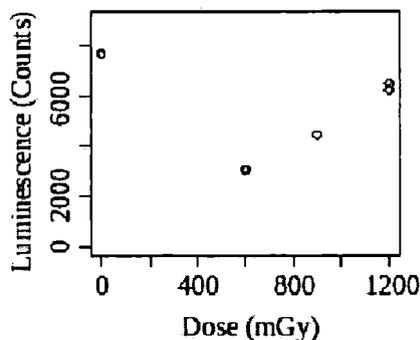


Figure 8.2: The initial preliminary experiment carried out by Bailiff for sample 318-2, Tattershall Castle.

Figure 8.2 shows the initial run of the SAR protocol, carried out at a preheat temperature of 220°C. The regenerative doses were chosen by estimating a possible palaeodose value using the documentary evidence to estimate age and dating experi-

ence to assign a likely value to the dose rate. In Figure 8.2, the natural luminescence value does not fall within the range of luminescence intensities produced by the laboratory irradiated doses, and so they are not optimal for use in the SAR protocol. However, the luminescence produced from the repeated doses are similar, and a linear trend is apparent, so there is no evidence here to suggest the luminescence properties of the sample would render it unsuitable for dating.

As this first range is a trial for application of the SAR protocol, further preliminary readings were taken. Several aliquots were used and after repeated measurements, the doses to be irradiated in the laboratory for the SAR procedure were chosen to be 1494, 1793 and 2091 mGy with the lowest and highest being repeated. The data are given in Appendix G.5.

8.1.2 Prior Elicitation

The prior judgements required for the combined aliquot model (Chapter 3) were based on discussions with Bailiff, and the results from the preliminary experiments. The magnitude of the regenerative laboratory doses were chosen as the expert believed that the palaeodose would be contained within that range. Therefore, the prior distribution for palaeodose was set as

$$x_R \sim N(1750, 150^2). \quad (8.1)$$

The dispersion of the aliquot estimates for palaeodose at each preheat temperature, x_{Rj} , was judged to be around 2-3%, and so γ_R was assigned a value of 50.

It is assumed here that the relationship between dose and luminescence counts is linear (i.e. saturation is not being approached), and that this line goes through the origin. So, the prior mean m_α of the intercept α was given a value of zero. To assign the prior mean of the gradient parameter β , the point on the line made by the natural luminescence intensity and the palaeodose was considered. The data for each aliquot have been normalised to a natural luminescence value of 10000 counts, and so, with $m_\alpha = 0$ and $\mu_R = 1750$, then $m_\beta = 5.7$, approximately $y_R = m_\alpha + m_\beta * \mu_R$.

From the preliminary experiments, a range of intercept values were observed. Based on this, the standard deviation σ_α of the intercept was assigned a value of

200. To consider the value of σ_β , the standard deviation of the gradient, then $m_\alpha \pm 2\sigma_\alpha$ and $\mu_R \pm 2\sigma_R$ were considered in the relation $y_R = \alpha + \beta x_R$ where for this purpose, the error term is ignored. This leads to a range for β of around 4-7 and so $\sigma_\beta = 1$ a priori. The linear coefficients α and β will be negatively correlated, but this correlation is not thought to be strong, so we let $\rho = -0.3$. This leads to the prior distribution

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 5.7 \end{pmatrix}, \begin{pmatrix} 100^2 & -0.3(100)(1) \\ -0.3(100)(1) & 1^2 \end{pmatrix} \right). \quad (8.2)$$

The spread $\gamma_\alpha, \gamma_\beta$ of the aliquot estimates for the linear coefficients, α_j, β_j were assigned values 25 and 5 respectively. The spread of the gradient estimates is expected to be less than that of the intercept estimates, based on the results from the preliminary experiments.

The remaining parameter to elicit a prior distribution for is σ^2 , the standard deviation of the residuals ϵ_{ij} , which are assumed to be independent and identically distributed, $\epsilon_{ij} \sim N(0, \sigma^2)$. Consider the mean of σ to be 50 counts. In the Gibbs sampler, the precision $\tau = 1/\sigma^2$ is used. So, the mean of τ was assigned a value of 0.0004 with a variance of 0.001. This leads to a gamma prior distribution for τ ,

$$\tau = 1/\sigma^2 \sim \Gamma(0.00016, 0.4). \quad (8.3)$$

The judgements made about the prior parameters above are summarised in Table 8.1.

x_R			α, β							$\tau = 1/\sigma^2$	
μ_R	σ_R	γ_R	m_α	σ_α	m_β	σ_β	ρ	γ_α	γ_β	d	a
1750	150	50	0	100	5.7	1	-0.3	25	5	0.00016	0.4

Table 8.1: Prior Parameters for the combined aliquot model to evaluate palaeodose for sample 318-2 from Tattershall Castle.

8.1.3 Posterior Distributions

At each preheat temperature in turn, the posterior distribution for palaeodose was found using the prior specifications given above with the combined aliquot model, the

code for which is given in Appendix H.1. The posterior distributions for palaeodose are shown in Figure 8.3 and their statistics given in Table 8.2.

Preheat	#Aliquots	Posterior Mean	Posterior SD
200	2	1690	50
210	1	1809	53
220	7	1840	28
240	3	1941	77

Table 8.2: Posterior palaeodose distributions at each preheat temperature for sample 318-2.

The posterior palaeodose distribution has lowest variance at a preheat temperature of 220°C, as here the density is based on 7 aliquots and so more information is available from the data. The spread of the posterior palaeodose distribution is largest for a preheat of 240°C. A possible explanation for this behaviour is that the higher preheat thermally erodes the luminescence signal, and so the signal strength is smaller. This means that there is greater error in the counting of the signal, which could lead to greater dispersion in the estimates of palaeodose from each aliquot.

There is some agreement in the palaeodose evaluation at the different preheat temperatures. This will be considered in Section 8.2 where the preheat plateau for palaeodose is identified.

Convergence of the Gibbs Sampler

A number of diagnostic tools have been described in Section 3.4 for assessing the convergence and stability of the Gibbs sampler. It is important to ensure that the sampler has converged before any inferences are made from the posterior distributions.

Here, first the trace plots of the sampler were viewed, the plot for the first 1000 iterations of the first chain for x_R with a preheat of 200°C is shown in Figure 8.4. This trace plot indicates that the sampler has converged as it is ‘spikey’, and remains in the same region throughout.

Table 8.3 shows the evaluation of the posterior palaeodose mean and standard

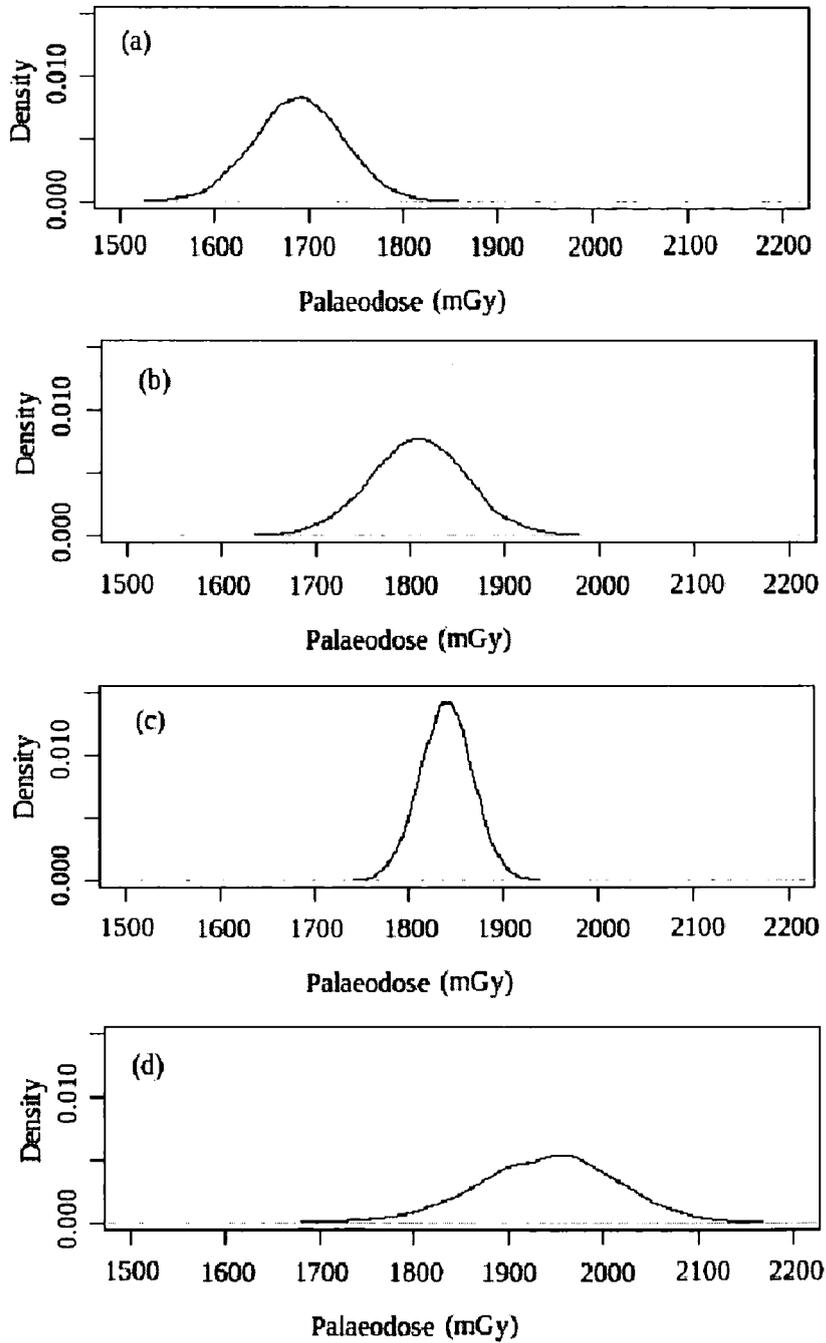


Figure 8.3: Posterior palaeodose distributions for aliquots from sample 318-2 with preheat temperatures (a) 200°C (b) 210°C (c) 220°C (d) 240°C.

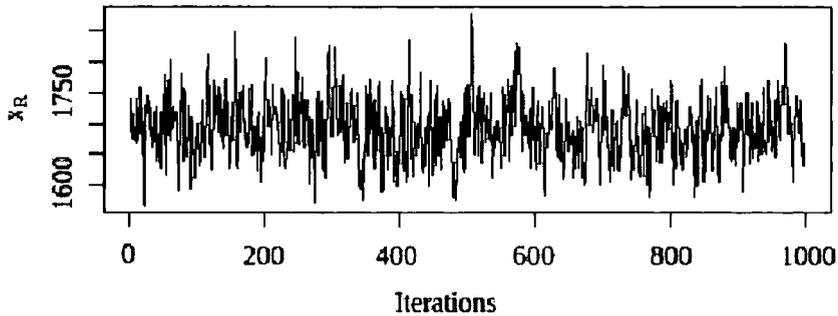


Figure 8.4: Raw trace plot of x_R simulations for the first 1000 iterations with a preheat of 200°C.

Burn-in	Posterior Mean	Posterior SD
0	1689.7	49.7
50	1689.8	49.5
1000	1689.8	49.5

Table 8.3: Posterior palaeodose mean and standard deviation at preheat 200°C evaluated with different burn-in lengths.

deviation with different burn-in periods. Here the length of the burn-in does not have an influence on the posterior mean and standard deviation evaluation. A burn-in of 1000 was chosen, to eliminate the possibility that convergence has not been reached by this point.

To further check the convergence of the sampler, the mean and standard deviation of posterior palaeodose at a preheat of 200°C was computed for different lengths of chain, and these are shown in Table 8.4. For a chain of only 1000 iterations, the posterior standard deviation is slightly higher. However, past this point there is no marked difference in the evaluation of posterior palaeodose, again suggesting that the sampler has reached convergence.

To achieve approximately independent draws from the posterior distribution, the chains are thinned. Table 8.5 gives the posterior palaeodose mean and standard deviation for different amounts of thinning, assuming that a burn-in of 1000 iterations is adopted. Again, the thinning level adopted does not have a large influence on the

# Iterations	Posterior Mean	Posterior SD
1000	1688.9	53.7
5000	1689.4	50.6
10000	1689.4	49.9

Table 8.4: Posterior palaeodose mean and standard deviation at preheat 200°C evaluated for different length of chains.

Thin	Posterior Mean	Posterior SD
1	1689.8	49.5
2	1689.8	49.5
5	1689.5	49.6
10	1689.2	49.4
20	1688.8	49.9

Table 8.5: Posterior palaeodose mean and standard deviation at preheat 200°C evaluated for different thin levels.

outcome of the posterior distribution for palaeodose, and so every 10th iteration was taken from the chains.

The Gelman and Rubin method for testing convergence (Section 3.4.3) looks at the variance between chains as, if convergence has been reached, the inferences from each of the chains should be similar. Here, for the iterations of palaeodose at a preheat of 200°C the test statistic $R_c = 1.000061$. $R_c \rightarrow 1$ as the number of iterations $n \rightarrow \infty$ and this provides evidence that the sampler has converged.

Bringing all these evaluations together, it appears that convergence has been achieved for this sampler. Similar tests were carried out for the simulations at the other preheat temperatures. The posterior distributions were evaluated based on the simulations from a 5-chain 20,000 iteration sampler with a burn-in of 1000 and thinned every 10.

Model Diagnostics

Both linear and Bayesian model diagnostics were carried out to check the suitability of the model and thus validate the posterior distributions for palaeodose achieved.

The residuals for the aliquots with a preheat temperature of 200°C are shown in Figure 8.5. There is no apparent trend in the residual plot, and the magnitude of the residuals is small in comparison to the fitted values, so this plot does not provide evidence that fitting linear model to the data is inappropriate here.

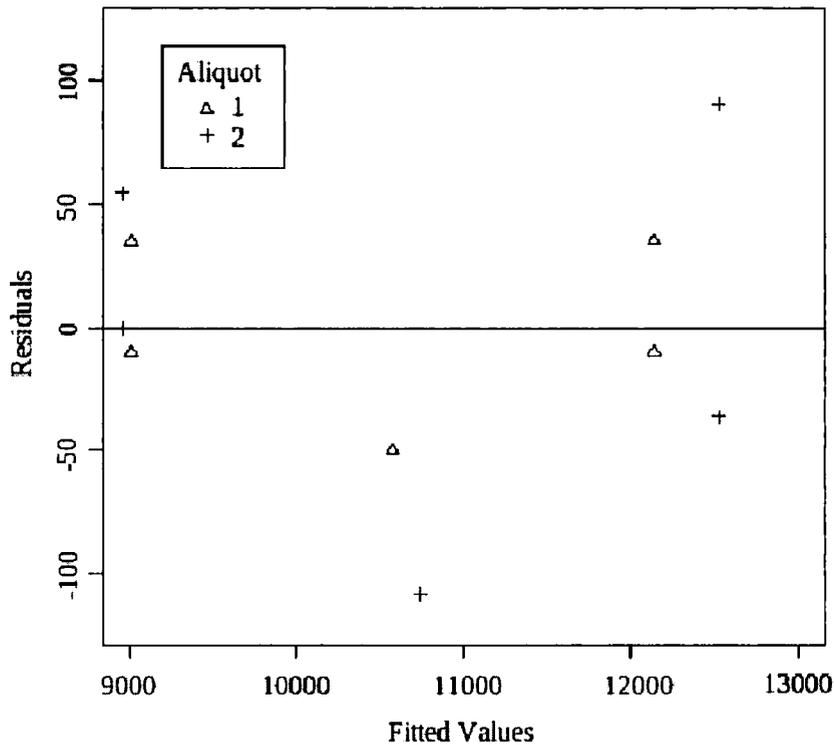


Figure 8.5: Residuals plotted against fitted values of the linear model for aliquots with a preheat of 200°C.

The correlation between the regenerative dose and the intensity of the luminescence response was also calculated for these aliquots. These were 0.9997 and 0.9990 respectively for aliquots 1 and 2 at a 200°C preheat, indicative of a good linear fit.

The residual standard error (RSE) was also computed, with regard to the assumption that the residuals are identically distributed, $\epsilon_{ij} \sim N(0, \sigma^2)$ over all aliquots using the same preheat temperature. The RSE for aliquot 1 is 41.5 and aliquot 2 90.1. Although these values are different, relative to the magnitude of the fitted values, this difference is not marked.

Bringing all of these calculations together, there is no evidence that fitting a

linear model to luminescence intensity and irradiated dose is inappropriate. This is coupled with the fact that the age of the building and the size of the palaeodose lead the expert to believe that the relationship between dose and luminescence is linear and saturation is not being approached.

To check the Bayesian aspect of the model, the prior and posterior means were calculated and these are presented in Table 8.6. Here the posterior mean values attained do not look unrealistic, compared to the prior. It can be noted here that for a preheat of 220°C the difference between the prior and posterior mean is large compared to the posterior standard deviation, with reference to this value calculated for the other preheat temperatures. At 220°C there is a large number of aliquots and so more data, which in turn are likely to reduce the standard deviation and give less weight to the prior.

Preheat	# Aliquots	Prior Mean	Posterior Mean	Difference	Posterior SD
200	2	1750	1689.5	49.6	60.5
210	1	1750	1808.6	58.6	52.7
220	7	1750	1839.5	89.5	28.0
240	3	1750	1940.9	190.9	76.7

Table 8.6: Prior and posterior palaeodose means at each preheat temperature for sample 318-2.

Since none of the diagnostic checks applied here provide any indication that there is a problem with the model, the posterior distributions achieved for palaeodose at each preheat temperature can be accepted and taken forward into the next stage of the analysis: locating the preheat plateau.

8.2 Preheat Plateau

The palaeodose for a sample is evaluated using data from all aliquots which lie on the preheat plateau (Section 2.3). The model to find the distribution for the starting point of the preheat plateau is given in Chapter 4, and uses the means of the posterior palaeodose distributions at each preheat temperature found previously as

data. These are plotted in Figure 8.6 for sample 318-2 along with their two standard deviation uncertainty bars.

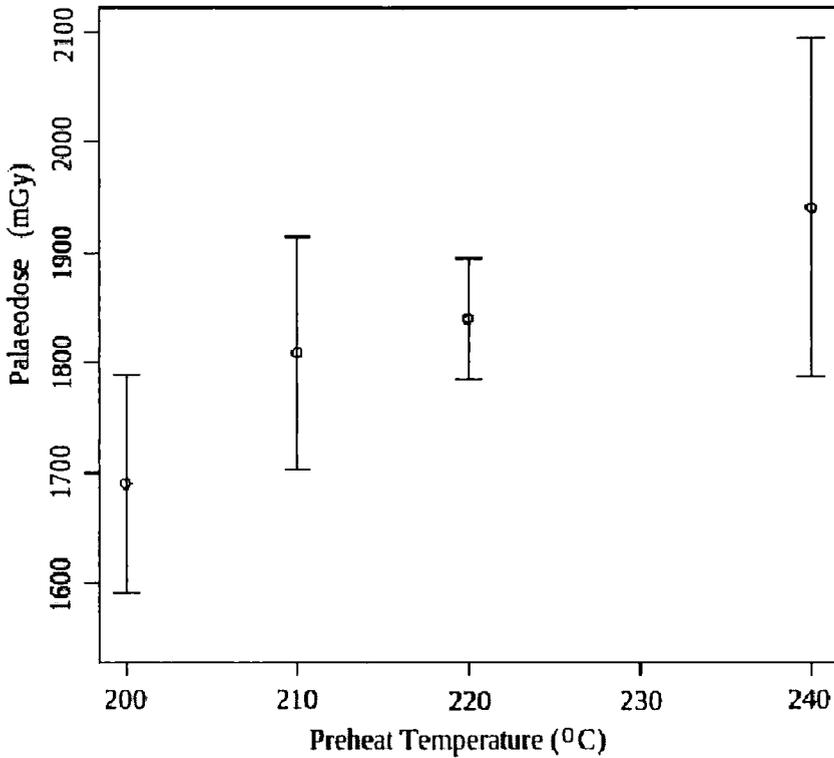


Figure 8.6: Posterior palaeodose means plotted against preheat temperature for sample 318-2, with two standard deviation uncertainty bars of the posterior distributions for palaeodose.

8.2.1 Prior Elicitation

The prior distributions assigned to the parameters in the plateau model (Section 4.2) are given in Table 8.7 and are assumed to be independent a priori.

The level of the plateau, x_{R*} , was assigned the same prior distribution as that of the palaeodose x_R at each of the preheat temperatures in the combined aliquot model. The parameters of the curve leading up to the plateau, η , κ , were assigned prior distributions to allow a wide range of curve shapes, as in the previous example (Section 4.3.1).

x_{R^*}	T_a	η	κ
$N(1750, 150^2)$	$N(200, 15^2)$	$N(0.001, 1^2)$	$N(0.001, 1^2)$

Table 8.7: Prior distributions assigned to the plateau model parameters for Sample 318-2.

The practitioner believes that the preheat plateau has begun before the lowest preheat, otherwise measurements would not have been made at that temperature (as they would not have contributed to the evaluation of the sample palaeodose). After discussions with Bailiff about his past dating experience and the unpredictable nature of luminescence (hence the need to locate the preheat plateau), a prior of mean 200, standard deviation 15°C was assigned to T_a to reflect this level of uncertainty.

8.2.2 Posterior Distribution

The posterior distribution for T_a , the temperature at which the preheat plateau begins, is given in Figure 8.7, and the code detailed in Appendix H.2.

This distribution assigns a probability of 0.0128 to the plateau starting after 200°C and before 210°C, 0.0008 to the plateau starting at a temperature higher than 210°C and lower than 220°, with the remaining majority of probability to the plateau starting before 200°C is reached.

The influence of the choice of prior distribution chosen for T_a on the posterior probability is examined in Table 8.8. This table shows the posterior probability of the plateau starting in particular temperature ranges.

Prior	Posterior Probability		
	$P[T_a < 200]$	$P[200 \leq T_a < 210]$	$P[T_a \geq 210]$
$N(200, 15^2)$	0.9864	0.0128	0.0008
$N(200, 30^2)$	0.9929	0.0061	0.0010
$N(190, 15^2)$	0.9936	0.0062	0.0002
$N(220, 15^2)$	0.9622	0.0289	0.0089

Table 8.8: Posterior probabilities for plateau starting temperature for different prior judgements

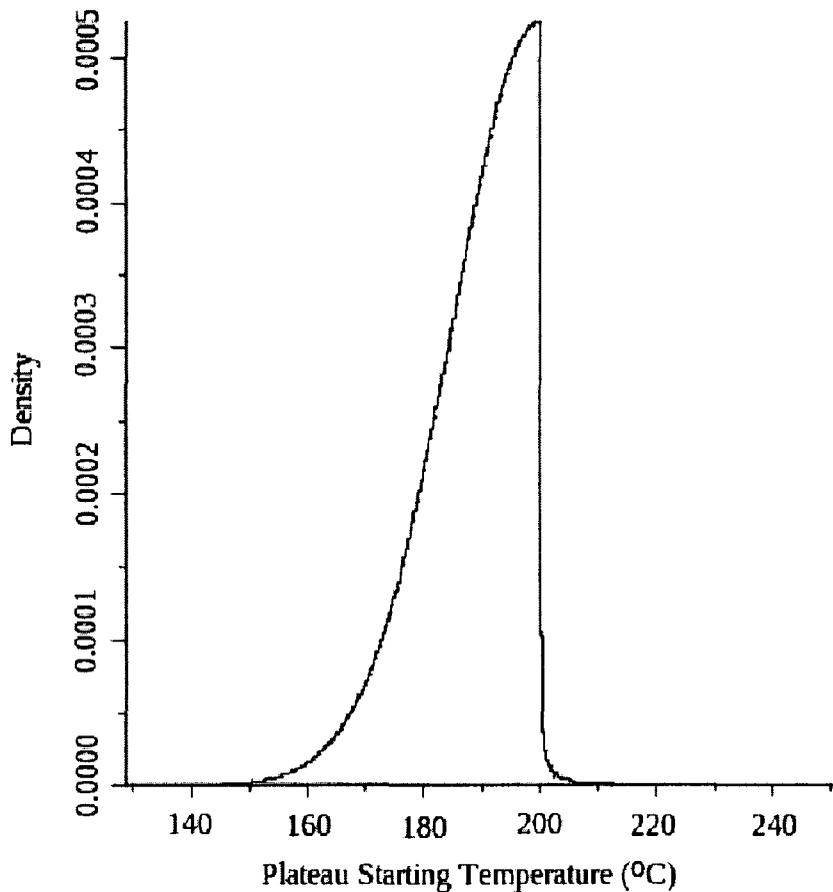


Figure 8.7: Posterior distribution of plateau starting temperature, T_a

The prior distribution for T_a has an effect on the posterior distribution, but for all choices it is most likely the plateau has begun by 200°C. The choice of prior is only influential on the tails of the distribution. When the prior for T_a is broad, or its prior mean is below 200°C, the probability that the plateau starts at a temperature higher than 200°C is less than 1%.

When the prior mean for T_a is chosen to be 220°C, a higher proportion of the probability is assigned to the plateau starting in the range 200-210°C, though this is still a relatively small probability ($\sim 3\%$). The probability of the plateau starting above 210°C is larger than when a smaller value is assigned to the prior mean, but not significant.

8.3 Sample Palaeodose

The palaeodose estimate for sample 318-2 is based on the aliquots which were judged to lie on the preheat plateau. Here, from Figure 8.7, the plateau is most likely to have started at a lower preheat than that used for any aliquot. So all aliquots lie on the preheat plateau and should be used to evaluate the palaeodose of 318-2. The model also assigns a small probability to the plateau starting between 200°C and 210°C, if this was the case the aliquots with a preheat of 200°C should not be included in the evaluation of sample palaeodose.

Here we evaluate the palaeodose for 318-2 using the combined aliquot model with all the aliquots of the sample, and then only with aliquots with a preheat of 210°C and above. A mixture of these two distributions weighted with the posterior probabilities of T_a will make up the palaeodose distribution.

The values assigned to the parameters of the combined aliquot model, used to evaluate palaeodose, are the same prior values which were used when the palaeodose was found at each preheat temperature (Table 8.1). The Gibbs sampler was run for 40,000 iterations over 5 chains, with a burn-in of 1000 iterations and thinning every 5 to produce approximately independent draws from the posterior distribution. After convergence analysis, the posterior palaeodose based on the 13 aliquots over all preheat temperatures has mean 1793 and standard deviation 29 mGy. If the 2 aliquots which have been preheated to 200°C are not included in the palaeodose evaluation, the posterior distribution has mean 1871 and standard deviation 28 mGy.

These posterior distributions for palaeodose, conditional on the interval in which the preheat plateau begins, are shown in Figure 8.8. Also presented is the mixture of two normal distributions with mean and standard deviation of the two alternative posterior sample palaeodose distribution. The mixture is weighted with the posterior probabilities of T_a : $N(1793, 29^2)$ with weight 0.987 and $N(1871, 28^2)$ with weight 0.013.

It is clear from Figure 8.8 that the inclusion of the two aliquots with a preheat of 200°C notably impacts the posterior palaeodose distribution. Although the posterior standard deviation is similar for both, the mean of the distribution is lower for the distribution based on the full set of data available. As the probability that the

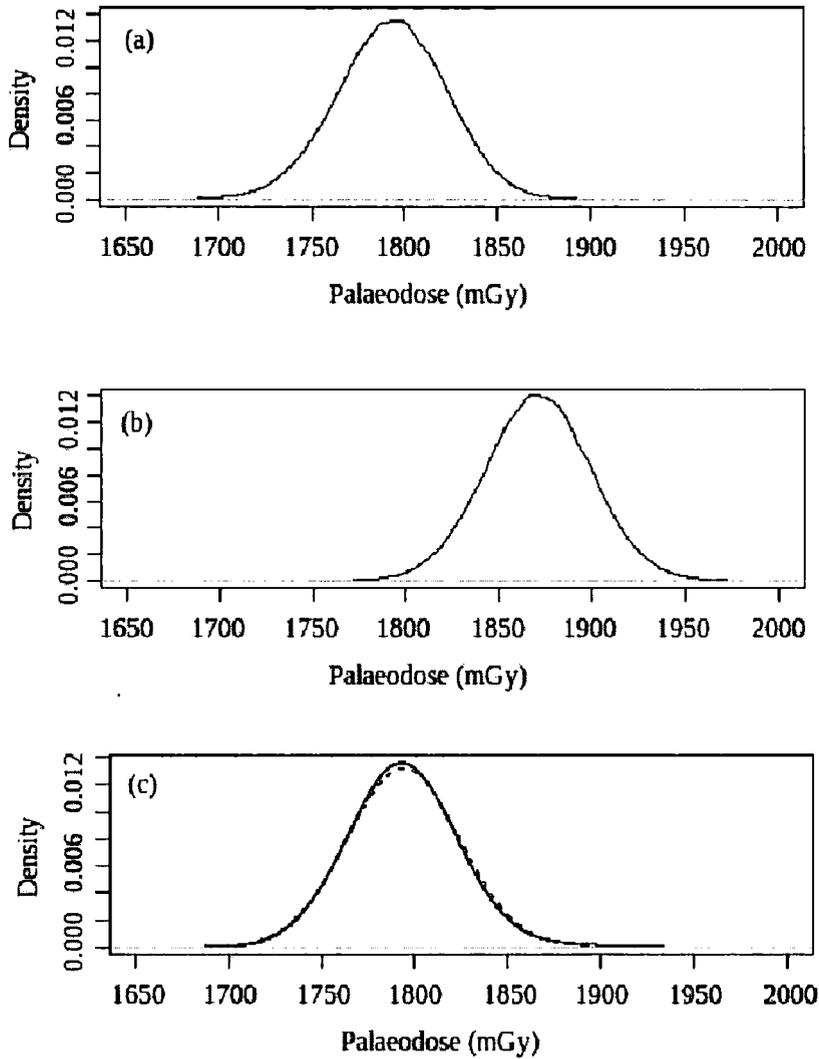


Figure 8.8: Posterior palaeodose distribution of 318-2.

The distributions are based on (a) all aliquots, (b) aliquots with preheat $\geq 210^\circ\text{C}$, (c) mixture of two normals with the same statistics of (a) and (b) weighted by the posterior distribution for plateau starting temperature. The dashed line is the density of a normal distribution with the same mean and variance as the mixture of normals.

preheat plateau for this sample begins above 200°C is small, the mixture of normals weighted using these probabilities is similar to the palaeodose distribution from the full aliquot set. Similarly, the normal distribution with the same statistics as the mixture (mean 1794, standard deviation 29 mGy) is close to both the mixture distribution and the posterior palaeodose distribution assuming the plateau has begun by 200°C.

The higher preheat temperature of 240°C results in a lower signal strength, and thus there is more dispersion through the aliquot estimates. The palaeodose was computed without the data from aliquots at a preheat of 240°C, and the posterior distribution is shown in Figure 8.9. This distribution has a mean of 1806.9 and a standard deviation of 22.9 mGy, which is notably smaller than the posterior standard deviations in Figure 8.8.

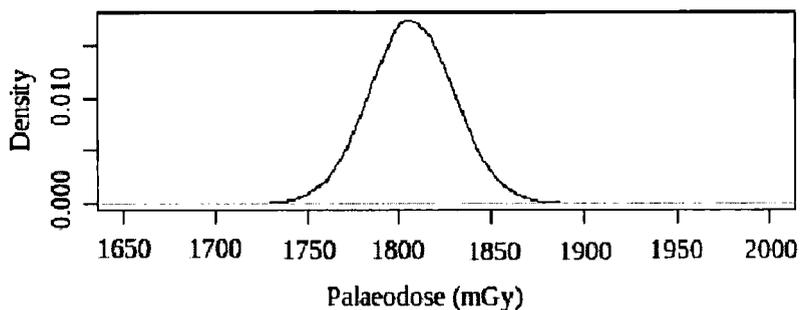


Figure 8.9: Posterior palaeodose distribution for 318-2 based on aliquots with preheat treatments < 240°C.

As the model for the relationship between palaeodose and preheat temperature off the plateau is not robust, further experimental measurements at low preheat temperatures could be useful here. This would provide more data on which to base judgements on where the preheat plateau begins, and as such which aliquot to include in the evaluation of the sample palaeodose.

Parameter	Distribution
\dot{D}_β	$N(2.61, ((0.025)(2.61))^2)$
\dot{D}_γ	$N(0.85, ((0.025)(0.85))^2)$
\dot{D}_c	$N(0.1, ((0.025)(0.1))^2)$
W	$N(1, ((0.025)(1))^2)$
b	$N(0.92, 0.05^2)$
H_β	$N(1.25, 0.1^2)$
G	$N(1.09, 0.01^2)$
F	$\beta(72.27, 2351.25)$

Table 8.9: Distributions assigned to parameters to evaluate the dose rate for sample 318-2, based on prior judgements.

8.4 Dose Rate

Rather than using the model for dose rate set out in Chapter 5, here we will adopt the adapted version of this model which was used during the conventional analysis of this sample [14]. Here the dose rate \dot{D} is computed using

$$\dot{D} = \frac{b}{1 + H_\beta W F} \dot{D}_\beta + G \dot{D}_\gamma + \dot{D}_c \quad (8.4)$$

where the coefficients of \dot{D}_β are the same as in Chapter 5.

The gamma dose rate was measured in situ using γ -TLD (Section 2.4.1), and the parameter G here corrects for the attenuation of the gamma radiation by the dosimeter capsule wall [14]. The value of G was estimated using Monte Carlo simulations of gamma radiation transport, and it was concluded that a correction of 9% increase in \dot{D}_γ was appropriate. The parameter G is equivalent to the coefficient $\frac{g}{1+H_\gamma W F}$ used in Chapter 5, and here was assigned a distribution of $N(1.09, 0.01^2)$.

The distribution for dose rate achieved after 20,000 iterations using the distributions given in Table 8.9 is illustrated in Figure 8.10. This distribution has mean 3.34 and standard deviation 0.16. This analysis assumes a uniform beta dose rate.

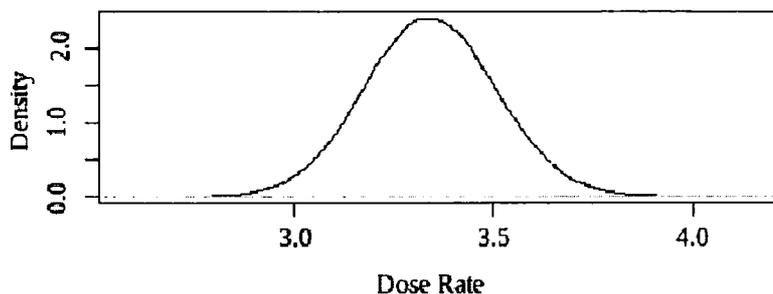


Figure 8.10: Posterior dose rate distribution for sample 318-2.

8.5 Age Ratio

In luminescence dating, the age of a sample is estimated using

$$A_E = \frac{\text{Palaeodose}}{\text{Dose Rate}} \quad (8.5)$$

and this is the final step in routine age analysis. Here we estimate this age ratio by combining draws from the sample palaeodose distribution with ones from the distribution for dose rate. For sample 318-2, the resulting distribution is shown in Figure 8.11. This distribution has mean 538.1, standard deviation 28.0 years, and is overlaid with the density of a normal distribution with the same mean and standard deviation (dashed line).

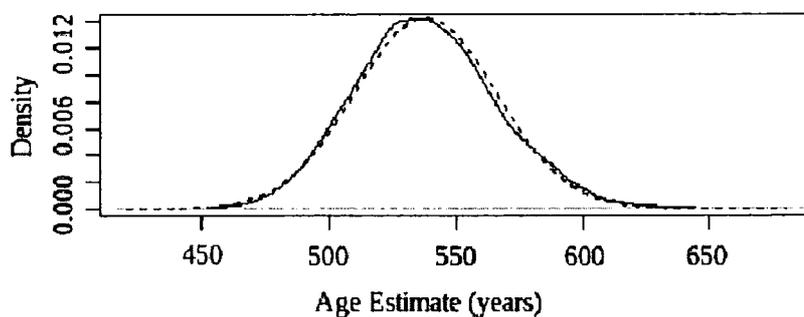


Figure 8.11: Posterior age ratio of sample 318-2, with the dashed line giving the density of the normal distribution with the same statistics.

8.6 Age

There is strong documentary evidence about the age of the building, which in part will be ignored to assigned a realistic prior distribution to age reflecting the level more common to prior beliefs in routine dating. The parameters used to compute the posterior distribution for the age of 318-2 are given in Table 8.10.

Data Input		Prior Judgements				
\bar{A}_E	ω_E^2	m_A	σ_A	σ_E	ρ_E	ρ_A
538.1	28.0	558	20	10	0.55	0.2

Table 8.10: Prior parameter values used to calculate posterior sample age using the model outlined in Chapter 6.

The building has been assigned a date range of 1445-1450 based on documentary evidence alone [14]. Here we take the mean age to be 558 years, where the age is considered as years before 2005 (when the laboratory measurements were made). This falls within the known date range, however we will apply a larger prior standard deviation, of 20 years, to the model, which is a more common uncertainty in brick dating.

The parameter σ_E represents the uncertainty in the age ratio A_E as an indicator of the sample age. Here $\sigma_E = 10$, based on expert judgements.

Using the values laid out in Table 8.10, the posterior age distribution is normal with mean 552 years, standard deviation 17. This leads to a date of 1453 ± 17 .

8.6.1 Sensitivity to Prior Specifications

The influence of the prior mean on the posterior mean age is shown in Figure 8.12. Here there is a direct relationship between the prior and posterior mean ages, with the effect being less marked when the prior standard deviation is high.

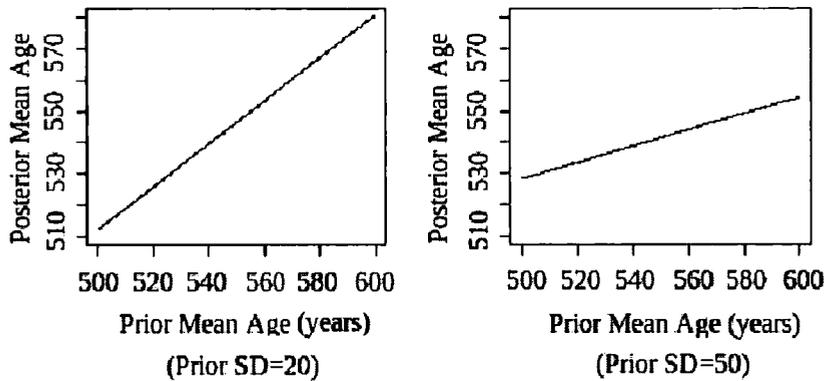


Figure 8.12: Influence of the prior mean on the posterior mean age with a prior age standard deviation of 20 and 50 years.

8.7 Sample 318-1

Sample 318-1 was taken from a lower level of the tower of Tattershall castle [14]. The same prior judgements were made for palaeodose as sample 318-2 (Section 8.1.2). The posterior distribution for palaeodose at each preheat temperature are presented in Figure 8.13, with their statistics given in Table 8.11. The posterior palaeodose standard deviation is smallest for the preheat 220°C as here the palaeodose evaluation is based on 8 aliquots, and so there is more information available than for temperatures 200 and 240°C at which there are only 2 aliquots each.

Preheat	# Aliquots	Posterior Mean (mGy)	Posterior SD
200°C	2	1711.5	89.5
220°C	8	1771.3	30.5
240°C	2	1885.2	66.5

Table 8.11: Posterior palaeodose mean and standard deviation of palaeodose evaluated at each preheat temperature of sample 318-1.

Analogous to routine dating analysis, the posterior palaeodose means were plotted against preheat temperature in Figure 8.14. Although there is an increasing trend in the posterior palaeodose means, given the posterior variances it would be hard to argue that a preheat plateau is not present. This proposition was supported

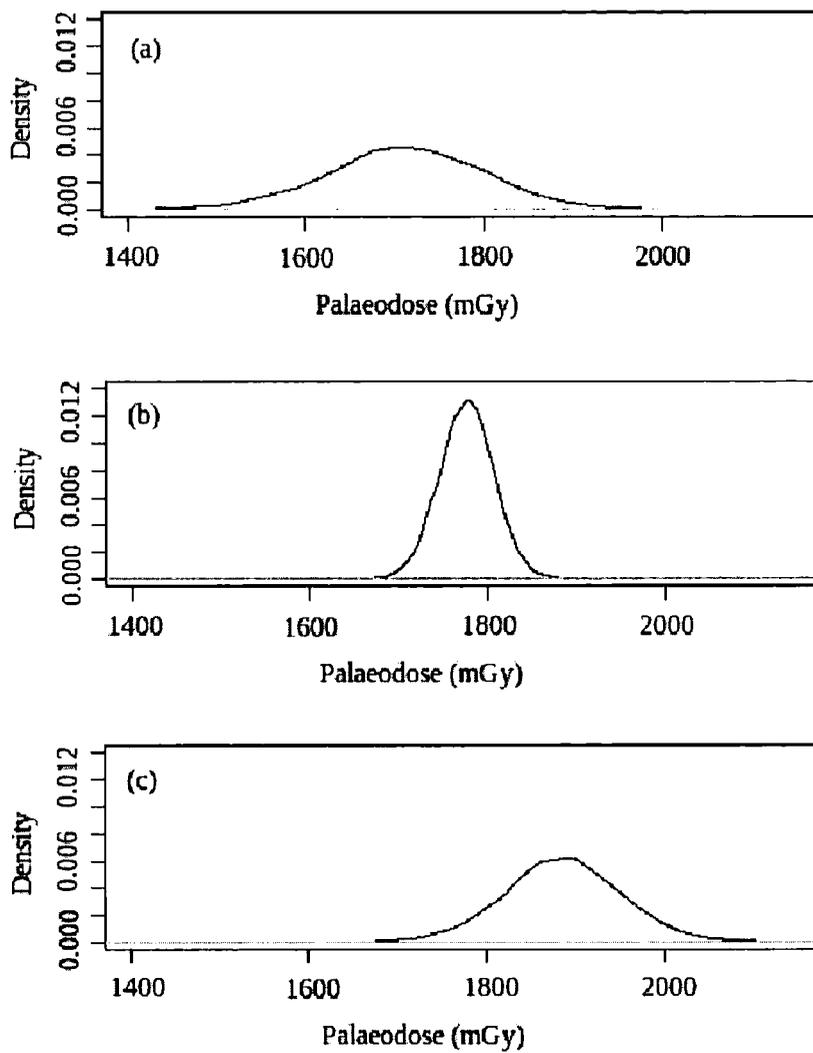


Figure 8.13: Posterior palaeodose distribution at preheat temperature (a) 200°C, (b) 220°C, (c) 240°C.

by the plateau model, that was run with the same prior distribution as elicited for sample 318-2 (Section 8.2.1) to evaluate the posterior distribution for the plateau starting temperature. This posterior distribution assigned a negligible probability to the plateau starting after 200°C under the current prior specifications.

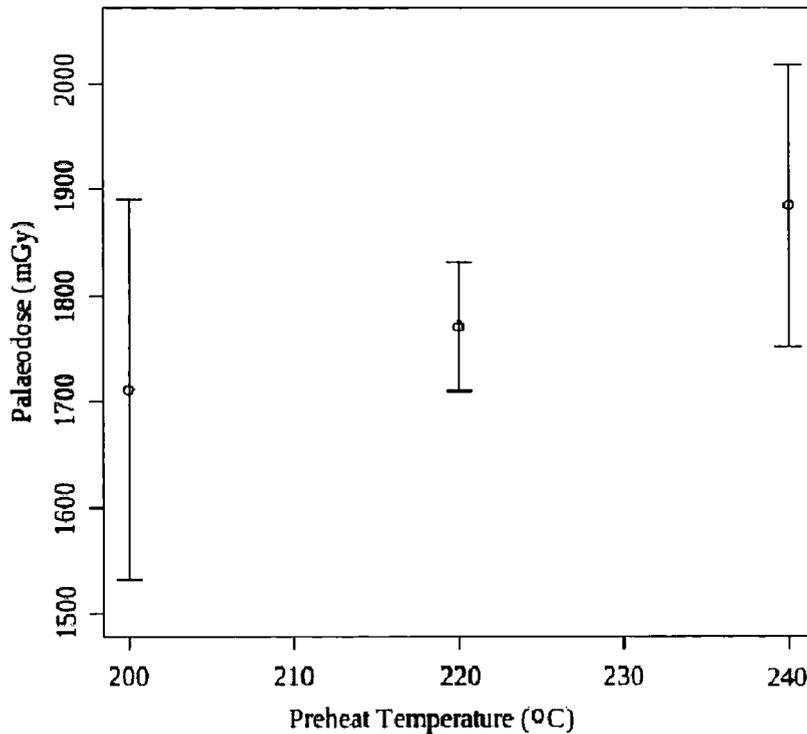


Figure 8.14: Posterior palaeodose means plotted against preheat temperature for sample 318-1, with two standard deviation uncertainty bars.

As all the aliquots are thought to lie on the preheat plateau, all aliquots of sample 318-1 were used to evaluate the sample palaeodose. The combined aliquot model was used with all 12 aliquots and the same prior distributions as when the palaeodose was evaluated at each preheat temperature (and hence the same as for sample 318-2). The sampler was run for 40,000 iterations over 5 chains. After appropriate convergence checks, the chains were thinned every 10th iteration, and a burn-in of 1000 was used. The resulting posterior distribution for the palaeodose of sample 318-1 is illustrated in Figure 8.15. This distribution has mean 1783.0 and standard deviation 24.0 mGy.

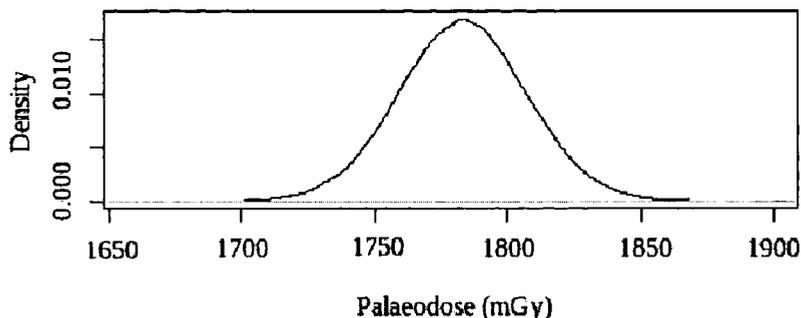


Figure 8.15: Posterior distribution for sample palaeodose for 318-1.

The next step in the age analysis is to evaluate the dose rate. Here, the same model for dose rate was used as for sample 318-2 above. Sample 318-1 was extracted from a brick in the basement of Tattershall Castle, in contrast to 318-2 which was sampled from ground level. The location of 318-1 was relatively damp, based on contemporary measurements, and as such the mean of the average fraction of saturation F was taken to be $5 \pm 1\%$ (compared with $3 \pm 1\%$ for 318-1). With $\dot{D}_\beta = 2.42$ and $\dot{D}_\gamma = 0.97$, but otherwise the same dose rate parameters as 318-2 (Table 8.9), the resulting dose rate distribution for sample 318-1 is shown in Figure 8.16. This distribution has mean 3.52 and standard deviation 0.16.

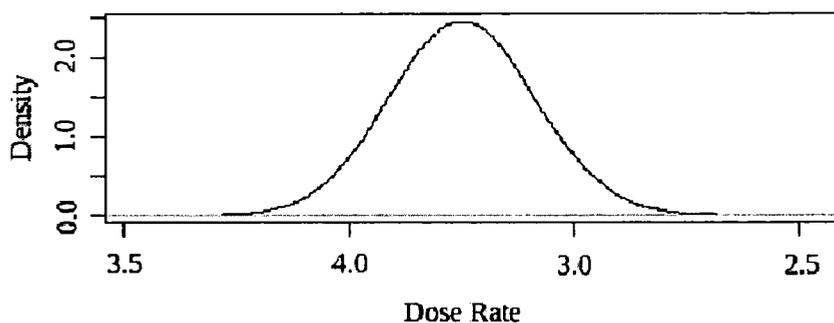


Figure 8.16: Dose rate distribution for 318-1.

The age ratio of 318-1 was then evaluated using the posterior distribution for sample palaeodose and the dose rate distribution. The distribution for the age ratio is shown in Figure 8.17, along with a normal density with the same statistics ($A_E | \bar{A}_E$). The similarity between these two densities suggests that modelling the

age ratio in this way (Section 6.1) is not unsuitable. Here $\bar{A}_E = 549.6$ and $\omega = 28.4$, the mean and standard deviation of the age ratio.

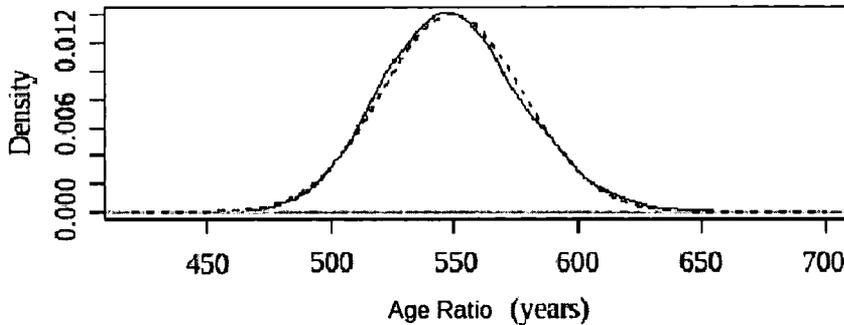


Figure 8.17: Age ratio of sample 318-1, with the dashed line giving the density of the normal distribution with the same statistics.

To evaluate the posterior age distribution, the same prior distribution for age was assigned as sample 318-2. Both these samples have been taken from the same building, and there is no evidence that they are different ages. So, with a priori $A \sim N(558, 20^2)$, and $\sigma_E = 5$, the posterior normal distribution for age for 318-1 has mean 555.3 and standard deviation 16.4 years, leading to a date of 1450 ± 16 .

8.8 Age Analysis

We compare the dates achieved for Tattershall Castle by applying different models for the data. The two samples can be considered independently, where the posterior age is evaluated for both individually. The coeval model is appropriate here, as there is documentary evidence that indicates the base of the tower of Tattershall Castle was built from one stock of bricks, and so the two samples have come from bricks manufactured in the same year. A further comparison was made with the dates achieved through applying the similar age model.

Table 8.12 shows the dates achieved under each model, and Figure 8.18 shows the posterior age densities in each case. As the same data input and analogous prior distributions are used in each model, the posterior age distributions are all similar.

Model	Prior	Date
Single Age (318-1)	$N(558, 20^2)$	1450 ± 16
Single Age (318-2)	$N(558, 20^2)$	1453 ± 17
Coeval Model	$N(558, 20^2)$	1453 ± 16
Similar Age Model	$N \left(\begin{pmatrix} 558 \\ 558 \end{pmatrix}, \begin{pmatrix} 20^2 & (0.8)20^2 \\ (0.8)20^2 & 20^2 \end{pmatrix} \right)$	1451 ± 16 (318-1) 1453 ± 16 (318-2)

Table 8.12: Comparison of different age models for 318-1, 318-2.

When the posterior age distribution is evaluated for 318-1 and 318-2 independently, the posterior age uncertainty is the same in each case, and the mean age is two years younger for 318-2. This difference is small compared with the posterior variance, and easily accounted for by random errors present in luminescence dating. Any practitioner would willingly conclude that these two samples have the same date.

However, assessing the ages independently ignores a significant piece of prior information: there is strong documentary evidence that suggests that one batch of 322,000 bricks was used to build the tower of Tattershall Castle [14]. The availability of this type of documentary evidence is very rare. If the samples have been taken from bricks in the same batch, then they will have been fired at the same time (or at most within a few weeks of each other). It is therefore appropriate to apply the coeval model to the age estimates here. Using the same prior for age as previously, the date assigned to the two samples is 1453 ± 16 AD.

It is notable that here the posterior standard deviation for age from the coeval model is not smaller than when the two ages were considered individually. As more data are available in the coeval model, a smaller variance would usually be expected. Here the posterior standard deviation is slightly smaller for the coeval model, as it has been rounded up to the nearest year to 16 years, whereas the single age evaluation of 318-2 the posterior standard deviation was rounded down to 16 years. In addition, the spread of two age ratio estimates ($\bar{A}_{E1} = 549.6$, $\bar{A}_{E2} = 538.1$) and the prior mean (558 years) is such that this posterior standard deviation seems reasonable.

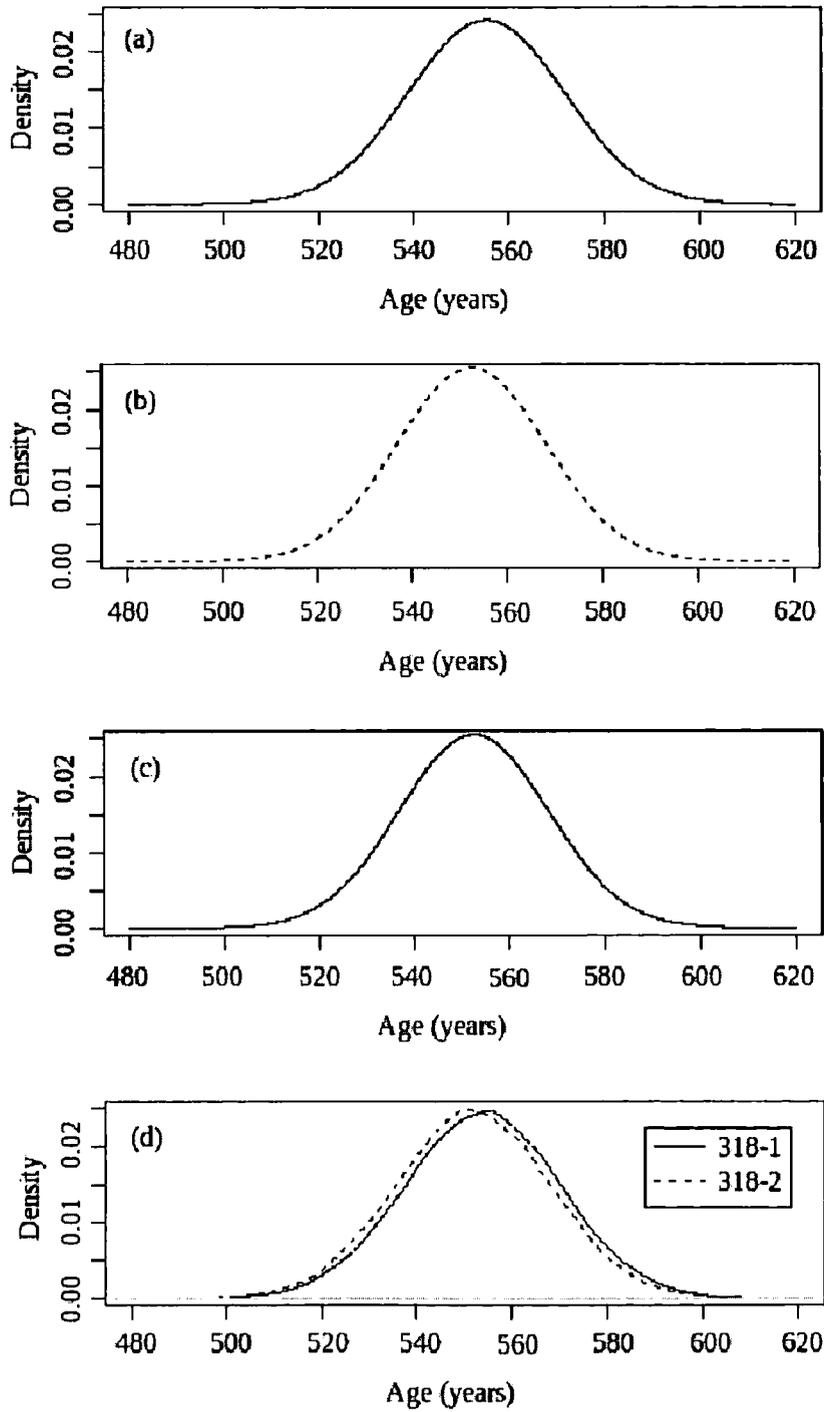


Figure 8.18: Posterior age distribution for (a) 318-1, (b) 318-2, (c) 318-1 and 318-2 under the coeval model, (d) 318-1 and 318-2 under the similar age model.

For comparison, the similar age model was also used to date the samples from Tattershall Castle. The prior assumption that the two samples are related (and the strength of the similarity induced by the prior distribution for age) is not as strong as using the coeval model. Here the same mean and standard deviation for age for both the samples was used, along with a correlation of 0.8. This allows there to be some discrepancy between the two dates. Again, the posterior distributions for the two ages are very similar, both to each other and the posterior densities from the other models.

The dates achieved here (Table 8.12) can be compared with those from the conventional analysis (1455 ± 15 , 1453 ± 15) and also that concluded from the documentary evidence alone (1445-1450). These dates are in agreement with those concluded through the Bayesian analysis, as the strength of the prior information was ignored here to replicate a routine dating situation. However the Bayesian model has scope to include the prior information which is very valuable in situations where the luminescence behaviour induces large errors in the analysis.

Chapter 9

Conclusions

The Bayesian model developed here was designed to establish a computational framework that incorporated the basic elements of a routine age analysis in luminescence dating, based on the application of a single aliquot OSL regenerative dating procedure. That is, the palaeodose evaluation is first carried out at each preheat temperature. Then, these distributions are used to identify the preheat plateau; the aliquots whose palaeodose estimates lie on the plateau are used to evaluate the posterior distribution of the palaeodose of the sample. The dose rate is evaluated using a separate experimental procedure, the distribution of which is combined with the palaeodose to find the age ratio using the age equation. An additional step is added to consider the relationship between the age ratio and the sample age. This allows the expert knowledge of the palaeodose, annual dose, preheat plateau and the uncertainties associated with their measurement to be incorporated in a way that reflects current practice.

By separating the age evaluation into its component parts and considering the posterior distribution at each step, the modularised model is very adaptable and enables individual sections to be tailored to experimental techniques specific to the laboratory without significant restructuring of the model. It can also be modified as experimental methods and theoretical knowledge is further advanced, without having to remodel the whole process.

From a pragmatic viewpoint, with consideration to the clearly defined steps of luminescence and the complexity of the errors within them, this model structure

seems appropriate as it allows the uncertainty resolved at each stage of the analysis to be compared with the empirical reckoning of the practitioner.

The analogous structure of the Bayesian model and routine age analysis increases the potential accessibility of it to luminescence dating practitioners. Although finding the age of the sample is the ultimate aim, the posterior distributions of the parameters at each stage are of significant interest to the experimenter. It also enables palaeodose distributions from different evaluations to be compared.

Considering the age analysis in a number of separate stages results in MCMC computations, in our experience, that are well behaved. As such, they could be implemented by a practitioner who is new to the field of Bayesian statistics without experiencing great convergence problems.

In line with routine dating methods, we have considered the preheat plateau as a diagnostic test for the aliquots which should be included in the evaluation of sample palaeodose, using the posterior palaeodose evaluations at each preheat temperature as the 'data' input to this section of the model. Once the appropriate aliquots have been selected, these are then used in the original 'combined aliquot model' to evaluate the posterior distribution for the palaeodose of the sample. An alternative approach to this problem would have been to construct a large single model which makes inferences about the sample palaeodose from the original data. However, the relationship between all of the parameters is not straightforward and hence it is difficult to write down their joint distributions.

It is recognised that the preheat plateau section of the model is in a preliminary stage of development, reflecting current understanding of luminescence behaviour in this region. Hence this component of the model is most suitable to test for the presence of a preheat plateau, and so indicate which aliquots should be used in the sample palaeodose evaluation. As in routine age analysis, such plateau analysis supports reasonable scientific judgement, rather than being used as a quantitative element in the determination of the palaeodose.

This approach to the structure of the analysis also corresponds well to the general sources of prior information. At each stage of the model, there is prior information available, from past dating experience, environment and, in the case of the palaeo-

dose, preliminary experiments. Considering the inference in a number of separate stages allows appropriate prior information to be utilised at each step and the influence of each prior judgement to be carefully monitored.

Putting practical limitations aside, the development of a model with a single calculation and corresponding complex MCMC calculations could lead to a ‘black-box’ approach to statistics for any luminescence practitioners wishing to apply Bayesian methods, which should be avoided [28]. In contrast, the approach developed in the thesis provides the basis for a transparent age analysis, where each step in the calculation can be easily tracked. It allows Bayesian methods to be accessible to luminescence practitioners with well-behaved MCMC calculations and an analogous method to their routine analysis.

There are several of areas in which this model can be extended to include a greater range of dating situations. As developed here the model is limited to a linear relationship between dose and luminescence. However, it would not be difficult to adapt the model to incorporate the non-linear case.

The examples chosen have been fired brick samples from buildings with good dating control where there is no doubt in the completeness of the resetting process. When sediment samples are considered, this is not always the case and partial resetting often occurs, resulting in a skewed distribution of palaeodose estimates across the sample. The model can be adapted to such situations, for example a log-normal distribution could be assigned to palaeodose a priori.

Single grain methods for palaeodose evaluation are becoming more widespread in luminescence dating. Currently, the model assumes that the SAR procedure has been used, but it could be extended to cover single grain protocols as well (which are usually based on a regenerative procedure). Using aliquots of grains means that any grain-to-grain variation in luminescence brightness is diluted, and so a greater dispersion in palaeodose values are observed in single grain dating, which could be accommodated. Such adaptations are facilitated by the structure of our model, allowing individual stages to be tailored to the particular dating situation.

A barrier that needs to be overcome for Bayesian analysis to reach its full potential in this application is the crossing of inter-disciplinary boundaries. Although the

basic principles of Bayesian statistics are relatively simple, the application to real-life situations are somewhat more complex. The need to have a full understanding of complex mathematical ideas in order to apply a Bayesian model may be somewhat daunting for practitioners of luminescence dating, and indeed the lack of true understanding of Bayesian principles led to some initial opposition in radiocarbon dating [90].

The potential of Bayesian statistics in luminescence dating has been noted [92], and the aim of this thesis was to realise some of this potential. The model has to be accessible to the luminescence dating community as a whole, and we feel that the intuitive approach to the modelling taken here will facilitate this. However, it will still be an area where Bayesian experts are required to work in collaboration with luminescence experts, as is emphasised by advocates of Bayesian methods in radiocarbon dating [98].

9.1 Future Work

The model developed in this thesis is in the preliminary stages of development, and there are a number of areas that would benefit from further work. A number of these are outlined below.

Incomplete resetting of the luminescence signal on deposition can be a common problem when dating sediments. We currently have not considered this in the model, but with work it could be adapted into the model for palaeodose. It is likely that such a development would be of particular interest to many luminescence dating practitioners. In the model here, we consider the palaeodose to follow a normal distribution a priori; this would not be the case when the sample was incompletely bleached. A possible line of research in this area would be to consider Galbraith's models for such situations [47] and adapt them into a Bayesian framework. A straight forward adaption of the current combined aliquot model would be to let the palaeodose have a log normal distribution a priori; that is use the Gibbs sampler detailed in Chapter 3 but consider the log of the palaeodose.

In the combined aliquot model, used to evaluate palaeodose and detailed in

Chapter 3, we consider the relationship between luminescence intensity and applied dose to be linear (after appropriate sensitivity and background signal corrections). This is an approximation to the true saturating exponential relationship, but is appropriate for the samples considered. However, to make the model more widely applicable, this relationship could be adopted.

Another area of the combined aliquot model which has potential for further research is the use of a normal distribution for the errors around the fitted dose response line. The data here is photon counts, and so a poisson distribution would be suitable.

It is noted in Chapter 4 that the model given for the preheat plateau is in the early stages of development. Although a useful diagnostic tool, it would benefit from further work as currently it does not consider the behaviour of palaeodose estimates beyond the plateau, or the possibility that a ‘false plateau’ could be observed.

There is room for development in the dose rate model (Chapter 5), as here we have assumed that the material is homogeneous. Heterogeneity in the beta dose rate can be a source of error here, so it would be beneficial if scope for this, where appropriate, could be included in the model.

At present, we consider the age of the sample to be normal a priori. The model would be more adaptable if it could be applied for a range of different prior distributions for age. We also considered inference from a number of samples, including those which the relative chronology is known. The construction of chronologies using Bayesian methods has been researched and applied in depth for other dating methods, particularly carbon dating [30], and it would be interesting to see if this could be included in the model here for luminescence dating.

Eliciting the expert’s views on the prior parameters is a challenge for all Bayesian practitioners. Further work looking into this process in the context of this model would be very valuable.

As the model develops it would become applicable to a wider number of dating situations and so make it more appealing to dating practitioners. Such developments would be more valuable if the computational side of the model was more user-friendly, with software being written to provide an accessible user interface. If this

work was carried out then the luminescence dating community would be enthusiastic to adopt Bayesian philosophies..

Bibliography

- [1] Adamiec, G., Aitken, M.J., 1998. Dose-rate conversion factors: update. *Ancient TL*, 16, 37-50.
- [2] Aitken, M.J., 1974. *Physics and Archaeology*. Second Edition. Clarendon Press, Oxford.
- [3] Aitken, M.J., 1985. *Thermoluminescence Dating*. Academic Press.
- [4] Aitken, M.J., 1994. Optical dating: a non-specialist review. *Quaternary Science Reviews*, 13, 503-508.
- [5] Aitken, M.J., 1998. *An Introduction to Optical Dating*. Oxford University Press.
- [6] Aitken, M.J., Smith, B.W., 1988. Optical dating: recuperation after bleaching. *Quaternary Science Reviews*, 7, 387-393.
- [7] Aitken, M.J., Xie, J., 1990. Moisture correction for annual gamma dose. *Ancient TL*, 8, 6-9.
- [8] Ballarini, M., Wallinga, J., Wintle, A.G., Bos, A.J.J., 2007. Analysis of equivalent-dose distributions for single grains of quartz from modern deposits. *Quaternary Geochronology*, 2, 77-82.
- [9] Bailey, R.M., Smith, B.W., Rhodes, E.J., 1997. Partial bleaching and the decay form characteristics of quartz OSL. *Radiation Measurements*, 27, 123-136.
- [10] Bailey, S.D., Wintle, A.G., Duller, G.A.T., Bristow, C.S., 2000. Sand deposition during the last millennium at Aberffraw, Anglesey, North Wales as

- determined by OSL dating of quartz. *Quaternary Geochronology (Quaternary Science Reviews)* 20, 701-704.
- [11] Bailey, R.M., 2000. The slow component of quartz OSL. *Radiation Measurements*, 32, 233-246.
- [12] Bailey, R.M., 2001. Towards a general kinetic model for optically and thermally stimulated luminescence of quartz. *Radiation Measurements*, 33, 17-45.
- [13] Bailey, R.M., Arnold, L.J., 2006 Statistical modelling of single grain quartz D_e distributions and an assessment of procedures for estimating burial dose. *Quaternary Science Reviews*, 25, 2475-2502.
- [14] Bailiff, I.K., 2007. Methodological developments in the luminescence dating of brick from English late-Medieval and post-Medieval buildings. *Archaeometry* 49 (4), 827-851.
- [15] Bateman, M.D., Frederick, C.D., Jaiswal, M.K., Singhvi, A.K., 2003. Investigations into the potential effects of pedoturbation on luminescence dating. *Quaternary Science Reviews*, 22, 1169-1176.
- [16] Bateman, M.D., Boulter, C.H., Carr, A.S., Frederick, C.D., Peter, D., Wilder, M., 2007. Detecting post-depositional sediment disturbance in sandy deposits using optical luminescence. *Quaternary Geochronology*, 2, 57-64.
- [17] Bevington, P.R., Robinson, D.K., 2003. *Data reduction and Error Analysis for the Physical Sciences*. McGraw-Hill.
- [18] Bøtter-Jensen, L., Mejdahl, V., 1988. Assessment of beta dose-rate using a GM multiscaler system. *Nuclear Tracks and Radiation Measurements*, 14, 187-191.
- [19] Bøtter-Jensen, L., Duller, G.A.T., Popton, N.R.J., 1994. Excitation and emission spectrometry of stimulated luminescence from quartz and feldspars. *Radiation Measurements*, 23, 613-616.

- [20] Bøtter-Jensen, L., Agersnap Larsen, N., Markey, B.G., McKeever, S.W.S., 1997. Al₂O₃:C as a sensitive OSL dosimeter for rapid assessment of environmental photon dose rates. *Radiation Measurements*, 27, 295-298.
- [21] Bøtter-Jensen, L., Bulur, E., Duller, G.A.T., Murray, A.S., 2000. Advances in luminescence instrument systems. *Radiation Measurements*, 32, 523-528.
- [22] Bøtter-Jensen, L., Andersen, C.E., Duller, G.A.T., Murray, A.S., 2003. Developments in radiation, stimulation and observation facilities in luminescence measurements. *Radiation Measurements*, 37, 535-541.
- [23] Bronk Ramsey, C., 1995. Radiocarbon calibration and analysis of stratigraphy: the OxCal program. *Radiocarbon*, 37, 425-430.
- [24] Bronk Ramsey, C., 1998. Probability and Dating. *Radiocarbon*, 40, 461-474.
- [25] Brooks, S.P., 1998. Quantitative convergence diagnosis for MCMC via CUSUMs. *Statistics and Computing*, 8, 267-274.
- [26] Brown, P.J., 1982. Multivariate Calibration. *Journal of the Royal Statistical Society. Series B (Methodological)*, 44(3), 287-321.
- [27] Brown, P.J., 1993. *Measurement, Regression and Calibration*. Clarendon Press, Oxford.
- [28] Buck, C.E., Cavanagh, W.G., Littan, C.D., 1996. *Bayesian Approach to Interpreting Archaeological Data*. John Wiley and Sons, Chichester.
- [29] Buck, C.E., Christen, J.A., James, G.N., 1999. BCal: an on-line Bayesian radiocarbon calibration tool. *Internet Archaeology*, 7.
- [30] Buck, C.E., Millard, A.R., 2004. (eds.) *Tools for Constructing Chronologies*. Springer Verlag, London.
- [31] Bürgi, A., Flisch, M., 1991. Cosmic ray dose rate determination using a portable gamma-ray spectrometer. *Ancient TL*, 9, 1-5.

- [32] Chen, G., Murray, A.S., Li, S-H., 2001. Effect of heating in the quartz dose-response curve. *Radiation Measurements*, 33, 59-63.
- [33] Daniels, F., Boyd, C.A., Saunders, D.F., 1953. Thermoluminescence as a research tool. *Science*, 117, 343-9.
- [34] Davey, N., 1961. *A History of Building Materials*. Phoenix House, London.
- [35] Duller, G.A.T., 1991. Equivalent dose determination using single aliquots. *Nuclear Tracks and Radiation Measurements*, 18, 371-378.
- [36] Duller, G.A.T., 1997. Behavioural studies of stimulated luminescence from feldspars. *Radiation Measurements*, 27, 663-694.
- [37] Duller, G.A.T., Bøtter-Jensen, L., Murray, A.S., 2000. Optical dating of single sand-sized grains of quartz: sources of variability. *Radiation Measurements*, 32, 453-457.
- [38] Duller, G.A.T., 2004. Luminescence dating of Quaternary sediments: recent advances. *Journal of Quaternary Science* 19, 183-192.
- [39] Duller, G.A.T., 2007. Assessing the error on equivalent dose estimates derived from the single aliquot regenerative dose measurements. *Ancient TL*, 25, 15-24.
- [40] Duller, G.A.T., 2007. *Luminescence dating. Guidelines on using luminescence dating in archaeology*. Produced for English Heritage.
- [41] Dütsch, C., Krbetschek, M.R., 1997. New methods for a better internal ^{40}K dose rate determination. *Radiation Measurements*, 27, 377-381.
- [42] Dunsmore, I.R., 1968. A Bayesian approach to calibration. *Journal of the Royal Statistical Society. Series B (Methodological)*, 30 (2), 396-405.
- [43] El Adlouni, S., Favre, A., Bobée, B., 2006. Comparison of methodologies to assess the convergence of Markov Chain Monte Carlo methods. *Computational Statistics and Data Analysis*, 50, 2685-2701.

- [44] Eisenhart, C., 1939. The interpretation of certain regression methods and their use in biological and industrial research. *Annals of Mathematical Statistics*, 10, 162-186.
- [45] Fattahi, M., Stokes, S., 2003. Red luminescence from potassium feldspar for dating applications: a study of some properties relevant for dating. *Radiation Measurements*, 37, 647-660.
- [46] Feathers, J.K., 2003. Use of luminescence dating in archaeology. *Measurement Science and Technology*, 14, 1493-1509.
- [47] Galbraith, R.F., Roberts, R.G., Laslett, G.M., Yoshid, H., Olley, J.M., 1999. Optical dating of single and multiple grains of quartz from Jinmium rock shelter, northern Australia. Part I: experimental design and statistical models. *Archaeometry* 41, 339-364.
- [48] Galbraith, R., 2002. A note on the variance of a background-corrected OSL count. *Ancient TL*, 20, 49-51.
- [49] Gelman, A., Rubin, D.B., 1992. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457-472.
- [50] Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2003. *Bayesian Data Analysis*. Chapman and Hall.
- [51] Geweke, J., 1992. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. *Bayesian Statistics 4* (J.M. Bernardo, J.O. Berger, A.P. Dawid, A.F.M. Smith, eds.), Oxford University Press, 169-193.
- [52] Gilks, W.R., Richardson, S., Spiegelhalter D.J., 1996. *Markov Chain Monte Carlo in Practice*. Chapman and Hall, London.
- [53] Hinkley, D.V., 1969. On the ratio of two correlated normal random variables. *Biometrika*, 56, 635-639.
- [54] Hunter, W.G., Lamboy, W.F., 1981. A Bayesian analysis of the linear calibration problem. *Technometrics*, 23(4), 323-328.

- [55] Huntley, D.J., Godfrey-Smith, D.I., Thewalt, M.L.W., 1985. Optical dating of sediments. *Nature*, 313, 105-107.
- [56] Hoadley, B., 1970. A Bayesian look at inverse linear regression. *Journal of the American Statistical Association*, 65, 356-369.
- [57] Jacobs, Z., Wintle, A.G., Duller, G.A.T., 2003. Optical dating of dune sand from Blombos Cave, South Africa: I-multiple grain data. *Journal of Human Evolution*, 44, 599-612.
- [58] Jain, M., Murray, A.S., Bøtter-Jensen, L., 2003. Characterisation of blue-light stimulated luminescence components in different quartz samples: implications for dose measurement. *Radiation Measurements*, 37, 441-449.
- [59] Jain, M., Choi, J.H., Thomas, P.J., 2008. The ultrafast OSL component in quartz: Origins and implications. *Radiation Measurements in press*.
- [60] Kacker, R., Toman, B., Huang, D., 2005. Comparison of ISO-GUM, draft GUM supplement 1 and Bayesian statistics using simple linear calibration. *Metrologica*, 43, 167-177.
- [61] Kalchgruber, R., Fuchs, M., Murray, A.S., Wagner, G.A., 2003. Evaluating dose-rate distributions in natural sediments using α -Al₂O₃:C grains. *Radiation Measurements*, 37, 293-297.
- [62] Kiyak, N.G., Polymeris, G.S., Kitis, G., 2007. Component resolved OSL dose response and sensitization of various sedimentary quartz samples. *Radiation Measurements*, 42, 144-155.
- [63] Krutchkoff R.G., 1967. Classical and inverse methods of calibration. *Technometrics*, 9, 425-439.
- [64] Lamothe, M., Balescu, S., Auclair, M., 1994. Natural IRSL intensities and apparent luminescence ages of single feldspar grains extracted from partially bleached sediments. *Radiation Measurements*, 23, 555-562.

- [65] Lamothe, M., Auclair, M., Hamzaoui, C., Huot, S., 2003. Towards a prediction of long-term anomalous fading of feldspar IRSL. *Radiation Measurements*, 37, 493-498.
- [66] Lanos, P., Le Goff, M., Kovacheva, M., Schnepf, E., 2005. Hierarchical modelling of archaeomagnetic data and curve estimation by moving average technique. *Geophysical Journal International*, 160, 440-476.
- [67] Li, B., Li, S., 2006. Comparison of D_e estimates using the fast component and medium component of quartz OSL. *Radiation Measurements*, 41, 125-136.
- [68] Lian, O.B., Roberts, R.G., 2006. Dating the Quaternary: progress in luminescence dating of sediments. *Quaternary Science Reviews*, 25, 2449-2468.
- [69] Madsen, A.T., Murray, A.S., Andersen, T.J., Pejrup, M., 2007. Optical dating of young tidal sediments in the Danish Wadden Sea. *Quaternary Geochronology*, 2, 89-94.
- [70] Marsaglia, G., 1965. Ratios of normal variables and ratios of sums of uniform variables. *Journal of the American Statistical Association*, 60, 193-204.
- [71] McKeever, S.W.S., Chen, R., 1997. Luminescence Models. *Radiation Measurements*, 27, 625-661.
- [72] Millard, A.R., 2006. Bayesian analysis of ESR dates, with application to Border Cave. *Quaternary Geochronology*, 1, 159-166.
- [73] Millard, A.R., 2006. Bayesian analysis of Pleistocene chronometric methods. *Archaeometry*, 48, 359-375.
- [74] Munyikwa, K., 2000. Cosmic ray contribution to environmental dose rates with varying overburden thickness. *Ancient TL*, 18, 27-34.
- [75] Murray, A.S., Olley, J.M., Caitcheon, G.G., 1995. Measurement of equivalent dose in quartz from contemporary water-lain sediments using optically stimulated luminescence. *Quaternary Science Reviews*, 14, 365-371.

- [76] Murray, A.S., Mejdahl, V., 1999. Comparison of regenerative-dose single aliquot and multiple-aliquot (SARA) protocols using heated quartz from archaeological sites. *Quaternary Geochronology*, 18, 223-229.
- [77] Murray, A.S., Roberts, R.G., Wintle, A.G., 1997. Equivalent dose measurement using a single aliquot of quartz. *Radiation Measurements*, 27, 171-184.
- [78] Murray, A.S., Roberts, R.G., 1997. Determining the burial time of single grains of quartz using optically stimulated luminescence. *Earth and Planetary Science Letters*, 152, 163-180.
- [79] Murray, A.S., Wintle, A.G., 1998. Factors controlling the shape of the OSL decay curve in quartz. *Radiation Measurements*, 29, 65-79.
- [80] Murray, A.S., Wintle, A.G., 1999. Isothermal decay of optically stimulated luminescence in quartz. *Radiation Measurements*, 30, 119-125.
- [81] Murray, A.S., Wintle, A.G., 2000. Luminescence Dating of quartz using an improved single-aliquot regenerative-dose protocol. *Radiation Measurements*, 32, 57-73.
- [82] Murray, A.S., Wintle, A.G., 2003. The single-aliquot regenerative dose protocol: potential for improvements in reliability. *Radiation Measurements* 37, 377-381.
- [83] Naes, T., Isaksson, T., Fearn, T., Davies, T., 2002. A user-friendly guide to multivariate calibration and classification. NIR Publications, Chichester.
- [84] Nathan, R.P., Thomas, P.J., Jain, M., Murray, A.S., Rhodes, E.J., 2003. Environmental dose rate heterogeneity of beta radiation and its implications for luminescence dating: Monte Carlo modelling and experimental validation. *Radiation Measurements*, 37, 305-313.
- [85] Nicholls, G., Jones, M., 2001. Radiocarbon dating with temporal order constraints. *Applied Statistics*, 50, 503-521.

- [86] Olley, J.M., Caitcheon, C.G., Roberts, R.G., 1999. The origin of dose distributions in fluvial sediments, and the prospect of dating single grains from fluvial deposits using optically stimulated luminescence. *Radiation Measurements*, 30, 207-217.
- [87] Osbourne, C.S., 1991. Statistical Calibration: A review. *International Statistical Review*, 59(3), 309-336.
- [88] Pevsner, N., Harris, J., revised by Antram, N., 1989. *The Buildings of England, Lincolnshire*. Penguin Books.
- [89] Prescott, J.R., Robertson, G.B., 1997. Sediment dating by luminescence: a review. *Radiation Measurements* 27, 893-822.
- [90] Reece, R., 1994. Are Bayesian statistics useful to archaeological reasoning? *Antiquity*, 68, 848-50.
- [91] Rhodes, E.J., 1988. Methodological considerations in the optical dating of quartz. *Quaternary Science Reviews*, 7, 395-400.
- [92] Rhodes, E.J., Bronk Ramsey, C., Outram, Z., Batt, C., Willis, L., Dockrill, S., Bond, J., 2003. Bayesian methods applied to the interpretation of multiple OSL dates: high precision sediment ages from Old Scatness Broch excavations, Shetland Isles. *Quaternary Science Reviews*, 22, 1231-1244.
- [93] Roberts, R.G., Galbraith, R.F., Yoshida, H., Laslett, G.M., Olley, J.M., 2000. Distinguishing dose populations in sediment mixtures: a test of single-grain optical dating procedures using mixtures of laboratory-dosed quartz. *Radiation Measurements*, 32, 459-465.
- [94] Roberts, R.G., Bird, M., Olley, J.M., Galbraith, R.F., Lawson, E., Laslett, G.M., Yoshida, H., Jones, R., Fullagar, R., Jacobsen, G., Hua, Q., 1998. Optical and radiocarbon dating at Jinmium rock shelter in northern Australia. *Nature*, 393, 353-362.
- [95] Roberts, R.G., Galbraith, R.F., Olley, J.M., Yoshida, H., Laslett, G.M., 1999. Optical dating of single and multiple grains of quartz from Jinmium rock

- shelter, North Australia: Part II, results and implications. *Archaeometry*, 41, 365-395.
- [96] Rice, J.A., 1995. *Mathematical Statistics and Data Analysis*. Duxbury Press.
- [97] Salzman, F.S.A., 1952. *Building in England down to 1540*. Clarendon Press, Oxford.
- [98] Scott, M., 2000. Bayesian methods: what can we gain and at what cost? *Radiocarbon*, 42, 181.
- [99] Singarayer, J.S., Bailey, R.M., 2003. Further investigations of the quartz optically stimulated luminescence components using linear modulation. *Radiation Measurements*, 37, 451-458.
- [100] Singarayer, J.S., Bailey, R.M., Ward, S., Stokes, S., 2005. Assessing the completeness of optical resetting of quartz in the natural environment. *Radiation Measurements*, 40, 13-25.
- [101] Sivia, D.S., Burbridge, C., Roberts, R.G., Bailey, R.M., 2004. A Bayesian approach to the evaluation of equivalent doses in sediment mixtures for luminescence dating. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, AIP Conference Proceedings, 305-311.
- [102] Sjostrand, H., Prescott, J.R., 2002. Thick source alpha counting: the measurement of thorium. *Ancient TL*, 20, 7-10.
- [103] Smith, B.W., Aitken, M.J., Rhodes, E.J., Robinson, P.D., Geldard, D.M., 1986. Optical dating; Methodological aspects. *Radiation Protection Dosimetry*, 17, 229-233.
- [104] Smith, B.W., Rhodes, E.J., 1994. Charge movements in quartz and their relevance to optical dating. *Radiation Measurements*, 23, 329-333.
- [105] Steier, P., Rom, W., 2000. The use of Bayesian statistics for ^{14}C dates of chronologically ordered samples: a critical analysis. *Radiocarbon*, 42, 183-198.

- [106] Stein, C., 1965. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 197-206.
- [107] Stokes, S., 1994. The timing of OSL sensitivity changes in a natural quartz. Radiation Measurements, 23, 601-605.
- [108] Thomsen, K.J., Bøtter-Jensen, L., Denby, P., Moska, P., Murray, A.S., 2006. Developments in luminescence measurement techniques. Radiation Measurements, 41, 768-773.
- [109] Thomsen, K.J., Murray, A.S., Bøtter-Jensen, L., Kinahan, J., 2007. Determination of burial dose in incompletely bleached fluvial samples using single grains of quartz. Radiation Measurements, 42, 307-379.
- [110] Tsukamoto, S., Denby, P.M., Murray, A.S., Bøtter-Jensen, L., 2006. Time-resolved luminescence from feldspars: New insight into fading. Radiation Measurements, 41, 790-795.
- [111] Vandenberghe, D., De Corte, F., Buylaert, J.-P., Kučera, J., Van de Haute, P., 2008. On the internal radioactivity in quartz. Radiation Measurements *in press*.
- [112] Vartanian, E., Guibert, P., Roque, C., Bechtel, F., Schvoerer, M., 2000. Changes in OSL properties of quartz by preheating: an interpretation. Radiation Measurements, 32, 647-652.
- [113] Wallinga, J., Murray, J.S., Duller, G.A.T., Törnqvist, T.E., 2001. Testing optically stimulated luminescence dating of sand-sized quartz and feldspar from fluvial deposits. Earth and Planetary Science Letters, 193, 617-603.
- [114] Wallinga, J., Bos, A.J.J, Duller, G.A.T, 2008. On the separation of quartz OSL signal components using different stimulation modes. Radiation Measurements *in press*.
- [115] Weisberg, S., 1985. Applied Linear Regression, Second Edition. Wiley Series in Probability and Mathematical Statistics.

- [116] Wintle, A.G., 1973. Anomalous fading of thermoluminescence in mineral samples. *Nature*, 245, 143-144.
- [117] Wintle, A.G., Murray, A.S., 1998. Towards the development of a preheat procedure for OSL dating of quartz. *Radiation Measurements*, 29, 81-94.
- [118] Wintle, A.G., Murray, A.S., 1999. Luminescence sensitivity changes in quartz. *Radiation Measurements*, 30, 107-118.
- [119] Wintle, A.G., Murray, A.S., 2000. Quartz OSL: Effects of thermal treatment and their relevance to laboratory dating procedures. *Radiation Measurements*, 32, 387-400.
- [120] Wintle, A.G., Murray, A.S., 2006. A review of optically stimulated luminescence characteristics and their relevance in single-aliquot regeneration dating protocols. *Radiation Measurements*, 41, 369-391.
- [121] Yoshida, H., Roberts, R.G., Olley, J.M., Laslett, G.M., Galbraith, R.F., 2000. Extending the age range of optical dating using single 'supergrains' of quartz. *Radiation Measurements*, 32, 439-446.
- [122] Yu, B., Mykland, P., 1998. Looking at Markov samplers through cusum path plots: a simple diagnostic idea. *Statistics and Computing*, 8, 275-286.

Appendix A

Notation

A.1 Notation for Palaeodose Evaluation

x_{ij}	i^{th} regenerative dose applied in the SAR procedure to aliquot j , $i = 1, \dots, n_j$
\mathcal{Y}_{Rj}	Natural luminescence of aliquot j
y_{Rj}	Standardised natural luminescence of aliquot j
\mathcal{Y}_{ij}	Sensitivity and background corrected luminescence response of aliquot j by regenerative dose i , $i = 1, \dots, n_j$
y_{ij}	Sensitivity and background corrected luminescence response of aliquot j by regenerative dose i , $i = 1, \dots, n_j$, normalised to natural luminescence
n_j	Number of regenerative doses applied to aliquot j
x_R	Palaeodose
x_{Rj}	Palaeodose estimate from aliquot j
J	Number of aliquots used for the palaeodose estimate (either with the same preheat temperature (Chapter 3), or those which lie on the preheat plateau (Chapter 4))
γ_R	Error between aliquot palaeodose estimates x_{Rj} and palaeodose x_R
α, β	Coefficients of linear relationship between luminescence and dose
α_j	Estimate of α from aliquot j
β_j	Estimate of β from aliquot j
ρ	Correlation between α and β , and α_j, β_j
σ^2	Variance of the independent and identically distributed errors of (x_{ij}, y_{ij})

γ_α	Error between α_j and α
γ_β	Error between β_j and β
μ_R, σ_R^2	Prior mean and variance for palaeodose x_R
$m_\alpha, \sigma_\alpha^2$	Prior mean and variance for α
m_β, σ_β^2	Prior mean and variance for β

A.2 Notation for Preheat Plateau Model

x_{RT}	Mean palaeodose at preheat temperature T
x_{R*}	Palaeodose value on the preheat plateau
T_a	Starting temperature of the preheat plateau
η, κ	Curve parameters for the relationship between palaeodose and preheat temperature before the plateau has begun.
\bar{x}_{RT_i}	Mean of the posterior distribution for palaeodose at temperature T_i , $E[x_{RT_i}]$.
σ_{RT_i}	Posterior standard deviation of x_{RT_i} .

A.3 Notation For Dose Rate Model

\dot{D}_β	Beta dose rate
\dot{D}_γ	Gamma dose rate
\dot{D}_c	Cosmic radiation dose rate
W	Water content at saturation
F	Average fractional saturation level
H_β	Water attenuation factor for beta particles
H_γ	Water attenuation factor for gamma rays
h	Attenuation factor in the surrounding material for beta particles
g	Attenuation factor in the surrounding material for gamma rays

A.4 Notation for Age evaluation

- A_E Ratio of palaeodose to annual dose used to estimate age, the age ratio.
- \bar{A}_E Ratio estimate, the posterior mean of the distribution computed for A_E .
- ω_E^2 Variance of the ratio estimate around the age ratio, which takes the value of the variance of the posterior distribution of A_E .
- A Sample age
- σ_E^2 Variance of the age ratio around the true sample age (specified by the expert).

Appendix B

Convergence Diagnostics

A number of diagnostic tools were used to investigate the convergence of the Gibbs sampler used in the combined aliquot model (Chapter 3). Here we provide the details of two of these methods.

B.1 Independence of Starting Values

To look at the convergence of the Gibbs sampler (Section 3.3), starting values for the parameters far from their expected values were chosen. The raw trace plots (Figure B.1) show the first 1000 iterations for x_R , α and β , and indicate that the chains converge quickly from distant starting values. These iterations are based on the simulated data set detailed in Section 3.4.

B.2 Gelman and Rubin Method

For a sampler with m chains, each with n iterations, then for parameter θ let θ_i^t denote the t^{th} iteration from chain i , compute

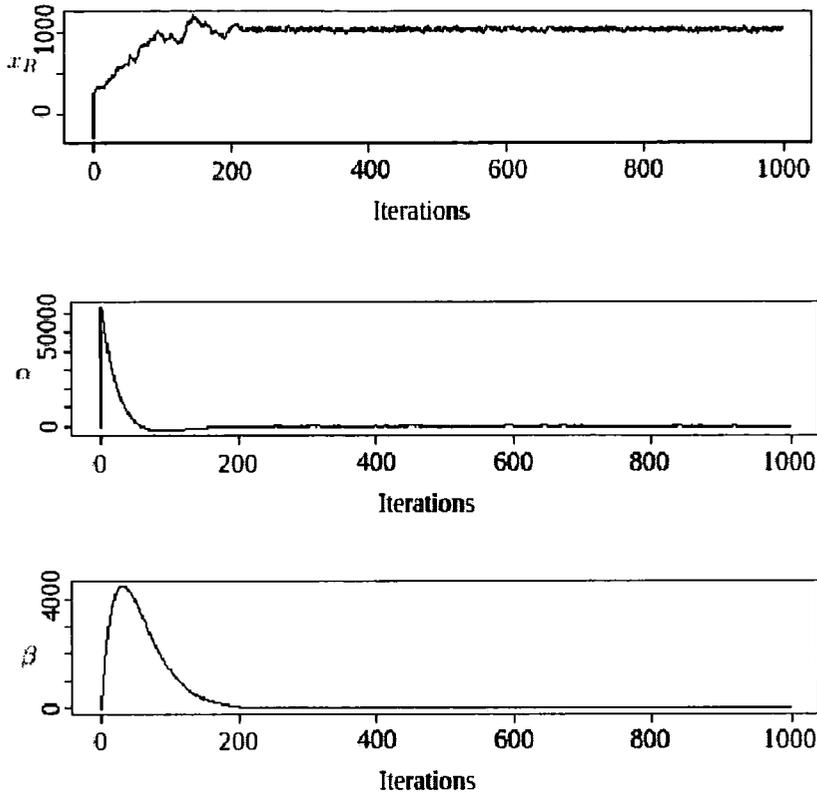


Figure B.1: Raw trace plots from the Gibbs sampler for x_R , α , β , showing convergence independent of the starting values chosen.

$$\bar{\theta}_i = \frac{1}{n} \sum_{t=1}^n \theta_i^t \tag{B.2.1}$$

$$\bar{\theta} = \frac{1}{m} \sum_{i=1}^m \bar{\theta}_i \tag{B.2.2}$$

$$s_i^2 = \frac{1}{n-1} \sum_{t=1}^n (\theta_i^t - \bar{\theta}_i)^2 \tag{B.2.3}$$

The variance between the means of each chain is calculated,

$$\frac{B}{n} = \frac{1}{m-1} \sum_{i=1}^m (\bar{\theta}_i - \bar{\theta})^2 \tag{B.2.4}$$

and the mean of the m variances within each chain,

$$W = \frac{1}{m} \sum_{i=1}^m s_i^2 \tag{B.2.5}$$

The variance across all the simulations is estimated by \hat{V} ,

$$\hat{V} = \frac{n-1}{n}W + \left(1 + \frac{1}{m}\right) \frac{B}{n} \tag{B.2.6}$$

and used to compute the ratio

$$R_c = \left(\frac{d+3}{d+1}\right) \frac{\hat{V}}{W} \tag{B.2.7}$$

where $d = \frac{2\hat{\sigma}^2}{\widehat{Var}[\hat{V}]}$ and

$$\begin{aligned} \widehat{Var}[\hat{V}] = & \left(\frac{n-1}{n}\right)^2 \frac{1}{m} Var[s_i^2] + \left(\frac{m+1}{mn}\right)^2 \frac{2}{m-1} B^2 \\ & + 2 \frac{(m+1)(n-1)}{mn^2} \frac{m}{m} (cov(s_i^2, \bar{\theta}_i^2) - 2\bar{\theta} cov(s_i^2, \bar{\theta}_i)) \end{aligned} \tag{B.2.8}$$

Appendix C

Conditional Posterior

Distributions for the Combined

Aliquot Model

The combined aliquot model is used to evaluate the palaeodose based on a collection of aliquots (Chapter 3). Consider $i = 1, \dots, n_j$ doses applied to $j = 1, \dots, J$ aliquots which have all undergone a preheat treatment at temperature T . The conditional posterior distributions of the combined aliquot model parameters are below, and are used to implement the Gibbs Sampler and thus produce a posterior distribution for palaeodose.

Let Θ be the set of all parameters,

$$\{x_R, \alpha, \beta, \sigma^2, x_{Rj}, \alpha_j, \beta_j \text{ for } j = 1, \dots, J\} \quad (\text{C.0.1})$$

and $\Theta \setminus x_R$, for example, represent this set with x_R removed. Here y_{0j} is the natural luminescence y_{Rj} and $x_{0j} = x_{Rj}$, the j th aliquot palaeodose parameter.

C.1 Conditional Posterior Distribution for Palaeodose

$$\begin{aligned}
P[x_R|D, \Theta \setminus x_R] &\propto P[D|\Theta]P[x_R|\Theta \setminus x_R] \\
&\propto P[D|\Theta]P[x_R|x_{R1}, \dots, x_{RJ}] \\
&\propto \left(\prod_{j=1}^J P[x_{Rj}|x_R] \right) P[x_R] \\
&\propto \exp \left\{ -\frac{1}{2\sigma_R^2} [x_R - \mu_R]^2 \right\} \prod_{j=1}^J \exp \left\{ -\frac{1}{2\gamma_R^2} [x_{Rj} - x_R]^2 \right\} \\
&\propto \exp \left\{ -\frac{1}{2\sigma_R^2} [x_R - \mu_R]^2 \right\} \exp \left\{ -\frac{1}{2\gamma_R^2} \sum_{j=1}^J [x_{Rj} - x_R]^2 \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{1}{\sigma_R^2} + \frac{J}{\gamma_R^2} \right] x_R^2 - 2 \left[\frac{\mu_R}{\sigma_R^2} + \frac{\sum_{j=1}^J x_{Rj}}{\gamma_R^2} \right] x_R \right) \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2} \right) \left[x_R - \frac{\frac{\sum_{j=1}^J x_{Rj}}{\gamma_R^2} + \frac{\mu_R}{\sigma_R^2}}{\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2}} \right]^2 \right\}
\end{aligned}$$

so that

$$x_R|\Theta \setminus x_R, D \sim N \left(\frac{\frac{\sum_{j=1}^J x_{Rj}}{\gamma_R^2} + \frac{\mu_R}{\sigma_R^2}}{\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2}}, \left(\frac{J}{\gamma_R^2} + \frac{1}{\sigma_R^2} \right)^{-1} \right) \quad (\text{C.1.2})$$

C.2 Conditional Posterior Distribution for Aliquot Palaeodose Estimates

$$\begin{aligned}
P[x_{Rj}|D, \Theta \setminus x_{Rj}] &\propto P[D|\Theta]P[x_{Rj}|\Theta \setminus x_{Rj}] \\
&\propto P[D|\Theta]P[x_{Rj}|x_R] \\
&\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \exp \left\{ -\frac{1}{2\gamma_R^2} (x_{Rj} - x_R)^2 \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma_R^2} \right] x_{Rj}^2 - 2 \left[\frac{\beta_j (y_{Rj} \alpha_j)}{\sigma^2} + \frac{x_R}{\gamma_R^2} \right] x_{Rj} \right) \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma_R^2} \right) \left[x_{Rj} - \frac{\frac{\beta_j (y_{Rj} \alpha_j)}{\sigma^2} + \frac{x_R}{\gamma_R^2}}{\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma_R^2}} \right]^2 \right\}
\end{aligned}$$

So

$$x_{Rj}|D, \Theta \setminus x_{Rj} \sim N \left(\frac{\beta_j(y_{Rj} - \alpha_j) + \frac{x_{Rj}}{\gamma_R}}{\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma_R^2}}, \left(\frac{\beta_j^2}{\sigma^2} + \frac{1}{\gamma_R^2} \right)^{-1} \right). \quad (\text{C.2.3})$$

C.3 Conditional Posterior Distributions for the Regression Coefficients

$$\begin{aligned} P[\alpha|D, \Theta \setminus \alpha] &\propto P[D|\Theta]P[\alpha|\Theta \setminus \alpha] \\ &\propto P[D|\Theta]P[\alpha|\beta, \alpha_1, \dots, \alpha_J, \beta_1, \dots, \beta_J] \\ &\propto P[D|\Theta]P[\alpha_1, \dots, \alpha_J|\alpha, \beta, \beta_1, \dots, \beta_J]P[\alpha|\beta, \beta_1, \dots, \beta_J] \\ &\propto P[D|\Theta]P[\alpha_1, \dots, \alpha_J|\alpha, \beta, \beta_1, \dots, \beta_J]P[\alpha|\beta] \\ &\propto P[D|\Theta] \left(\prod_{j=1}^J (P[\alpha_j|\alpha, \beta, \beta_j]) \right) P[\alpha|\beta] \\ &\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \\ &\quad \prod_{j=1}^J \left[\exp \left\{ -\frac{1}{2\gamma_\alpha^2(1-\rho^2)} \left(\alpha_j - \alpha - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta} \right)^2 \right\} \right] \\ &\quad \exp \left\{ -\frac{1}{2\sigma_\alpha^2(1-\rho^2)} \left(\alpha - m_\alpha - \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta} \right)^2 \right\} \\ &\propto \exp \left\{ -\frac{1}{2\gamma_\alpha^2(1-\rho^2)} \left(\sum_{j=1}^J \left(\alpha_j - \alpha - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta} \right) \right)^2 \right\} \\ &\quad \exp \left\{ -\frac{1}{2\sigma_\alpha^2(1-\rho^2)} \left(\alpha - m_\alpha - \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta} \right)^2 \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{1}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)} \right] \alpha^2 \right. \right. \\ &\quad \left. \left. - 2 \left[\frac{\sum_{j=1}^J \left(\alpha_j - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta} \right) + \frac{m_\alpha - \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta}}{\sigma_\alpha^2(1-\rho^2)}}{\gamma_\alpha^2(1-\rho^2)} \right] \alpha \right) \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)} \right) \right. \\ &\quad \left. \left[\alpha - \frac{\frac{\sum_{j=1}^J \left(\alpha_j - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta} \right) + \left(m_\alpha + \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta} \right)}{\gamma_\alpha^2(1-\rho^2)} + \frac{m_\alpha + \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta}}{\sigma_\alpha^2(1-\rho^2)}}{\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)}} \right]^2 \right\} \end{aligned}$$

C.3. Conditional Posterior Distributions for the Regression Coefficients

so that

$$\alpha|D, \Theta \setminus \alpha \sim N \left(\frac{\frac{\sum_{j=1}^J \left(\alpha_j - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta} \right)}{\gamma_\alpha^2(1-\rho^2)} + \frac{\left(m_\alpha + \rho(\beta - m_\beta) \frac{\sigma_\alpha}{\sigma_\beta} \right)}{\sigma_\alpha^2(1-\rho^2)}}{\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)}}, \left(\frac{J}{\gamma_\alpha^2(1-\rho^2)} + \frac{1}{\sigma_\alpha^2(1-\rho^2)} \right)^{-1} \right) \quad (\text{C.3.4})$$

$$\begin{aligned} P[\alpha_j|D, \Theta \setminus \alpha_j] &\propto P[D|\Theta]P[\alpha_j|\Theta \setminus \alpha_j] \\ &\propto P[D|\Theta]P[\alpha_j|\alpha, \beta, \beta_j] \\ &\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \\ &\quad \exp \left\{ -\frac{1}{2\gamma_\alpha^2(1-\rho^2)} \left(\alpha_j - \alpha - \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta} \right)^2 \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{n_j+1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)} \right] \alpha_j^2 \right. \right. \\ &\quad \left. \left. - 2 \left[\frac{\sum_{i=1}^{n_j} (y_{ij} - \beta_j x_{ij})}{\sigma^2} + \frac{\alpha + \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta}}{\gamma_\alpha^2(1-\rho^2)} \right] \alpha_j \right) \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\frac{n_j+1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)} \right) \left[\alpha_j - \frac{\frac{\sum_{i=1}^{n_j} (y_{ij} - \beta_j x_{ij})}{\sigma^2} + \frac{\alpha + \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta}}{\lambda^2(1-\rho^2)}}{\frac{n_j+1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)}} \right]^2 \right\} \end{aligned}$$

leading to

$$\alpha_j|D, \Theta \setminus \alpha_j \sim N \left(\frac{\frac{\sum_{i=1}^{n_j} (y_{ij} - \beta_j x_{ij})}{\sigma^2} + \frac{\alpha + \rho(\beta_j - \beta) \frac{\gamma_\alpha}{\gamma_\beta}}{\lambda^2(1-\rho^2)}}{\frac{n_j+1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)}}, \left(\frac{n_j+1}{\sigma^2} + \frac{1}{\gamma_\alpha^2(1-\rho^2)} \right)^{-1} \right) \quad (\text{C.3.5})$$

$$\begin{aligned}
 P[\beta|D, \Theta \setminus \beta] &\propto P[D|\Theta]P[\beta|\Theta \setminus \beta] \\
 &\propto P[D|\Theta]P[\beta|\alpha, \alpha_1, \dots, \alpha_J, \beta_1, \dots, \beta_J] \\
 &\propto P[D|\Theta]P[\beta_1, \dots, \beta_J|\beta, \alpha, \alpha_1, \dots, \alpha_J]P[\beta|\alpha, \alpha_1, \dots, \alpha_J] \\
 &\propto P[D|\Theta]P[\beta_1, \dots, \beta_J|\beta, \alpha, \alpha_1, \dots, \alpha_J]P[\beta|\alpha] \\
 &\propto P[D|\Theta] \left(\prod_{j=1}^J P[\beta_j|\alpha, \beta, \beta_j] \right) P[\beta|\alpha] \\
 &\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \\
 &\quad \prod_{j=1}^J \exp \left\{ -\frac{1}{2\gamma_\beta^2} \left(\beta_j - \beta - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha} \right)^2 \right\} \\
 &\quad \exp \left\{ -\frac{1}{2\sigma_\beta^2} \left(\beta - m_\beta - \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha} \right)^2 \right\} \\
 &\propto \exp \left\{ -\frac{1}{2\gamma_\beta^2} \sum_{j=1}^J \left(\beta_j - \beta - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha} \right)^2 \right\} \\
 &\quad \exp \left\{ -\frac{1}{2\sigma_\beta^2} \left(\beta - m_\beta - \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha} \right)^2 \right\} \\
 &\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)} \right] \beta^2 \right. \right. \\
 &\quad \left. \left. - 2 \left[\frac{\sum_{j=1}^J (\beta_j - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha})}{\gamma_\beta^2(1-\rho^2)} + \frac{m_\beta + \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha}}{\sigma_\beta^2(1-\rho^2)} \right] \beta \right) \right\} \\
 &\propto \exp \left\{ -\frac{1}{2} \left(\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)} \right) \right. \\
 &\quad \left. \left[\beta - \frac{\frac{\sum_{j=1}^J (\beta_j - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha})}{\gamma_\beta^2(1-\rho^2)} + \frac{(m_\beta + \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha})}{\sigma_\beta^2(1-\rho^2)}}{\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)}} \right]^2 \right\}
 \end{aligned}$$

so

$$\beta|D, \Theta \setminus \beta \sim N \left(\frac{\frac{\sum_{j=1}^J (\beta_j - \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha})}{\gamma_\beta^2(1-\rho^2)} + \frac{(m_\beta + \rho(\alpha - m_\alpha) \frac{\sigma_\beta}{\sigma_\alpha})}{\sigma_\beta^2(1-\rho^2)}}{\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)}}, \left(\frac{J}{\gamma_\beta^2(1-\rho^2)} + \frac{1}{\sigma_\beta^2(1-\rho^2)} \right)^{-1} \right) \quad (\text{C.3.6})$$

$$\begin{aligned}
P[\beta_j|D, \Theta \setminus \beta_j] &\propto P[D|\Theta]P[\beta_j|\Theta \setminus \beta_j] \\
&\propto P[D|\Theta]P[\beta_j|\alpha_j, \alpha, \beta] \\
&\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \\
&\quad \exp \left\{ -\frac{1}{2\gamma_\beta^2(1-\rho^2)} \left(\beta_j - \beta - \rho(\alpha_j - \alpha) \frac{\gamma_\alpha}{\gamma_\beta} \right)^2 \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\left[\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)} \right] \beta_j^2 \right. \right. \\
&\quad \left. \left. - 2 \left[\frac{\sum_{i=1}^{n_j} (x_{ij} y_{ij} - \alpha_j x_{ij})}{\sigma^2} + \frac{\beta + \rho(\alpha - \alpha_j) \frac{\gamma_\beta}{\gamma_\alpha}}{\gamma_\beta^2(1-\rho^2)} \right] \beta_j \right) \right\} \\
&\propto \exp \left\{ -\frac{1}{2} \left(\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)} \right) \right. \\
&\quad \left. \left[\beta_j - \frac{\frac{\sum_{i=1}^{n_j} (x_{ij} y_{ij} - \alpha_j x_{ij})}{\sigma^2} + \frac{\beta + \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha}}{\gamma_\beta^2(1-\rho^2)}}{\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)}} \right]^2 \right\} \\
\beta_j|D, \Theta \setminus \beta_j &\sim N \left(\frac{\frac{\sum_{i=1}^{n_j} (x_{ij} y_{ij} - \alpha_j x_{ij})}{\sigma^2} + \frac{\beta + \rho(\alpha_j - \alpha) \frac{\gamma_\beta}{\gamma_\alpha}}{\gamma_\beta^2(1-\rho^2)}}{\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)}}, \left(\frac{\sum_{i=1}^{n_j} x_{ij}^2}{\sigma^2} + \frac{1}{\gamma_\beta^2(1-\rho^2)} \right)^{-1} \right)
\end{aligned} \tag{C.3.7}$$

C.4 Conditional Posterior Distribution for σ^2

$$\begin{aligned}
P[\sigma^2|D, \Theta \setminus \sigma^2] &\propto P[D|\Theta]P[\sigma^2|\Theta \setminus \sigma^2] \\
&\propto P[D|\Theta]P[\sigma^2] \\
&\propto \frac{1}{(2\pi\sigma^2)^{(\sum_{j=1}^J n_j)/2}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 \right\} \\
&\quad \frac{(a/2)^{d/2}}{\Gamma(d/2)} \frac{1}{\sigma^{2(d/2+1)}} \exp \left\{ -\frac{a}{2\sigma^2} \right\} \\
&\propto \frac{1}{(\sigma^2)^{[(d+\sum_{j=1}^J n_j)/2-1]}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{j=1}^J \sum_{i=0}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 + \frac{a}{2} \right\}
\end{aligned}$$

$$\sigma^2|D, \Theta \setminus \sigma^2 \sim \text{Inv}\Gamma\left(\frac{d + \sum_{j=1}^J n_j}{2}, \frac{1}{2} \left(\sum_{j=1}^J \sum_{i=1}^{n_j} (y_{ij} - \alpha_j - \beta_j x_{ij})^2 + a \right)\right) \quad (\text{C.4.8})$$

C.5 Conditional Posterior Distribution for γ_R^2

If the parameter γ_R^2 is considered to be unknown (rather than a fixed constant), then it can be added to the Gibbs Sampler. If a priori $\gamma_R^2 \sim \text{Inv}\Gamma(b/2, c/2)$ then its conditional posterior distribution is

$$\begin{aligned} P[\gamma_R^2|D, \Theta] &\propto P[D|\Theta, \gamma_R^2]P[\gamma_R^2|\Theta] \\ &\propto P[D|\Theta, \gamma_R^2]P[\gamma_R^2|x_R, x_{R1}, \dots, x_{RJ}] \\ &\propto \left(\prod_{j=1}^J P[x_{Rj}|x_R, \gamma_R^2] \right) P[\gamma_R^2] \\ &\propto \left(\prod_{j=1}^J \frac{1}{\sqrt{2\pi\gamma_R^2}} \exp\left\{-\frac{1}{2\gamma_R^2} (x_{Rj} - x_R)^2\right\} \right) \\ &\quad \frac{(c/2)^{(b/2)}}{\Gamma(b/2)} \frac{1}{\gamma_R^{2(b/2+1)}} \exp\left\{-\frac{c}{2\gamma_R^2}\right\} \\ &\propto \frac{1}{(\gamma_R^2)^{[(J+b)/2+1]}} \exp\left\{-\frac{1}{2\gamma_R^2} \sum_{j=1}^J (x_{Rj} - x_R)^2 + \frac{c}{2}\right\} \\ \gamma_R^2|D, \Theta &\sim \text{Inv}\Gamma\left(\frac{J+b}{2}, \frac{1}{2} \left(\sum_{j=1}^J (x_{Rj} - x_R)^2 + c \right)\right) \end{aligned} \quad (\text{C.5.9})$$

Appendix D

Simulated Data used to Test Convergence and Stability of the Gibbs Sampler

A data set was simulated to investigate the convergence properties of the Gibbs Sampler used in the combined aliquot model (Chapter 3). The true values of the parameters were chosen (given below), and then these were used to simulate a data set. This data was used in the combined aliquot model to see if the parameter values were returned. The code used to simulate the data, along with the data set achieved, follows. The table below shows the values chosen for the parameters, on which the data was based.

x_R	γ_R	α	γ_α	β	γ_β
500	5	0	5	10	5

R Code to simulate a data set

```
library(MASS)
```

```
J<-5 #Number of aliquots
```

```
xR<-500 #Palaeodose
```

```
gamR<-5
xRj<-c(rnorm(J,xR,gam)) #Simulate the aliquot palaeodose values
yRj<-c(5000,5000,5000)
x<-c(250,250,400,650,650)
alpha<-0
beta<-10
gambet<-5
gamalp<-5
rho<--0.3
aljbetj<-array(mvrnorm(J,c(alpha,beta),
matrix(c(gamalp^2,rho*gamalp*gambet,rho*gamalp*gambet,gambet^2),2,2)),c(J,2))

mtau<-0.01
vartau<-0.00015
a<-mtau^2/vartau
b<-mtau/vartau
tau<-rgamma(length(x)*J,a,b)
sig<-1/sqrt(tau)

ep<-array(rnorm(length(x)*J,0,sig),c(J,length(x)))

y<-array(NA,c(J,length(x)))
for (i in 1:length(x)){
for (j in 1:J){
y[j,i]<-aljbetj[j,1]+aljbetj[j,2]*x[i]+ep[j,i]
}
}

#Use the known palaeodose and the fitted line to find the natural
luminescence values.
```

```

yRj<-aljbetj[,1]+aljbetj[,2]*xRj

#Standardise the luminescence response against a natural luminescence
  value of 5000 counts.
stany<-array(NA,c(J,length(x)))
for (i in 1:length(x)){
for (j in 1:J){
stany[j,i]<-y[j,i]*5000/yRj[j]
}
}

```

The simulated data set used in Section 3.4 was produced using the code above and is shown in Table D.

Applied Dose (mGy)				
250	250	400	650	650
2488.3	2486.3	3987.2	6494.6	6461.5
2547.1	2482.7	4033.9	6559.5	6575.7
2477.5	2474.4	3956.7	6423.8	6420.5

Table D.1: The simulated luminescence response values for the given doses, standardised to a natural luminescence value of 5000 counts

Appendix E

Sensitivity of the Combined Aliquot Model to Prior Judgements

Section 3.6 looked at the sensitivity of the posterior distributions to the prior judgements made in the combined aliquot model for evaluating palaeodose. Here we provide further details and results of the analysis carried out.

Table E.1 shows the influence of the choice of σ_R , the prior palaeodose standard deviation, when the prior mean μ_R is 900 mGy. Figure E.1 shows a selection of the corresponding posterior distributions for palaeodose. Similarly, Table E.2 illustrates the influence of σ_R on posterior palaeodose when $\mu_R = 1000$ mGy, and Table E.3 and Figure E.2 for $\mu_R = 1100$ mGy.

As discussed in Section 3.6, the sensitivity of the posterior palaeodose distribution to the choice of prior mean is dependent on the choice of prior standard deviation. If the value of σ_R is small, then the posterior palaeodose mean is similar to the prior.

Prior	Posterior			
	σ_R	Mean	SD	Mode
5	902.2	5.0	902.2	6.8
10	909.1	10.3	909.2	13.9
15	922.5	16.1	922.4	21.8
20	943.3	21.4	943.7	29.4
25	967.7	23.0	969.9	31.0
30	984.9	20.9	987.1	27.3
50	1007.5	17.0	1000.8	21.6
100	1017.4	16.7	1017.5	21.2
500	1020.3	16.7	1020.2	21.1

Table E.1: Influence of σ_R on posterior palaeodose when $\mu_R = 900\text{mGy}$

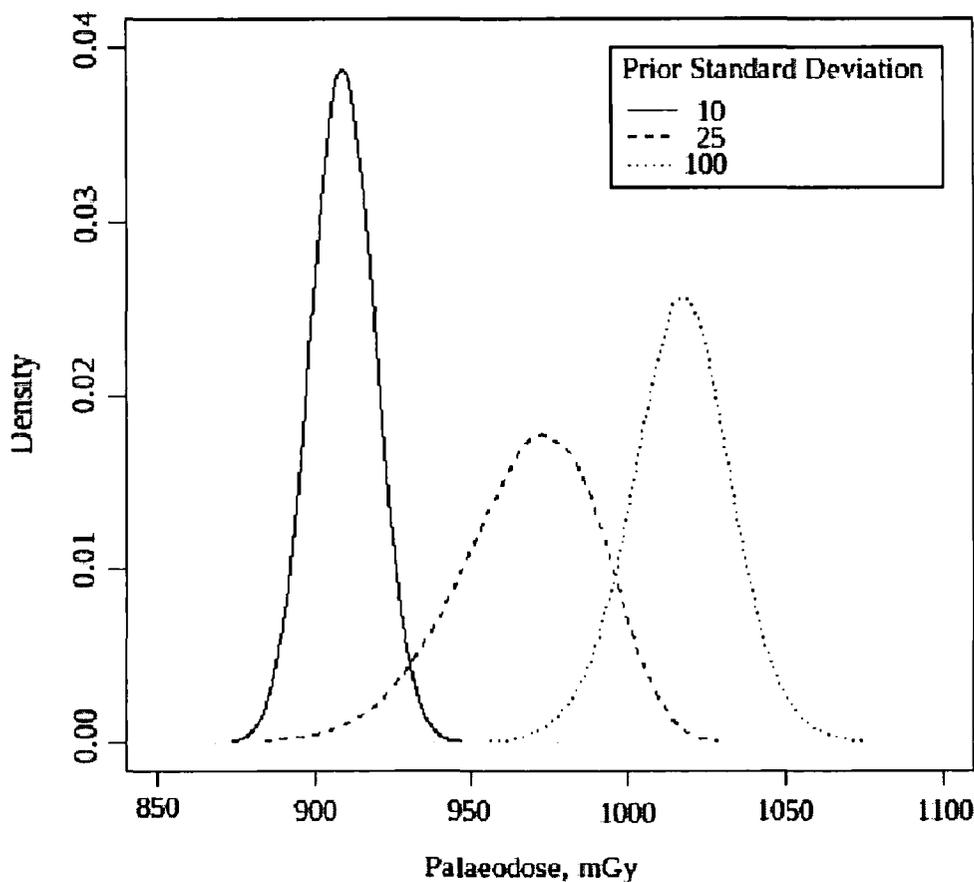


Figure E.1: Influence of σ_R on posterior palaeodose distribution when $\mu_R = 900\text{ mGy}$.

Prior	Posterior			
σ_R	Mean	SD	Mode	IQ range
5	1001.8	4.8	1001.8	6.4
10	1005.9	8.6	1006.0	11.6
15	1009.6	11.3	1009.8	15.0
20	1012.4	12.7	1012.6	16.7
25	1014.5	13.5	1014.7	17.6
30	1015.6	14.5	1016.0	18.9
50	1017.7	16.0	1017.9	20.2
100	1019.8	16.7	1020.0	21.3
500	1020.5	16.5	1020.4	21.0

Table E.2: Influence of σ_R on posterior palaeodose distribution when $\mu_R = 1000$.

Prior	Posterior			
σ_R	Mean	SD	Mode	IQ range
5	1097.4	501	1097.3	69
10	1089.0	10.2	1089.0	13.8
15	1074.6	14.8	1074.3	20.4
20	1058.2	17.0	1057.1	22.7
25	1047.6	17.0	1046.1	22.3
30	1039.6	16.4	1038.6	21.0
50	1028.3	16.4	1027.6	20.6
100	1021.8	16.0	1021.5	20.5
500	1020.2	16.9	1020.0	21.2

Table E.3: Influence of σ_R on posterior palaeodose distribution when $\mu_R = 1100$ mGy.

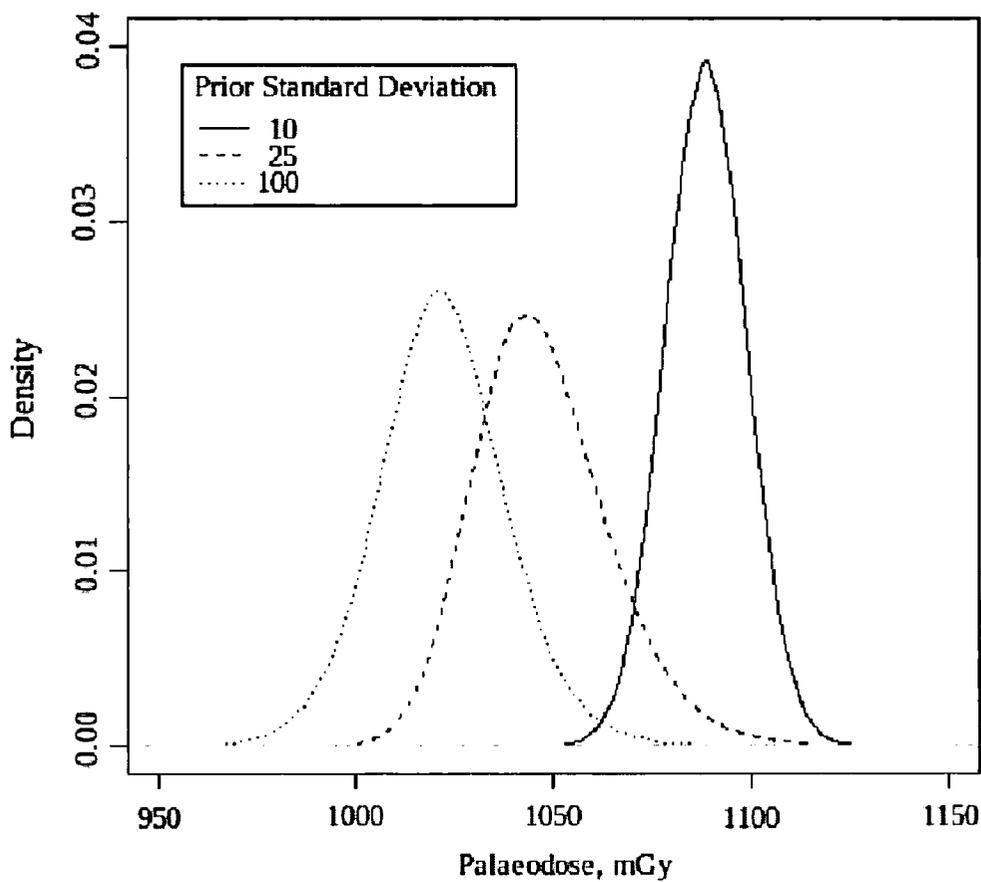


Figure E.2: Influence of σ_R on posterior palaeodose distribution when $\mu_R = 1100$ mGy.

Prior			Posterior	
μ_R	σ_R	γ_R	Mean	SD
900	100	50	1013.8	28.6
900	50	50	991.1	25.9
900	25	50	950.7	19.3
900	10	50	912.2	9.6
1000	100	50	1021.4	28.2
1000	50	50	1017.8	25.4
1000	25	50	1009.9	19.1
1000	10	50	1002.4	9.4
1100	100	50	1029.9	28.7
1100	50	50	1043.7	25.8
1100	25	50	1068.4	19.3
1100	10	50	1092.3	9.5

Table E.4: Influence of μ_R , σ_R on palaeodose when $\gamma_R = 50$.

The influence of γ_R , the measure of spread of the aliquot palaeodose estimates, on the posterior palaeodose distribution was also considered in Section 3.6. Here we provide further details of this investigation.

Table E.4, E.5, E.6 and E.7 show how the prior mean and standard deviation influence the posterior distribution for palaeodose for different values of γ_R . When γ_R is large, this indicates less confidence in the aliquot estimates evaluating the palaeodose, and so the prior mean and standard deviation for palaeodose have more influence. This is particularly true when the value of γ_R is large in comparison to σ_R . The reverse trend is also observed: when γ_R is small in comparison to σ_R , then the prior specifications have less impact on the posterior palaeodose distribution.

Prior			Posterior	
μ_R	σ_R	γ_R	Mean	SD
900	100	25	1019.4	16.6
900	50	25	1010.1	16.2
900	25	25	980.9	16.7
900	10	25	911.2	10.6
1000	100	25	1021.8	16.7
1000	50	25	1020.4	15.8
1000	25	25	1015.7	14.1
1000	10	25	1006.0	8.6
1100	100	25	1024.9	16.5
1100	50	25	103.5	16.0
1100	25	25	1047.0	14.3
1100	10	25	1082.8	9.8

Table E.5: Influence of μ_R , σ_R on palaeodose when $\gamma_R = 25$.

Prior			Posterior	
μ_R	σ_R	γ_R	Mean	SD
900	100	10	1016.6	16.3
900	50	10	1007.5	16.8
900	25	10	969.3	22.3
900	10	10	909.4	10.3
1000	100	10	1019.8	16.7
1000	50	10	1020.4	15.8
1000	25	10	1015.7	14.1
1000	10	10	1006.0	8.6
1100	100	10	1021.8	16.4
1100	50	10	1028.4	16.3
1100	25	10	1047.4	17.0
1100	10	10	1088.7	10.2

Table E.6: Influence of μ_R , σ_R on palaeodose when $\gamma_R = 10$.

Prior			Posterior	
μ_R	σ_R	γ_R	Mean	SD
900	100	5	1015.5	16.7
900	50	5	1005.8	17.2
900	25	5	964.2	23.0
900	10	5	908.8	10.4
1000	100	5	1019.0	17.5
1000	50	5	1017.1	15.9
1000	25	5	1013.3	14.2
1000	10	5	1005.6	8.5
1100	100	5	1088.7	10.2
1100	50	5	1027.9	16.1
1100	25	5	1029.1	16.6
1100	10	5	1089.4	10.1

Table E.7: Influence of μ_R , σ_R on palaeodose when $\gamma_R = 5$.

Appendix F

Influence of Prior Parameters in the Plateau Model

The plateau model evaluates the posterior distribution for the plateau, based on the palaeodose estimates at each preheat temperature. This model is detailed in Chapter 4, and in Section 4.3.3 the influence of the prior judgements on the posterior distribution for T_a , the plateau starting temperature, for the example of Fydell House, 311 – 6. The statistics of these distributions have been presented in Table 4.3 and here we show the corresponding posterior distributions.

Figure F.1 shows how the posterior distribution for T_a changes with prior mean, with the prior standard deviation set at 10°C. Figure F.2 and F.3 also show how the posterior T_a distribution is influenced by the prior mean, but with the prior standard deviation set at 30°C and 50°C respectively.

F.1 Influence of curve parameters

In the plateau model, the parameters η , κ control the shape of the curve before the plateau has begun on a plot of palaeodose against preheat temperature. We assume that a monotone continuous function leads to the plateau, which starts at temperature T_a at palaeodose level x_{R^*} , so that for palaeodose estimate x_{RT} with a

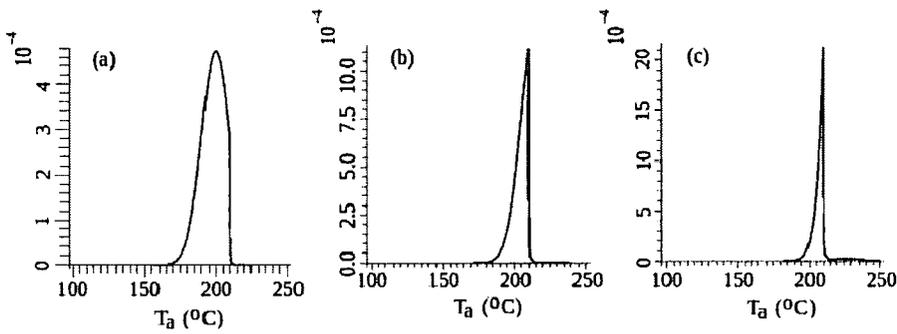


Figure F.1: Posterior T_a distributions with prior (a) $T_a \sim N(200, 10^2)$, (b) $T_a \sim N(215, 10^2)$, (c) $T_a \sim N(230, 10^2)$.

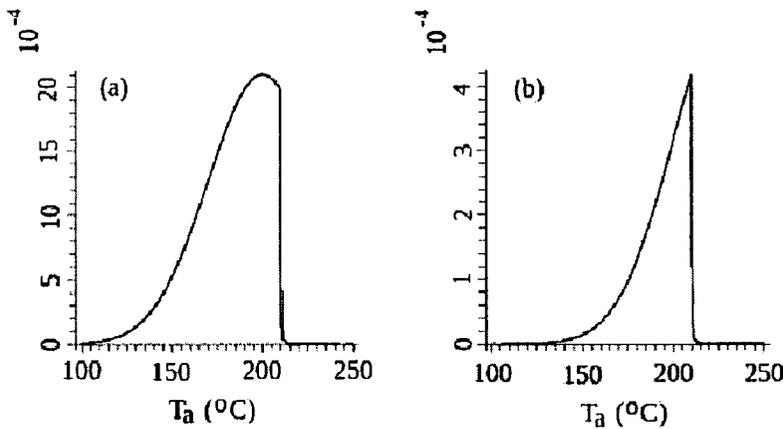


Figure F.2: Posterior T_a distributions with prior (a) $T_a \sim N(200, 30^2)$, (b) $T_a \sim N(230, 30^2)$.

preheat T ,

$$x_{RT} = \begin{cases} x_{R*} \left(\frac{(1 - \exp\{-\eta T\}) \exp\{-\kappa T\}}{(1 - \exp\{-\eta T_a\}) \exp\{-\kappa T_a\}} \right) & T < T_a \\ x_{R*} & T \geq T_a \end{cases} \quad (\text{F.1.1})$$

In the example in Section 4.3 these parameters are assigned prior distributions that allow for a wide range of curve shapes. Here we look at how the prior judgements made about these parameters influences the posterior distribution for T_a .

Table F.1 gives the posterior probability that, for sample 311-6 from Fydell House, the preheat plateau starts above 210°C for different prior judgements about the curve parameters. If prior standard deviation of κ , σ_κ , is small then the probability that the plateau starts after 210°C is increased. The prior standard deviation of η is not as influential.

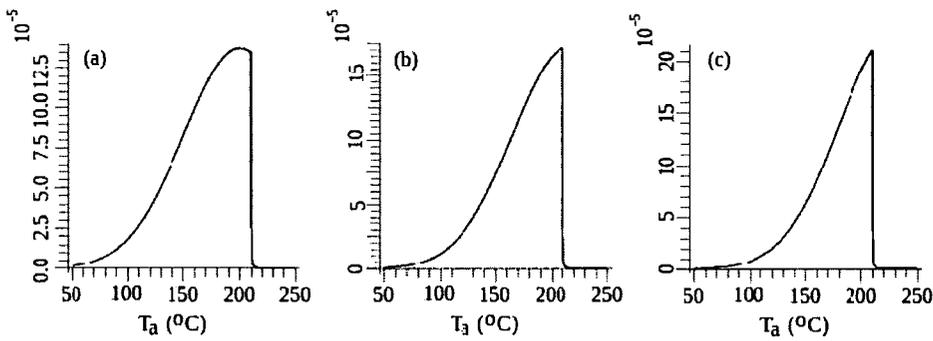


Figure F.3: Posterior T_a distributions with prior (a) $T_a \sim N(200, 50^2)$, (b) $T_a \sim N(215, 50^2)$, (c) $T_a \sim N(230, 50^2)$.

The prior means of η and κ are not influential if they are small, i.e. they represent reasonable shapes of curves. When $\sigma_\kappa, \sigma_\eta$ are high with m_κ, m_η set to 1, then the posterior probability for the plateau starting after 210°C increases. However, these statistics do not represent an expected curve shape, and as such should not be used in the analysis.

Prior				Posterior
m_η	σ_η	m_κ	σ_κ	$P[T_a > 210 data]$
0.001	1	0.001	1	0.008
0.001	10	0.001	10	0.007
0.001	0.1	0.001	0.1	0.035
0.001	0.1	0.001	10	0.007
0.001	10	0.001	0.1	0.034
0.005	1	0.005	1	0.008
0.01	1	0.01	1	0.008
0.1	1	0.1	1	0.007
1	1	1	1	0.007
1	10	1	10	0.019

Table F.1: Posterior probability for the preheat plateau starting above 210°C, for different prior judgements about the curve parameters

Appendix G

Data

G.1 Fydell House, 311-6

This data has been produced by Bailiff [14] using the single aliquot regeneration protocol [81]. The tables gives the sensitivity corrected luminescence signal strength (counts),with the background count subtracted as well as this data after standardising against the natural luminescence values, for each of the regenerative doses applied.

Preheat 210°C

aliquot	Natural	Laboratory irradiated dose (mGy)				
reference	luminescence	603	904	1206	603	1206
22031	68570	40732	61606	81469	40536	80127
	10000	5940	8984	11881	5912	11685
22032	54227	30516	45592	59952	30365	59642
	10000	5627	8408	11056	5600	10998
22033	52077	32427	48737	64561	32391	65714
	10000	6227	9359	12397	6220	12619

Preheat 220°C

aliquot	Natural	Laboratory irradiated dose (mGy)				
reference	luminescence	603	904	1206	603	1206
21031	44548	27299	40872	53995	27375	52748
	10000	6128	9175	12121	6145	11841
21032	73517	45790	69388	93210	45768	91847
	10000	6228	9438	12679	6226	12493
21033	36849	23154	34966	45951	23090	45261
	10000	6283	9489	12470	6266	12283
21034	25405	16620	24919	32789	16621	33105
	10000	6542	9809	12906	6542	13031
21035	29860	19194	28913	38062	19150	38390
	10000	6428	9683	12747	6413	12857
21036	43424	26477	39411	52806	26542	52074
	10000	6097	9076	12161	6112	11992

Preheat 230°C

Aliquot	Natural	Laboratory irradiated dose (mGy)				
reference	luminescence	603	904	1206	603	1206
22034	59783	38173	57645	75799	38080	74945
	10000	6385	9642	12679	6369	12536
22035	20231	12838	19114	25455	12874	25038
	10000	6346	9448	12582	6364	12376
22036	49640	32444	49024	64903	32298	64599
	10000	6536	9876	13075	6507	13013

Preheat 240°C

Aliquot	Natural	Laboratory irradiated dose (mGy)				
reference	luminescence	603	904	1206	603	1206
05041	28458	16398	24573	32479	16317	32310
	10000	5762	8635	11413	5734	11354
05042	59757	35540	53043	69669	35416	69983
	10000	5947	8876	11659	5927	11711
05043	31616	19208	29163	38600	19142	38363
	10000	6076	9224	12209	6055	12134
05044	36121	22508	34169	44999	22460	45409
	10000	6231	9459	12458	6218	12571
05046	40620	25215	37643	50012	25094	49640
	10000	6208	9267	12312	6178	12220

Preheat 250°C

Aliquot	Natural	Laboratory irradiated dose (mGy)				
reference	luminescence	603	904	1206	603	1206
22037	52295	32619	48572	64549	32539	65486
	10000	6238	9288	12343	6222	12522
22038	42999	26954	40087	53205	26959	53719
	10000	6268	9323	12373	6270	12493
22039	15347	9381	14003	18446	9368	18751
	10000	6113	9124	12019	6104	12218

G.2 Data at Lower Preheat Temperatures

Further experimental data was collected from sample 311 – 6 in an extension to routine dating procedure. Measurements were taken from aliquots with low preheat temperatures used, to investigate the behaviour of palaeodose before the preheat plateau has begin.

The data below gives both the sensitivity corrected luminescence values (counts), and the values after standardisation by natural luminescence (set to a value of 10000

counts).

Preheat 140°C

Aliquot	Laboratory irradiated dose (mGy)						
ref	0	750	923	1096	750	923	1096
01071	8264	7902	9624	11466	8066	9652	11785
	10000	9562	11646	13875	9760	11680	14261
01072	9848	8619	10149	12349	8447	10302	12186
	10000	8752	10306	12540	8577	10461	12374

Preheat 160°C

Aliquot	Laboratory irradiated dose (mGy)						
ref	0	750	923	1096	750	923	1096
01073	19467	15654	19104	22613	15890	19056	22467
	10000	8041	9814	11616	8162	9789	11541
01074	7701	6661	7820	9457	6561	7956	9561
	10000	8650	10155	12280	8520	10370	12454

Preheat 180°C

Aliquot	Laboratory irradiated dose (mGy)						
ref	0	750	923	1096	750	923	1096
01075	4546	3668	4399	5322	3597	4422	5398
	10000	8069	9677	11707	7912	9727	11874
01076	11608	9591	12049	13784	9746	11914	14296
	10000	8262	10380	11875	8396	10264	12346

Preheat 200°C

Aliquot	Laboratory irradiated dose (mGy)						
ref	0	750	923	1096	750	923	1096
01707	6494	5019	6634	7675	5121	6316	7308
	10000	7729	10216	11819	7886	9726	11253
01078	7135	5529	6795	7932	5391	6815	8012
	10000	7735	9523	11117	7556	9552	11229

G.3 Fydell House 311-2

This data is from aliquots sampled from a different part of Fydell House, Lincolnshire [14]. The brick is thought to be part of the original building. The tables gives the sensitivity corrected luminescence signal strength (counts), with the background count subtracted as well as this data after standardising against the natural luminescence values, for each of the regenerative doses applied.

Preheat 200°C

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	914	1097	1280	914
3112.2001	4395	4251	5038	5549	4116
	10000	9672	11463	12626	9365

Preheat 220°C

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	914	1097	1280	914
3112.2201	3860	3270	3891	4663	3351
	10000	8472	10080	12080	8861
3112.2202	1975	1729	1893	2234	1628
	10000	8754	9484	11311	8243

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	616	1231	2462	616
3112.2203	1882	1212	2229	1521	1216
	10000	6440	11844	24022	6461

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	865	1081	1298	865
3112.2204	1130	1036	1257	1418	1037
	10000	9168	11123	12589	9177
3112.2205	567	517	600	813	510
	10000	9167	10638	14415	9043

Preheat 230°C

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	865	1081	1298	865
3112.2301	1127	928	1114	1363	928
	10000	9233	9777	12094	8237

Preheat 240°C

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	865	1081	1298	865
3112.2401	1246	952	1204	1439	951
	10000	7637	9666	11549	7630
3112.2402	829	697	934	1110	697
	10000	8413	11269	13394	8409

Preheat 250°C

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	865	1081	1298	865
3112.2501	720	487	802	834	
	10000	6764	11139	11583	
3112.2502	425	293	452	514	292
	10000	6899	10630	12093	6871

Preheat 260°C

Aliquot	Laboratory irradiated dose (mGy)			
ref	0	865	1081	1298
3112.2601	838	809	1026	1557
	10000	9654	12243	18580

G.4 Fydell House 311-4

This data is from aliquots sampled from a different part of Fydell House, Lincolnshire [14]. The brick is thought to be part of the original building. The tables gives the sensitivity corrected luminescence signal strength (counts), with the background count subtracted as well as this data after standardising against the natural luminescence values, for each of the regenerative doses applied.

Preheat 200°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	882	1176	1470	882	1470
3114.2001	1164	1108	1365	1747	1073	1671
	10000	9519	11728	15007	9221	14357
3114.2002	750	550	737	890	555	934
	10000	7334	9826	11868	7404	12451
3114.2003	552	350	414	356	463	
	10000	6340	7498	6455	8383	

Preheat 210°C

Aliquot	Laboratory irradiated dose (mGy)					
	ref	0	882	1176	1470	882
3114.2101	1589	1087	1347	1870	1085	1762
	10000	6843	8476	11770	6827	11089
3114.2102	955	922	1162	1524	911	1692
	10000	9657	12166	15954	9542	17720
3114.2103	1344	912	1338	1446	917	1478
	10000	6785	9553	10755	6819	10996

Preheat 220°C

Aliquot	Laboratory irradiated dose (mGy)				
	ref	0	882	1176	1470
3114.2201	1123	866	1180	1430	872
	10000	7707	10504	12736	7797
3114.2202	929	580	654	959	614
	10000	6240	7036	10318	6607

Aliquot	Laboratory irradiated dose (mGy)					
	ref	0	1003	1239	1475	1003
3114.2203	1235	713	877	978	709	972
	10000	5777	7104	7916	5743	7870
3114.2204	898	777	959	1102	783	1080
	10000	8655	10685	12272	8716	12027

Preheat 230°C

Aliquot	Laboratory irradiated dose (mGy)					
	ref	0	1211	1513	1816	1211
3114.2301	1852	1968	2239	2508	1976	2460
	10000	10627	12092	13542	10671	13283

Preheat 240°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1176	1470	1764	1176	1764
3114.2401	1192	952	1224	1463	988	1374
	10000	7987	10269	12269	8289	11527

Aliquot	Laboratory irradiated dose (mGy)				
ref	0	882	1470	882	1470
3114.2402	473	339	611	345	613
	10000	7168	12920	7298	12987

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	882	1179	1470	882	1470
3114.2403	808	535	764	954	537	878
	10000	6625	9459	11811	6650	10871

G.5 Tattershall Castle 318-2

The samples taken from Tattershall Castle are part of the same project on brick dating using OSL [14]. This example was used to illustrate the Bayesian model for age analysis in Chapter 8. The tables gives the sensitivity corrected luminescence signal strength (counts), with the background count subtracted as well as this data after standardising against the natural luminescence values, for each of the regenerative doses applied.

Preheat 200°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1494	1793	2091	1494	2091
3182.2001	10488	9443	11045	12774	9490	12726
	10000	9003	10531	12179	9048	12134
3182.2002	3811	3412	4054	4810	3434	4762
	10000	8954	10637	12622	9009	12495

Preheat 210°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1494	1793	2091	1494	2091
3182.2101	8641	7186	8455	9960	7166	9895
	10000	8316	9785	11527	8293	11452

Preheat 220°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1496	1796	2095	1496	2095
3182.2201	6285	5156	6109	7309	5213	7336
	10000	8204	9720	11630	8294	11672
3182.2202	8299	7132	8432	9782	7064	9800
	10000	8594	10161	11787	8512	11809
3182.2203	4449	4366	4824	5197	4652	5018
	10000	9814	10843	11682	10457	11279
3182.2204	7147	6701	7125	7960	6178	7879
	10000	9376	9969	11138	9399	11024
3182.2205	7921	7909	8870	9477	8189	9211
	10000	9985	11198	11965	10338	11629

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1317	1497	1676	1317	1676
3182.2206	5083	3612	4015	4422	3637	4343
	10000	7107	7898	8700	7155	8545

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1483	1839	2195	1483	2195
3182.2207	3810	3138	3799	4402	3174	4361
	1000	8236	9972	11555	8332	11447

Preheat 240°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1494	1793	2091	1494	2091
3182.2401	2926	2523	2887	3447	2500	3483
	10000	8624	9867	11779	8545	11903
3182.2402	3206	2040	2750	2724	2033	2836
	10000	6363	8579	8496	6340	8846

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1483	1839	2195	1483	2195
3182.2403	2243	1673	2065	2386	1593	2437
	10000	7458	9206	10637	7104	10866

G.6 Tattershall Castle 318-1

This sample was taken from Tattershall Castle [14] and is used to illustrate the model for inference from a number of samples in Chapter 8. The tables gives the sensitivity corrected luminescence signal strength (counts), with the background count subtracted as well as this data after standardising against the natural luminescence values, for each of the regenerative doses applied.

Preheat 200°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1495	1794	2093	1495	2093
3181.2001	21573	18304	21673	25092	18185	25497
	10000	8485	10047	11631	8429	11819
3181.2002	21231	19787	24003	27524	19661	27521
	10000	9320	11306	12964	9261	12962

Preheat 220°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1496	1796	2095	1496	2095
3181.2201	16387	13671	16190	19169	13748	18616
	10000	8342	9880	11698	8390	11360
3181.2202	32089	29092	34556	40415	29094	41145
	10000	9342	9880	11698	8390	11360

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	2095	2394	2693	2095	2693
3181.2203	11739	13734	15103	16921	13582	16951
	10000	11700	12865	14414	11570	14440
3181.2204	16890	19908	23095	25724	20060	25532
	10000	11801	13690	15248	11891	15134
3181.2205	10141	12076	15756	15841	12446	15585
	10000	11908	15534	15620	12273	15369

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	600	899	1199	600	1199
3181.2206	15811	5131	7511	10279	5072	10352
	10000	3245	4750	6501	3208	6548

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1424	1780	2136	1424	2136
3181.2207	6826	5076	6525	7794	5277	7834
	10000	7436	9559	11418	7730	11477
3181.2208	8228	6561	8292	9997	6621	9491
	10000	7974	10077	12150	8047	11535

Preheat 240°C

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1424	1780	2136	1424	2136
3181.2401	3074	2260	2957	3473	2286	3502
	10000	7352	9620	11298	7436	11391

Aliquot	Laboratory irradiated dose (mGy)					
ref	0	1499	1799	2098	1499	2098
3181.2402	6634	5108	6322	6975	5274	6959
	10000	7700	9530	10513	7949	10490

Appendix H

Code

Here we present the code used to programme the various stages of the model.

H.1 Code for the Combined Aliquot Model

The combined aliquot model is detailed in Chapter 3, and is used to evaluate the palaeodose at a single preheat temperature. It is used later in the analysis, once the preheat plateau has been located, to estimate the posterior palaeodose distribution for the sample. The code is written for R.

```
##Code for the Combined Aliquot Model to estimate Palaeodose
#####

#Evaluates the posterior palaeodose distribution based on J aliquots
#i.e. J aliquots with the same preheat temperature
#or J aliquots which lie on the preheat plateau

#####

#It assumes that the single aliquot regeneration procedure has been used
#and there is a linear relationship between dose and luminescence intensity
```

```
#####
```

```
#PRIOR SPECIFICATIONS
```

```
muR<-1000 #Prior mean palaeodose (xR)
```

```
sigmaR<-100 #prior palaeodose SD
```

```
mA<-0 #prior mean for alpha, the intercept coefficient of the dose response
```

```
mB<-10 #prior mean for beta, the gradient coefficient of the dose response
```

```
sigmaA<-50 #prior SD for alpha
```

```
sigmaB<-20 #prior SD for beta
```

```
gamR<-50 #SD of the aliquot estimates xRj around true palaeodose xR
```

```
gamalp<-20 #SD of the aliquot estimates alphaj around true alpha
```

```
gambet<-5 #SD of the aliquot estimates betaj around true beta
```

```
mtau<-0.01 #prior mean of tau, the precision of the residuals
```

```
vartau<-0.0015 #prior SD of tau
```

```
a<-2*mtau/vartau
```

```
d<-a*mtau #hyperparameters of the gamma distribution for tau a priori
```

```
rho<--0.3 #correlation between alpha and beta a priori, and alphaj and betaj
```

```
conditional on alpha and beta
```

```
#####
```

```
#DATA INPUT
```

```
yR<-10000 #Standardised value of natural luminescence used for all
```

```
#aliquots
```

```
dat1<-read.table("/filepath/aliquot1.txt", header=T)
dat2<-read.table("/filepath/aliquot2.txt", header=T)
dat3<-read.table("/filepath/aliquot3.txt", header=T)

#The data should consist of two columns for each aliquot, the
#regenerative dose values 'x' and the background and sensitivity
#corrected luminescence response, 'y', standardised against yR.

#Here the data should be saved in a plain text document,
#one for each aliquot, with the columns labelled appropriately.

n<-length(dat1$x) #number of regenerative doses applied

J<-3 #number of aliquots which are to be used to evaluate palaeodose

#Write all the data in a single array

dat<-array(NA,c(n,2,J))
for (i in 1:n){
for (j in 1:2){
dat[i,j,1]<-dat1[i,j]
dat[i,j,2]<-dat2[i,j]
dat[i,j,3]<-dat3[i,j]
}
}
```

```
#Find the required sums for the conditional distributions
```

```
sumy<-array(NA,J)
sumxy<-array(NA,J)
sumx<-array(NA,J)
sumx2<-array(NA,J)
```

```
for (k in 1:J){
sumy[k]<-sum(dat[,2,k])
sumxy[k]<-sum(dat[,1,k]*dat[,2,k])
sumx[k]<-sum(dat[,1,k])
sumx2[k]<-sum(dat[,1,k]^2)
}
```

```
#####
```

```
#CONDITIONAL POSTERIOR DISTRIBUTIONS FOR THE GIBBS SAMPLER
```

```
xR.update<-function(){
v.xR<-1/(J/gamR^2+1/sigmaR^2)
xR.hat<-((sum(xRj)/gamR^2)+muR/sigmaR^2)*v.xR
rnorm(1,xR.hat,sqrt(v.xR))
}
```

```
alpha.update<-function(){
v.alpha<-1/(J/(gamalp^2*(1-rho^2))+1/(sigmaA^2*(1-rho^2)))
alpha.hat<-((sum((alphaj-rho*(betaj-beta)*(gamalp/gambet))/
(gamalp^2*(1-rho^2))+(mA-rho*(beta-mB)*(sigmaA/sigmaB))/
(sigmaA^2*(1-rho^2))))*v.alpha
rnorm(1,alpha.hat,sqrt(v.alpha))
}
```

```
xRj.update<-function(){
v.xRj<-1/(betaj^2*tau+1/gamR^2)
xRj.hat<-(betaj*(yR-alpha)*tau+xR/gamR^2)*v.xRj
rnorm(J,xRj.hat,sqrt(v.xRj))
}
```

```
alphaj.update<-function(){
v.alphaj<-1/((n+1)*tau+1/(gamalp^2*(1-rho^2)))
alphaj.hat<-((sumy+yR-betaj*(sumx+xRj))*tau+
(alpha+rho*(betaj-beta)*(gamalp/gambet))/
(gamalp^2*(1-rho^2)))*v.alphaj
rnorm(J,alphaj.hat,sqrt(v.alphaj))
}
```

```
betaj.update<-function(){
v.betaj<-1/((sumx2+xRj^2)*tau+1/(gambet^2*(1-rho^2)))
betaj.hat<-((sumxy+yR*xRj-alpha*(sumx+xRj))*tau+
(beta+rho*(alphaj-alpha)*(gambet/gamalp))/
(gambet^2*(1-rho^2)))*v.betaj
rnorm(J,betaj.hat,sqrt(v.betaj))
}
```

```
beta.update<-function(){
v.beta<-1/(J/(gambet^2*(1-rho^2))+1/(sigmaB^2*(1-rho^2)))
beta.hat<-(((sum(betaj-rho*(alphaj-alpha)*(gambet/gamalp)))/
(gambet^2*(1-rho^2)))+(mB-rho*(beta-mB)*(sigmaB/sigmaA))/
(sigmaB^2*(1-rho^2)))*v.beta
rnorm(1,beta.hat,sqrt(v.beta))
}
```

```

tau.update<-function(){
tau.shape<-(d+J*n)/2
resid<-array(NA,J)
for (j in 1:J){
resid[j]<-sum((dat[,2,j]-alphaj[j]-betaj[j]*dat[,1,j])^2)+
  (yR-alphaj[j]-betaj[j]*xRj[j])^2
}
tau.scale<-(sum(resid)+a)/2
rgamma(1,tau.shape,tau.scale)
}

```

```

#####
#GIBBS SAMPLER
n.chains<-5
n.iter<-20000
lations<-array(NA,c(n.iter,n.chains,J*3+4))
dimnames(lations)<-list(NULL,NULL,c("xR","alpha","beta",
  paste("xRj[",1:J,"]",sep=""),paste("alphaj[",1:J,"]",sep=""),
  paste("betaj[",1:J,"]",sep=""),"tau"))

for(m in 1:n.chains){
alpha<-mA
alphaj<-rep(mA,J)
beta<-mB
betaj<-rep(mB,J)
xR<-muR
xRj<-rep(muR,J)
tau<-0.1

```

```
for (t in 1:n.iter){
xR<-xR.update()
alpha<-alpha.update()
beta<-beta.update()
xRj<-array(xRj.update(),J)
alphaj<-array(alphaj.update(),J)
betaj<-array(betaj.update(),J)
tau<-tau.update()

lations[t,m,]<-c(xR,alpha,beta,xRj,alphaj,betaj,tau)
}
}

#####

nos<-seq(1000,n.iter,4) #Thin the chains, with a burn-in period

mean(lations[nos,,"xR"]) #Posterior mean for palaeodose

sd(as.vector(lations[nos,,"xR"])) #Posterior SD for palaeodose

plot(density(lations[nos,,"xR"])) #Posterior palaeodose density
```

H.2 Code for Plateau Model

Chapter 4 details a model to compute the posterior distribution of the starting temperature of the preheat plateau. Below is the code used in Maple to find this distribution for sample 311-6, Fydell House.

```
with(Statistics):

with(plots):

post:=proc(x,k,g,t)

#Data Input (posterior palaeodose means at each preheat)
X:=[1019.0,984.0,955.6,1002.2,990.6];
#Preheat temperatures
T:=[210,220,230,240,250];
#Variance in X
sigma:=[27,31,34.7,26.9,44.3];

#Prior mean and SD for kappa, eta, curve parameters
mK:=0.001;
sigK:=1;
me:=0.001;
sige:=1;
#Prior mean and SD for xR*, palaeodose on the plateau
mx:=1000;
sigx:=100;
#Prior mean and SD for Ta, temperature at which plateau starts
mT:=215;
sigT:=30;
```

```

#Compute likelihood

for i from 1 to nops(X) do
  xR:=proc(x,k,e,t,i)
    if T[i]<t then
      (x*(1-exp(-e*T[i]))*(exp(-k*T[i])))/((1-exp(-e*t))*(exp(-k*t)))
    else x
  end if;
end proc;

  S[i]:=(1/sqrt(2*Pi*sigma[i]^2))*exp(-(1/(2*sigma[i]^2))*(xR(x,k,e,t,i)-x)^2);
end do;

lik:=mul(S[i],i=1..nops(X));

kap:=RandomVariable(Normal(mK,sigK)):
ka:=proc(k) PDF(kap,k) end proc;
eta:=RandomVariable(Normal(me,sige)):
et:=proc(g) PDF(eta,e) end proc;
xRss:=RandomVariable(Normal(mx,sigx)):
xRs:=proc(x) PDF(xRss,x) end proc;
Tss:=RandomVariable(Normal(mT,sigT)):
Ts:=proc(t) PDF(Tss,t) end proc;
post:=lik*ka(k)*ga(g)*xRs(x)*Ts(t);

end proc;

```

```
#Estimate the posterior distribution for Ta using numerical
#integration

L1:=[]:
for i from 100 to 250 by 0.01 do
  L1:=[op(L1),evalf(Int(Int(Int(post(x,k,e,i),k=0..1),e=0..1),x=600..1500))]
od:

L2:=[]:
for j from 100 to 250 by 0.01 do
  L2:=[op(L2),j]
od:

L3:=[]:
for k from 1 to nops(L2) do
  L3:=[op(L3),[L2[k],L1[k]]]
od:

pointplot(L3,connect=true);

#This is not a true density, but proportional to it. If all the
#probability lies in this region, then they can be normalised to
#find the probability of the plateau starting in certain preheat
#regions.
```

H.3 Code for Dose Rate

The model for evaluating the dose rate distribution is detailed in Chapter 5. Here we present the code used in R to compute the dose rate distribution for a particular example.

```
#####
#Measured beta dose rate
mB<-2.21
#Measured gamma dose rate
mgam<-1.30
#Cosmic dose rate
mc<-0.2

#Water saturation content
mW<-0.033
#Mean average fractional water content (F)
mF<-0.15
#Standard deviation of average fractional water content
sigF<-0.2
#Calculation of hyperparameters of beta distribution for F
alphaF<-(mF*(1-mF)/sigF^2)-mF
betaF<-alphaF*(1-mF)/mF

#####

num<-200000 #Number of iterations

#Dose rate based on the dose rate equation
```

```
Annual<-function(){  
  b<-rnorm(1,0.92,0.05)  
  g<-rnorm(1,0.93,0.1)  
  Hg<-rnorm(1,1.25,0.1)  
  Hb<-rnorm(1,1.14,0.1)  
  beta<-rnorm(1,mB,0.025*mB)  
  gam<-rnorm(1,mgam,0.025*mgam)  
  c<-rnorm(1,mc,0.025*mc)  
  W<-rnorm(1,mW,0.025*mW)  
  F<-rbeta(1,alphaF,betaF)  
  (b/(1+Hb*W*F))*beta+(g/(1+Hg*W*F))*gam+c  
}
```

```
dose<-array(NA,num)  
for (p in 1:num){  
  dose[p]<-Annual()  
}
```

```
mean(dose)  
sd(dose)  
plot(density(dose))
```

H.4 Code for Age Evaluation

The model for sample age is given in Chapter 6. The age ratio is estimated by taking draws from the simulated palaeodose distribution and dividing each one by a draw from the dose rate distribution. The mean and standard deviation of these are used as the data input in the evaluation of the sample age. The R code used for sample age is shown below, assuming a normal distribution for age a priori.

```
#####
#Data Input
AE<-549.6 #mean of the ratio estimate
omE<-28.4 #SD of the ratio estimate

#####
#Prior Specifications
mA<-558 #prior age mean
sigA<-20 #prior age SD

sigE<-5 #uncertainty in age ratio as an estimate for age

#####
#Posterior

meanage<-(AE/(sigE^2+omE^2)+mA/sigA^2)/(1/(sigE^2+omE^2)+1/sigA^2)
sdage<-sqrt(1/(1/(sigE^2+omE^2)+1/sigA^2))

meanage
sdage

#Date of the sample (given measurements taken in 2005)
2005-round(meanage)
```

H.5 Coeval Model

The coeval model is developed in Section 7.1 for inference from a number of samples which are the same age. The R code for this model for two coeval samples is laid out below.

```
#####
#Data Input

AE1<-260.9 #Mean of the ratio estimate 1
AE2<-273.5 #Mean of the ratio estimate 2
w1<-15 #SD of ratio estimate 1
w2<-18 #SD of ratio estimate 2

#####
#Prior Specifications

#Correlation between the ratio estimates given the age ratios
rhoE<-0.55
#Correlation between the age ratios given the age
rhoA<-0.2

mA<-280 #Prior mean age
sigA<-25 #Prior SD age
sigE<-5 #Uncertainty in the age ratios estimating the age

#####

sig1<-sqrt(w1^2+sigE^2)
sig2<-sqrt(w2^2+sigE^2)
```

```
rho<- (rhoE*w1*w2+rhoA*sigE^2)/(sig1*sig2)

agesigma<- ((1/(1-rho^2))*(1/sig1^2-2*rho/(sig1*sig2)+1/sig2^2)
+1/sigA^2)^(-1)

age<- ((1/(1-rho^2))*(Ae1/sig1^2-rho*(Ae1+Ae2)/(sig1*sig2)+
Ae2/sig2^2)+mA/sigA^2)*sigma

mu          #Posterior mean age
sqrt(sigma) #Posterior SD age

2005-round(mu) #Date (given measurements taken in 2005)

#####
```

H.6 Similar Age Model

The similar age model is applicable to a number of samples which are similar in age. The similarity is expressed through the prior specifications made. The model is set out in Section 7.2 and the R code for two such samples below.

```
#####
```

```
#Data Input
```

```
AE1<-260.87 #Mean of ratio estimate 1
```

```
AE2<-273.53 #Mean of ratio estimate 2
```

```
w1<-15.09 #SD of ratio estimate 1
```

```
w2<-18.98 #SD of ratio estimate 2
```

```
#####
```

```
#Prior Specifications
```

```
mA1<-280 #Prior mean of age 1
```

```
mA2<-290 #Prior mean of age 2
```

```
sigmaA1<-10 #Prior SD of age 1
```

```
sigmaA2<-10 #Prior SD of age 2
```

```
rhop<-0.2 #Prior correlation between age 1 and age 2
```

```
#Correlation between the ratio estimates given the age ratios
```

```
rhoE<-0.55
```

```

#Correlation between the age ratios given the age
rhoA<-0.2

#Uncertainty in the age ratios as an estimate for age
sigmaE<-5

#####

#Posterior Conditional Distributions

sigma1<-sqrt(w1^2+sigmaE^2)
sigma2<-sqrt(w2^2+sigmaE^2)
rho<-(rhoE*w1*w2+rhoA*sigmaE^2)/(sigma1*sigma2)

A1.update<-function(){
V.A1<-(1/(sigma1^2*(1-rho^2))+1/(sigmaA1^2*(1-rhop^2)))^(-1)
A1.hat<-(AE1/(sigma1^2*(1-rho^2))+(rho*(A2-AE2))
  /(sigma1*sigma2*(1-rho^2))+(mA1+rhop*(sigmaA1/sigmaA2)*(A2-mA2))
  /(sigmaA1^2*(1-rhop^2)))*V.A1
rnorm(1,A1.hat,sqrt(V.A1))
}

A2.update<-function(){
V.A2<-(1/(sigma2^2*(1-rho^2))+1/(sigmaA2^2*(1-rhop^2)))^(-1)
A2.hat<-(AE2/(sigma2^2*(1-rho^2))+rho*(A1-AE1)
  /(sigma2*sigma1*(1-rho^2))+(mA2+rhop*(sigmaA1/sigmaA2)*(A1-mA1))
  /(sigmaA1^2*(1-rhop^2)))*V.A2
rnorm(1,A2.hat,sqrt(V.A2))
}

#####

```

```
#Gibbs Sampler

n.iter<-20000
n.chains<-3

ages<-array(NA,c(n.iter,n.chains,2))
dimnames(ages)<-list(NULL,NULL,c("A1","A2"))
for (m in 1:n.chains){
  A1<-mA1
  A2<-mA2
  for (t in 1:n.iter){
    A1<-A1.update()
    A2<-A2.update()
    ages[t,m,]<-c(A1,A2)
  }
}

nos<-seq(1000,n.iter,6) #Thin chains with a burn-in

mean(ages[,nos,"A1"]) #Posterior mean age 1
mean(ages[,nos,"A2"]) #Posterior mean age 2
sd(as.vector(ages[,nos,"A1"])) #Posterior SD age 1
sd(as.vector(ages[,nos,"A2"])) #Posterior SD age 2

2005-round(mean(ages[nos,,"A1"])) #Date 1
2005-round(mean(ages[nos,,"A2"])) #Date 2
```

H.7 Ordered Age Model

This model considers samples which have a known relative chronology (Section 7.3) which implies a priori that, say, $A_1 < A_2$. The R code for such a model is given below.

```
#####
#Data Input

AE1<-260.87 #Mean of ratio estimate 1
AE2<-273.53 #Mean of ratio estimate 2

w1<-15.09 #SD of ratio estimate 1
w2<-18.98 #SD of ratio estimate 2

#####
#Prior Specifications

A_1<A_2
mA1<-280 #Prior mean of age 1
mA2<-290 #Prior mean of age 2

sigmaA1<-10 #Prior SD of age 1
sigmaA2<-10 #Prior SD of age 2

rhoE<-0.2 #Prior correlation between age 1 and age 2

#Correlation between the ratio estimates given the age ratios
rhoE<-0.55

#Correlation between the age ratios given the age
```

```

rhoA<-0.2

#Uncertainty in the age ratios as an estimate for age
sigmaE<-5

#####
#Posterior Conditional Distributions with rejection criteria

sigma1<-sqrt(w1^2+sigmaE^2)
sigma2<-sqrt(w2^2+sigmaE^2)
rho<-(rhoE*w1*w2+rhoA*sigmaE^2)/(sigma1*sigma2)

A1.update<-function(){
V.A1<-(1/(sigma1^2*(1-rho^2))+1/(sigmaA1^2*(1-rhop^2)))^(-1)
A1.hat<-(AE1/(sigma1^2*(1-rho^2))+(rho*(A2-AE2))
  /((sigma1*sigma2*(1-rho^2))+(mA1+rhop*(sigmaA1/sigmaA2)*(A2-mA2))
  /((sigmaA1^2*(1-rhop^2))))*V.A1
rnorm(1,A1.hat,sqrt(V.A1))
}

A2.update<-function(){
V.A2<-(1/(sigma2^2*(1-rho^2))+1/(sigmaA2^2*(1-rhop^2)))^(-1)
A2.hat<-(AE2/(sigma2^2*(1-rho^2))+rho*(A1-AE1)
  /((sigma2*sigma1*(1-rho^2))+(mA2+rhop*(sigmaA1/sigmaA2)*(A1-mA1))
  /((sigmaA1^2*(1-rhop^2))))*V.A2
repeat{
Ato<-rnorm(1,A2.hat,sqrt(V.A2))
if(Ato>A1) return (Ato)
}
}

```

```
#####
```

```
#Gibbs sampler
```

```
n.iter<-20000
```

```
n.chains<-3
```

```
ages<-array(NA,c(n.iter,n.chains,2))
```

```
dimnames(ages)<-list(NULL,NULL,c("A1","A2"))
```

```
for (m in 1:n.chains){
```

```
  A1<-mA1
```

```
  A2<-mA2
```

```
  for (t in 1:n.iter){
```

```
    A1<-A1.update()
```

```
    A2<-A2.update()
```

```
    ages[t,m,]<-c(A1,A2)
```

```
  }
```

```
}
```

```
mean(ages[,,"A1"]) #Posterior mean age 1
```

```
mean(ages[,nos,"A2"]) #Posterior mean age 2
```

```
sd(as.vector(ages[,nos,"A1"])) #Posterior SD age 1
```

```
sd(as.vector(ages[,nos,"A2"])) #Posterior SD age 2
```

```
2005-round(mean(ages[nos,,"A1"])) #Date 1
```

```
2005-round(mean(ages[nos,,"A2"])) #Date 2
```

Appendix I

Analysis of 311-2, 311-4 Fydell House

Samples 311-2, 311-4 are taken from bricks in Fydell House, as was sample 311-6 which has been used as an illustrative example throughout the thesis. Here we present the evaluation of the age ratios of samples 311-2, 311-4 which are used as an example in Chapter 7. The data for these samples is given in Appendix G.

I.1 Sample 311-2

First, the palaeodose was estimated at each preheat temperature using the combined aliquot model. The prior distributions used were the same as in the analysis of 311-6, as they were taken from the same building. The prior hyperparameters are given in Table I.1, and were chosen using the reasoning in Section 3.5.1.

μ_R	σ_R	γ_R	m_α	σ_α	m_β	σ_β	ρ	γ_α	γ_β	d	a
1000	100	50	0	50	10	20	-0.3	20	5	$\frac{5}{7}$	$13\frac{1}{3}$

Table I.1: Values assigned to the prior hyperparameters for the combined aliquot model, when it is applied to

The Gibbs sampler was run for 20,000 iterations for 5 chains, and after appropriate convergence diagnostics a burn-in of 1000 iterations was used and they were thinned every 10. The posterior distributions for palaeodose at each preheat

temperature are summarised in Table I.2 and given in Figure I.1.

Preheat Temperature (°C)	# Aliquots	Posterior Mean (mGy)	Posterior SD
200	1	987	62
220	5	1021	34
230	1	1022	76
240	2	1042	45
250	2	1036	78
260	1	987	62

Table I.2: Posterior mean and standard deviation of palaeodose at each preheat temperature, for sample 311-2.

The posterior mean and standard deviations given in Table I.2 were used as data input for the preheat plateau model, to estimate the location of the preheat plateau. These values are plotted against preheat temperature in Figure I.2 for sample 311-6.

The prior distributions used in the plateau model to evaluate the posterior distribution of the plateau starting temperature are given in Table I.3; these values were elicited using the reasons laid out in Section 4.3.1. This resulted in a posterior probability that the plateau started after 200°C of 0.0001.

Parameter	T_a	x_{R*}	η	κ
Prior	$N(215, 30^2)$	$N(1000, 100^2)$	$N(0.003, 1^2)$	$N(0.003, 1^2)$

Table I.3: Prior distributions used for Sample 311-6, Fydell House, in the plateau model.

This indicates that it is likely that the preheat plateau starts before 200°C, and so all aliquots from sample 311-2 lie on the plateau and hence can be used to evaluate the sample palaeodose. So the combined aliquot model was applied to all the aliquots, with prior distributions as Table I.1. This resulted in a posterior palaeodose distribution for the sample with mean 1013 mGy, standard deviation 27.

Next, the dose rate was evaluated. The distributions used in the dose rate model are given in Table I.4.

After 100,000 iterations this resulted in a dose rate distribution with mean 3.88

Parameter	Distribution
\dot{D}_β	$N(2.8, (0.025 \times 2.81)^2)$
\dot{D}_γ	$N(1.22, (0.025 \times 1.22)^2)$
\dot{D}_c	$N(0.2, (0.025 \times 0.2)^2)$
W	$N(0.033, (0.025 \times 0.033)^2)$
b	$N(0.92, 0.05^2)$
g	$N(0.93, 0.1^2)$
H_β	$N(1.25, 0.1^2)$
H_γ	$N(1.14, 0.1^2)$
F	$\beta(3.0375, 17, 2125)$

Table I.4: Distributions assigned to the parameters in the model for dose rate for sample 311-2 of Fydell House.

mGy/a and standard deviation 0.20. This was used, along with the sample palaeodose distribution, to estimate the age ratio (Chapter 6), which had mean 260.9 and standard deviation 15 years.

I.2 Sample 311-4

Similarly, the luminescence measurements from sample 311-4 were analysed. The same prior distributions for the palaeodose were used, as the sample was taken from the same building. The posterior distributions for palaeodose at each preheat temperature are summarised in Table I.5 and shown in Figure I.3.

The palaeodose estimates are plotted against preheat temperature in Figure I.4. These are used as data input in the plateau model, to find the posterior distribution for the temperature at which the plateau starts. Using the same prior distributions as for sample 311-2 (Table I.3), the posterior probability that the preheat plateau starts above 200°C is 0.001. This indicates that all the palaeodose estimates lie on the preheat plateau, and so all aliquots should be used to evaluate the sample palaeodose.

The combined aliquot model was applied to all 14 aliquots of 311-4, with the same

Preheat Temperature (°C)	# Aliquots	Posterior Mean (mGy)	Posterior SD
200	3	1068	79
210	3	1070	62
220	4	1173	77
230	1	1063	94
240	3	1223	53

Table I.5: Posterior mean and standard deviation of palaeodose at each preheat temperature, for sample 311-4.

prior hyperparameters as previously (Table I.1). The posterior distribution of the sample was evaluated to have a mean 1084 mGy, standard deviation 52.

The dose rate for sample 311-4 was evaluated using the model set out in Chapter 5. The distributions used for this calculation are given in Table I.6, and the resulting dose rate distribution has mean 3.97 mGy/a, standard deviation 0.20.

Parameter	Distribution
\dot{D}_β	$N(2.91, (0.025 \times 2.91)^2)$
\dot{D}_γ	$N(1.2, (0.025 \times 1.2)^2)$
\dot{D}_c	$N(0.2, (0.025 \times 0.2)^2)$
W	$N(0.033, (0.025 \times 0.033)^2)$
b	$N(0.92, 0.05^2)$
g	$N(0.93, 0.1^2)$
H_β	$N(1.25, 0.1^2)$
H_γ	$N(1.14, 0.1^2)$
F	$\beta(3.0375, 17, 2125)$

Table I.6: Distributions assigned to the parameters in the model for dose rate for sample 311-4 of Fydell House.

The palaeodose and dose rate distributions are used to evaluate the age ratio using the age equation. For sample 311-4, the age ratio was estimated to have a mean of 273.5 years with standard deviation 18.

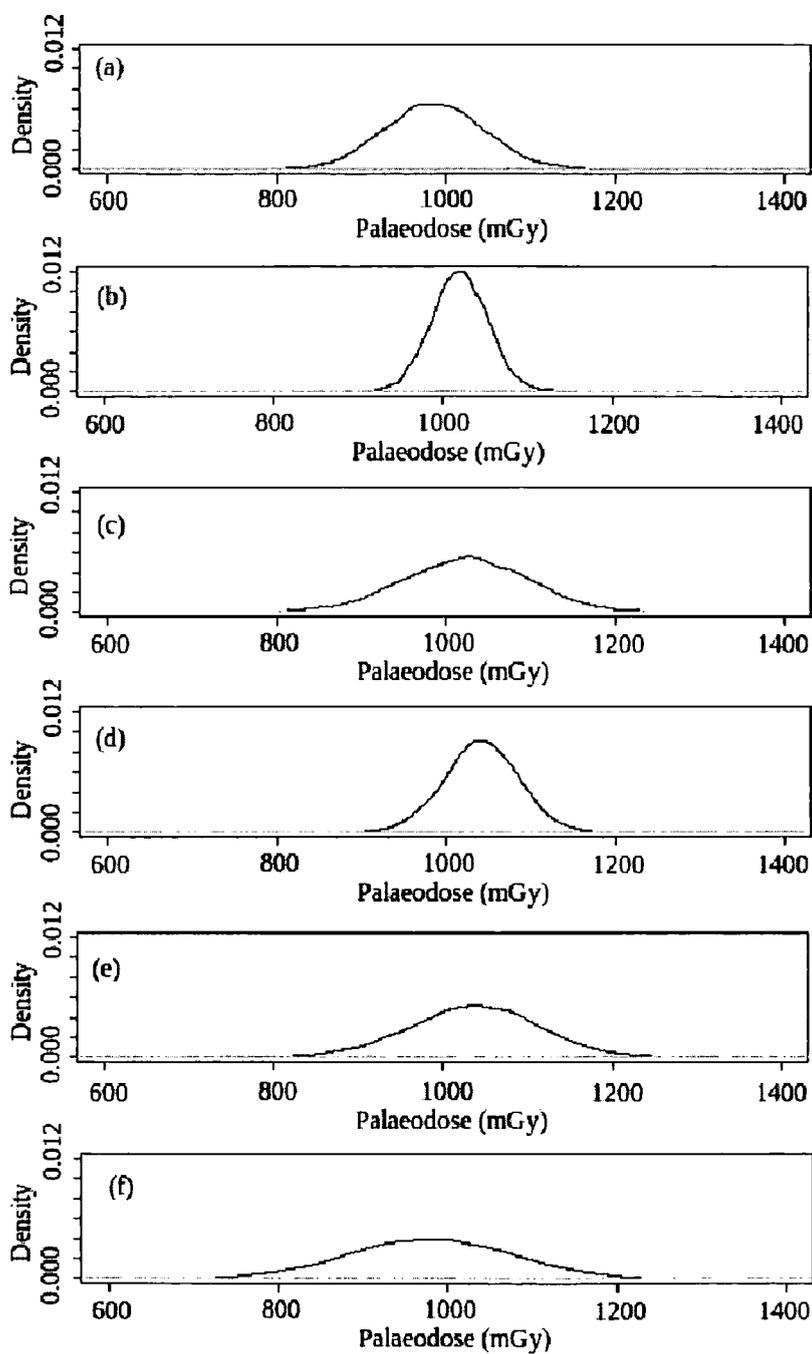


Figure I.1: Posterior Palaeodose distribution for sample 311-2 at preheats (a) 200°C, (b) 220°C, (c) 230°C, (d) 240°C, (e) 250°C, (f) 260°C.

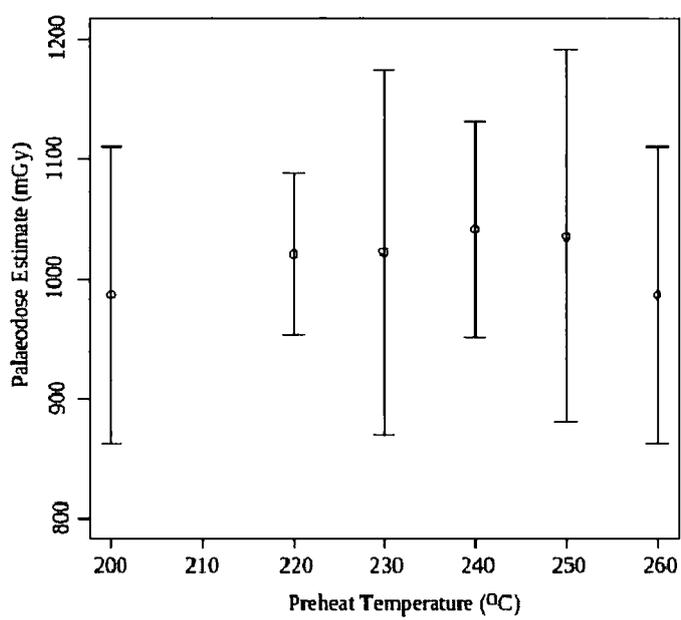


Figure I.2: Palaeodose estimates plotted against preheat temperature for 311-2, with two standard deviation uncertainty bars.

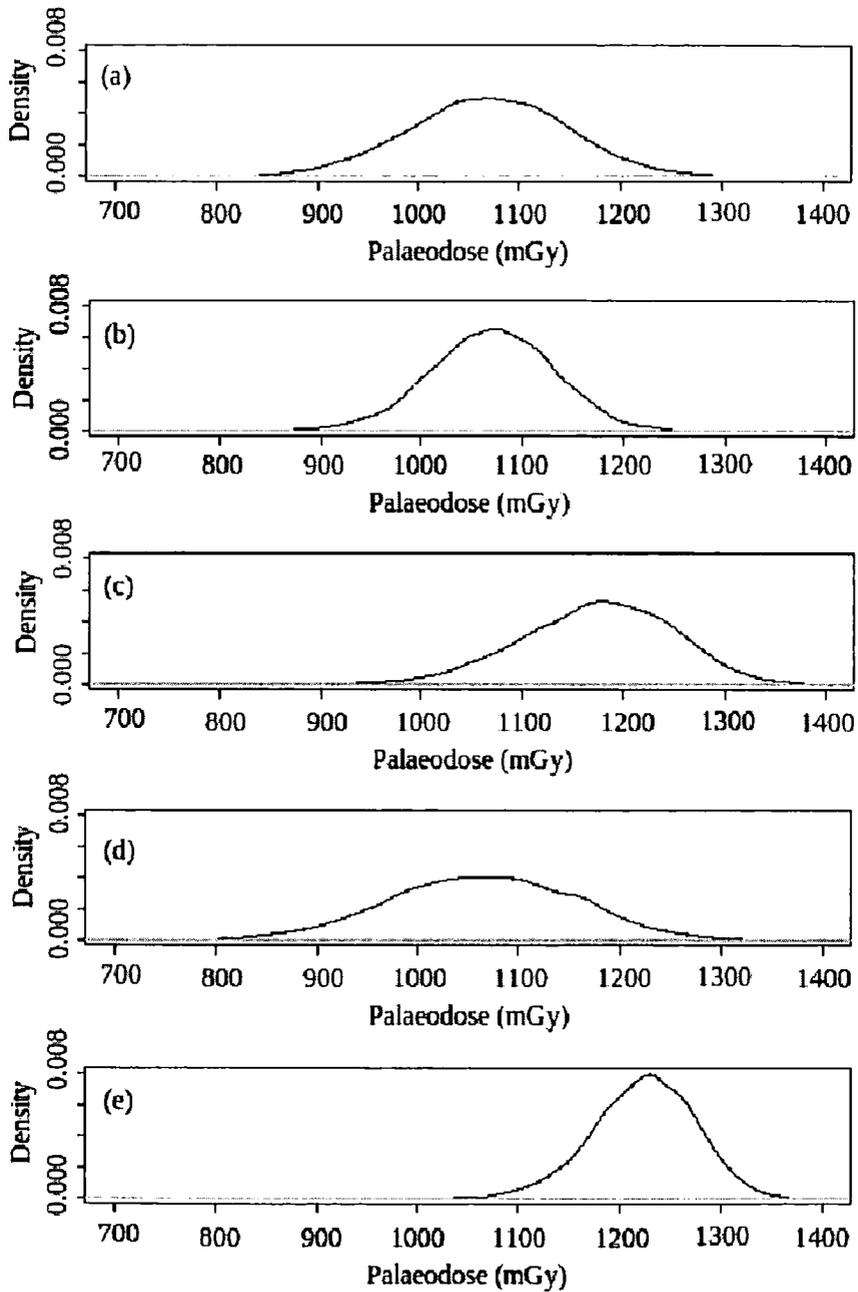


Figure I.3: Posterior Palaeodose distribution for sample 311-4 at preheats (a) 200°C, (b) 210°C, (c) 220°C, (d) 230°C, (e) 240°C.

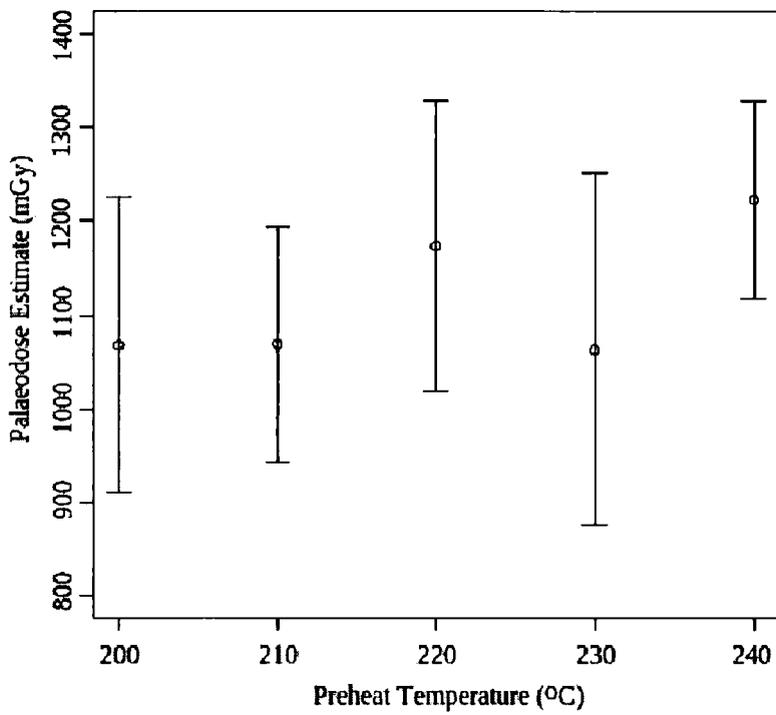


Figure I.4: Palaeodose estimates plotted against preheat temperature for 311-4, with two standard deviation uncertainty bars.

