

ROBUST BAYESIAN ANALYSIS, AN ATTEMPT TO IMPROVE BAYESIAN SEQUENCING

Franz Weninger¹ • Peter Steier • Walter Kutschera • Eva Maria Wild

Vienna Environmental Research Accelerator (VERA), Faculty of Physics, Isotope Research, University of Vienna, Währinger Straße 17, A-1090 Vienna, Austria.

ABSTRACT. Bayesian sequencing of radiocarbon dates deals with the problem that in most cases there does not exist an unambiguous way to define the so-called prior function, which represents information in addition to the result of the ¹⁴C measurements alone. However, a random choice of a particular prior function can lead to biased results. In this paper, “robust Bayesian analysis,” which uses a whole set of prior functions, is introduced as a more reliable method. The most important aspects of the mathematical foundation and of the practical realization of the method are described. As a general result, robust Bayesian analysis leads to higher accuracy, but paid for with reduced precision. Our investigations indicate that it seems possible to establish robust analysis for practical applications.

INTRODUCTION

Motivation

Bayesian sequencing has become a generally accepted and very successful tool to reduce the uncertainty of calibrated radiocarbon ages, when additional information (i.e. “prior” information) about the temporal relationship within a series of ¹⁴C dates from an archaeological stratigraphy exist. However, there is a well-known fundamental problem in applying Bayesian statistics: The prior information, which is the archaeological evidence in this application, is often not sufficient to unambiguously define a prior function. Nevertheless, in the standard approach, a specific prior function is selected in a canonical way. The gaps in the prior knowledge are filled with assumed information, which is intended to have no influence on the final result. This may, however, not necessarily be the case. A possible way to overcome this uncertainty is the use of various prior functions that are all in agreement with the archaeological information, and then unify the results; this is called “robust Bayesian analysis.”

After introducing the basics of Bayesian sequencing below, the idea of robust analysis will be described in detail. However, the actual realization of the method is not so simple as the principle idea suggests. Two main problems will be discussed in the paper: the need of discarding corrupt priors and how to deal with an infinite set of prior functions. A comparison of robust analysis with the usual sequencing is performed with the help of 2 specific examples.

A Remark to the Used Notation

The basic method of Bayesian sequencing will be illustrated by the help of 3 different probability densities, denoted as prior function, likelihood function, and posterior function, as described in detail below. In full notation, these densities can be distinguished by their arguments, because the prior function is a probability density of the real ages $p(t_1, \dots, t_n)$, the likelihood function is a conditional density of the ¹⁴C ages (denoted x_1, \dots, x_n) at given real ages $p(x_1, \dots, x_n | t_1, \dots, t_n)$, and the posterior function is a conditional density of the real ages at given ¹⁴C ages $p(t_1, \dots, t_n | x_1, \dots, x_n)$. However, for clarity, sometimes it is useful to simplify the arguments or omit them completely. To distinguish the different densities independently of their arguments, we use the letter a to denote the

¹Corresponding author. Email: franz.weninger@univie.ac.at.

prior function (from “a priori,” which is the origin of “prior”), l for the likelihood function, and p for the posterior function.

Basic Description of Bayesian Sequencing

The basic single-sample calibration in ^{14}C dating has a serious drawback. Depending on the shape of the calibration curve (e.g. for sections with large wiggles), the procedure may produce calibrated ages with high uncertainty. Fortunately, in many cases a series of samples with various known relations between their ages deduced from the excavation site is available. Including this archaeological information within the calibration process by means of Bayesian statistics can improve the accuracy of the resulting ages. The procedure of Bayesian sequencing (or “multi-sample calibration” more generally spoken) is given in the following in its basic form. (Since the calculations are performed for the particular determined set of ^{14}C ages, a simplified notation without indicating the ^{14}C ages is used.) The method was introduced by Buck et al. (1991, 1992) and is described in detail in Buck et al. (1996); a description focusing on the practical realization is given by Bronk Ramsey (2009). Weninger et al. (2006) may be an adequate basic description for readers that are not familiar with the field.

For a given set of samples, first—as in single-sample calibration—each determined ^{14}C age is transformed into a probability density distribution for the real sample age. This distribution is called the single-sample likelihood function and characterizes the probability distribution for the real sample age, based only on the information from the measurements so far. (This is only true under the assumption that any real age is previously equally probable. Actually, as mentioned above, the likelihood function is the conditional probability density to obtain a particular ^{14}C age if a sample of given real age is measured.) In its most simplified form (neglecting the error of the calibration curve and without taking account of standardization), the single-sample likelihood function l_i is given by Equation 1:

$$l_i(t) \propto \exp\left(-\frac{(x_i - c(t))^2}{2 \cdot \sigma_i^2}\right) \quad (1)$$

where x_i is the determined ^{14}C age of the i -th sample and σ_i its uncertainty; t is the unknown real age of the sample; and $c(t)$ gives the calibration curve.

The single-sample likelihood functions are combined to a multidimensional likelihood function $l(t_1, \dots, t_n)$ that gives the probability density for any particular combination of the real sample ages t_1, \dots, t_n , still based on the measurements only; see Equation 2:

$$l(t_1, \dots, t_n) \propto l_1(t_1) \cdot \dots \cdot l_n(t_n) \quad (2)$$

The available *a priori* information from archaeology is introduced as a further n -dimensional function, the prior function $a(t_1, \dots, t_n)$, which is the probability density considering exclusively the archaeological information.

The product of both gives the so-called (multidimensional) posterior function $p(t_1, \dots, t_n)$ (see Equation 3), which finally gives the probability density for any particular sample age combination, now considering both the information from the ^{14}C measurement and the archaeological information (deduced e.g. from stratigraphy).

$$p(t_1, \dots, t_n) \propto l(t_1, \dots, t_n) \cdot a(t_1, \dots, t_n) \quad (3)$$

The probability distributions of the individual real sample ages, the so-called marginal posterior densities $p_i(t_i)$, are calculated by projecting the posterior probability to the individual sample coordinates, integrating over all other coordinates (excluding that of the sample actually evaluated), as shown in Equation 4.

$$p_i(t_i) \propto \int_{-\infty}^{+\infty} \dots \int_{-\infty}^{+\infty} p(t_1, \dots, t_n) dt_1 \dots dt_{i-1} dt_{i+1} \dots dt_n \tag{4}$$

These resulting probability distributions can be, similar to the case of single-sample calibration, reduced to highest posterior density ranges (hpd ranges) by collecting the most probable years until the required confidence level (e.g. 95%) is reached.

Equations 1 to 4 are sufficient to perform Bayesian sequencing in its straightforward form. Although the numerical integration imposes some difficulties, because there are too many points in the multi-dimensional space even for a very low-resolution grid, it can be performed well with the help of Monte Carlo methods.

Frequently, additional statistical parameters beyond the sample ages t_i are used, mainly boundaries of archaeological phases. This can be implemented without difficulty, but will not be shown in this brief introduction.

To perform our investigations, we developed a program based on the mathematical language MATLAB (The MathWorks, Inc., Natick, Massachusetts, USA), which runs the calculations by Gibbs sampling, which is a very basic Markov-chain Monte Carlo method; see e.g. Gilks et al. (1996) or Krause (1994).

The simplified notation used here to demonstrate the method most clearly somewhat hides Bayes’ theorem on which the method is based. The Bayes’ theorem—expressed above by Equation 3—looks as follows when given in exact notation (Equation 5):

$$p(t_1, \dots, t_n | x_1, \dots, x_n) = \frac{l(x_1, \dots, x_n | t_1, \dots, t_n) \cdot a(t_1, \dots, t_n)}{\int_{-\infty}^{+\infty} \dots \int_{-\infty}^{+\infty} l(x_1, \dots, x_n | t_1, \dots, t_n) \cdot a(t_1, \dots, t_n) dt_1 \dots dt_n} \tag{5}$$

In this notation, one can see that the likelihood function, which is the conditional probability density to get the set of ¹⁴C ages for a given set of real ages, is converted into a conditional probability density for a set of real ages for given ¹⁴C ages—what we are looking for—with the help of the prior probability density. Detailed descriptions of Bayes’ theorem and its statistical foundations can be found e.g. in Jeffreys (1961) (theoretical background) and Sivia (1996) (tutorial).

THE AMBIGUITY PROBLEM AND THE BASIC IDEA OF ROBUST BAYESIAN ANALYSIS

A critical step within the procedure of Bayesian multi-sample calibration is the transformation of the archaeological information into a prior function, which has to be a probability density in the real age space $a(t_1, \dots, t_n)$. Usually, the archaeological facts do not determine the shape of this function in an unambiguous way. This means, for example, that a given set of time relations deduced from stratigraphy can be described by various differently shaped prior functions. A simple example can illustrate the nature of the problem: The knowledge that, for example, sample B is older than sample A

is described correctly by both functions displayed in Figure 1, which shows the probability densities for the age difference $\text{age}(B) - \text{age}(A)$.

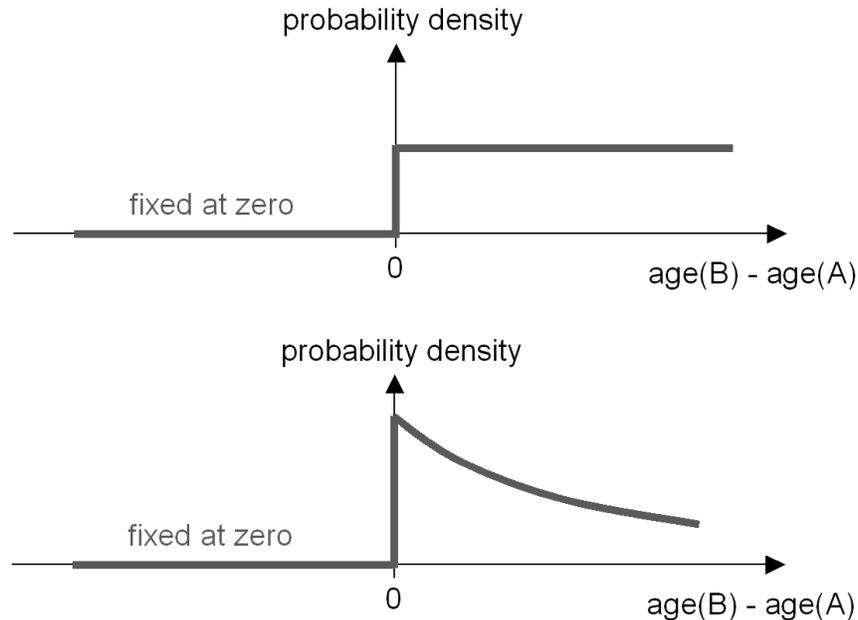


Figure 1 Both prior functions are in full agreement with the prior information that sample B is older than sample A. There is no way to prefer one of them based only on this information.

One knows that age B must not be younger than age A, so the prior function has to be zero on the negative left part of the axis. But any shape of the function one chooses on the positive right side defines a particular probability density for any given age difference. This choice cannot be based on available information, because it is only known that sample B is older than A and nothing more. However, the procedure of Bayesian sequencing as given above requires the use of one particular prior function. That means that Bayesian modeling has to assume information that is actually without foundation. The problem is discussed in previous studies (e.g. Bronk Ramsey 2000; Steier and Rom 2000; Steier et al. 2001).

It should be briefly mentioned here that there are methods to find priors that bias the result as little as possible by reducing their unwanted information content. These priors are then called “non-informative” priors. One way to minimize unwanted information within a prior function is to characterize its non-informative behavior by an entropy measure (principle of maximum entropy; see e.g. Sivia 1996). However, we do not follow this approach in the present work.

A very general approach to deal with the ambiguity problem is “robust Bayesian analysis.” Robustness in Bayesian analysis has become a theoretically complex field. The theory (summarized by Berger 1994; Rios Insua and Ruggeri 2000) is not finally settled, as there are ongoing developments and discussions (e.g. Berger 2006). Thus, we do not try to analyze the problem in a theoretical way, but only introduce the following particular concept to the archaeological application.

The basic idea is to use a (theoretically infinite) set of prior functions including all possible shapes consistent with the available information (Figure 2). The Bayesian calculations are performed with

each prior function individually and the final result is the union of the individual results, that is, the union of the highest posterior density ranges (hpd ranges) to a chosen confidence level. By this, one gets hpd ranges for the sample ages that are valid for all possible prior functions and, therefore, independent of a subjective choice of a particular function. The meaning of these unified hpd ranges is then a bit different from the usual hpd ranges: The statement “the age falls into the interval with a probability of 95%” is changed to “the age falls into the interval with a probability of at least 95%.” Of course, the first statement for the usual interval is only true if the correct prior was used; the second statement for the interval from robust analysis is true if the correct prior was among the used set.

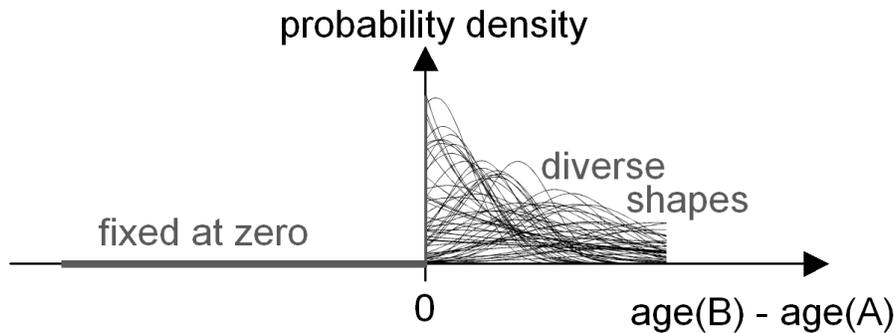


Figure 2 The principle of “robust Bayesian analysis” using all possible shapes of priors that are consistent with the available information: sample B is older than sample A.

Although robust analysis is very simple in principle, there are some difficulties in the mathematical concept and also in the numerical methods that will be discussed in this work. Before entering into methodical and technical details, a comparison between usual sequencing and robust analysis by the use of a simple example shall clarify the reason why there is a need for improvement.

AN ILLUSTRATIVE EXAMPLE: DATING THE ICEMAN AND HIS AXE

In 1991, the famous, very well-preserved body of an Early Bronze Age man was found in the European Alps on the Italian side of the Austrian-Italian border in the Ötztaler Alps, released by a melting ice shield. Samples from the body were ^{14}C dated at the accelerator mass spectrometry (AMS) laboratories in Oxford and Zürich (see Hedges et al. 1992; Bonani et al. 1994). The actual value used in our work (see Figure 3) is taken from Kutschera and Rom (2000), which is a combined age based on data including tissue as well as bone samples. (A possible effect of inbuilt age in the bone samples is neglected for this example.) Nearby the Iceman, various parts of his equipment were found and many samples were dated at the AMS lab in Vienna (see Rom et al. 1999; Kutschera and Müller 2003).

To demonstrate the differences of usual sequencing and robust analysis, the direct dates of the Iceman and the dates of the wooden shaft of his axe are evaluated with a simple model. The determined mean ^{14}C dates and their unmodeled calibration are shown in Figure 3.

If one accepts that the axe found close to the Iceman was actually used by him, then the wood of the axe shaft cannot be younger than the Iceman himself. Maybe one could get more detailed information by further archaeological analysis, which could then be used to construct an improved prior function: e.g. that an axe shaft typically may not be produced of very young wood, or similar considerations. But as such investigations have not been carried out by a qualified expert in this case,

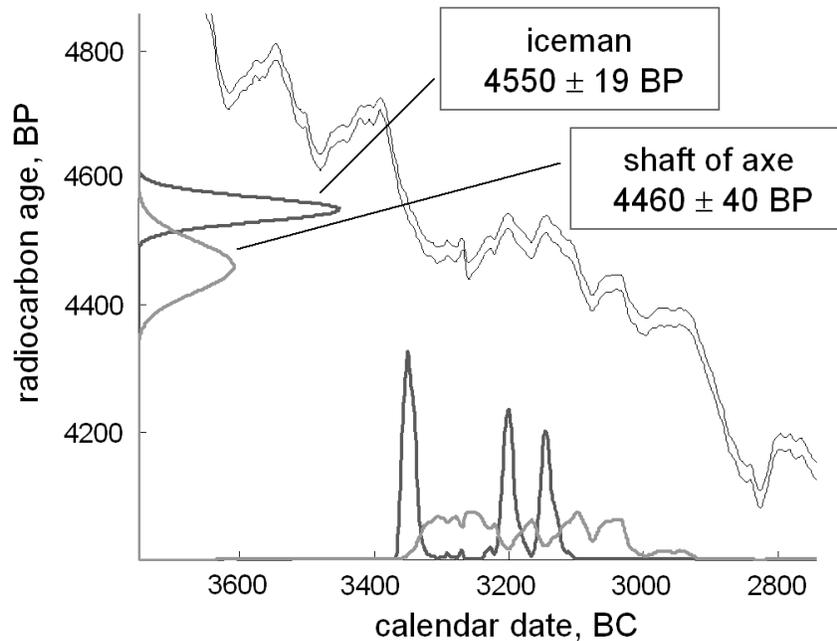


Figure 3 Mean ¹⁴C ages and corresponding unmodeled calibrations of the Iceman’s body and of the wooden shaft of the axe. The calibration curve (Reimer et al. 2004) is given with 2 lines indicating the 1-σ precision band.

the only reliable prior information is the fact that the shaft has to be older than or of equal age as the Iceman.

Uniform Prior

The simplest way to formulate the knowledge from above in functional form is the “uniform prior,” which is zero for age combinations with a wrong chronological order, and constant for all allowed cases:

$$a_{uniform}(t_{man}, t_{axe}) \propto \begin{cases} 1 & \text{if } t_{axe} \geq t_{man} \\ 0 & \text{if } t_{axe} < t_{man} \end{cases}$$

It can be shown that in this case where only 2 ages are involved, the uniform prior implies also equal initial probabilities for all age differences $t_{axe} - t_{man}$, where $t_{axe} - t_{man} \geq 0$ (correct chronological order). Assuming equal probability for any age difference seems reasonable (except for very large age differences) as there is no knowledge whether the shaft is made from young or old wood. However, there are theoretical aspects to favor decreasing functions (exponential, $1/x$) as non-informative priors for positive definite numbers (see e.g. Jeffreys 1961).

Uniform Span Prior

An alternative possible prior function results from the assumption that the ages of man and axe lie within a particular possible time span with unknown length. This time span is modeled by 2 outer boundaries (t_{old} and t_{young}) and the prior function is constructed in such a way that any age difference between these boundaries has the same initial probability. This prior function, denoted as the uni-

form span prior in the following, is mentioned because in a generalized form, it is the commonly used prior for conventional archaeological sequences (see Bronk Ramsey 2001). It can be shown that its functional form is this:

$$a_{unif.span}(t_{man}, t_{axe}, t_{old}, t_{young}) \propto \begin{cases} 1/(t_{old} - t_{young})^2 & \text{if } t_{old} \geq t_{axe} \geq t_{man} \geq t_{young} \\ 0 & \text{otherwise} \end{cases}$$

Prior Set

Our preferred approach is robust analysis. The finite prior set used as an approximation for the infinite set of all possible functions is of the following simple form:

$$a_{robust,\alpha}(t_{man}, t_{axe}) \propto \begin{cases} e^{(t_{axe} - t_{man})/\alpha} & \text{if } t_{axe} \geq t_{man} \\ 0 & \text{otherwise} \end{cases}$$

where the parameter α steps through negative values simulating probabilities decreasing with increasing age difference with various slopes, and also through positive values simulating increasing probabilities. The uniform prior ($\alpha = \infty$) is also included.

Figure 4 gives a comparison of the outcome using the 2 different single priors with that of robust analysis. Always, the hpd ranges at 95.4% confidence level are shown.

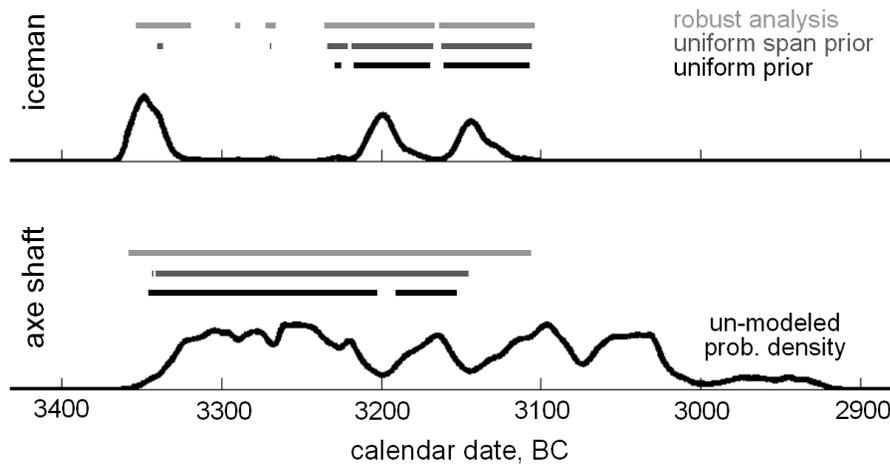


Figure 4 Comparison of the hpd ranges at a 95.4% confidence level resulting from 2 different single prior models (uniform prior and uniform span prior) and from robust Bayesian analysis. The unmodeled calibrations are given as probability density distributions.

The difference between robust Bayesian analysis and the use of a single prior can be seen most clearly when we consider the possibility that the axe shaft could be of the same age or only a few years older than the Iceman. (This is included in our prior set for robust analysis by priors with $\alpha \approx 10$ yr or less.) In that case, Iceman ages at the oldest peak of the single-sample calibration of the Iceman are possible, because the unmodeled probability densities overlap considerably. In full accordance with this fact, robust analysis includes the oldest peak of the Iceman's probability density within the 95.4% hpd range. In contrast, when using the uniform prior, Iceman ages at the oldest

peak are misleadingly excluded, although the uniform prior seems a reasonable choice for this case as mentioned above. The result with the uniform span prior is a bit closer to that of robust analysis; the oldest peak is not excluded totally as with the uniform prior. However, the uniform span prior is still only 1 particular choice out of an infinite number of possible shapes, and so it can still exclude possible age bands, as indicated by comparison with the result of robust analysis.

Looking at these different results, one could get the impression that the results of Bayesian sequencing are in some way arbitrary and depend only on the choice of the prior. This was a reason for us to introduce robust analysis. But on the other hand, Figure 4 also shows results that are independent from the choice of the prior: ages for the axe shaft younger than about 3100 BC are not possible. Robust analysis preserves this common improvement of all Bayesian methods over single-sample calibration.

Hpd-Range Envelopes Instead of Probability Densities

The final result of robust Bayesian analysis is only available in the form of highest posterior density ranges (hpd ranges). There are no resulting marginal posterior probability densities calculated, only the hpd ranges for the various different priors are unified. However, density curves bear valuable information, and they are appreciated by the users of the method. We therefore have developed analogue curves for robust analysis. The unified hpd ranges can be calculated for all confidence levels. Enveloping all these hpd ranges results in a continuous function again, which offers information comparable with the marginal posterior density. The “hpd-range envelopes” are shown in Figure 5 for the current example.

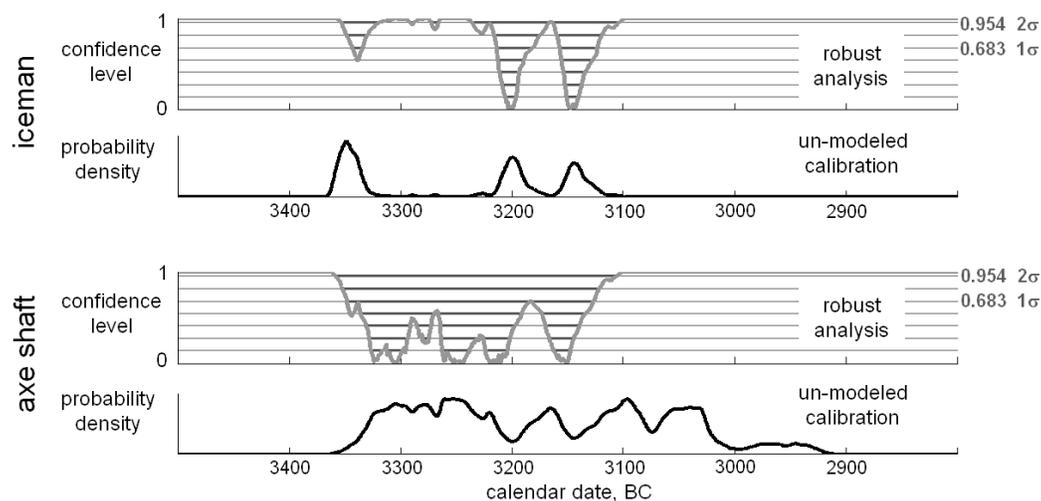


Figure 5 Resulting “hpd-range envelopes” from robust analysis of the Iceman example. These curves (light gray) envelope the hpd ranges at all different confidence levels. For clarification, hpd ranges to some particular levels (thin lines) are also given (thick lines). The black curves give again the un-modeled probability densities as in Figure 4. To avoid misunderstandings, the hpd-range envelopes are not probability density distributions; they only show the limits of the hpd ranges to a corresponding confidence level.

DISCARDING INCOMPATIBLE PRIOR FUNCTIONS

The idea of robust analysis in principle is the use of all shapes of priors that are consistent with the available information. The problem of this theoretical concept is that one can always find extremely

shaped prior functions that, although still consistent with the given information, can “damage” the result by producing posterior probabilities that are in disagreement with the measurements. These are usually priors that have a strong “bias” for a specific result, essentially the opposite of the “non-informative” priors mentioned above. If we take the Iceman example, and assume a prior function that gives a very high probability density for age differences between 1000 and 1100 yr (axe older than man) and very low probability for all other age differences, this prior forces the posterior of either the Iceman or the axe shaft to lie far apart, completely out of the range obtained by unmodeled calibration. But still this prior is not in disagreement with the (only used) prior information that the axe shaft has to be older (or of equal age) than the Iceman. Less artificial, even a prior decreasing exponentially with increasing age differences (as used in the example above) can damage the result if the slope is much steeper as consistent with the real, but unknown, historical situation.

The literature suggests “prior elicitation,” i.e. manual identification and rejection of such priors from the prior sets (see e.g. Berger 1994). However, we think that this is an error-prone and cumbersome procedure, which would prevent broad acceptance of the method. The manual rejection of priors could be used to “tune” the result towards the expectations of the user, and thus introduces a unwanted subjectivity. We are looking for a more objective and automatic way to identify corrupt prior functions, focusing on the agreement of the prior with the measured ^{14}C ages (assuming that the ^{14}C measurements are correct). As a prerequisite, one has to find a good measure for the agreement of model and data. This will typically be a so-called agreement index, i.e. a number calculated by some mathematical formalism, and a threshold for this number, below which the prior is rejected.

Discarding corrupt prior functions is essential for using robust Bayesian sequencing in practice. However, it must not be ignored that this could be seen as a mutilation of the pure theoretical concept: the measured ^{14}C data now enters the mathematical procedure from 2 sides, through the reduction of the prior set and through the likelihood functions. Additionally, the choice of the agreement criterion and the threshold level could introduce again a subjectivity similar to the choice of the prior for the “non-robust” multisample calibration. We thus have put significant effort into the development of this criterion.

Measuring the Agreement of Model and Data

The most fundamental measure for the agreement of model (i.e. prior) and determined ^{14}C ages is the so-called prior predictive probability distribution, denoted as “prior prediction” in the following. It occurs as a “standardization term” in the denominator in Equation 5—the Bayes’ theorem—and is given by the multidimensional volume (therefore, denoted v) of the product of likelihood and prior function:

$$v(x_1, \dots, x_n) = \int_{-\infty}^{+\infty} \dots \int_{-\infty}^{+\infty} l(x_1, \dots, x_n | t_1, \dots, t_n) \cdot a(t_1, \dots, t_n) dt_1 \dots dt_n \quad (6)$$

The prior prediction $v(x_1, \dots, x_n)$ is the probability density to determine a particular set of ^{14}C ages for a given prior function $a(t_1, \dots, t_n)$. In the specific application here, one is only interested in 1 particular value of v , that of the actually measured set of ^{14}C values. This value can be compared for different priors used. The ratio of the v values gives the ratio of the probabilities to get the measured data for different *a priori* probability densities for the real ages. This ratio is known as the Bayes’ factor and established for model comparison (see e.g. Garcia-Donato and Chen 2005), and has been already used for archaeological applications (Jones and Nicholls 2002). Priors with a higher infor-

mation content (i.e. more stratigraphic information) are preferred against unspecific ones, since they are less likely to fit to the measured data by chance.

Before we give the formalism to discard prior functions that do not agree with the data with the help of the Bayes' factor, an additional problem has to be solved: In practical applications of Bayesian sequencing, technical difficulties arise. Often, priors are used that are significantly different from zero on an unrestricted domain and thus cannot be standardized (they have an infinite integral). Examples are both the uniform prior and the uniform span prior mentioned above. From a theoretical point of view, these priors are forbidden. However, the infinite domain does not affect the usual multisample calibration procedure, since the tails of the likelihood functions used in ¹⁴C calibration decrease very fast to zero (essentially proportional to exp(-x²)). Therefore, unrestricted priors are used in almost all examples in the literature, and we have to be able to cope with them. However, a proper definition of their prior prediction is impossible (the standard definition results in a factor of zero, i.e. they do not fit to any measured data). To circumvent the problem of standardization, the integration could be confined to a restricted domain. Unfortunately, the prior prediction obtained this way will depend (in indirect proportion) on the volume of this arbitrarily chosen domain.

Our (disputable) approach to calculate the prior predictions is to transform the unrestricted priors into standardizable ones, by multiplying them with a standardizable weighting function. This can be understood also as restricting the integration to a domain, but with gradual boundaries instead of sharp edges. The function should be significantly different from zero in the region where the ¹⁴C measurements are situated, but should not extend much farther, since this will reduce the prior prediction by increasing the "domain." Consequently, we use the sample ages measured, and their respective likelihood functions, to define the weighting function. The straightforward approach is to simply add all likelihood functions of the individual samples and use this as the weighting function for all coordinates. The total *n*-dimensional weighting function λ(t₁, ..., t_n) is given by Equation 7 (in the following, the simplified notation for the fixed set of determined ¹⁴C ages, not indicating the ¹⁴C ages explicitly, is used again):

$$\lambda(t_1, \dots, t_n) \propto \prod_{j=1}^n \left(\sum_{i=1}^n \left(\frac{l_i(t_j)}{\int_{-\infty}^{+\infty} l_i(t) dt} \right) \right) \tag{7}$$

where *l_i* are the single-sample likelihood functions as defined in Equation 1.

Since all priors can be standardized now and the prior predictions can be calculated, it is possible to define an agreement factor *B* (which is a Bayes' factor as mentioned above) that describes quantitatively the agreement of the model and data regarding to a reference prior *a_{ref}*, see Equation 8:

$$B = \frac{v^{(\lambda)}}{v_{ref}^{(\lambda)}} = \frac{\int_{vol} l \cdot \frac{a \cdot \lambda}{\int_{vol} a \cdot \lambda dt'} dt}{\int_{vol} l \cdot \frac{a_{ref} \cdot \lambda}{\int_{vol} a_{ref} \cdot \lambda dt'} dt} = \frac{\int_{vol} l \cdot a \cdot \lambda dt / \int_{vol} a \cdot \lambda dt}{\int_{vol} l \cdot a_{ref} \cdot \lambda dt / \int_{vol} a_{ref} \cdot \lambda dt} \tag{8}$$

where ∫_{vol} ... dt is a volume integration over all age dimensions (we generally denote vectors with bold letters); *a*, *a_{ref}*, and λ denote non-standardized functions here; the needed standardizations are explicitly given in the equation. One can see that the prior functions *a* and *a_{ref}* need not to be stan-

standardizable any more; it is sufficient if $a \cdot \lambda$ and $a_{ref} \cdot \lambda$ have finite integrals, which is usually granted by the fast decrease of the likelihood functions.

When using a constant reference prior, the agreement factor B is the Bayes' factor that compares the prior prediction $v^{(\lambda)}$, calculated with a prior that combines the tested prior function with the weighting function, with the prior prediction $v^{(\lambda)}_{ref}$, using a prior that is the weighting function only. Slightly simplified, the value gives the increase ($B > 1$) or decrease ($B < 1$) of the probability to get the determined set of ^{14}C ages resulting from the use of the prior information. Therefore, B is a good quantitative measure of the agreement of the prior function with the ^{14}C ages.

To illustrate the potential of the agreement factor B , we look at an ordered sequence of n samples, which are separated sufficiently in time so that the likelihood functions do not overlap. Again, the constant function is used as reference prior, so for the constant prior B is equal to 1. If one tests the uniform prior for a sequence of samples (which is 1 for age combinations that are in the order given by the prior information, and zero elsewhere), then 2 results are possible: If the assumed prior order is in disagreement with the order of the ^{14}C ages, the agreement factor would be zero (for non-overlapping likelihoods). If the assumed prior order is in agreement with the order of the ^{14}C ages, the resulting value is the faculty of n . This result is reasonable, because $1/n!$ is the probability to get the right order by chance.

It has to be acknowledged that the restriction of the domain by the function λ as given by Equation 7 is not the only possible way. A different definition (e.g. including also all regions between determined ^{14}C ages) would lead to different results for the agreement factor B . However, we see no way to compare priors with finite and infinite integrals in an unambiguous way.

It should be noted that there is a relation of the agreement factor B to an agreement measure used in the well-known program OxCal (Bronk Ramsey 1995, 2001, 2009): If one would use the normal multidimensional likelihood function as a weighting function, by replacing λ with l (and still using a constant a_{ref}), B would become the same as F_{model} in OxCal.

Finding a Reasonable Threshold Level

The threshold level of B for rejecting a model should be chosen relative to that of a model that is in good agreement with the real ages. A “perfect” model per definition is one that knows the real ages of the samples and constrains the prior function sharply around them. Within this “perfect” model, only the statistical variances of the measurements reduce the agreement factor. Thus, one can use this case to get an estimate for the threshold of B so that the model is always accepted, except for a allowed percentage of statistical outliers. We estimate the threshold for the most simplest situation considering n samples, all with the same real age and measured with the same accuracy, generating a n -dimensional Gaussian likelihood function (assuming a linear calibration curve). The posterior function resulting from the “perfect” model is a n -dimensional δ function at the position of the set of the real sample ages. For this simple case, the corresponding agreement factor B , assuming a constant reference prior, can be analytically calculated and is denoted $B_{perfect}$:

$$B_{perfect} = 2^{\frac{n}{2}} \cdot e^{-\frac{1}{2} \cdot \sum_i \left(\frac{x_i - x^{(0)}}{\sigma} \right)^2} \tag{9}$$

where x_i are the measured ^{14}C ages with their common value of uncertainty σ ; $x^{(0)}$ is the ^{14}C age corresponding to the common real age of the samples when assuming no measurement error. The sum in Equation 9 is the well-known chi-squared (χ^2) distribution. Analogous to a conventional χ^2 test,

for a certain confidence level P (e.g. 95.4%), we obtain a corresponding threshold level. The threshold depends on the number of samples n :

$$B_{threshold} = 2^{\frac{n}{2}} \cdot e^{-\frac{1}{2} \cdot \chi^2(P, n)} \quad (10)$$

So with $B_{threshold}$, a characteristic threshold level has been estimated that accepts models that agree with the real sample ages, in most cases of possible measured ^{14}C data sets.

The Actual Applied Agreement Factor

The power of using a Bayes' factor as an agreement measure is the fact that it compares prior predictions that are absolute measures for the quality of a prior. Unfortunately, this absolute character can be a disadvantage in practical applications. More often than not, there are outliers (beyond the statistical variance) in the ^{14}C measurements and/or mistakes in the model definition (e.g. a sample is misplaced in the stratigraphy). Both would lead to a general decrease of the agreement factors (when defined with a constant reference prior) for all individual priors used, a behavior that is for sure correct. However, in our method, now more priors would be excluded (in the worst case, all), leading to an artificially small prior set. To bypass this problem, one has to use a reference prior in the definition of the agreement factor (Equation 8) that is as well inflected of the described faults but represents a model that is for sure not extreme, so that it should be discarded itself. Our approach is the use of the uniform prior (a_{unif}) as a reference prior, which achieves the mentioned requirements. So the actual applied agreement factor is denoted $B^{(unif)}$ and is given by Equation 11 (according to the definition in Equation 8):

$$B^{(unif)} = \frac{\int_{vol} l \cdot a \cdot \lambda \, dt / \int_{vol} a \cdot \lambda \, dt}{\int_{vol} l \cdot a_{unif} \cdot \lambda \, dt / \int_{vol} a_{unif} \cdot \lambda \, dt} \quad (11)$$

The threshold defined in Equation 10 is used unchanged for $B^{(unif)}$. This is reliable, because in the considerations to estimate a threshold level, the "perfect model" was compared with the "neutral model" represented by the constant reference prior. In the actual applications, all priors in the set have to agree with the given stratigraphic constraints, so that the uniform prior can be seen as a "neutral prior." Naturally, this is just a rough estimate, because as pointed out earlier in this work, since neither the constant nor the uniform prior are really "neutral" in the sense of carrying no information or only the information of the constraints, respectively. However, the choice of the threshold cannot be objective anyway; we can only define a consistent procedure.

Equations 11, 10, and 7 establish the needed system for discarding corrupt prior functions that would damage the result of robust Bayesian analysis. It should be noted that the definition of the agreement factor can be generalized for the use of additional statistical parameters like phase boundaries as well.

An additional aspect that has to be mentioned shortly is the fact that numerical evaluation of the multidimensional volume integrals in Equation 11 is not trivial. Actually, we were not able to find an applicable method in the literature. We thus developed a method based on a comparison of the volume to be estimated with a reference volume, executed by a modified kind of Gibbs sampling. The method is sufficient up to about 30 samples or dimensions, but leads to convergence problems beyond that number.

A Tentative Procedure for Large Sequences

Due to the numerical problems mentioned above, for large sequences we tentatively base the criterion to discard corrupt priors on the single-sample agreement indices. For the 2-dimensional case, l and λ are the same (using Bayes’ theorem; Equation 5), Equation 11 can be expressed as

$$B_{onedim}^{(unif)} = \frac{\int_{-\infty}^{\infty} l \cdot p \, dt}{\int_{-\infty}^{\infty} l \cdot p_{unif} \, dt}$$

where p and p_{unif} are the posterior functions. So for the 1-dimensional case, the degrees of overlap of likelihood and posterior function are compared. For the multidimensional case, individual single-sample agreement indices are defined analogously by using the single-sample likelihood functions l_i and the marginal posterior densities p_i (see Equations 1 and 4):

$$B_i^{(unif)} = \frac{\int_{-\infty}^{\infty} l_i \cdot p_i \, dt}{\int_{-\infty}^{\infty} l_i \cdot p_{i,unif} \, dt}$$

$\int l_i p_i dt / \int l_i dt$ is the usual definition of the single-sample agreement index A_i , see Bronk Ramsey 1995; the second part is canceled out in the equation above. Similar considerations as above show that also for these single-sample agreement indices, the threshold from Equation 10, evaluated for $n = 1$, is reliable (i.e. $\sqrt{2}/e^2$ for $P = 0.954$ “ 2σ ”). We accept a prior if none of the $B_i^{(unif)}$ drops under that level. This has turned out to be more efficient than using an “overall index” deduced from the individual agreement indices. However, the use of a single-sample-based agreement criterion is not what we finally intend, and it should be replaced by the multidimensional agreement factor (Equation 11), as soon as we are able to calculate the latter for large sequences.

PARAMETRIC PRIORS AND MODEL AVERAGING

To perform robust Bayesian analysis, calculations have to be performed individually for each function from a complete set of prior functions. Theoretically, there is an infinite number of possible prior functions, but one can find ways to reduce the set to a practical manageable number (see section “Second Example: A Large Archaeological Sequence”). However, it would be advantageous if there was a method that could make the use of an infinite prior set actually possible. The obvious idea is to use a set of prior functions of the same mathematical form, but generated by a set of free parameters. These parameters are then treated as additional dimensions within the formalism, simulating various shapes of prior functions by only a single evaluation of the posteriors. For example, a prior set with exponential decreasing probability densities for an age difference (a similar set with decreasing and increasing functions was used in the Iceman example above) is generated by the following parameterization:

$$prior = \frac{1}{\alpha} \cdot e^{-(t_B - t_A)/\alpha} \text{ for } t_B \geq t_A, \text{ and } 0 \text{ otherwise; } (\alpha > 0)$$

It can be shown that Gibbs sampling with this parametric prior function leads to a summation of the posterior probability densities for all possible shapes of the prior function due to various parameter

values, automatically weighted by the prior prediction. Therefore, the method would have the potential to suppress corrupt shaped priors intrinsically by weighting them low. However, there are fundamental differences compared to the independent calculations for the various different prior functions.

First, the generating parameters now need a prior function, too. This brings along all the problems of the choice of a proper prior, e.g. that the result now depends on the scaling of the parameters, which changes with the form of the parameterization, for example:

$$\frac{1}{\alpha} \cdot e^{-(t_B-t_A)/\alpha} \text{ or } \beta \cdot e^{-(t_B-t_A) \cdot \beta}$$

describe exactly the same set of functions, but give different results, when assuming a constant prior for the parameter α or β , respectively.

Second, this approach is mathematically equivalent to the calculation with a certain, single prior, which can be obtained by integrating first over all prior parameters. While one could argue that this effective prior might now be better—less informative—than the commonly used priors, the search for non-informative priors is not the target of our research. The equivalence to a single prior can be described as follows.

The final posterior probability density $p(\mathbf{t})$ when using a set of free prior parameters α , results from integrating the posterior density including the parameters $p(\mathbf{t}, \alpha)$ (a probability density within the combined vector space of sample ages and prior parameters) along all parameter dimensions (projection to the sample age subspace), see Equation 12. (Once again, the simplified notation without indicating the set of ^{14}C ages is used.)

$$p(\mathbf{t}) = \int_{vol} p(\mathbf{t}, \alpha) d\alpha \tag{12}$$

Where $p(\mathbf{t}, \alpha)$ results from Bayes' theorem:

$$p(\mathbf{t}, \alpha) = \frac{l(\mathbf{t}) \cdot a(\mathbf{t}, \alpha)}{\int_{vol} l(\mathbf{t}) \cdot a(\mathbf{t}, \alpha) dt d\alpha} \tag{13}$$

Thus, the final posterior function $p(\mathbf{t})$ when integrating over the parameters results in:

$$p(\mathbf{t}) = \int_{vol} \frac{l(\mathbf{t}) \cdot a(\mathbf{t}, \alpha)}{\int_{vol} l(\mathbf{t}) \cdot a(\mathbf{t}, \alpha') dt d\alpha'} d\alpha = \frac{l(\mathbf{t}) \cdot \int_{vol} a(\mathbf{t}, \alpha) d\alpha}{\int_{vol} l(\mathbf{t}) \cdot (\int_{vol} a(\mathbf{t}, \alpha') d\alpha') dt} \tag{14}$$

Denoting the integral of the prior function along the parameters—which is a particular function in \mathbf{t} again—as effective prior function $a_{eff}(\mathbf{t})$, the final posterior results from a single application of Bayes' theorem again:

$$p(\mathbf{t}) = \frac{l(\mathbf{t}) \cdot a_{eff}(\mathbf{t})}{\int_{vol} l(\mathbf{t}) \cdot a_{eff}(\mathbf{t}) dt} \tag{15}$$

Calculation with parametric priors is therefore equivalent to calculating with a single effective prior, and its shape depends only on the chosen kind of parameterization.

Equation 12 can also be interpreted as a weighted sum of posteriors within the normal space of sample ages (when the parameters are not treated as statistical variables for Bayes' theorem), denoted as $p^{(t)}(\mathbf{t}, \boldsymbol{\alpha})$. The weighting factors are the prior predictions within the normal sample age space, which can be seen in the following.

In the sample age space, Bayes' theorem can be expressed as:

$$p^{(t)}(\mathbf{t}, \boldsymbol{\alpha}) = \frac{l(\mathbf{t}) \cdot a(\mathbf{t}, \boldsymbol{\alpha})}{\int_{vol} l(\mathbf{t}) \cdot a(\mathbf{t}, \boldsymbol{\alpha}) dt} \tag{16}$$

Unlike Equation 13, there is no integration over the parameters, and the prior prediction remains parameter-dependent. Combining Equations 13 and 16 results in:

$$p(\mathbf{t}, \boldsymbol{\alpha}) = p^{(t)}(\mathbf{t}, \boldsymbol{\alpha}) \cdot \frac{\int_{vol} l(\mathbf{t}) \cdot a(\mathbf{t}, \boldsymbol{\alpha}) dt}{\int_{vol} l(\mathbf{t}) \cdot a(\mathbf{t}, \boldsymbol{\alpha}) dt d\boldsymbol{\alpha}} = p^{(t)}(\mathbf{t}, \boldsymbol{\alpha}) \cdot \frac{v^{(t)}(\boldsymbol{\alpha})}{v} \tag{17}$$

where $v^{(t)}(\boldsymbol{\alpha})$ is the parameter-dependent prior prediction within the normal age space and v is the prior prediction in the combined age and parameter space, which is a constant independent of $\boldsymbol{\alpha}$. So putting $p(\mathbf{t}, \boldsymbol{\alpha})$ into Equation 12, one gets Equation 18, an integration over the individual posteriors weighted with $v^{(t)}(\boldsymbol{\alpha})$:

$$p(\mathbf{t}) = \int_{vol} p^{(t)}(\mathbf{t}, \boldsymbol{\alpha}) \cdot \frac{v^{(t)}(\boldsymbol{\alpha})}{v} d\boldsymbol{\alpha} \tag{18}$$

It should be remarked here that this kind of weighted summation of the posteriors for different priors is in principle a method already used and denoted as "model averaging" (see e.g. Hoeting et al. 1999). It can be seen as a kind of robust analysis; however, it has to cope with the a significant problem: Equation 18 shows clearly that the weighting depends unfortunately not only on the fixed value of $v^{(t)}(\boldsymbol{\alpha})$ for a particular parameter set $\boldsymbol{\alpha}$, but also on the arbitrary definable scaling of the parameters, mentioned above. A change to another parameter set $\boldsymbol{\beta}(\boldsymbol{\alpha})$ would require to introduce factors like $\partial\alpha_i/\partial\beta_j$ to keep the integration constant; however, no such factors will be present if the set $\boldsymbol{\beta}$ is used from the beginning.

Thus, robust Bayesian analysis that wants to check actually all possible prior functions in our understanding has to calculate each prior function individually. We therefore have to reduce the theoretical infinite prior set to a manageable finite one, which is discussed within the archaeological example demonstrated below.

SECOND EXAMPLE: A LARGE ARCHAEOLOGICAL SEQUENCE

In this section, robust Bayesian sequencing and the usual Bayesian method will be compared using a sequence taken from Manning et al. (2006). This large sequence, which links together related contexts of various sites, is very well documented and therefore highly suitable as an example. The left side of Figure 6 gives the structure of the sequence in a slightly simplified illustration. The sequence

consists mainly of 4 phases; some of them include substructures. Two tree-ring wiggle-matches are used as *terminus post quem* (TPQ). The sequence includes about 100 ¹⁴C data; many of the shown ages of contexts are based on averages of several ¹⁴C data.

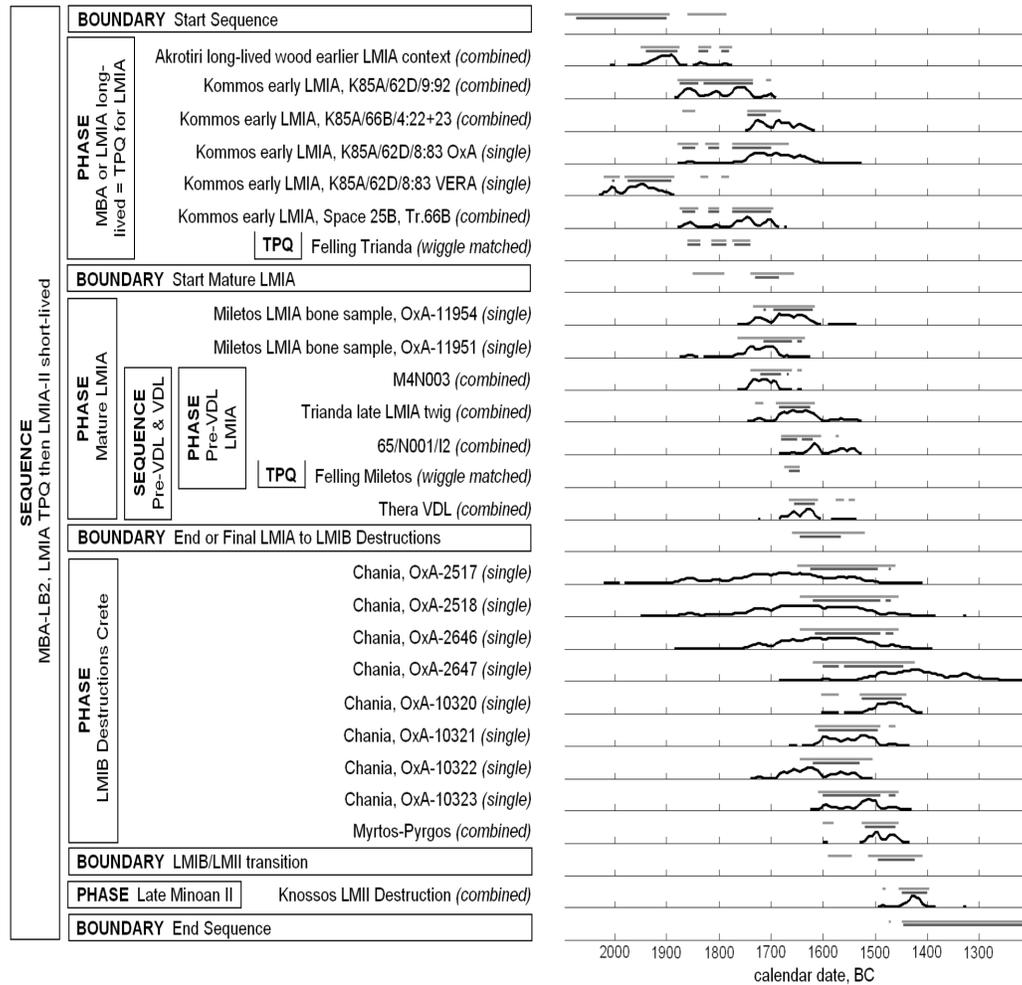


Figure 6 Second example: a large archaeological sequence taken from Manning et al. (2006). A comparison of the hpd ranges at a 95.4% confidence level resulting from robust Bayesian analysis (light gray bars) and from the usual uniform span prior (dark gray bars) is given. Unmodeled calibrations are given as probability density distributions (black curves).

One has to find a finite set of priors that is a good approximation to the theoretical infinite set of differently shaped prior functions. Although we yet do not have a general rule to find such a good and small prior set, the challenge seems to be manageable in practice. Often, critical prior shapes that influence the result considerably are obvious. We want to note that in principle the final result of robust analysis is only a set of *n* highest posterior density ranges for *n* samples and parameters. Each interval limit is finally determined by only a single “extreme,” but acceptable prior. If we ignore split hpd ranges, *2n* such extreme priors account for the complete result. Our task is thus not to calculate as many different, but redundant priors as possible, but to find these extreme priors.

Considering the main structure of the sequence, we put priors on the lengths of the phases. Similar to the Iceman example above, exponential decreasing and increasing probability densities were used. It turned out to be sufficient to use only 2 different slopes of decreasing and 2 different slopes of increasing functions, and additionally the constant function for each phase length. A complete n -dimensional prior for the whole sequence is an arbitrary combination of these individual single-phase priors. So, for 5 different slopes and 4 phases, one would get $5^4 = 625$ different total prior functions. Fortunately, in practice it turns out that this number can be decreased considerably, because phases that are far from each other seem to behave fairly independently; thus, there is no need to try all respective combinations of their priors. This is at least the case for the 2 outer phases in this example. Additionally, most of the effect of robust analysis apparently can be achieved by focusing on extreme combinations, as combining a prior with decreasing probability for a phase length with one with increasing density of the neighboring phase. With these considerations, the prior set was reduced to 45 differently shaped prior functions, making the calculation time manageable, although the computer program developed by us is based on the mathematical programming language MATLAB and not optimized with regard to runtime, and therefore considerably slower than technically possible. Additionally, the commonly used uniform span prior was also included in the prior set.

As mentioned earlier, the priors are sieved for their agreement with the measured ^{14}C ages. In this example, 26 of the initial 46 priors were excluded by the agreement criterion. It is no disadvantage that a considerable number of prior shapes is discarded. If no prior would fail the agreement criterion, this could be an indication that the prior set does not cover the full range of possible prior shapes, and the resulting hpd ranges could be too small. The resulting hpd ranges at 95.4% confidence level for robust Bayesian analysis are given in Figure 6 along with hpd ranges for the common uniform span prior.

When comparing the results of robust analysis with that of the uniform span prior alone, one can see that there are no considerable systematic shifts; the positions of the hpd ranges are not changed in principle. This indicates that the choice of the uniform span prior does not produce a misleading result in general for this example. However, it is also obvious that most of the intervals become noticeable and sometimes considerably wider. For some contexts (e.g. “Trianda late LMIA twig” or “Thera VDL”), minor peaks within the unmodeled calibration, which are excluded with the uniform span prior, are included within the 95.4% interval by robust analysis. The hpd ranges of phase boundaries are generally changed more than that of contexts, because they are not so directly related to measured ^{14}C ages (e.g. “Start Mature LMIA” or “LMIB/LMII transition”).

Even though the structure of the sequence is not too simple, the change of some hpd ranges by robust analysis can be understood qualitatively. For example, the Thera VDL has to be younger than the contexts in the preceding phase “Pre-VDL LMIA.” Especially the context 65/N001/I2 within this phase is very young and so tends to shift the Thera VDL date to younger ages. On the other hand, the subsequent phase “LMIB Destructions Crete,” which starts (with 95.4% confidence) around 1650 BC at the earliest when using the uniform span prior, forces the Thera VDL in the opposite direction towards older dates, which tends to compensate the shift from 65/N001/I2. Robust analysis tries to consider all possible priors, and so also the case where the probability increases highly for very short lengths of the “LMIB Destructions Crete” phase. Thus, a very short “LMIB Destructions Crete” phase has to start and end somewhere around 1530 BC where all unmodeled calibrations of the contexts within the phase overlap. But if the phase starts that late, it does not tend to shift the Thera VDL age towards older ages any more. Only the shift to younger ages remains and enhances the small tail—in detail, 2 small peaks—at the young side of the unmod-

eled Thera VDL probability distribution, so that they are included in the resulting 95.4% hpd range. Thus, when focusing only on the 95.4% (2σ) confidence level, the result for the Thera VDL from robust analysis differs remarkably from this using the uniform span prior. However, at the 68.3% (1σ) confidence level (we calculate the intervals at all levels of confidence simultaneously), the minor peaks are excluded again by both methods.

Nevertheless, to avoid premature interpretations of this or any other detailed result of the sequence, one should be aware of the following fact: When evaluating the sequence with robust analysis, we use in principle all possible prior functions that are in agreement with the constraints as given in simplified form on the left side of Figure 6, without asking whether archaeological reasons could be found further that exclude a part of the particularly used prior functions within the used set. If we actually could exclude priors by further archaeological arguments (e.g. expectation values on some phase length or maxima for some time spans, etc.), this would in general reduce the hpd ranges. Therefore, robust Bayesian analysis forces the user to collect and actually use all available archaeological information, to avoid getting wider hpd ranges than necessary.

DISCUSSION

Surely, we do not see the procedures introduced by this work as a completed method; we rather want to introduce ideas that are hopefully discussed within the community. Our view of what we consider to be the most critical aspects is summarized in the following.

What is the Goal of Our “Robust Analysis,” and What is Not?

The idea is to avoid (or at least to reduce) the subjectivity of the chosen prior function for a given archaeological setting. This means we assume that the archaeologists came to a final description of the archaeological facts, which allows one to define a complete set of possible different prior functions that are all in agreement with this constraints. Our results are subsequently deduced from this set as described in detail in this work. Of course, this is an idealization of the real scenario, where the archaeological facts can hardly be fixed in an unchangeable form. Archaeologists will rather discuss the stratigraphy and find different possible interpretations and perspectives. They are able to assess different assumptions but cannot assign quantitative probability values. To handle this complexity remains with the archaeologists. The goal of our method is not to deal with this part of the problem but to avoid the need of selecting a particular prior function, as there remain an infinite number of different possible functions, even when the archaeological constraints are fixed.

However, the calculation of Bayes' factors to compare models in a quantitative way based on their agreement with the measured data could also support the selection of the archaeological model (the fixation of the archaeological constraints). Although, in our opinion, the archaeological constraints should better be defined based on archaeological information only, without considering the measurements.

What is the Rigorous Meaning of the Resulting Unified Highest Posterior Density Ranges?

Usually in Bayesian statistics, the result is expressed as a posterior probability density, which is not possible any more within the method introduced here. Instead, we use the unified highest posterior density ranges (unified hpd ranges). Since these are not commonly used, it is important to elaborate their rigorous meaning very carefully: For given constraints based on the archaeological evidence, there is an infinite number of consistent realizations of the prior function, each resulting in different marginal posterior distributions. If we assume that the formulated archaeological constraints are correct, which means that there is no statement included that is not historically true, then the prior

set will definitely include a function that correctly represents the actual probability to find samples of particular ages. This prior function would be the ideal choice to get an unbiased posterior density. To unify the hpd ranges of the posterior marginals for all priors guarantees that the hpd ranges for the ideal prior are included. As the real sample ages lie within the hpd ranges for the ideal prior with the corresponding probability (e.g. 95%), they lie within the unified ranges with at least the corresponding probability. Thus, the meaning of the unified ranges is this: One can be sure that the real sample ages lie within those intervals with at least the specified probability. The conditions therefore are just that the archaeological constraints and the used accuracies of the measurements are correct. A similar statement is not possible if one uses a single particular prior, because one does not know how strong the result is biased by the prior. (The deviation from this ideal concept caused by the need of discarding “corrupt” priors does not change the considerations here in principle.)

The kind of robust analysis used by us is based on Berger (1994; sections 1.3 and 4.1): the minimum and the maximum of a “posterior quantity of interest” is calculated for a prior set to get a robust conclusion. When dealing with limits of hpd ranges as in our case, we think that the union of all hpd ranges is the most straightforward implementation of this method.

Finally, it has to be noted that the use of hpd ranges instead of probability densities has a small disadvantage: If the marginal posterior density is very flat, the excluded parts can have densities nearly as high as the included ones. However, this problem is not specific to robust analysis; it is present whenever hpd ranges are used. Regardless of this, hpd ranges are commonly accepted to characterize the results of Bayesian analysis.

Is this Emphasis on “Safe” Results Really Necessary?

It is disputable whether it is really necessary that the results have to be “correct” in a rigorous sense as provided by robust analysis. Or are results even meaningful although they may be biased by the choice of a particular prior function? Naturally, it is reasonable to look at results calculated with the best assumption for the prior function, but one has always to be aware that these results are only accurate if the prior function reflects the unbiased functional representation of the constraints given by the archaeological facts. The meaning of “an unbiased functional representation” can be illustrated by the Iceman example from above. Imagine one would know the real ages (times of death) of a large number of Early Bronze Age men and the ages of the wood of the axes carried by them, too. One could then plot a point for each pair of ages within a 2-dimensional coordinate system and fit the points with a corresponding probability density function. This density would represent an unbiased prior function, because now the prediction on the probability of a particular age difference of man and axe is the “right” one. In reality this density is not known; therefore, the use of a particular prior function usually biases the result. So deciding to use a particular single prior or robust analysis instead is the decision between a high precision but also a higher risk of error and lower precision but a lower risk of error. It should be noted that the perception expressed by the example is not valid in general, because the actual probability for an event does not have to be based on a distribution of events (imaginary or not) in every case.

The Loss of a Particular Prior or “Model”

It can be seen as drawback of the used method that one only focuses on the union of all priors and does not analyze the influence of the individual priors—which are usually seen as different models—on the result. First of all, it should be remarked that the calculation performed by us provides all individual marginal posteriors to any prior used within the prior set. So it is possible to analyze the influence of individual priors on the result manually. Of course, if there are hundreds of different prior functions used—which can be reasonable—this will become cumbersome. However, our per-

ception of this topic is the following: We think what should be actually seen as a “model” for this considerations is not the individual prior function, but the available archaeological constraints or information. And, as pointed out earlier in this paper, this information can usually not be expressed by a single function, but by a set of prior functions. The entirety of the prior functions characterizes the model; therefore, the result for this entirety has a more fundamental meaning than the results for the individual prior functions.

The Subjectivity of the Finite Prior Set

In fact, it is a disadvantage of robust Bayesian analysis, as performed by us, that the theoretical infinite set of possible prior functions has to be approximated by a finite set. This could lead to the impression that the choice of the actual used priors brings in the subjectivity again, which the method tries to avoid. Fortunately, the problem turns out to be solvable in practice, because one is allowed to include every prior one can think about within the set, if just consistent with the archaeological constraints. Neither redundant priors nor meaningless priors will (in principle) degrade the result. Naturally, one can never fully exclude the possibility that priors or classes of priors are missed that would influence the result, but even then the risk of getting an incorrect result has been reduced.

The Arbitrary Procedure for Prior Standardization

There could be legitimate criticism that we use an arbitrary method to standardize prior functions with originally infinite integrals, to be able to calculate Bayes’ factors. The problem is not easy to solve, because it is a fundamental one. It arises from the use of unrestricted priors that have no defined probability density values (their density would be zero anywhere, if standardized on the infinite domain). On the other hand, it would not make sense to discard unrestricted priors generally, because many archaeological constraints are most directly realized by these kind of priors. Thus, to do a quantitative model comparison, one has to restrict these priors in an appropriate way. This restriction will always have to be based on the available calibrated ^{14}C ages or on the likelihood function. The particular way to realize this restriction is arbitrary in fact. However, it does not make a difference in principle if one defines a sharp domain around the likelihood or uses a “graduated domain” derived from the sum of all single-sample calibrations, as we do in this work.

Although it is not possible to avoid this arbitrariness, it is important to make clear that it only influences the process of discarding “corrupt” priors (that are in serious disagreement with the data). Calculation of the posterior densities for the different priors used can be performed without standardization of the prior function. This is the benefit of our method where only the hpd ranges are unified rather than the densities are summarized, which makes the result independent of a weighting of the individual priors.

The Problem with “Model Averaging”

We presented a detailed discussion in our work on defining a “parametric” prior set and handle the parameters as additional dimensions within Bayesian calculation, which is a possible technical realization of an approach that is known as model averaging. Of course, this method has the great advantage that it can be performed fully embedded in the Bayesian framework, resulting in marginal posterior densities for the original parameters (sample ages, boundaries, etc.) and for the prior parameters, too. When starting our work, we assumed that we will realize our goal with this method, which is mathematically well-defined and clear. It was not easy to accept a concept with discrete prior functions that have to approximate an infinite set of functions. However, “model averaging” also has a serious drawback: The different functions have to be weighted now by a prior probability

density for the parameters. In our opinion, it is very difficult to assign well-grounded probabilities for differently shaped prior functions. This problem introduces serious arbitrariness again and it is not possible to solve it pragmatically by postulating that all priors should have the same probability, because the chosen kind of parameterization (scaling) is an intrinsically arbitrary weighting anyhow. Aside from this complication, the method delivers an effective (or “average”) prior function, which again can be “right” or “wrong” in the sense that it can approximate well the unbiased functional representation of the constraints given by the archaeological facts or it does not.

However, we do not want to deny that model averaging is a meaningful method. The method assists powerfully the process of finding an appropriate average prior function and the posterior density of the prior parameters shows furthermore to which extent the individual prior shapes contribute to the average prior. So, in our understanding, model averaging is a kind of robust analysis that is based on the principle of finding a non-informative prior function; an approach that is different in principle from that we favor in this work, but well established in other fields where Bayesian mathematics are applied, and well worth further investigation.

CONCLUSION

It seems obvious that there are applications where usual Bayesian sequencing suppresses possible ranges for the real sample ages. These results from the use of a single, particular prior function, which is commonly not fully defined by the available archaeological information, and therefore is actually only one of various (infinitely) different possible functional realizations of the prior information. Robust Bayesian analysis, which in principle considers all possible differently shaped prior functions, could be a way to avoid these artifacts, thus leading to a higher reliability, even though it is paid for with a reduced precision.

However, there are difficulties that have to be investigated further. Not completely solved is the problem that the use of “corrupt” priors can destroy the result. Discarding the priors with respect to their agreement with the ^{14}C measurements by means of a threshold level is possible, but not finally satisfying, because this technique breaks the absolute objectivity of robust analysis again. A further difficulty is the fact that the calculations have to be done for each prior individually, so that only a finite number of prior functions can be considered. Fortunately, it seems to be not difficult in practice to find a finite prior set that is a sufficient approximation for the theoretical infinite set of prior functions.

In conclusion, we think that there are no unsolvable obstacles to establish robust Bayesian analysis as a safe sequencing method for ^{14}C dates that are related through archaeological evidence.

ACKNOWLEDGMENTS

This work is supported by “SCIEM2000—Synchronization of Civilizations in the Eastern Mediterranean in the 2nd Millennium BC,” a Special Research Program (SFB) of the Austrian Academy at the Austrian Science Fund and under Patronage of the UNESCO. We thank 2 anonymous reviewers for critical and helpful comments.

REFERENCES

- Berger JO. 1994. An overview of robust Bayesian analysis. *Test* 3(1):5–124.
- Berger JO. 2006. The case for objective Bayesian analysis. *Bayesian Analysis* 1(3):385–402.
- Bonani G, Ivy S, Hajdas I, Niklaus TR, Suter M. 1994. AMS ^{14}C age determinations of tissue, bone and grass samples from the Ötztal Ice Man. *Radiocarbon* 36(2): 247–50.

- Bronk Ramsey C. 1995. Radiocarbon calibration and analysis of stratigraphy: the OxCal program. *Radiocarbon* 37(2):425–30.
- Bronk Ramsey C. 2000. Comment on ‘The use of Bayesian statistics for ^{14}C dates of chronologically ordered samples: a critical analysis.’ *Radiocarbon* 42(2):199–202.
- Bronk Ramsey C. 2001. Development of the radiocarbon calibration program. *Radiocarbon* 43(2A):355–63.
- Bronk Ramsey C. 2009. Bayesian analysis of radiocarbon dates. *Radiocarbon* 51(1):337–60.
- Buck CE, Kenworthy JB, Litton CD, Smith AFM. 1991. Combining archaeological and radiocarbon information: a Bayesian approach to calibration. *Antiquity* 65(249):808–21.
- Buck CE, Litton CD, Smith AFM. 1992. Calibration of radiocarbon results pertaining to related archaeological results. *Journal of Archaeological Science* 19(5):497–512.
- Buck CE, Cavanagh WG, Litton CD. 1996. *Bayesian Approach to Interpreting Archaeological Data*. Chichester: John Wiley & Sons Ltd.
- Garcia-Donato G, Chen M-H. 2005. Calibrating Bayes factor under prior predictive distributions. *Statistica Sinica* 15:359–80.
- Gilks WR, Richardson S, Spiegelhalter DJ, editors. 1996. *Markov Chain Monte Carlo in Practice*. London: Chapman & Hall.
- Hedges REM, Housley RA, Bronk CR, van Klinken GJ. 1992. Radiocarbon dates from the Oxford AMS system: Archaeometry Datelist 15. *Archaeometry* 34(2):337–57.
- Hoeting JA, Madigan D, Raftery AE, Volinsky CT. 1999. Bayesian model averaging: a tutorial. *Statistical Science* 14(4):382–401.
- Jeffreys H. 1961. *Theory of Probability*. London: Oxford University Press.
- Jones MD, Nicholls GK. 2002. New radiocarbon calibration software. *Radiocarbon* 44(3):663–74.
- Krause A. 1994. *Computerintensive statistische Methoden: Gibbs sampling in Regressionsmodellen*. Stuttgart: Gustav Fischer Verlag. In German.
- Kutschera W, Müller W. 2003. “Isotope language” of the Alpine Iceman investigated with AMS and MS. *Nuclear Instruments and Methods in Physics Research B* 204:705–19.
- Kutschera W, Rom W. 2000. Ötzi, the prehistoric Iceman. *Nuclear Instruments and Methods in Physics Research B* 164–165:12–22.
- Manning SW, Bronk Ramsey C, Kutschera W, Higham T, Kromer B, Steier P, Wild EM. 2006. Chronology for the Aegean Late Bronze Age 1700–1400 B.C. *Science* 312(5773):565–9.
- Reimer PJ, Baillie MGL, Bard E, Bayliss A, Beck JW, Bertrand CJH, Blackwell PG, Buck CE, Burr GS, Cutler KB, Damon PE, Edwards RL, Fairbanks RG, Friedrich M, Guilderson TP, Hogg AG, Hughen KA, Kromer B, McCormac G, Manning S, Bronk Ramsey C, Reimer RW, Remmele S, Southon JR, Stuiver M, Talamo S, Taylor FW, van der Plicht J, Weyhenmeyer CE. 2004. IntCal04 terrestrial radiocarbon age calibration, 0–26 cal kyr BP. *Radiocarbon* 46(3):1029–58.
- Rios Insua D, Ruggeri F. 2000. *Robust Bayesian Analysis*. New York: Springer.
- Rom W, Golser R, Kutschera W, Priller A, Steier P, Wild EM. 1999. AMS ^{14}C dating of equipment from the Iceman and of spruce logs from the prehistoric salt mines of Hallstatt. *Radiocarbon* 41(2):183–97.
- Sivia DS. 1996. *Data Analysis—A Bayesian Tutorial*. New York: Oxford University Press.
- Steier P, Rom W. 2000. The use of Bayesian statistics for ^{14}C dates of chronologically ordered samples: a critical analysis. *Radiocarbon* 42(2):183–98.
- Steier P, Rom W, Puchegger S. 2001. New methods and critical aspects in Bayesian mathematics for ^{14}C calibration. *Radiocarbon* 43(2A):373–80.
- Weninger F, Steier P, Kutschera W, Wild EM. 2006. The Principle of the Bayesian Method. *Egypt and the Levant* 16:317–24.