Corbelled Domes in Two and Three Dimensions: The Treasury of Atreus

Nicole A. Lazar ^{*} Joseph B. Kadane [†] Fang Chen [‡]

William G. Cavanagh [§] Clifford D. Litton [¶]

Address of Corresponding Author:

Nicole A. Lazar
Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213
Email: nlazar@stat.cmu.edu
Phone: (412) 268-7830
Fax: (412) 268-7828

^{*}Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.

[†]Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.

[‡]Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.

[§]Department of Archaeology, University of Nottingham, University Park, Nottingham, NG7 2RD,

UK

[¶]School of Mathematical Sciences, University of Nottingham, University Park, Nottingham, NG7

2RD, UK

Summary

Before the development of the true dome, many ancient cultures used the technique of corbelling to roof spaces. Recently, a series of related statistical models have been proposed in the literature for explaining how corbelled domes might have been constructed. The most sophisticated of these models is based on a piecewise linear structure, with an unknown number of changepoints, to guide the building process. This model is analyzed by the reversible jump Markov Chain Monte Carlo (MCMC) technique. All models considered to date have been two-dimensional, that is, they have taken a single cross section through the dome; even when more extensive data, in the form of measurements on multiple slices through the dome, have been available, these have been averaged together for the purposes of analysis. In this paper, we extend the two-dimensional analysis to a three-dimensional analysis, that takes full advantage of the data collected by the archaeologists and of the rotational symmetries inherent in the structure. We also explore ways of graphically presenting the results from a complex, reversible jump MCMC implementation, in order to check convergence, good mixing, and appropriate exploration of the (high dimensional and varying dimension) parameter space. The model and the graphical techniques are demonstrated on the Treasury of Atreus in Mycenae, Greece, one of the finest extant examples

of the corbelling method.

*Key words: Archaeology, Bayesian Data Analysis, Graphical Representation, Multiple
Changepoints, Reversible Jump MCMC, Statistical Modeling*

1 Introduction

Corbelling is an ancient technique, used before the development of the true dome, for spanning or roofing spaces. Examples of this construction method have been found in a number of prehistoric societies, ranging from tombs in Iberia and the British Isles, to buildings in Italy, to passages and chambers of the Great Pyramid of Kheops in Giza, Egypt. Most of these utilize corbelling to span relatively small distances, 2 to 3 meters at most, with a conservative slope.

One of the most impressive uses of the technique, which involves laying stones on top of each other, with successive layers being slightly offset (think of using rectangular Lego pieces to build an arch; see Figure 1), comes from Mycenae in Greece, where archaeologists have found corbelled structures spanning 8 to 14 meters. As noted by Cavanagh and Laxton (1981), only the invention of the true dome enabled builders to bridge larger distances without internal support.

Due to the prevalence of corbelling in the ancient world, and the fact that many of the structures which were built with this spanning method are still standing today, thousands of years later, archaeologists have become interested in studying this technique more closely. In particular, interest has focused on finding models that adequately describe the shape of corbelled domes, and perhaps give insight into how these structures were built in practice. Much has been written about this problem in the last twenty years, by both archaeologists and statisticians. Especially interesting from the statistical point of view is that the data lend themselves to a variety of surprisingly advanced analysis

Figure 1: Shape of the inner face of an idealized corbelled structure. Reprinted from Buck, Litton and Stephens (1993) with permission from the Royal Statistical Society.

techniques, such as multiple changepoint detection and reversible jump Markov Chain Monte Carlo. In this paper, we discuss in detail the archaeological, mathematical and statistical developments that surround the analysis of corbelled domes. We also extend the current existing models to take advantage of the three dimensional nature of the structures, which to date has been ignored. The outline of the rest of the paper is as follows. Section 2 traces the development of the model, from a deterministic least squares line fit, to a multiple changepoint model cast in the Bayesian setting and analyzed using the most advanced MCMC techniques. These models are based on measurements from multiple cross-sections of the dome, that are averaged together to get a single, two-dimensional “slice.” Section 3 introduces an extension of the existing models to include the third dimension of the structure. In other words, the measurements along the various cross-sections are no longer averaged together. Rotational symmetries are imposed on the individual slices, however, to model accurately the constraints that would have been placed on the builders. This three dimensional analysis is demonstrated, in some detail, on the Treasury of Atreus. Section 4 gives a discussion of the results and general conclusions.

2 Development of a Statistical Model

2.1 Classical Analysis

In the first mathematical study, Cavanagh and Laxton (1981) considered the structural mechanics of the Mycenaean *tholos* tombs. By examining the structure of these tombs and the factors necessary for the vault to remain standing, they came up with the following, deterministic, model:

$$\text{radius} = \alpha \times \text{depth}^\beta,$$

where *depth* is the depth below the top of the tomb, *radius* is the radius of the tomb at the measured depth, and α and β are constants to be determined for each tomb individually. Taking logarithms of each side of the equation, a simple linear equation results, and it is possible to find values for the two coefficients using least squares. Cavanagh and Laxton examined five tombs that varied in location, age and building style, and found that all five had a value of β of around $2/3$. Values of α varied, as this parameter depends on the size (specifically, the height) of each tomb. One of the tombs, the Treasury of Atreus, is almost twice the size of the other tombs in this study, and it had $\alpha = 2.7$, compared to $\alpha \approx 2$ for the other four. Their model, while providing a starting point for the analysis of the Mycenaean tombs, was lacking in two important aspects, addressed in Cavanagh, Laxton and Litton (1985). First, many of the Mycenaean *tholoi* were built partly below ground. As described by Cavanagh, Laxton and Litton (1985), the floor of such tombs was made by cutting a cylindrical cavity out of the bedrock, or

from the edge of a hillside. Thus, the Mycenaean tholoi were often either sunk into a hillside or consisted of a subterranean burial chamber with a corbelled roof; by way of contrast, the earlier Minoan circular tombs on the island of Crete were built above ground (see also http://projects.dartmouth.edu/history/bronze_age/ for historical details and comparisons). The walls in the lower portion of the tomb may have been constructed according to a different plan than the upper part. For instance, they may have been straighter, because of the constraints of the rock-cut “cylinder” or in order to accommodate the doorway. Instead of one value of α and one value of β sufficing to describe the entire tomb, there might be a point at which the building regime, and the parameters guiding construction, change. The changepoint, if it exists, would correspond to where the dome and the wall meet. The first simple model considered by Cavanagh and Laxton was based only on measurements from the dome, and not the entire tomb. Thus, whereas their simple form was the basic one required for stability of the dome, it did not necessarily describe adequately the Mycenaean tombs of mainland Greece. Second, it might not be possible to measure from the apex of the dome, that is, at depth=0. There are two reasons for this: the top of the tomb might have collapsed, or the corbelling wasn’t originally continued to the top of the arch, instead a large slab was placed down on the stones, to close off the dome. Under these two modifications, and introducing a stochastic element to account for errors, the proposed model became

$$\log(r_i) = \begin{cases} \log(\alpha_1) + \beta_1 \log(d_i + \delta) + \epsilon_i & d_i \leq \gamma \\ \log(\alpha_2) + \beta_2 \log(d_i + \delta) + \epsilon_i & d_i \geq \gamma \end{cases}$$

Here, r_i is the radius measured at depth d_i , γ is the location of the changepoint, and the ϵ_i are assumed to be normally distributed, with mean 0 and variance σ^2 . Physical considerations require that $\beta_1 > \beta_2 > 0$. At depth γ , the two lines are required to intersect, that is, $\alpha_1(\gamma+\delta)^{\beta_1} = \alpha_2(\gamma+\delta)^{\beta_2}$, so that there are only six free parameters in the problem, instead of seven. In particular, β_2 can be expressed as $\beta_1 + [\log(\alpha_1/\alpha_2)/\log(\gamma + \delta)]$.

The authors used maximum likelihood techniques developed by Hinkley (1969, 1971) for the analysis of changepoint problems, to compare the parameters of Greek tholoi and Sardinian *nuraghi* (Bronze Age stone towers which served as defensive structures). Based on the parameter estimates, it was possible to distinguish between the tholoi and the nuraghi. For example, for the Sardinian structures, no changepoints were apparent. It should be noted that the nuraghi, unlike the tholoi, are completely above ground, so one of the reasons for proposing a changepoint model doesn't hold, and this seemed to be confirmed by the data. For the tholoi under consideration in this study, δ was taken to be zero (that is, known and fixed), whereas for the nuraghi, it was allowed to vary, as an unknown parameter of the model. The method was also able to pick out one of the Mycenaean tombs, that at Karditsa, as different from the others in its group, since it appeared to have no changepoint. As pointed out by the authors, this tomb was unusual in two ways. One of the distinguishing features of the Karditsa tholos is that it is set entirely within a knoll. Cavanagh *et al.* speculated that the builders of the tomb were able to take advantage of this fact and adapted the construction regime accordingly. In

addition, the tomb at Karditsa had been a source of some dispute among archaeologists, as to whether it belonged to the Bronze or Iron Age. The authors suggested that the statistical analysis could provide another, independent, criterion for dating this structure.

2.2 The Bayesian Perspective

An important step forward was taken by Buck, Litton and Stephens (1993), who used methods developed in Stephens (1994), for the analysis of changepoint problems from a Bayesian point of view. The authors applied MCMC methods to a Minoan tholos tomb at Stylos, Crete. Their base model was

$$\log(r_i) = \begin{cases} \log(\alpha_1) + \beta_1 \log(d_i + \delta) + \epsilon_i & d_i \leq \gamma \\ \log(\alpha_2) + \beta_2 \log(d_i + \delta) + \epsilon_i & d_i \geq \gamma \end{cases}$$

as in previous work. The conditions on the model were also as described above. With the assumption of normal errors, the likelihood is easily specified, although it should be noted that because of the parameter constraints and the piecewise nature of the model, sampling from the posterior is not simple, regardless of the prior specification. Buck and colleagues chose non-informative (vague) priors for $\log \alpha_1$, $\log \alpha_2$, β_1 , σ^2 and γ ; the prior for δ was uniform between 0 and 0.34, to account for the thickness of the capping stone on this tomb, which was 0.34 meters.

Under this specification, the Markov Chain Monte Carlo sampling scheme can be carried out using Gibbs sampling, which draws successive observations from the full conditional distributions, that is, the posterior distribution of each parameter given all of

the others (Casella and George, 1992). The three parameters $\log \alpha_1$, $\log \alpha_2$ and β_1 have conditional posterior distributions in the normal family, while σ^2 is conditionally inverse gamma. These are easily obtained. On the other hand, the full conditional densities for δ and γ are both of non-standard form, hence are sampled from using *rejection sampling* (see Buck, Litton and Stephens, 1993 and Stephens, 1994 for details).

A number of advantages of the Bayesian approach are apparent for this problem. Instead of just point estimates of the parameters of interest, entire posterior distributions are obtained. This gives archaeologists more flexibility in studying and interpreting the results of the analysis. The classical analysis did not incorporate any uncertainty about δ into the inference process, whereas this was an outcome of the Bayesian analysis by its very nature. Furthermore, information about the main parameter of interest, the location of the changepoint, was refined. The 95% confidence interval reported for γ arising from the classical analysis (1.64–1.84), was apparently too short; the 95% highest posterior density interval from the Bayesian analysis was 1.48–2.44. As pointed out by Buck *et al.*, this latter interval just includes the depth at which the lintel is below the capping stone of the tomb, 2.44 meters. This finding supported a supposition by archaeologists that the placing of the lintel would cause a break in the process of corbelling, with different structural forces coming in to play above and below that point. Using the classical approach, there was no support for this intuition, as the changepoint was placed too high.

This argument, however, may put too much weight on exactly where the top end

of a 95% credible interval falls. The number 0.95 has nothing intrinsic to do with the problem. Its use is merely traditional; it is “treat[ed] . . . with the same superstitious awe that is usually reserved for the number 13” (Raiffa and Schlaifer, 2000, pg. vii).

Regardless of whether a single changepoint model points to a change at the level of the lintel, these considerations lead to the possibility of models with two or more changepoints, perhaps one at the lintel and one at the ground level at the time of construction. That is, are there features of a corbelled structure that might impose an adjustment in the parameters guiding construction? One could consider expanding the model described above to account for more than one changepoint. However, the appropriate number of changepoints is not known, and so it would be desirable to include this as a parameter in a new model. While seemingly a minor modification of the existing framework, in fact letting the number of building segments vary introduces nontrivial complications. The problem is that, as the number of changepoints varies across models, interpretation of the model parameters also changes. In addition, the practical issue of how to estimate the true dimension of the model, by sampling over the space of models with differing numbers of changepoints, is not an easy one to resolve.

Fortunately, at around the time we became interested in this question, a solution was proposed in the Bayesian literature, in the form of *Reversible Jump MCMC* (Green, 1995; Richardson and Green, 1997). The reversible jump technique designs a Markov chain that jumps between models of different dimensions, while still preserving the conditions needed to eventually reach a stationary distribution. Implementation of the method can be

complicated and the details tend to be unique to each application. Denison, Mallick and Smith (1998) derived a reversible jump algorithm for dealing with multiple changepoints, however they made certain simplifying assumptions which are not appropriate for the problem we are working on.

More recently, Fan and Brooks (2000) developed methodology specifically for the analysis of corbelled structures. They extended the basic model to allow for two generalizations. First, δ was not restricted to be the same in each linear piece. Second, at lower depths, the wall of the tomb might be nearly vertical, that is, governed by a single parameter. An example of a candidate two changepoint model is

$$\log(r_i) = \begin{cases} \log(\alpha_1) + \beta_1 \log(d_i + \delta_1) + \epsilon_i & d_i \leq \gamma_1 \\ \log(\alpha_2) + \beta_2 \log(d_i + \delta_2) + \epsilon_i & \gamma_1 \leq d_i \leq \gamma_2 \\ \log(\alpha_3) + \epsilon_i & d_i \geq \gamma_2 \end{cases}$$

The number of segments was allowed to increase to a maximum of four, by permitting as many as three changepoints. The reversible jump algorithm implemented by Fan and Brooks also explored the possibility that δ was in fact the same for all segments, and the possibility that there was no constant segment at the bottom of the tomb. In addition to reanalyzing the data from the tomb at Stylos in Crete, Fan and Brooks compared the results of their reversible jump MCMC algorithm for another Minoan tholos to those for a Mycenaean tholos and a Sardinian nuraghe. They found that for each, the method assigned slightly different orderings of the posterior probabilities, that is, indicated different underlying models as having the highest posterior probabilities.

The estimated values of the β parameters in particular were similar enough in all three of the buildings considered in this comparison, that the authors concluded that this information alone was not enough to distinguish among structures of differing origin. It is worth pointing out, however, that interpretation of the β parameters in the comparison is complicated by the fact that different models were fit to each of the structures – a two changepoint model with a constant segment, a one changepoint model with no constant segment, and a two changepoint model with no constant segment for the Minoan tholos, the Mycenaean tholos and the nuraghe, respectively. Since the models are different, the interpretation of β in each case is also different, making an overall conclusion such as that drawn by Fan and Brooks (2000) tentative at best.

3 Three Dimensional Data and the Treasury of Atreus

3.1 Three Dimensional Data

Corbelled domes are three dimensional structures, and many, like the tholos tombs, were built rotationally symmetric, or nearly so, around a center line (though later settlement, whilst still leaving the structure stable, could distort slightly the original symmetry). For this reason, the analyses reviewed above are in fact based on averages of measurements taken on several slices starting at the apex. The model discussed below is the first, to our knowledge, to consider each measurement slice separately. Since a lintel and doorway intercept one of the slices, simple averaging, as was performed previously, is neither

appropriate nor sufficient. The slice that is intersected by the lintel and the doorway has a different number of measurements than the other slices, and some of the measurements were taken at different places, as explained below. Note that under the assumption that all measurement slices contain complete data taken at the same depths, a simple sufficiency argument shows that no information is lost by averaging. However when one or more slices has missing data or has measurements taken at different locations, averaging makes little sense and results in the loss of information. From the practical point of view, furthermore, the model we develop here, which can easily handle missing data, gives the archaeologist more flexibility, since difficulties in collecting complete data along all cross-sections is not an impediment to statistical analysis.

3.2 Background and Data

The Treasury of Atreus, also known as the Tomb of Agamemnon, is a tholos tomb in Mycenae, Greece. It is considered to be the finest tomb to have survived and dates to about 1300 BCE. It is far from typical of the majority of tholos tombs because of the unusual masonry used in its construction. An interesting feature of the Treasury is that the lintel is very long, large and heavy, with a relieving triangle above. The role of this relieving triangle is to concentrate the weight of the masonry on to the jambs rather than on the lintel. This helps prevent the lintel from cracking and breaking.

A detailed architectural survey was carried out and accurate plans made. From the plans, it is possible to determine the radius of the tholos at a given depth. We have

data on three separate cross-sections of the Treasury of Atreus, as shown in Figure 2. Cross-section G-H cuts through the entrance of the tomb, and thus the relieving triangle as well. The center of the tomb is denoted O. For each cross-section, the two radii were measured at regular depths from the apex of the dome. Note that for slice OH, there are no measurements for the relieving triangle, nor are there any for below the lintel. Radii OC, OD, OE, OF, OG and OH have 46, 45, 46, 45, 45 and 21 measurements respectively. The data are given in Table 1.

3.3 The Statistical Model

We used the base model proposed by Fan and Brooks (2000), with the exception that for the Treasury of Atreus, the δ parameter is identically zero, since the apex of the structure is known (very unusually, the form of the last few centimetres of the dome were carved out of the capstone of the Treasury of Atreus, giving a very precise location of where the builders placed the apex of the vault). More specifically, the data are in the form of $(d_i, \mathbf{y}_i)$, where d_i , with $i = 1, \dots, n_j$, denotes the depth below the apex of the tomb, and $\mathbf{y}_i = \{y_{ij}, j = 1, \dots, 6\}$ are corresponding radius measures at that depth, for the five slices with complete data, and one with incomplete data. The radii y_{ij} are modelled as normal centered at r_i , with unknown variance τ_k^2 , and k indicates the segment in the piecewise linear model to which y_{ij} belongs. An individual observation therefore belongs to both a *slice*, that is, a particular cross-section of the tomb, and to a *segment* of the piecewise linear fit.

Figure 2: Overview of the Treasury of Atreus, showing the cross-sections along which the data were acquired.

Depth	OC	OD	OE	OF	OG	OH	Depth	OC	OD	OE	OF	OG	OH
0.30	0.36	0.36	0.40	0.32	0.30	0.50	5.40	4.22	4.24	4.16	4.36	4.30	–
0.60	0.68	0.66	0.70	0.62	0.68	0.72	5.70	4.40	4.40	4.32	4.56	4.46	–
0.90	0.88	0.92	0.88	0.90	0.86	0.96	6.00	4.54	4.58	4.52	4.74	4.64	–
1.20	1.10	1.16	1.12	1.16	1.10	1.16	6.30	4.76	4.76	4.74	4.94	4.86	–
1.50	1.32	1.36	1.42	1.42	1.36	1.40	6.60	4.96	5.00	4.92	5.14	5.04	–
1.80	1.54	1.62	1.60	1.68	1.58	1.62	6.90	5.12	5.18	5.10	5.36	5.22	–
2.10	1.78	1.84	1.84	1.90	1.82	1.86	6.92	–	–	–	–	–	5.16
2.40	2.06	2.09	2.06	2.12	2.04	2.12	7.20	5.32	5.36	5.26	5.52	5.38	5.24
2.70	2.24	2.30	2.30	2.40	2.34	2.38	7.50	5.48	5.50	5.38	5.64	5.58	5.34
3.00	2.56	2.58	2.52	2.68	2.56	2.58	7.80	5.66	5.66	5.50	5.78	5.72	5.48
3.30	2.80	2.84	2.76	2.94	2.84	2.84	8.10	5.78	5.80	5.66	5.90	5.88	5.64
3.60	3.08	3.12	3.04	3.28	3.12	3.08	8.40	5.93	5.92	5.78	6.04	5.98	–
3.90	3.32	3.34	3.26	3.46	3.38	3.32	8.70	6.06	6.04	5.90	6.18	6.08	–
4.20	3.48	3.56	3.48	3.64	3.54	3.50	9.00	6.18	6.14	6.02	6.30	6.20	–
4.50	3.65	3.74	3.64	3.82	3.72	3.74	9.30	6.28	6.26	6.14	6.42	6.34	–
4.70	–	–	–	–	–	3.90	9.60	6.40	6.36	6.26	6.52	6.42	–
4.80	3.82	3.92	3.80	4.02	3.90	–	9.90	6.56	6.48	6.38	6.64	6.56	–
5.10	4.04	4.08	3.98	4.18	4.08	–	10.20	6.64	6.54	6.50	6.76	6.66	–

Depth	OC	OD	OE	OF	OG	OH
10.50	6.72	6.66	6.60	6.84	6.80	–
10.80	6.84	6.78	6.72	6.94	6.90	–
11.10	6.94	6.90	6.82	7.06	7.00	–
11.40	7.04	7.02	6.92	7.16	7.08	–
11.70	7.12	7.12	7.00	7.24	7.18	–
12.00	7.20	7.22	7.08	7.32	7.24	–
12.30	7.26	7.24	7.06	7.36	7.30	–
12.60	7.30	7.26	7.08	7.40	7.32	–
12.90	7.31	7.26	7.08	7.44	7.34	–
13.20	7.32	7.26	7.16	7.46	7.38	–
13.50	7.34	7.24	7.20	7.48	7.38	–
13.60	7.34	–	7.20	–	–	–

Table 1: **Measurements, in meters, of the depth from the apex on six radii from the Treasury of Atreus**

As before, d_i and r_i enjoy the following basic relationship:

$$\log(r_i) = \log(\alpha) + \beta \log(d_i) + \varepsilon_i \quad (1)$$

Notice that we now need to make a distinction between the measured radii, y_{ij} , and r_i , the radius at depth d_i which acts also as the theoretical mean of the observations at depth d_i . Since we are in a Bayesian framework, these r_i are themselves taken to be random quantities, drawn from a distribution with parameters of its own. Previous analysis did not need this hierarchy, because only a single radius measurement was used at each depth.

Given that we believe that the tomb is formed segmentwise with changepoints γ , equation (1) can likewise be extended to allow for multiple changepoints, for instance, two changepoints with a constant segment at the bottom:

$$\log(r_i) = \begin{cases} \log(\alpha_1) + \beta_1 \log(d_i) + \varepsilon_i & 0 \leq d_i < \gamma_1 \\ \log(\alpha_2) + \beta_2 \log(d_i) + \varepsilon_i & \gamma_1 \leq d_i < \gamma_2 \\ \log(\alpha_3) + \varepsilon_i & \gamma_2 \leq d_i \end{cases} \quad (2)$$

with continuity constraints

$$\alpha_1 \gamma_1^{\beta_1} = \alpha_2 \gamma_1^{\beta_2}$$

$$\alpha_2 \gamma_2^{\beta_2} = \alpha_3$$

where α_3 denotes the constant term. $\varepsilon_i \sim N(0, \sigma^2)$ with σ^2 unknown.

For a model that has K segments ($K - 1$ change points), the likelihood is

$$L(\mathbf{Y}|\mathbf{r}, \boldsymbol{\tau}^2, \boldsymbol{\gamma}) = \prod_{i=1}^{m_1} \prod_{j=1}^6 N(y_{ij}|r_i, \tau_1^2) \cdot \prod_{i=m_1+1}^{m_2} \prod_{j=1}^6 N(y_{ij}|r_i, \tau_2^2) \cdots \prod_{i=m_K+1}^{i=n} \prod_{j=1}^6 N(y_{ij}|r_i, \tau_{K+1}^2),$$

where τ_k^2 is the variance of y_{ij} within the k th segment. Furthermore

$$m_1 = \sum_{i=1}^n \mathbf{I}\{d_i < \gamma_1\}, m_2 = \sum_{i=1}^n \mathbf{I}\{\gamma_1 \leq d_i < \gamma_2\}, \dots, m_K = \sum_{i=1}^n \mathbf{I}\{\gamma_{K-1} \leq d_i\} \quad (3)$$

In the presence of missing data, we denote $\mathbf{Y} = \{\mathbf{Y}_{\text{obs}}, \mathbf{Y}_{\text{mis}}\}$, for “observed” and “missing” data respectively. $\mathbf{Y}_{\text{obs}}$, the data, are modelled via the likelihood as Normal centered at the corresponding $\mathbf{r}$ with variance τ^2 . $\mathbf{Y}_{\text{mis}}$ are treated as parameters, and samples from its predictive distribution $p(y_{ij}|\mathbf{Y})$ can be easily obtained: $p(y_{ij}|\mathbf{Y}) = \int p(y_{ij}|\mathbf{r}, \tau^2) p(\mathbf{r}, \tau^2|\mathbf{Y})$.

From a sampling point of view, we first draw samples from the posterior distribution of $(\mathbf{r}^{(h)}, \tau^{2(h)}|\mathbf{Y})$, where $\mathbf{Y} = \{\mathbf{Y}_{\text{obs}}, \mathbf{Y}_{\text{mis}}^{(h-1)}\}$, h is the current step in the Gibbs setting. Next we use the drawn samples of $\mathbf{r}$ and τ^2 to produce $\mathbf{Y}_{\text{mis}}^{(h)}$, where $y_{ij} \sim N(r_i, \tau_k^2)$ with the corresponding i (location), j (slice), and k (segment).

The second level of the hierarchical model is defined in the following way:

$$\begin{aligned} L(\mathbf{r}|\mathbf{d}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\sigma}^2) &= \prod_{i=1}^{m_1} N(\log(r_i)|\alpha_1, \beta_1, \sigma^2) \cdot \prod_{i=m_1+1}^{i=m_2} N(\log(r_i)|\alpha_2, \beta_2, \sigma^2) \cdot \\ &\quad \dots \prod_{i=m_K+1}^{i=n} N(\log(r_i)|\alpha_K, \beta_K, \sigma^2) \\ \log(L(\mathbf{r}|\mathbf{d}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\sigma}^2)) &\propto -n \log(\sigma^2) + \sum_{i=1}^{m_1} -(\log(r_i) - \log(\alpha_1) - \beta_1 \log(d_i))^2 / 2\sigma^2 + \dots \\ &\quad + \sum_{i=m_K+1}^{i=n} -(\log(r_i) - \log(\alpha_k) - \beta_k \log(d_i))^2 / 2\sigma^2 \end{aligned}$$

We chose priors to be both conjugate (hence convenient from a computational point of view, and consistent with previous analysis, such as that of Fan and Brooks) and

reasonably diffuse.

$$\log(\alpha_k) \sim N(0, 3), \forall k = 1, \dots, K$$

$$\beta_k \sim N(1, 2), \forall k$$

$$\gamma \sim \mathcal{U}(d_1, d_{k-1}), \gamma_i > \gamma_j, i > j$$

$$1/\sigma^2 \sim \Gamma(1, 0.01)$$

$$\log(r_i) \sim N(0, 10), \forall i$$

$$1/\tau_k^2 \sim \Gamma(1, 0.01), \forall k = 1, \dots, K$$

Hence, the posterior is proportional to

$$\begin{aligned} & \prod_{k=1}^K N(\log \alpha_k | 0, 3) \cdot \prod_{k=1}^K N(\beta_k | 1, 2) \cdot [\pi(\gamma)] \cdot \Gamma\left(\frac{1}{\sigma^2} | 1, 0.01\right) \\ & \cdot L(\mathbf{r} | \mathbf{d}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\sigma}^2) \cdot \prod_{i=1}^n N(\log r_i | 0, 10) \cdot \prod_{k=1}^K \Gamma\left(\frac{1}{\tau_k^2} | 1, 0.01\right) \cdot L(\mathbf{Y} | \mathbf{r}, \boldsymbol{\tau}^2, \boldsymbol{\gamma}) \end{aligned}$$

3.4 Results

The data were analyzed using reversible jump MCMC (Green, 1995; Richardson and Green, 1997). The implementation of the reversible jump method followed Fan and Brooks (2000). By using the reversible jump approach, we avoid the need to explicitly calculate Bayes factors to choose among the models; the reversible jump algorithm provides information about the size of the model, since one of the parameters is the number of changepoints. We thus obtain the posterior probabilities of each model from the chain. As pointed out by DiCiccio *et al.* (1997), a variety of techniques have been proposed over the years for dealing with this question. The methods are based on asymptotic

approximations (for instance, the Laplace approximation, Tierney and Kadane, 1986), simulation (marginal likelihood, Chib (1995), Chib and Jeliazkov (2001); bridge sampling, Meng and Wong (1996), Gelman and Meng (1998)), and a combination of them both (various modifications on existing methods proposed by DiCiccio *et al.*, 1997).

We selected five very different starting values, as can be seen in Figure 3. For each model, we ran 15×10^6 iterations. For every 5000th iteration, we recorded which of the models the chain was visiting at that iteration and its log(posterior) value.

In order to assess the stability and convergence of the chain, we looked at the traces of the average log(posterior) for each of the five runs, the cumulative probability of spending time in what turned out to be a much-visited model for each of the five runs, and the joint plot of these two measures. These are Figures 4, 5 and 6, respectively. Inspection of these Figures reveals that from all starting points the chains converged to similar places.

Since the results from the five runs, starting at five different locations, all converged to more or less the same place, we then ran a much longer chain, of 60 million iterations, discarded the first 10 million as burn in, and kept the results of every 5000th sample. Table 2 shows the proportion of time spent, out of the resulting 10,000 iterations, in each of the models. As can be seen from inspection of the Table, by far the most time was spent in models with 1 or 2 changepoints. Models without a constant segment at the end were overwhelmingly preferred.

Only two models were visited with any frequency – the one changepoint model with no constant segment at the bottom, and the two changepoint model with no constant

Figure 3: The five starting points used in the reversible jump MCMC simulation. Note that most of these starting points represent very poor fitting models.

segment. All other models were rarely visited. Table 3 gives the posterior summaries for the two most visited models.

The reversible jump analysis points to the model with a single changepoint and no constant segment at the bottom, as being by far the preferred model. This model matches

Figure 4: Trace plots of the average of the $\log(\text{posterior})$ for the MCMC runs from each of the starting points. From all five starting values, this measure converges to approximately the same place.

Figure 5: Trace plots of the probability of being in a particular model for the MCMC runs from each of the starting points. From all five starting values, this measure converges to approximately the same place.

Figure 6: Joint plot of the average $\log(\text{posterior})$ and the model probability, for each of the five MCMC runs. The big “O” is the starting value in each run, the little “o” is the ending value. The ending values are centred around the point $(0.8, 133)$.

k	Constant segment	
	Yes	No
0	0	0
1	0	0.760
2	0.024	0.196
3	0.003	0.017
4	0	0

Table 2: **Reversible Jump MCMC: Proportion of visits spent in each model, out of 10,000 iterations. k is the number of changepoints.**

somewhat with a structural interpretation. As described above, the OH slice has no measurements for the relieving triangle, nor below the lintel. The relieving triangle thus starts at depth 4.7 meters. The lintel starts at depth 8.1 meters. The posterior mean of γ_1 for the preferred model is 7.8 meters, near the start of the lintel. The other popular model places the two changepoints lower (at an average of 7 meters and 10.5 meters respectively) than the structural points, but the latter are included in the 95% HPD intervals.

In addition, the estimated radius measures for the missing data along the OH slice are very close to the observed values from the five complete slices. Figures 7 and 8 show the estimated variances at each observation, for the two popular models, along with 95% pointwise confidence intervals. The vertical lines give the locations of the changepoints.

	2 changepoints		1 changepoint	
	Mean	95% HPDI	Mean	95% HPDI
$\log(\alpha_1)$	0.021	$[-0.011, 0.055]$	0.022	$[-0.01, 0.055]$
$\log(\alpha_2)$	0.475	$[0.09, 0.89]$	0.767	$[0.478, 1.20]$
$\log(\alpha_3)$	1.098	$[0.584, 1.848]$	—	—
β_1	0.844	$[0.821, 0.867]$	0.841	$[0.818, 0.864]$
β_2	0.614	$[0.426, 0.801]$	0.481	$[0.30, 0.603]$
β_3	0.353	$[0.054, 0.558]$	—	—
γ_1	6.96	$[4.69, 8.77]$	7.8	$[6.42, 9.30]$
γ_2	10.48	$[7.70, 12.56]$	—	—
σ^2	0.001	$[0.0007, 0.0021]$	0.001	$[0.0007, 0.0023]$

Table 3: **Results for the two most visited models – a one changepoint model and a two changepoint model; in both cases, the bottom segment is not a constant.**

As can be seen from the figures, the variance increases at greater depths, that is, near the bottom of the tomb. Our supposition is that this is due to the fact that near the top, where the angles must be steeper in order to create the characteristic shape, and to close off the space, there is less opportunity for variability in placing the stones. At the bottom, this is not the case.

One of the challenges in understanding the results from a complex model–simulation combination such as the one presented here, is to find clear ways of representing the outcome. Elsewhere (Lazar and Kadane, 2002) we have suggested the use of a statistical movie, to look at the posterior distribution output from high parameter problems, where simple trace plots might not suffice. Figure 9 shows another way of visualizing the simulation output for our particular model. In the Figure, we have plotted a typical model sampled from the posterior distribution of α s, β s and γ s, for the model with one changepoint and no constant segment at the bottom. Also on this plot is the full posterior distribution for the γ parameter, that is, the location of the changepoint in this model. The two together give information on where the changepoints for these models tend to be located, how variable these locations are, and what sorts of fits might result from this particular placement of changepoints. While the plot doesn’t have the dynamic nature of the movie, it does provide much more information than either a simple movie snapshot or a trace plot on their own could. A similar plot for the two changepoint model with no constant appears in Figure 10.

Inspection of the Figures reveals some interesting features of the posterior distribu-

Figure 7: Estimated variance at each measurement depth, for the model with one changepoint and no constant. The location of the changepoint is given by the vertical line in the plot, denoted by “G”.

Figure 8: Estimated variance at each measurement depth, for the model with two changepoints and no constant. The locations of the changepoints are given by the vertical lines in the plot, denoted by “G”.

Figure 9: Snapshot of one simulation of the model with one changepoint and no constant segment at the bottom. Plotted along the x -axis is the posterior distribution of the changepoint location for all samples from models of this type. The “G” marks the placement of the changepoint for the simulation shown and corresponds to the $\diamond$ in the fit.

Figure 10: Snapshot of one simulation of the model with two changepoints and no constant segment at the bottom. Plotted along the x -axis are the posterior distributions of the changepoint locations for all samples from models of this type. The “G”s mark the placement of the changepoints for the simulation shown and correspond to the $\diamond$ s in the fit.

tions. Much of the effort is spent in placing changepoints near the bottom, indicating that more adjustments were required during the construction of the lower layers of the tomb than of the higher ones. However, the model also considers the possibility of placing a changepoint near the top of the tomb. This can be seen by the long tail in the posterior for the location of the first changepoint in Figure 10.

Finally, it is possible to plot the output in three dimensions, as in Figure 11, which shows the data, along with the fitted model from a random iteration in the chain. A movie of this three dimensional plot shows how the fitted shape of the tomb changes as the simulation progresses. A clip from such a movie, which lasts approximately one minute, can be viewed at <http://www.stat.cmu.edu/~nlazar/movie.mpeg>.

4 Discussion

The analysis of corbelled domes, as described here, demonstrates the way in which advances in statistical thinking and methodology can interact with a scientific discipline, to further understanding of some aspect of that discipline. Indeed, in the example developed in this article, it is clear that statistical methods needed to adapt to the nature of the data and the problem – stretching the initial, deterministic models to allow for a Bayesian changepoint analysis, for example, required the development of Stephens (1994). Going in the other direction, Fan and Brooks (2000) utilized the new statistical technique of reversible jump MCMC to introduce the possibility of multiple changepoints, but of an unknown number. Here, the methodology existed, although the application was

Model with One Change Point and No Constant

Figure 11: Snapshot of one iteration of the chain, plotted as a three dimensional structure. The observed data are the underlying points in the image and the lines are the fitted values of the piecewise linear components. The triangle indicates the location of the changepoint.

new. Finally, incorporating the true, three-dimensional nature of the structures, and the availability of such data, albeit partial in some cases by necessity, required an additional generalization of Fan and Brooks (2000). Indeed, as we have argued, the presence of missing data in the slice that passed through the doorway of the tomb meant that simply averaging the data from different cross-sections, as had been previously done, would have resulted in a loss of information. In other words, what has unfolded here is the ideal interplay between statistics and the demands of the scientific discipline, applications driving new theory, which is turned back again to the applied problem.

The process will not end at this point – nor should it. Corbelled structures are very common across the Mediterranean region, and elsewhere. Archaeologists are still uncertain as to whether the method developed independently in different eras and regions, or was spread via word of mouth and apprenticeships. While statistics alone is not likely to answer this question, the use of more complicated spatio-temporal models that incorporate the age and geography of the structures might help to narrow the possibilities down and provide further research directions for the archaeologists. This is the logical next step in the analysis of these data.

Résumé

Avant le développement du dôme vrai, beaucoup de cultures antiques ont employé la technique de “corbelling” pour couvrir les espaces. Récemment, on a proposé une série de modèles statistiques relatifs dans la littérature pour expliquer comment des dômes corbellés pourrait avoir été construit. Le plus sophistiqué de ces modèles est basé sur une structure linéaire, avec un nombre inconnu de changepoints, pour guider le processus de bâtiment. Ce modèle est analysé par la technique “Reversible Jump Markov Chain Monte Carlo” (RJMCMC). Tous les modèles considérés jusqu’ici ont été bidimensionnels, c.-à-d., ils ont pris une coupe simple par le dôme; même lorsque des données plus étendues, sous forme de mesures sur les tranches multiples par le dôme, ont été disponibles, ceux-ci ont été ramenés à une moyenne ensemble pour les buts de l’analyse. En cet article, nous prolongeons l’analyse bidimensionnelle à une analyse tridimensionnelle, qui profite pleinement des données rassemblées par les archéologues et des symétries de rotation inhérentes à la structure. Nous explorons également des manières de présenter en forme de graphiques les résultats d’une exécution complexe du RJMCMC, afin de vérifier la convergence, le bon mélange, et l’exploration appropriée de l’espace de paramètre (élevé et variable de dimension). Le modèle et les techniques graphiques sont démontrés sur le trésor d’Atreus dans Mycenae, Grèce, un des exemples existants les plus fins de la méthode de corbelling.

REFERENCES

- http://projects.dartmouth.edu/history/bronze_age/ *The Prehistoric Archaeology of the Aegean*, Lessons 6 and 19.
- Buck, C.E., Litton, C.D. and Stephens, D.A. (1993) Detecting a change in the shape of a prehistoric corbelled tomb. *The Statistician*, **42**, 483–490.
- Casella, G. and George, E.I. (1992) Explaining the Gibbs sampler. *The American Statistician*, **46**, 167–174.
- Cavanagh, W.G. and Laxton, R.R. (1981) The structural mechanics of the Mycenaean tholos tomb. *Annual of the British School at Athens*, **76**, 109–140.
- Cavanagh, W.G., Laxton, R.R. and Litton, C.D. (1985) An application of change-point analysis to the shape of prehistoric corbelled domes; I. The maximum likelihood method. In *To Pattern the Past*, A. Voorrips and S.H. Loving (Eds.), Part 11, III 4, 191–203.
- Chib, S. (1995) Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association*, **90**, 1313–1321.
- Chib, S. and Jeliazkov, I. (2001) Marginal likelihood from the Metropolis–Hastings output. *Journal of the American Statistical Association*, **96**, 270–281.
- Denison, D.G.T., Mallick, B.K. and Smith, A.F.M. (1998) Automatic Bayesian curve fitting. *Journal of the Royal Statistical Society, Series B*, **60**, 333–350.

- DiCiccio, T.J., Kass, R.E., Raftery, A. and Wasserman, L. (1997) Computing Bayes factors by combining simulation and asymptotic approximations. *Journal of the American Statistical Association*, **92**, 903–915.
- Fan, Y. and Brooks, S.P. (2000) Bayesian modelling of prehistoric corbelled tombs. *The Statistician*, **49**, 339–354.
- Gelman, A. and Meng, X.L. (1998) Simulation normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Science*, **13**, 163–185.
- Green, P.J. (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82**, 711–732.
- Hinkley, D.V. (1969) Inference about the intersection of two-phase regression. *Biometrika*, **56**, 495–504.
- Hinkley, D.V. (1971) Inference in two-phase regression. *Journal of the American Statistical Association*, **66**, 736–743.
- Lazar, N.A. and Kadane, J.B. (2002) Movies for the visualization of MCMC output. *Journal of Computational and Graphical Statistics*, **11**, 863–874.
- Meng, X.L. and Wong, W.H. (1996) Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica*, **6**, 831–860.
- Raiffa, H. and Schlaifer, R. (2000) *Applied Statistical Decision Theory*, New York: Wiley.
- Richardson, S. and Green, P.J. (1997) On Bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society, Series B*, **59**, 731–

792 (with discussion).

Stephens, D.A. (1994) Bayesian retrospective multiple- changepoint identification. *Applied Statistics*, **43**, 159–178.

Tierney, L. and Kadane, J. (1986) Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, **81**, 82–86.