ESTIMATING AGE-AT-DEATH DISTRIBUTIONS FROM SKELETAL SAMPLES:

A MULTIVARIATE LATENT TRAIT APPROACH

Darryl J. Holman[1]

James W. Wood[2]

Kathleen A. O'Connor[1]

[1]Department of Anthropology and
Center for Studies of Demography and Ecology
University of Washington
Seattle, WA 98195, USA

[2]Department of Anthropology and
Population Research Institute
Pennsylvania State University
University Park, PA 16802, USA

Introduction

Most approaches to age estimation currently used in paleodemography and forensic science are not based on formal (or even informal) statistical methods. Instead, various *ad hoc* procedures have been developed, frequently based on simple tabulations of skeletal markers by age. The classic methods of Todd (1920) and McKern and Stewart (1957), for example, involve a non-statistical assignment of a skeleton's age at death according to documented changes in the pubic symphysis. These methods produce either a non-statistical age range or a point estimate of age, without any assessment of the error structure of the estimate based on formal probability arguments. The individual ages produced in this way are then aggregated to estimate the age-at-death distribution for an entire sample. As discussed elsewhere in this volume, the age-at-death distribution produced by this procedure will usually be biased in the direction of the age distribution of whatever reference sample was used to generate the individual estimates in the first place. In addition, we are left with little understanding of the degree of estimation error involved, either in the individual age estimates or the estimate of the aggregate-level age-at-death distribution as a whole.

In this paper we explore some statistical methods for estimating age-at-death distributions from skeletal samples, with special emphasis on recovering the parameters of parametric models of the age-at-death distribution (see Wood et al., this volume). Only methods compatible with the Rostock protocol, described elsewhere in this book, are discussed. We begin by examining univariate methods – those that use a single skeletal age indicator – and then go on to examine multivariate methods. We introduce a new multivariate method for estimating a parametric age-at-death distribution from a skeletal sample. The method at least partially corrects for the correlations that almost inevitably exist among skeletal traits, and handles missing observations on particular traits.

Estimation of an age-at-death distribution

The data used for paleodemographic reconstruction of a population's age-at-death distribution are macro- and microscopic morphological indicators of age at death from individual skeletons. A considerable body of work has appeared over the past eighty years on the

identification and quantification of age-related morphological changes in the human skeleton for use as indicators of age at death. Despite this work, the correlation of skeletal indicators with true chronological age, and the accuracy and reliability of most indicators, remain far from ideal (Bocquet-Appel and Masset 1982, 1985; Buikstra and Konigsberg 1985; Jackes 1992). The limitations are partly biological, and, aside from developing new and more biologically informative indicators, little can be done to improve upon them. There is considerable room, however, for improvement in the statistical methods used for paleodemographic reconstruction.

Methodological advances are needed in at least three areas. First, methods are needed that produce a target age-at-death distribution that does not mimic the age-distribution of the reference sample.[1] This "age mimicry" bias was empirically demonstrated by Bocquet-Appel and Masset (1982) and mathematically explained by Konigsberg and Frankenberg (1992, 1997), who also proposed a statistical solution to the problem. We build on the methods of Konigsberg and Frankenberg in this paper.

The second area concerns how age estimates are produced for individual skeletons in an archaeological target sample. Traditionally, ages have been assigned to a skeleton directly from that individual's skeletal age indicators. As discussed by Müller and Love (this volume), ages produced in this way are usually biased. In most applications, accurate individual ages can be found only *after* the age-at-death distribution has been estimated for the entire target sample. And even then, the resulting age for each target skeleton should be reported as a *distribution* of probable ages, not just as a point estimate.

The third area is the development of multivariate aging methods that accommodate missing skeletal age indicators. The ideal method would allow multiple aging indicators to be combined in way that makes statistical sense. Clearly, one motivation for developing such a method is to wring as much reliable aging information as possible from every skeleton. Another

[1] As in the rest of this book, *reference* is used throughout this paper to indicate an individual skeleton or sample of skeletons of known age used to calibrate our age estimation procedure. *Target* refers to the archaeological or forensic skeleton(s) whose age(s) at death we wish to estimate. These usages follow Konigsberg and Frankenberg (1992). As emphasized by Usher (this volume) the "known" ages reported for many famous reference collections are often quite approximate. We ignore this problem and treat reference sample ages as if they were known without error.

compelling motivation is that, in real skeleton collections, most if not all skeletons will be missing one or more aging indicators for taphonomic reasons.

Konigsberg and Frankenberg (1992) proposed a multivariate method for estimating age-at-death distributions using continuous age indicators. An extension to discrete age indicators was given by Konigsberg and Holman (1999). Both methods estimate a series of means from a multivariate normal (or Gaussian) distribution for the joint distribution of all age indicators, along with the entire variance-covariance matrix among indicators. Using this method with a set of 10 indicators, all distributed as multivariate normals, would require us to estimate a total of 65 parameters: 10 means, 10 variances, and 45 covariances. As the number of age indicators increases the method becomes even more parameter-heavy, which, in turn, requires larger and larger sample sizes. In addition, numerically intensive methods must be used for multivariate integration, since the method always requires integration in one more dimension than there are age indicators. The strength of the method is that it does not require us to assume statistical independence among age indicators.

Boldsen et al. (this volume) propose a related method, called *transition analysis*, that generates an age-at-death distribution from the joint distributions of a series of skeletal age indicators. This approach makes the simplifying assumption that the indicators are independent of each other once they have been conditioned on chronological age. For 10 binary indicators, each with an independent distribution (normal, logistic, etc.), the method yields 20 parameters: 10 location parameters and 10 scale parameters. Boldsen's approach is considerably simpler than Konigsberg's, if only because no integration is necessary for estimating the parameters from the reference sample. Similarly, sample size is less of a problem because fewer parameters are estimated. But Boldsen's method comes at a price: we are required to make the possibly erroneous assumption that the indicators are independent of each other conditional on age.

In this paper, we develop an alternative approach to estimating a multivariate age-at-death distribution – an approach we call the "latent trait" method for reasons that will become clear presently. Our method represents something of a compromise between the two methods discussed above. Age indicators are not considered conditionally independent of each other as in Boldsen's transition analysis, but neither do they require estimation of the entire variance-covariance matrix as in Konigsberg's multivariate probit method. Our method also falls between

the others in numerical complexity: numerical integration is required in a single dimension for parameter estimation from a reference sample and in two dimensions for recovering the age-at-death distribution from a target sample. The advantages of this method are threefold: we do not need to assume that age indicators are independent of each other, the number of parameters to be estimated grows linearly (not exponentially) with the number of indicators, and the method is numerically tractable. In addition, our model is motivated by some simple biological principles, so that some parameters may be of genuine biological interest.

Statistical age-at-death estimation can be logically separated into two distinct stages. The first stage is the generation of one (or more) standard age distributions from a known-age reference sample. The second stage is the estimation of an age-at-death distribution from some target sample, making use of the reference distribution(s) found in the first stage. Throughout this paper, we explicitly divide every method into these two parts and provide the corresponding likelihoods for both. Maximum likelihood methods are then used to estimate parameters. The basic idea of maximum likelihood estimation is to compute a probability (or an individual likelihood) for each observation given some underlying probability model of the process. The overall likelihood of the model, given a series of independently sampled cases, is the product of the individual likelihoods. The parameter values that globally maximize the overall likelihood are the maximum likelihood estimates (MLEs). Useful introductions to maximum likelihood estimation are provided by Edwards (1972), Pickles (1985), and Eliason (1993).

Missing skeletal observations

Because of differential preservation and recovery, few skeletons display all possible indicators of age. In almost every collection of skeletal material, there will be missing indicators for at least some of the individual skeletons. As an example, Table 1 shows the distribution of multiple age-at-death indicators in human skeletons from the archaeological site of Tipu in Belize. In this case, most skeletons (318 of 532 juveniles and adults) could be aged by only one or two indicators. Only five could be aged by all six indicators. When using multivariate aging methods, missing observations for one or more indicators are likely to be the norm. Any serious multivariate aging method must be able to accommodate such missing data.

The particular adult age-at-death indicators used at Tipu are listed in Table 2. The skeletal material from this site is, comparatively speaking, reasonably well preserved. Nonetheless, while the pubic bone is one of the most common (and one of the best) indicators of adult age, only 33 of 255 adults had pubic bones in sufficiently good condition for aging purposes. For the Tipu collection, we would have four options for age-at-death reconstruction: we could drop the pubic bone from our suite of age indicators, we could base our age-at-death methods on only 33 skeletons (!), we could somehow combine multiple univariate methods for different indicators, or we could use a genuinely multivariate method that accommodates the fact that the pubic bone is missing in most skeletons.

Because individual skeletons vary in the age indicators they display – and because each indicator varies in its reliability and accuracy – skeletons will differ in the quality and quantity of information they contribute to any estimate of an age-at-death distribution. Most investigators have not addressed this problem except with *ad hoc* solutions. For example, summary age at death, a simple unweighted average of all indicators available for an individual skeleton, is a common method for combining multiple univariate age-at-death estimates to come up with a single point estimate of individual age (Acsádi and Nemeskéri 1970). Simple averaging is clearly improper for several reasons: (1) the different age indicators do not provide exactly the same amount of information, (2) the indicators may not be independent, (3) all information about the error structure of the individual age estimates is thrown away, and (4) the age-at-death distribution for a target sample must be estimated by aggregating the individual age estimates, introducing the risk of reference sample age mimicry (Konigsberg and Frankenberg 1992).

Some researchers have advocated weighting indicators, but there is no agreed-upon, statistically valid method currently available for selecting the weights. One popular method, "multifactorial aging" (Lovejoy et al. 1985), weights each age indicator according to its loading on the first principal component estimated from the correlation matrix of all indicators (on the assumption that the first principal component represents true chronological age). In theory, the principal components analysis is supposed to be performed on the target sample, not the reference sample. However, principal components analysis requires complete information for each individual, and the numbers of complete individuals in most skeletal samples are far too

small to support such an analysis. In the Tipu sample, for example, only five skeletons display all six indicators (Table 1).

The problem of missing data cannot be ignored for any real skeletal sample. It is essential, therefore, to develop a systematic multivariate method for handling missing skeletal data without resorting to *ad hoc* procedures and adjustments. In the multivariate methods we discuss below, we pay particular attention to dealing with missing data. Our general approach is to assume that data are *missing at random*, by which we mean that the parameters defining the probability of an indicator being missing are independent of the parameters for the age-at-death distribution itself. This assumption may not always be a good one for skeletal data. For example, postmortem preservation of skeletal age indicators may vary with age at death since the bones of very young and very old individuals often do not survive as well as those drawn from the middle portion of the age range (Walker et al. 1988). This differential preservation by age is potentially an important source of bias, since it is likely to result in a disproportionate number of missing observations for older adults and young children. But the assumption that data are missing at random still allows for a more satisfactory treatment of missing observations than has been possible in the past.

Univariate Methods

As already mentioned, traditional methods for using age standards derived from a reference sample to compute point estimates of age at death for individual skeletons in a target sample can be seriously biased. There are two circumstances, however, when the traditional methods actually work (Konigsberg and Frankenberg 1992). The first is when the age indicator is almost perfectly correlated with chronological age, as in the case of annual tree rings.[2] The second circumstance is when the skeletons making up the reference sample are uniformly distributed by age. If one of these conditions is not satisfied, then the more complex procedures described below must be used in order to avoid age mimicry – bias in the estimated target age-at-death distribution reflecting the age distribution of the reference sample.

[2] Cemental annulations are sometimes touted as such indicators by human osteologists, but – as the validation studies presented by **????** in this volume clearly show – their correlation with true age is actually much lower than that of tree rings.

The life-table method

One of the simplest procedures for generating a life-table age-at-death distribution was given by Konigsberg and Frankenberg (1992). In this section, we briefly review their method in order to set the stage for the more complicated methods that follow.

We begin with a reference sample made up of N_r skeletons whose ages at death are known and who have been scored for a single age indicator (the subscript *r* denotes the *reference* sample). The indicator might be pubic symphysis stage, osteon count, suture closure stage, or dental root development stage. The age indicator is assumed to have *m* non-overlapping states. For an indicator such as suture closure, *m* might be two, closed or unclosed. For features of the pubic symphysis, *m* will usually be some larger number (e.g. six for one component of the Gilbert and McKern pubic system).

Our goal is to use this known-age reference sample to estimate an age-at-death distribution for a target sample of N_t individuals whose ages are unknown but for whom we have scored the relevant age indicator. The result of our analysis, in this particular case, will be an age-at-death distribution in the form of a life table with *w* discrete age intervals.

Estimating the parameters of the reference distribution – We begin by computing each element p_{ia} of the matrix **P** as the relative frequency in the reference sample of individuals in indicator state *i* given age *a* (where *a*, in this case, denotes a single age interval). This array is constructed from simple cross-tabulations. The resulting estimated elements, called $\hat{p}_{ia}$, are the probabilities of observing indicator stage *i* for some age *a* in the reference sample. We use carets (^) over parameters to denote values estimates from a sample – which differs from the way in which Konigsberg and Frankenberg (1992) use carets.

Estimating the target age-at-death distribution – The probability of someone in the target population dying in the *a*-th life-table age interval is denoted d_a. Initially we do not know the value of each d_a. The goal is to estimate the age-at-death distribution $\hat{\mathbf{d}} = \hat{d}_1, ..., \hat{d}_w$ subject to the constraint

(1) $$\sum_{a=1}^{w} d_a = 1, \quad 0 \le d_a \le 1.$$

As shown by Konigsberg and Frankenberg (1992), we can find maximum likelihood estimates of $\hat{\mathbf{d}}$ as follows. Given $\hat{p}_{ia}$, we can compute p_i, the probability of observing indicator stage i assuming a target age distribution. For convenience, we define

(2) $$p_i = \sum_{a=1}^{w} \hat{p}_{ia} d_a \ .$$

Then, the likelihood function for a given $\hat{p}_{ia}$ and some target distribution d_a is

(3) $$L = \prod_{i=1}^{N_t} p_1^{\delta_{1i}} p_2^{\delta_{2i}} p_3^{\delta_{3i}} \cdots p_n^{\delta_{ni}} \ ,$$

where δ_{ij} is an indicator variable that is equal to one if the j-th target individual is in stage i and zero otherwise. We can rewrite this likelihood as

(4) $$L = \prod_{j=1}^{N_t} \prod_{i=1}^{n} \left[\sum_{a=1}^{w} \hat{p}_{ia} d_a \right]^{\delta_{ji}} .$$

All the life-table probabilities in **d** are estimated simultaneously as the set of numerical values that maximizes the likelihood in equation (3) or (4). Additional discussion of this method, along with paleodemographic examples, can be found in Konigsberg and Frankenberg (1992) and O'Connor (1995).

As outlined by Wood et al. in this volume, describing an age at death distribution with a life table is not ideal. Most human mortality distributions can be well described using five or fewer parameters, so that a parsimonious parametric model should be used in place of life tables. In what follows, we discuss methods that are fully parametric, both for the distribution of age indicators as well as for the age at death distribution.

A parametric univariate method

In this section we discuss univariate methods for estimating a parametric age-at-death distribution when age (a) is continuous. For simplicity, we initially focus on age indicators that undergo a single transition (e.g. a suture that makes a transition from opened to closed). Discussion of the more complicated "staged" indicators (which can be viewed as multivariate data) is postponed for a later section.

Estimating parameters for the reference distribution – As before, we begin with some reference sample of N_r individuals for whom exact ages at death are known. For each skeleton in the reference sample, we observe an indicator state. In what follows, the indicator can either be present or it can be absent. For example, we might have recorded whether a particular suture is opened (absent) or closed (present) in a known-age reference sample of skeletons.

Following Boldsen et al. (this volume), we will refer to the age at which the indicator went from absent to present as the *transition age*. Let $f(a|\mu,\sigma)$ denote the probability density function (PDF) for the age at which the transition occurs in all human populations – the assumption of invariance discussed by Müller and Love in this volume. It is often reasonable to assume that $f(a|\ \mu,\sigma)$ is either a normal, lognormal, or logistic distribution, but it could be any parametric probability density function – preferably one that somehow mimics the underlying biological processes.[3] And we will make frequent use of both the cumulative distribution function (CDF) associated with $f(a|\ \mu,\sigma)$:

$$F(a\,|\,\mu,\sigma) = \int_0^a f(x\,|\,\mu,\sigma)dx, \tag{5}$$

and the corresponding survival function:

[3] Throughout this paper μ and σ are used to representing the location and scale parameters of a two-parameter distribution. The distribution may have more than two parameters as well. Although not strictly necessary, many of the likelihoods that follow assume that the PDF is 0 for negative ages.

$$S(a \mid \mu, \sigma) = 1 - F(a \mid \mu, \sigma) = \int_a^{\infty} f(x \mid \mu, \sigma) dx. \tag{6}$$

The goal at this point is to find $\hat{\mu}$ and $\hat{\sigma}$, the estimates of μ and σ from the reference sample. These two parameters completely describe the distribution of transition ages.

Reference samples are usually observed cross-sectionally. That is, the aging indicator is observed only once, at the age at death.[4] Based on the state of the indicator, the skeleton has, at the time of death, either made the transition or has not. The likelihood for a skeleton that has made the transition is constructed by specifying the probability that reference individual j aged a_j made the transition to the indicator state at some unknown age between birth and a_j. This probability is given by the entire area under the PDF to the left of age a_j, equal to $F(a|\mu,\sigma)$, the cumulative distribution at age a.

For a reference skeleton that did not make the transition by observation age a, the likelihood is the area under the PDF from age a to infinity, that is, the survival function at a, $S(a|\mu,\sigma)$. We will assume that all individuals who live long enough will eventually make the transition (an assumption that can be relaxed if needed; see Holman and Jones 1998).

An overall likelihood can be computed from a sample of N_r cross-sectionally sampled reference individuals, some who have made the transition ($\delta_j = 1$) and some who have not ($\delta_j = 0$) made the transition by the age at which they are observed, a_j. Taking the product of the individual likelihoods, we get

[4] In fact, this need not be the case. Depending on the specific age indicators being used, it might be possible to observe a living sample longitudinally. Methods for finding reference parameters from mixtures of interval-censored, right-censored, and cross-sectionally observed reference individuals are given in Wood et al. (1992) and Holman and Jones (1998).

$$L = \prod_{j=1}^{N_r} \left\{ \left[\int_{a_j}^{\infty} f(x \mid \mu, \sigma) dx \right]^{1-\delta_j} \left[\int_{0}^{a_j} f(x \mid \mu, \sigma) dx \right]^{\delta_j} \right\}. \tag{7}$$
$$= \prod_{j=1}^{N_r} \left[S(a_j \mid \mu, \sigma)^{1-\delta_j} F(a_j \mid \mu, \sigma)^{\delta_j} \right]$$

Estimating the target age-at-death distribution – Now assume that we have already found parameters $\hat{\mu}$ and $\hat{\sigma}$ for $f(a|\mu,\sigma)$ from an appropriate reference sample. For a target sample of N_t individuals, we want to estimate the parameters of a continuous age-at-death distribution, $g_d(a|\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is a vector of parameters (for a review of parametric age-at-death distributions, see the chapter by Wood et al. in this volume). Assume that only cross-sectional observations are made on indicators of the target sample, and that these indicators denote the state at death. When $\delta_j = 0$ the j-th target subject has not yet made the transition (the trait is absent), and when $\delta_j = 1$ the j-th target subject has completed the transition (the trait is present). We can rewrite equation (4) for this continuous case with indicator states "absent" and "present" as

$$L = \prod_{j=1}^{N} \left[\int_{0}^{\infty} F(a \mid \hat{\mu}, \hat{\sigma})^{\delta_j} S(a \mid \hat{\mu}, \hat{\sigma})^{1-\delta_j} g_d(a \mid \boldsymbol{\theta}) da \right]. \tag{8}$$

Maximizing equation (8) over $\boldsymbol{\theta}$ yields maximum likelihood estimates, $\hat{\boldsymbol{\theta}}$.

Parametric Methods With "Stage" or "Phase" Data

Many traditional aging methods are based on a series of stages or phases rather than single transitions. While these methods are useful in the non-statistical context for which they were developed, they add serious complications when we adapt them for the statistical methods discussed here. When the indicator of interest is not a present/absent indicator but has ordinal states (as in pubic symphysis stages), the above parametric method must be modified.

Most adult (senescent) age-at-death indicators are based on phases or stages, including the pubic symphysis (Todd 1920; McKern and Stewart 1957; Gilbert and McKern 1973), the auricular surface (Lovejoy et al. 1985), ectocranial suture closure (Meindl and Lovejoy 1985),

and morphological changes in the ribs (Iscan et al. 1984, 1985), proximal femur (Walker and Lovejoy 1985), or clavicle (Walker and Lovejoy 1985). Todd (1920) was the first to study the relationship between chronological age and metamorphosis of the articular face of the pubic symphysis in a systematic way. Using the Todd collection,[5] Todd described 10 modal phases for adults between the ages of 18 and 50 years, each phase corresponding to a specific age range (Table 3). Each Todd phase is defined by several different features of the pubic symphysis scored in combination, including the dorsal plateau, ventral rampart, symphyseal face, symphyseal rim, furrows, pitting, and so on. Todd eliminated any pubic symphysis in the sample that did not conform to what he considered the "normal" modal phases of development. Brooks (1955) modified Todd's phases by shifting the age range for the phases covering 26-45 years downward three years to correct for a tendency to over-estimate age.

McKern and Stewart (1957) used a reference sample of American soldiers killed in the Korean War to develop a three-component system for estimating age from the pubic symphysis. Although reported ages are much more accurate in this sample than in the Todd collection, the McKern-Stewart collection has a much more restricted age distribution, with few skeletons over the age of 30.

Each of the three McKern-Stewart components has six stages. To estimate age, each component is ranked on a scale of 0-5; then a sum of scores for the three components is totaled and compared with a table of scores and associated chronological ages (Table 4). Although a large number of combinations is theoretically possible, only 21 of these occur with any frequency. Because of the restricted age distribution, this system does not work well for ages beyond 30 or 40 years (O'Connor 1995).

It is important to realize that most "stage" or "phase" methods are trying to code for multiple morphological changes in a variety of structures. This is clear from reading the description of the first two Todd stages:

---

[5] Consisting of the skeletons of 465 indigenes from the Cleveland, Ohio, area (306 white males, 47 white females, 90 black males, and 22 black females). While this collection covers a broad range of reported ages, ages are often poorly known and display abundant heaping at years ending in zero or five (see Usher, this volume).

> (1) Symphysial surface rugged, traversed by horizontal ridges separated by well marked grooves; no ossific nodules fusing with the surface; no definite delimiting margin; no definition of extremities.
> (2) Symphysial surface still rugged, traversed by horizontal ridges, the grooves between which are, however, becoming filled near the dorsal limit.... Ossific nodules fusing with the upper symphysial face may occur; dorsal limiting margin begins to develop, no delimitation of extremities; foreshadowing of ventral bevel (Bass 1971:155).

Clearly, this is a multifactorial indicator involving many different types of surface remodeling on the pubic symphysis, with an occasional nod toward future changes. Similarly, the McKern-Stewart method uses multiple components to assign a score. For example, scores 3 and 4 of component three are described as follows:

> (3) The symphyseal rim is complete. The enclosed symphyseal surface is finely grained in texture and irregular or undulating in appearance.
> (4) The rim begins to break down. The face becomes smooth and flat and the rim is no longer round but sharply defined. There is some evidence of lipping on the ventral edge (Stewart 1979:163).

These stages do *not* represent a single biological trait that changes with age in a straightforward way. Rather, the stages are based on an entire *suite* of traits that are packaged together for descriptive convenience. Unfortunately, the convenience evaporates when we try to develop formal multivariate parametric methods for age-at-death estimation. We strongly recommend that new reference age indicators be developed that are based only on single transitions (or continuous measures for continuous indicators). Our reasons are twofold. First, it makes the math much easier and more logical. But more importantly, it discourages us from developing methods that treat complex traits as if they resulted from a single process.

Estimating parameters of the reference distribution using staged indicators – For each skeleton in the reference sample, we observe an indicator variable that includes $m > 2$ ordered stages. As before, the reference sample comprises N_r skeletons of known age at death. One way to treat these data is to define $m - 1$ transitions that occur from one phase to the next, and use $f_1(a|\,\mu_1, \sigma_1)$, $f_2(a|\,\mu_2, \sigma_2), \ldots, f_{m-1}(a|\,\mu_{m-1}, \sigma_{m-1})$ to denote the PDFs for the ages at which each transition occurs in the population. The multivariate methods described in the following sections can then be used to estimate all $m - 1$ distributions.

To estimate the distributions, we define a set of $m - 1$ stage transition variables T_2 to T_m (or T_1 to T_{m-1} if the phases are numbered from 0) that are set equal to one if the transition has

been made for that stage and zero if not. Consider, for example, one particular McKern-Stewart pubic symphysis component. Table 5 shows how to convert component phases 0 to 5 into the five transition variables T_1 to T_5. We can find the $\hat{\mu}$s and $\hat{\sigma}$s corresponding to these transition variables by maximum likelihood using one of the multivariate methods described later in this paper. The results can then be used to estimate the target age-at-death distribution using the corresponding multivariate method.

Alternative methods for staged indicators – Alternative methods for handling staged data that can be found in the statistical literature include the ordered probit or logit method, which treats the entire set of phases as a single ordered process. The model was introduced by McKelvey and Zavoina (1975) and McCullagh (1980), and is commonly used in the social sciences (Long 1997). An example of the method applied to the estimation of an age-at-death distribution can be found in the chapters by Hermann and Konigsberg and by Boldsen and colleagues.

Multivariate Methods

Independent age indicators

Now consider the case in which multiple age indicators are observed. Each individual has one or more of these indicators. Initially, we assume these indicators are completely independent of each other. Later we will discuss ways of treating non-independence among indicators.

Estimating parameters of the reference distribution – For n independent age indicators scored in our reference sample, we can simply take the product of the individual likelihoods for each indicator. A sample of N_r reference individuals thus yields the likelihood:

$$
\begin{aligned}
L &= \prod_{j=1}^{N_r}\left\{\prod_{i=1}^{n}\left[\int_{a_j}^{\infty} f(x\,|\,\mu_i,\sigma_i)dx\right]^{(1-\delta_{ij})\varepsilon_{ij}}\left[\int_{0}^{a_j} f(x\,|\,\mu_i,\sigma_i)dx\right]^{\delta_{ij}\varepsilon_{ij}}\right\}, \\
&= \prod_{j=1}^{N_r}\left[\prod_{i=1}^{n} S(a_j\,|\,\mu_i,\sigma_i)^{(1-\delta_{ij})\varepsilon_{ij}}\, F(a_j\,|\,\mu_i,\sigma_i)^{\delta_{ij}\varepsilon_{ij}}\right]
\end{aligned}
\tag{9}
$$

where δ_{ij} is an indicator variable that is equal to one if the *j*-th reference individual for the *i*-th age indicator is present and 0 if it is absent. The indicator variable ε_{ij} denotes that the age indicator is available for scoring. It is set to zero when the *i*-th age indicator is missing for the *j*-th individual, and one if it is not. This, in effect, yields a likelihood of one for each missing observation so that it makes no contribution to the overall likelihood.

It is important to realize that equation (9) assumes that the probability of each transition is independent of all other transitions in the same individual. If some age indicators are correlated, then estimates for the μs and σs could be biased. The degree of bias is an empirical question for any combination of age indicators and reference sample.

Estimating the target age-at-death distribution – The extension for estimating the age-at-death distribution in the target sample is straightforward. Again we assume indicators are independent, and we have estimated the parameters $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\sigma}}$ for the distributions $f_i(a|\mu_i,\sigma_i)$, i = 1..*n*, from the reference sample. The transition indicator δ_{ij} is now equal to one if the *j*-th target individual has made transition for age indicator *i*, and zero otherwise. The likelihood is a straightforward multivariate extension of equation (8), in which we place all the independent reference distributions inside the integral:

$$
L = \prod_{j=1}^{N_t}\int_{0}^{\infty} g_d(a\,|\,\boldsymbol{\theta})\prod_{i=1}^{n}\left[S_i(a\,|\,\hat{\mu}_i,\hat{\sigma}_i)^{(1-\delta_{ij})\varepsilon_{ij}}\, F_i(a\,|\,\hat{\mu}_i,\hat{\sigma}_i)^{\delta_{ij}\varepsilon_{ij}}\right]da\,. \tag{10}
$$

Missing age indicators are handled by setting ε_{ij} to zero.

Non-independent indicators: The full multivariate method

Estimating parameters of the reference distribution – The full method for handling multiple age indicators is to treat them as following some multivariate distribution that includes

all covariance terms. If, for example, the indicators are assumed to be multivariate lognormal, then we use a multivariate lognormal distribution, including all covariances. For n independent age indicators, let $\mathbf{m}$ be the array of n means ($\mu_1, \mu_2, \ldots, \mu_n$) and $\mathbf{V}$ the $n \times n$ variance-covariance matrix. For a sample of N_r reference individuals, the likelihood is

$$L = \prod_{j=1}^{N_t} \int_{a_j(\mathbf{1}-\mathbf{d}_j)\boldsymbol{\varepsilon}_j}^{a_j/(\mathbf{d}_j\boldsymbol{\varepsilon}_j)} \mathbf{f}(\mathbf{x}|\mathbf{m},\mathbf{V})d\mathbf{x}\,. \tag{11}$$

Note that this likelihood requires us to integrate the array of ages over all dimensions of the multivariate distribution $\mathbf{f}(\mathbf{a}|\mathbf{m},\mathbf{V})$. The array $\mathbf{d}_j$ consists of n indicator variables for the j-th individual: element δ_{ij} is equal to one if the i-th transition has occurred in that individual, and zero if it has not occurred. Then $\mathbf{1}\text{-}\mathbf{d}_j$ is the array of complements of $\mathbf{d}_j$ and $\boldsymbol{\varepsilon}_j$ is a vector of n indicators denoting missing observations (ε_{ij} is 0 if missing, 1 if not). The upper limit of integration is set to infinity whenever an observation is missing *or* the age indicator is absent (but not missing) in any given dimension. The lower limit of the integral goes to zero for a missing indicator or when the age indicator is present in a dimension. In this way, a missing age indicator results in a marginal likelihood of one (integration from 0 to infinity in that dimension). When the indicator is not missing, the limits of integration in one dimension will be from a_j to infinity if the transition has not been made, and from 0 to a_j if the transition has been made. This method of estimating $\hat{\mathbf{m}}$ and $\hat{\mathbf{V}}$ is identical to multivariate probit analysis (Bock and Gibbons 1996; Chib and Greenberg 1998; Konigsberg and Holman 1999).

There are two practical difficulties with this full multivariate method. First, it is extremely numerically intensive. Multivariate integration over more than about five dimensions takes a great deal of computing time, even using very fast computers. To get around this difficulty, one of several methods of stochastic integration can be used, such as the Gibbs sampler or the Markov chain Monte Carlo method (see Konigsberg and Holman 1999; Hermann and Konigsberg, this volume). These methods make it feasible to integrate multivariate integrals to fairly high dimensions.

The second difficulty is posed by the number of parameters that must be estimated. As the number of age indicators increases, the number of parameters that must be estimated grows

as $\left(n^2+3n\right)/2$. For example, with two indicators we estimate two means, two variances, and one covariance term – a total of five parameters. For five indicators we must estimate 20 parameters (five means, five variances, ten covariances). And for ten indicators there are 64 parameters to estimate. If we wanted to use 20 indicators (for example, by observing the emergence of all the deciduous teeth), we would need to estimate 230 parameters! Alas, as the number of indicators grows, the reference sample size needed to estimate the parameters with any certainty increases. In response, we might be tempted to reduce the number of indicators we use for age estimation by throwing out data – not an appealing strategy.

Estimating the target age-at-death distribution – Estimation of the target age-at-death distribution for the multivariate case uses all the information from the full multivariate distribution $\mathbf{f}(\mathbf{a}|\mathbf{m},\mathbf{V})$. If we have already estimated $\hat{\mathbf{m}}$ (the n means) and $\hat{\mathbf{V}}$ (the $n \times n$ variance-covariance matrix) from the reference sample, we can find the likelihood for the N_t target individuals as a simple multivariate extension of equation (8):

$$L = \prod_{j=1}^{N_t} \int_0^{\infty} g_d(a|\boldsymbol{\theta}) \int_{a(\mathbf{1}-\mathbf{d}_j)\boldsymbol{\varepsilon}_j}^{a/(\mathbf{d}_j\boldsymbol{\varepsilon}_j)} \mathbf{f}(\mathbf{x}|\hat{\mathbf{m}},\hat{\mathbf{V}}) d\mathbf{x} da \,. \tag{12}$$

Non-independent indicators: The latent trait method

The latent trait method is intended as a compromise between the extremes of assuming that age indicators are independent of each other and trying to estimate the full multivariate distribution. The method is based on a model of a particular type of non-independence among age indicators, one that is of much lower dimensionality than the full variance-covariance matrix. Integration is required over only one dimension for finding parameters from the reference sample, and two dimensions for finding the age-at-death distribution from the target sample. The major advantage of this model compared to the full multivariate model discussed in the previous section is that the number of parameters increases linearly with the number of age indicators, not as the square.

The method is based upon a simple biological model for the underlying developmental or senescent process that affects the skeletal indicators of interest. For simplicity, we discuss the

method assuming that the indicators are developmental (growth-related) rather than senescent in nature. For concreteness, we use the first emergence of various deciduous teeth as our example of the skeletal indicators of interest. The principles apply equally well to senescent traits.

The method supposes that each child has its own individual growth rate *z*, and that the value of *z* acts to accelerate or decelerate emergence of all the child's teeth simultaneously (Figure 1). In a child with a low value of *z* – and thus a slow underlying growth trajectory – all teeth will emerge later, on average, then in a child whose *z* value is high. Under this model, the correlations among the various emergence times *within* a child reflect both the child's age *and* the value of its underlying growth parameter z_j.

The effect of *z* can be different for each tooth. For some teeth *z* may have almost no effect, for others the effect may be strong. The different effects of *z* can be seen as different slopes across *z* for the teeth in Figure 1. We use a series of parameters β_{zi} to describe the strength of association between latent trait *z* and age indicator *i*.

Although the model assumes that each child has its own unique growth trait, we do not attempt to measure the value of *z* for each child. This value is not directly observable, but rather is concealed or *latent*.[6] We assume that the trait has a particular parametric distribution among children – a distribution whose parameters are initially unknown. The lower panel of Figure 1 shows a hypothetical distribution of *z* among children in a population. Even though we cannot measure the *z* value for each child, we can estimate the entire distribution of *z* values among children, as well as the average effect of the latent trait on the emergence of each tooth.

The method controls for correlations among age indicators in a way similar to that of models of shared frailty and some random effects models (e.g. Hougaard 1986; Klein et al. 1999). The effect of *z* on the PDF $f_i(a|\mu_i, \sigma_i, z, \beta_{zi})$ or the survival function $S_i(a|\mu_i, \sigma_i, z, \beta_{zi})$ of transition times for the *i*-th aging indicator, can be modeled in one of two standard ways. The first is by using an accelerated failure-time model, in which the effect of *z* is either to accelerate or decelerate the time to the transition (Klein et al. 1999). One common specification for an

[6] We use *latent trait* in a biological sense to denote a continuous unmeasurable trait that affects a series of binary indicators. Konigsberg et al. (Chapter 10) use the term *latent variable* in the statistical sense to denote an underlying continuous variable that is revealed as a binary or staged indicator. Thus, the method we discuss here is a latent variable model (for a series of binary indicators) for which each indicator is also affected by an additional latent trait, *z*.

accelerated-failure time model is $f_i(a|\mu_i, \sigma_i, z, \beta_{zi}) = f_i[a|\mu_i \exp(z\beta_{zi}), \sigma_i]$, in which β_{zi} simply shifts the mean time to emergence up or down without changing the variance (Figure 2, top panel). A second standard way to model the effects of z is to assume that it increases or decreases the hazard of making the transition at each age. If a proportional hazards model is specified, the effect of z on the PDF of transition times is $f_i(a \mid \mu_i, \sigma_i, z, \beta_{zi}) = f_i(a \mid \mu_i, \sigma_i) S_i(a \mid \mu_i, \sigma_i)^{\exp(z\beta_{zi})-1} e^{z\beta_{zi}}$ and the effect of z on the SDF is $S_i(a \mid \mu_i, \sigma_i, z, \beta_{zi}) = S_i(a \mid \mu_i, \sigma_i)^{\exp(z\beta_{zi})}$. Under this specification both the mean and the variance of times to emergence change with different values of β_{zi} (Figure 2, lower panel).

The distribution of z must be specified parametrically – for example, as a gamma or normal distribution, both of which are often used in this kind of analysis. In the examples presented below, we use a normal distribution for z, which we denote $g_z(z|\mu_z,\sigma_z)$. The parameter μ_z is constrained to equal zero, and we estimate the σ_z parameter along with the arrays $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ for the n age indicators. When more than two age indicators are used, an array of n–1 β_{zi} parameters is found as well, each β_{zi} telling us something about the strength of association between z and the i-th age indicator. The value of β_{z1} is constrained to equal one, so that the other β parameters model the effect of z on the corresponding age indicators *relative* to its effect on the first age indicator.

Estimating parameters of the reference distribution – For a sample of N_r reference individuals and n age indicators, we need to modify equation (9), which assumed independence among all the aging indicators. We now want to estimate the function $g_z(z|0,\sigma_z)$ that describes how z varies among individuals. In addition, the distribution for each age indicator, $f_i(a|\mu_i,\sigma_i)$ or $S_i(a|\mu_i,\sigma_i)$, has a new parameter β_{zi} that describes how strongly the indicator is affected by the individual's underlying growth trajectory z. The necessary likelihood is

$$
\begin{aligned}
L &= \prod_{j=1}^{N_r} \int_{-\infty}^{\infty} g_z(z \mid 0, \sigma_z) \prod_{i=1}^{n} \left[\int_{a_i}^{\infty} f_i(x \mid \mu_i, \sigma_i, z, \beta_{z_i}) dx \right]^{(1-\delta_{ij})\varepsilon_{ij}} \left[\int_{0}^{a_j} f_i(x \mid \mu_i, \sigma_i, z, \beta_{z_i}) dx \right]^{\delta_{ij}\varepsilon_{ij}} dz \\
&= \prod_{j=1}^{N_r} \int_{-\infty}^{\infty} g_z(z \mid 0, \sigma_z) \prod_{i=1}^{n} \left[\left(S_i(a_j \mid \mu_i, \sigma_i, z, \beta_{z_i}) \right)^{(1-\delta_{ij})\varepsilon_{ij}} F_i(a_j \mid \mu_i, \sigma_i, z, \beta_{z_i})^{\delta_{ij}\varepsilon_{ij}} \right] dz
\end{aligned} \tag{13}
$$

The latent trait method can be used for either developmental or senescent traits. When both types of trait are available in the reference sample, it is conceivable that distributions for two separate latent traits (one for growth and one for senescence) can be estimated.

Estimating the target age-at-death distribution – We assume that the parameters $\hat{\boldsymbol{\mu}}$, $\hat{\boldsymbol{\sigma}}$, $\hat{\boldsymbol{\beta}}_{\boldsymbol{z}}$, and $\hat{\sigma}_z$ have already been estimated from the reference sample. The likelihood for the target sample is then an extension of equation (10), to which we add integration over the distribution of *z*. The likelihood for N_t target individuals is

$$
L = \prod_{j=1}^{N_t} \int_{0}^{\infty} g_d(a \mid \boldsymbol{\theta}) \int_{-\infty}^{\infty} \hat{g}_z(z) \prod_{i=1}^{n} \left[S_i(a \mid z, \hat{\beta}_{z_i}, \hat{\mu}_i, \hat{\sigma}_i)^{(1-\delta_{ij})\varepsilon_{ij}} F_i(a \mid z, \hat{\beta}_{z_i}, \hat{\mu}_i, \hat{\sigma}_i)^{\delta_{ij}\varepsilon_{ij}} \right] dz da . \tag{14}
$$

Application

In this section, we present an illustrative analyses using the latent trait method. The data set was provided by Lyle Konigsberg, who used it in Chapter 10. Reference and target distributions were by partitioning out of a sample of 737 known age males, each scored by the Suchey system. A target sample of 149 target individuals was drawn according to Gompertz-Makeham distribution with parameters $\alpha_1 = 0.01$, $\alpha_2 = 0.001$, and $\beta_g = 0.1$, and the reference distribution encompassed the remaining 588 individuals.

In our attempts to retrieve parameters for the age at death distribution of the data set we treated the six pubic phases as a series of five transitions representing five correlated age indicators, modeled as being log-normally distributed. Maximum likelihood estimates of the reference and target parameters were found by the latent trait method using equations (13) and (14), and for comparison we estimated the corresponding models assuming independence among traits by equations (9) and (10). A proportional hazards specification was used to model the

effect of *z* on the age indicator distributions. Parameters were estimated by numerically maximizing the loglikelihood using *mle* version 2.0 software (Holman 2000). Numerical integration was performed by 30-point trapezoidal approximations. Estimates of the standard errors were found by the method of Nelson (1982), which involves inverting a numerical approximation of Fisher's information matrix.

The latent trait model used to estimate the multivariate reference distribution has five μ and σ parameters, four β parameters, and one σ_z parameter. The resulting parameter estimates are given in **Table 6**. The σ_z parameter was not well estimated for the reference sample and the β parameters were not significantly different from zero. It appears that the transition times between different phases are relatively independent. To further explore this issue, we also fit the ten-parameter multivariate-independent model given by equation (9) obtaining log likelihoods of -671.83 for the reference sample. The Akaike Information Criterion (AIC) can be used to select the between the two models (Akaike 1973, 1992; Burnham and Anderson 1998). The difference in AIC is 117, suggesting the latent-trait model does provides a better fit to the data.

The parameter estimates derived from the reference sample by the latent-trait model were used, in turn, to estimate the parameters of a Gompertz-Makeham age-at-death distribution using data from the target sample. The Gompertz-Makeham model has three parameters, α_1, α_2, and β_g (see Wood et al., this volume, for details). Parameter estimates for the resulting age-at-death distribution are given in Table 7. Additionally, target ages were provided for the target sample, so we could estimate the parameters of the Gompertz-Makeham directly from the known target ages (Table 7). The parameters recovered by the latent trait model were very close to the parameters used for the simulation as well as the parameters estimated from the known ages of the target sample. Clearly, the β_g parameter was not well estimated by the latent trait method for the target sample. Nevertheless, the difference in AIC between the latent trait model and the multivariate independent model was 14.52 indicating that latent-trait model fits somewhat better than the model assuming independence.

Age-at-death distributions estimated from known ages and by the latent trait method are shown in Figure 3. The distributions recovered from the known ages and by the latent trait

method are not significantly different, but we note that the standard errors recovered by the latent trait method are quite large.

We conclude that the latent-trait model does a reasonable, though not perfect, job of recovering the parameters in these simulated age-at-death distributions. One of the difficulties of the estimates presented here is that the standard errors may have been poorly estimated by errors introduced in the numerical integration. Methods that use Markov Chain Monte Carlo for the integration and bootstrapped estimates of parameter uncertainty would be useful refinements of the method.

Conclusions

We have presented a method for estimating an age-at-death distribution from multivariate skeletal data with possible missing values. The method adheres rigorously to the Rostock protocol outlined in other parts of this book. Thus, the method complements those used by Hermann and Konigsberg (Chapter 10) and Müller and Love (Chapter ??). It is also consistent with methods found elsewhere in the recent paleodemographic literature (Konigsberg and Frankenberg 1992; Konigsberg and Holman 1999; O'Connor 1995).

We would argue that skeletal data, by their nature, absolutely demand multivariate treatment. But three additional criteria must be met for any multivariate method to be of use to the practicing paleodemographer. First, the method must not assume that traits are statistically independent within an individual (for biological reasons, skeletal traits are unlikely to meet this assumption). Second, the number of parameters to be estimated must not grow as an exponential function of the number of traits (paleodemographic samples are too small to support estimation of a large numbers of parameters). Third, the method must be able to accommodate missing data for some skeletons (imperfect preservation almost inevitably results in missing data). We have developed the latent trait-method in response to these demands. Although it is a method of intermediate computational complexity – two nested integrals appear in the likelihood for the parameters of the age-at-death distribution – even this degree of complexity may require stochastic methods of integration such as the Markov chain Monte Carlo methods used in the chapter by Hermann and Konigsberg. Nonetheless, the latent trait approach represents a major gain in practicality over methods that estimate the full variance-covariance matrix among age

indicators – and a major gain in realism over methods that assume that indicators are independent.

The usefulness of this method (or any other multivariate method) rests, in part, on the development of true multivariate reference samples. The ideal reference sample would include numerous binary and continuous indicators from many parts of the skeleton, maximizing the chance that at least one indicator would be available for any skeleton. We eschew the notion of developing any stage or phase indicators — as we argued earlier, staged traits are likely to reflect multiple semi-independent processes that would be better coded as a series of binary traits or as a continuous trait.

Acknowledgments

We thank Jesper Boldsen, Rob Hoppa, Lyle Konigsberg, Brad Love, George Milner, Hans-Georg Müller, Bethany Usher, and Jim Vaupel for helpful comments and discussion. This research was supported in part by F32-HD07994 and NICHD Population Center Grant 1-HD28263.

Literature Cited

Acsádi G and Nemeskéri J (1970) *History of Human Life Span and Mortality*. Budapest: Akademiai Kaido.

Akaike H (1973) Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory*, ed. B.N. Petrov and F. Csaki, pp. 268-281. Budapest: Hungarian Academy of Sciences.

Akaike H (1992) Information theory and an extension of the maximum likelihood principle. In *Breakthroughs in Statistics*, ed. S. Kotz and N. Johnson, pp. 610-624. New York: Springer Verlag.

Bass WM (1971) *Human Osteology: A Laboratory and Field Manual of the Human Skeleton*. Columbia, Mo.: Missouri Archaeological Society.

Bock RD and Gibbons RD (1996) High-dimensional multivariate probit analysis. *Biometrics*, **52**, 1183-1194.

Bocquet-Appel JP and Masset C (1982) Farewell to paleodemography. *Journal of Human Evolution*, **11**, 321-333.

Bocquet-Appel JP and Masset C (1985) Paleodemography: Resurrection or ghost? *Journal of Human Evolution*, **14**, 107-111.

Brooks ST (1955) Skeletal age at death: The reliability of cranial and pubic age indicators. *American Journal of Physical Anthropology*, **13**, 567-597.

Buikstra JE and Konigsberg LW (1985) Paleodemography: Critiques and controversies. *American Anthropologist*, **87**, 316-333.

Burnham KP and Anderson DR (1998) *Model Selection and Inference: A Practical Information-Theoretic Approach*. New York: Springer Verlag.

Chib S and Greenberg E (1998) Analysis of multivariate probit models. *Biometrika*, **85**, 347-361.

Edwards AWF (1972) *Likelihood*. London: Cambridge University Press.

Eliason SR (1993) *Maximum Likelihood Estimation: Logic and Practice*. Newbury Park, Ca.: Sage Publications.

Gilbert BM and McKern TW (1973) A method for aging the female Os Pubis. *American Journal of Physical Anthropology*, **38**, 31-38.

Holman DJ and Jones RE (1998) Longitudinal analysis of deciduous tooth emergence II: Parametric survival analysis in Bangladeshi, Guatemalan, Japanese and Javanese children. *American Journal of Physical Anthropology*, **105**, 209-230.

Holman DJ (2000) *mle*: *A Programming Language for Building Likelihood Models.* Version 2 (Software and manual). http://faculty.washington.edu/~djholman/mle/

Hougaard P (1986) A class of multivariate failure time distributions. *Biometrika*, **73**, 671-678.

Iscan MY, Loth SR and Wright RK (1984) Age estimation from the rib by phase analysis: White males. *Journal of Forensic Science*, **29**, 1094-1104.

Iscan MY, Loth SR and Wright RK (1985) Age estimation from the rib by phase analysis: White females. *Journal of Forensic Science*, **30**, 853-863.

Jackes MK (1992) Paleodemography: Problems and techniques. In *Skeletal Biology of Past Peoples: Research Methods*, ed S.R. Saunders and M.A. Katzenberg, pp. 189-224. New York: Wiley-Liss.

Klein JP, Pelz C and Zhang M (1999) Modeling random effects for censored data by a multivariate normal regression model. *Biometrics*, **55**, 497-506.

Konigsberg LW and Frankenberg SR (1992) Estimation of age structure in anthropological demography. *American Journal of Physical Anthropology*, **89**, 235-256.

Konigsberg LW, Frankenberg SR and Walker RB (1997) Regress what on what? Paleodemographic age estimation as a calibration problem. In *Integrating Archaeological Demography: Multidisciplinary Approaches to Prehistoric Population*, ed. R.R. Paine, pp. 64-88. Carbondale, Ill.: Southern Illinois University Press.

Konigsberg L and Holman DJ (1999) Estimation of age at death from dental emergence and implications for studies of prehistoric somatic growth. In *Human Growth in the Past: Studies from Bones and Teeth*, ed. R.D. Hoppa and C.M. Fitzgerald, pp. 264-289. Cambridge: Cambridge University Press.

Long JS (1997) *Regression Models for Categorical and Limited Dependent Variables*. Thousand Oaks, CA: Sage.

Lovejoy CO, Meindl RS, Pryzbeck TR and Mensforth RP (1985) Chronological metamorphosis of the auricular surface of the ilium: A new method for the determination of adult skeletal age at death. *American Journal of Physical Anthropology*, **68**, 15-28.

McCullagh P (1980) Regression models for ordinal data. *Journal of the Royal Statistical Society*, **42**, 109-142.

McKern T and Stewart TD (1957) *Skeletal Age Changes in Young American Males*. Natick, Mass.: US Army Quartermaster Research and Development Center.

McKelvey RD and Zavoina W (1975) A statistical model for the analysis of ordinal level dependent variables. *Journal of Mathematical Sociology*, **4**, 103-120.

Meindl RS and Lovejoy CO (1985) Ectocranial suture closure: A revised method for the determination of skeletal age at death based on the lateral-anterior sutures. *American Journal of Physical Anthropology*, **68**, 57-66.

Nelson W (1982) *Applied Life Data Analysis.* New York: John Wiley and Sons.

Walker RA and Lovejoy CO (1985) Radiographic changes in the clavicle and proximal femur and their use in the determination of skeletal age at death. *American Journal of Physical Anthropology*, **68**, 67-78.

O'Connor KA (1995) *The Age Pattern of Mortality: A Micro-analysis of Tipu and a Meta-analysis of Twenty-nine Paleodemographic Samples.* Doctoral Dissertation, Department of Anthropology, State University of New York, Albany, N.Y.

Pickles A (1985) *An Introduction to Likelihood Analysis*. Norwich: Geo Books.

Snow CC (1983) Equations for estimating age at death from the pubic symphysis: A modification of the McKern-Stewart method. *Journal of Forensic Sciences*, **28**, 864-870.

Stewart TD (1979) *Essentials of Forensic Anthropology, Especially as Developed in the United States.* Springfield, Ill.: Charles C. Thomas.

Todd TW (1920) Age changes in the pubic bone. I. The male White pubis. *American Journal of Physical Anthropology*, **3**, 285-339.

Walker PL, Johnson JR and Lambert PM (1988) Age and sex biases in the preservation of human skeletal remains. *American Journal of Physical Anthropology*, **76**, 183-188.

Wood JW, Holman DJ, Weiss KM, Buchanan AV and LeFor B (1992) Hazards models for human population biology. *Yearbook of Physical Anthropology*, **35**, 43-87.

Table 1 Number of individuals in the Tipu collection by number of age-at-death indicators (O'Connor 1995).

Number of age indicators	Number of individuals[1]	%
1	166	31
2	152	29
3	149	28
4	45	8
5	15	3
6	5	1

[1]*N* = 532 juveniles and adults. Indicators for adults are given in Table 2; indicators for subadults include tooth development and eruption, epipysial union, diaphyseal length, and tooth wear.

Table 2 Number of adult individuals with particular aging indicators in the Tipu skeletal collection (O'Connor 1995).

Age indicator	**N**
cranial suture closure	143
tooth wear	139
auricular surface	128
cemental annulation	37
pubic symphysis	33
vertebral osteophytosis	33
Total	255

Table 3 The Todd pubic phases (Todd 1920).

Modal Phase	Age Range
I	18-19
II	20-21
III	22-24
IV	25-26
V	27-30
VI	30-35
VII	35-39
VIII	39-44
IX	45-50
X	50+

Table 4 McKern and Stewart sum of three pubic component scores (Snow 1983).

Total Score	Age range	Mean age
0	17	17.2
1-2	17-20	19.04
3	18-21	19.7
4-5	18-23	20.8
6-7	20-24	22.4
8-9	22-28	24.1
10	23-28	26.1
11-13	23-39	29.2
14	29+	35.8
15	36+	41.0

Table 5 Stage-specific transition variables (T_1 to T_5) defined for a 6-phase marker such as a McKern-Stewart pubic symphysis component.

Stage	T_1	T_2	T_3	T_4	T_5
Phase 0	0	0	0	0	0
Phase 1	1	0	0	0	0
Phase 2	1	1	0	0	0
Phase 3	1	1	1	0	0
Phase 4	1	1	1	1	0
Phase 5	1	1	1	1	1

Table 6 Parameter estimates for the reference distribution of 587 individuals found by the latent trait method. The loglikelihood is -608.09.

Parameter	Estimate	SE
σ_z	0.13	0.28
μ_1	3.06	0.01
μ_2	3.21	0.03
μ_3	3.29	0.02
μ_4	3.63	0.03
μ_5	4.38	0.11
σ_1	0.10	0.04
σ_2	0.10	0.04
σ_3	0.11	0.03
σ_4	0.20	0.05
σ_5	0.33	0.09
β_2	-10.7	20.4
β_3	-13.2	26.2
β_4	-10.5	22.1
β_5	-15.0	33.8

Table 7 Parameter values used to simulate the example target sample, and parameter estimates for the target age-at-death distributions for the target sample. Log likelihoods for the known-age model and the latent trait model were -607.77 and -222.61 respectively.

Parameter name	Simulation parameter[a]	known-age estimates[b] (SE)	Latent trait estimates[c] (SE)
α_1	0.01	0.013	0.012
		(0.003)	(0.006)
α_2	0.001	0.00009	0.00007
		(0.00012)	(0.0006)
β_g	0.1	0.11	0.13
		(0.02)	(0.15)

a. Parameter value used to simulate the target sample (see Chapter 10)

b. Parameters recovered by direct estimation of known ages in the simulated target sample.

c. Parameters recovered by the latent-trait method.

Figure Captions

Figure 1 (*Upper panel*) The relationship between latent growth rate variable *z* and the average age at which a child will emerge five deciduous teeth. (*Lower panel*) The distribution of *z* among children in the population.

Figure 2 The effect of $z\beta_z$ on the distribution of times to transition. The distribution of the latent trait *z* is normal with $\mu = 18$ and $\sigma = 4$. (*Upper panel*) Accelerated failure time model in which $z\beta_z$ shifts the distribution rigidly to the left or right. (*Lower panel*) Proportional hazards model, in which $z\beta_z$ changes both the mean and the variance of the distribution.

Figure 3 Age-at-death distribution for the target sample based on estimated parameters in Table 7. The solid line (± one standard error bars) shows the distribution recovered directly from the known ages of the target sample. The dashed line is the target distribution recovered by the latent trait method using pubic symphysis indicators, and dotted lines are ± one standard error.

FIGURE 1

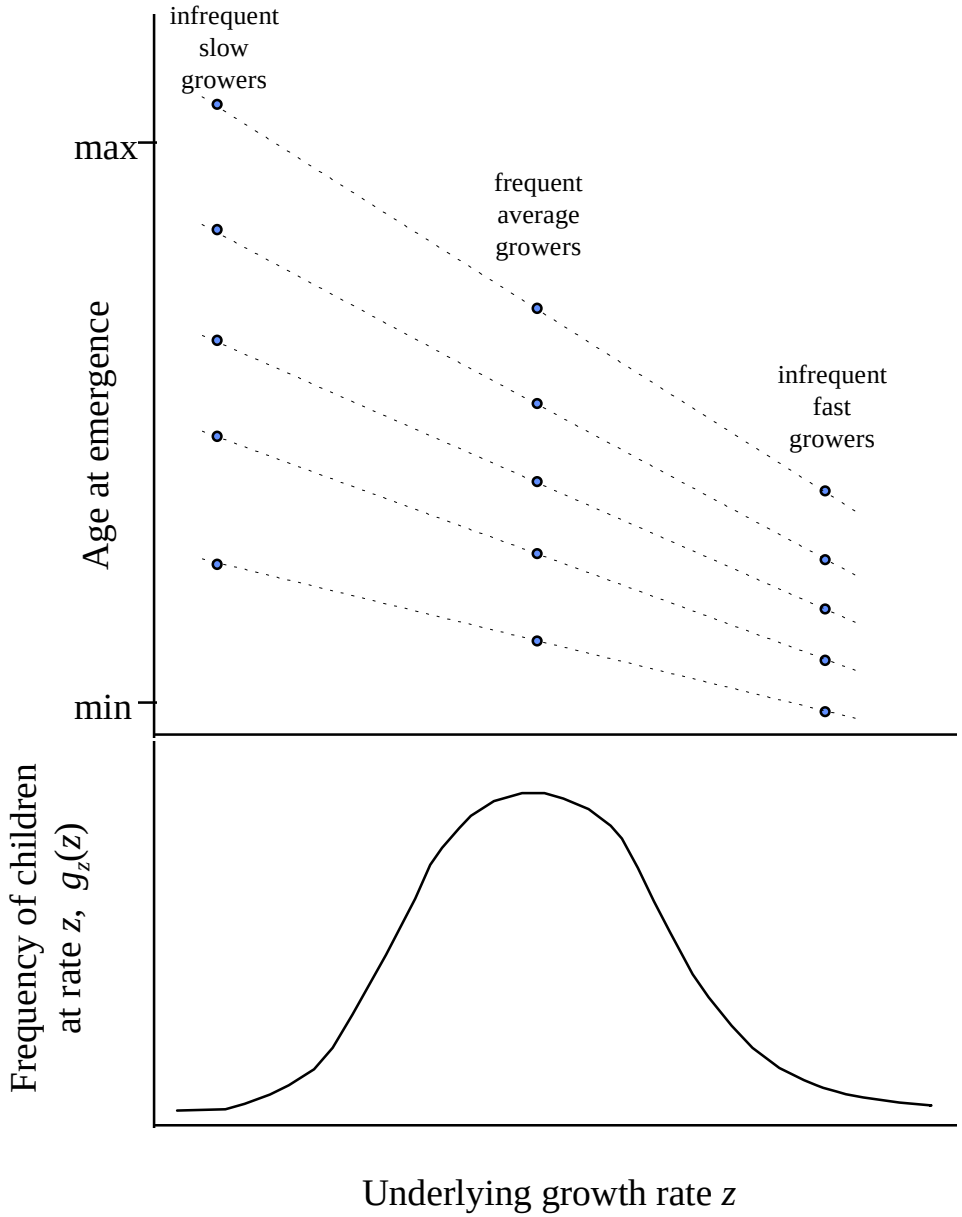
infrequent
slow
growers
max
frequent
average
growers
infrequent
fast
growers
Age at emergence
min
Frequency of children
at rate z, $g_z(z)$
Underlying growth rate z

FIGURE 2

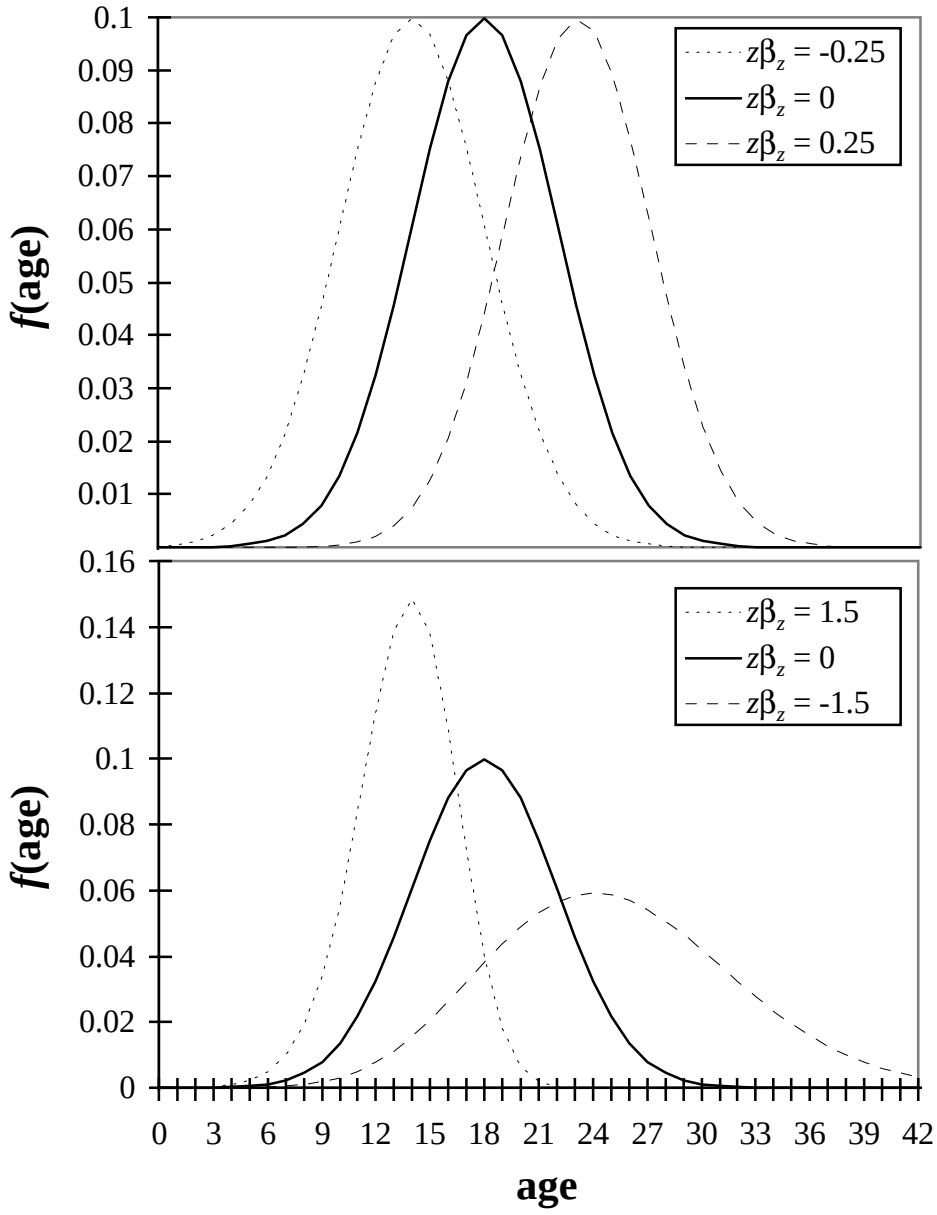
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
f(age)
zβz = -0.25
zβz = 0
zβz = 0.25
0.16
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
f(age)
zβz = 1.5
zβz = 0
zβz = -1.5
0 3 6 9 12 15 18 21 24 27 30 33 36 39 42
age

FIGURE 3