

MORTALITY MODELS FOR PALEODEMOGRAPHY

James W. Wood¹
Darryl J. Holman²
Kathleen A. O'Connor²
Rebecca J. Ferrell¹

¹Department of Anthropology and Population Research Institute
Pennsylvania State University, University Park, PA, USA

²Department of Anthropology and Center for Studies in Demography and Ecology
University of Washington, Seattle, WA, USA

Introduction

Population scientists concerned with long-term trends in human mortality ought to be interested in skeletal samples from extinct communities. Such samples are, in principle, the only possible source of information for most preindustrial populations lacking written records – by far the most common kind of human community that has ever existed. Samples of skeletons provide two broad classes of information of potential interest to demographers and other population specialists: frequency counts of bony lesions that may reveal something about pathological processes active in the population, and data on ages at death from which age patterns of mortality may be inferred. Of these, the latter class of information has generally been deemed to be the less problematic. It has been assumed that skeletal age at death can be estimated well enough, albeit with some inevitable degree of error, to support a few crude but revealing statistics such as mean age at death, life expectancies, and age-specific mortality rates. And so for decades it has been considered perfectly acceptable to use skeletal data to compute *life tables*, the traditional demographic tool for investigating age patterns of mortality. All that is needed, in this view, are a few simple modifications of standard life-table techniques, modifications that were laid down thirty years ago by Acsádi and Nemeskéri (1970:60-65).

Over the years, paleodemographers have computed innumerable life tables, and they continue to do so to this day (for a few examples, see Green et al. 1974; Lovejoy et al. 1977; Greene et al. 1986; Lanphear 1989; Mensforth 1990; Benedictow 1996:36-41; Alesan et al. 1999). But the life-table approach, so long the mainstay of paleodemographic mortality analysis, is open to criticism on several grounds (Sattenspiel and Harpending 1983; Konigsberg and Frankenberg 1992, 1994; Milner et al. 2000). First, paleodemographic studies do not produce the kinds of data needed to compute life-table mortality rates using standard methods – specifically, the numbers of deaths among people at each (known) age and the number of person-years of exposure to the risk of death at that age during some well-defined reference period. Instead, paleodemographers have been forced to work with fuzzily-defined, error-prone distributions of purported ages at death, which can, under restrictive circumstances, be used to generate life tables –

if, that is, one is willing to use methods whose statistical properties are poorly characterized.

Second, the life-table approach assumes that the target population being studied was *stationary* in the technical demographic sense of the term. That is, it assumes that the population was closed to migration and had an intrinsic rate of increase equal to zero, age-specific schedules of fertility and mortality that were unchanging over time, and an equilibrium age distribution induced by those age-specific birth and death rates (Lotka 1922). Only in this special (and not necessarily realistic) case is the empirical age distribution of skeletons expected to have a simple, straightforward relationship to the cohort age-at-death column in the life table. This problem was recognized by one of the earliest advocates of the paleodemographic life table (Angel 1969) and has been discussed in several more recent treatments (see, for example, Moore et al. 1975; Sattenspiel and Harpending 1983; Johannson and Horowitz 1986; Wood et al. 1992b; Konigsberg and Frankenberg 1994).

Third, the use of fixed age intervals in the life table implies that the ages of all skeletons are known within the same margin of error, including those of fragmentary skeletons that exhibit only a few, unreliable indicators of age. Thus, the life-table approach is unacceptably procrustean: it tries to force the complicated error structure of paleodemographic age estimates into a rigid framework of a few discrete age intervals.

Fourth, and perhaps most seriously, the life table is a wasteful way to use the small samples typical of paleodemographic studies – samples that are often on the order of a few dozen or, at best, a few hundred skeletons. In computing a life table we need to estimate one parameter (an age-specific mortality rate) for each and every age interval in the table, often requiring ten or more separate parameters to be estimated. Few paleodemographic samples will support a method with such a gargantuan appetite for data.

For the past three decades, paleodemographers have attempted to circumvent some of these problems by using so-called *model* life tables (UN 1955, 1956; Coale and Demeny 1966; Weiss 1973). In this approach, the investigator searches through published tabulations of theoretical age-specific mortality patterns to find an age-at-death distribution that appears to mimic the empirical distribution being studied. In theory, this

approach allows the assumption of stationarity to be relaxed (Paine 1989). In practice, however, the methods for fitting model life tables have been *ad hoc* and informal, and the results are only good if the published tabulations happen to include a table that corresponds closely to the population under study – something that is inherently untestable.

The Rostock protocol outlined by Love and Müller (present volume) – and the earlier work of Königsberg and Frankenberg (1992), which anticipates it (see also Königsberg et al. 1997) – represents a major advance in our thinking about how to estimate mortality statistics from skeletal samples. Under the Rostock protocol, we never compute a life table – although, as we show below, we can eventually compute something that looks like a life table if we so desire. Indeed, we do not begin by classifying skeletons by age at all, as we would have to do in the life-table approach. Instead, we directly estimate the age pattern of death from the total sample of skeletons unclassified by age. Using c to indicate a vector of observed skeletal traits that provide information about age at death, the probability of observing a particular c value – say, c_i – out of the sample as a whole is the marginal density of c_i :

$$\Pr(c_i) = \int_0^{\infty} \Pr(c_i | a) \Pr(a) da . \quad (1)$$

Since $\Pr(c_i)$ is the *likelihood* of observing a skeleton with characteristics c_i in our sample, the likelihood function for the entire sample of n skeletons is

$$L = \prod_{i=1}^n \Pr(c_i) = \prod_{i=1}^n \int_0^{\infty} \Pr^*(c_i | a) \Pr(a) da , \quad (2)$$

where the asterisk (*) denotes an empirical estimate from a reference sample of skeletons whose ages at death are known (see Usher, this volume). The function $\Pr(a)$ is the age-at-death distribution in the target sample whose mortality pattern we wish to estimate. It is $\Pr(a)$ that tells us what we want to learn about mortality in the past. And maximization of Equation (2) provides the basis for maximum likelihood estimates of the $\Pr(a)$ function from the target sample.

If the Rostock protocol is to be used in paleodemographic research, we need to find a suitable *parametric model* for the age-at-death distribution $\text{Pr}(a)$. In other words, we need to boil all the complexities of age-specific mortality down to a single, more-or-less simple set of equations – equations containing constants (known as *parameters*) whose values we hope to estimate from skeletal data. Although some paleodemographers might balk at the notion of reducing all the manifold variability in human mortality to naked math, the parametric approach actually has a number of virtues for paleodemographic analysis. As we show below, it allows us to correct for the confounding effects of non-stationarity – population growth or decline – on the age-at-death distribution. It also permits us to compare mortality patterns across populations in a straightforward way by examining parameter estimates and their associated standard errors. And if we construct our parametric model wisely, it may even reveal something interesting about the biological processes underlying the human mortality curve.

The parametric approach does, however, have one profound limitation: it is only as good as the model chosen for the age-at-death distribution. In this paper, we review parametric models of human mortality with an eye toward identifying models that may be of use in paleodemographic estimation. A secondary (but important) goal is to find models that facilitate *etiologic* ways of thinking about paleodemographic mortality profiles – that is, models that allow for some kind of meaningful biological interpretation and insight. We examine the etiologic foundations of current models and develop extensions that provide insights into the mortality processes experienced by past populations. Finally, we discuss some important issues, including heterogeneity in the risk of death, non-stationarity, and the sex differential in mortality, that must be considered in reconstructing the demographic past.

Before we go into the details of the alternative model specifications, it is worth asking what we are trying to accomplish in paleodemographic mortality analysis. We also need to be honest about what we can *never* accomplish, even with the best skeletal samples imaginable. Mainstream demographers often have the luxuries of huge samples, known ages, and information about specific causes of death (both primary and contributory). They can justify using some very complicated models that at once require such data and take advantage of them (see, for example, Schoen 1975; Manton and

Stallard 1988; Nam 1990). As a result, they can examine the fine details of human mortality with comparative ease. Paleodemographers do not have – and never will have – any of these luxuries. Paleodemographic samples will almost always be small and subject to a number of unavoidable taphonomic biases (Gordon and Buikstra 1981; Waldron 1987; Walker et al. 1988; Mays 1992). It is unreasonable, therefore, to expect that paleodemographers will ever be able to reconstruct the fine details of any set of mortality rates. At best, we can hope to learn something about the overall level and age pattern of death in the distant past – and perhaps something about the gross differences in material conditions that led to variation in level and age pattern. This fact places a limit on the kinds of models worthy of consideration by paleodemographers. In general, simple models that reveal overall patterns are to be preferred over complicated models that purport to tell us about the detailed squiggles and bumps of the age-at-death curve. It is on such simple models that we concentrate in this paper.

What Exactly Do We Need to Model?

To implement the Rostock approach, we need to model $\Pr(a)$, the *age-at-death distribution* of the past population under study. But what exactly is this distribution? And what is its relationship to the underlying age pattern of mortality? Intuitively, it might seem as if the relationship has to be simple. In fact it is complicated, and we need to be clear about it if we are to avoid going wrong.

For simplicity, imagine that we observe all the deaths that occurred in a well-defined population during some specified period of time, and that we know the exact age at which each and every death took place. (Needless to say, we never have it so good in paleodemography; but for the moment we are interested in theory, not reality.) How can we best characterize the age-specific mortality pattern of our ideal population in a formal statistical sense? And how can we model that pattern mathematically? Conceptually, if not computationally, it is simplest to begin with $\mu(a)$, the *age-specific mortality rate* at exact age a (normally measured in years). If we treat age as a continuously varying quantity – and throughout this paper we will – then $\mu(a)$ is called the *force of mortality* (Keyfitz 1968:5) and is defined as

$$\mu(a) = \lim_{\Delta a \rightarrow 0} \left(\frac{\text{number of deaths at age } [a, a + \Delta a]}{\text{person} - \text{years of exposure at age } [a, a + \Delta a]} \right). \quad (3)$$

This function defines a rate that is strictly non-negative. It can be thought of as the continuous-time analogue of the central mortality rate, the usual starting point for calculation of the life table. But we cannot compute $\mu(a)$ directly from a paleodemographic age-at-death distribution (even if we know that distribution perfectly), so it behooves us to define some related functions. One of these is the *survival function*, $S(a)$, derived from the age-specific mortality function as

$$S(a) = e^{-\int_0^a \mu(x) dx}. \quad (4)$$

$S(a)$ is the probability that an individual survives from birth to *at least* age a . Since a cannot take on negative values, it follows that $S(0) = 1$. In addition, $S(a)$ is monotonically non-increasing with a , i.e. it can only go down (or remain the same) as age increases. As $a \rightarrow \infty$, $S(a)$ approaches zero. Thus, $S(a)$ is analogous to the survivorship column in the life table in all its particulars, save that age is reckoned continuously rather than in discrete intervals.

We now inch our way toward something that starts to look like the paleodemographic age-at-death distribution $\text{Pr}(a)$ – but, in most circumstances, is not equivalent to it. This is the probability density function (PDF) of ages at death in a *birth cohort* of individuals subjected to the mortality function $\mu(a)$ at each age. We will write this PDF as $f_0(a)$. (The reason for the zero subscript will become clear presently.) It can be derived from $S(a)$ as

$$f_0(a) = -dS(a) / da. \quad (5)$$

If we were dealing with skeletons from a single cohort, $f_0(a)$ would indeed be equivalent to $\text{Pr}(a)$. But such is never the case in paleodemography – and if, by some miracle, it were the case, we would never know it.

According to some basic results from renewal theory (Cox 1962), the hazard, density, and survival functions are related to each other in the following ways:

$$\mu(a) = -\frac{d \ln S(a)}{da} = -\frac{1}{S(a)} \frac{dS(a)}{da} = \frac{f_0(a)}{S(a)}, \quad (6)$$

$$S(a) = \int_a^{\infty} f_0(x) dx, \quad (7)$$

$$f_0(a) = \frac{\mu(a)S(a)}{\int_0^{\infty} \mu(x)S(x) dx}. \quad (8)$$

The denominator in Equation (8) rescales $f_0(a)$ so that it behaves like a proper PDF and integrates to one. These relationships will be useful at several points in the following discussion. Because of these mathematical relationships, once we know one of these three functions, we can immediately determine the other two.

It is important to emphasize the parallels that exist between $\mu(a)$, $S(a)$, and $f_0(a)$, on the one hand, and certain columns in the classic life table on the other. We have already mentioned that $\mu(a)$ is analogous to the life-table central mortality rate, and $S(a)$ to the survivorship schedule. Similarly, $f_0(a)$ is analogous to the life-table (cohort) distribution of ages at death. Other “life-table-like” functions can be derived from $\mu(a)$, $S(a)$, or $f_0(a)$. For example, $\mu(a)$ can be converted into an age-specific *probability* of death, $q(a)$, during some small subinterval $[a - \frac{1}{2}\Delta a, a + \frac{1}{2}\Delta a]$ around a by solving

$$q(a) = 1 - \exp\left[-\int_{a-\frac{1}{2}\Delta a}^{a+\frac{1}{2}\Delta a} \mu(x) dx\right]. \quad (9)$$

For $\mu(a)$ in the interval $[0, \infty)$, this expression constrains $q(a)$ to fall between zero and one. Another quantity related to $\mu(a)$, $f_0(a)$, and $S(a)$ is the *life expectancy* or expected remaining life time for an individual alive at age a ,

$$e(a) = \frac{\int_a^{\infty} S(x)dx}{S(a)}. \quad (10)$$

As these last two equations show, $\mu(a)$, $f_0(a)$, or $S(a)$ can be used to derive all the information we would normally hope to learn from an old-fashioned life table without ever requiring us to compute one. Or, rather, they *would* if only we could estimate them.

Which brings us back to the age-at-death distribution $\text{Pr}(a)$ – the nearest we can get theoretically to paleodemographic data on skeletal age at death. We have hinted that there is a close (if complicated) relationship between $\text{Pr}(a)$ and $f_0(a)$, and it is now time to make that relationship explicit.

As already noted, $f_0(a)$ is the age-at-death distribution of a single birth cohort exposed to the mortality function $\mu(a)$. As it happens, it is also the expected age-at-death distribution for all the deaths occurring in a stationary population over some delimitable period of time – for example, the time span during which skeletons are deposited in a cemetery (see Appendix). If we were sure that the population was stationary during the entire period of deposition, we could substitute Equation (8) into our likelihood function (Equation 2) and – once we have specified a parametric model for $\mu(a)$ and $S(a)$ – maximize it to obtain parameter estimates. But what if our target population was *not* stationary? What, for example, if it was changing in size, no matter how slowly? Then $f_0(a)$ is not the same as $\text{Pr}(a)$, and we cannot use Equation (8) in our likelihood. What do we do?

Even if we cannot take it for granted that our target population was stationary, it may still be reasonable to assume that it was *stable*. In other words, we may be able to make all the assumptions listed above for the stationary population, except allowing for the possibility of a non-zero growth rate. (Note, by this logic, that the stationary population is simply a special case of the more general stable population.) As decades' worth of demographic analysis has shown, the assumption of stability is much less restrictive than the assumption of stationarity; even when fertility and mortality rates are changing and migration is occurring, most human populations still closely approximate a stable age distribution at any given time (Keyfitz 1968:89-94; Parlett 1970; Bourgeois-

Pichat 1971; Coale 1972:117-61). This property, known as *weak ergodicity* (Lopez 1961:66-68), ensures that stable population models almost always fit well, unless the populations to which they are being fit have been subjected to unusually rapid, cataclysmic change.

In a stable but non-stationary population, the age-at-death distribution is only partly a function of age-specific mortality; it is also influenced by the number of living individuals at risk of death at each age, which is influenced in turn by population growth. More precisely, the number of deaths at age a is proportional to the product of the force of mortality, $\mu(a)$, and the fraction of the total population that is age a , conventionally labeled $c(a)$. In a stationary population, $c(a)$ is proportional to $S(a)$, the probability of surviving from birth to age a , which makes the age-at-death distribution a reflection of mortality alone – but *only* in that special case. In a stable population with a non-zero growth rate equal to r , the value of $c(a)$ is proportional to $S(a)e^{-ra}$. The quantity e^{-ra} corrects for the fact that the absolute number of newborns entering the population each year is changing as a result of population growth, thus distorting the age distribution that would have been expected under conditions of stationarity. For a positive growth rate, for example, there are more individuals born this year than, say, ten years ago: if B babies are born this year into a stable population, then $B \times e^{-10r}$ babies must have been born ten years ago.

This change in the number of individuals entering the population at $a = 0$ means that the number of people dying at each subsequent age must be a function not only of the force of mortality, but of the growth rate as well. The number of people surviving to each age is proportional to $S(a)e^{-ra}$; those survivors are then exposed to the age-specific mortality rate $\mu(a)$. Thus, the probability density function for deaths in a stable population with growth rate r is

$$f_r(a) = \frac{\mu(a)S(a)e^{-ra}}{\int_0^{\infty} \mu(x)S(x)e^{-rx} dx} = \frac{f_0(a)e^{-ra}}{\int_0^{\infty} f_0(x)e^{-rx} dx}. \quad (11)$$

(Compare Equation 8. It should now be clear why we mark $f_0(a)$ with a subscript zero: it represents $\text{Pr}(a)$ only if the population's growth rate is zero – or in the profoundly

unlikely event that we are dealing with a single cohort.) As shown in the Appendix, this same expression applies to all the skeletons accumulated by a stable population over some more or less protracted span of time. In principle, then, we can treat $f_r(a)$ as the $\text{Pr}(a)$ function in our likelihood (Equation 2) and estimate r as an additional parameter of the model – if we can assume that the population was stable. And if it was *not* stable, at least approximately, we have probably reached the outer limits of what we can ever hope to learn about age-specific mortality from skeletal samples.

This correction for non-stationarity still requires us to specify a parametric model for the age pattern of mortality. In other words, we still need to write down an equation for either $\mu(a)$, $f_0(a)$, or $S(a)$. And we should try hard to choose an equation that is flexible enough to approximate all *known* human mortality distributions in order to be reasonably confident that the model will accommodate the *unknown* mortality distribution we are trying to reconstruct. At the same time, the model must be sufficiently bounded that growth rates are uniquely identifiable since identifiability of the growth rate is not guaranteed for some possible parametric models (Holman et al. 1997, 1998). So we need the simplest possible model that is still complicated enough to capture most of what we know about human mortality patterns. Which immediately raises the question: what *do* we know about human mortality patterns, including their common features and their range of variation?

What Does the Human Mortality Curve Look Like?

Mortality trends and patterns have been well characterized for many contemporary human populations and some historical ones (mostly European, mostly confined to the past four centuries) (Coale and Demeny 1966; Keyfitz and Flieger 1968, 1990; Preston 1976; Gage 1990). Much less is known about mortality conditions among the types of populations typically studied by anthropologists: the small foraging or horticultural societies characteristic of most of human existence. Nonetheless, work to date suggests that the mortality profiles of these populations tend to conform to a generalized human pattern, although often at a level of mortality near the upper end of the range typically observed in national and historic populations (Weiss 1973; Gage 1988). It thus seems meaningful to talk about the “common” age pattern of human mortality.

The basic pattern of the age-specific force of mortality is, in some respects, strikingly similar across a wide range of human populations, whether characterized by high mortality or low (Figure 1). The general pattern appears to be one of excess mortality at the youngest ages of the life span, with a rapid, monotonic decline to a lifetime low at around 10-15 years of age. This low point is followed by an accelerating rise in mortality at later ages, a rise that appears to be roughly exponential. Because this age pattern of mortality looks rather like the cross-section of an old-fashioned clawfoot bathtub, it is sometimes referred to as the *bathtub curve*. Figure 2 shows the survival function and the cohort PDF associated with the bathtub curve.

The principal variations on this common theme that are observed in historical and modern populations include wide differentials in the excess mortality occurring at the youngest and oldest ages and, in some populations, marked differences in the timing of the decline in juvenile mortality or the rise in adult mortality (Coale and Demeny 1966; Keyfitz and Flieger 1968, 1990; Preston 1976). All these phenomena are illustrated in Figure 1. These are, we suggest, the minimal kinds of variation we should expect our model to be able to capture.

Types of variation in the age pattern of human mortality that are *less* commonly observed – perhaps because they are of much smaller magnitude and thus require uncommonly good data to show through – include the so-called “accident hump” at late juvenile and early adult ages and an apparent slowing down of the rate of increase of mortality among the oldest of the old. The accident hump, as Gage and Mode (1993) have noted, is most clearly observed in males from European-derived populations with low mortality (most notably the U.S., Canada, and Australia). Luder (1993) has suggested that it also occurs in non-human primates, although inadequate data make this claim difficult to evaluate. Even if the accident hump is a widespread phenomenon in human populations, the actual magnitude of the mortality rise associated with it appears to be miniscule, a point rightly emphasized by Gage and Mode (1993).

The deceleration of mortality among the oldest old is sometimes observed in populations for which exceptionally good data on the elderly are available (Horiuchi and Coale 1990; Kannisto 1994; Thatcher et al. 1997; Vaupel 1997; Vaupel et al. 1998). One possible explanation for this deceleration of mortality at the oldest ages is selective

mortality, which might be expected to eliminate all but the least vulnerable individuals by the time the oldest segments of the life span are reached (Vaupel et al. 1979; Brooks et al. 1994; Himes 1994). Recent work on other organisms also highlights the possibility that the deceleration in mortality is real at the individual level, and not just an artifact of selectivity (Carey et al. 1992; Fukui et al. 1993; Vaupel et al. 1994). From a paleodemographic perspective, these issues seem moot for the simple reason that the deceleration of mortality, whatever its cause, is only observed at ages so advanced (after, say, 90 years of age) that it cannot have been an important feature of mortality in any preindustrial population.

In our opinion, then, the accident hump and the senescent deceleration in mortality exemplify just the sorts of “bumps and squiggles” in the mortality curve that paleodemographers will never be able to resurrect with any credibility. It would seem sufficiently challenging to try to reconstruct the general shape and level of the bathtub curve.

Ways of Modeling Mortality

As the previous section suggests, the mortality patterns of human populations can all be regarded as variations on a common, species-wide theme – where both the variations and the commonalities are of interest. The challenge in modeling mortality consists in capturing the underlying “universal” age structure of death while allowing for at least the principal kinds of variation in its detailed realization observed in the real world. Past attempts to model mortality can be classified in several different ways; one way that is especially telling in the present context is to subdivide them into *semi-parametric* (or perhaps *semi-empirical*) and *fully parametric* forms. Semi-parametric models start with empirically observed mortality schedules and generalize them, usually by subjecting them to some form of regression analysis. For example, the pioneering work on model life tables, published by the United Nations (1955, 1956), involved regressing estimates of the infant mortality rate on the rest of the age schedule of mortality across 24 different populations. No attempt was made, beyond the regression model itself, to reduce all the empirical complexities to a simple mathematical form. But reduction to a simple mathematical form is precisely what the fully parametric approach

seeks to do. In this approach, empirical data are examined rather informally to get a sense of what the age pattern of mortality ought to look like, and then an equation is found that mimics that pattern to some acceptable degree of approximation.

It might be thought that the semi-parametric approach is always preferable because, from its very outset, it hews more closely to real data. But, as we detail in the rest of this paper, this is far from being the case. Particularly when parametric models are simple and allow some etiologic interpretation, they can be much more enlightening about real-world processes affecting mortality.

In the following sections, we discuss one semi-parametric model and several fully parametric ones. The semi-parametric model we have chosen is one of several known as *relational models* (Zaba 1979, 1981; Heligman and Pollard 1980; Ewbank et al. 1983; Aalen 1989), so-called because they are all based on statistical relationships among empirical mortality patterns. Relational models, in some respects, represent a compromise between traditional life tables and fully parametric models – hence our description of them as “semi-parametric”.

Relational Models

The development of relational models was originally inspired by a quest to find the minimal number of parameters needed to capture all the variation in the level and shape of the human curve of age-specific mortality. A preliminary solution to this problem was provided by Ledermann and Breas (1959) who performed a factor analysis of estimated age-specific mortality rates from a large number of populations, showing that two latent factors (apart from sex) accounted for more than half of the observed variation in mortality. This result inspired Brass (1971) to develop a two-parameter model of mortality, one that underlies what has come to be called the Brass (or logit) approach to mortality estimation. The Brass model is the prototype for all later relational models (e.g. Zaba 1979; Ewbank et al. 1983), and it can be used to exemplify the approach as a whole.

The logic of the Brass system starts with the theoretical survival function $S(a)$. Imagine for the moment that two populations (denoted by the subscripts 1 and 2) differ

only in the level of mortality, so that $\mu_1(a) = \kappa\mu_2(a)$ for all a , where κ is a constant. From Equation (6) it follows that

$$\frac{1}{S_1(a)} \frac{dS_1(a)}{da} = \frac{\kappa}{S_2(a)} \frac{dS_2(a)}{da}. \quad (12)$$

By inspecting a large number of empirical mortality schedules, Brass discovered that the scalar κ relating different schedules is not in fact a constant, but declines toward unity with advancing age. For example, in one extreme comparison $\mu(a)$ was more than 16 times higher in one population than in another in the age interval 1 to 4 years, but dropped to about 1.5 times higher at ages 75 to 79 (Brass 1971). Brass found that this pattern could be closely approximated by a function of the form

$$\frac{d S_1(a)/da}{S_1(a)[1 - S_1(a)]} = \frac{\kappa dS_2(a)/da}{S_2(a)[1 - S_2(a)]}. \quad (13)$$

Solving for $S_i(a)$,

$$\ln \left[\frac{1 - S_1(a)}{S_1(a)} \right] = \alpha + \beta \ln \left[\frac{1 - S_2(a)}{S_2(a)} \right], \quad (14)$$

where α and β are new constants.

If the number x lies between 0 and 1, then $\ln [x/(1 - x)]$ is known as the *logit transform* of x , often written $\text{logit}(x)$. Thus, we can rewrite Equation (14) as $\text{logit} [1 - S_1(a)] = \alpha + \beta \text{logit} [1 - S_2(a)]$. This equation is the basis of the Brass relational model, and α and β are its two parameters. Roughly speaking, a choice of α sets the overall level of mortality (as reflected in, say, the life expectancy at birth) while β sets the “tilt” of mortality curve 1 compared to curve 2.

Now suppose that “population 2” is a well-studied reference population whose survival schedule has been estimated properly from high quality data, and “population 1”

is some target population whose survival schedule is only poorly known. Then a linear regression of logit $[1 - S(a)]$ from the target population on that of the reference population, in the form of Equation (14), can be used to smooth the target population's mortality curve and fill in any gaps (see Brass 1975 for technical details). In this way, information on part of the target population's survival schedule can be used to generate the entire schedule.

In his original paper on the subject, Brass (1971) showed that the logit approach is reasonably flexible and provides plausible results when applied to data from a wide variety of national populations. In the same paper, Brass provided a reference life table that has proven useful in analyses of mortality data from Africa and Asia (see Brass and Coale 1968; Carrier and Hobcraft 1971). It is important to emphasize, however, that neither the logit approach in general nor the Brass reference table in particular has been able to cover all known human mortality patterns, and both may be especially bad for the small, high-mortality populations commonly studied by anthropologists (Wood 1987a). In addition, the form of Equation (13), and hence (14), was not derived from theoretical considerations, but is purely empirical. Nonetheless, relational models provide a simple system for mortality estimation that is flexible enough to warrant more attention by paleodemographers than they have hitherto received.

Fully Parametric Models

An alternative to model life tables and relational models are parametric models of the age pattern of mortality (Wood et al. 1992a). If constructed properly, these models reduce the numerous life-table age classes into a small number of biologically meaningful parameters that can all be estimated from data on the target population being studied (Gage and Dyke 1986; Gage 1989; Gavrilov and Gavrilova 1991). Parametric models have only begun to be widely applied in demographic research in the last two decades as advances in computer technology have facilitated the development, testing, and application of complex statistical models (Mode and Busby 1982; Mode and Jacobsen 1984; Gage 1988, 1989; Wood et al. 1992a; Gage and Mode 1993). These models are extremely promising for use with small paleodemographic samples because of their

parsimony in describing mortality patterns with the smallest possible number of parameters.

Like relational models, fully parametric models of mortality can be used to smooth and correct inadequate mortality data. But they can be much more flexible than relational models. All mortality models impose a certain amount of *a priori* age structure onto the data being examined, but good parametric models make the fewest assumptions about what the detailed age pattern of mortality ought to be. In theory, this permits us to come closer to the “true” underlying age structure of mortality in the population being studied – assuming that we have selected the *right* parametric model.

A number of parametric models of the age patterns of mortality have been developed over the years, as attempts have been made to formulate a general “law of mortality” applicable to all human populations (for reviews, see Mode 1985:35-74; Gage 1989; Gavrilov and Gavrilova 1991; Wood et al. 1992a). In the following sections, we discuss models that we consider especially promising for paleodemography. Since paleodemographic cause-of-death analysis is (and will probably remain) poorly developed, all the models we consider deal with mortality from all causes simultaneously.

Weibull, Rayleigh, and Bi-Weibull Models

The two-parameter *Weibull model* (Weibull 1951) is widely used in industrial reliability testing, mainly to model the effects of accumulated damage on product breakage (Thompson 1988). By analogy, it may provide a reasonable model for human aging, which is a kind of “wear-out” process. The force of mortality in the standard Weibull model is

$$\mu(a) = \beta a^{\beta-1} / \eta^\beta . \tag{15}$$

The associated survival function is

$$S(a) = \exp[-(a/\eta)^\beta] , \tag{16}$$

and the cohort PDF of ages at death is

$$f_0(a) = (\beta a^{\beta-1} / \eta^\beta) \exp[-(a/\eta)^\beta]. \quad (17)$$

As Nordling (1953) first noted, the two-parameter Weibull specification can be used to model so-call *multi-hit* or *multi-stage* processes, in which a fixed number of insults or disease stages must be experienced before death ensues. Examples for which such models may be relevant include cancer, in which two or more somatic mutations must occur before a cell line becomes malignant and metastatic; diabetic nephropathy, which is preceded by a fairly regular sequence of diabetic stages; and the formation of arterial plaques, for which multiple, sequential lesions in the arterial wall appear to provide a starting point (Whittemore and Keller 1978; Andersen 1988; Weiss and Chakraborty 1990).

Recently, a special case of the Weibull has been used for paleodemographic mortality analysis (Konigsberg and Herrmann 2000). This is the *Rayleigh model*, which is obtained from the Weibull by setting $\beta = 2$. By fixing one parameter, this model gains some efficiency in estimation, albeit at the cost of a corresponding loss in generality and flexibility.

A related model that has found some application in reliability testing is the *bi-Weibull model* (Evans et al. 2000:199-200). Reliability specialists have used this model to capture complex processes with both “burn-in” and “wear-out” stages, roughly paralleling the maturation and senescent phases of the human life span. The bi-Weibull is formed by adding together two Weibull mortality functions: the first a two-parameter Weibull that applies to all ages and the second a three-parameter Weibull that is added to the baseline hazard after $a = \gamma$, the earliest age at which wear-out affects the risk of death. To specify the force of mortality in the bi-Weibull, we need two separate equations:

$$\mu(a) = \lambda\theta(\lambda a)^{\theta-1}, \quad 0 \leq a < \gamma, \quad (18)$$

and

$$\mu(a) = \lambda\theta(\lambda a)^{\theta-1} + \left(\frac{\beta}{\eta}\right)\left(\frac{a-\gamma}{\eta}\right)^{\beta-1}, \quad a \geq \gamma. \quad (19)$$

The corresponding survival function is given by

$$S(a) = e^{-(\lambda a)^\theta}, \quad 0 \leq a < \gamma, \quad (20)$$

and

$$S(a) = \exp\left(-\left\{(\lambda a)^\theta + [(a-\gamma)/\eta]^\beta\right\}\right), \quad a \geq \gamma. \quad (21)$$

The bi-Weibull model does a quite decent job of mimicking the bathtub curve of human mortality (Figure 3). So far as we know, however, it has never been used in paleodemography – or any other branch of demography that we are aware of. If we were willing to rely on evolutionary theory that suggests that senescent causes of death do not begin to be important until about the time of sexual maturation (Hamilton 1966), we could reduce the standard bi-Weibull specification to a four-parameter model by setting γ equal to, say, fifteen years of age. One unfortunate feature of the bi-Weibull, incidentally, is that its force of mortality may be undefined at age zero if $\theta < 1$ (because it involves division by zero), making it impossible to estimate neonatal mortality.

The Gompertz Model

The very first attempt to develop a parametric model of mortality was that of Gompertz (1825). Gompertz modeled the aging or senescent component of mortality with two parameters: a positive scale parameter α that sets the overall level of adult mortality, and a positive shape parameter β that determines how the risk of death accelerates with advancing age. The force of mortality in the Gompertz model is

$$\mu(a) = \alpha e^{\beta a}. \quad (22)$$

The corresponding cohort PDF is

$$f_0(a) = \alpha \exp\left[\beta a + \frac{\alpha}{\beta}(1 - e^{\beta a})\right], \quad (23)$$

and the survival function is

$$S(a) = \exp\left[\frac{\alpha}{\beta}(1 - e^{\beta a})\right]. \quad (24)$$

Gompertz, who was concerned exclusively with mortality associated with aging across the adult life span, assumed that the observed increase in adult mortality with age is a result of a negative exponential decline in physiological capacity (Gage 1989). A variety of other parametric models of aging have since been developed, some of them based on different assumptions about the aging process (e.g. linear rather than exponential decline in physiological capacity with age, or models of accumulated damage with age). Most of these ultimately reduce to or approximate the Gompertz equation (Wood et al. 1994; for reviews of these models see Mode 1985; Gage 1989; Gavrilov and Gavrilova 1991).

The Gompertz-Makeham Model

The earliest modification to the Gompertz model, proposed by Makeham (1860), involves adding a single parameter to capture age-independent adult mortality. This parameter represents mortality resulting from causes, such as accidents or sexually transmitted diseases, unrelated to either maturation or senescence. The Gompertz-Makeham model specifies the force of mortality as

$$\mu(a) = \alpha_1 + \alpha_2 e^{\beta a}. \quad (25)$$

The α_1 parameter in this expression represents the constant, age-independent component of mortality; the $\alpha_2 \exp(\beta a)$ term is just a Gompertz function describing the senescent component.

The cohort PDF for the Gompertz-Makeham model is

$$f_0(a) = (\alpha_1 + \alpha_2 e^{\beta a}) \exp\left[-\alpha_1 a + \frac{\alpha_2}{\beta} (1 - e^{\beta a})\right], \quad (26)$$

and the Gompertz-Makeham survival function is

$$S(a) = \exp\left[-\alpha_1 a + \frac{\alpha_2}{\beta} (1 - e^{\beta a})\right]. \quad (27)$$

The Gompertz-Makeham model fits well to empirical mortality distributions between the ages of 30 and 85 years (Finch 1990). Nearly all subsequent models of the age pattern of mortality have been extensions of the Gompertz-Makeham model, primarily intended to cover the rest of the life span – for example, by allowing for an early-adult accident hump (Thiele 1871; Heligman and Pollard 1980; Mode and Busby 1982; Mode and Jacobsen 1984; Gage 1989; Gavrilov and Gavrilova 1991). As we have already suggested, it is probably pointless for paleodemographers to concern themselves with a detail as small as the accident hump.

The Siler Model

One of the most parsimonious parametric models of mortality across the *entire* life span, including pre-adult ages, is the Siler competing hazards model (Siler 1979, 1983). This model fits as well as or better than most other models to human mortality data (Gage and Dyke 1986; Gage and Mode 1993). Siler added a third component to the Gompertz-Makeham model to represent the earliest segment of life, when the risk of death often starts out high but then declines rapidly. The force of mortality in Siler's model is

$$\mu(a) = \alpha_1 e^{-\beta_1 a} + \alpha_2 + \alpha_3 e^{\beta_3 a}. \quad (28)$$

Note that the parameters of Equation (25) have been renumbered here. Now α_1 is the level of neonatal mortality and β_1 is the rate of decline in early mortality with age. The second term is the constant (Makeham) component of the model, and the third term the senescent (Gompertz) component. The structure of the Siler model invites a simple interpretation of mortality as the sum of three components:

$$\mu(a) = \mu_1(a) + \mu_2 + \mu_3(a), \quad (29)$$

where each μ represents a distinct set of competing causes of death. Indeed, Siler (1979) called his model a competing hazards model precisely because he interpreted its three components as sets of risks that compete simultaneously throughout life. Because of the β parameters, however, the first component is unimportant after the earliest juvenile years, and the third component does not become dominant until adulthood.

The cohort PDF and survival function of the Siler model are

$$f_0(a) = (\alpha_1 e^{-\beta_1 a} + \alpha_2 + \alpha_3 e^{\beta_3 a}) \exp \left[-\frac{\alpha_1}{\beta_1} (1 - e^{-\beta_1 a}) - \alpha_2 a + \frac{\alpha_3}{\beta_3} (1 - e^{\beta_3 a}) \right] \quad (30)$$

and

$$S(a) = \exp \left[-\frac{\alpha_1}{\beta_1} (1 - e^{-\beta_1 a}) - \alpha_2 a + \frac{\alpha_3}{\beta_3} (1 - e^{\beta_3 a}) \right]. \quad (31)$$

Figure 4 shows an example of the Siler model with parameters chosen to reflect a typical human mortality pattern. Despite the fact that the Siler model does not include an accident hump, it still fits reasonably well to human populations, including those that do have this feature (Gage and Mode 1993).

The three components of the Siler model – immature, age-independent, senescent – are assumed to be competing but non-interacting causes of death (or, somewhat more realistically, clusters of distinct causes of death). That is, individuals who survive one set of potential causes (for example, age-independent ones) are just as susceptible as all other individuals to other causes (say, senescent ones).

Although the Siler model was not originally developed with much detailed etiology in mind (especially with regard to its immature component), Gage (1991) has shown empirically that the model has considerable etiologic coherence. For example, mortality attributable to infectious diseases such as pneumonia and diarrhea is highly correlated with the immature and senescent components; degenerative causes of death are primarily associated with the senescent component; and accidents and maternal mortality fall largely into the age-independent component of the Siler model (Gage 1991). Maternal mortality is not, of course, truly age-independent; it simply is not associated with either immaturity or advanced age.

Gage, who pioneered the application of parametric mortality models in anthropological demography (Gage and Dyke 1986; Gage 1988, 1989), has used the Siler competing hazards model extensively for both empirical and theoretical investigations of the age patterns of mortality. He has examined international variation in human mortality (Gage 1990), the relationship of covariates to this variation (Gage 1994; Gage and O'Connor 1994), the age pattern of mortality in anthropological populations (Gage 1988, 1989), and even the age pattern of mortality in non-human primates (Gage and Dyke 1988; Dyke et al. 1993; Gage 1998). He has also examined hypotheses regarding the underlying etiology and epidemiology of disease processes and their relationships to the age pattern of mortality (Gage 1991, 1994). And one of his former students has used the Siler model extensively in paleodemographic analysis (O'Connor 1995; O'Connor et al., 1997).

Although the Siler model is unquestionably useful for investigating human mortality, it does have some limitations. Its immature component, for example, is often difficult to estimate and interpret. There are two distinct reasons for this difficulty. First, although the negative exponential specification of juvenile mortality fits most human data fairly well, it is not etiologically derived (Gage 1989). Second, juvenile mortality is difficult to estimate reliably from small samples because information on it comes primarily from a tiny subset of the data – those from the first five years or so of life. In most populations, mortality during this segment of the life span is high but declines rapidly with small increments in age; thus, almost all the information about juvenile mortality must be extracted from an extremely narrow age range. With small samples,

the scale of the juvenile component can sometimes be estimated reasonably well, but capturing the shape is more problematic (Gage 1989). This problem is worsened in paleodemography because of the common under-representation of infants and young children in skeletal samples owing to differential preservation (Gordon and Buikstra 1981; Waldron 1987; Walker et al. 1988). This whole issue is important because early juvenile deaths make up a large fraction of all deaths in most human populations, especially preindustrial ones, and much of the variation in mortality among human populations falls in infancy and early childhood (see, for example, Figure 1). For these reasons a better theoretical model of mortality at juvenile ages would be useful.

A second limitation of the Siler model is that it assumes individuals in a population to be homogenous with respect to their genetic, physiological, environmental, and behavioral risks of death (Gage and Dyke 1986; Gage 1989). Variation in risk factors among individuals or subgroups in a population may influence the age pattern of mortality in ways that make comparative analyses difficult to interpret (Vaupel et al. 1979; Vaupel and Yashin 1985a,b; Gage 1989; Wood et al. 1992a,b; Himes 1994). We expand on this point in the next section.

Interpreting Competing Hazards Models When Mortality is Heterogeneous

In this section, we show that the competing hazard interpretation of the Siler model implies that the population being studied is homogenous in mortality risk – that is, the population is made up of individuals who are all subject to exactly the same causes of death and are equally susceptible to them. In the presence of heterogeneity, the model's parameters cannot be interpreted in the conventional way proposed by Siler (1979). In other words, models like the Gompertz-Makeham and Siler are implicitly models of homogenous risks. With heterogeneity among individuals in the risk of death, interpretation of the μ 's on the right-hand side of Equation (28) is not possible except under some not-very-plausible circumstances, as shown below.

Generally speaking, we do not believe that the members of any natural population, human or non-human, have exactly the same age-specific risks of death (for a discussion of this point, see Milner et al. 2000). For example, some individuals may be constitutionally frailer than others, a subset of individuals may engage in risky behavior,

or some individuals may simply live in riskier environments such as those associated with poverty. A number of recent advances in statistical methodology provide a framework for modeling heterogeneity among individuals in a population (see, for example, Heckman and Singer 1984; Manton et al. 1992). Another section of this paper discusses several different methods for incorporating heterogeneity into parametric mortality models in ways that improve our ability to interpret the estimated parameters.

The simplest imaginable form of heterogeneity is one in which variation in risk comes packaged in the form of two distinct subgroups – but we observe only the mixture of the two. As we show in the Appendix, the Siler model cannot be interpreted under this scenario as representing independent competing hazards. We will refer to models that are mixtures of two or more non-overlapping subgroups as *mixed hazards models* to distinguish them from *competing hazards models*. Are there any possible two-component mixed hazards models that can be interpreted in terms of independent competing causes? The Appendix shows that there are (see Equations A.11-A.12), but also that they make little if any biological sense. If we believe that heterogeneity was likely to have existed in our target population – and it has almost certainly existed in every human population – then we should probably abandon models that purport to be competing hazards models and replace them with specifications that can be interpreted explicitly in terms of mixtures of heterogeneous subgroups.

The Mixed-Makeham Model

In this section we develop a mixed hazard model of human mortality that fits as well as the Siler model – and, just as importantly, has no more parameters. Consider a population made up of two subgroups, and assume that each subgroup has a different constant (Makeham) hazard but that both subgroups have the same senescent (Gompertz) hazard. The model is thus a mixture of two Gompertz-Makeham models, but constrained so that the two senescent components are identical. Accordingly, we call it the *mixed-Makeham model*. The force of mortality in the mixture as a whole is

$$\begin{aligned}
\mu(a) &= p(a)(\alpha_1 + \alpha_3 e^{\beta_3 a}) + [1 - p(a)](\alpha_2 + \alpha_3 e^{\beta_3 a}) \\
&= p(a)\alpha_1 + [1 - p(a)]\alpha_2 + \alpha_3 e^{\beta_3 a}
\end{aligned} \tag{32}$$

where α_1 now represents the constant hazard in the first, high-risk subgroup and α_2 represents the constant hazard in the second, low-risk subgroup. The term $p(a)$ is the fraction of high-risk individuals among all the individuals alive at age a , given by

$$\begin{aligned}
p(a) &= \frac{pS_1(a)}{pS_1(a) + (1-p)S_2(a)} \\
&= \frac{p \exp\left[-\alpha_1 a + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right]}{p \exp\left[-\alpha_1 a + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right] + (1-p) \exp\left[-\alpha_2 a + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right]}
\end{aligned} \tag{33}$$

where p is the *initial* fraction of individuals in the high-risk subgroup (i.e. the fraction at birth). From the starting point p , $p(a)$ declines as high-risk individuals are selectively removed by death. As a result, the $p(a)\alpha_1$ term in Equation (32) gets smaller and smaller with age, and the overall force of mortality declines accordingly. Thus, even though there is no distinct juvenile component of mortality under the mixed-Makeham model, mortality declines during the early years of life as the aggregate mixture comes more and more to reflect the low-risk portion of the population (Figure 5). After a while, however, the shared senescent component begins to dominate the overall force of mortality, and the risk of death increases correspondingly at later ages in both subgroups.

The cohort PDF and survival function of the mixed-Makeham model are

$$\begin{aligned}
f_0(a) &= p \exp\left[-\alpha_1 a + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right] [\alpha_1 + \alpha_3 e^{\beta_3 a}] \\
&+ (1-p) \exp\left[-\alpha_2 a + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right] [\alpha_2 + \alpha_3 e^{\beta_3 a}]
\end{aligned} \tag{34}$$

and

$$S(a) = p \exp\left[-\alpha_1 a + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right] + (1 - p) \exp\left[-\alpha_2 t + \frac{\alpha_3}{\beta_3}(1 - e^{\beta_3 a})\right]. \quad (35)$$

We have shown elsewhere that this model usually fits paleodemographic mortality profiles at least as well as the Siler model (O'Connor 1995; Holman et al. 1997, 1998; O'Connor et al. 1997). The difference between this model and the Siler is that its parameters are easier to interpret and may provide clues about the existence of important forms of intra-population variation in material conditions that affect the risk of death. In addition, the α_1 and α_2 parameters of the mixed-Makeham model are estimated from observations drawn from the entire life span and are thus less sensitive to deficiencies in data on the very young than are the Siler parameters α_1 and β_1 .

A More General Approach to Modeling Heterogeneity

The above discussion has focused on discrete heterogeneity in which individuals can be assigned to one of two subgroups, each subgroup differing from the other but containing members who all share common mortality risks. This approach can be extended to any number of discrete subgroups. Subgroups may have risks of death that are all drawn from the same distribution but with different parameter values, or they may have different distributions altogether. Mixtures of different distributions have been used for a number of models in demography, for example, to describe the postpartum resumption of menses (Ford and Kim 1987) and pregnancy loss (Wood 1989; Holman 1996). Recently, Louzada-Neto (1999) has proposed a “polyhazard” mortality model along these same lines.

A reasonable question to ask is whether we can justify adding risk groups without limit. Additional subgroups are perfectly easy to handle mathematically, but parameter estimation becomes increasingly difficult with each latent subgroup thrown into the pot. Most paleodemographic samples would not be able to cope with more than two or three subgroups. With many subgroups, moreover, we begin to lose the straightforward interpretation associated with the simple two-subgroup model. And once we forfeit the ability to interpret parameters, we descend from etiologic modeling to empirical curve fitting.

Rather than blindly subdivide the population into hypothetical subgroups, we might consider a population that consists of such a large number of subgroups that the risk of death appears to vary continuously among individuals. We can then think in terms of a continuous probability density function of underlying risk rather than proportions falling into discrete categories of risk. If z is the individual-level component of the risk of death – that is, the part of the risk that varies among members of the population – then we can write $g(z)$ for the continuous distribution of risk. The age-at-death distribution that we observe is then the expectation over all values of z :

$$\bar{f}_0(a) = \int_{-\infty}^{\infty} g(z) f_0(a|z) dz. \quad (36)$$

We can specify z in $f_0(a|z)$ in a number of ways, including as a covariate on a particular parameter or on the force of mortality as a whole, as in the proportional hazards model (Cox 1972; Manton et al. 1986).

Several researchers have shown that parameter estimates can be disturbingly sensitive to the precise choice of equations for $g(z)$ or $f_0(a|z)$ (Heckman and Walker 1987; Manton et al. 1992; Moreno 1994; Rodríguez 1994). Consequently, specification of these terms should be based, whenever possible, on some theory about the underlying mechanisms that generate the heterogeneity in risk (see, for example, Weiss 1990; Wood 1998). The gamma, beta, and lognormal distributions are frequently used to model heterogeneity. For example, Gage (1989) has explored the behavior of the Siler model with gamma-distributed heterogeneity. But these specifications are often based on mathematical convenience rather than any established biological principles.

Capturing the Sex Differential

There is one form of heterogeneity that can reasonably be captured by a simple dichotomous model: the difference between males and females. Since the one thing we know about human mortality is that it always differs between the sexes – and sometimes markedly so – it makes little sense to apply a single, homogeneous parametric model of mortality to a combined sample of male and female skeletons. Moreover, sex differences in mortality are interesting in their own right, and we would like to be able to say something about them. But to examine these differences using the Rostock protocol in its

present form we must, in effect, cut our sample size in half by applying the method to the two sexes separately. Even worse, we have to throw out some important biological constraints on our parameter values: no matter what the difference in mortality between the sexes, the male and female segments of the population have to be growing or declining at the same rate r (Coale 1972:53-58), and meiosis re-establishes a sex ratio at birth that is always close to one-half. In addition, male and female age-specific mortality rates are not *completely* unrelated to each other, but differ in quite limited and specific ways (Keyfitz 1985:54-76). It would be a fine thing if we could make use of these universal constraints in estimating our model.

Imagine that our vector of skeletal traits c contains measures that provide information about sex as well as age in a sample made up both male and female skeletons. What is the probability, over the sample as a whole, of observing a particular c value – say, c_i ? It is just the marginal density of c_i :

$$\begin{aligned} \Pr(c_i) &= \sum_{k \in (0,1)} \int_0^{\infty} \Pr(c_i | a, k) f_r(a | k) \Pr(k) da \\ &= \sum_{k \in (0,1)} \int_0^{\infty} \Pr(c_i | a, k) f_r(a, k) da, \end{aligned} \tag{37}$$

where k is an indicator variable for sex ($k = 0$ for females, 1 for males), and $f_r(a, k)$ is the joint distribution of deaths by age and sex in the target population. (The fact that this distribution is subscripted with an r indicates that it has been corrected for non-stationarity.) The likelihood function for the whole sample of n skeletons is $\prod_{i=1}^n \Pr(c_i)$.

To use this likelihood we need two new quantities: an estimate of $\Pr(c_i | a, k)$ from a reference sample in which both age and sex are known, and a parametric expression for $f_r(a, k)$. The first is a purely statistical problem, and we ignore it here. We focus instead on finding an expression for $f_r(a, k)$.

We begin with the elementary relationship $\Pr(a, k) = \Pr(k | a) \Pr(a)$. It follows that

$$f_r(a, k) = \begin{cases} [1 - \rho(a)]f_r(a), & \text{if } k = 0 \\ \rho(a)f_r(a), & \text{if } k = 1 \end{cases}, \quad (38)$$

where $\rho(a)$ is the proportion of surviving people at age a for whom $k = 1$. That is, $\rho(a) = \Pr(k = 1 | a)$. This quantity can be found as

$$\rho(a) = \frac{\rho(0)S_0(a)}{\rho(0)S_0(a) + [1 - \rho(0)]S_1(a)}, \quad (39)$$

where $\rho(0)$ is the sex ratio at birth expressed as a proportion and the subscripts 0 and 1 refer to females and males respectively. Now, everyone knows that $\rho(0)$ is not *exactly* equal to $\frac{1}{2}$, but it never strays very far from it (in some populations it soars to 0.51, in others it plunges to 0.49). So we assume from now on that $\rho(0) = \frac{1}{2}$. Thus, Equation (39) reduces to

$$\rho(a) = \frac{S_0(a)}{S_0(a) + S_1(a)}. \quad (40)$$

Since the numbers of both sexes in a stable population must be changing at the same constant rate r , it must be the case that

$$f_r(a) = \frac{\bar{\mu}(a)\bar{S}(a)e^{-ra}}{\int_0^{\infty} \bar{\mu}(x)\bar{S}(x)e^{-rx} dx}. \quad (41)$$

In Equation (41), the bars denote weighted averages over the two sexes. That is,

$$\bar{\mu}(a) = [1 - \rho(a)]\mu_0(a) + \rho(a)\mu_1(a) \quad (42)$$

and

$$\begin{aligned}\bar{S}(a) &= [1 - \rho(0)]S_0(a) + \rho(0)S_1(a) \\ &= \frac{1}{2}[S_0(a) + S_1(a)]\end{aligned}\tag{43}$$

if $\rho(0) = \frac{1}{2}$.

How should we model $\mu_k(a)$ and $S_k(a)$ themselves? As a general strategy, we propose treating the mortality of one sex as a baseline and letting the other sex differ from it in what might be called “quasi-proportional” fashion:

$$\mu_0(a) = \textit{baseline hazard} ,\tag{44}$$

and

$$\mu_1(a) = \mu_0(a)e^{\delta(a)} ,\tag{45}$$

where $\delta(a)$ is some function of age that models the sex differential $\ln[\mu_1(a)/\mu_0(a)]$. We have tested several specifications of $\delta(a)$ against data from the empirical life tables compiled by Keyfitz and Flieger (1968, 1990). Although we find that $\delta(a)$ is positive at all ages in almost all human populations (male mortality is almost always greater than female mortality), neither a constant difference $\delta(a) = \kappa$ nor the linear function $\delta(a) = \alpha + \beta a$ captures the real age pattern of the sex differential. Instead, the empirical differential is typically bimodal by age, peaking at ages 20-25 years and again at 55-65 years (sometimes one mode is higher than the other, sometimes a mode is missing). If our quasi-proportional model is to be implemented, we will eventually need to identify a simple function that duplicates this pattern.

What happens from this point on depends in its details on the precise way we decide to specify $\delta(a)$. For the present, we will assume that $\delta(a)$ acts as a true proportional hazard so that we can sketch out the rest of the method as simply as possible. If we can also assume that $\rho(0) = \frac{1}{2}$, then

$$\bar{\mu}(a) = \frac{\mu_0(a)}{2\bar{S}(a)} [S_0(a) + S_1(a)e^{\delta(a)}] .\tag{46}$$

Consequently,

$$\bar{\mu}(a)\bar{S}(a) = \frac{1}{2}\mu_0(a)[S_0(a) + S_1(a)e^{\delta(a)}], \quad (47)$$

which is what we need to substitute in Equation (41).

Combining all these results and rearranging, the likelihood of a set of observed age- and sex-related traits in a sample of n skeletons is

$$\begin{aligned} L &= \prod_{i=1}^n \sum_{k \in (0,1)} \int_0^{\infty} \text{Pr}^*(c_i | a, k) f_r(a, k) da \\ &= \prod_{i=1}^n \sum_{k \in (0,1)} \int_0^{\infty} \text{Pr}^*(c_i | a, k) \left[\frac{S_k(a)}{S_0(a) + S_1(a)} \right] f_r(a) da, \end{aligned} \quad (48)$$

where, as before, the asterisk (*) denotes an estimate made from data on a reference sample, and

$$f_r(a) = \frac{\mu_0(a)[S_0(a) + S_1(a)e^{\delta(a)}]e^{-ra}}{\int_0^{\infty} \mu_0(x)[S_0(x) + S_1(x)e^{\delta(x)}]e^{-rx} dx}. \quad (49)$$

Once we specify parametric models for $\mu_0(a)$ and $\delta(a)$, we can use Equations (48) and (49) to estimate $f_r(a, k)$ from the target sample by maximum likelihood. And once we have that in hand, we have both the mortality profile (by sex) and the population growth rate of the target population. We can also say something about how old an individual skeleton in the target population is likely to be and what its probable sex is – things we would like to know for paleopathological purposes – by plugging our estimates of $f_r(a, k)$ into the multivariate generalization of Bayes's theorem:

$$\hat{\text{Pr}}(a, k | c_i) = \frac{\text{Pr}^*(c_i | a, k) \hat{f}_r(a, k)}{\sum_{y \in (0,1)} \int_0^\infty \text{Pr}^*(c_i | x, y) \hat{f}_r(x, y) dx} = \frac{\text{Pr}^*(c_i | a, k) \hat{f}_r(a, k)}{\hat{\text{Pr}}(c_i)}, \quad (50)$$

where the hats (^) denote maximum likelihood estimates from the target sample. This expression, which is a straightforward extension of the original Rostock protocol, ought to provide us with the proper error structure for both our age estimates and our classifications by sex.

Discussion

In this paper, we have reviewed several parametric models of mortality processes that can be used in conjunction with the Rostock approach to paleodemographic mortality analysis. Since paleodemographers will never be able to fit complicated models to their skeletal data, we have emphasized simple models that still do a reasonable job of capturing the main features of the human mortality curve. (The fact that the equations describing these models often look dauntingly complicated should not obscure their underlying simplicity.) At the same time, we have tried to focus most of our attention on models that support at least a certain amount of etiologic interpretation, so that we may actually stand to learn something interesting from our skeletal samples instead of just fitting meaningless curves to them.

On biological grounds, we believe that within-population heterogeneity in health and the risk of death ought to be a central theoretical concern of paleodemography (see Wood et al. 1992b; Wood 1998; Milner et al. 2000). Accordingly, we have spent a fair amount of effort in exploring the implications of heterogeneity for etiologic models of mortality. One form of heterogeneity that is always with us – *viva la hétérogénéité!* – is the difference between males and females. As it happens, sex is also one of the fundamental dimensions along which we would like to be able to classify our skeletons. Therefore, we have proposed an extension of the Rostock approach that estimates the sex differential at the same time that it probabilistically assigns age and sex to our skeletons. One of the challenges in applying this extension will be to find a simple parameterization

of the sex differential in risk of death. In other words, we need even more parametric models, not fewer.

Appendix

The age-at-death distribution for skeletons deposited over time

In a series of famous papers, Lotka (1907, 1922, 1931) worked out the characteristics of the stationary and stable population at any instant in time. In examining skeletons from archaeological sites, however, we are never dealing with a single instant of time, but rather with some more or less prolonged (and usually unknown) period during which skeletons are laid down. How do we go from the stable or stationary age-at-death distribution at one time to the corresponding distribution over the entire period of deposition?

If skeletons are accumulated over a span of time equal to ω , then

$$\Pr(a) = \kappa \int_0^{\omega} \frac{f_r(a,t)}{\int_0^{\infty} f_r(x,t) dx} dt, \quad (\text{A.1})$$

where $f_r(a,t)$ is the age-at-death distribution (corrected for non-stationarity) at time t , and κ is a normalizing constant ensuring that $\Pr(a)$ integrates to one. If the population is stationary, $r = 0$ and

$$f_0(a,t) = \frac{\mu(a)S(a)}{\int_0^{\infty} \mu(x)S(x) dx} \quad (\text{A.2})$$

(Lotka 1907). Since nothing on the right-hand side of Equation (A.2) varies with t , its integral is simply a constant equal to $f_0(a,t)$ itself:

$$\Pr(a) = \kappa \int_0^{\omega} \frac{f_0(a,t)}{\int_0^{\infty} f_0(x,t) dx} dt = \frac{\mu(a)S(a)}{\int_0^{\infty} \mu(x)S(x) dx}. \quad (\text{A.3})$$

Note that the right-hand portion of this expression does not contain ω . Thus, the fact that we usually do not know the exact period over which skeletons were deposited is of no concern.

If the population is stable but not stationary ($r \neq 0$), we must take into account the fact that the number of skeletons being deposited each year changes in proportion to

population growth or decline. In general, the number of deaths age a at time t is $n(a,t)\mu(a)$, where $n(a,t)$ is the number of *living* individuals age a at risk of death at time t . But since a stable population is closed to migration, $n(a,t) = n(0,t-a)S(a)$. And since $f_0(a) = \mu(a)S(a)$, the number of deaths at a in t , $n(a,t)\mu(a)$, becomes $n(0,t-a)f_0(a)$. Lotka (1907) showed that the number of births changes exponentially in a stable population. Thus, the number of deaths age a at t can be rewritten as

$$n(0,0)e^{r(t-a)} f_0(a) = n(0,0)e^{rt} f_0(a)e^{-ra}. \quad (\text{A.4})$$

Substituting in (A.1), we have

$$\Pr(a) = \kappa \int_0^{\omega} \frac{f_r(a,t)}{\int_0^{\infty} f_r(x,t)} dt = \kappa \int_0^{\omega} \frac{n(0,0)e^{rt} f_0(a,t)e^{-ra}}{\int_0^{\infty} n(0,0)e^{rt} f_0(x,t)e^{-rx} dx} dt. \quad (\text{A.5})$$

Since the term $n(0,0)e^{rt}$ in the denominator of the right-hand side of Equation (A.5) does not vary with x , we can pull it out of the inner integral and cancel it from the numerator and denominator. We are left with the relation

$$\Pr(a) = \frac{\mu(a)S(a)e^{-ra}}{\int_0^{\infty} \mu(x)S(x)e^{-rx} dx} \quad (\text{A.6})$$

in a stable population with growth rate r . Again, this expression does not contain ω , so we do not need to know its value. If $r = 0$, this equation reduces to Equation (A.3), which highlights the fact that the stationary population is just a special case of the stable population. It also shows that we can use Equation (A.6) in our likelihood function to estimate r whether it is zero or non-zero, positive or negative.

The implications of heterogeneity for competing hazards models

Can models such as the Siler model, which are normally interpreted as models of competing hazards, support such an interpretation when the population involved is

heterogeneous in the risk of death? Consider a population for which there are only two types of individuals. Individuals of type 1 are all at hazard $\mu_1(a)$ and individuals of type 2 are all at hazard $\mu_2(a)$. Assume that within the two subgroups individuals are homogenous for mortality risk, and let the proportion of newborns in group 1 be p and in group 2 be $(1 - p)$. Since $f_0(a)$ and $S(a)$ are probabilities, they can be found for the mixture of the two groups by using the law of total probability:

$$S(a) = pS_1(a) + (1 - p)S_2(a) \quad (\text{A.7})$$

and

$$f_0(a) = pf_{0,1}(a) + (1 - p)f_{0,2}(a). \quad (\text{A.8})$$

Using Equation (6), we can now write the mortality function for the entire population as

$$\mu(a) = \frac{f_0(a)}{S(a)} = \frac{pf_{0,1}(a) + (1 - p)f_{0,2}(a)}{pS_1(a) + (1 - p)S_2(a)}. \quad (\text{A.9})$$

Clearly, Equation (A.9) does not take the form $\mu_1(a) + \mu_2(a)$. It does not even take the form of a simple weighted average of the two subgroups: $p\mu_1(a) + (1 - p)\mu_2(a)$. The proper total hazard in terms of both the sub-component hazards is $\mu(a) = p(a)\mu_1(a) + [1 - p(a)]\mu_2(a)$. In this expression, $p(a)$ is the fraction of those individuals surviving to age a who belong to group 1, equal to

$$p(a) = \frac{pS_1(a)}{S(a)} = \frac{pS_1(a)}{pS_1(a) + (1 - p)S_2(a)}. \quad (\text{A.10})$$

The numerator is the fraction of survivors in group 1 at age a and the denominator is the fraction of all survivors at age a .

The above exercise shows that interpreting the individual components of a “competing hazards” model as if they really were independent competing causes of death may be inappropriate when the population consists of two subgroups. Can the parameters

of a two-subgroup mixed hazards model *ever* be interpreted as a competing hazards model? Some algebra reveals that this is permissible if

$$\begin{aligned}\mu_1(a) + \mu_2(a) &= \frac{pf_{0,1}(a) + (1-p)f_{0,2}(a)}{pS_1(a) + (1-p)S_2(a)} \\ \frac{f_{0,1}(a)}{S_1(a)} + \frac{f_{0,2}(a)}{S_2(a)} &= \frac{pf_{0,1}(a) + (1-p)f_{0,2}(a)}{pS_1(a) + (1-p)S_2(a)} \quad . \quad (\text{A.11}) \\ \frac{f_{0,1}(a)S_2 + f_{0,2}(a)S_1}{S_1(a)S_2(a)} &= \frac{pf_{0,1}(a) + (1-p)f_{0,2}(a)}{pS_1(a) + (1-p)S_2(a)}\end{aligned}$$

Since $f_0(a) = -dS(a)/da$, the equivalencies in Equation (A.11) hold when

$$p = \frac{f_{0,1}(a)S_2(a)^2}{f_{0,1}(a)S_2(a)^2 + f_{0,2}(a)S_1(a)^2} \quad (\text{A.12})$$

or in the trivial case in which each subgroup experiences exactly the same risk, $\mu_1(a) = \mu_2(a)$. If one of these conditions – the first of which is completely arbitrary and the second not a model of heterogeneity at all – is not met, competing hazards models such as the Siler model are inappropriate and cannot be interpreted properly. It can be shown that similar conditions hold when more than two heterogeneous subgroups exist in the population.

Acknowledgments

This paper has benefited enormously from the general discussion that took place at the meeting on paleodemographic age estimation held at the Max Planck Institute for Demographic Research, Rostock, Germany, in the summer of 2000. We thank all the meeting participants for their lively contributions to that discussion. We are especially grateful to Jesper Boldsen, Rob Hoppa, Lyle Konigsberg, Brad Love, George Milner, Bethany Usher, and Jim Vaupel for helpful suggestions, which we have not always been wise enough to follow.

References Cited

- Aalen OOA (1989) A linear regression model for the analysis of life times. *Statistics in Medicine*, **8**, 907-925.
- Acsádi G and Nemeskéri J (1970) *History of Human Life Span and Mortality*. Budapest: Akadémiai Kiadó.
- Alesan A, Malgosa A and Simó C (1999) Looking into the demography of an Iron Age population in the western Mediterranean. I. Mortality. *American Journal of Physical Anthropology*, **110**, 285-301.
- Andersen PK (1988) Multistate models in survival analysis: A study of nephropathy and mortality in diabetes. *Statistics in Medicine*, **7**, 661-670.
- Angel JL (1969) The bases of paleodemography. *American Journal of Physical Anthropology*, **30**, 427-438.
- Benedictow OJ (1996) *The Medieval Demographic System of the Nordic Countries (revised edition)*. Oslo: Middelalderforlaget.
- Bourgeois-Pichat J (1971) Stable, semi-stable populations and growth potential. *Population Studies*, **25**, 235-254.
- Brass W (1971) On the scale of mortality. In *Biological Aspects of Demography*, ed. W. Brass, pp. 69-110. London: Taylor and Francis.

- Brass W (1975) *Methods for Estimating Fertility and Mortality from Limited and Defective Data*. Chapel Hill, N.C.: Carolina Population Center, University of North Carolina.
- Brass W and Coale AJ (1968) Methods of analysis and estimation. In Brass W (ed.): *The Demography of Tropical Africa*, ed. W. Brass, pp. 88-150. Princeton, N.J.: Office of Population Research, Princeton University.
- Brooks A, Lithgow GJ and Johnson TE (1994) Mortality rates in a genetically heterogeneous population of *Caenorhabditis elegans*. *Science*, **263**, 668-671.
- Carey JR, Liedo P, Orozco D and Vaupel JW (1992) Slowing of mortality rates at older ages in large medfly cohorts. *Science*, **258**, 457-461.
- Carrier NH and Hobcraft J (1971) *Demographic Estimation for Developing Societies*. London: Population Investigation Committee, London School of Economics.
- Chowdhury MK, Becker S, Razzaque A, Sarder AM, Shaikh K and Chen LC (1981) *Demographic Surveillance System – Matlab: Vital Events and Migration 1978*. Scientific Report No. 47. Dhaka, Bangladesh: International Centre for Diarrheal Disease Research.
- Coale AJ (1972) *The Growth and Structure of Human Populations: A Mathematical Investigation*. Princeton, N.J.: Princeton University Press.
- Coale AJ and Demeny P (1966) *Regional Model Life Tables and Stable Populations*. New York: Academic Press.
- Cox DR (1962) *Renewal Theory*. London: Chapman and Hall.
- Cox DR (1972) Regression models and life-tables (with discussion). *Journal of the Royal Statistical Society (Series B)*, **34**, 187-220.
- Dyke B, Gage TB, Ballou JD, Petto AJ, Tardif AD and Williams LE (1993) Model life tables for the smaller New World monkeys. *American Journal of Primatology*, **29**, 269-285.
- Evans M, Hastings N and Peacock B (2000) *Statistical Distributions (third edition)*. New York: John Wiley and Sons.

- Ewbank DC, Gomez de Leon JC and Stoto MA (1983) A reducible four-parameter system of model life tables. *Population Studies*, **37**, 105-127.
- Finch CE (1990) *Longevity, Senescence, and the Genome*. Chicago: University of Chicago Press.
- Ford K and Kim Y (1987) Distributions of postpartum amenorrhea: Some new evidence. *Demography*, **24**, 413-430.
- Fukui HH, Xiu L and Curtsinger JW (1993) Slowing of age-specific mortality rates in *Drosophila melanogaster*. *Experimental Gerontology*, **28**, 585-599.
- Gage TB (1988) Mathematical hazard models of mortality: An alternative to model life tables. *American Journal of Physical Anthropology*, **76**, 429-441.
- Gage TB (1989) Bio-mathematical approaches to the study of human variation in mortality. *Yearbook of Physical Anthropology*, **32**, 185-214.
- Gage TB (1990) Variation and classification of human age patterns of mortality: Analysis using competing hazards models. *Human Biology*, **62**, 589-614.
- Gage TB (1991) Causes of death and the components of mortality: Testing the biological interpretations of a competing hazards model. *American Journal of Human Biology*, **3**, 289-300.
- Gage TB (1994) Population variation in cause of death: Level, gender, and period effects. *Demography*, **31**, 271-296.
- Gage TB (1998) The comparative demography of primates, with some comments on the evolution of life histories. *Annual Review of Anthropology*, **27**, 197-221.
- Gage TB and Dyke B (1986) Parameterizing abridged mortality tables: The Siler three-component hazard model. *Human Biology*, **58**, 275-291.
- Gage TB and Dyke B (1988) Model life tables for the larger Old World monkeys. *American Journal of Primatology*, **16**, 305-320.
- Gage TB and Mode CJ (1993) Some laws of mortality: How well do they fit? *Human Biology*, **65**, 445-461.

- Gage TB and O'Connor KA (1994) Nutrition and the variation in level and age patterns of mortality. *Human Biology*, **66**, 77-103.
- Gavrilov LA and Gavrilova NS (1991) *The Biology of Life Span: A Quantitative Approach*. London: Harwood Academic Publishers.
- Gompertz B (1825) On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. *Philosophical Transactions of the Royal Society of London (Series A)*, **115**, 513-585.
- Gordon CC and Buikstra JE (1981) Soil pH, bone preservation, and sampling bias at mortuary sites. *American Antiquity*, **46**, 566-571.
- Green S, Green S and Armelagos GJ (1974) Settlement and mortality of the Christian site (1050 A.D. - 1300 A.D.) of Meinarti (Sudan). *Journal of Human Evolution*, **3**, 297-316.
- Greene DL, Van Gerven DP and Armelagos GJ (1986) Life and death in ancient populations: Bones of contention in paleodemography. *Human Evolution*, **1**, 193-207.
- Hamilton WD (1966) The moulding of senescence by natural selection. *Journal of Theoretical Biology*, **12**, 12-45.
- Heckman JJ and Singer B (1984) Econometric duration analysis. *Journal of Econometrics*, **24**, 63-132.
- Heckman JJ and Walker JR (1987) Using goodness of fit and other criteria to choose among competing duration models: A case study of Hutterite data. In *Sociological Methodology 1987*, ed. C. Clogg, pp. 247-307. New York: American Sociological Association.
- Heligman L and Pollard JH (1980) The age pattern of mortality. *Journal of the Institute of Actuaries*, **107**, 49-80.
- Himes CL (1994) Age patterns of mortality and cause-of-death structures in Sweden, Japan, and the United States. *Demography*, **31**, 633-650.
- Holman DJ (1996) *Fecundability and Total Fetal Loss in Bangladesh*. Doctoral dissertation, Pennsylvania State University, University Park, Pa.

- Holman DJ, O'Connor KA, Wood JW and Boldsen JL (1997) Correcting for nonstationarity in paleodemographic mortality models [abstract]. *American Journal of Physical Anthropology, Supplement*, **24**, 132.
- Holman DJ, O'Connor KA, Wood JW and Boldsen J (1998) Estimating population growth rates from skeletal samples. Paper presented at the annual meeting of the American Anthropological Association, Philadelphia, Pa.
- Horiuchi S and Coale AJ (1990) Age patterns of mortality for older women: An analysis using age-specific rate of mortality change with age. *Mathematical Population Studies*, **2**, 245-267.
- Kannisto V (1994) *Development of Oldest-Old Mortality, 1950-1990: Evidence from 28 Developed Countries*. Odense, Denmark: Odense University Press.
- Johansson SR and Horowitz S (1986) Estimating mortality in skeletal populations: Influence of the growth rate on the interpretation of levels and trends during the transition to agriculture. *American Journal of Physical Anthropology*, **71**, 233-250.
- Keyfitz N (1968) *Introduction to the Mathematics of Population*. Reading, Mass.: Addison-Wesley.
- Keyfitz N (1985) *Applied Mathematical Demography (second edition)*. Berlin: Springer-Verlag.
- Keyfitz N and Flieger W (1968) *World Population: An Analysis of Vital Data*. Chicago: University of Chicago Press.
- Keyfitz N and Flieger W (1990) *World Population Growth and Aging: Demographic Trends in the Late Twentieth Century*. Chicago: University of Chicago Press.
- Konigsberg LW and Frankenberg SR (1992) Estimation of age structure in anthropological demography. *American Journal of Physical Anthropology*, **89**, 235-256.
- Konigsberg LW and Frankenberg SR (1994) Paleodemography: "Not quite dead". *Evolutionary Anthropology*, **3**, 92-105.

- 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.: Center for Archaeological Investigations, Southern Illinois University.
- Konigsberg LW and Herrmann NP (2000) Estimation of age-at-death distributions: Further examples from Indian Knoll. Presented at the Workshop on Paleodemographic Age Estimation, Max Planck Institute for Demographic Research, Rostock, Germany.
- Lanphear KM (1989) Testing the value of skeletal samples in demographic research: A comparison with vital registration samples. *International Journal of Anthropology*, **4**, 185-193.
- Ledermann S and Breas J (1959) Les dimensions de la mortalité. *Population*, **14**, 637-682.
- Lopez A (1961) *Problems in Stable Population Theory*. Princeton, N.J.: Office of Population Research, Princeton University.
- Lotka AJ (1907) Relation between birth rates and death rates. *Science*, **26**, 21-22.
- Lotka AJ (1922) The stability of the normal age distribution. *Proceedings of the National Academy of Sciences*, **8**, 339-345.
- Lotka AJ (1931) The structure of a growing population. *Human Biology*, **3**, 459-493.
- Louzada-Neto F (1999) Polyhazard models for lifetime data. *Biometrics*, **55**, 1281-1285.
- Lovejoy CO, Meindl RS, Pryzbeck TR, Barton TS, Heiple KG and Kottling D (1977) Paleodemography of the Libben site, Ottawa County, Ohio. *Science*, **198**, 291-293.
- Luder HU (1993) Hazard rates and causes of death in a captive group of crab-eating monkeys (*Macaca fascicularis*). *American Journal of Primatology*, **39**, 40-50.
- Makeham WM (1860) On the law of mortality. *Journal of the Institute of Actuaries*, **13**, 325-358.

- Manton KG, Singer B and Woodbury MA (1992) Some issues in the quantitative characterization of heterogeneous populations. In *Demographic Applications of Event History Analysis*, ed. J. Trussell, R. Hankinson and J. Tilton, pp. 9-37. Oxford: Oxford University Press.
- Manton KG and Stallard E (1988) *Chronic Disease Modelling: Measurement and Evaluation of the Risks of Chronic Disease Processes*. London: Charles Griffin.
- Manton KG, Stallard E and Vaupel JW (1986) Alternative models for the heterogeneity of mortality risks among the aged. *Journal of the American Statistical Association*, **81**, 635-644.
- Mays S (1992) Taphonomic factors in a human skeletal assemblage. *Circaea*, **9**, 54-58.
- Mensforth RP (1990) Paleodemography of the Carlston Annis (Bt-5) Late Archaic skeletal population. *American Journal of Physical Anthropology*, **82**, 81-99.
- Milner GR, Wood JW and Boldsen JL (2000) Paleodemography. In *Biological Anthropology of the Human Skeleton*, ed. S.R. Saunders and M.A. Katzenberg, pp. 467-497. New York: Wiley-Liss.
- Mode CJ (1985) *Stochastic Processes in Demography and Their Computer Implementation*. Berlin: Springer-Verlag.
- Mode CJ and Busby RC (1982) An eight-parameter model of human mortality: The single decrement case. *Bulletin of Mathematical Biology*, **44**, 647-659.
- Mode CJ and Jacobsen ME (1984) A parametric algorithm for computing model period and cohort human survival functions. *International Journal of Bio-Medical Computing*, **15**, 341-356.
- Moore JA, Swedlund AC and Armelagos GJ (1975) The use of life tables in paleodemography. *American Antiquity*, **40**, 57-70.
- Moreno L (1994) Frailty selection in bivariate survival models: A cautionary note. *Mathematical Population Studies*, **4**, 225-233.

- Nam CB (1990) Mortality differentials from a multiple-cause-of-death perspective. In *Measurement and Analysis of Mortality: New Approaches*, ed. J. Vallin, S. D'Souza and A. Palloni, pp. 328-342. Oxford: Oxford University Press.
- Nordling CO (1953) A new theory on the cancer inducing mechanism. *British Journal of Cancer*, **7**, 68-72.
- 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, State University of New York, Albany, N.Y.
- O'Connor KA, Holman DJ, Boldsen J, Wood JW and Gage TB (1997) Competing and mixed hazards models of mortality: Osteological applications [abstract]. *American Journal of Physical Anthropology, Supplement*, **24**, 180.
- Paine RR (1989) Model life table fitting by maximum likelihood estimation: A procedure to reconstruct paleodemographic characteristics from skeletal age distributions. *American Journal of Physical Anthropology*, **79**, 51-61.
- Parlett B (1970) Ergodic properties of populations. I. The one sex model. *Theoretical Population Biology*, **1**, 191-207.
- Preston SH (1976) *Mortality Patterns in National Populations, with Special Reference to Recorded Causes of Death*. New York: Academic Press.
- Rodríguez G (1994) Statistical issues in the analysis of reproductive histories using hazard models. *Annals of the New York Academy of Sciences*, **709**, 266-279.
- Sattenspiel L and Harpending H (1983) Stable populations and skeletal age. *American Antiquity*, **48**, 489-498.
- Schoen R (1975) Constructing increment-decrement life tables. *Demography*, **12**, 313-324.
- Siler W (1979) A competing-risk model for animal mortality. *Ecology*, **60**, 750-757.
- Siler W (1983) Parameters of mortality in human populations with widely varying life spans. *Statistics in Medicine*, **2**, 373-380.

- Thatcher AR, Kannisto V and Vaupel JW (1997) *The Force of Mortality from Age 80 to 120*. Odense, Denmark: Odense University Press.
- Thiele PN (1871) On a mathematical formula to express the rate of mortality throughout the whole of life. *Journal of the Institute of Actuaries*, **16**, 313-329.
- Thompson WA (1988) *Point Process Models with Applications to Safety and Reliability*. London: Chapman and Hall.
- United Nations (1955) *Age and Sex Patterns of Mortality: Model Life Tables for Underdeveloped Countries*. Population Studies, No. 22. New York: United Nations.
- United Nations (1956) *Methods for Population Projections by Age and Sex*. Population Studies, No. 25. New York: United Nations.
- Vaupel JW (1997) Trajectories of mortality at advanced ages. In *Between Zeus and the Salmon: The Biodemography of Longevity*, ed. K.W. Wachter and C.E. Finch, pp. 17-37. Washington, D.C.: National Academy Press.
- Vaupel JW, Carey JR, Christensen K, Johnson TE, Yashin AI, Holm NV, Iachine IA, Kannisto V, Khazaeli AA, Liedo P, Longo VD, Zeng Y, Manton KG and Curtsinger JW (1998) Biodemographic trajectories of longevity. *Science*, **280**, 855-860.
- Vaupel JW, Johnson TE and Lithgow GJ (1994) Rates of mortality in populations of *Caenorhabditis elegans*. *Science*, **266**, 826.
- Vaupel JW, Manton KG and Stallard E (1979) The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography*, **16**, 439-454.
- Vaupel JW and Yashin AI (1985a) Heterogeneity's ruses: Some surprising effects of selection on population dynamics. *American Statistician*, **39**, 176-185.
- Vaupel JW and Yashin AI (1985b) The deviant dynamics of death in heterogeneous populations. In *Sociological Methodology 1985*, ed. N.B. Tuma, pp. 179-211. San Francisco: Jossey-Bass.
- Waldron T (1987) The relative survival of the human skeleton: Implications for paleopathology. In *Death, Decay, and Reconstruction*, ed. A. Boddington, A.N.

- Garland and R.C. Janaway, pp. 55-64. Manchester, U.K.: Manchester University Press.
- 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.
- Weibull W (1951) A statistical distribution function of wide applicability. *Journal of Applied Mechanics*, **18**, 293-297.
- Weiss KM (1973) *Demographic Models for Anthropology*. Washington, D.C.: Society for American Archaeology.
- Weiss KM (1990) The biodemography of variation in human frailty. *Demography*, **27**, 185-206.
- Weiss KM and Chakraborty R (1990) Multistage models and the age patterns of cancer: Does the statistical analogy imply genetic homology? In *Familial Adenomatous Polyposis*, ed. L. Herrera, pp. 77-89. New York: Wiley-Liss.
- Whittemore AS and Keller J (1978) Quantitative theories of carcinogenesis. *SIAM Review*, **20**, 1-30.
- Wood JW (1987a) Problems of applying model fertility and mortality schedules to demographic data from Papua New Guinea. In *The Survey Under Difficult Conditions: Population Data Collection and Analysis in Papua New Guinea*, ed. T.M. McDevitt, pp. 27-53. New Haven, Conn.: Human Relations Area File Press.
- Wood JW (1987b) The genetic demography of the Gainj of Papua New Guinea. 2. Determinants of effective population size. *The American Naturalist*, **129**, 165-187.
- Wood JW (1989) Fecundity and natural fertility in humans. *Oxford Review of Reproductive Biology*, **11**, 61-109.
- Wood JW (1998) A theory of preindustrial population dynamics: Demography, economy, and well-being in Malthusian systems. *Current Anthropology*, **39**, 99-135.
- Wood JW, Holman DJ, Weiss KM, Buchanan AV and LeFor B (1992a) Hazards models for human population biology. *Yearbook of Physical Anthropology*, **35**, 43-87.

- Wood JW, Milner GR, Harpending H and Weiss KM (1992b) The osteological paradox: Problems of inferring prehistoric health from skeletal samples. *Current Anthropology*, **33**, 343-370.
- Wood JW, Weeks SC, Bentley GR and Weiss KM (1994) Human population biology and the evolution of aging. In *Biological Anthropology and Aging: Perspectives on Human Variation Over the Life Span*, ed. D.E. Crews and R.M. Garruto, pp. 19-75. Oxford: Oxford University Press.
- Zaba B (1979) The four-parameter logit life table system. *Population Studies*, **33**, 79-100.
- Zaba B (1981) Use of the relational Gompertz model in analysing data collected in retrospective surveys. *Centre for Population Studies Working Paper 81-2*. London: London School of Hygiene and Tropical Medicine.

Figure Captions

Figure 1. Age-specific force of mortality in four human populations with widely differing levels of mortality: Sweden 1985, females (Keyfitz and Flieger 1990); El Salvador 1950, males (Keyfitz and Flieger 1968); Bangladesh 1978, both sexes (Chowdhury et al. 1981); Gainj (highland New Guinea) 1970-1977, males (Wood 1987b). Note that the Gainj, a small horticultural group, was the only one of the four without regular access to modern medical care at the time of data collection. In addition, the Gainj curve is based on a small sample (< 150 deaths) and therefore appears somewhat more “jagged” than the other examples.

Figure 2. Survival, force of mortality, and cohort PDF curves associated with the “typical” human age pattern of mortality.

Figure 3. An example of the bi-Weibull model of human mortality.

Figure 4. An example of the Siler model of human mortality.

Figure 5. An example of the mixed-Makeham model of human mortality. The broken and dotted curves show the force of mortality in the low- and high-risk subgroups, respectively, whereas the solid curve shows the aggregate-level force of mortality in the mixture as a whole. Although neither subgroup curve has a distinct juvenile component, the aggregate curve displays a decline in juvenile mortality reflecting selective mortality against the high-risk subgroup. As high-risk individuals are selected out of the population, the aggregate curve converges on the low-risk pattern.

FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4




FIGURE 5

