HAZARD CURVES AND LIFESPAN PROSPECTS

Kenneth W. Wachter
15 April 2003
Population and Development Review
Volume 29 (Supplement) 270-291.

When we see centenarians among us, singing, swimming, driving, and hear of supercentenarians, ten years or more beyond a hundred, it seems to auger well for the hope that the end of steady progress against mortality is not yet in sight. Can the fact that most people in developed societies die in their seventies and eighties be altogether resistant to change, if some people are living so much longer?

One may ask about specific causes, changing environments, social support, medical advances, but one must also ask about underlying biological potential. Humans are a special case, but also an example of the general case, molded like other species by principles of evolution, whose application to longevity is imperfectly understood. Comparing human survival at extreme ages with survival in other species, biodemography has been fostering a new view of biological potential.

Early achievements of biodemography are recounted in the volume *Between Zeus and the Salmon,* edited by Wachter and Finch (1997). The papers in this special issue review what has been learned since then, over the last five years, and what we stand to learn today. This essay looks at those questions through the prism of connections between the forms of hazard curves and the prospects for future progress against old-age mortality. Hazard curves are mathematical summaries of the dependence of mortality on age. Understanding how hazard curves are shaped into their shapes helps tell us how much resistance to extended survival may be programmed into our biology. The mathematics of hazard modeling intertwines with evolutionary theory and empirical research.

It is not the intention here to provide a synopsis of the contributions to this special issue. Rather, the aim is to take up one theme central to biodemography and draw selectively from the papers, reacting to suggestions that they offer and bringing out cross-cutting ideas and ties among them.

The opening contribution to this special issue, by Jim Carey (2003), sets the stage for our discussion as for the rest of the collection. In particular, Carey outlines the context within which lifespan comes to be viewed as an evolutionary adaptation. He reviews the ecological correlates of lifespan and the roles of the elderly in social species, and he gives a roster of behavioral and scientific developments likely to propel the biological future of human lifespans, under the four headings of healthful living, disease prevention and cure, organ replacement and repair, and aging arrest and rejuvenation. Such specifics provide an essential backdrop for the general consideration of constraints and opportunities in this essay. It is important to keep in mind that the "lifespan prospects" under discussion here are prospects for gradual, incremental extensions of lifespans across large populations, progress of the kind seen around the world over the last century. Drastic bio-engineering and visionary advances like those surveyed in Michio Kaku's (1997) *Visions,* however enticing, are beyond our chosen scope.

The mathematical representation of mortality to which we mainly refer, the "hazard function" or "hazard curve", is a graph of the force of age-specific mortality as a continuous function of age. Equal to the downward slope of the logarithm of the survivorship curve, it is discussed in detail in Horiuchi's (2003) contribution. Hazard curves that rise very steeply at advanced ages seem to suggest formidable obstacles to broad-based extensions of life. Hazard curves whose steepening tapers off at advanced ages seem to suggest more manageable challenges. The evolutionary theory of senescence, developed by biologists over the last half-century, encourages the idea of a strong connection between the shapes of hazard curves by age and their resistance to reduction over time. The theory and the connection will be described shortly. What is important to say first is that under this theory the same general principles operate to shape the hazard curves of other species as to shape our own.

In this essay, the word "tapering" is used as shorthand for reductions in the slope of the logarithm of the hazard curve at advanced ages. Taking logarithms puts measurements of slopes on a proportional basis and turns exponential curves into straight lines. Tapering hazard curves are often called mortality "plateaus", but hazards may taper so far as to bend downward, stretching the analogy with terrestrial plateaus.

So long as it was still widely believed that hazards in other species typ-

ically rose exponentially with age, free of tapering, such beliefs fostered the view that diminishing returns in the struggle against human mortality at older ages could not be far away. Then came the discoveries initiated by Carey et al. (1992) and Curtsinger et al. (1992) of tapering hazards at extreme ages in a number of other species, first Mediterranean fruit flies and *Drosophila,* later nematode worms and yeast. Data may be found in Vaupel et al. (1998) and in Figures 1, 2, and 3 of Horiuchi's (2003) contribution. Robine and Vaupel (2001) and the contribution of Robine and Saito (2003) show that the record human lifespans confirmed over the last few years demand human hazard functions that also taper beyond age 95 and level out or drop beyond 108.

Guiding forecasts

It is fitting that the meeting which inspired the papers in this special issue took place on the Greek island of Santorini, a home of ancient augury. Biodemography speaks to demographers in their role as augurers.

The regularity of past trends in itself can scarcely indicate whether extrapolation is justified or whether a change point, abrupt or gradual, lies around the corner. Structural understanding has to set the context. This truism applies palpably to predicting the course of mortality over the next half-century. Official forecasts almost all impose assumptions of diminishing returns to progress against mortality at older ages. In contrast, researchers impressed by the long-sustained pace of progress extrapolate unabating gains. An airing of the issues may be found in *Beyond Six Billion,* edited by Bongaarts and Bulatao (2000). The policy ramifications are obvious, as the different forecasts put very different constraints on the political decisions required now to preserve the solvency of social insurance systems and the economic health of nation states.

Choice between such forms of forecasts is not a technical issue. It is a matter of judgement, context, and plausibility, as are the ultimately subjective uncertainty bounds to be attached to all the forecasts. Here, in making possibilities plausible, biodemography can have its say.

The relevance of underlying biology to future trends in mortality, however, is not incontestable. It may turn out that features of the human environment, configurations of feasible technology and behaviors unique to humans will

govern trends over the next century, in ways that are hardly clarified by understanding our genetic heritage and our commonalities with other species. It may also be that concerns about pushing up against biological limits to longevity will come to seem like *fin de siècle* luxuries, if recurring acts of terrorism in the new century elevate mortality and unhinge the security and liberty on which our long run of progress has depended.

If civilization is successfully defended in the future, as it has been, often but not always, in the past, then several kinds of evidence do argue in favor of attention to fundamental biology. Notwithstanding all the human paraphernalia of houses, supermarkets, and hospitals, over a broad range the functional form of adult human mortality as a function of age across the globe in 2003 is the same as the functional form for worms, flies, deer, and most of the complex species that we know. From middle adulthood up to moderately old ages, exponential functions enshrined in the models of Benjamin Gompertz (1825) give serviceable approximations to hazard curves in species after species, as Finch (1990:13-23) describes. Exponential curves have straight lines for their logarithms, that is, lines with unchanging slopes. Progress over time has shifted the curves but preserved their shape, strongly suggesting that the general biological principles behind the common shape are still operative for humans.

The salience of underlying biology is also supported by the steadiness of the gains in developed countries over several generations. Oeppen and Vaupel (2002) have shown that "best-practice" life expectancy – e_0 in the country leading the world in e_0 in each given year – has been increasing quite nearly linearly at a rate of about three months per year over 160 years. Limits on life expectancy announced by demographers have been broken shortly after or even before they have been stated. Steady progress in the face of enormous differences in the character of advances during different stages of this long period argues for some deeply-rooted structure of opportunity.

One might have expected heroic medical breakthroughs like the introduction of antibiotics to show up sharply in the historical record. One might also have expected fitful starts and stops in the gains at different ages, as different diseases and conditions fatal at different ages successively yield to medical treatment and health-enhancing behaviors. In fact, however, age-specific mortality rates for different age groups have been declining in tandem, each with its own exponential rate of decline, well represented by Lee and Carter's

(1992) model. For all the "G-7" industrialized countries since at least 1945 (Tuljapurkar, Li, and Boe, 2000), these rates of decline have stayed quite nearly constant. The exponential functional form for the changes over time recapitulates Gompertz' exponential for changes over age. Age groups hold on to their shares in the broad march forward.

The regularities just described are features of a broad-brush picture. With close scrutiny, deviations can be seen. Horiuchi (2003) directs attention at upward bending that often appears in the logarithms of human hazard curves before tapering sets in, consistent with demographers' longstanding use of Makeham models in place of Gompertz models for refined work. Change points appear when Lee-Carter models are fit to longer stretches of historical data. Even after 1945 the shares of age groups in overall mortality reduction are not strictly constant. The trend in maximum recorded lifespan and the burgeoning number of centenarians described in the contributions by Wilmoth and Robine (2003) and by Robine and Saito (2003) signal an expanding share in overall progress for the oldest old. Reflecting developments over the last five years, many of the papers in this special issue press beyond broad regularities to examine nuanced differences in hazard curves across species and across time and age. But the existence of the regularities remains as a strong argument that long-term mortality forecasting is not solely a matter of second-guessing particular medical and health advances, but also a matter of understanding the dynamics of resource deployment and biological potential.

How many battles?

Classical evolutionary theories of senescence are recounted in this special issue by Orzack (2003:pp.2-5), who introduces "mutation accumulation", "antagonistic pleiotropy", and the "disposable soma". Overviews accessible to demographers may also be found in Wachter and Finch (1997), Finch (1990:36-42), Charlesworth (1992), and Rose (1991). Here we focus on aspects bearing on connections between hazard curves and potential progress. All the variant theories have the same starting point: The less individuals of a certain age contribute to reproducing and sustaining the next generation, the less the force of Darwinian natural selection will work to suppress mortality at that age. "A corollary", as Finch (1990:38) puts it, " to the proposition that the force of selection diminishes with age ... is that genet-

ically based manifestations of senescence should vary widely among species and within species."

This idea has pessimistic implications, when one asks how many separate battles will have to be fought by medicine or lifestyle modification to repeat at ever older ages the victories over mortality won at medium-old age. The implications are particularly striking for simpler forms of mutation accumulation theory. These posit the existence of large numbers of mutations each with small deleterious effects on survival at specific ages or over bounded ranges of ages. Mutations which affect young ages are weeded out of the population quickly by natural selection. Those affecting ages which reproduce less or contribute less or nothing to the next generation are weeded out less rapidly, or not at all.

All the while the weeding proceeds, more small mutations are being introduced into the population. Mutations accumulate to the point where there is a balance between the force of mutation and the force of selection. Mortality rates rise with age in proportion as the force of selection weakens with age. At any time, mutations with specific age effects are to be found scattered independently across the individuals in a population, combining in large numbers to dictate each individual's internal susceptibilities to causes of death.

In this picture, many, many things go wrong more and more at older ages. Countless battles loom against modes of failure dispersed across the population in different forms more and more numerous in later life. Hard-won progress at younger ages only presages harder challenges at older ages. For humans, further increases in life expectancy depend on progress at ages well beyond the ages which could plausibly have been contributing to reproductive fitness over evolutionary time. Without further provisos, one would expect mutation accumulation to drive mortality rates toward infinity at post-reproductive, post-nurturant ages. Such considerations abetted pessimistic mortality projections several decades ago. The discovery of hazard curves tapering at extreme ages brings the simple theory into doubt. As discussed below, current thinking focuses on whether revising the theory to accommodate tapering does or does not mean abandoning its essentially pessimistic implications.

Antagonistic pleiotropy is a process which could in principle drive up hazard curves at older ages even more dramatically than mutation accumu-

lation and spell even harder challenges to human lifespan extension. The theory posits mutations which promote reproduction or survival at some ages but incur a penalty at other ages, depressing reproduction or survival. Under natural selection, survival at old ages is readily sacrificed for benefits at younger ages. Genes with such pleiotropic actions could reflect the tradeoffs between investment in reproduction and investment in repair and maintenance described in Kirkwood's theory of disposable somata reviewed in Orzack's (2003:pp.2) contribution.

Theories of antagonistic pleiotropy, however, need not carry all the pessimistic implications of mutation accumulation theory. Antagonistic pleiotropy might act on hazard functions through the dynamics of mutation-selection balance as just described, or it might act through the gradual fixation of mutations in populations where the net effect on fitness of the mix of positives and negatives were positive. In the latter setting, heterogeneity could arise mainly from heterogeneous interactions between genes and the environment and life histories of individuals rather than from differing individual portfolios of mutations. A limited number of mutations with substantial effects rather than a large number of mutations with small effects might be determinative, softening the prospect of an endless array of future battles against mortality at ever older ages.

The large issue with which biodemographers are struggling today concerns the conclusions to be drawn from the observations of tapering hazard curves and the experimental results that have come in their wake. Are the new discoveries consistent with versions of the classical theories, or do they demand revisions which remove the main pessimistic implications?

Some biodemographers think that the tapering primarily reflects compositional change of the group of survivors in the presence of persistent heterogeneous frailty, without the need to appeal to detailed age-specific programs or far-reaching revisions to classical theories. Other biodemographers think that the tapering of hazard functions at the population level reflects programs of age-based change at the individual level which do need to be grounded in critically revised theories. The leading contender for a compositional explanation will be called "simple culling" in this essay. We discuss it first, and afterwards take up versions of revisionism.

Simple culling

The compositional explanations for tapering involve demographic selection in the presence of heterogeneity in probabilities of survival. In considering these explanations, it is important to distinguish age-free heterogeneity from age-bound heterogeneity. The frailty models of Vaupel, Manton, and Stallard (1979) and Vaupel and Yashin (1985) which have guided hazard analysis over the last twenty years posit "age-free" heterogeneity: differences from individual to individual are represented by a personal multiplier which operates independently of age, remaining fixed over the whole relevant segment of the life course. The variation of hazards with age, that is, the shape of the age profile of mortality, is the same for all individuals. The power of this framework is its simplicity.

One can also envision models for age-bound or age-structured heterogeneity in which individuals differ in the age profile of expected mortality. For example, individuals may differ in their inherent vulnerability to certain causes of death, and causes of death may each have a characteristic age pattern. Another example would be models in which the Gompertz slope parameter or Horiuchi's (2003:pp.3-6) Lifetable Aging Rate are significantly heterogeneous. In models with age-free and age-bound heterogeneity alike, more vulnerable individuals tend to die earlier, survivors are progressively selected from the less and less vulnerable, and this culling causes hazard rates to taper. However, unlike the age-free models, in the age-bound models culling is not the whole story: the shapes of the age profiles in the mixture modulate whatever tapering occurs, and prospects for reductions in mortality are interwoven with the determinants of the causes and profiles.

By "simple culling" we mean survival of the less vulnerable within the age-free, fixed frailty framework. If nothing but simple culling is behind the observed tapering of hazard rates in species studied so far, then the classical accounts of mutation accumulation and antagonistic pleiotropy, with their pessimistic overtones, are easier to maintain. No one doubts that something like simple culling plays some role. The question is whether simple culling suffices to explain all the tapering.

Data underlying this debate are found in Figures 1, 2, 3, and 6 of Horiuchi's (2003) contribution and in Vaupel et al. (1998). Tapering shows up in the graphs of the logarithms of hazard curves versus age as they bend away from straight lines. The log-hazard curves for medflies and yeast actually

bend over and drop downwards beyond the domains shown in Horiuchi's figures. For humans, the record lifespan of Madame Jeanne Calment, at 122, may be a hint that the human curve also ultimately drops.

The models in mind when simple culling is under discussion posit a Gompertz baseline hazard function with the same slope parameter for each individual. The frailty multiplier scales the level of the hazard function up or down by the same factor at all ages. The distribution of the frailty multiplier is most often taken from the gamma family of distributions, partly because this choice produces convenient closed-form formulas, but also because simple culling keeps the frailty distribution within the gamma family. However, gamma frailty with Gompertz hazards leads to a logistic model for aggregate survival, which forbids any drop in extreme age hazards. To account for drops observed with medflies and yeast and possibly with humans, a simple culling model has to have a small concentration of individuals more robust than any distribution from the gamma family would allow.

The early work in biodemography surveyed in *Between Zeus and the Salmon* established that tapering is reduced but not eliminated when mixed populations of flies or worms are replaced by strains that are (nearly) genetically identical. Experiments which attempt to control environmental heterogeneity have comparable outcomes. Ewbank and Jones (2001) citing and building on work by Yashin, DeBenedictis, and Vaupel, have been able with humans to make the variance in frailty an identifiable parameter within the framework of fixed-frailty models, and they estimate moderate values.

It takes a lot of heterogeneity to make hazard functions visibly bend or drop. When one fits a frailty model to a data set, one can compute how low the lifelong hazards have to be among the least frail few percent of individuals compared to the hazards for the median individual or the most frail few percent of individuals, if simple culling is to explain all the observed tapering. Many authors find these ratios too extreme to be plausible, in comparison with the measured effects of covariates or interventions. Mueller et al. (2002), for instance, obtain a ratio of 1/162 for the 2.5-th percentile compared to the the 97.5-th percentile in a recent experimental population of *Drosophila,* excessive, in their view, in comparison to measured effects of diet change.

Horiuchi (2003), in the crucial footnote 13 of his contribution, argues the opposite point of view. He likens the hazard ratios for percentiles from

frailty fits for a cohort of Swedish males to hazard ratios calculated by multiplying up marginal effects for common covariates estimated from fitting a proportional hazard model to recent U.S. sample data. However, multiplying marginal effects estimated from a model without a full set of interaction terms is probably producing overestimates of the ratios among the groups with the largest contrasts. The ratio of 1/32.8 which Horiuchi calculates from the 0.2-th and 99.8-th percentiles of the Swedish frailty fit seems rather to support the implausibility of simple culling as a full explanation of tapering. Further investigation should bring this question to resolution.

Reliability and Vitality

Horiuchi observes that fits of frailty models reveal much less heterogeneity on a relative basis for humans and other vertebrates than for insects and other invertebrates. He interprets this contrast, with good reason, as reflecting a greater degree of quality control in the formation of physiological components for organisms whose body plans and lifecourse strategies involve greater investments in individual offspring. Quality control is one aspect of reliability theory, whose importance for biodemography has been emphasized most especially by Gavrilov and Gavrilova (2001). Vaupel (2003), in his contribution, joins Horiuchi in turning to reliability theory to make sense of contrasts. The contrasts of interest are relative. Absolute levels of hazard curves only tell us that a day in the life of an insect is something like a year in the life of a human. Horiuchi standardizes across species by expressing time in units of one-hundredth of the modal lifespan, the age at which the greatest number of adults die. Vaupel standardizes across species by taking ratios of certain percentiles from the tails of the lifespan distributions. Both approaches bring out clearly the features of low spread and short upper tails that distinguish the lifespan distributions of vertebrates in general and humans in particular.

Vaupel (2003:pp.9) implicates three processes contributing to these patterns:

> "The combination of redundancy, repair, and low variability among individuals might be referred to as the 'reliability' of a species. Humans are a reliable species in a steady environment; medflies are an unreliable species in an uncertain environment ..."

Vaupel has a handy way of thinking about redundancy and its impact on hazard functions. Start with any survival curve. Draw two death dates independently at random from the survival curve, and build in redundancy by letting an individual die only at the latter of the two dates. This process generates a new survival curve out of the original one. Vaupel only treats the case of constant hazards, but his main finding is a generic property. This transformation reduces the ratio between the ages when 1% and when 10% of the population still survive. In other words, on a ratio basis, it shortens the upper tail. In general the effect is not independent of mortality level, as Vaupel shows it to be in the constant hazard case, but renormalizations of level can be shown to preserve the property under reasonable conditions.

What this redundancy transform does, however, is to shift the logarithm of the survival curve upward by an amount that is essentially constant out in the upper tail. It leaves the hazard function, which is the downward slope of this logarithm, unchanged in the tail. It necessarily does so, because any death age found far out in the tail will be, with high probability, the larger of any pair of independently drawn death ages. The same point holds for three-fold, four-fold, or higher-fold redundancy transforms. The ratio measure for tail percentiles is reduced, not because of effects at later ages, beyond the direct influence of natural selection, but rather because of reductions in early-age mortality. All three of the reliability processes that Vaupel highlights would seem, on the face of it, to operate primarily on early ages rather than on the upper tails themselves. Theories of senescence do stand in need of something to account for the prolongation of patterns from reproductive and early post-reproductive ages out into later post-reproductive ages, and considerations of reliability engineering do seem like logical elements for the purpose. But the pieces do not yet fit together.

Gavrilov and Gavrilova (2001) have developed specific formal models grounded in reliability theory to generate hazard curves Gompertzian over a range of ages and tapering toward a flat asymptotic plateau at extreme ages. Independent, identically distributed components with constant probabilities of failure over time are arranged together in blocks with built-in redundancy. A block fails only when all its components fail, like electric circuits wired in parallel. The organism dies when any one block fails, like electric circuits wired in series. In the most intriguing of their models, Gavrilov and Gavrilova take the numbers of components in the blocks to be independent Poisson random variables.

Evans and Steinsaltz (2003) correct the derivation of the hazard function in Gavrilov and Gavrilova (2001), give a closed-form formula, and check the implied shapes of hazard curves for a variety of parameter settings. Evans and Steinsaltz are sceptical of the model, because it generates a stretch of Gompertzian hazards only for quite specific choices of parameter values. Most choices lead to wholly non-Gompertzian shapes. As with string theory in physics, the lack of a canonical version is a sobering objection to models of this type. However, demographic reasons could be given for preferring certain ranges for the parameters, independent of the shapes to which they lead. The models are more ambitious than frailty models in seeking to explain rather than assume the involvement of exponential functions. The questions surrounding them deserve further review.

The connected component models of Gavrilov and Gavrilova are special cases of a general class of stochastic models for hazard functions which will be called "changing vitality models" in this essay. The prototype for this class is the Brownian motion model brought to the attention of demographers by Weitz and Fraser (2001) and also put forward by Anderson (2000). An idealized measure of a person's vitality is assumed to change randomly across the person's lifetime like a random walk or like its continuous analogue, a Brownian motion, drifting haphazardly downward. When this random path first hits or crosses zero, the person dies. The framework is reminiscent of Strehler and Mildvan (1960), but unlike Strehler and Mildvan's model, an exponential response function is not arbitrarily imposed.

The models of Gavrilov and Gavrilova take on a similar structure, when the counts of intact components in each block of the system are treated as a multivariate version of the idealized vitality measure. Evans and Steinsaltz (2003) define a broad class of such models in which vitality changes according to a Markov Process. They prove that tapering hazard functions converging to constants at extreme ages are a generic phenomenon common to all models in this class. The hazard function levels out because the distribution of vitalities among survivors converges to a fixed distribution, for much the same mathematical reason as convergence of age structures in stable population theory. As Evans and Steinsaltz (2003:3) say, these models help us envision an alternative form of selectivity:

> " ... the mortality rate stops increasing, not because we have selected out an exceptional subset of the population, but because

> the condition of the survivors is reflective of their being survivors, even though they started out the same as everyone else. ”

These changing vitality models have strong inbuilt tendencies toward producing tapering hazards. The challenge for modelers is to construct special cases with hazard functions that do not taper too early in the age span. The purpose of all these stylized models, at the present stage of inquiry, is to serve as guides to thinking. Although many variations are possible, the underlying logic of changing vitality models has optimistic overtones, as far as opportunities for future extensions in lifespan are concerned. Improvements in mortality at medium old ages would be evidence of upward shifts in the vitalities of individuals which would be expected to confer benefits on later ages as well.

Revising theories of senescence

If compositional explanations do not account fully for the observed tapering in hazard functions, then classical evolutionary theories of senescence need to be revised. The theories rest on assumptions about the existence and prominence of genes with certain kinds of actions on survival at older ages, and they also rest on mathematical models for translating specifications of gene effects into predictions of hazard functions. Early evidence about kinds of genes that actually occur will be discussed in the concluding section of this essay. This section focuses on the interplay between specifications of theory and shapes of hazard curves.

Revising theories of senescence is a tricky matter, if the many successes of the theories are not be thrown out in the process of repairing the limitations. Most deaths after all occur in the exponentially rising segment of the hazard curve, long treated by the classical theories, not in the tapering segment, where the need for revisions appears.

One strategy for explaining tapering is to argue that mutations with effects exclusively at a range of older ages are rare and aberrant. Pletcher and Curtsinger (1998) give an illuminating discussion. The large number and variety of such genes that are required to drive mutation accumulation theory may be a chimera. Few genes may be expressed only first in later life, and few genes that act on later survival may fail to have some role in development or in somatic maintenance during reproductive adulthood. Debate over the

nature of the genes that might figure in mutation accumulation goes all the way back to Haldane, as discussed by Finch (1990:37), and remains an active subject for research.

It is easy to see that tapering could be explained, if mutations with distinctively bad effects on survival at late ages typically always had at least some bad effects at early ages. Such reinforcing pleiotropic effects would produce tapering, inasmuch as the force of selection would be buoyed up by the early effects and would not fall as rapidly as it otherwise would at older ages. The reverse of antagonistic pleiotropy, reinforcing effects are called "positive pleiotropy" because correlations are positive, but this technical term has misleading connotations when the correlated effects themselves are negative. The non-technical discussion here will keep reinforcement in the foreground. Reinforcing pleiotropic effects essentially introduce an extra penalty term into calculations of fitness, a plausible ingredient of a wide range of formal models for tapering, as in Wachter (1999).

Charlesworth (2001), aligning himself with the idea of reinforcing pleiotropic effects, regards the idea as a minor revision to mutation accumulation theory. It is, however, major in its demographic consequences. In contrast to classical mutation accumulation, an assumption of possibly small but ubiquitous reinforcing pleiotropy would offer escape from the picture of an endless succession of new battles looming for future progress against old-age mortality. Evolution would have cooperated with the goals of medicine, weeding out contributors to late-age failures thanks to their younger-age concomitants.

The more pervasive any patterns of reinforcing pleiotropy, the more strongly interconnected should be the challenges of mortality at medium-old and at older-old ages. A picture of large numbers of largely independent accumulating mutations could be retained, if the reinforcing pleiotropic effects were small and diffuse. Stronger and more systematic reinforcing pleiotropy would bolster the case for optimism about lifespan extension. A few good steps forward might have many good consequences across the age range.

Antagonistic pleiotropy competes with mutation accumulation for the allegiance of experimentalists. The weight of evidence from selection studies like Rose et al. (2002) swings to and fro. Theories of antagonistic pleiotropy are currently handicapped by the lack of a viable mathematical model to connect the genetic processes with hazard curves of the forms actually observed. The hope would be to represent the hazard curve and an associated age-

specific fertility schedule for a population as the limiting equilibrium state of a Markovian stochastic process. The steps or transitions for the process would correspond to mutations going to fixation in a population in the face of genetic drift and natural selection. Mueller and Rose (1996) put forward versions of such a model devised to account for mortality plateaus, that is, for tapering hazard curves. The proof in Wachter (1999) shows that the proposed models fail. The models do not have the intended limiting states.

Mueller et al. (2003:25) take the position that it is sufficient for their model to produce the intended tapering hazard curves on a purely temporary basis, as transient states. The model is started with a stylized initial state, in their case with a flat adult hazard curve, and after a while it passes through states of the intended form before going on to diverge from the intended form and converge to a limiting state. The problem with this approach is the initial state. For the transient states to be meaningful, evolution would have to have established the specific stylized initial state as the hazard curve at a specific time back in the past. An account would therefore be needed of how evolution arranged such a specific starting state in the first place and why the pleiotropic model is meant to apply only after this instant of time and not before. It would also be necessary to explain, for any species to which the model is meant to apply, how we happen to be observing the species just at the right period of time after the establishment of the specific initial state.

Mueller, Rose, and their coauthors contend that their transient states last for millions of generations. The parameters of the model can indeed be set to make the whole process run arbitrarily slowly, slowly arriving at the transient states of interest from other transient states, and slowly departing from them toward other transient states on the way to the limiting states. The parameters can also be set to make the process run quickly. The point of importance is that the transient states of interest are ephemeral relative to the timescale of the process. For Markovian models of pleiotropy, recourse to transient states is not a tenable position.

It is therefore a continuing priority to construct new Markovian models incorporating antagonistic and reinforcing pleiotropy which can successfully generate tapering hazard curves as limiting states. Ideas for the design of such models are set out in Wachter (1999:10547).

As Mueller et al. (2003:25) go on to point out, environmental conditions could change over time periods that are shorter than the timescales for the

establishment of equilibria in models of this kind. For arguments based on transient states, this situation would redouble the difficulties associated with the establishment of a specific initial state at a specific initial time. It would also mean that a model incorporating environmental fluctuations would have to be spelled out in order to derive the consequences of pleiotropic processes.

The broader need for ambitious modeling of environmental fluctuations is taken up by Orzack (2003) in his contribution. One set of effects, which Orzack discusses in detail, arise from the way in which symmetric random variability in components of a population projection matrix produce non-symmetric distributions for rates of growth and associated measures of fitness. More complex effects necessarily arise when temporally correlated fluctuations interact with generational renewal.

From a long-term perspective, the inherent randomness of demographic events in combination with random variations in resources and environments can drive populations to extinction. Orzack advocates extensions of the traditional reckoning of fitness to take account of such impacts on the gene pool. From a short-term perspective, the environments in which organisms find themselves are partly of their own making and choosing. In his contribution, Mangel (2003) presents a remarkable example in which Pacific rockfish adjust their metabolic rates and exposure to oxidative damage by choices of ocean depth across their lifetimes. The subjects of both these contributions take them well beyond the scope of the present essay, but they represent initiatives likely to figure prominently in the future of biodemography.

Mutation Accumulation

Charlesworth's (2001) formulation of mutation accumulation theory which has already been mentioned is of particular demographic interest, going beyond the issue of reinforcing pleiotropy and its role in tapering. He spells out a specific explanation for the Gompertzian shape of hazard curves with a wealth of testable consequences.

Charlesworth posits high background levels of extrinsic, age-independent hazards in the wild. The constancy of hazards over age, in a context of near-zero population growth, produces stable age pyramids with sizes of older age groups decreasing exponentially with age at a rate equal to the constant hazard rate. The simple version of the theory makes the assumption (which

can be relaxed to some extent) that there is a steady contribution over age for relevant age groups to the production of the next generation either through reproduction or through provisioning, rearing, and protection. For each given age, mutations are assumed which raise hazards only beyond that age, by small amounts, either in a window of ages starting at the given age or at all ages from the given age upward. The density of mutations per unit time is small and uniform over age.

The mutations tend to add a small perturbation of extra mortality onto the high constant background mortality. If the perturbation is small enough, the age pyramid nearly retains its exponential shape. According to the mathematics of mutation-selection balance, the exponential shape of the age pyramid is impressed onto the shape of the additional term in the hazard function. When populations are taken out of the wild into experimental settings, laboratories, zoos, or, as with humans, into the modern world, the constant background contribution to the hazard function is largely removed, and what remains is the Gompertzian addition.

Mutation accumulation theory in this form has power. It offers to account at a stroke for the ubiquity of Gompertzian hazards across a wide range of species, body plans, lifespans, and circumstances. It provides for the exponential function to enter directly from first principles rather than from *ad hoc* assumptions. It reaffirms the association between the level of extrinsic mortality and the pace of senescent mortality, a widely-verified general prediction of the evolutionary theory of senescence.

This special issue offers one of the first opportunities to test Charlesworth's proposal by confronting predictions with data. The contribution of Gaillard and coauthors (2003) reports estimates from a dozen wild populations of deer, sheep, moose, and other large herbivorous mammals. Gompertz hazard models have been fit to data on lengths of life collected from long-term monitoring of individually marked animals. Senescent mortality as measured by the Gompertz slope parameter is substantial. After the onset of adulthood around age 2, the yearly percentage rise in female hazard curves ranges from 11% to 40%. Initial mortality is low enough to allow many individuals to survive to ages when senescent mortality is dominant. In Gompertz models, the hazard rate at the modal age of death equals the Gompertz slope parameter. In all but one of these populations, at the modal age of death, the background level of mortality comes out to be less than half the overall

level of mortality, and in four cases less than one part in six.

The direct evidence reported by Gaillard and his co-authors for large herbivorous mammals showing a pervasive role for senecent mortality in the wild among large mammels is consistent with circumstantial evidence for birds reviewed by Ricklefs and Scheuerlein (2003:pp.8) in their contribution. It also accords with the discovery by Bronikowski et al. (2001) of substantial and comparable values of the Gompertz slope parameter for wild (Kenyan) and captive (U.S.) populations of the same species of baboon, set in context by Altmann and Alberts (2003).

The most appealing feature of Charlesworth's proposal, the explanation of the exponential functional form, rests on the claim that mortality in the wild is dominated by extrinsic, age-independent background mortality. This feature is in conflict with these data. The theory further predicts that the background level of extrinsic mortality in the wild over evolutionary time can be read off from the present value of the Gompertz slope parameter in captive or protected populations. For humans one would be talking about values of e_1 on the order of a dozen years. For the mammals described by Gaillard and his co-authors, in most cases the predicted levels of background mortality would not only be well above the observed levels of background mortality, but above the observed levels of overall mortality under present-day conditions in the wild.

It may be that "the wild" is not what it used to be. Under Charlesworth's proposal, the time periods for which the background level of mortality is relevant are constrained by the appeal to mutation-selection balance. The genetic effects responsible for the Gompertzian shape have to have been continually and (as evolutionary time goes) relatively recently renewed. But this constraint is not a tight one. The analysis so far is restricted to a handful of vertebrate species, and sample sizes are not very large. Although practical difficulties stand in the way of measuring lifetimes for insects or other invertebrates in nature, a broader collection of examples would be invaluable. On the whole, however, the early evidence is not very encouraging for this theoretically appealing version of mutation accumulation theory.

Social Support

Jim Carey (2003:pp.5) lists two clusters of factors favoring extended longevity of species. Resource-based factors are scarcity and uncertainty.

Kin-based factors are parental care and sociality. For the evolution of hominid lifespans out of primate lifespans, social support and exchange must have played large roles. Informally, it has always been recognized that survival into age groups which contribute to the provisioning and nurturing as well as the procreating of the next generation will be favored by natural selection. Recent formal models seek to quantify such components of selective pressure.

In their contribution Kaplan and Robson (2003) present a model for evolutionary tradeoffs in social species in which transfers of resources occur between age groups within a population. Mathematical proofs are given in Kaplan and Robson (2002). In unpublished work, Ronald Lee has developed further related ideas. In Kaplan and Robson's model, the growth rate for a stable population is chosen, via a Support Equation, by requiring that the difference between age-specific production and age-specific consumption averaged over all the age groups in the stable population must equal zero at all times. Fertility rates are assumed to adjust in some way to make Lotka's Equation hold. The demands of dependent age groups are covered by excess net production from any other age group. These demands comprise the physiological investments made in creating progeny as well as rearing them, including a cost reckoned at age zero of investment in a brain of a given size.

In this model, the age-specific production function and the initial costs of progeny depend explicitly on a parameter K and the age-specific survival function depends implicitly on this parameter. The authors interpret K as an index of investment in embodied capital including the size of the brain which governs long-range returns on skills through the balance of production and consumption. Some parametric assumptions and an optimization argument enable them to derive a U-shaped profile for the hazard curve from infancy to old age and to prove that growth is maximized at an optimal intermediate value of K associated with lower than maximal fertility.

Kaplan's and Robson's approach has appealing features. Their Support Equation puts calculations of selective advantage into a context of constrained growth. It sets up a pattern of tradeoffs between juvenile and senior survival that generate relationships between the shapes of hazard functions at both ends, suggesting evolutionary interconnections between them.

Kaplan and Robson primarily seem to envision mutations going to fixation very gradually fixing K at its optimum and fixing the associated shape of the

survival curve and the intrinsic rate of growth. Insofar as K is taken to govern secular changes in brain size and somatic organization, this interpretation is natural. However, in their formalism, K is a free parameter that can be adjusting all sorts of strategic allocations. The tuning of K could easily be partly behavioral, operating rapidly and enabling populations of primates or hominids to adjust their somatic investments and resulting age schedules of survival to prevailing resource constraints. Natural selection, for its part, could be gradually transforming the shape of the dependence on K through mutations which affect, for instance, the costliness of lowered death rates. The response profiles rather than just the optimal values would be evolving. With this interpretation, calculation of selective advantage within the model give a rationale for certain kinds of pleiotropy.

The Kaplan-Robson framework, through the Support Equation, is tied to stable populations, and this feature places limits on the range of evolutionary contexts to which it applies. The framework does not help to account for the shapes of hazard curves at extreme ages. Significant provisioning does not go on indefinitely. However, in its broad outlines, the approach does lend support to an optimistic view about the biological potential for future progress against old-age mortality. If evolution equipped social species like our own with the ability to shift across a continuum of survival schedules in response to nutritional and environmental conditions, then gains along a few dimensions could imply mortality reductions over many ages. The picture of endless looming battles tends to fade away. Accounts that emphasize the co-evolution of earlier-age and late-age mortality increase the hope that the successes already achieved at earlier ages may be repeatable at later ages too.

Genes for instance

Along with fresh thinking about the mechanisms that must shape the statistical regularities of hazard functions, the bounty of experimental knowledge about the genomes of model organisms with respect to longevity has been growing. In this issue Larry Harshman (2003) gives a comprehensive guide to findings for *Drosophila melanogaster* treating developments since Finch's (1990) authoritative survey of the broader field.

The first genes associated with longevity that are coming into view are a little like the first galaxies glimpsed in small telescopes. They serve as

instances. There are now instances of genes with effects of certain shapes, as regards hazard functions, just as once there were instances of barred spirals or giant elliptical galaxies. A morphology of effects will take time, and the early discoveries may be atypical, easier to spot than cases that may prove to predominate. It is too early to try to constrain models for hazard curves on the basis of identified mutations, selection studies, and analyses of quantitative trait loci for lifespan. However, the roster of instances is likely to expand.

Recently discovered induced mutations in genes in *Drosophila* discussed in detail by Harshman (2003:pp.18-20) make an intriguing contrast. The mutations can increase flies' lifespans substantially. The *Indy* gene (Rogina et al., 2000) is involved in metabolic processes, and Harshman (2003) sums up by saying that for the *Indy* gene "the mechanism of life span extension could be caloric restriction". On the other hand, the *InR* gene (Tatar et al., 2001) is involved in insulin reception and synthesis of juvenile hormones, and the mutations at issue, while reducing size and making female flies sterile, produce no measured change in metabolic rate. Enhanced resistance to various kinds of stress is likely part of their advantage.

It seems harder to draw analogies with humans for genes like *Indy* than for genes like *InR.* Whatever the benefits of caloric restriction may be for some people today, over the course of human history vast increases in caloric intake have been closely coupled to gains in longevity, as Robert Fogel (cf. Fogel and Costa, 1997) has shown in rich detail. Human bodies appear to have been pre-adapted to make ready use of diets much richer than can have prevailed over evolutionary time. In itself, that is something of a puzzle. The puzzle would be worse, if inbuilt genetic tuning mechanisms were for turning down longevity in conjunction with turning up caloric consumption. On the other hand, it is reasonable to expect genetic tuning mechanisms which turn up longevity in conjunction with turning up capacities for resisting stress and privation and outlasting environmental fluctuations. In a context of homeostatic population regulation, partial reductions in fecundity might have had little cost and even some benefit.

Many questions remain open pending further data. As Harshman (2003:pp.7-8) says, some experimenters find selection primarily affecting the Gompertz level parameter, others the Gompertz slope parameter. The early quantitative trait loci for longevity along with the selection experiments suggest that

there is wide variability keyed to environmental conditions.

An instance of genetic effects particularly pertinent to demography comes from the experiments of Sgro and Partridge (1999) discussed by Harshman (2003:pp.9-10) with lines of *Drosophila* artificially selected for longevity. For females, the hazard curve for the selected lines show both a Gompertzian exponential segment and a tapering late-age segment shifted to the right toward later ages as compared to the hazard curve for a control population. The selected lines also show lower early-age reproduction. When reproduction is suppressed in both selected and control lines, by each of two methods, the shift vanishes. Sgro and Partridge see a delayed direct cost of reproduction likely reflected in the hazard curves. As discussed by Wachter (1999), such push-back mechanisms are the sorts of pleiotropic effects needed in mathematical models for generating hazard curves consistent with observations.

A particular gene in a particular model organism may tell us something general about feasible genetic effects on hazard functions, or it may tell us something specific about the adaptations of that organism. Experimental results have to take their place within our broad knowledge of natural history. In their contribution, Ricklefs and Scheuerlein (2003:pp.17-18) give an analysis of variance, partialling out variability in lifespan by taxonomic levels. The largest entries in the analysis of variance table for fitted hazard curve parameters are found at the level of genera, rather than at the broader level of orders and families. The same body plan can accommodate large ranges of variation, suggesting that it is relatively easy for evolutionary change to move along a continuum between shorter and longer lives. That would suggest genetic organization with some high-level control and synchronization of effects on hazard curves.

It is important, as Hillard Kaplan has emphasized, not to conflate such genetic capacity for moving along a continuum with abilities of members of a species to adapt less or more flexibly to a range of environments. The mathematical models with which researchers are working probably put excessive emphasis on a picture of hazard curves as outputs of genetic programs, rather than on a picture of individual members of species forging their own hazards through their lives in the rough and tumble of environmental challenges. Although distinct phenomena are involved, both genetic plasticity and behavioral flexibility support prospects for coordinated progress.

Onward

Many of the questions addressed in this special issue remain open, and research is in progress to settle them. The next few years should bring more clarity. We shall be learning whether human hazard functions do drop at extreme ages. We shall be building consensus about the plausible efficacy of simple culling and the extent of tapering it may explain. We shall be sorting out the generic implications of reliability theory from the special features of specific models. We shall be seeing how far one can take ideas from mutation accumulation theory and mutation selection balance while remaining faithful to empirical evidence. We shall be coming to grips with the potential and the limitations of formal models for social support. We shall be giving fuller attention to the evolutionary implications of varying environments, population extinction, and adaptive lifecourse learning. We shall be finding out how to design viable mathematical models for the shaping of hazard functions by processes of antagonistic and reinforcing pleiotropy.

One major need is for gathering more data for more species on aging in the wild. The patient investments in long-term monitoring of populations of primates and some other large mammals are paying off handsomely for aging studies. Stylized "facts" which have long channeled thinking are turning out to be non-facts. But our information is very limited in terms of the range of organisms observed and measured under natural conditions. Such data are incredibly difficult to obtain, but they are becoming indispensable to the interpretation of laboratory experiments and the articulation of evolutionary theory.

Many points of contact exist between the biodemography of longevity and the biodemography of fertility. New developments in that field are summarized in *Offspring: Human Fertility Behavior in Biodemographic Perspective* edited by Wachter and Bulatao (2003). For humans, extended longevity opens up alternatives for the restructuring of traditional lifecourse stages, as discussed in the contribution of Lee and Goldstein (2003), blending the concerns of biodemography with economic and social demography. Cooperation between specialties should facilitate a fuller integration of life-course perspectives.

We may expect the next few years to bring a large collection of genetic studies with model organisms and with humans, including systematic studies of age profiles of gene expression. The picture is likely to be very complex,

and it may be quite a while before generalizations emerge. Among humans, one promising strategy is to concentrate on exceptional longevity. As mortality progressively takes its toll on each human cohort and leaves a rarer and rarer subset of survivors, it may distill away some of the complexity. It may concentrate alleles of certain genes to the point where their roles are easier to identify. It may accentuate dimensions of heterogeneity that can be subject to measurement and observation. The more we learn about the upper tails of the lifetime distribution, the better we can judge whether the oldest-old among us today are truly harbingers of widespread future gains.

Many disciplines and many strands of research come together in the biodemography of longevity. They are well represented in this special issue. Broad regularities seen in patterns of human mortality over age and time plausibly reflect structures of biological potential honed by evolution and shared with other species. Many of the new developments in data and theory examined in this essay appear to have optimistic overtones, with respect to the permissiveness of our biological heritage. A wide range of ideas and possibilities, however, remain in play. The work reviewed here suggests but does not yet establish connections between the shapes of hazard curves and the prospects for continued progress, leaving us, as scientists and mortals, in suspense.

References

Altmann, Jeanne, and S. Alberts. 2003. "Fertility, offspring care, and ontogeny in primates," in K. W. Wachter and R. Bulatao (eds.), *Offspring: Human Fertility Behavior in Biodemographic Perspective.* Washington, D.C.: National Academy Press.

Anderson, James J. 2000. "A vitality-based model relating stressors and environmental properties to organism suvival," *Ecological Monographs* 70: 445-470.

Bongaarts, John, and R. Bulatao. 2000. *Beyond Six Billion: Forecasting the World's Population.* Washington, D.C.: National Academy Press.

Bronikowski, Ann M., S.C. Alberts, J. Altmann, C. Packer, K.D. Carey, and M. Tatar. 2002. "The Aging Baboon: comparative demography in a nonhuman primate," *Proceedings of the National Academy of Sciences* 99: 9591-9595.

Carey, James R. 2003. "Lifespan: A Conceptual Overview," *Population and Development Review,* to appear.

Carey, James R., P. Liedo, D. Orozco, and J. W. Vaupel. 1992. "Slowing of mortality rates at older ages in large medfly cohorts," *Science* 258: 457-461.

Charlesworth, Brian. 1994. *Evolution in Age-Structured Populations.* Second Edition. Cambridge: Cambridge University Press.

Charlesworth, Brian. 2001. "Patterns of Age-specific Means and Genetic Variances of Mortality Rates Predicted by the Mutation-Accumulating Theory of Ageing," *Journal of Theoretical Biology* 210: 47-65.

Curtsinger, James W., H. Fukui, D. Townsend, and J.W. Vaupel. 1992. "Demography of genotypes: failure of the limited life-span paradigm in *Drosophila melanogaster," Science* 258: 461-463.

Evans, Steve, and David Steinsaltz. 2003. "Markov Mortality Models: Implications of Quasistationarity and Varying Initial Distributions." Working Paper, U.C. Berkeley, www.demog.berkeley.edu.

Ewbank, Douglas, and N. Jones. 2001. "Preliminary Observations on the Variations in the Risk of Death," paper presented at the Annual Meetings of the Population Association of America, Washington, D.C.

Finch, Caleb. 1990. *Longevity, Senescence, and the Genome.* Chicago: University of Chicago Press.

Finch, Caleb, and T. Kirkwood. 2000. *Chance, Development, and Aging.* Oxford: Oxford University Press.

Fogel, Robert, and D. L. Costa. 1997. "A theory of technophysio-evolution, with some implications for forecasting population, health care costs, and pension costs," *Demography* 34: 49-66.

Gaillard, Jean-Michel, A. Loison, M. Festa-Bianchet, N. Yoccoz, and E. Solberg. 2003. "Ecological Correlates of Life Span in Populations of Large Herbivorous Mammals," *Population and Development Review,* to appear.

Gavrilov, Leonid, and N. Gavrilova. 2001. "The Reliability Theory of Aging and Longevity," *Journal of Theoretical Biology* 213: 527-545.

Gompertz, Benjamin. 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* 115: 513-585.

Harshman, Lawrence G. 2003. "Life Span Extension of *Drosophila melanogaster:* Genetic Approaches," *Population and Development Review,* to appear.

Horiuchi, Shiro. 2003. "Interspecies Differences in the Life Span Distribution: Humans versus Invertebrates," *Population and Development Review,* to appear.

Kaku, Michio. 1997. *Visions.* New York: Anchor Books.

Kaplan, Hillard, K. Hill, J. Lancaster, and A.M. Hurtado. 2000. "A Theory of Human Life History Evolution: Diet, Intelligence, and Longevity," *Evolutionary Anthropology* 9: 156-184.

Kaplan, Hillard, K. Hill, A.M. Hurtado, J. Lancaster, and A. Robson. 2003. "Embodied Capital and the Evolutionary Economics of the Human Lifespan," *Population and Development Review,* to appear.

Kaplan, Hillard, and A. Robson. 2002. "The Emergence of Humans: The Coevolution of Intelligence and Longevity with Intergenerational Transfers," *Proceedings of the National Academy of Sciences, USA,* 99: 10221-10226.

Lee, Ronald D. and L. Carter. 1992. "Modeling and forecasting the time series of U.S. mortality," *Journal of the American Statistical Association* 87: 659-671.

Lee, Ronald D., and J. R. Goldstein. 2003. "Rescaling the Life Cycle: Longevity and Proportionality," *Population and Development Review,* to appear.

Mangel, Marc. 2003. "Environment and Longevity: The Demography of the Growth Rate," *Population and Development Review,* to appear.

Mueller, Laurence, and Michael Rose. 1996. "Evolutionary theory predicts late-life mortality plateaus," *Proceedings of the National Academy of Sciences, USA* 93: 15249-15253.

Mueller, Laurence, Mark Drapeau, Curtis Adams, Christopher Hammerle, Kristy Doyal, Ali Jazayeri, Tuan Ly, Suhail Beuwala, Avi Mamimdi,

and Michael R. Rose, "Statistical tests of demographic heterogeneity theories," *Experimental Gerontology,* in press.

Oeppen, James, and James W. Vaupel. 2002. "Broken Limits to Life Expectancy," *Science* 296: 1029-1031.

Orzack, Stephen. 2003. "How and Why a Species Lifespan Evolves," *Population and Development Review,* to appear.

Pletcher, Scott, and J. Curtsinger. 1998. "Mortality plateaus and the evolution of senescence: Why are old-age mortality rates so low?" *Evolution* 52: 454-464.

Promislow, Daniel, M. Tatar, A. Khazaeli, and J. W. Curtsinger. 1996. "Age-specific patterns of genetic variance in *Drosophila melanogaster.* I. Mortality," *Genetics* 143: 839-848,

Ricklefs, Robert, and A. Scheuerlein. 2003. "Life Span as a Component of Avian Life Histories," *Population and Development Review,* to appear.

Robine, Jean-Marie, and J. W. Vaupel. 2001. "Supercentenarians, slower ageing individuals or senile elderly?" *Experimental Gerontology* 36: 915-930.

Robine, Jean-Marie, and Y. Saito. 2003 "Acceleration of Survival Beyond Age 100: The Case of Japan," *Population and Development Review,* to appear.

Rogina, Blanka, R. Reenan, S. Nilsen, and S. Helfand. 2000. "Extended life-span conferred by cotransporter gene mutations in *Drosophila," Science* 290: 2137-2140.

Rose, Michael R. 1991. *The Evolutionary Biology of Aging.* Oxford: Oxford University Press.

Rose, Michael R., Mark Drappeau, Puya Yazdi, Kandarp Shah, Diana Moise, Rena Thakar, Casandra Rauser, and Laurence D. Mueller. 2002. "Evolution of Late-Life Mortality in *Drosophila melanogaster," Evolution.* 56: 1982-1991.

Sgro, Carla, and L. Partridge. 1999. "A delayed wave of death from reproduction in *Drosophila," Science* 286: 2521-2524.

Strehler, B., and A. Mildvan. 1960. "General theory of mortality and aging." *Science* 132: 14-21.

Tatar, Marc, A. Kopelman, D. Epstein, M.-P. Tu, C.-M. Yin, and R.S. Garofalo. 2001. "Mutant *Drosophila* Insulin Receptor Homolog That Extends Life-Span and Impairs Neuroendocrine Function," *Science* 292: 107-110.

Tuljapurkar, S., N. Li, and C. Boe. 2000. "A universal pattern of mortality decline in the G-7 countries," *Nature* 405: 789-792.

Vaupel, James W. 2003. "Post-Darwinian Longevity," *Population and Development Review,* to appear.

Vaupel, James W., J. Carey, K. Christensen, T. Johnson, A. Yashin, N. Holm, I. Iachine, A. Khazaeli, P. Liedo, V. Longo, Zeng Yi, K. Manton, and J. Curtsinger. 1998. "Biodemographic Trajectories of Longevity," *Science* 280: 855-860.

Vaupel, James W., K. Manton and E. Stallard. 1979. "The impact of heterogeneity in individual frailty on the dynamics of mortality," *Demography* 16: 439-454.

Vaupel, James W., and A. Yashin. 1985. "The deviant dynamics of death in heterogeneous populations," *Sociological Methodology* 15: 179-211.

Wachter, Kenneth W. 1999. "Evolutionary demographic models for mortality plateaus," *Proceedings of the National Academy of Sciences, USA* 96: 10544 - 10547.

Wachter, Kenneth W. and Caleb Finch (eds.) 1997. *Between Zeus and the Salmon: The Biodemography of Longevity.* Washington, D.C.: National Academy Press.

Wachter, Kenneth W., and R. Bulatao (eds.) 2003. *Offspring: Human Fertility Behavior in Biodemographic Perspective.* Washington, D.C.: National Academy Press.

Weitz, Joshua, and H. Fraser. 2001. "Explaining Mortality Rate Plateaus," *Proceedings of the National Academy of Sciences, USA* 98: 15383 - 15386.

Wilmoth, John, and J.M. Robine. 2003 "The World Trend in Maximum Lifespan," *Population and Development Review,* to appear.