Pleiotropic Theories of Aging and their Empirical Foundation II: Demography

Joshua Mitteldorf

Dept of Biology, Leidy Laboratory
University of Pennsylvania
Philadelphia, PA 19104-6018
josh@mathforum.org

Abstract

Pleiotropic theories for the evolutionary origins of senescence have been ascendant for forty years, and it is not surprising that interpreters of demographic data seek to frame their results in this context. But without alternative hypotheses, it can sometimes be unclear to what extent the same data might equally well have been viewed as antagonistic. We cite as examples two prominent studies that have been taken as support for the prevailing theories, and show that the same data may be open to other interpretations.

1. Introduction

Senescence has posed a dilemma to evolutionary theorists for more than a century (Weismann 1889; Medawar 1952; Williams 1957; Hayflick 1994; Clark 1999) A trait that is so ubiquitous and which has been conserved over evolutionary eons would likely be regarded as an adaptation; but natural selection works for traits that contribute positively to fitness, and senescence, by definition, has only a negative impact on individual fitness. Senescence may contribute positively to the fitness of populations or entire ecosystems; but natural selection on the basis of such dispersed benefits is regarded with suspicion by most evolutionary theorists (Williams, 1966; Maynard Smith 1976). Hence, the only viable evolutionary theories today regard senescence as an evolutionary epiphenomenon. The most popular theories postulate that senescence is a side-effect of the strength of natural selection for fertility.

1.1 Evolutionary theories of aging

Pleiotropy refers to one gene that has multiple effects. Thus pleiotropic theories of aging postulate the existence of fitness benefits, typically fertility or acceleration of the reproductive schedule, that are coupled to life span limits and decline of vigor with age. These theories have their origin in Medawar's (1952) observation that the force of natural selection declines with age for two reasons: first, that offspring produced late in life contribute less strongly to the exponential rate of increase (or reproductive value, r); and second because incidental mortality

causes the probability of realizing age-specific reproductive potential to decline monotonically. Alleles associated with senescence may be adopted into the genome as part of the mutational load (because selective force associated with advanced age is weak) or as a side-effect of strategies that enhance fertility or survival early in life. The former idea has been codified as the theory of Mutation Accumulation (MA) (Edney and Gill 1968), and the latter as the theory of Antagonistic Pleiotropy (AP). Within the second category, two flavors can be distinguished: In explicit pleiotropy, as originally formulated by Williams (1957), there are genes that have survival value to the young individual, or that enhance fecundity directly, but incur a cost by contributing to the rate of senescence. The second flavor of pleiotropy, called the Disposable Soma Theory (DS) (Kirkwood 1977) derives from the observation that metabolism may be optimized under genetic control to allocate scarce resources in a way that maximizes fitness, which is assumed to depend both on present reproduction and somatic repair and maintenance. This will always lead to an imperfect performance of the latter tasks, which leads to an accumulation of damage over the lifetime of the organism, the outward manifestation of which is the constellation of symptoms we identify as senescence.

We will say nothing further about the non-pleiotropic MA theory, because it has been contradicted by genetic evidence in recent years (Kenyon, 2001; Promislow et al, 1996).

1.2 Relevance of demographic evidence

If the evolutionary provenance of senescence is indeed connected to alleles that enhance fertility, then these alleles should have effects that are observable in experimental studies and demographic statistics. The fact that senescence is so common a part of our daily experience, both with humans and with animals, led at one time to an expectation that these genetic tradeoffs might be comparably prominent and ubiquitous. Surprisingly, their observation has proved elusive and controversial. Both in the laboratory and in demographic studies, evidence for negative correlations between fitness traits early and late in life has been weaker and less general than expected. We discuss experimental results in several areas in a companion paper (Mitteldorf 2003), and focus here on evidence from human demography.

1.3 Different kinds of tradeoffs from AP and DS

There is a subtle difference between the kinds of tradeoffs predicted by the two flavors of pleiotropic theory: In theories involving explicit pleiotropy (AP theories), pleiotropic alleles may be differentially distributed through the population, leading to an inverse correlation between inborn fecundity and potential longevity, which can best be observed when there is no extraneous interference with conception. Search for such an effect must rely on historic data, or on modern populations that practice no birth control, active or passive. Theories of resource allocation (e.g. the DS theory), in contrast, predict an inverse relationship between longevity and the actual number of children born to a woman. It follows that this relationship should be observable in any era, as it cannot be masked by contraceptive measures.

In the pages that follow, we critically examine one study of each type. Both were widely

publicized when they came out, and both have been interpreted as supportive of the pleiotropic theories. The first demographic study is a statistical analysis by Westendorp and Kirkwood (1998) of a historic database of women from the British aristocracy. Although the results were reported to support a negative correlation between fertility and longevity, we note that the data on their face seem to attest to a positive correlation. The second is a survey by Perls et. al. (1997) of centenarian women in the Boston metropolitan area, showing that these women bore children later in life than a group of demographically matched controls who died in their seventies. The authors interpret this data as evidence of a correlation in the individual variations of two biological aging rates; we argue that the data are just as consistent with the hypothesis that bearing a child after age 40 actually confers protection on a woman, enhancing her prospect for longevity, a conclusion that is at odds with the dominant pleiotropic theories.

2 The Peerage Database

2.1 Results of Westendorp and Kirkwood

Westendorp and Kirkwood take for their source a published database with family histories of British aristocracy. Their choice of an upper-class sample was made to lower the noise in the data from accidental deaths related to poverty or hardship; historical rather than present data was preferred because the relationship they sought concerned inborn fecundity, and could be masked by contraceptive practices that have become common in the twentieth century. Their analysis was limited to women because the reproductive physiology of males is less directly related to the actual number of children sired. The authors further limit analysis to married women for whom there is complete data, who were born before 1875, and who died after menopause (taken alternatively as 50, 55 or 60 years of age); they seek to correlate the number of children each woman bore with her age at death.

The thrust of their result is announced in the title, “Human longevity at the cost of reproductive success.” Elsewhere, they couch their conclusions in more guarded language. “The relation between age at death and progeny number in British aristocratic women (and men) is consistent with the hypothesis that the longest lived individuals have reduced fertility compared with the majority of the population. This effect is seen most strongly in women born before 1700, for whom the number of children was larger than for women born between 1700 and 1875.” (Westendorp and Kirkwood 1998, p. 745)

The sophisticated statistical methods which these authors bring to bear on the data uncover an inverse relationship between fertility and longevity. But it is also true that a less sophisticated analysis points to the opposite conclusion.

2.2 Alternative analysis with linear regression

Linear regression of age at death vs. number of children produces a correlation coefficient ($r=-0.008\pm0.018$) that is negative but not significant. The clearest signals in the data are that

over the centuries, longevity increases (r=0.247±0.018) while fertility decreases (r=−0.126±0.018). An approach based on analysis of variance suggests a two-variable regression, age at death vs number of children and year of birth, to segregate this effect. When this is done, the result is a marginally significant *positive* relationship between number of children and age at death (r=0.025±0.019). Westendorp (2001) has suggested that a more sensitive correction for social trends in fertility may be made if from the number of children each woman has borne is subtracted the average number of children borne by a sliding group of her contemporaries. When this is done, the positive correlation becomes significant (r=0.060±0.018). (The AP predictions ought to apply equally well to women and to men. The corresponding analysis for men in the database again yields a significant positive correlation between fertility and longevity, r=0.021±0.011).

Fig. 1 displays age at death for women, averaged by the number of their children. There is no clear downward trend in the chart, even without correction for year of birth. (Data for the chart and regression calculations derive from a later edition of the same source (Bloore 2000) used by Westendorp and Kirkwood, a CD ROM listing family records of British Peerage. Included are the 2919 women born before 1906 for whom exact birth dates and death dates are available, and who survived to their 55th birthdays.)

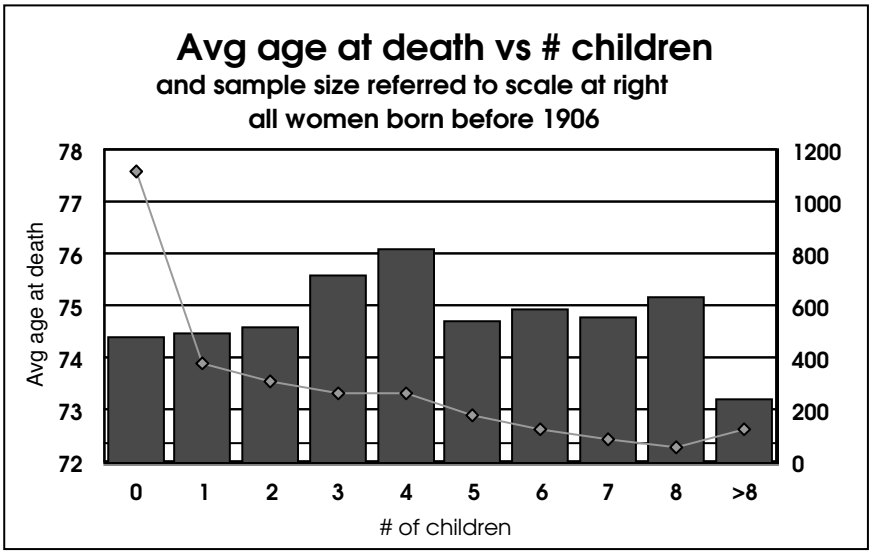

Fig. 1. Age at death for women, averaged by the number of their children. There is no clear downward trend in the chart, even without correction for year of birth. The line indicates sample size, referred to the scale at the right.

2.3 Why the disparity in conclusions?

Westendorp's only significant results arise when the sample is limited to women born before 1700. This subsample is relatively small, numbering only 241 out of the full 2919 in the database that meet other inclusion criteria. The group in which they claimed to detect a

measurable deficit in fertility comprises the longest-lived subset of those early women, who attained the age of 80 or more. There were 26 such women, and their mean number of progeny was 2.65±0.76, compared with 3.26±0.28 for women who lived 55 to 79 years. (Uncertainties here are computed as standard deviation of the mean.)

But the small sample size does not fully explain the disparity between the present results and Westendorp's: a linear regression limited to women born before 1700 produces an insignificant result ($r=-0.14\pm0.46$), and a bar chart limited to these women shows no clear trend (Fig. 2).

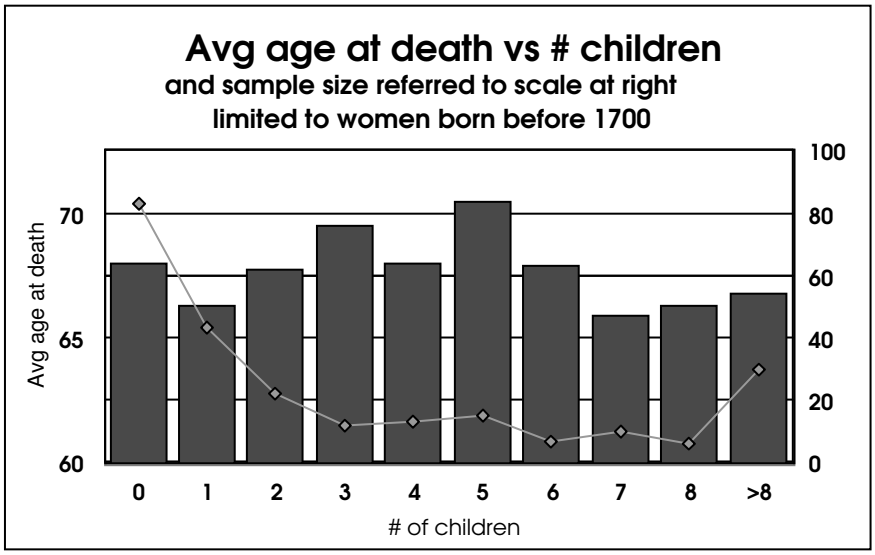

Fig 2. Same as figure 1, limited to women born before 1700. It is in these women that Westendorp and Kirkwood claim to discern a downward trend at the right side of the chart. Note the small sample size, indicated in the scale at the right.

A more important explanation for the disparity in results is that we use linear regression while Westendorp and Kirkwood use a less familiar tool, the Poisson regression. Poisson regression can be used to measure an association between any independent variable that is continuous, and a dependent variable that is Poisson distributed. The Poisson character of the dependent variable is an essential precondition; when the dependent variable is not Poisson distributed, the method can produce anomalous results. A Poisson distribution is controlled by a single parameter, and the core of Poisson regression analysis is to optimize over that parameter in such a way that likelihood of the observed distribution is a maximum.

For the Westendorp analysis, the calculation would proceed thus: Assume that the probability that a given woman will bear n children derives from a Poisson probability P(n,x), where the Poisson parameter x is a linear function of her age at death, x=A*age+B. Try different values for A and B, and calculate a product of probabilities for each woman that she would have borne the number of children that she did in fact bear. Home in on those values of A and B that maximize the probability product. Westendorp's result is that A is barely negative for women in their sample born before 1875, and is significantly negative when the sample is limited to women born before 1700. They report these results after allowance was made for trends over time in both

fertility and longevity.

It is possible that a major reason for the disparity between Westendorp's Poisson result and the linear regression analysis is that number of children born to a woman is not a Poisson variable. Fig. 3 shows the actual distribution of children for each woman in the sample, with a mean-adjusted Poisson curve overlaid, where x=2.36 was the actual mean number of children per woman in the sample. Compared to the Poisson curve, there is an excess for n=0 births, and a deficiency for n=1 through 4.

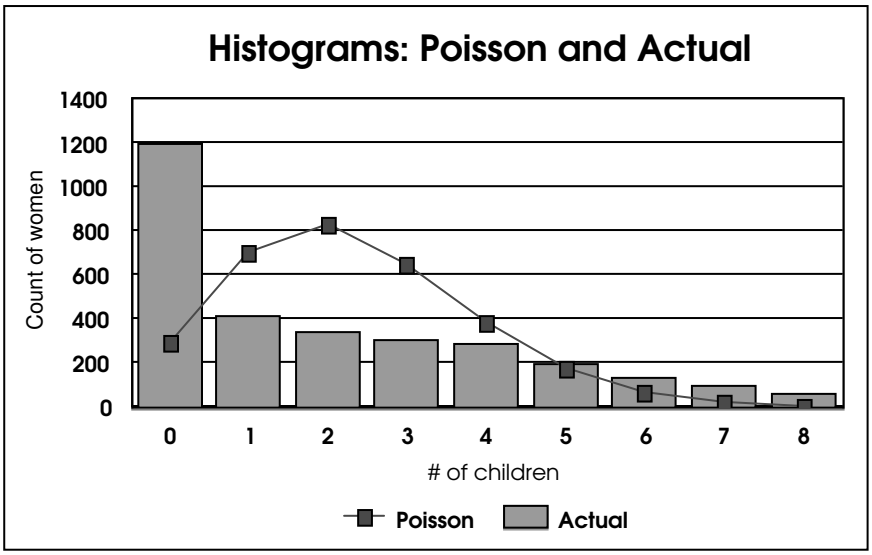

Fig 3. For Poisson regression to be a valid trend indicator, the underlying data must be Poisson distributed. The figure compares the distribution of women in the full sample by number of children with a Poisson distribution of the same mean.

More important to the Poisson regression calculation is the excess of women with large numbers of children. The Poisson probability becomes exponentially small for high n, but there are substantial numbers of women in the sample with more than 10 children. For example, in the sample of 2919, the most prolific woman bore 18 children. The Poisson probability that there should be such a woman is 10^{-7}. There are 9 women with 15 or more children, while the Poisson function predicts 0.0001 such women.

These observations suggest a reason for the difference between the Poisson and linear results. The Poisson probability that is maximized in Westendorp's regression calculation is inordinately influenced by the women at the very end of the distribution. Even though their numbers are small, the contribution that they make to the aggregate probability is dominant. If these few women should happen to have short lives, that would be sufficient to tilt the Poisson regression result to a negative coefficient. As it happened, the average age at death of the 9 women with the most children was 67.6, compared to 74.7 for the sample as a whole, a disparity that is well

explained by the average birth year of the 9: 1677, compared to 1830 for the whole sample. If these 9 women are removed from the sample of 2919, the Poisson regression no longer generates a negative coefficient. 5 of the 9 women with the most children were born before 1700. The conclusion of Westendorp and Kirkwood depends critically on this handful of women; if they are removed from the analysis, Poisson regression on the remaining subjects produces an insignificant result, even for the pre-1700 data set.

3. The Centenarian Study

Perls et al (1997) report that 15 of 78 centenarian women in the Boston area (19.2%) had borne children after their fortieth birthdays, compared to 3 of 54 (5.5%) in a control group of women who died at age 73. Women in both groups were born in 1894-96, and data on the centenarian women was obtained from interviewing them in 1995-96, while data on the controls were derived from hospital records from the late 1960's when these women died.

Perls suggests as an explanation that the effect derives from a correlation in the individual variations of two biological aging rates: the overall senescence rate and the rate of fecundity decline. Thus the subset of women whose reproductive aging rates are sufficiently slow as to permit them to bear children after age 40 overlaps strongly with the subset whose metabolic aging rate is such as to permit them to live 100 years. This hypothesis is consistent with prevailing theories for the evolution of senescence. He does not mention the alternative hypothesis that the act of bearing a child after 40, or of rearing a child thereafter, confers protection on the mother, enhancing her probability of survival to a great age. This latter explanation is presumably ruled out on theoretical grounds.

But, in fact, the centenarian results may be combined with previously compiled actuarial tables and data on fecundity decline to yield a fair comparison of the two hypotheses. For actuarial data, we rely on a cohort life table (Bell et al 1992) for women born in 1900. Standard mortality tables are tabulated as a snapshot in time, so that the data for each age apply to persons who are currently that age; a cohort life table is appropriate for our purpose, because it applies across time to all persons born in a single year. Perls's subjects were born in 1894-96, and 1900 was the closest year for which such a table was available from the National Office of Vital Statistics. For fecundity data, we have consulted studies from the 1980's by Menken et al (1986; Mencken and Larsen 1986) of fecundity decline with age. Finally, we have consulted fertility tables (Whelpton and Campbell 1976) from the US Dept. of HEW for women in the 1896 birth cohort.

In Menken's tables, 29% of women are reported to be infecund by age 40. In a separate category, which Menken et al label "impaired fecundity", they estimate that 55% of women 35-44 experience some "difficulty in conceiving or delivering a baby". In the HEW birth cohort data, the total fertility rate during the era 1936-1945 (a time when fertility was suppressed by non-biological factors of economic depression and war) was 7.2%

3.1 Analysis

Using these numbers, it is possible to assess the likelihood of Perls's preferred explanation: that infertility at 40 is a biomarker identifying women who are less likely to enjoy long lifespans. Among Perls's controls, who lived an average lifespan, 3 of 54 women bore children after age 40 (5.5%). For this small sample, the result is consistent with the national 7.2% rate from the HEW data (p=0.47). The higher rate among centenarians (15 of 78 = 19.2%) stands out enough to command an explanation, despite the small sample size (p= 0.0004). This is the probability of choosing 15 or more from a sample of 78, when each has a chance 7.2% of being chosen.

The fault we find with the Perls hypothesis is that it is inadequate to explain the strength of the observed association, even if the hypothesized correlation of aging rates is as strong as it can logically be. If this correlation is presumed to be unity, then the entire group of centenarians would be drawn not from the sample of all women born in 1896, but from the 71% of that group who remained fecund at age 40. But even this extreme assumption is insufficient to explain the data that Perls observed. Limiting the population entirely to the fecund 71% raises the expected number of late-childbearing women only by a factor of 100%/ 71%= 1.4; we should expect 1.4 times the 7.2% rate, which is 10.1% compared to the 19.2% observed. The likelihood of the observed result is p=0.011, which is the probability of choosing 15 or more from a sample of 78 when each individually has a 10.1% chance of being chosen.)

With the added extreme assumption that none of the 55% of women with "impaired fertility" were either among those who conceived a child after 40 or those who survived to age 100, this probability rises only to p=0.16.

The greatest uncertainty in this analysis derives from assumptions about the decline of fecundity with age. The Menken studies on which we rely are state-of-the-art, however he revises downward earlier estimates of the rate at which senescence erodes fecundity. In our simple model, fecundity is either on or off; allowing for a gradual decline could weaken or strengthen the statistical case against the Perls hypothesis. If the 29% figure quoted corresponds to an extreme of absolute sterility, then accounting for the impaired fecundity of some of the remaining 71% could help reconcile the hypothesis with the observations. On the other hand, if some women with impaired fecundity were included in the 29%, then compensatory measures that they took to achieve pregnancy could skew the results in the opposite direction. In all events, we believe the analysis using 55% infecundity is overly generous to the hypothesis.

We have used 7.2% total fertility 1936-1945 as a measure of the number of women in the 1896 cohort who actually bore children after age 40; this figure is a small overestimate, since it includes women who were born late in 1896 but who bore children early in 1936, and it double-counts those few women who had two children during this time. This overestimation can only make our conclusions more conservative.

For Perls's subjects, all drawn from the Boston area, the national fertility statistic which we employed (7.2%) may not have been appropriate. We have assumed that the difference between his control group (3/54 = 5.5%) and the national average is due to sampling error; but

alternatively if the 5.5%figure is taken as representative of the Boston sample, then our conclusion is strengthened from the range p=0.011 - 0.16 reported above to the range p=0.0009 - 0.05.

3.2 An alternative analysis

We consider here the alternative hypothesis that the act of bearing a child late in life confers some protection against aging.

A standard technique for modeling any influence on the aging process is to apply setback or setforward years to the mortality table. A mortality table that applies to a control population may be modified so that it applies approximately to a test group by interpreting all ages in the table as if they were *s* years younger or older, where *s* is the setback or setforward. (For fractional *s*, linear interpolation may be applied to create the setback or setforward table.)

Modeling the alternative hypothesis, we applied an actuarial setback *s* to the 7.2%of women in the cohort who actually bore children past age 40. A corresponding setforward *t* was applied to all others; this adjustment was necessary to reproduce the standard table for the composite statistics from the two groups together. Perls's data were best reproduced for $s = 3.3$ years, corresponding to $t = 0.5$ years. With these parameters, both the 15/78 in the centenarian group and the 3/54 in the control group were reproduced with (maximized) composite probability p=0.41 for the data space $\geq$15/78 for test and $\leq$3/54 for control subjects. Thus our model predicts that women who bear a child past age 40 live an average 3.8 years beyond the life expectancy for women who do not.

If extended life is indeed a response to late childbirth, the effect might be an evolutionary adaptation permitting women to survive long enough to nurture their children and grandchildren. The phenomenon would be akin to the reason that women's lifespans extend beyond menopause, and the reason that, among primate species, those in which the male assumes a nurturing role for the young are also those in which the male's lifespan compares favorably to the female's (Allman et al 1998).

3.3 Comparison of the Two Models

The explanation offered by Perls and the alternative suggested here have each been modeled with one free parameter. In the case of the Perls model, the observed result was still found to be unlikely, even when pushing that parameter to its logical limit (perfect association between fecundity decline and mortality increase with age). The alternative model handily explains the data with the assumption that late childbearing women add 3.8 years to their lives.

4. Other demographic and laboratory evidence

These results should be viewed in the context of a growing body of animal and human data, some of which supports the theoretical notion of a cost of reproduction but much of which suggests the opposite. In his encyclopedic review of the literature, Finch (1990) reported that there was considerable animal evidence for an immediate mortality cost of reproduction, but little evidence that reproduction affects the rate of aging. Stearns (1992) has appended a survey of animal tradeoff studies to his text on life histories, cataloguing evidence for and against the existence of a tradeoff. Genetic experiments with nematodes suggests that genes affecting the rate of aging are not necessarily linked to fertility (Partridge & Gems, 2002); while work with fruitflies has been interpreted as supportive of tradeoffs (Partridge & Gems, 2002; but see Mitteldorf, 2003 for a contrary interpretation). Reznick et al (2000) review evidence that longevity and fertility in water fleas appear to be independently variable.

Besides the Westendorp study, a number of parallel attempts have failed to detect a negative association between fertility and longevity in historic data on human females (Korpelainen 2000; Lycett et al 1999; Le Bourg et al 1993). These studies and other demographic data are reviewed by Le Bourg (2001). In a contemporary study of human males, Davey Smith et al (1997) report that sexual activity lengthens lifespan, and they note the consistency of their result with other studies of females (Abramov 1976) as well as males (Palrnore 1982; Kaplan 1988).

5. Conclusion

Demographic science has a valuable role to play in verification of evolutionary hypotheses concerning aging. This role is best served when competing theories are confronted with population data for comparison, and the conflicts are reported along with the consonance.

Evidence concerning the predicted tradeoffs between fecundity and longevity has thus far been equivocal. If indeed these tradeoffs have been the evolutionary basis for selection of senescence, one might have expected tradeoffs to be observed more cleanly. Thus demographic data on fertility and longevity should not be reported as a triumph for the theories, but rather as an invitation to further analysis.

References

ABRAMOV, L. A. 1976. Sexual life and frigidity among women developing acute myocardial infarction. *Psychosom Med*. **38**:418-425.

ALLMAN, J, A. ROSIN, R. KUMAR AND A. HASENSTAUB. 1998. Parenting and survial in anthropoid primates: Caretakers live longer. *Proc. Natl. Acad. Sci.* **95**:6866-6869.

BELL F. C, WADE A. H. AND GOSS S. C. 1992. Cohort Life Tables for US Social Security Area by Year of Birth and Sex (Women born in 1900). Life Tables for the United States, Actuarial Study #107, U.S. Social Security Admin, Dept of HHS

BLOORE, J ed. 2000. Peerage CD-ROM, 13th edition, S&N Genealogical Supplies, Salisbury, UK

CLARK, W.R. 1999. A means to an end: the biological basis of aging and death. Oxford University Press. New York

DAVEY SMITH, G, FRANKEL S. AND YARNELL J. 1997. Sex and death: are they related? Findings from the Caerphilly cohort study**.** *Brit Med J* **315**

EDNEY, E. B. AND GILL R. W. 1968. Evolution of senescence and specific longevity. Nature **220**:281-282

FINCH, C.E. 1990. Longevity, Senescence and the Genome. University of Chicago Press, Chicago, IL.

HAYFLICK, L. 1994. How and why we age. Ballantine books. New York.

KAPLAN, S D. 1988. Retrospective cohort mortality study of Roman Catholic priests. *Prev Med.* **17**:335-343.

KENYON, C. 2001. A conserved regulatory system of aging. Cell. **105**:165-168.

KIRKWOOD, T. B. L. 1977. Evolution of aging. *Nature*. **270**:301–304.

KORPELAINEN, H. 2000. Fitness, reproduction and longevity among European aristocratic and rural Finnish families in the 1700s and 1800s. *Proc R Soc Lond* **B 267**:345-348

LE BOURG, E, THON B, LEGARE J, DESJARDINS B. AND CHARBONNEAU H. 1993. Reproductive life of French-Canadians in the 17-18th centuries: a search for a tradeoff between early fecundity and longevity. *Exp Geront*. **28**:217-232

LE BOURG, E. 2001. A mini-review of the evolutionary theories of aging. Is it the time to accept them? *Demog Res.* **4**:1-28

LYCETT, J.E, DUNBAR R. I. M. AND VOLAND E. 1999. Longevity and the cost of reproduction in a historical human population. *Proc Soc Lond. **267**:31-35*

MAYNARD SMITH J. 1976. Group Selection. Quart Rev Biol. **51**:277-283

MEDAWAR, P. B. 1952. An unsolved problem of biology. H. K. Lewis. London.

MENKEN J, TRUSSELL J. AND LARSEN U. 1986. Age and Infertility. *Science*. **233**:1389-1394

MENKEN J. AND LARSEN U. 1986b In: Mastroianni L. & Paulsen C.A, ed. Aging, Reproduction and the Climacteric, New York, NY: Plenum Publishing; 1986; 147

MITTELDORF J. 2003. Pleiotropic theories of aging and their empirical foundation. I: Laboratory experiments. Forthcoming.

PALRNORE E. B. 1982. Predictors of the longevity difference: a 25-year follow-up. *Gerontologist*. **6**:513-518.

PARTRIDGE L. AND Gems D. 2002. Mechanisms of ageing: public or private? *Nature Genetics*. **3**:165-175.

PERLS T, L. ALPERT AND R. C. FRETTS. 1997. Middle-aged mothers live longer. Nature. **389**:133

PROMISLOW D., TATAR M., KHAZAELI A.A. AND CURTSINGER J.W. 1996. Age specific patterns of genetic variance in *D. melanogastar*. I. Mortality. *Genetics* **143**:839-848

REZNICK D, L. NUNNEY AND A. TESSIER. 2000. Big houses, big cars, superfleas and the costs of reproduction. *TREE* **15**:421-425.

STEARNS, S. 1992. The evolution of life histories. Oxford Press, New York.

WEISMANN A. 1889. Essays upon heredity and kindred biological problems. Clarendon Press. London.

WESTENDORP, R. G. J. 2001. Personal communication.

WESTENDORP, R. G. J. AND T. B. L. KIRKWOOD. 1998. Human longevity at the cost of reproductive success. *Nature*. **398**:743-746.

WHELPTON, P.K. AND CAMPBELL A. 1976. Fertility tables for birth cohorts of American women, US Dept of HEW, National Office of Vital Statistics.

WILLIAMS, G. 1957. Pleiotropy, natural selection, and the evolution of senescence. *Evolution*. **11**:398-411.

_____ 1966. Adaptation and natural selection. Princeton University Press. Princeton, NJ.