COLD SPRING HARBOR SYMPOSIA ON QUANTITATIVE BIOLOGY

VOLUME XXII

Population Studies: Animal Ecology and Demography

THE BIOLOGICAL LABORATORY
COLD SPRING HARBOR, L. I., NEW YORK
1957

Reprinted from COLD SPRING HARBOR SYMPOSIA ON QUANTITATIVE BIOLOGY
Volume XXII, 1957
Printed in U.S.A.

How the Age Distribution of a Human Population is Determined

ANSLEY J. COALE

Office of Population Research, Princeton University, Princeton, New Jersey

This analysis of age distribution will be restricted to closed populations of human females; restricted to closed populations because to consider the effects of migration would be unduly complicated; restricted to human populations because the author is a demographer rather than a biologist; and restricted to females because differences in age composition and in age at parenthood between the sexes make a combined treatment awkward. If data were equally available, a completely similar analysis could be made for males, with only one major reservation: The special effects of war mortality on males of military age would require special analysis.

I. STABLE POPULATIONS

The forces affecting the shape of the age distribution are most readily visualized in the special case of an unchanging distribution. Lotka (Dublin and Lotka, 1925; Lotka, 1939) has shown that constant mortality and fertility schedules will ultimately produce a constant age distribution, and a constant rate of growth. This constant age distribution can be calculated as the product of a factor representing relative number of survivors to each age from birth (the life table) and a factor reflecting the continuously growing (or shrinking) number of births. Specifically:

$$c(a) = be^{-ra}p(a) \tag{1}$$

where $c(a)$ is the fraction at any given age a, b is the birth rate (which Lotka shows is constant under the specified conditions) r is the growth rate, and $p(a)$ is the fraction who survive from birth to age a. No matter what the initial age distribution (provided the childbearing ages are amply represented) a constant age schedule of fertility ($m(a)$) and probability of surviving ($p(a)$) eventually establish the distribution given in equation (1). Actually, a period of 50 to 100 years is adequate to produce a very close approximation to the stable population.

The basis for equation (1) can be simply explained. The birth rate (b) is the proportion at age zero (set $a = 0$, and the other factors are each unity). The term e^{-ra} relates the size at birth of the cohort now aged a to the current birth cohort, and $p(a)$ allows for the attrition of mortality.

Since $b = \dfrac{1}{\int_0^\omega e^{-ra}p(a)\,da}$, (where ω is the highest age attained) the stable age distribution is wholly determined by the growth rate r, and the survivorship function $p(a)$. The value of r in turn can be calculated from the second fundamental equation of stable population theory.

$$\int_0^\omega e^{-ra}p(a)m(a)\,da = 1 \tag{2}$$

The real root of this integral equation is the stable growth rate, while the complex roots determine how the stable population is approached from arbitrary initial conditions.

We now turn to the role of fertility and mortality in determining the shape of the stable age distribution. The role of mortality will be described first.

Mortality and the stable age distribution

Two rather surprising conclusions emerge when the effect of mortality schedules on the stable age distribution is examined:

(a) The effect of alternative mortality schedules is relatively minor. Roughly similar age distributions result from a given fertility schedule in conjunction with a very high mortality life table on the one hand, or a very low mortality life table on the other. Figure 1 illustrates this point. The life table of Sweden for 1860, with a life expectancy of 45.4 years, and that for 1946 to 50, with a life expectancy of 71.6 years produce about the same stable age distributions when combined with the same fertility schedule.

(b) Among life tables reflecting experience so far recorded, it is nearly universally true that a more favorable mortality schedule—with a higher life expectancy—will yield a stable population with a higher proportion under 15. In a vast majority of contrasting life tables, the lower mortality life table will produce a distribution with a lower average age; and at least half the time, lower mortality will produce a smaller fraction over 65.

These conclusions are partly analytical and partly empirical. A full analysis of the contrasting stable age distributions associated with different mortality schedules is somewhat laborious (Coale, 1956). A simple qualitative argument is enough to dispel the common belief that lower mortality inevitably means an older population. If in one life table the probability of surviving for a year at each age exceeded the corresponding proba-

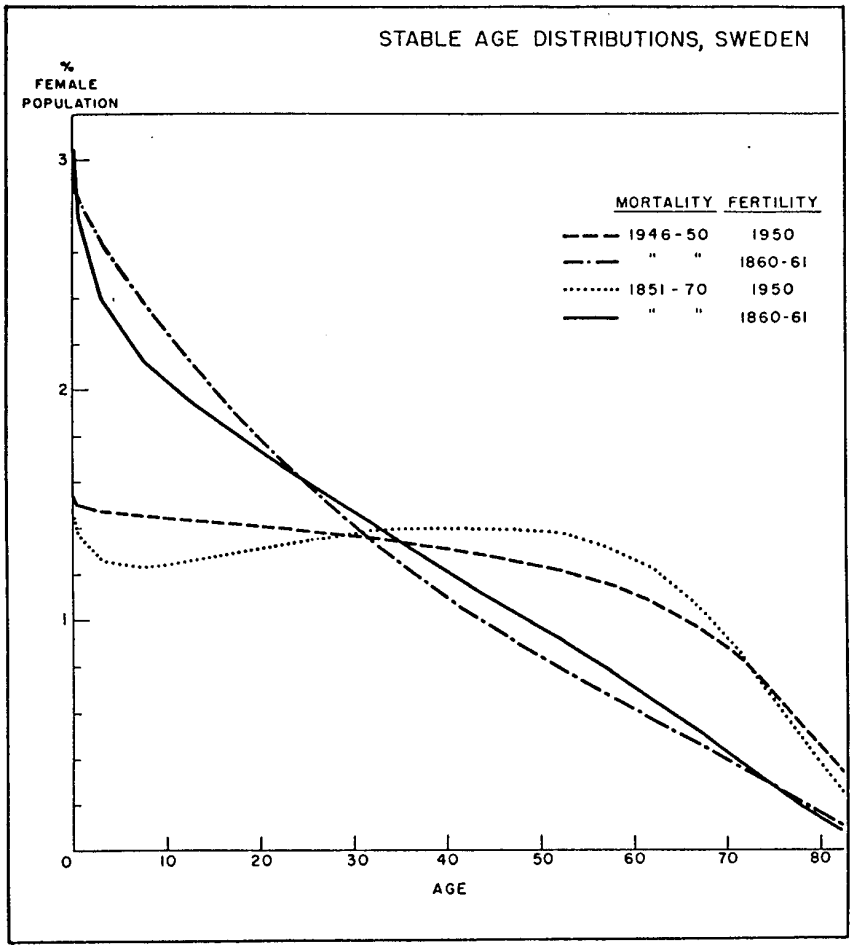

FIGURE 1.

bility in another life table by a fixed proportion (for example, 1%) the two life tables would produce precisely the same age distributions. The tendency toward more survivors with advancing age would be exactly offset by a faster rate of growth that tends to make each cohort *smaller* than the next younger.

This point can be proved by assuming a sudden shift to a life table with a one per cent higher probability of surviving at each age. The year following this change, there would be one per cent more one-year olds surviving from birth, one per cent more two-year olds, etc. Improved survivorship would produce one per cent more persons at every age above 0. But since the increase would yield one per cent more persons at every childbearing age, with constant fertility there would be one per cent more births—persons at age 0—as well. The whole population would be one per cent larger; and the age distribution would be unaffected.

The common view that lower mortality means an older population takes account of only part of the effect of lower death rates. It is immediately clear that lower death rates produce more old people. However, lower death rates also produce more parents, more births, and more children. Whether the dominant effect on the stable age distribution of a lower level of mortality is to enlarge the upper end of the age distribution through higher survivorship, or to tilt the age distribution more steeply through more rapid growth depends on the relative age pattern of mortality in the two life tables (Coale, 1956).

A particularly illuminating way of comparing the age pattern of mortality is to compute the proportionate difference in the probability of surviving for one year at each age. Let the probability of surviving from age a to $a + 1$ be $\pi'(a)$ in the better life table, and $\pi(a)$ in the higher mortality table, and consider the ratio $\frac{\pi' - \pi}{\pi}(a)$. We have already shown that a constant value of $\frac{\pi' - \pi}{\pi}(a)$ implies the same age distributions. We can subtract the minimum value of $\frac{\pi' - \pi}{\pi}$ from the value at other ages—the subtracted portion is equiva-

lent to no difference in mortality—and consider only the residual. A large excess of $\frac{\pi' - \pi}{\pi}$ above the minimum at ages below childbearing means there is a large difference in the *growth rates* of the two stable distributions. Moreover, since a large excess $\frac{\pi' - \pi}{\pi}$ at the youngest ages implies more survivors at *all* subsequent ages, the *proportion* of survivors above—say—age 50 in the better life table would not be much larger than in the poorer one. In short, an excess of $\frac{\pi' - \pi}{\pi}$ above the minimum at young ages means that the better mortality schedule tends to have a *younger* age distribution.

On the other hand, an above-the-minimum $\frac{\pi' - \pi}{\pi}(a)$ at ages over 50 implies no difference in the long term growth rate (since the reproductive ages are not involved) but does mean more old age survivors in the better life table. The net result is a higher fraction at ages above 50.

The effects of three special instances of percent difference in the probability of surviving have been described in a non-rigorous fashion. These particular patterns have been emphasized because differences among the great majority of recorded life tables can be approximated as the sum of three components—a minimum per cent difference in the probability of surviving from age 5 to 50, above minimum differences below age 5, and above minimum differences above age 50 (see Fig. 2). The general pattern of $\frac{\pi' - \pi}{\pi}(a)$ is roughly U-shaped, declining from a high at age zero to near its minimum by age 5. It is relatively constant until age 50 or 60 where it frequently but not universally rises. The central portion (from 5 to 50) contributes nothing to the differences in stable age distribution, serving only to diminish the effect of the differences represented by the two legs of the U, and to raise the growth rate of the population.

The two legs of the U-shaped pattern work more or less in opposition in causing differences in age distribution. If there were no right leg, the age distribution with the better life table could be approximated at all ages except those under 5 by a sort of pivoting of the higher mortality distribution on its average age, with increased fractions at less than the average age, and reduced fractions above. If there were no left leg, rising per cent differences in survivorship over age 50 would mean higher fractions above 50, and lower fractions below. When the two legs operate together, the net effect is to reduce the differences each would cause alone. If the left leg is large, and the right leg is small, the effect of growth rates will dominate; and the low mortality population may have smaller proportions at all ages above the average. This domination is typical when a life table of high mortality is compared with a moderate or low mortality table. However, when mortality of 20 years ago in the most advanced areas is compared with current mortality, the right leg—improved survivorship at the older ages—takes on more importance. There is not much room for a left leg in a comparison of future life tables with the best current ones. A two or three per cent further improvement in the probability of surviving to age 5 in Sweden would raise this probability to one.

There is no logical necessity that mortality differences should produce only minor differences in stable age distributions. This result arises from the age structure of differences among life tables observed to date. However, a contributing factor in the typical U-shaped age structure of differences in the probability of surviving is that large differences can last only when there is room for them. Differences of the order of those observed for age 0 could not occur at those ages where even with high general mortality levels the risk of death is moderate.

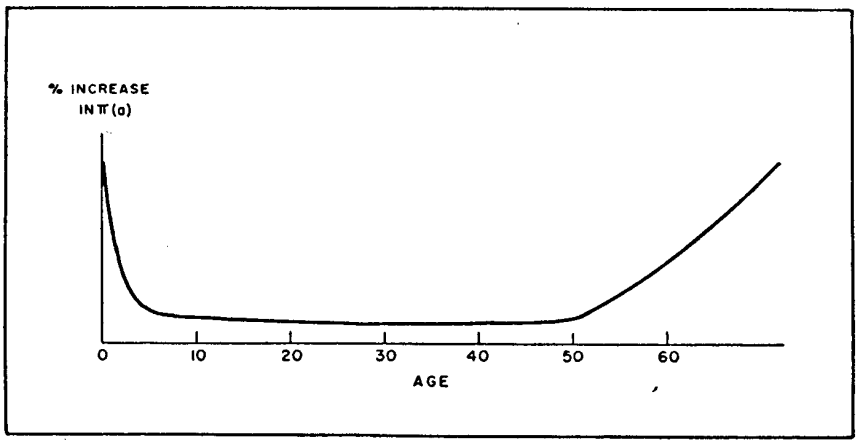

FIGURE 2.

Fertility levels and the stable age distribution

The role of fertility in shaping the stable age distribution is at once simpler than that of mortality, and quantitatively much more important, when account is taken of the range of mortality and fertility observed in the world. The simplicity and the quantitative importance have a common source: the fact that differences in fertility operate in a single direction in affecting the relative shape of the age distribution. While higher probabilities of surviving simultaneously flatten out $p(a)$ in equation (1) and make e^{-ra} steeper, higher probabilities of giving birth affect only the growth rate—making it larger, of course, and making the exponential factor in equation (1) contribute to a more rapid taper in the distribution.

If a high fertility stable age distribution is compared with a low fertility distribution with the same mortality, the two will be found to intersect at the mean of the average ages. The higher fertility population will have higher proportions at ages below the average, of course.

Figure 1 makes it clear that fertility differences can produce profoundly different stable age distributions. This fact together with the relatively slight influence of mortality on the stable age distribution means that a schedule of fertility by age is sufficient to give at least a fair approximation to the stable age distribution even if mortality rates by age are not known. One would simply use whatever female life table was lying around in calculating the stable population. It must be admitted that there is a much better chance of a close fit if some hint about infant mortality is visible.

One final comment on a common sense basis for understanding the powerful influence of fertility on the stable age distribution. The general fertility rate establishes the ratio of the area in the age distribution from ages 17 to 44 to the zero ordinate of the distribution. If fertility is twice as high, this ratio must be cut in half irrespective of mortality. The exact inverse relation of this ratio to fertility is a consequence of the fact that the zero ordinate is proportional to births, and the area from 17 to 44 to the number of women of childbearing age. Thus the fertility level clamps a vice on the relation of the beginning of the distribution to an area near the middle.

Several straightforward conclusions emerge from this consideration of stable age distributions:

(1) Sustained high fertility (average completed size of family, 6 or more children) produces a young population with a median age well below 25 years, more than 40 per cent of the population under 15, and no more than three or four per cent over 65.

(2) Sustained low fertility (average completed size of family below 2.5 children) produces an old population with a median age above 35, no more than 20 per cent of the population under 14, and at least 15 per cent over 65.

(3) In general, the approximate form of the age distribution is determined by the level of fertility. The level of mortality has more or less second order effects on the distribution. The general quality of these effects has been to date that low mortality yields a slightly larger fraction at ages up to at least 15 but as high as the average age, and somewhat smaller fractions either at all higher ages or until age 50 or higher. More often than not, sustained low mortality yields a slightly lower average age.

II. Varying Fertility and Mortality and the Age Distribution

We turn now to a brief consideration of how an age distribution is determined under a regime of continuously changing mortality and fertility rates.

The general equations corresponding to equation (1) are:

(3) $n(a, t) = B(t - a)p(a, t)$ where $B(t - a)$ is the number of births a years before time t, $n(a, t)$ is the number of persons at age a at time t, and $p(a, t)$ is the proportion of those born at $(t - a)$ who survive to achieve age a at time t, and

$$(4)\quad c(a, t) = \frac{n(a, t)}{\int_0^\omega n(a, t)\, da}\ .$$

The equation corresponding to (2) is:

$$(5)\quad B(t) = \int_0^\omega n(a, t)m(a, t)\, da$$ where $m(a, t)$ is the probability of bearing a female child at age a and time t.

These equations of course do not lead to Lotka's tidy solutions since we here permit mortality and fertility to vary with time. In fact, they do not even give much clue to the role of fertility and mortality in determining age distributions, since equation (3) requires us to know $B(t - a)$, and equation (5) tells us that $B(t - a)$ depends on the age distribution at time $(t - a)$ as well as on fertility rates at time $(t - a)$. It would appear that to account for the present age distribution, one needs to know:

(a) an age distribution at some past date, and

(b) schedules of mortality and fertility since that date.

However, as the date of the past age distribution is made more remote, its form makes less and less difference to the shape of the current age distribution. It seems intuitively plausible, in fact, that if the course of fertility and mortality were known since $t - \infty$, the proportionate age distribution would be wholly determined at time t no matter what the distribution at $t - \infty$, assuming of course that the initial distribution was not one—for example with no one under 50—headed for extinction. Lotka (1939) shows this statement to hold in the special case of endlessly unchanging fertility and mortality. But the same factors that cause the transient effects of a particular initial

age distribution to disappear from the stable population would also operate for *any* observed time path of fertility and mortality rate. After a a suitably long period the effect of an initial age distribution is swamped by the cumulative effect of the time pattern of vital rates. To put the point more concretely, any age distribution with persons at every age could be assumed for 100 years ago in place of the actual distribution. If such an assumed population were projected to the present with observed fertility and mortality rates, the proportionate age distribution would differ negligibly from the actual. In short, the age distribution of a closed population is determined by the mortality and fertility rates of recent history.

Another result from stable population analysis is suggestive for our general case—the conclusion that mortality differences have only second order effects on the age distribution. An increase or decrease in mortality tends to decrease or increase all cohorts, implying a small effect of mortality on the immediate age distribution as well as the stable. The short run transient age distribution effect of a mortality change may differ somewhat in form from the long run effects as expressed by stable age distributions, but the magnitude should be small in both instances.

We shall proceed with the provisional hypothesis that changes over time in mortality schedules do not have major effects on the age distribution. If this hypothesis is valid, a current age distribution could be closely approximated by calculating what the distribution would have been had observed fertility rates and *unchanging* mortality rates prevailed for the last 80 or 90 years. I have tried such a calculation using Swedish data. The results are shown in Figure 3. The proportionate distribution based on mortality unchanged for 90 years indeed does come close to the actual distribution. Moreover, the differences between the census population and the hypothetical are nearly identical to the differences between two stable populations with the same fertility, one based on an 1860 Swedish life table, and the other on a life table for 1946 to 50 (compare Fig. 1). The principal effect of projecting with unchanged mortality is to produce smaller fractions at ages under 25 and over 70. The reason is that since 1860 there have been disproportionate improvements in the probability of surviving in infancy and in the older

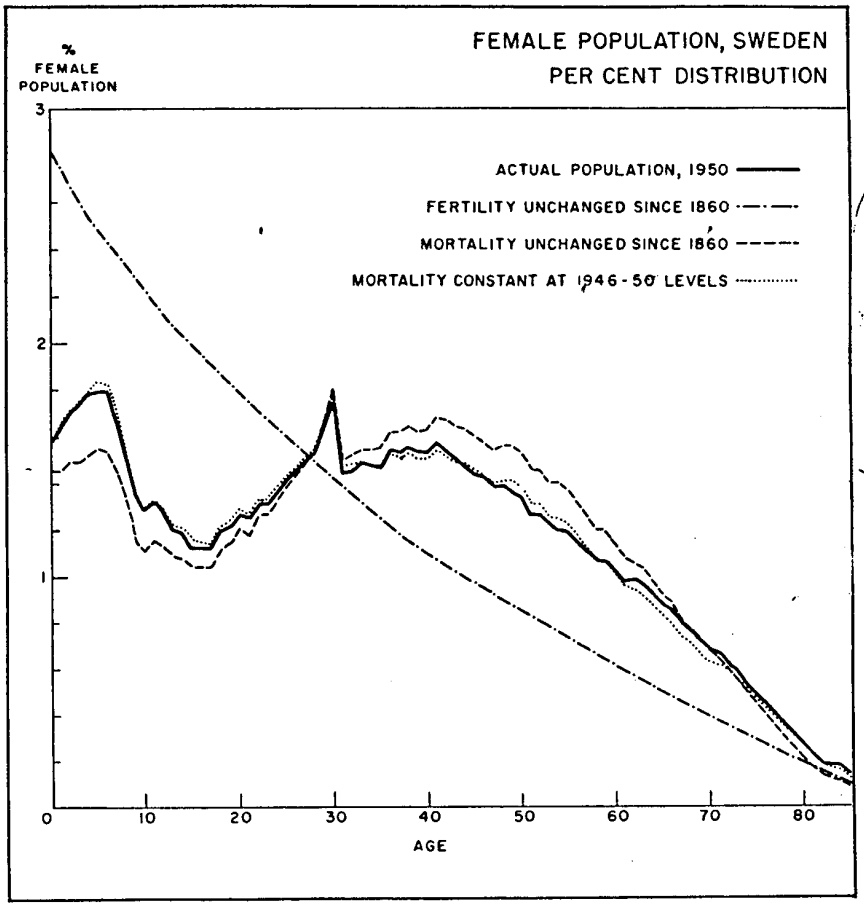

FIGURE 3.

ages. If the projection from 1860 is calculated with mortality unchanging at 1946 to 50 levels, the result is a distribution virtually indistinguishable from the census distribution. If in the 90 years before 1950 mortality had remained as high as in 1860 the result would have been a much smaller female population with an age distribution similar to the actual; if in these 90 years mortality had always been as low as in 1950, the result would have been a much *larger* female population with an age distribution nearly identical to the actual.

Figure 3 also shows the distribution that would have resulted with observed mortality risks during the 90 years before 1950, and with fertility assumed constant at 1860 levels. This figure makes clear what is the major determinant of an age distribution—the course of fertility. I wish I were now able to present a short, simple explanation of the effects of fertility on the age distribution. It is clear that fertility determines the age distribution, but an attempt to explain the relationship precisely soon runs afoul of major complications.

The basic difficulty is that when high fertility produces a brief series of unusually large cohorts, for example, these cohorts not only exceed their neighbors through their lifetime (this is a simple effect), but also when they pass through the childbearing ages they produce a diffused "echo" of larger birth cohorts. It is easy to give instructions about how to compute these consequences, but impossible to describe them simply in terms of the resultant age distribution.

We can make these observations about the effects of fertility on age distribution:

(1) Current and recent high fertility produce a younger population than would low fertility, and *vice versa*.

(2) Transitory waves of unusually high or low fertility create humps and hollows that move out through the age distribution as the cohorts move through life. In fact all of the notches and knobs in the Swedish female age distribution of 1950 can be traced back to unusual birth crops.

(3) A long period of high fertility, or a period of rising fertility creates a section of the age distribution that tapers rapidly with age. Conversely a long period of low fertility, or a period of falling fertility creates a relatively flat (or even rising) section of the age distribution.

(4) When a cohort of unusual size reaches the childbearing ages, it sets up an attenuated and flattened out "echo" in the number of births. It is this second generation effect that makes the analysis of the relation of the age distribution to the course of fertility so complex.

Finally I turn to one or two practical or at least worldly observations arising from this anslysis. First, the rising fraction of the aged in western countries has not resulted from lowered death rates but almost wholly from a long history of declining fertility. However, future improvements in mortality may make this statement obsolete. Second, improved mortality will not reduce the "burden of dependency" imposed on low income areas because of their age distributions. To the contrary, mortality reduction can be expected to lead to a somewhat higher fraction in dependent ages. If the burden is to be reduced, fertility must be lowered. Third, the remarkable variation in fertility in the U. S. during the past 20 years has produced a very irregular age distribution. The succession of variously sized cohorts will doubtless have interesting implications in the next half century. Thus in the next 20 years the number of persons 20 to 24 is due to increase by at least 75 per cent, while the number 40 to 44 can be expected to decline by about three per cent.

References

Coale, Ansley, J., 1956, The effects of changes in mortality and fertility on age composition. Milbank Mem. Fund Quart. *34:* 79–114.

Dublin, Louis I., and Lotka, Alfred J., 1925, On the true rate of natural increase. J. Amer. Statis. Assoc. *20:* 305–327.

Lorimer, Frank, 1951, Dynamics of age structure in a population with initially high fertility and mortality. Population Bull. U. N. *1:* 31–41.

Lotka, Alfred J., 1939, Théorie analytique des associations biologiques. Deuxième partie: Analyse démographique avec application particulière à l'espèce humaine. Paris, Hermann et Cie.

DISCUSSION

G. J. Stolnitz, Indiana University, Bloomington, Indiana: Several aspects of Dr. Coale's paper can be usefully supplemented. First, stable-age models of the kind discussed depend upon highly limiting assumptions concerning the nature of the initial population, the types of changes in vital rates, and the length of time needed for the full consequences of such changes to evolve. With respect to mortality changes at least, and probably with respect to fertility as well, simpler yet more general methods can be applied to determine the effects on age composition in actual populations, both for short-run and long-rung periods. The major importance of Dr. Coale's results is that all of his main conclusions can in fact be shown to hold in a more realistic way than he has indicated.

Secondly, the restriction of his analysis to females is unnecessary in dealing with mortality and arises largely from his preoccupation with stable-age situations. The effects of mortality changes on male age composition are closely related to the effects for females. This follows from the near constancy of the sex ratio at birth; the similarity of mortality changes among males and females up to the advanced ages, roughly 50 to to 60; and from the fact that it is sufficient to consider only mortality changes among females under 50 in determining the effects of male and female mortality trends on numbers of births. Male-female differences in age-specific fertility are important from a stable-age viewpoint, since they may imply absurd eventual sex ratios. In contrast,

they are of little importance in considering actual populations when we study the effects of given mortality movements.

Thirdly, it needs to be emphasized that the limited over-all effects found for both stable-age and actual populations result from the fact that the appropriate measures for analyzing age composition and mortality are proportionate changes in survival rates and not either proportionate or absolute changes in mortality rates proper. Fourthly, the effects of a succession of mortality changes can be sufficiently well approximated in a much easier fashion than has been indicated, by techniques involving the addition of the effects of one-time changes. Finally, the statement that the mortality changes we have seen in the past typically make populations younger deals primarily with the effects on average age. It does not cover some leading senses in which students have used the concept of population "ageing". Mean age of population may remain the same even while the proportion in the older ages, for example 65 or over, increases by substantial margins.

H. V. MUHSAM, United Nations, New York: There can be little doubt that fertility is in fact the principal factor determining the age distribution. This statement requires however, qualification: It is true as long as mortality and fertility are permitted to vary within the range actually observed—or at the most, imaginable—in the human species.

This range is limited by two types of circumstances, those which stem from the very character of the phenomena, birth and death, which we are concerned with, and those which are peculiar to the human species. These latter lead us to consider "total fertility rates" of say 2 to 8, and life expectation at birth of say 30 to 75 years. It thus appears as if fertility could vary much more than mortality, at a ratio of 1:4 as against of 1:2.5. But this comparison is almost meaningless. A meaningful comparison of the variability of mortality and fertility can only be made if both are measured in an equivalent manner, that is, if fertility is measured in terms of the total fertility rate, mortality should be measured in terms of "total mortality rate". This rate is obviously unity, that is, the total mortality rate is "one per person". Variations in mortality are therefore only due to differences in timing of deaths, not in the intensity of mortality. If variations of fertility were also restricted to differences in timing—which is quite imaginable for other than human species—it should be expected that mortality would have the stronger effect on the age structure of the population, because death can occur at any age (in man, between birth and ω years), while the timing of fertility is limited to the much narrower range of reproductive ages (between, say, 15 and 50 years). Under these circumstances, mortality would probably have the first order effect on the age structure, and fertility, which is in any event affected by mortality, owing to the latter's effect on the proportion of a cohort reaching—and completing—the fertile age span, would be considered of second order importance.

But, in fact, at least in *Homo sapiens*, fertility varies not only in timing but also in intensity, while mortality necessarily varies in timing only. This seems to be the main reason why fertility has apparently a higher effect on the age structure than mortality.