Averaging Life Expectancy

Andreev, E., Lutz, W. & Scherbov, S.

IIASA Working Paper

WP-89-035

June 1989

Andreev E, Lutz W, & Scherbov S (1989). Averaging Life Expectancy. IIASA Working Paper. IIASA, Laxenburg, Austria: WP-89-035 Copyright © 1989 by the author(s). http://pure.iiasa.ac.at/id/eprint/3309/

WORKING PAPER

AVERAGING LIFE EXPECTANCY

Evgenii Andreev
Wolfgang Lutz
Sergei Scherbov

June 1989
WP-89-35

IIASA
International Institute
for Applied Systems Analysis

AVERAGING LIFE EXPECTANCY

Evgenii Andreev
Wolfgang Lutz
Sergei Scherbov

June 1989
WP-89-35

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
A-2361 Laxenburg, Austria

Foreword

More than 50 years ago a peculiar feature of averages was pointed out by Cohen and Nagel (1934, p. 449) in their well-known book on logic. Their example happened to be mortality from tuberculosis; Blacks in Richmond, Virginia, had a lower rate than Blacks in New York; Whites in Richmond had a lower rate than Whites in New York; yet the overall rate for Richmond was higher than that in New York. Many other examples of the paradox have subsequently been pointed out, and some general theory has been presented by Colin Blyth.

If this were only an arithmetical curiosity no one would care much about it, but in fact its very possibility is a troubling consideration for all numerical comparisons. If recognizing Blacks and Whites reverses the standing of New York and Richmond, how do we know that recognizing some further breakdown will not reverse the standing once again?

This paper presents a different paradox, but one that is also threatening to the drawing of conclusions from numerical data. To follow Cohen and Nagel's example but disregarding the distinction between Blacks and Whites, the arithmetic average mortality (expressed as a death rate) of Richmond and New York combined will always fall between the rate for New York and that for Richmond. The result of a linear averaging process cannot fall outside the units averaged.

This is no longer true when a non-linear form of average is used—the present paper shows a hypothetical example for a harmonic mean. There are many questions that require non-linear averaging. One such is life expectancy, that is a weighted function of the usual (age-specific) rates, but the weighting is nonlinear. The authors came on this paradox in studying life expectancy for women in the Soviet Union as projected to the year 2020. All of the republics fall between 77.713 and 78.026, but the figure for the USSR comes out to 77.632.

This is not an error due to rounding; it is not due to Simpson's paradox that would result from internal heterogeneity in the several republics; it is due to the nonlinear weighting implicit in the calculation of the life expectancy.

Nathan Keyfitz
Leader, Population Program

Contents

1. When the Average is Lower than Either Value 1
2. Two Stable Populations 3
3. Two Certain Statements 4
4. Estimating the Proportion of Cases where the Aggregate Life Expectancy Lies Outside the Range by Simulation 6
5. Sensitivity to the Growth Rate of the Stable Population 9
6. Two Regions with Migration 11
7. Alternative Approaches to Life Expectancy 13
References 15

AVERAGING LIFE EXPECTANCY

*Evgenii Andreev**, *Wolfgang Lutz***, *Sergei Scherbov***,

1. WHEN THE AVERAGE IS LOWER THAN EITHER VALUE

In a recent paper Scherbov and Lutz (1989) calculate scenarios on future mortality trends in the republics of the Soviet Union. The latest available information for 1986 indicates that female life expectancy was lowest in Turkmenistan with 67.8 years and highest in Lithuania with 75.9 years. The all-union average for female life expectancy lies at 73.8 years. One scenario assumed that mortality curves would be shifted up and down to converge to the same level of life expectancy in all republics by 2020. The results for 2019 (see Table 1) show that the difference between the republics is already minimal[1] but most surprisingly showed that the average life expectancy for the Soviet Union was lower than the life expectancy in any of the 15 republics.

A first analysis of this phenomenon reveals the following: it is not a mistake in the calculations but a quite possible result of the fact that different shapes of the force of mortality function in the various republics are aggregated using different age structural weights. In the Central Asian Republics infant and child mortality is very high relative to mortality at older ages, whereas in most European republics old age mortality is higher; this differential in the age pattern of mortality together with the fact that the age structure in Asia is much younger than in Europe (see Table 1) results in an aggregate life expectancy (calculated by adding up deaths and risk populations for each age group) that is lower than the life expectancy in the individual republics.

The above described case is certainly not the only empirical case in which the aggregate life expectancy lies outside the range of its constituents. It becomes intuitively clear that always when mortality curves cross and the age structure of the two populations is

---

*Institute of Statistics, USSR State Committee of Sciences, Moscow, USR
**Population Program, IIASA, A-2361 Laxenburg, Austria

[1] Variations are partly due to the iteration procedure used.

Table 1. Projected population distribution for the republics of the Soviet Union under the assumption of life expectancy converging by 2020; women, year: 2019.

Republic	Population (thousands)	% Total	Mean age	Life expectancy
RSFSR	83964.898	47.171	40.175	77.831
UkrSSR	29354.686	16.491	40.186	77.962
BelSSR	5796.118	3.256	39.261	77.980
UzbSSR	18526.979	10.408	28.699	77.750
KazSSR	11711.602	6.579	34.637	77.855
GrSSR	2818.267	1.583	39.631	77.933
AzSSR	4572.162	2.569	34.492	77.816
LitSSR	2143.815	1.204	39.698	77.976
MSSR	2820.524	1.585	35.628	77.781
LatSSR	1475.366	0.829	41.032	77.811
KirSSR	3431.611	1.928	30.183	77.713
TadSSR	4712.997	2.648	27.706	77.790
ArmSSR	2304.101	1.294	36.657	78.026
TurkmSSR	3435.693	1.930	28.017	77.730
ESSR	933.012	0.524	39.896	77.825
USSR	178001.812	100.000	37.559	77.632

sufficiently different, the phenomenon may occur. The phenomenon is also not restricted to life expectancy but might occur with a large number of non-linear averaging functions. To illustrate this point, consider the harmonic mean as a simple example of non-linearity. Table 2 lists two regions, each of them consisting of two groups with different intensities of a given event. While the harmonic mean of the two groups in Region 1 is .181 and in Region 2 .080 and aggregation of the two regions yields a mean of .217 which is clearly greater than any of the two regional means.

Table 2. Example of a simple case where the harmonic mean of the aggregate lies outside the range of the two individual means.

	Group A			Group B			
	Size	Events	Rate	Size	Events	Rate	Harmonic Mean
Region 1	100	95	.95	300	30	.10	.181
Region 2	200	10	.05	400	80	.20	.080
Both Regions	300	105	.35	700	110	.16	.217

Despite the possible generalization of the phenomenon to an infinite number of non-linear functions we will restrict the analysis in this research note to the concrete questions that arise when averaging life expectancy. We will ask for the conditions under which, generally, the average life expectancy calculated by aggregation of the age groups is different from the arithmetic mean of the individual life expectancies weighted by the proportions of births and, in particular, under which conditions the aggregate life expectancy lies outside the range of its components.

The question will be approached in several steps: first, a comparison between different stable populations is given; next, two theorems are proven for specific cases; to estimate the probability of the phenomenon for real populations the parameters of the Brass logit life table are modified step by step by simulation to cover all possible regional mortality patterns; finally, a stable multi-state case with migration is considered and conclusions are drawn.

2. TWO STABLE POPULATIONS

Illustrating the phenomenon in the case of two stable populations might make the underlying dynamics clearer than in the heuristic case of the Soviet republics described above. Assume a stable shrinking population with an intrinsic growth rate of -0.005 (population 1) and another with a growth rate of +0.010 (population 2). Population 1 which might stay for an aged European society has a relatively higher mortality above age 20 and lower under age 20 than population 2 which might resemble an Asian pattern. The mortality schedules plotted in Figure 1 result from the Brass logit life table with the following parameters: α = -1.100 and β = 1.162 in the case of population 1 and α = -0.627 and β = 0.230 in the case of population 2. Figure 2 gives the resulting stable age distributions.

In this example the life expectancy for population 1 is 71.14 years and that for population 2 is 73.08 years while the joint life expectancy of both populations lies with 69.26 years far below the expectancies of both individual populations. Loosely spoken, we may say that the combined population suffers from bad features of both sub-populations: below age 20 the joint force of mortality function is closer to the high child mortality in population 2 and above age 20 it is closer to the higher adult mortality of population 1. This follows directly from the age-distributional weights plotted in Figure 2 which are greater for population 2 at the young age and for population 1 at ages 25 to 75.

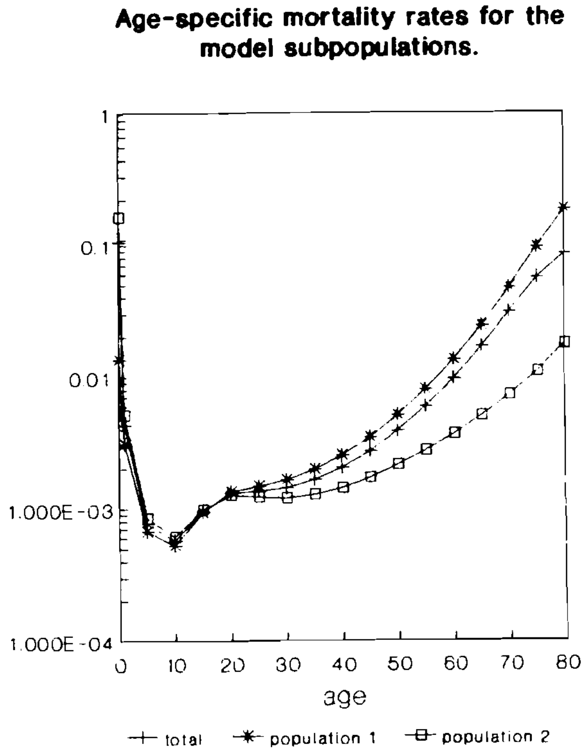

Figure 1. Two selected mortality schedules and their average.

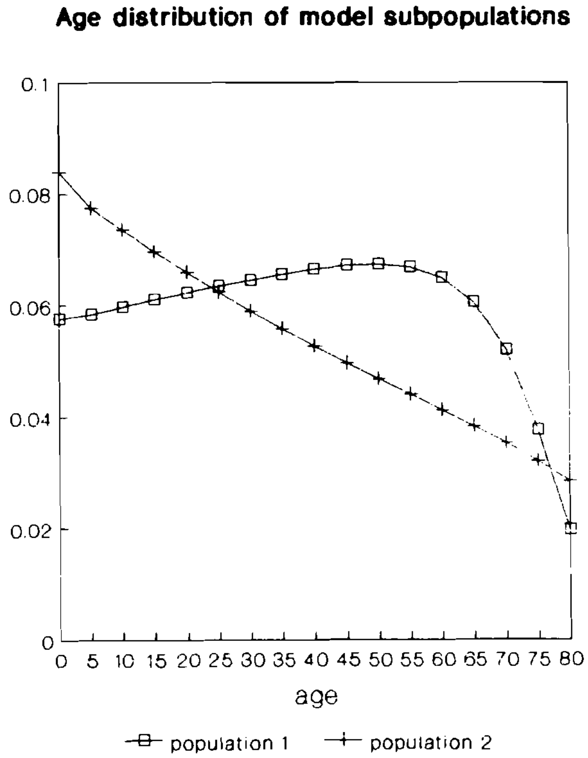

Figure 2. Stable age distributions for the two populations.

3. TWO CERTAIN STATEMENTS

Suppose that a population is a sum of two sub-populations with different mortality curves that are crossing. We assume that the size of the combined population in each age group is greater than zero. Let y_a denote the proportion of the first sub-population of the total population of age a. We shall consider the discrete mortality model with $a = 0,1,\dots,w$; where w is the highest age of survivors. Suppose that y_a is a random variable with linear distribution function $0 \le y_a \le 1$, and $Y = (y_0,\dots,y_w)$ is a random vector, whose coordinates are independent.

In the following theorem we shall show that in the case that expectations of life at birth are equal at some time interval, the probability that the aggregate life expectancy has the same value is zero.

A: Let $m_a^{(i)}$ denote the age-specific mortality rate in the whole population, first and second sub-population correspondingly where $a = 0,1,\dots,w$ and $i = 0,1,2$. Then

$$m_a^{(0)} = m_a^{(1)} \cdot y_a + m_a^{(2)}(1 - y_a) \qquad (1)$$

and there exists an age $\tilde{a}$, where

$$\frac{dm_{\tilde{a}}^{(0)}}{dy_{\tilde{a}}} = m_{\tilde{a}}^{(1)} - m_{\tilde{a}}^{(2)} \neq 0 \qquad (2)$$

because mortality curves are different.

The life expectancy is a function of age-specific death rates. We can assume that this function has continuous negative first partial derivatives. These assumptions follow from well-known life expectations calculation procedure and mathematical demography formulas (reference). So we can write for the aggregate population

$$e_0^{(0)} = f(m_0^{(0)},...,m_w^{(0)}) \qquad (3)$$

and

$$\frac{\partial f(m_0^{(0)} \cdots m_w^{(0)})}{\partial m_a^{(0)}} < 0 \qquad (4)$$

for each value a and $m_a^{(0)}$.

Substituting (1) to (3) and (2) in (3) to (4) we obtain that

$$e_0^{(0)} = f(Y) \quad \text{and} \quad \frac{\partial e_0^{(0)}}{\partial y_{\tilde{a}}} \neq 0$$

when $0 \leq y_{\tilde{a}} \leq 1$. Consequently, the set $\{Y, f(Y) = \mathit{const}\}$ is a surface, whose dimension is less than w. This means that the probability of life expectancies being equal is zero.

B: Now, suppose that expectations of life of sub-populations correspondingly are $e_0^{(1)}$ and $e_0^{(2)}$, $e_0^{(1)} < e_0^{(2)}$, but mortality curves are crossing. Then a set of ages exist, where $m_a^{(1)} < m_a^{(2)}$. We shall show that the probability of $e_0^{(1)} < e_0^{(0)} < e_0^{(2)}$ is less than 1.

Let $\tilde{y}_a = 1$ if $m_0^{(1)} < m_0^{(2)}$ and $\tilde{y}_a = 0$, if $m_0^{(2)} \leq m_0^{(1)}$, then $\tilde{m}_a^{(0)} \leq m_0^{(2)}$ for each a, and such an age $\hat{a}$ exists that $\tilde{m}_{\hat{a}}^{(0)} < m_{\hat{a}}^{(2)}$. Using (4) we obtain

$$\tilde{e}_0^{(0)} = f(\tilde{Y}) > e_0^{(2)} \quad .$$

Since $f(Y)$ is a continuous function there exists a $\epsilon > 0$, that $f(Y) > e_0^{(2)}$ if $|y_a - \tilde{y}_a| < \epsilon$. This means that the probability of $e_0^{(1)} \leq e_0^{(0)} \leq e_0^{(2)}$ is not greater than $1 - \epsilon^w$.

The analytical approach does not allow to estimate this probability more precisely and we do not know conditions under which $e_0^{(2)} < e_0^{(0)}$ (or $e_0^{(0)} < e_0^{(1)}$) and the variable Y takes a more realistic value than "0" or "1". Such an estimate, however, might be found by simulation.

4. ESTIMATING THE PROPORTION OF CASES WHERE THE AGGREGATE LIFE EXPECTANCY LIES OUTSIDE THE RANGE BY SIMULATION

The above approach could prove two important theorems but did not yield a realistic estimate of Y other than the fact that it must be smaller than 1. Another approach is through simulation. For this purpose we first have to define the possible range of mortality patterns to be considered. We chose to do this by using the Brass logit life table approach and expressing the complete range of regional model life tables at different mortality levels in terms of the two parameters α and β in the Brass model (Brass et al., 1968).

Figure 3 graphically presents the space defined by the four regional types of model life tables defined by Coale-Demeny (Coale and Demeny, 1966) and the five regional tables defined by the UN (United Nations, 1982) with life expectancy levels ranging from 60 to 80 in steps of 2.5 years. All together this results in 81 data points that spread over a range from β-values .85 to 1.4 and α-values between -1.5 to -0.5, with a concentration of points in the lower left corner.

Based on these 81 possible mortality patterns, couples of non-identical mortality schedules were randomly chosen and if the difference in life expectancy did not exceed 3.0 years they were combined with a set of two intrinsic growth rates randomly chosen from within the range -0.005 to 0.02. The three-year limit and the range of intrinsic growth rate were chosen in order to avoid cases unlikely to occur in reality. 100,000 such pairs of populations were randomly generated. Their distribution with respect to the relative position of the aggregate life expectancy is plotted in Figure 4. The x-axis gives the difference between the aggregate life expectancy (e_{tot}) and the lower life expectancy of the two given population ($e_{\min}$) divided by the difference between the two ($e_{\max} - e_{\min}$) (this indicator is denoted in the figure as e). On the y-axis the distribution of the 100,000 pairs is given on a logarithmic scale. The distribution is heavily concentrated around the arithmetic

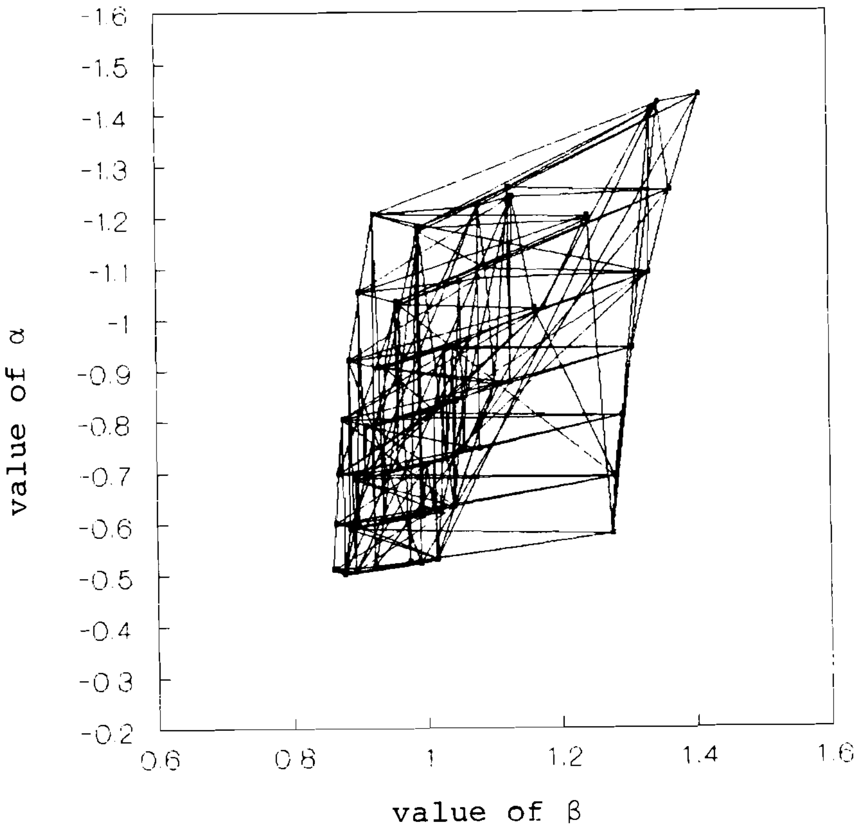

Figure 3. Space defined by the four Coale-Demeny and five UN regional life tables with life expectancy ranging from 60 to 80 in steps of 2.5 years and expressed in terms of the parameters α and β from the Brass logit life table.

mean of the two life expectancies. In other words, in the majority of likely empirical cases the aggregate life expectancy lies close to the center of the distance between the two life expectancies. In 80.5% of all simulated cases the aggregate mean lies within the range given by the two life expectancies. But in almost one fifth of all cases it lies outside the range.

This numerical estimate only gives a very rough indication of the empirical likelihood of the phenomenon that the aggregate life expectancy may lie outside the range of the life expectancies in the populations that constitute the total. It is subject to the above stated model assumptions and limitations. But the simulations show that the phenomenon initially described for the Soviet Union is not a marginal event that might be neglected empirically.

A further point of interest is the question to what extent the difference between life expectancies within a given couple of populations determines the likelihood of the mean to

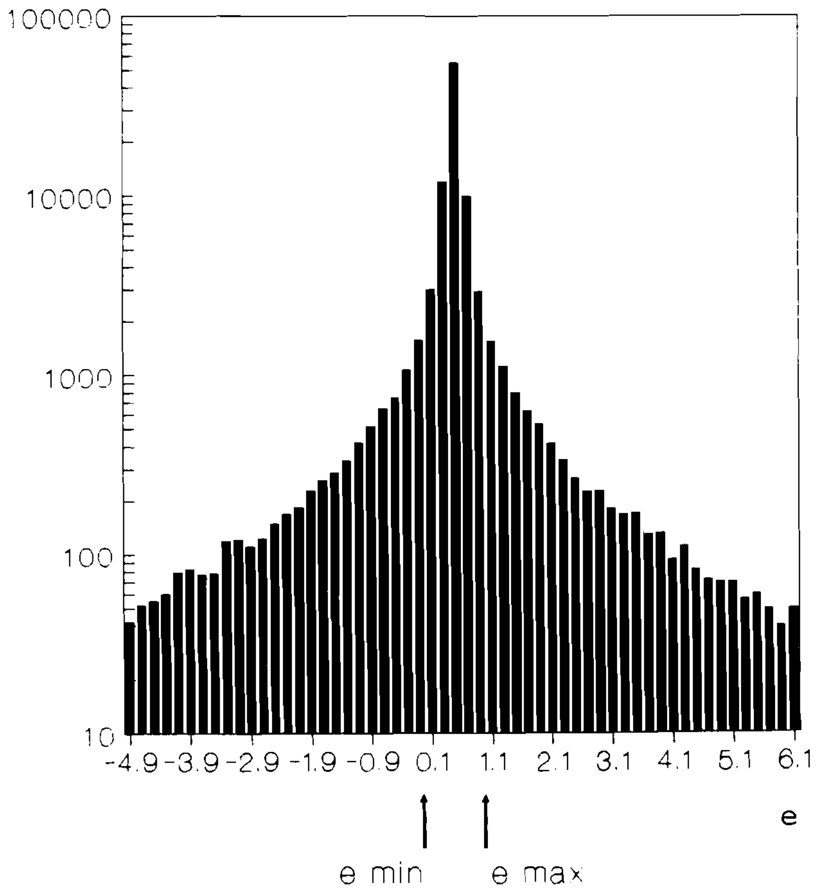

Figure 4. Distribution of the mean life expectancy of 100,000 randomly chosen pairs based on the 81 points shown in Figure 3 with a maximum difference in e_0 of 3 years.

lie outside the range. For this purpose a different set of simulations was carried out that reveals a very smooth and monotonically declining association between the differences in life expectancies and the frequency of the mean lying outside the range (see Figure 5).

In the case of equal life expectancies (difference zero) theorem A above showed that the aggregate expectancy must be different (with probability one) if the force of mortality functions are not identical. For a difference of half a year between the two life expectancies in about half of the 10,000 simulated cases the aggregate life expectancy lies outside the range of that in the two sub-populations. For a difference of two years this happens only in 20% of the cases. For greater differences this frequency seems to converge to a level of around 5% (for 10 years 0.0535). In Figure 5 each data point is based on a simulation of 10,000 cases.

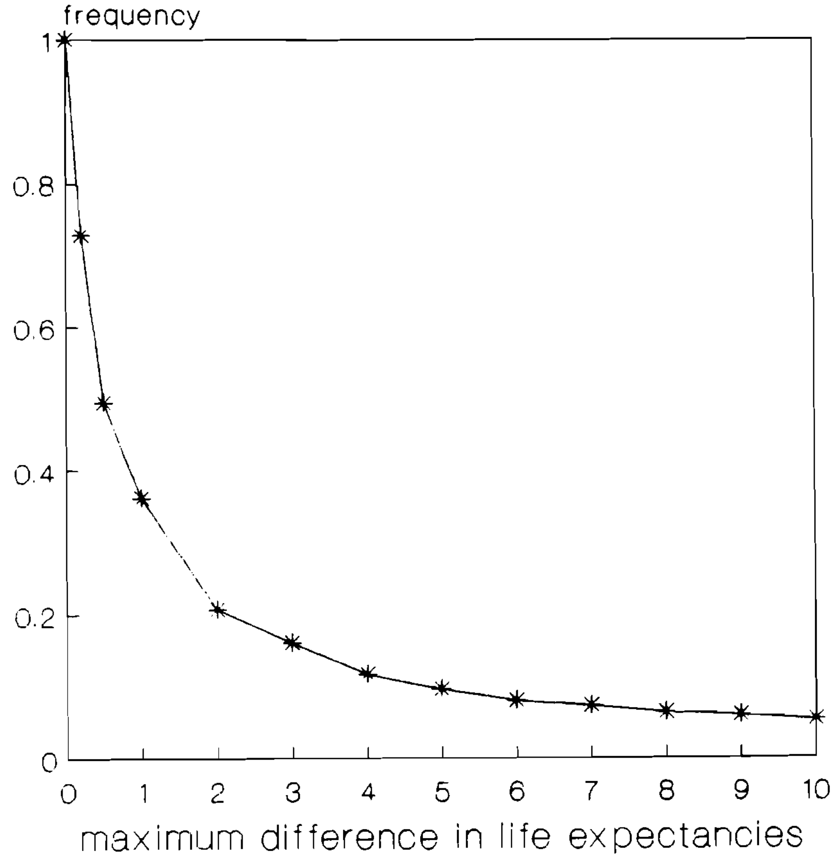

Figure 5. Dependence of the frequency of the mean lying outside the range on the difference between the two life expectancies considered (each point based on 10,000 simulated cases).

From the above described simulations we may conclude that the probability for the aggregate life expectancy to lie outside the range of its components is clearly a function of the difference in life expectancies. Of the empirically plausible cases with a difference of three or less years the frequency of lying outside the range is about 20%.

5. SENSITIVITY TO THE GROWTH RATE OF THE STABLE POPULATION

In the cases described above the central point of interest was in the difference between underlying mortality schedules. Assumptions on the intrinsic growth rates of the stable populations considered had to be made but were not studied systematically. In this section we want to study the sensitivity of the phenomenon with respect to changes in the intrinsic growth rates.

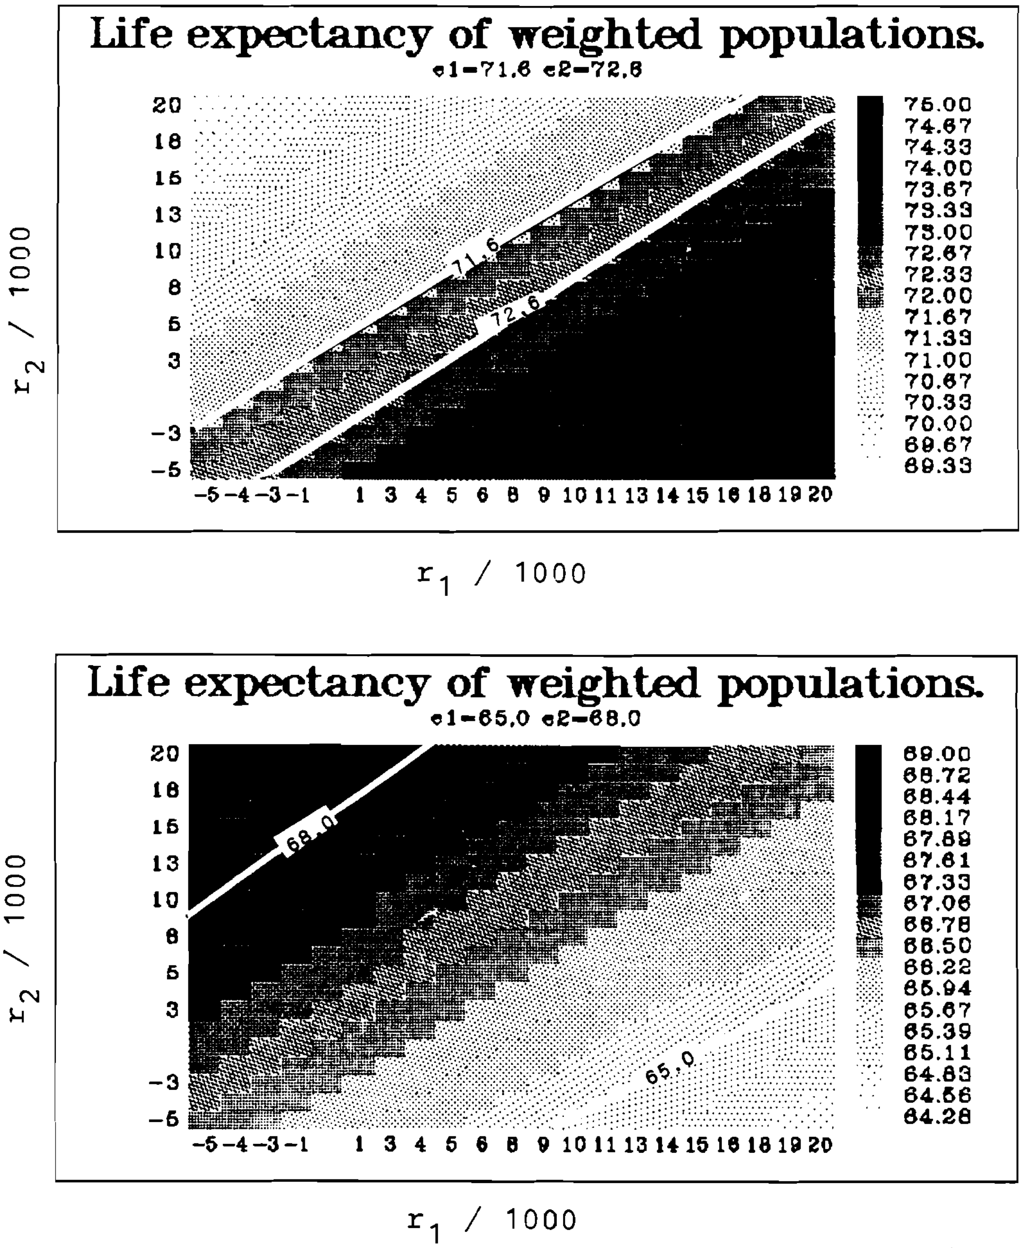

Figure 6. Shaded contour maps of average life expectancies in dependence on the growth rates of the stable populations considered for two selected pairs of life expectancy.

Figures 6a and 6b give shaded contour maps on the dependence of the aggregate life expectancy on the intrinsic growth rates r_1 and r_2 of stable sub-populations with randomly chosen mortality schedules (stated in the legend of Figures 6a and 6b). The intrinsic rates considered range from -.005 to +.020. The graph clearly indicates that only in the cases when the growth rates are similar to each other the aggregate life expectancy lies within the range of the expectancies of the two sub-populations (indicated in the graph by the flat area). In all other cases the aggregate expectancy lies outside the range.

Since it has been indicated in the previous section that the probability of the mean to lie outside the range depends on the difference between the life expectancies considered the pictures are given for differences of one year (Figure 6a) and three years (Figure 6b). Naturally the proportion of means lying outside the range is greater in the first case.

6. TWO REGIONS WITH MIGRATION

Are such cases where mortality schedules and stable age distributions cross likely to occur in neighbouring or otherwise related populations that are natural candidates for aggregation or is the described phenomenon without practical consequences? In the case of migration between two related regions the above described conditions may easily become true. Think of two provinces of a country or even urban and rural areas within the same province where it might well be the case that in one area child mortality is higher whereas in the other adult mortality is greater. And due to an unbalanced migration pattern on top of fertility differentials the population age structures are different. What will the averaging be like in this case?

Consider population 1 (urban) with a life expectancy of 70.69 and population 2 (rural) with 70.06. Fertility is somewhat higher in population 2 and there is a stream of migrants from 2 to 1 with an age profile that peaks at ages 20–25. The resulting population age structures are plotted in Figure 8. This might be typical for the situation of any town that draws from the surrounding rural areas.

Table 3 presents the life expectancies for periods and cohorts resulting from different ways of averaging. Aggregating the deaths and the risk populations from both sub-populations and calculating life expectancy based on the resulting age-specific death rates yields a joint life expectancy of 69.92 that is smaller than the expectancy in any of the two populations. An alternative and widely-used method is to average the life expectancies for the two populations by weighting them with the proportions of newborn. In this case the life expectancy must clearly lie within the range and comes to 70.15 in our case.

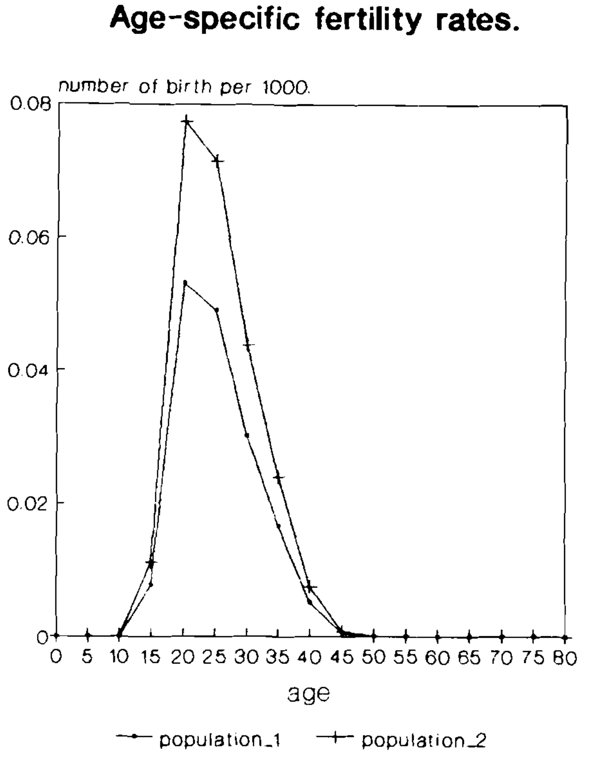

Figure 7a. Age-specific fertility rates in example considered.

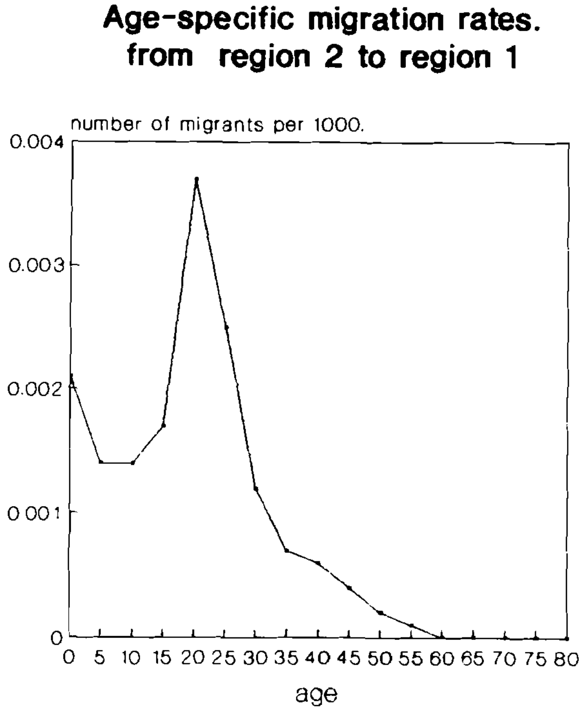

Figure 7b. Age-specific migration rates from region 2 to region 1 in example considered.

Table 3. Simulation of survivorship in a two-region population system.

Life expectancy	Total population	Population 1	Population 2
based on age-specific death rates	69.92	70.69	70.06
based on regional proportion of newborns	70.15	70.69	70.06
for real cohorts:			
total	69.92	70.69	69.79
in region 1	13.67	70.69	3.67
in region 2	56.25	0.00	66.13

Which of these two approaches that both seem to be straightforward and plausible is the "right one"? One criterion to find a decision is the question: which one resembles better the experience of a real cohort? A multi-state model for real cohorts results in a life expectancy for the joint population that is equal to that resulting from the first approach based on summing up age-group wise. However, for a cohort the aggregate life expectancy clearly must lie between the two cohort life expectancies for populations 1 and 2 because there may be no influence of the age structure. Indeed, Table 3 shows that this is the case. The difference to the period pattern described in the upper half of the table is that for population 2 the cohort life expectancy at birth is lower than the period life expectancy. This is a result of the migration from 2 to 1. Part of the population born in region 2 will not die in region 2 but in region 1. But since old age mortality is higher in region 1 than in region 2 the migrants will live shorter in region 2 than they would have in region 1. (This of course assumes that each person is exposed to the mortality risk of the region in which he lives.) On the average people born in region 2 live 66.13 years in that region and 3.67 years in region 1. Together this makes only 69.79 years. For region 1 there is no difference between cohort and period life expectancy because all people born in region 1 stay there until they die.

7. ALTERNATIVE APPROACHES TO LIFE EXPECTANCY

Some authors consider the calendar year life expectancy as a cumulative indicator of the mortality level; other authors see the period life expectancy as analogous to the cohort and consequently calculate the total life expectancy as a weighted mean of the sub-population's life expectancies weighted by the proportions of birth.

In the case of a stable system of closed sub-populations and stable mortality curves such period life expectancy really is equal to the life expectancy of the cohort. This is not true, however, in the case of migration between sub-populations. In a stable system of sub-populations with migration the cohort life expectancy is equal to the period aggregate life expectancy calculated by adding up deaths and risk populations for each age group in the way we did it in this paper. And in this case it may occur that the cohort life expectancy (being identical to the above described average period life expectancy) lies outside the range of the period life expectancies of the stable sub-populations such as in the example given in Table 3. This is in contradiction with the attempt to interpret period life expectancy in terms of real cohorts. Hence, weighting two period life expectancies in a

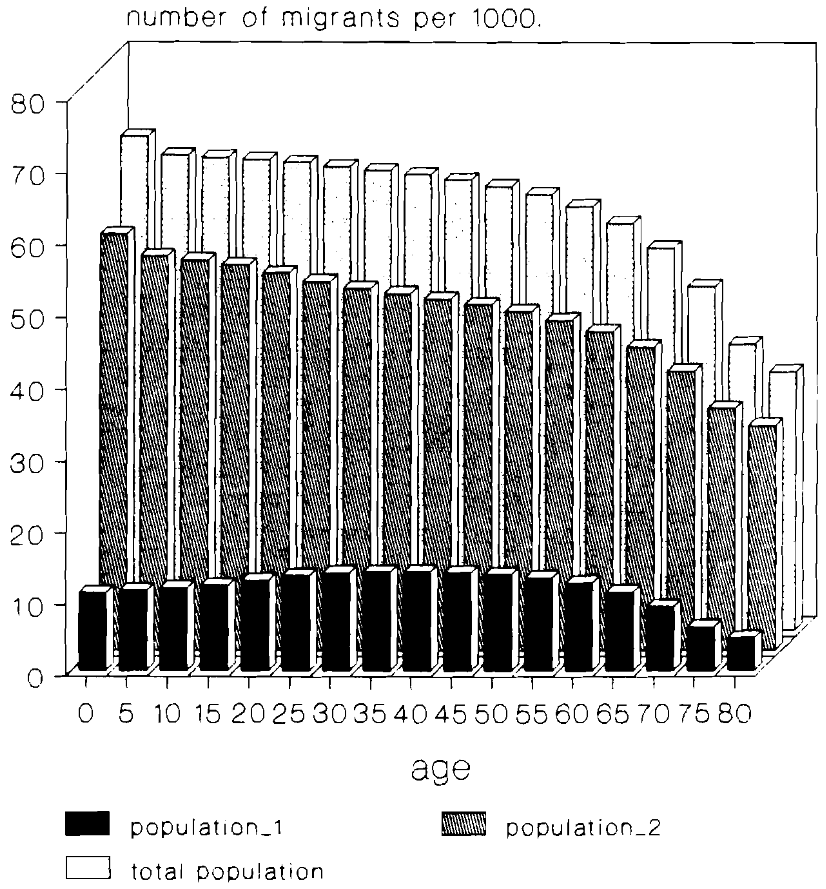

Figure 8. Population age structure in example considered (total population makes 1000).

cohort-like manner with proportions at birth is not correct in the case of migration.

For averaging life expectancies of closed populations (such as averaging male and female life expectancy) the situation is less clear. Whether to aggregate the populations age group wise or weight the expectancies by proportions at birth remains a more philosophical question depending on whether period life expectancy is just viewed as a summary indicator of mortality or as an analogon to real cohort life expectancy.

A possible next step of analysis is to apply the results of this paper also to models of cause-specific mortality analysis. In the cases of cause-deleted mortality tables, cause-specific death probabilities and mean ages at death from a certain cause, similar averaging problems appear.

REFERENCES

Brass, William and others (1968) *The Demography of Tropical Africa.* Princeton, New Jersey: Princeton University Press.

Coale, A. and P. Demeny (1966) *Regional Model Life Tables and Stable Populations.* Princeton, New Jersey: Princeton University Press.

United Nations (1982) *Model Life Tables for Developing Countries.* United Nations, Sales No. E.81.XIII.7.

Recent Working Papers Produced in IIASA's Population Program

Copies may be obtained at a cost of US$ 5.00 each from IIASA's Publications Department.

WP-88-10, *On the Concentration of Childbearing in China, 1955–1981* by W. Lutz. February 1988.

WP-88-13, *Beyond "The Average American Family": U.S. Cohort Parity Distributions and Fertility Concentration* by M. King and W. Lutz. March 1988.

WP-88-23, *Understanding Medical and Demographic Trends with MEDDAS* by M. Rusnak and S. Scherbov. April 1988.

WP-88-32, *Kinship Patterns and Household Composition of the Elderly: Hungarian Women, 1984* by D. Wolf. April 1988.

WP-88-36, *"DIAL" – A System for Modeling Multidimensional Demographic Processes* by S. Scherbov and V. Grechucha. May 1988.

WP-88-44, *Kin Availability and the Living Arrangements of Older Unmarried Women: Canada, 1985* by D. Wolf, T. Burch, and B. Matthews. June 1988.

WP-88-46, *Population Futures for Europe: An Analysis of Alternative Scenarios*, by D. Wolf, B. Wils, W. Lutz, and S. Scherbov. June 1988.

WP-88-90, *Comparative analysis of Completed Parity Distributions: A Global WFS-Perspective*, by W. Lutz. October 1988.

WP-88-104, *Future Regional Population Patterns in the Soviet Union: Scenarios to the Year 2050*, by S. Scherbov and W. Lutz. November 1988.

WP-88-120, *AIDS and HIV Surveillance in Europe*, by M. Artzrouni and G. Heilig. December 1988.

WP-88-124, *DDMSLT: A Computer Program for Estimating the Duration-Dependent Multistate Life Table Model*, by C. Calhoun. December 1988.

WP-89-05, *Multi-State Analysis of Family Dynamics in Austria: Scenarios to the Year 2030*, by W. Lutz and S. Scherbov. January 1989.

WP-89-06, *The Demographic Dimensions of Divorce: The Case of Finland*, by W. Lutz, B. Wils, and M. Nieminen. January 1989.

WP-89-18, *Markets as Queues, with an Application to Education*, by R. Boylan. February 1989.

WP-89-19, *Living Arrangements and Family Networks of Older Women in Italy*, by D. Wolf and A. Pinnelli. February 1989.

WP-89-27, *Reconciling Economic and Ecological Theory on Population*, by Nathan Keyfitz. March 1989.

WP-89-28, *Multistate Life Table with Duration-Dependence: An Application to Hungarian Female Marital History*, by Alain Belanger. April 1989.

WP-89-34, *Recent Trends in Living Arrangements in Fourteen Industrialized Countries*, by J.-P. Gonnot and G. Vukovich. May 1989.