

Decomposing the Evolution of Frailty in the China Multi-Generational Panel Dataset, 1749-1909:

Jim Oeppen*

Session 126: Historical Mortality Patterns

Annual Conference of the Population Association of America
March 31st – April 2nd, 2011. Washington, D.C.

KEYWORDS: Frailty; Mortality; Decomposition; China; Longitudinal; Historical Demography; Selection; Missingness; Attrition; Recruitment; Kinship; Microsimulation.

Abstract

A method of ‘scoring’ correlated frailty for a household is defined by comparing observed mortality with the expected mortality of its members as determined from an aggregate life-table. Changes in average frailty over time at the population level are decomposed into three components: change experienced by households measured on two consecutive occasions; a selection effect if the frailty of new households is consistently different from the established group; and a corresponding effect if the households that drop out differ from the ones that continue. Frailty scoring and decomposition are applied to households in the China Multi-Generational Panel Dataset, 1789-1885. The results suggest that the apparent stability of mortality in the population conceals a consistent evolutionary dynamic at the household level.

*Max Planck Institute for Demographic Research, Rostock, Germany. email: oeppen@demogr.mpg.de

1 Frailty and Mortality Concentration.

Individuals differ in their ability to withstand aging and mortality. In the 1970s demographers realised that populations at risk to die were not made sufficiently homogeneous by classifying them by age and sex, or additional factors, and that unobserved heterogeneity had serious consequences for the understanding of mortality change across age and time (Vaupel et al., 1979; Vaupel and Yashin, 1985). The authors conceptualised frailty as a person-specific positive multiplier applied to the average force of mortality. In addition, it was realised that this frailty parameter might be correlated across individuals who shared certain genetic or environmental characteristics.

One of the problems in measuring frailty is that mortality is not a repeated event. Frailty in reproduction had been conceptualised and quantified by Gini in the 1920s, since it is possible to determine from fertility behaviour if a couple has above- or below-average fecundability. To make progress with frailty, it is usually assumed that individual frailty comes from a known distribution, and that each parameter is age-invariant for an individual.

In this paper we assume that it is reasonable to ‘score’ correlated frailty for a household by comparing its observed mortality with the expected mortality of its members determined from an aggregate life-table. Effectively frailty becomes a residual from a model.

The household has always been considered as an environment for production and reproduction (Becker, 1981), but it has also been recognised as a setting for nuclear hardship and shared mortality factors (Derosas and Oris, 2002). Quantitative analyses of mortality concentration within sibling groups have been conducted by Das Gupta (1990); Zaba and David (1996) and later authors. Despite these efforts, our understanding of the household as a factor in mortality is weaker than the inferences we draw from age, sex, marital status, occupation, and other well-known covariates. A related interest in the inter-generational association between the longevity of family members has also had a long history within bio-gerontology, from Karl Pearson through Raymond Pearl, but these studies typically ignore the household dimension. In general the within- and between-generation associations have been studied separately, although Vandezande et al. (2010) is an exception. Criticisms of the field might be that the approach is static, often limited to specific age-groups such as infants, and concerned with a simple linkage, such as mother-child, or grandmother-grandchild (see Sabo and Chaganty, 2010, for more complex linkages).

Households experience multiple forms of change. Some that might be observed like formation, extinction, fusion and fission, and others that may be unobservable, but all may contribute to changing frailty. If we have longitudinal data we can assess the dynamics of frailty, but it will be useful to try to decompose change into important components. To do this we adapt the Price Theorem from evolutionary biology and decompose change in average frailty into three components: that experienced by households measured on two

consecutive occasions; a selection effect if the frailty of new households is consistently different from the established group; and a corresponding effect if the households that drop out differ from the ones that continue.

If the dynamics of frailty can be measured effectively, it presents several opportunities. We conceptualise the mortality transition within an Age–Period–Cohort framework. Did the frailty distribution shift as a whole, or did the selection process against frail households intensify, with the additional factor of lower fertility, so that the frail end of the distribution disappeared? The debate about the Rose hypothesis concerning intervention in morbidity and mortality processes is really a question about dynamic frailty (Rose, 2001). Do we have to shift the whole distribution, or is there a meaningful target group? Can a social planner reinforce a benign dynamic that already exists?

2 Scoring Frailty.

Frailty is the difference between the mortality experience of an individual and the mortality of a reference population. To define a suitable measure, we assume that we can count the number of deaths occurring over a fixed period between cross sections, such as censuses or surveys, which is known as panel-count data. For panel count data where the interval is n years, mortality is defined as ${}_nQ_x = 1 - \frac{L_{x+n}}{L_x}$, the complement of inter-censal survival.

Figure 1 shows a schematic version of the mortality scoring principle adopted. The N individuals in a household or other aggregate are coded as $y_i = 1$ if they died in the interval and zero otherwise. A corresponding expected value $\hat{\mu}_i$ can be obtained by selecting from the life table the appropriate ${}_nQ_x$ matched by age x and sex.

Denote the observed number of deaths in the group as $D = \sum y_i$. Let the expected number of deaths \hat{E} be the sum of the individual-based expectations $\hat{E} = \sum \hat{\mu}_i$. Although we have derived the expected counts of deaths from conventional life tables classified by age and sex, other factors such as marital status could be used. Covariates could also be incorporated using parameters estimated from survival analysis, such as the proportional hazard model, to estimate the \hat{E}_i .

The discrepancy $D - \hat{E}$, which we are interpreting as an index of frailty, is effectively the residual from the model, but we cannot assume its properties are known. Cameron and Trivedi (1998, p.141) wrote that ‘For count data there is no one residual that has zero mean, constant variance, and symmetric distribution.’

With larger aggregates it would be conventional to consider the Standardised Mortality Ratio, $SMR = \frac{D}{\hat{E}}$, but for panel counts with a low risk and a short period the majority of the numerators are zero. One possible measure of frailty is the Pearson residual r . Let the mortality risk for the group be $y = \frac{D}{N}$ and the expected value $\hat{p} = \frac{\hat{E}}{N}$, then the

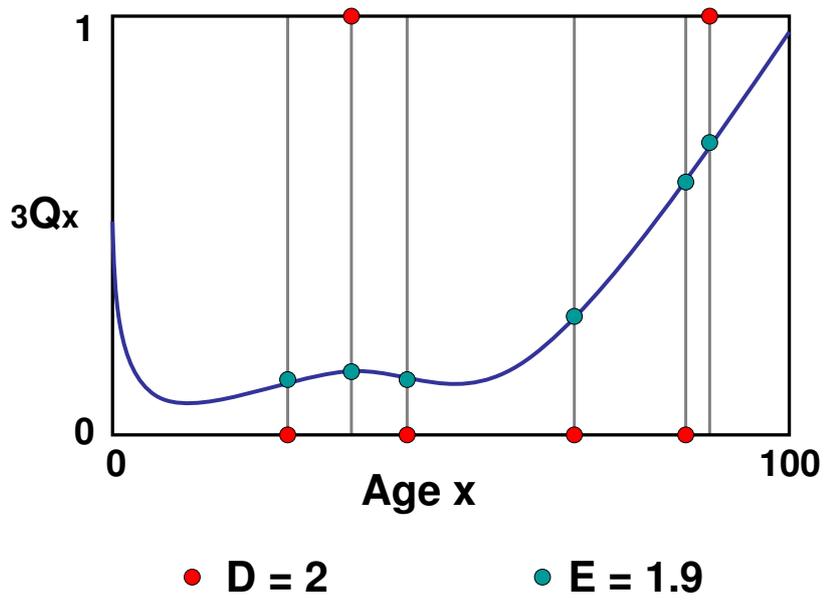


Figure 1. Panel Count of Observed and Expected Deaths for a Population Aggregate.

residual is

$$r = \frac{y - \hat{p}}{\sqrt{\hat{p} * (1 - \hat{p})}} \quad (1)$$

However, this measure of the ‘surprise’ at the gap between y and \hat{p} is independent of N and more appropriate for binary data. An alternative is to consider the Poisson *deviance residual* (Cameron and Trivedi, 1998, Eqn. 5.4)

$$\hat{d}_i = \text{sign}(D - \hat{E}) \sqrt{2D \ln(D/\hat{E}) - (D - \hat{E})} \quad (2)$$

when by definition $0 \log(0) \equiv 0$. If we apply this to the simulated data mortality described in the next section, \hat{d}_i has the useful properties that it is approximately normally distributed and centred on zero. A very similar result is achieved with the Poisson Anscombe residual (Cameron and Trivedi, 1998, Eqn. 5.6)

$$\hat{a}_i = \frac{1.5(D^{2/3} - \hat{E}^{2/3})}{\hat{E}^{1/6}} \quad (3)$$

However, these properties are lost when N is small. The distribution becomes bimodal with an extreme negative peak corresponding to $D = 0$ and a flatter positive peak when $D > 0$.

Another possibility is to recall the equation for R-squared

$$R^2 = 1 - \frac{SS_{err}}{SS_{Total}} \quad (4)$$

and consider how these error terms in the numerator are defined in the estimation of a pseudo R-squared for Poisson regression models for count data (Cameron and Windmeijer, 1996; Mittlböck and Heinzl, 2001).

$$\hat{F} = \text{sign}(D - \hat{E}) \begin{cases} [D * \log(D) - D] - [D * \log(\hat{E}) - \hat{E}] & \text{if } D > 0 \\ \hat{E} & \text{if } D = 0 \end{cases} \quad (5)$$

When no deaths are observed, $D = 0$, so that the estimated frailty \hat{F} of the group is simply $-\hat{E}$.

Figure 2 shows the behaviour of \hat{F} under conditions of homogeneous mortality with ${}_3Q_x = 0.05$. The lowest contour shows the ‘evidence’ for frailty, or ‘surprise’ at the outcome, if $D = 0$. When mortality is low and $N = 1$, observing zero deaths gives little evidence for negative frailty. However, if $N = 100$, zero deaths is evidence of robustness. When $D = 1$ and $N = 20$, $\hat{F} = 0$, as expected with this level of mortality. If $N = D = 10$, the leftmost point of the highest contour, this is very strong evidence of frailty.

The figure suggests that this frailty indicator will be centred on zero and unimodal, but it will be asymmetric. To explore its properties further we generate simulated data.

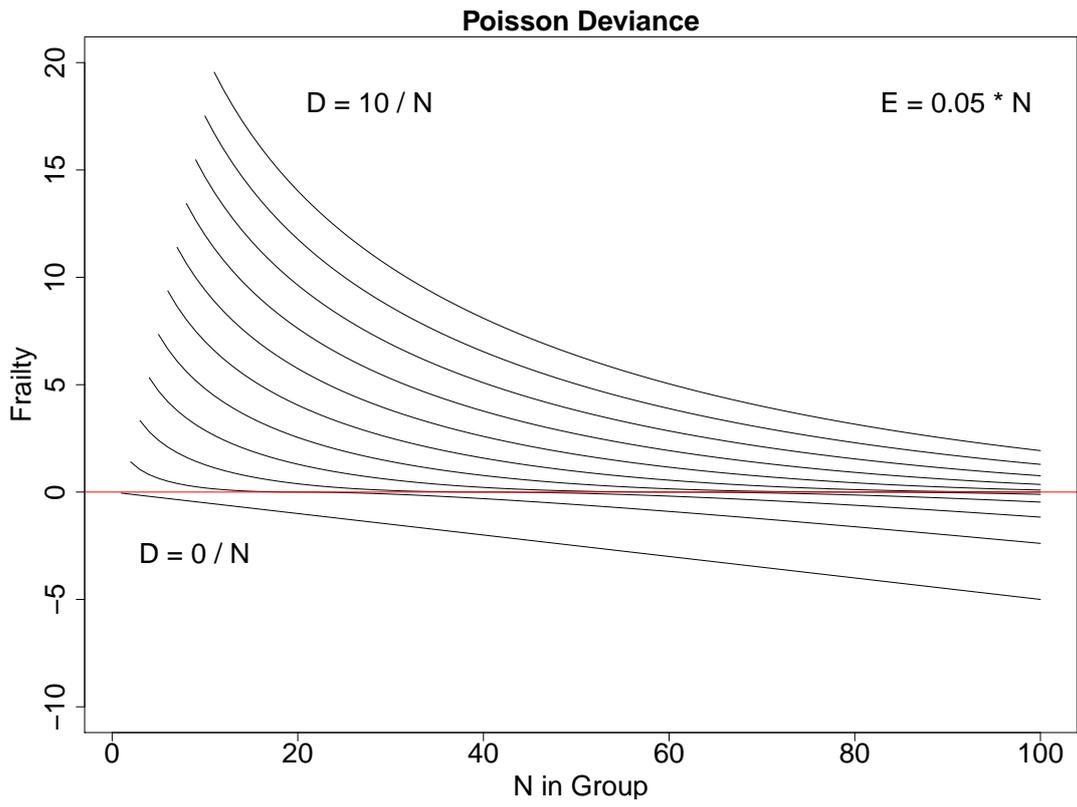


Figure 2. Contours of frailty \hat{F} for combinations of group size N , observed deaths D and expected deaths E .

2.1 Simulation Results.

Demographic micro-simulations were conducted using the CAMSIM program (Smith and Oeppen, 1993). Male egos were simulated with two ascending and two descending generations. This allows us to generate panels of the kinship structure experienced over ego's lifetime. Only one descendant generation was simulated from each ascendant one. This means that, using European kinship terminology, first cousins and their spouses are included but higher order cousins and their families are excluded. The demographic parameters for the simulation were taken from English family reconstitution results for the period 1800–1837 (Wrigley et al., 1997). A zero risk of permanent celibacy was assumed for males, although they could die before marriage as a result of the competing risk of mortality. Kinship panels were constructed for the kin-set of each ego at a randomly chosen birthday between ages 0 and 3 inclusive, and every three years thereafter up to his death. Egos that died before their first panel were excluded.

There is no left-truncation effect on mortality. Ego's parents have unconditional mortality after the birth of ego, and spouses are only included in the scoring process after they marry into the kinset.

The first experiment employed homogeneous mortality with respect to both age and sex $Q = 0.04$, corresponding to $e_0 = 24.2$. Figure 3 contrasts the distribution of Pearson residuals with the Poisson deviance. It is clear that the former does not match the objectives of zero-centred symmetry. The Poisson deviance is not quite symmetric either, and the slight bulge, at -1 in this case, is observed in other experiments.

To add realistic mortality, Coale and Demeny (1966) Model North was chosen because it has the strongest empirical base at high mortality, although the tables are largely extrapolations for life expectancies below 35. The mortality schedules were derived using MORTPAK (United Nations Population Division, 2003). The application MATCH allows the selection of a sex-specific life-table with a given value for an age-specific parameter such as e_x from a 'family' of tables. UNABR takes a table of q_x values for conventional age-groups and 'unabridges' the table to single-year age-groups.

Figure 4 shows the effect of scoring the same simulated data with lifetables with lower and higher life expectancy. The relocation on the frailty axis is as expected: if a population with $e_0 = 40$ is scored with a life-table having $e_0 = 30$ fewer groups should be regarded as frail. However, it is unfortunate that this shift is also associated with a change in the shape of the distribution.

Heterogeneity can be conceptualised as the effect of a mixture of two or more homogeneous distributions. Figure 5 shows the score distribution arising from equal numbers of egos having kinset mortality derived from two life-tables.¹ The expected scores for the mixed population are derived from the life-table for all individuals. In this case the zero-centring is preserved although there is a slight distortion of symmetry.

¹Although the ego numbers and therefore the number of scores are equal, the sizes of the kinsets differ because of the greater survival in the one group.

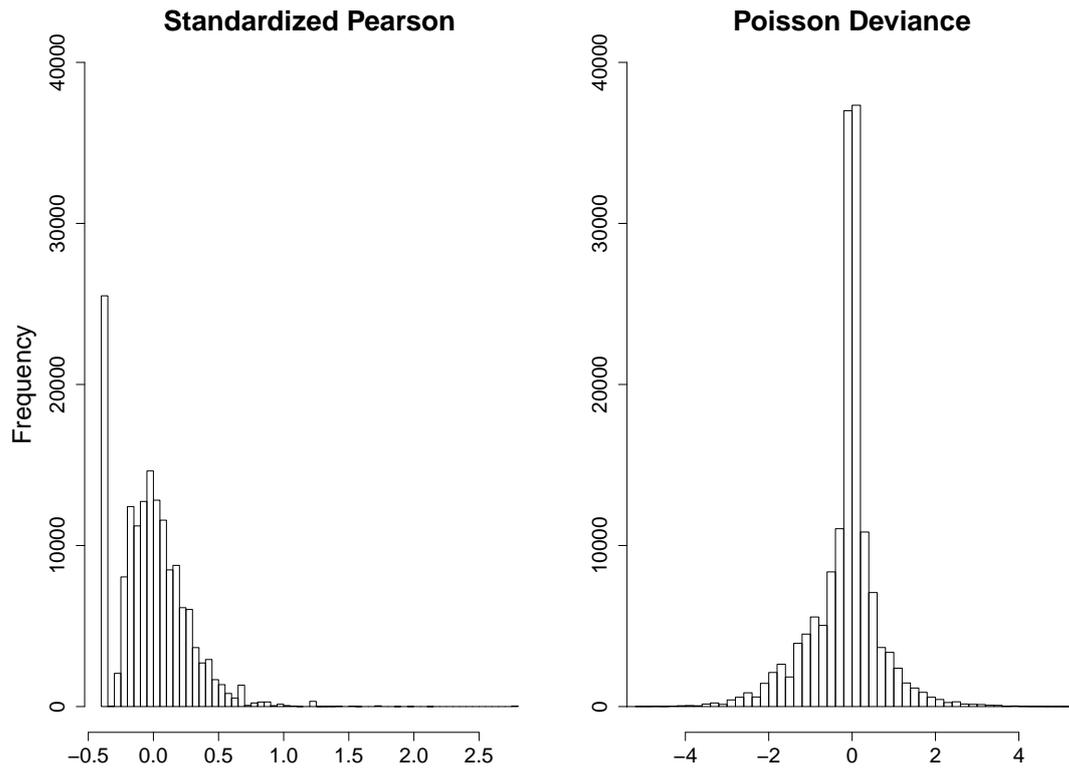


Figure 3. Frailty Indicators from Simulated Panels with Homogeneous Mortality.

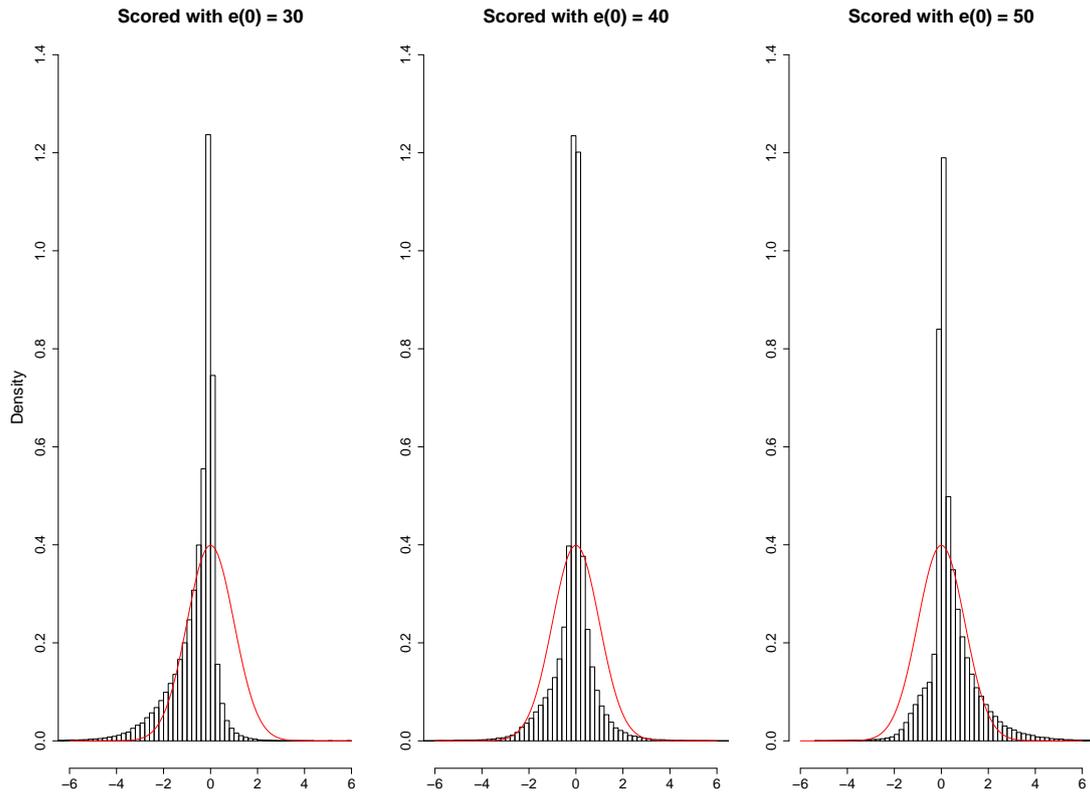


Figure 4. Effect of Life-Table used for Scoring; simulated panels with $e(0) = 40$. Unit normal shown in red for visual comparison

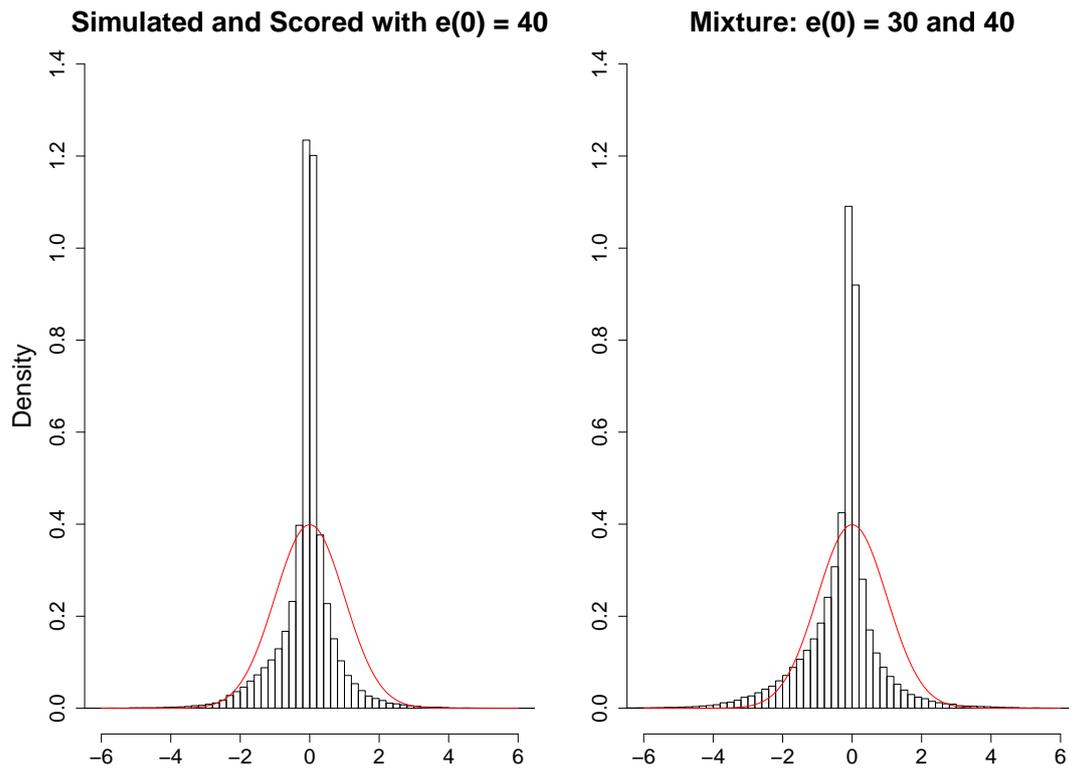


Figure 5. Frailty Mixture. Left panel simulated and scored with $e(0) = 40$. Right panel: simulated with mixture of $e(0) = 30$ and 50 ; scored with average life-table.

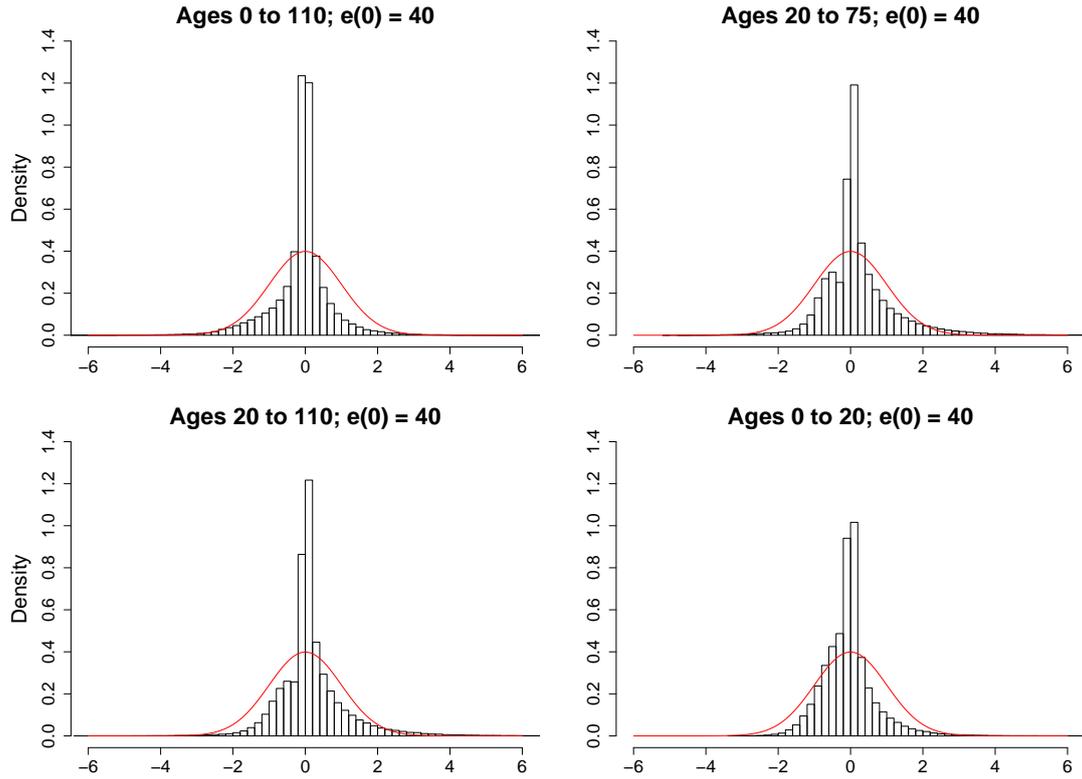


Figure 6. The effect of limiting the age-range analysed.

Finally, we calculate the effect of limiting the analysis to a part of the age-range. In Figure 6, all three age restrictions seem to produce a slight shift towards greater robustness in the negative half of the distribution.

3 Decomposing Change.

With the expectation that non-random selection is acting on heterogeneous frailty, we examine a method to establish the scale of the selection effect. We consider first the case where there is attrition of households but no recruitment. Change at the aggregate population level P can be exactly decomposed into average change for the households staying in the study s plus a compositional change due to selective disappearance of certain households d :

$$P = s + d. \tag{6}$$

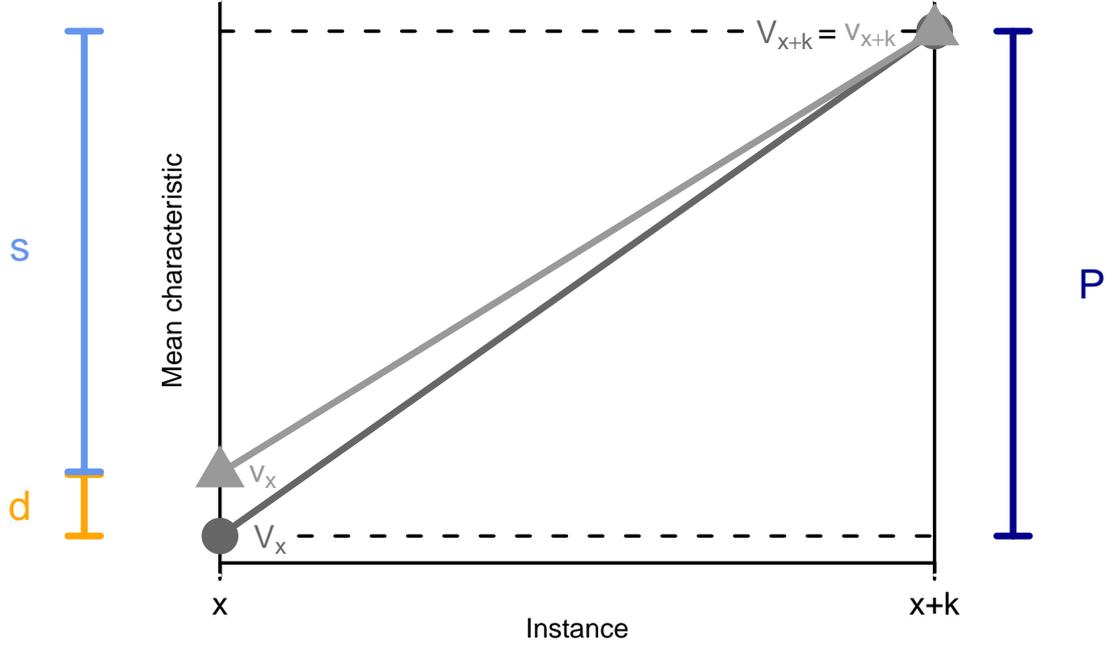


Figure 7. Adapted from Rebke et al. (2010, Fig. 1). Graphical representation of equation (6). Refer to the text for an explanation of the symbols.

In this equation the population change is given by $P = V_{x+k} - V_x$, the difference between the average V_x of a considered characteristic of all households at one instance x and the average V_{x+k} at the next instance $x+k$, where k denotes the interval between the two instances. The averages can also be scaled in which case they represent proportions. The average change for the households staying in the study represented by $s = v_{x+k} - v_x$, is the respective difference for these staying households at one instance (v_x) and the next instance (v_{x+k}). The term measuring change due to selective disappearance is given by $d = v_x - V_x$, which is the difference between the two averages at the first instance x , i.e. the average for the staying households v_x and the average for all households V_x . The decomposition was originally shown in Rebke et al. (2010).

Figure 7 is a graphical illustration of equation (6). In this example, household change s between x and $x+k$ is less than the population change P . The remaining increase d is accounted for by the loss of households whose average frailty at time x was below the population mean.

Equation (6) can be extended to account for new households entering the study population at a later occasion (Rebke et al., 2010). This might lead to another change in the composition of the population due to selective appearance a of households. This change is captured by adding the term $a = V_{x+k} - v_{x+k}$ to equation (6):

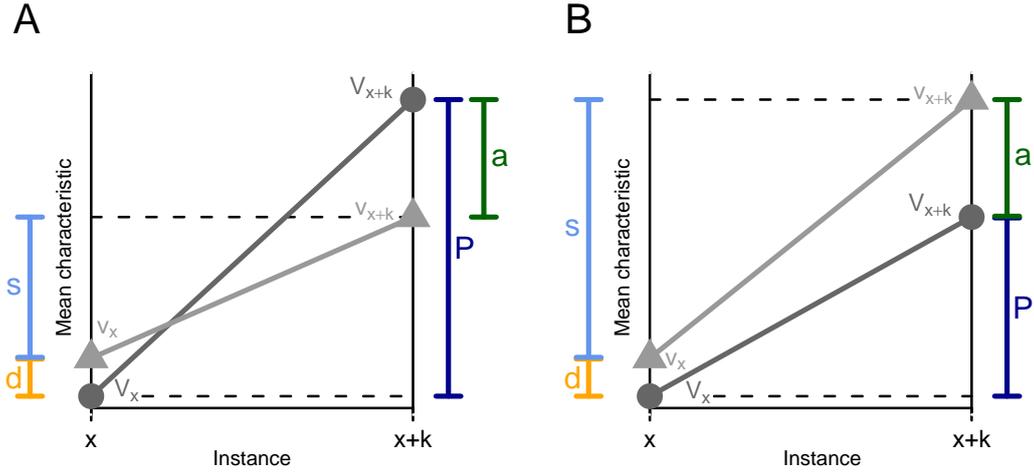


Figure 8. Adapted from Rebke et al. (2010, Fig. S1). Graphical representation of equation (7). In panel (A) newly entering households have a higher mean than the households already in the study and still participating at instance $x+k$. In panel (B) newly entering households have a lower mean than the households already in the study and still participating at instance $x+k$. Refer to the text for an explanation of the symbols.

$$P = s + d + a. \quad (7)$$

Graphical illustrations of this extension are given in Figure 8. In both panels, s is positive showing that the ‘stayer’ population raises its average frailty, and the effect of attrition d is positive. Negative values of d are possible, which means that low scores are inversely related to the probability of dropping-out. In panel A, the recruits have a higher average at $x+k$ than the retained households.² In panel B, the recruits have a lower average score than the ‘stayers’ so a is negative.

The theoretical distributions of these quantities are unknown. To give a representation of statistical significance, approximate 95% confidence intervals have been calculated for the components of change via bootstrapping, using the *boot* package (Canty and Ripley, 2009; Davison and Hinkley, 1997) in the R software environment (R Development Core Team, 2009).

Although we have provided a graphical representation of the decomposition, it can be related formally to the Price theorem, one of the fundamental representations of selection effects in population genetics (see Price, 1970; Gardner, 2008). Treating each panel as analogous to a generation, equation (7) is the Price equation when reproduction is a binary outcome (see Coulson and Tuljapurkar, 2008, eqn. 10). Inter-generational

²This is the manager’s dream: to get rid of below-average performers, improve the retained staff, and find recruits that raise the average skill level.

change in a trait (P) is decomposed into the effects of selection (d) and transmissibility (s). The derivation is explained in detail in Rebke (2010). Kerr and Godfrey-Smith (2009, eqn. 1) extend the formal treatment to handle recruits.

The decomposition can be used for both continuous variables and proportions defined at fixed points, which do not have to be regularly spaced. Normally these would be time points, durations or ages, but they can also be ordinal stages, such as pregnancy order. It is also suitable for chain-event, current status and panel-count data, and rates defined over a preceding interval. It is not suitable for irregularly observed data, unless they can be meaningfully converted to cross-sectional data or rates. It is an important feature of the method that it is suitable for data with non-monotone missingness.

The Price Theorem is exact, completely general and can be applied to any quantity. For example, the variability of the data within each panel or census may also be a quantity of interest that is subject to selection effects.

4 The China Multi-Generational Panel Dataset.

Large-scale longitudinal datasets are rare, so the China Multi-Generational Panel Dataset is exceptional in having 1.5 million observations for three-year panels over the period 1749 to 1909: recording over a quarter of a million individuals (Lee and Campbell, 2010).³ Households and individuals have unique identifiers, and the household head is identified, so that households can be linked between panels (a method is described below).

The principal disadvantages of the data are the lack of fertility recording, the under-registration of young children, and of unmarried females, and the complex missingness patterns (described below). Figure 9 shows the population at risk for mortality by age. The under-recording of younger people is obvious, and is overlain by a systematic pattern of age-misstatement for males.

4.1 Data Selection.

For mortality analysis there are 859,674 person-panels where `AT_RISK_DIE == 1` with 39% being female. About 61% of all records meet the `AT_RISK_DIE` criterion and this is the same for both sexes and over the age range 15–80 years. Under age 15 the proportion is somewhat lower for both sexes and falls to about 59% for girls. The average proportion at risk in each household is 68% and it is reassuring that the number at risk seems to be linearly related to the total number in the household over the full range of the latter.

We have made a number of exclusions:

³The website contains an extensive codebook and a long list of publications including Lee et al. (1995); Lee and Campbell (1997).

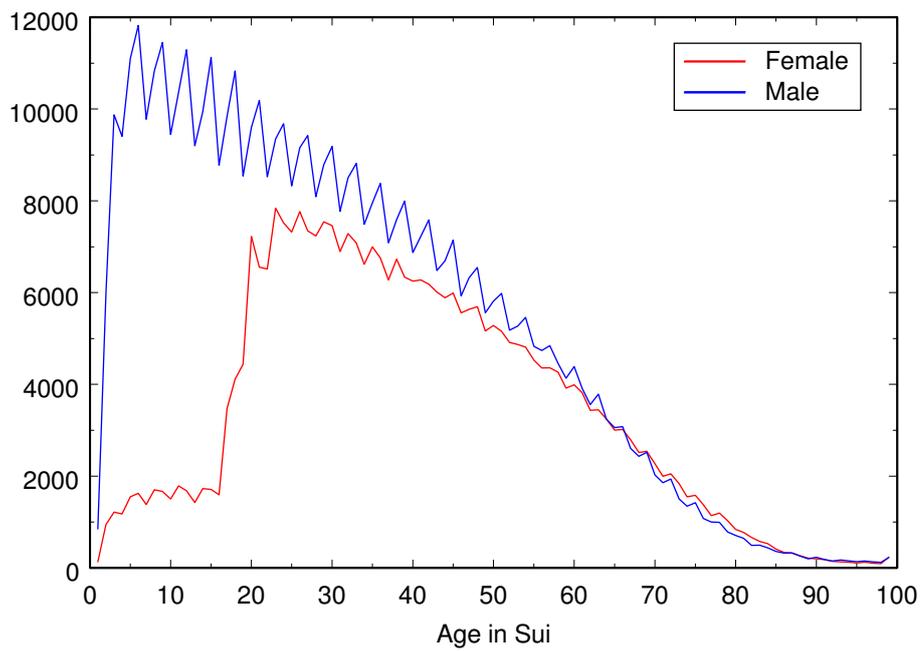


Figure 9. Number of Persons At Risk of Mortality.

1. Exclusions made at person-panel level:
 - `AT_RISK_DIE` $\neq 1$
2. Exclusions made at Household level:
 - Record 1 in household not an ego. [87]
 - Sequence error in household numbering. [38]
 - Multiple egos in household. [3]
3. Exclusions made at Person level:
 - Person reaches age 99 or greater but no death recorded. [715]
4. Right censoring of Personal history:
 - First abscond date. [2752]
 - First occurrence of sex change. [366]
 - `NEXT_DIE` `== 1` but later `AT_RISK_DIE` `== 1`. [18]

4.2 Comparison with Model Life Tables.

Figure 10 contrasts the observed ${}_3Q_x$ with model patterns derived via MORTPAK using MATCH and UNABR (United Nations Population Division, 2003). The selected model tables were not optimised for pattern or $e(1)$, but merely chosen to approximate the observed patterns over the age range 20 to 75. The U.N. South Asian pattern for males with $e(1) = 40$ corresponds to $e(0) = 30.6$. Other U.N. patterns for males show a much more marked ‘accident hump’ than the Chinese data. The U.N. General pattern for females with $e(1) = 45$ corresponds to $e(0) = 39.3$. The Female South Asian pattern at the same level also gives a close approximation with $e(0) = 37.3$. The dangers of assuming patterns for far eastern mortality are discussed in Zhao (2003) and the comments on that paper.

It is clear that the mortality patterns are implausible at older ages. Excluding apparent ‘immortals’ does not raise ${}_3Q_x$ to plausible levels. Two issues we have not addressed in the exclusion list are inconsistent age-reporting across the lifetime of an individual, and inconsistency of live ages with the death age. These show systematic and different patterns. The application of these patterns of age misreporting to model life-tables does result in a decline in the rise of ${}_3Q_x$ at higher ages. However, even at very high levels such age-independent shifts of numerator and denominator are insufficient to explain the observed decline in ${}_3Q_x$ because gains and losses at each age are largely offsetting.

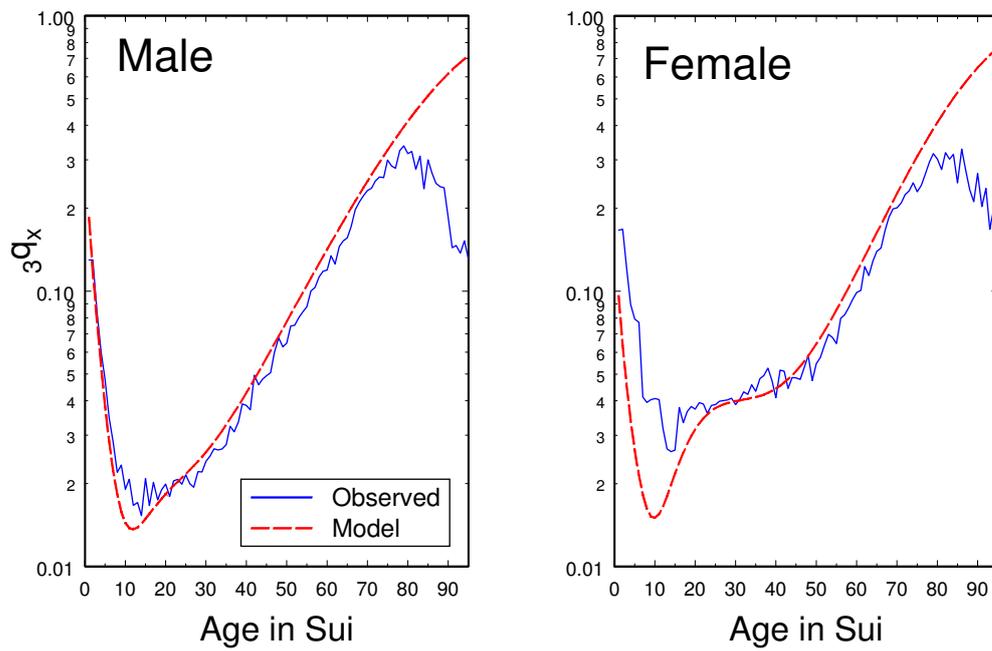


Figure 10. Conditional Mortality Risk ${}_3Q_x$. Left panel: UN South Asia Model, $e(1) = 40$; Right panel: UN General Model, $e(1) = 45$.

4.3 Linking Households between panels.

Identifiers to link households or household groups between successive panels are not provided in the downloadable data. To link panels a unique ‘chain id’ for the household was created when an ego id appeared for the first time. Household chains were constructed when ego, the person with relationship ‘e’, could be linked across consecutive panels. No attempt was made to span gaps in the data as the decomposition method used below makes no distinction between temporary and permanent right censoring.

If ‘ego’ continued and any member of the household who was not ego became an ego in the next panel, the latter were allocated new chain id’s but the original chain continued. If ego died, new chains were started for any members who were egos in the subsequent panel. The mean length of household chain was 5.26 panels with a median of 4. If a new household had been assumed for each new ego, these figures would fall to 3.80 and 3 respectively. These chains are left-truncated and right-censored by missingness patterns that make it impossible to estimate the true extinction times or half-lives of households.

4.4 Missingness Patterns.

Table 1 shows the nine most frequent pattern mixtures of observation using the village as the entity of interest. When a village is observed in a panel it is marked with 1 and · represents a missing observation. The row summaries have been repeated to help readability over a wide table. The most frequently observed pattern over the period 1789 to 1885 is experienced by 15 villages and contributes 300 village-panels. Unfortunately, this pattern only represents 2.6% of the villages or 3.8% of the village-panels, and the panels do not form a continuous temporal sequence. In fact, the missingness structure at the village level is extremely varied and the top 9 patterns only represent about 19% of all village-panels. Table 2 is a summary of the pattern mixtures at the household-chain level. The most frequent chain comprises the 1789, 1792 and 1795 panels, but it only represents 3% of the household-chains or 1.7% of the household-panel observations. As at the village level, the observational patterns are extremely diverse. It would be unrealistic to adopt one of the strategies from longitudinal data analysis: to consider each pattern-mixture as a stratum and fit a separate model.

The extreme missingness of the data also operates at lower levels of the hierarchy. As an example, Figure 11 plots the frequency of households for each panel within our time limits experienced by village 186. The village seems to switch between two recording patterns. A visual inspection of a village by panel contingency table suggests that this is extreme, but fluctuating counts are experienced to a lesser degree by many villages.

Incomplete data are also observed at the level of the household, where the average number of individuals at risk to die as a proportion of all household members is 0.68. At all levels in the data hierarchy, the principal problem that missingness presents is that we cannot distinguish attrition and recruitment from the processes that make the

Table 1. Leading Pattern Mixtures of Observations at the Village Level.

Villages	1789	1792	1795	1798	1801	1804	1807	1810	1813	1816	1819	1822	1825	1828	1831	1834	1837	Obs.	Perc.
1	15	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	300	3.8
2	11	1.9	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	253	3.2
3	12	2.1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	216	2.7
4	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	156	2
5	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	130	1.6
6	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	120	1.5
7	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	120	1.5
8	5	0.9	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	120	1.5
9	7	1.2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	112	1.4
Other	476	83.8	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	6447	80.9
Total	568	100	231	276	210	217	250	71	183	214	269	286	317	358	350	329	273	7974	100

Villages	1840	1843	1846	1849	1852	1855	1858	1861	1864	1867	1870	1873	1876	1879	1882	1885	Obs.	Perc.
1	15	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	300	3.8
2	11	1.9	1	1	1	1	1	1	1	1	1	1	1	1	1	1	253	3.2
3	12	2.1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	216	2.7
4	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	156	2
5	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	130	1.6
6	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	120	1.5
7	10	1.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	120	1.5
8	5	0.9	1	1	1	1	1	1	1	1	1	1	1	1	1	1	120	1.5
9	7	1.2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	112	1.4
Other	476	83.8	X	X	X	X	X	X	X	X	X	X	X	X	X	X	6447	80.9
Total	568	100	260	278	276	275	253	203	249	243	281	273	163	266	316	152	7974	100

Table 2. Leading Pattern Mixtures of Observations at the Household Chain Level.

	Households	Perc.	1789	1792	1795	1798	1801	1804	1807	1810	1813	1816	1819	1822	1825	1828	1831	1834	1837	Obs.	Perc.
1	624	3.0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1872	1.7
2	66	0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1518	1.4
3	155	0.7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1395	1.3
4	71	0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1136	1.0
5	161	0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1127	1.0
6	111	0.5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	999	0.9
7	52	0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	988	0.9
8	160	0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	960	0.9
9	81	0.4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	891	0.8
Other	19237	92.9	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	97904	90
Total	20718	100	2292	3009	2864	2391	2713	1103	676	834	2083	2833	4379	4897	4942	4694	4329	4370	3762	108790	100

	Households	Perc.	1840	1843	1846	1849	1852	1855	1858	1861	1864	1867	1870	1873	1876	1879	1882	1885	Obs.	Perc.	
1	624	3.0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1872	1.7
2	66	0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1518	1.4
3	155	0.7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1395	1.3
4	71	0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1136	1.0
5	161	0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1127	1.0
6	111	0.5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	999	0.9
7	52	0.3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	988	0.9
8	160	0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	960	0.9
9	81	0.4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	891	0.8
Other	19237	92.9	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	97904	90
Total	20718	100	3568	3640	3768	3806	2982	3076	4065	3737	3303	3217	3681	3949	2447	3826	5085	2469	108790	100	

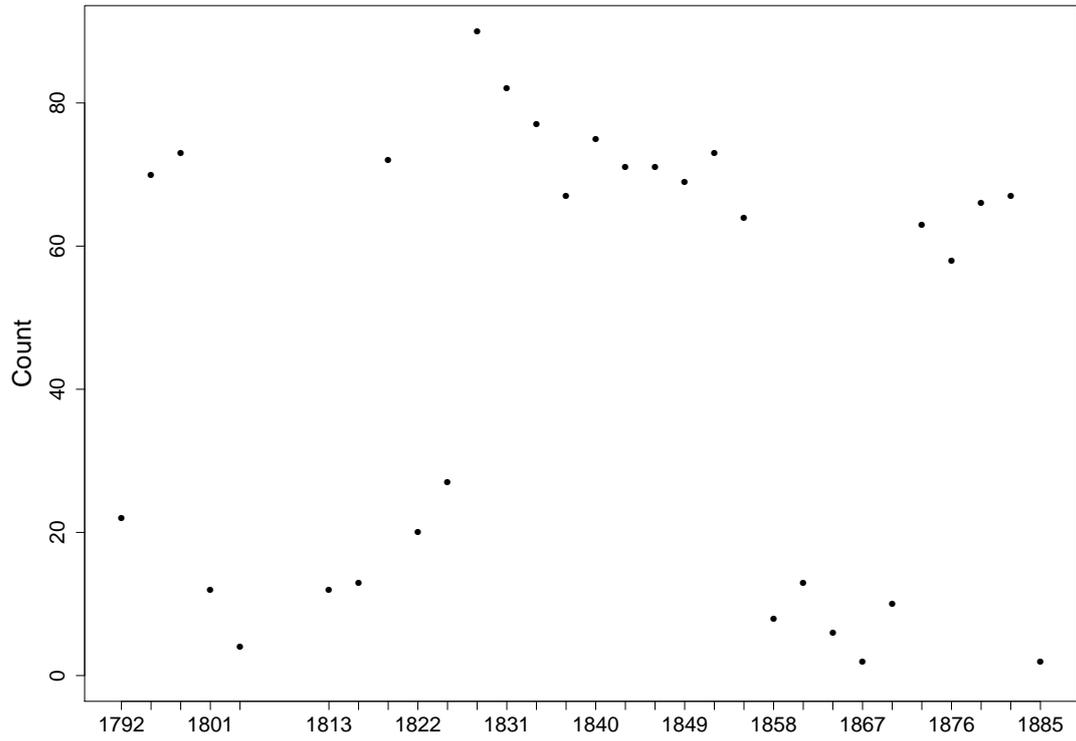


Figure 11. Count of households by panel for village 186.

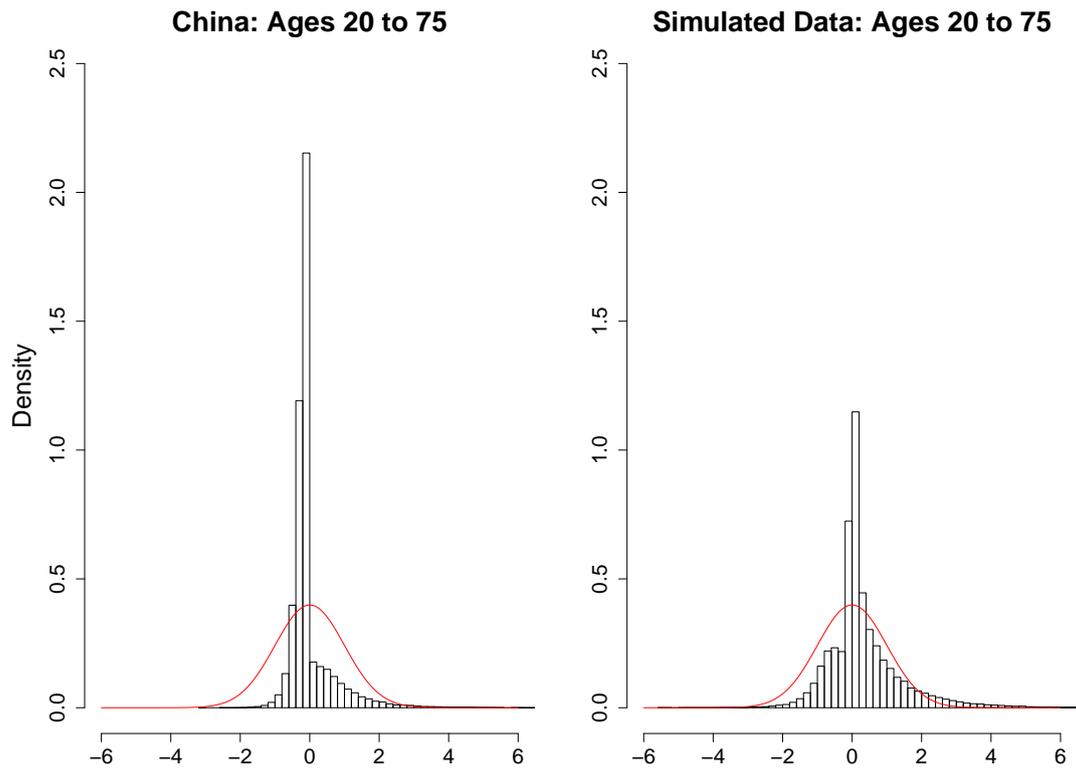


Figure 12. Frailty Distribution for Chinese Households and Simulated Data.

written records unavailable.

5 Results.

Figure 12 shows the frailty distribution for the Liaoning data, contrasted with the frailty of simulated kin-sets experiencing homogeneous mortality from the model life-tables shown in Figure 10. The relevance of the simulated frailty should be treated with caution for a number of reasons. Although the marriage-age and celibacy parameters were taken from Lee and Campbell (1997), the fertility model was based on English family reconstitutions. Secondly, the simulated frailty is based on kin-sets, not co-resident kin, and these are much larger on average. Ignoring these caveats, the surprise is that the empirical distribution seems to be under-dispersed relative to the homogeneous one. Conventional views about mortality disparity based on socio-economic and ecological differences, and on infectious processes, lead us to expect over-dispersion. One explanation might be that the death of an individual in a household increases the survival

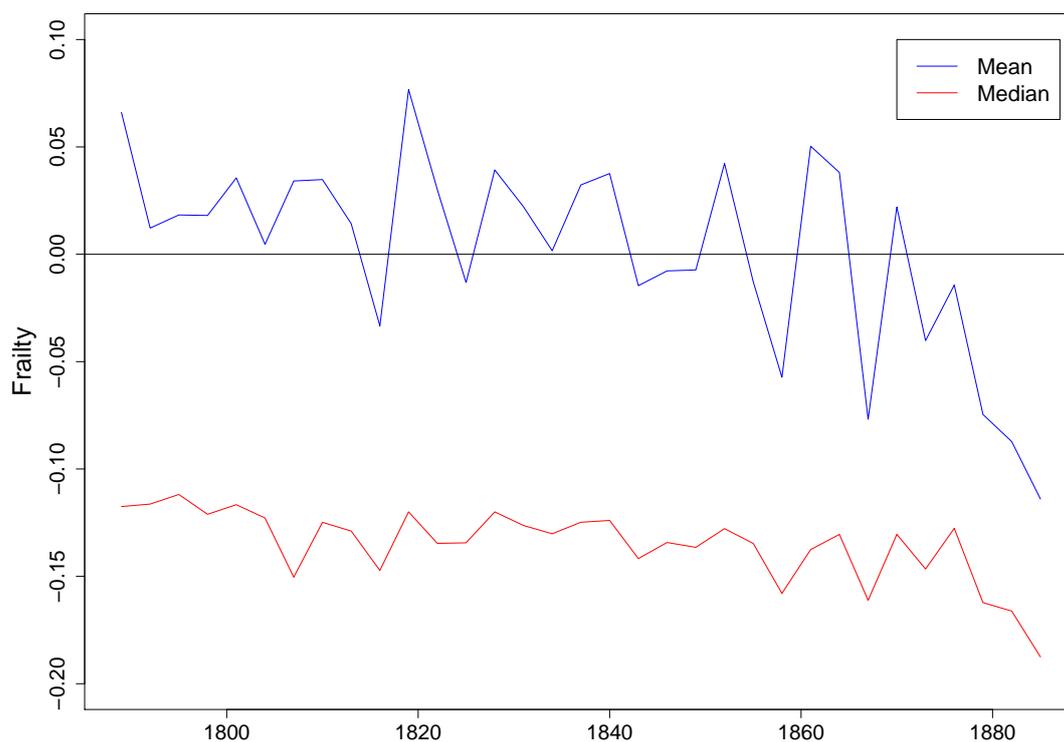


Figure 13. Time Series of Mean and Median Frailty for Chinese Households.

chances of the other members. A more likely reason is that selection is acting more strongly on the tails of the distribution, although it is not obvious what this process might be for the robust households.

Two indicators of the location of the frailty distribution over time are shown in Figure 13. The parameters are relatively stable, although there is a suggestion that frailty declines in the later panels. Before 1860 the asymmetry indicated by the difference between the median and the mean is also stable. Dispersion, as measured by the coefficient of variation in each panel, follows a downward trend but there is a massive increase in diversity at the 1834 panel.

5.1 Village level

The degree of missingness identified above makes it important to understand the nature of its possible effects. As an illustration of what might be done at different levels in the data hierarchy, we decompose frailty at the village level under the assumption

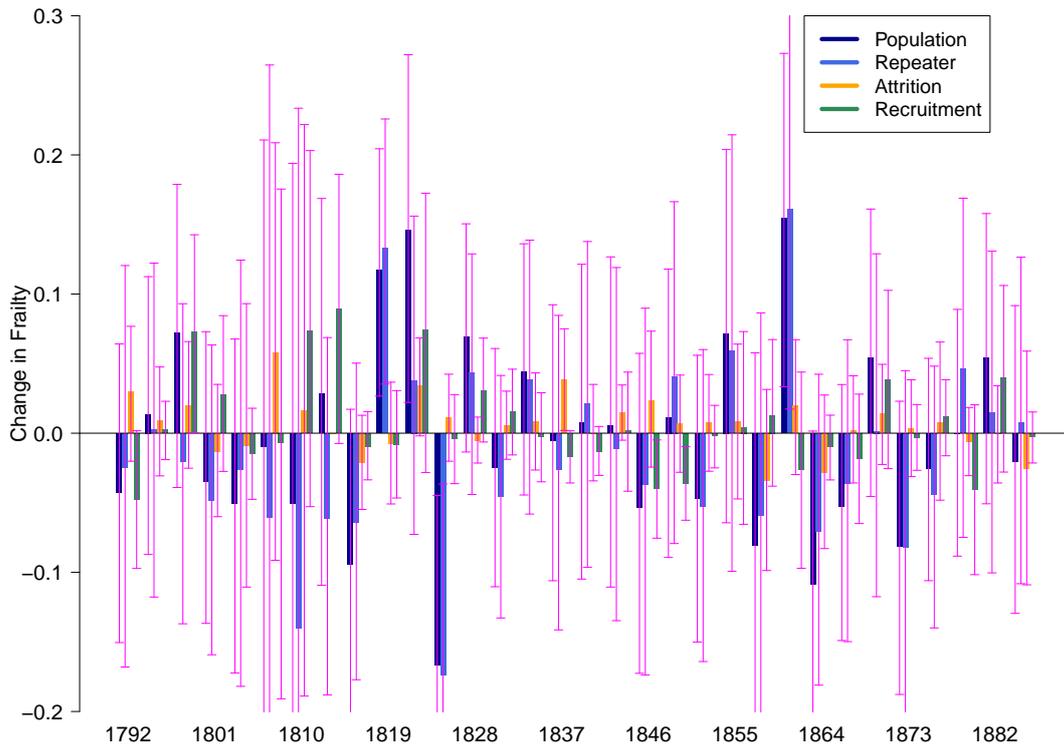


Figure 14. Decomposition of frailty change for villages.

that missingness for villages indicates data unavailability rather than true recruitment or attrition. In this case D and \hat{E} in equation 5 are observed and expected death counts for the village. Figure 14 shows the results. The dark blue bars are the change in the population average. This change is decomposed into three components that sum to the population value. Change in average frailty for villages measured in two consecutive panel counts is shown as the light-blue bars. The selection effect caused by the attrition of villages that can only be measured on the opening panel are coloured gold. The selection effect caused by the recruitment to the closing panel is shown in green. Bootstrapped confidence intervals are shown in pink, based on 999 iterations. Very few of the changes can be considered as significant, although this is not proof that the missingness can be classified as ‘missing at random’.

The synthetic accumulation of the effects is shown in Figure 15, the starting point being the average frailty across the data. Villages with repeated measurements before 1820 became less frail, but this was generally offset by a tendency for the villages that stop and start recording to be more frail, so that the net effect is close to zero over time. This is encouraging, but further examination is required.

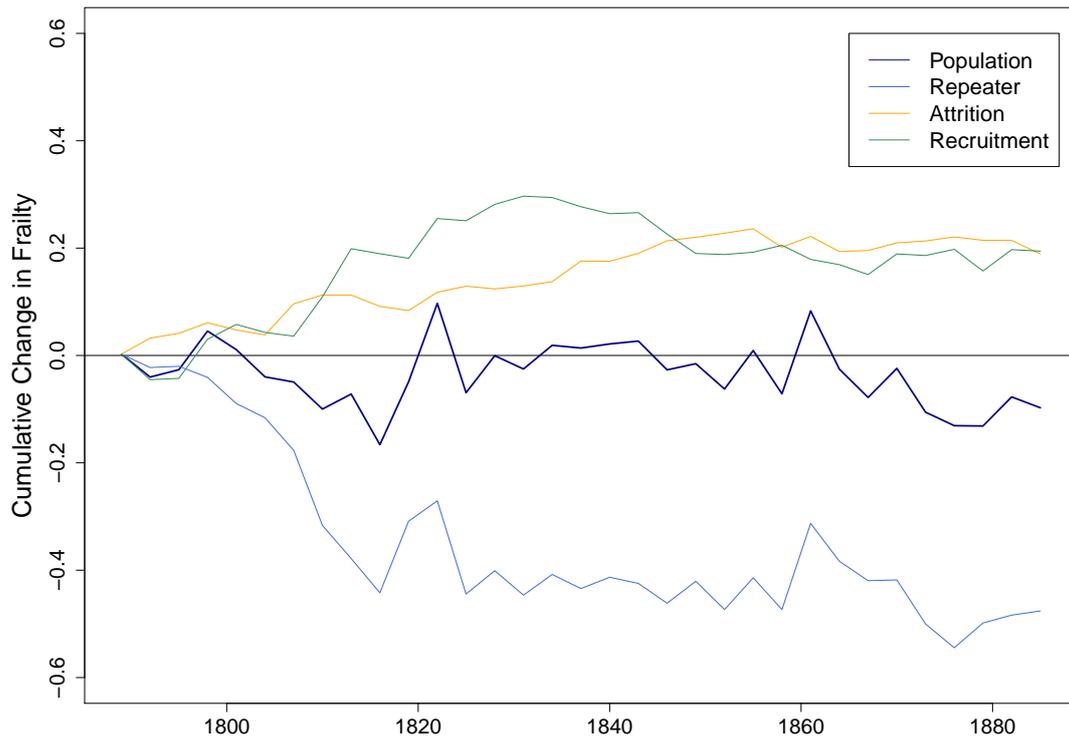


Figure 15. Cumulative change in frailty for villages.

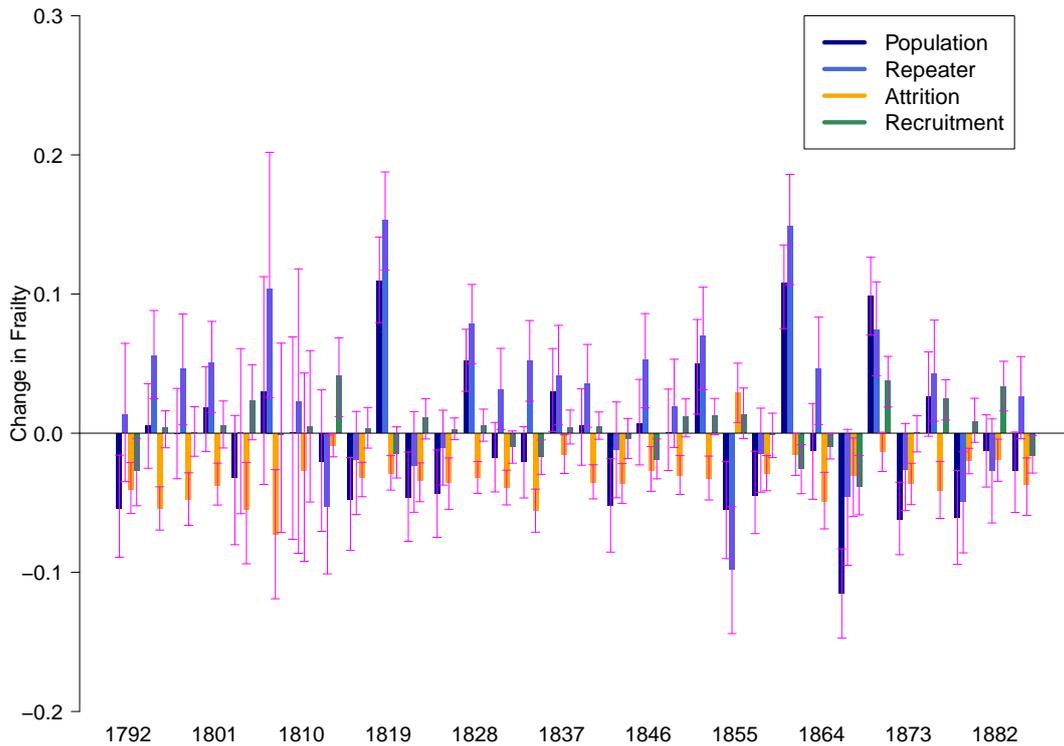


Figure 16. Decomposition of frailty change for households.

5.2 Household level

Figure 16 shows the results of decomposing changes in average frailty between panels. The colour scheme is the same as in Figure 14 but the entities experiencing change in \hat{F} are households rather than villages, and the bootstrap intervals are only based on 99 iterations.

There seems to be no clear trend in the population-level changes in frailty, but it is striking that the attrition effect (gold) is almost uniformly negative. Attrition is selective for higher frailty households with the effect that the remaining population becomes more robust after their removal. It might be thought that this is a natural result of a household chain ending observation because of the death of ego, but this result remains if egos are omitted from the frailty scoring for the households. The change in the average frailty of repeater households is generally positive. We expect epidemics to result in ‘frail’ panels, but it is interesting that there are two panel counts, 1855 – 1858 and 1867 – 1870, when frailty is low.

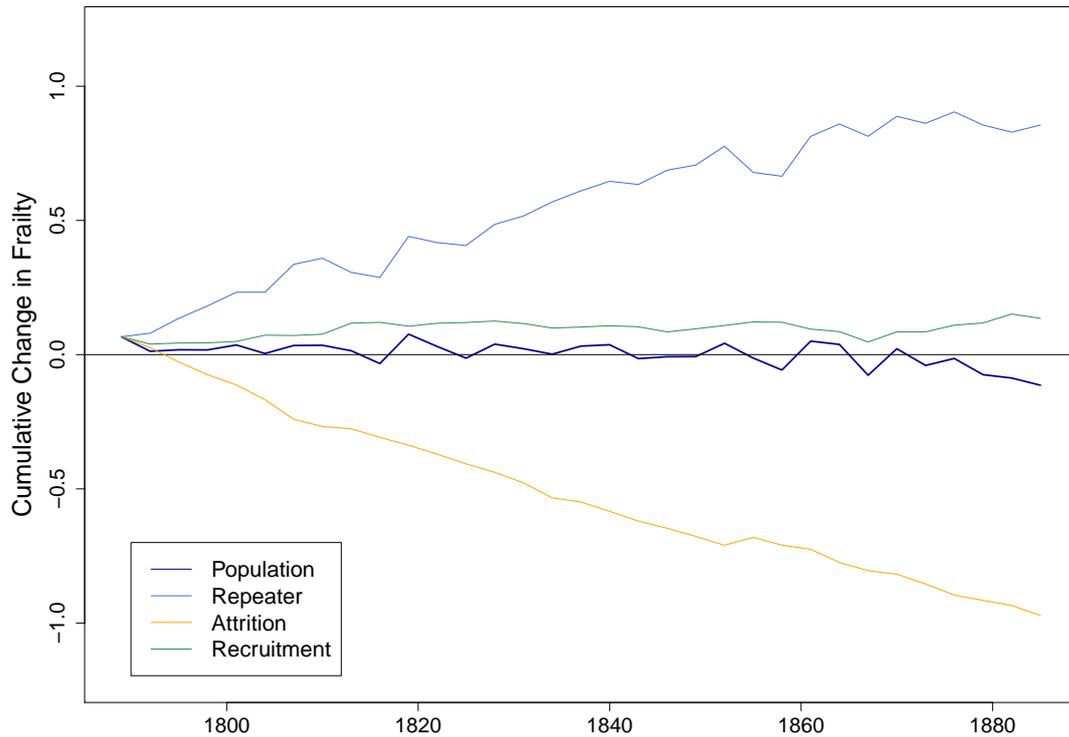


Figure 17. Cumulative change in frailty for households.

To clarify and summarise these results, each component is cumulated in Figure 17, using the average frailty across all panel counts as the starting position. This synthetic calculation makes it clear that the neutral evolution of frailty at the population level, which would suggest that a single life-table is representative of the whole period, conceals offsetting dynamics for repeaters and drop-outs.

A schematic interpretation of the frailty dynamics from Figure 17 is shown in Figure 18. Recruitment is non-selective, ‘repeater’ households experience increasing frailty, and attrition is selective for higher frailty. The time series of mean and median in Figure 13 show that the net-balance of these processes is close to zero over most of the observational period.

5.3 Household Size

At this stage, we do not know why repeater households increase their frailty, but one possibility is that the dynamics of frailty are positively associated with the dynamics of

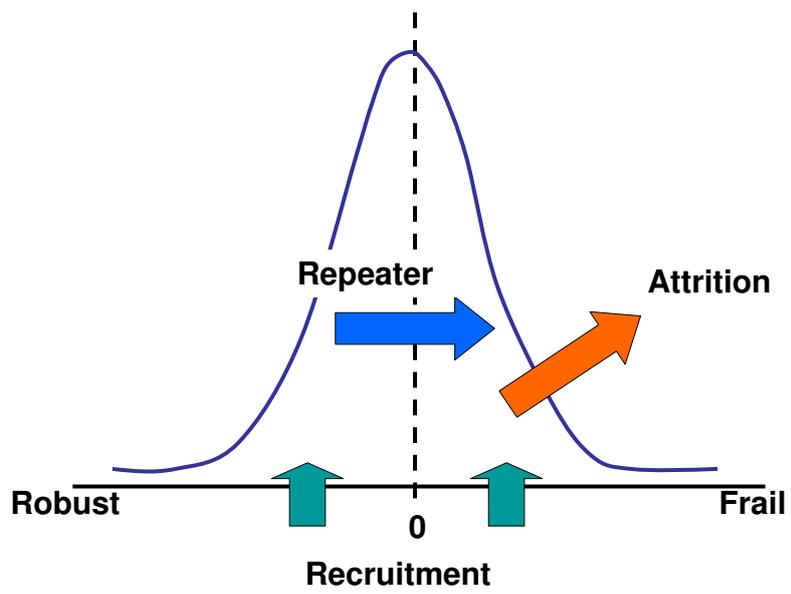


Figure 18. Dynamics of the Frailty Distribution.

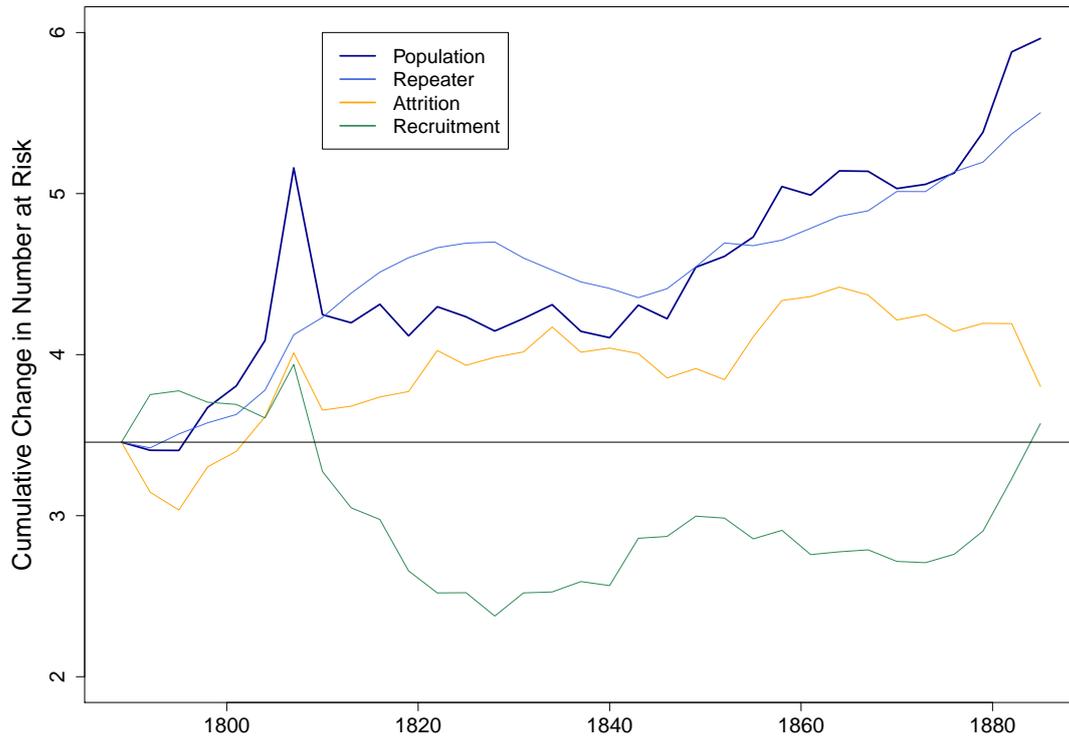


Figure 19. Cumulative change in number at risk in each household.

household size. Figure 19 shows the cumulants of the factors affecting the change in the average number at risk in each household. Both new and dropout households tend to be smaller than average, but continuing households grow in size and the population tends towards larger households. Figure 20 interprets the dynamics of changes in the number at risk distribution across households. Both recruitment and attrition are selective for smaller households. Repeater households increase in size and this, combined with selective attrition, is enough to lead to larger average numbers at risk in each household. It is possible that this is evidence of density-dependence for household mortality. One of the important virtues of considering the dynamics is that a cross-sectional plot of frailty against the number at risk does not show this positive association – if anything it is negative.

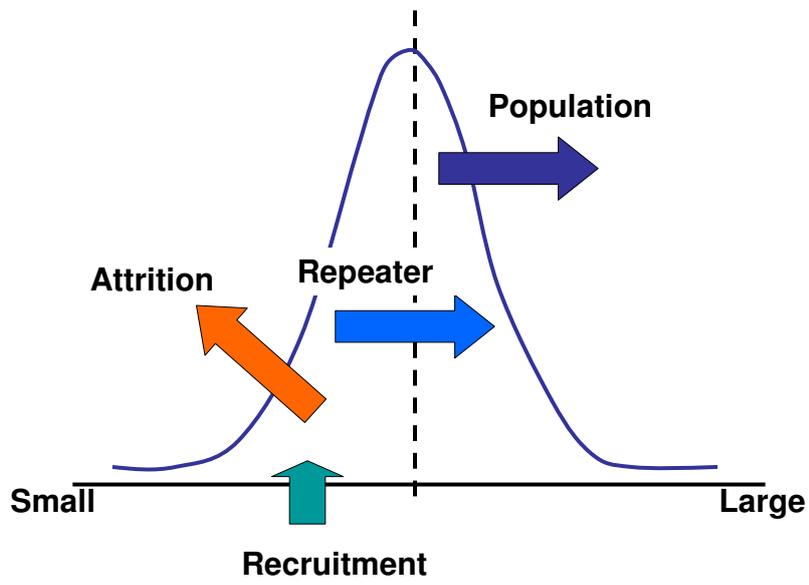


Figure 20. Dynamics of the Distribution of the Number at Risk.

6 Discussion.

One of the principal issues with the decomposition is the problem that missing data are indistinguishable from true recruitment and attrition of households, and this can apply at other levels of the data hierarchy. While one would like a method that separates them, it is likely that this is more than a technical issue and may be conceptually impossible.

The Price Theorem is a first-order Markovian analysis of change. It is conceivable that there may be more complex, but characteristic, cycles of household frailty. Seebohm Rowntree's ideas of a household poverty cycle for the nuclear family could have implications for mortality.

Although we have borrowed some ideas from evolutionary biology, the relatedness between individuals is known for many of the individuals in this dataset. This means that one could go further and consider measures of heritability and possibly quantitative genetics. However, the restricted age-range of the 'reliable' mortality data and the lack of fertility recording places limits on the study of the biological 'fitness' of this population. A wider approach to the biology of dynamic heterogeneity can be found in the recent papers of Tuljapurkar and co-authors (Tuljapurkar et al., 2009; Tuljapurkar and Steiner, 2010).

This study of dynamic frailty was partly motivated by a wish to develop tools that would inform strategic objectives. A social planner trying to improve longevity might assume that newly established households need the most help, but Figure 18 suggests that their frailty is broadly distributed. It seems that established households with long durations might have been a more susceptible target group for intervention.

These results have implications for modelling survival, such as fitting proportional hazard models. Figure 18 suggests that left-truncation (recruitment) is not informative about frailty, but that it is dangerous to make the assumption that right-censoring (attrition) is uninformative. In addition, household 'duration' seems to be positively associated with frailty, and we should take prior 'frailty scores' into account, although there will be a problem defining these for new households.

We have been largely concerned with household averages and their dynamics, but the scoring of frailty is conducted at the individual level and is not conditioned on a death occurring. This means that it would be possible to analyse associations of frailty between kin categories, such as sibling sets, or parent-child groups, and measure how these linkages change over the life-course.

7 Conclusion.

This paper has offered an exploration of two methods. The results suggest that it is possible to score frailty at the household level and other levels of the data hierarchy. The

decomposition method used to analyse dynamic frailty is well-known in evolutionary biology, but is applied to data with extremely complicated ‘missingness’ patterns. If the results are justified, they show that the apparent stability of mortality at the population level is maintained by a balance of forces at the household level. Persistent households appear to suffer increasing frailty over time, but this effect is almost exactly offset by the effects of household attrition, with recruitment having a neutral role. No satisfying explanation is offered, but we note that persistent households become progressively larger, which could suggest density-dependent dynamics for household mortality.

Acknowledgements.

I should like to thank my Maren Rebke and Jutta Gampe for their help and advice. Cameron Campbell provided a rapid response service to queries about the data, but we accept full responsibility for any errors we may have made. Fernando Colchero identified the link between this research and the papers of Tuljapurkar and his co-authors.

References

- Becker, G. S. (1981). *A treatise on the family*. Cambridge, Mass.: Harvard University Press.
- Cameron, A. C. and P. K. Trivedi (1998). *Regression analysis of count data*. Cambridge, UK; New York, NY, USA: Cambridge University Press.
- Cameron, A. C. and F. A. G. Windmeijer (1996). R-squared measures for count data regression models with applications to health-care utilization. *Journal of Business & Economic Statistics : a publication of the American Statistical Association*. 14(2), 209–.
- Canty, A. and B. D. Ripley (2009). *boot: Bootstrap R (S-Plus) Functions*. R package version 1.2-41.
- Coale, A. J. and P. G. Demeny (1966). *Regional model life tables and stable populations*. Princeton, N.J.: Princeton University Press.
- Coulson, T. and S. Tuljapurkar (2008). The dynamics of a quantitative trait in an age-structured population living in a variable environment. *The American Naturalist* 172(5), 599–612.
- Das Gupta, M. (1990). Death clustering, mothers’ education and the determinants of child mortality in rural Punjab, India. *Population Studies* 44(3), 489–505.
- Davison, A. C. and D. V. Hinkley (1997). *Bootstrap Methods and Their Applications*. Cambridge: Cambridge University Press. ISBN 0-521-57391-2.

- Derosas, R. and M. Oris (2002). When dad died : individuals and families coping with family stress in past societies. Bern; New York, pp. -. P. Lang.
- Gardner, A. (2008). The Price equation. *Current Biology* 18(5), 198–202.
- Kerr, B. and P. Godfrey-Smith (2009). Generalization of the Price equation for evolutionary change. *Evolution* 63(2), 531–536.
- Lee, J. Z. and C. D. Campbell (1997). *Fate and fortune in rural China : social organization and population behavior in Liaoning, 1774-1873*. New York: Cambridge University Press.
- Lee, J. Z. and C. D. Campbell (2010). China multi-generational panel dataset, Liaoning (cmgpd-ln), 1749-1909 [computer file].
- Lee, J. Z., C. D. Campbell, and L. Anthony (1995). A century of mortality in rural Liaoning, 1774–1873. In S. Harrell (Ed.), *Chinese historical microdemography*, pp. 163–182. University of California Press.
- Mittlböck, M. and H. Heinzl (2001). A note on R^2 measures for poisson and logistic regression models when both models are applicable. *Journal of Clinical Epidemiology*. 54(1), 99–103.
- Price, G. (1970). Selection and covariance. *Nature* 227(5257), 520–1.
- R Development Core Team (2009). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.
- Rebke, M. (2010). From the Price equation to a decomposition of population change. *Journal of Ornithology*, (available online).
- Rebke, M., T. Coulson, P. Becker, and J. Vaupel (2010). Reproductive improvement and senescence in a long-lived bird. *Proceedings - National Academy of Sciences USA* 107(17), 7841–7846.
- Rose, G. (2001). Sick individuals and sick populations. *International Journal of Epidemiology* 30(3), 427–.
- Sabo, R. and N. Chaganty (2010). Hypothesis testing for various familial dependence structures. *Communications in Statistics: Simulation and Computation* 39(1), 207–219.
- Smith, J. and J. Oeppen (1993). Estimating numbers of kin in historical England using demographic microsimulation. In D. S. Reher and R. Schofield (Eds.), *Old and new methods in historical demography*, Oxford [England]; New York, pp. 280–317. Clarendon Press; Oxford University Press.

- Tuljapurkar, S. and U. K. Steiner (2010). Dynamic heterogeneity and life histories. *Annals - New York Academy of Sciences* 1204, 65–72.
- Tuljapurkar, S., U. K. Steiner, and S. H. Orzack (2009). Review and synthesis: Dynamic heterogeneity in life histories. *Ecology Letters* 12(1), 93–106.
- United Nations Population Division (2003). MORTPAK for Windows, version 4.0.
- Vandezande, M., S. Moreels, and K. Matthijs (2010). Explaining death clustering. intergenerational patterns in infant mortality, Antwerp, 1846-1905. Working paper of the scientific research community historical demography vol:wog/hd/2010-13, Centre for Sociological Research, Katholieke Universiteit Leuven.
- Vaupel, J. W., K. G. Manton, and E. Stallard (1979). The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography* 16(3), 439–454.
- Vaupel, J. W. and A. I. Yashin (1985). Heterogeneity's ruses: some surprising effects of selection on population dynamics. *American Statistician* 39(3), 176–185.
- Wrigley, E. A., R. S. Davies, J. E. Oeppen, and R. S. Schofield (1997). *English population history from family reconstitution, 1580-1837*. Cambridge studies in population, economy, and society in past time, 32. Cambridge; New York: Cambridge University Press.
- Zaba, B. and P. H. David (1996). Fertility and the distribution of child mortality risk among women: An illustrative analysis. *Population Studies* 50(2), 263–278.
- Zhao, Z. (2003). On the Far Eastern pattern of mortality. *Population Studies* 57(2), 131–147.