Current Biology, Volume *23*

Supplemental Information

The Demographic Transition Influences Variance in Fitness and Selection on Height and BMI in Rural Gambia

Alexandre Courtiol, Ian J. Rickard, Virpi Lummaa, Andrew M. Prentice, Anthony J.C. Fulford, and Stephen C. Stearns

Supplemental Experimental Procedures

Demographic data: From 1950 until present, residents of three villages (Keneba, Kantong Kunda, Manduar) in the West Kiang district of The Gambia have been continuously treated and studied by the UK Medical Research Council. This arrangement was initiated by Sir Ian McGregor, when the population of the largest village, Keneba, numbered approximately 700 inhabitants. Local residents, who are mostly of Mandinka ethnicity, have traditionally subsisted largely on crops of rice, sorghum and millet, with a single cash crop of groundnuts. Over the study period, there has been incremental improvement in the provision of the healthcare to the local population, most notably in the form of increased antenatal and natal care since the 1970s. Overall under-five mortality rate declined from more than 40% prior to 1970 to less than 10% in the present day [3]. Fertility has also historically been high, with women giving birth to a total of around seven children over their lifetime on average [19]. The society is highly polygynous, and women spend virtually all of their adult lives married, beginning reproduction at about 18 years of age.

Our analysis was conducted at the 'person-year' level, which means that each line of data corresponded not to an individual person, but to a year in which a unique individual was present in the population, a 'person-year'. To derive such a data set, one must have a reliable longitudinal information resource that monitors individuals repeatedly over the study period. This has been made possible in West Kiang through a combination of approximately-annual health surveys (1950 to 1980), clinic attendance (1974 to present) and annual demographic surveys (2004 to present). Resident individuals were given unique ID numbers that identified them as residing in one of the villages. All research projects conducted at MRC Keneba (including secondary data analysis) are approved by the Gambian Government/MRC Ethics Committee. We restricted our analysis to individuals identified as being from Keneba and Manduar, and considered only years from 1956 onwards, because data collected before this point in these two villages may not be reliable, and the reliable data series for Kanton Kunda, the smallest village, would be much shorter [18]. Data from both villages were pooled since their inhabitants can be considered as being part of a single human population. We used only female individuals, for between the end of the annual health surveys (1980) and the start of the demographic surveys (2004), the focus of data collection on adults was on women and children, via clinics. From the entire complement of data, we derived a 'last contact year'. Individuals were scored as being 'present' in every year beginning with their birth year until their year of death or last contact year. We removed a small number of individuals (N=5) with inconsistency between the years of birth, death and last contact; they likely had incorrect recorded ages of maternity. Another 791 individuals were not included in the present analysis because of the lack of information about their anthropometrics. In sum, we assembled data on 51,909 person-years representing 2818 individuals.

Fitness measures: We estimated fitness at the person-year level and defined total fitness W_{total} as the genetic contribution that each individual present in year n made to the population in year n + 1. Thus, in a given year individuals scored 1 for surviving, or zero for dying, plus 1 for each offspring produced. As individuals of different ages are pooled, correlations across ages in survival/fertility are implicitly considered on a cross-sectional basis. From our total fitness estimate, we distinguished two fitness components: i) pre-adult fitness W_{child} that represents the component of W_{total} that is driven by mortality before 15y and that we computed as a binary (0 for death before 15y, 1 otherwise); and ii) adult fitness $W_{fertility}$ that represents the component of W_{total} that is driven by fertility from 16y onward and that we computed as the number of offspring produced by a given individual in a given year. From total fitness and its components, we derived the corresponding relative measures, w_{total}, w_{child} and $w_{fertility}$, by dividing each value by their associated mean over all person-years of the year in question.

Variance in relative annual fitness: From the relative fitness-related measures, w_{total}, w_{child} and $w_{fertility}$, we derived three corresponding variances: i) variance in relative total fitness, var(w_{total}); ii) variance in relative child survival, var(w_{child}); iii) variance in relative fertility, var($w_{fertility}$). Note that for any relative fitness-related measure w, var(w) could also be computed from the absolute fitness values W as var(W)/(mean(W)^2). Those two formulations are indeed mathematically identical but we present both for clarification since studies on opportunity for selection using aggregated data tend to use the latter formulation.

Height and BMI centiles: During the annual surveys (1950-1980) and clinic visits (1975- present), measurements of weight/height have been recorded. The ages at which they have been taken varies with respect to time, but most individuals have been weighed several times, usually from birth onwards (Figure S2). However, since data collection only began in 1950, limiting the analysis to individuals who had been weighed at a particular age would bias the data series with respect to time (e.g. if birth measurements were used we would only be able to include individuals born from 1950 onwards) or toward individuals who survived until a particular age (e.g. if we used measurements taken in adulthood).

To deal with these problems and to maximize information use, we devised a technique that allowed all individuals to be included in analysis, provided they were measured for the metric in question at least once. For each metric, we first produced a population-level LOWESS curve describing all measures as a function of the age at which an individual was measured, *i.e.* growth charts (Figure S2). Using these curves, we then transformed measures into residuals corresponding to a relative measure of a given trait that is independent from growth-induced variations. We then assigned each residual data point a rank relative to the 100 other data points that were closest in terms of age at measurement. To avoid cohort effects, we restricted this ranking procedure so that the date of birth of the individuals to whom the other data points corresponded were a maximum of five years before or after that of the date of birth of the focal individual. Therefore, each measure was associated with a rank from 1 to 100, *i.e.* a percentile, representing the height/BMI of an individual that could be considered as independent from age and cohort effects. As a result, each individual is represented by a pool of ranks, one for each measure on record. From this pool, the median value was chosen and used in analysis. This allowed us to obtain one metric for height and one for BMI for all individuals, irrespective of how many times they have been measured. In addition, computing the median reduces the potential influence of misreported measurements.

Considering the median value of several percentiles that have been computed at different ages for each individual implies the assumption that the growth trajectories of individuals are parallel. This assumption seems reasonable given our data, for the ratio of the variation in percentile scores within/between individuals is much smaller than expected in the absence of parallel growth curve (using individuals with at least three measurements and with at least 20 years between the first and last measurements: N=812; observed variance ratio for height = 0.25, observed variance ratio for BMI = 0.72; variance ratio when measurements for both traits are permuted among the same individuals = 17). We also verified that median ranks did not systematically vary according to cohort or age, and thus that values from individuals born in different time periods or measured at different ages were comparable with one another. Height and BMI were only very weakly correlated with one another ($\rho = 0.09$ from our calculated scores), and therefore can be considered to largely capture independent aspects of individual anthropometry. In contrast, height and weight were strongly correlated ($\rho = 0.63$), and would therefore not have been appropriate for separate analysis.

Standardized selection gradients: We investigated phenotypic selection on height and BMI for our three fitness-related measures using the traditional framework introduced by Lande and Arnold [27]. For each fitness-related measure, we estimated the standardized linear selection gradients β as derived from the following linear regression models:

$$w_{i,y} = 1 + \beta_{h_y}\text{score\_height}_i + \beta_{\text{BMI}_y}\text{score\_BMI}_i + \varepsilon_{i,y}$$

In this regression model, $w_{i,y}$ is the relative fitness-related measure considered for a given person i in a given year y, and score_height_i and score_BMI_i are the centile scores for height and BMI of a given individual i, converted into z-scores through standardization within each year. The intercept is constrained to 1 because fitness-related measures are mean-centered within each year. The term $\varepsilon_{i,y}$ corresponds to the residual relative fitness, the component of fitness which is not captured by our covariates. The framework of Lande and Arnold [27] has been specifically developed to model the selection gradient on each trait despite the correlation between traits (as long as this correlation is not too high). Consequently, our estimates of selection on BMI are not confounded by the selection acting on height but capture the effect of weight on fitness, controlling for height.

Trend analyses of selection gradients: Each selection gradient estimate is associated with a certain level of uncertainty (measurement and sampling error) and neglecting this uncertainty has been shown to compromise analysis of trends in selection over time [28]. In particular, because the level of uncertainty varies between years (Figure 1), biases could be particularly large in an analysis that ignored this uncertainty. Therefore, we decided to test the relationship between selection gradients and time using gradients estimated within each of 500 datasets simulated by bootstrapping individuals from the original data, rather than using direct estimates of selection gradients. From each simulated dataset, selection gradients were correlated with time for both height and BMI using the Spearman non-parametric correlation test, after rescaling both fitness-related measures and anthropometric scores within each year [27]. P-values were computed by comparing the average of the 500 correlation coefficients to the distribution of this statistic under the null hypothesis. This latter was approximated by averaging the 500 corresponding correlation coefficients produced by bootstraps for each of 500 datasets for which years were randomly shuffled (so 250,000 correlations were performed for each gradient/fitness-measure combination). As observed estimates and the null hypothesis both considered the non-independence between data points from different years involving the same individuals, the test is not biased by this non-independence. To analyze the robustness of our test, we checked that the distribution of p-values under the null hypothesis was uniform.

Furthermore, the uncertainty in selection gradient estimates also potentially depends on the uncertainty in variance in relative fitness estimates. Therefore, an artifactual relationship can emerge when trying to predict one from the other, which is nonetheless required to disentangle the effect of changes in variances in relative fitness on selection, from other changes happening through time. Consequently, we also estimated the relationship between selection gradients and variance in relative total fitness and its components within each of 500 datasets simulated by bootstrapping individuals from the original data, and we present r-squared values averaged over all simulations (Table 2).

Figure S1. Related to the Results and Discussion

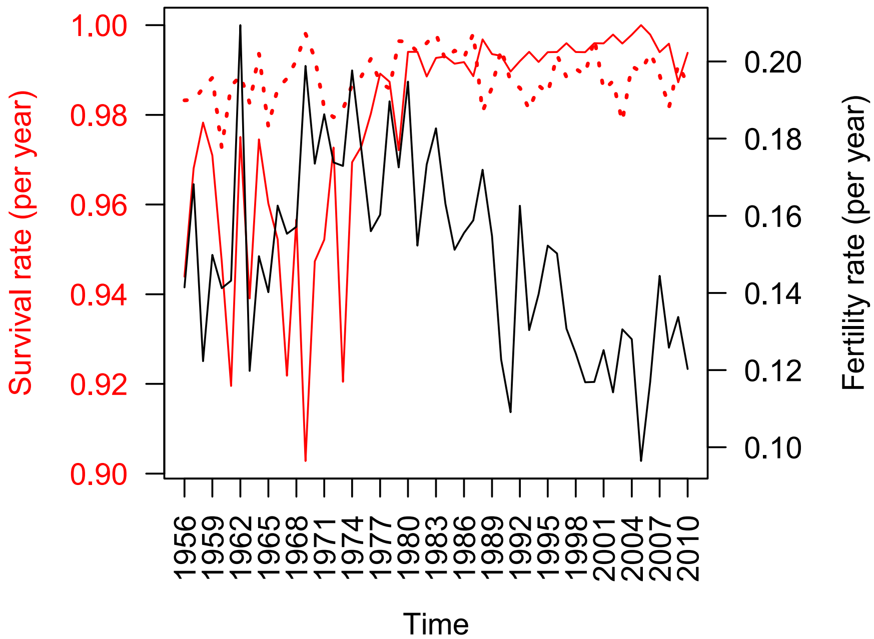

Temporal change in vital rates across a demographic transition in rural Gambia. Solid red line represents the mean annual survival rate for children up to the age of 15 years. The dotted red line represents the mean annual survival rate for adults (over 15 years). For both, survival rates are indicated by the y-axis on the left of the plot. Mean annual fertility rate is represented by the black line; corresponding values are indicated by the y-axis on the right of the plot.

Figure S2. Related to the Supplemental Experimental Procedures

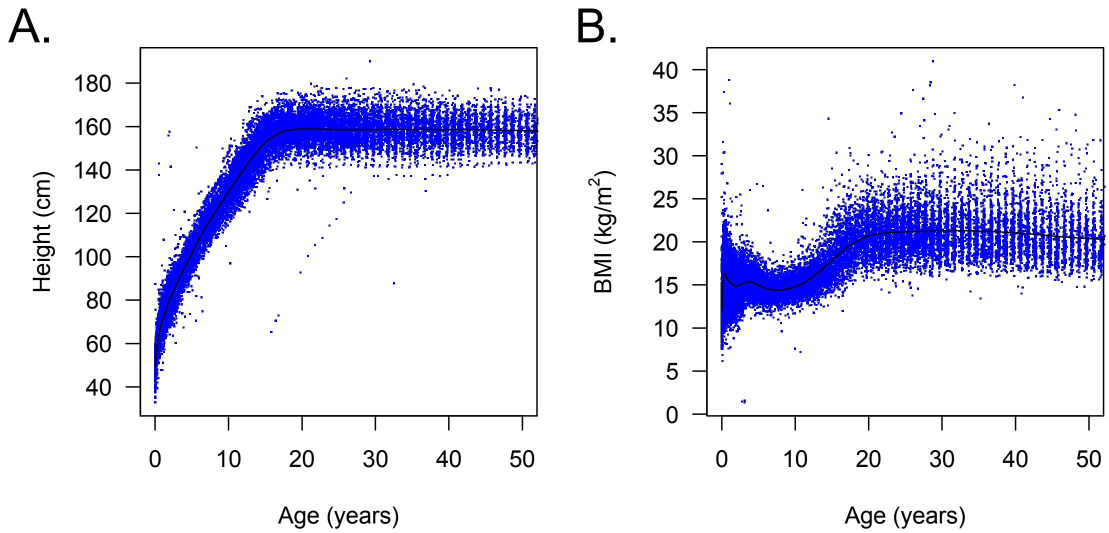

Growth charts for (A) height and (B) BMI in the rural Gambian sample. Anthropometric measures are displayed as a function of the age at which an individual was measured and are symbolised by small dots. Curves correspond to population-level Lowess smoothing used to compute the centile scores used in selection analyses. Fits were produced using the function *loess()* in R, with a span parameter set at 0.15 for both traits.