Aging is favored by natural selection in a changing environment

T. Fink* and M. Kotter†
*London Institute for Mathematical Sciences, Royal Institution, 22 Albermarle St, London W1S 4BS, UK
†Bit.Bio, Dorothy Hodgkin Building, Babraham Research Campus, Cambridge CB22 3FH, UK

Aging is thought to be a consequence of an accumulation of errors in the storage of genetic information. But mounting experimental evidence suggests that aging can be slowed or even stopped. To help resolve this mystery, we present a mathematical framework for understanding the evolutionary benefits of aging. We derive a mortality equation which governs the transition matrix of an evolving population. We find that its largest eigenvalue increases with the maximum age, but the spectral gap decreases. Remarkably, this is independent of the choice of fitness function. As the maximum age of a population decreases, the population grows slower but converges faster. Thus aging confers no benefit in a constant environment, but in a changing one can be favored by natural selection.

Why do we get old? Darwin's theory tells us that evolution is the result of mutation, inheritance and selection. It doesn't refer to death explicitly, but dying is a consequence of competition for finite resources, or when one life form becomes the resource for another carnivorous one. Some mutants, ill equipped to cope with their environment or changes to it, simply break down.

This accounts for the pervasiveness of death, but it tells us nothing about the universality of aging. The standard explanation is that aging is a consequence of an accumulation of errors in life's information storage system—as inevitable as the increase of entropy. But why should the second law of thermodynamics apply to information transmission, especially via machinery that is far from equilibrium? Error correcting codes are effective in life as well as in technology. Even if there were some inevitable rate of decay, it would be minuscule compared to that which is posited to explain aging over lifetimes.

Aging is something with which we have become so familiar with that we don't recognize its strangeness. So the question arises, rather than being inevitable, could aging simply be an evolutionary bolt-on, favored by natural selection? This question is of supreme importance, because 'if aging is programmed', as Werfel et al. [1] surmise, 'rather than a collection of secondary breakdowns or genetic tradeoffs, then effective health and life extensions through dietary, pharmacological, or genetic interventions are likely to be possible'.

One of the limitations of theoretical work in the study of aging is that it largely relies on narratives and computer simulation, rather than mathematical proof. In a detailed critique of mathematical and computational models of programmed death, Smith identified shortcomings ranging from hard-coded mechanisms of kin selection to the absence of genotypic storage. We set out to provide a more rigorous understanding of the dynamics of an aging population, and to what extent is it favored by natural selection.

In this Letter we do four things. First, we derive a mortality equation which relates the growth of a population with maximum age a to that with maximum age 1. This equation is universal in that it applies to any fitness function. Second, we show how to solve the mortality equation, and relate the spectrum of eigenvalues of the population transition matrix for maximum age a to that of maximum age 1. We provide a geometric characterization of the shift in eigenvalues as a increases. Third, we show that the growth rate of the population, set by the largest eigenvalue, increases with the maximum age, but the convergence rate, set by the spectral gap, decreases. Thus in a constant environment, immortality wins, but in a shifting environment, mortality can have the upper hand. Fourth, we test our predictions on populations evolving according to three different fitness functions—constant, Hamming and overlap—confirming our theory.

Mutation

In our model, we have a population of individuals with binary genotypes of length n. The population vector is given by $\mathbf{p} = (p_1, \ldots, p_{2^n})$, which is the size of the population with genotypes $1, \ldots, 2^n$. Offspring are identical to their parents apart from a single point mutation in the genotype.

The process of mutation in the population can be captured by the mutation matrix $\mathbf{M}$. The rows of $\mathbf{M}$ correspond to genotype $00\ldots0$ at the top to genotype $11\ldots1$ at the bottom, in lexicographical order. The same applies to the columns from left to right. The 1s in each row indicate the different parents that can give birth to that row's genotype, keeping in mind that an offspring differs from its parent by a single bit. For example, for $n = 4$, the genotype 1100 can arise from mutations in the parents 1101, 1110, 1000 and 0100.

The first eight mutation matrices $\mathbf{M}$ are shown in Figure 1. The mutation matrix $\mathbf{M}$ can be defined recursively in block form:

$$\mathbf{M}_{n+1} = \begin{pmatrix} \mathbf{M}_n & \mathbf{I}_n \\ \mathbf{I}_n & \mathbf{M}_n \end{pmatrix}, \tag{1}$$

where $\mathbf{I}_n$ is the $2^n \times 2^n$ identity matrix.

Let $p_{\mathbf{M}_n}$ be the characteristic polynomial of $\mathbf{M}_n$, the roots of which are the eigenvalues of the matrix. Then

$$\begin{aligned} p_{\mathbf{M}_{n+1}}(\lambda) &= \det \begin{pmatrix} \lambda\mathbf{I}_n - \mathbf{M}_n & -\mathbf{I}_n \\ -\mathbf{I}_n & \lambda\mathbf{I}_n - \mathbf{M}_n \end{pmatrix} \\ &= \det(\mathbf{M}_n - (\lambda-1)\mathbf{I}_n) \cdot \det(\mathbf{M}_n - (\lambda+1)\mathbf{I}_n) \\ &= p_{\mathbf{M}_n}(\lambda)|_{\lambda=\lambda-1} \cdot p_{\mathbf{M}_n}(\lambda)|_{\lambda=\lambda+1}, \end{aligned}$$

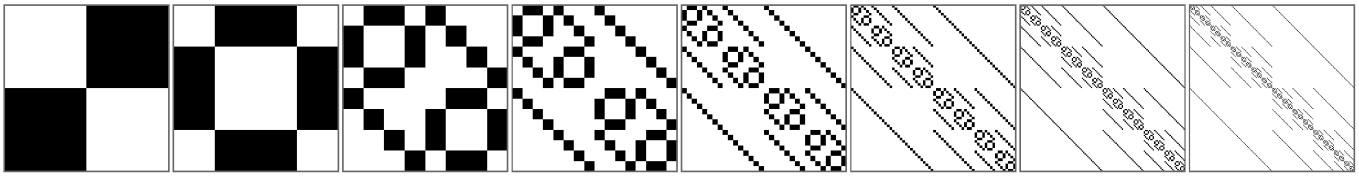

FIG. 1: The $\mathbf{M}$ matrices, for $n = 1$ to $n = 8$, which are of size 2×2 to 256×256, where 0 is white and 1 is black.

where $p_{\mathbf{M}_0} = \lambda$. Thus we see that the characteristic polynomial $p_{\mathbf{M}_{n+1}}$ is the product of $p_{\mathbf{M}_n}$ evaluated at $\lambda - 1$ and $p_{\mathbf{M}_n}$ evaluated at $\lambda+1$. This recursive step can be understood visually though a Pascal's triangle of the terms. The product of the terms in row n is the characteristic polynomial:

$$\begin{array}{ccccccccc} & & & & \lambda & & & & \\ & & & \lambda-1 & & \lambda+1 & & & \\ & & \lambda-2 & & \lambda^2 & & \lambda+2 & & \\ & \lambda-3 & & (\lambda-1)^3 & & (\lambda+1)^3 & & \lambda+3 & \\ \lambda-4 & & (\lambda-2)^4 & & \lambda^6 & & (\lambda+2)^4 & & \lambda+4. \end{array} \tag{2}$$

We immediately see that the eigenvalues of $\mathbf{M}_n$ are $n, n-2, n-4, \ldots, -n$, with degeneracies $\binom{n}{0}, \binom{n}{1}, \ldots, \binom{n}{n}$. The principal eigenvector of $\mathbf{M}_n$ is $(1, 1, 1, \ldots)$.

Selection

An environment amounts to an assignment of a fitness to each genotype. To give preferential treatment to certain genotypes, we need to define a fitness function. The fitness of each individual is determined by the distance of the genotype g from some target genotype $\tilde{g}$ which is optimally suited to the environment. The closer g is to the optimum $\tilde{g}$, the higher is the individual's fitness, and therefore reproduction rate. As the population repeatedly reproduces and mutates, it drifts towards the vicinity of this optimum $\tilde{g}$.

Let $\mathbf{f} = (f_1, \ldots, f_{2^n})$ be the vector of fitnesses for the 2^n genotypes, and let $\mathbf{F}$ be the matrix in which the diagonal is $\mathbf{f}$ and all other elements are zero. If we assume that every individual reproduces once before dying, then the distribution of the population at time t gets transformed to an updated distribution at time $t+1$ according to the transition matrix is $\mathbf{MF}$:

$$\mathbf{p}(t+1) = \mathbf{MF}\mathbf{p}(t). \tag{3}$$

While the results in this Letter are independent of the the fitness function $\mathbf{F}$, it is illustrative to consider an actual example. The Hamming fitness is the Hamming distance between g and some target genotype $\tilde{g}$, optimally suited to the environment, that is, the number of bits by which g and $\tilde{g}$ differ. For example, if $\tilde{g}$ is 11111 and g is 01101, then $h = 2$. The matrices $\mathbf{M}_5$, $\mathbf{F}_5$ and $\mathbf{M}_5\mathbf{F}_5$ for Hamming fitness are in the top of Fig. 2.

Starting from a given initial population at time $t = 0$, we can determine the population distribution at time t by repeatedly apply the matrix $\mathbf{MF}$, or just raising it to a power:

$$\mathbf{p}(t) = (\mathbf{MF})^t\mathbf{p}(0).$$

However, to find the shape of the steady state distribution, we don't actually need to raise the matrix to a power. This is given by the principal eigenvector of the matrix $\mathbf{MF}$, and the long term growth rate is given by the largest eigenvalue of $\mathbf{MF}$.

Age

The matrix $\mathbf{MF}$ tells us how a population with maximum age $a = 1$ evolves. Now let's consider a greater than 1.

We start with maximum age $a = 2$. Let $\mathbf{x} = (x_1, \ldots, x_{2^n})$ be the size of the population with age 1 and genotypes $1, \ldots, 2^n$, and let $\mathbf{y} = (y_1, \ldots, y_{2^n})$ be the size of the population with age 2 and genotypes $1, \ldots, 2^n$. Individuals with ages 1 and 2 can give birth, but all of the offspring are born with age 1. Let $\mathbf{p} = \mathbf{x}+\mathbf{y}$ be the total population size, of all ages. Then

$$\mathbf{x}(t+1) = \mathbf{MF}\mathbf{p}(t) \quad \text{and} \quad \mathbf{y}(t+1) = \mathbf{x}(t). \tag{4}$$

Inserting these into $\mathbf{p}(t+1) = \mathbf{x}(t+1) + \mathbf{y}(t+1)$ gives

$$\mathbf{p}(t+1) = \mathbf{MF}\mathbf{p}(t) + \mathbf{x}(t).$$

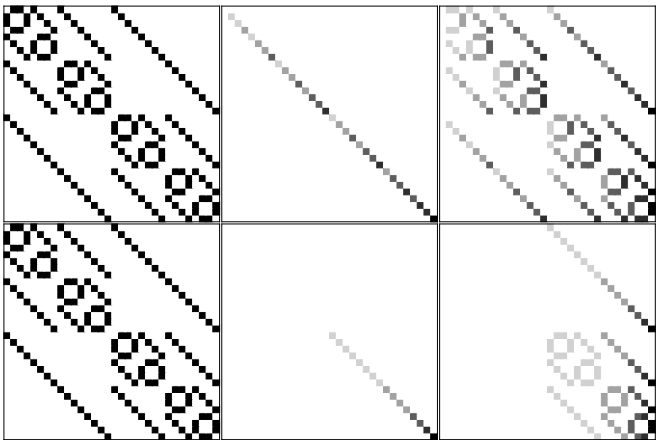

FIG. 2: The top row shows, for genome length $n = 5$, the mutation matrix $\mathbf{M}$, the Hamming fitness function $\mathbf{F}$, and their product $\mathbf{MF}$. The bottom row shows the mutation matrix $\mathbf{M}$, the overlap fitness function $\mathbf{F}$, and their product $\mathbf{MF}$. For maximum age $a = 1$, $\mathbf{Q} = \mathbf{MF}$, but for higher a, we must solve a matrix polynomial for $\mathbf{Q}$.

Incrementing both sides by one time step and again applying eq. (4) gives

$$\mathbf{p}(t+2) = \mathbf{MF}\mathbf{p}(t+1) + \mathbf{MF}\mathbf{p}(t). \tag{5}$$

Our goal is to obtain the transition matrix $\mathbf{Q}$ for which $\mathbf{p}(t+1) = \mathbf{Q}\mathbf{p}(t)$. Writing (5) in terms of $\mathbf{Q}\mathbf{p}(t)$, we find $\mathbf{Q}^2\mathbf{p}(t) = \mathbf{MF}(\mathbf{Q}+\mathbf{I})\mathbf{p}(t)$, and so $\mathbf{Q}$ satisfies

$$\mathbf{Q}^2 = \mathbf{MF}(\mathbf{I}+\mathbf{Q}).$$

This is a quadratic equation in terms of matrices, where $\mathbf{MF}$ is constant and we are solving for $\mathbf{Q}$.

We can take a similar approach for general maximum age a. Now we need to keep track of a population vectors, with ages $1, \ldots, a$, respectively. Instead of $\mathbf{x}$ and $\mathbf{y}$ that we used above, now we use $\mathbf{x_1}, \ldots, \mathbf{x_a}$ to indicate the vectors of populations with different ages. Individuals of all ages can give birth, but all of the offspring are born with age 1. Let

$$\mathbf{p}(t) = \mathbf{x_1}(t) + \ldots + \mathbf{x_a}(t) \tag{6}$$

be the total population size, of all ages. Then

$$\mathbf{x_1}(t+1) = \mathbf{MF}\mathbf{p}(t) \quad \text{and} \quad \mathbf{x_{i+1}}(t+1) = \mathbf{x_i}(t). \tag{7}$$

Inserting these into eq. (6) evaluated at time $t+1$ gives

$$\mathbf{p}(t+1) = \mathbf{MF}\mathbf{p}(t) + \mathbf{x_1}(t) + \ldots \mathbf{x_{a-1}}(t).$$

Incrementing the time in both sides by 1 and again applying eq. (7) gives

$$\mathbf{p}(t+2) = \mathbf{MF}\mathbf{p}(t+1) + \mathbf{MF}\mathbf{p}(t) + \mathbf{x_1}(t) + \ldots \mathbf{x_{a-2}}(t).$$

Repeating this process until all of the $\mathbf{x}$s are converted to $\mathbf{p}$s, we obtain

$$\mathbf{p}(t+a) = \mathbf{MF}(\mathbf{p}(t+a-1) + \ldots + \mathbf{MF}(\mathbf{p}(t).$$

Then, with $\mathbf{p}(t+1) = \mathbf{Q}\mathbf{p}(t)$, $\mathbf{Q}$ satisfies

$$\mathbf{Q}^a = \mathbf{MF}(\mathbf{I} + \mathbf{Q} + \ldots + \mathbf{Q}^{a-1}). \tag{8}$$

Since, for a general matrix $\mathbf{A}$,

$$\mathbf{I} + \mathbf{A} + \ldots + \mathbf{A}^{a-1} = (\mathbf{I}-\mathbf{A})^{-1}(\mathbf{I}-\mathbf{A}^a),$$

we can write $(\mathbf{I}-\mathbf{Q})\mathbf{Q}^a = \mathbf{MF}\,(\mathbf{I}-\mathbf{Q}^a)$ or, equally,

$$\mathbf{Q}^a(\mathbf{I} + \mathbf{MF} - \mathbf{Q}) = \mathbf{MF}. \tag{9}$$

We call this the mortality equation, and it is one of the main results of this Letter. Its concision belies its power. It gives the transition matrix of a population with maximum age a in terms of that with maximum age 1, for any fitness function $\mathbf{F}$. Notice that while the compact eq. (9) has degree $a+1$, it can always be reduced to degree a by dividing through by $\mathbf{I}-\mathbf{Q}$ to give eq. (8).

Solving the mortality equation

With our mortality equation in hand, what kind of meaning can we extract from it? It is unusual in that it is a polynomial equation in term of matrices rather than numbers. However, matrix polynomials are subject to most of the usual machinery for solving ordinary polynomials, with some extra care required for taking roots.

Taking roots of a matrix works as follows. A matrix $\mathbf{A}$ is diagonalizable if there exists an invertible matrix $\mathbf{P}$ and a diagonal matrix $\mathbf{D}$ such that $\mathbf{A} = \mathbf{PDP}^{-1}$. The roots of such a matrix $\mathbf{A}$ are then $\sqrt[n]{\mathbf{A}} = \mathbf{P}\sqrt[n]{\mathbf{D}}\mathbf{P}^{-1}$, where $\sqrt[n]{\mathbf{D}}$ just amounts to taking the roots of the diagonal elements, bearing in mind to take all n roots of each.

To provide some intuition for our mortality equation, we first consider the case of constant fitness across all genotypes, that is, we set $\mathbf{F} = \mathbf{I}$. Applying $\mathbf{M}$ to a population has the effect of diffusing it over the hypercube, smoothing it out towards a uniform population. Setting $a = 2$, eq. (9) becomes

$$\mathbf{Q}^2 = \mathbf{M}(\mathbf{Q}+\mathbf{I}). \tag{10}$$

Eq. (10) can be solved like an ordinary quadratic:

$$\mathbf{Q} = \mathbf{M}/2 \pm \sqrt{(\mathbf{M}/2)^2 + \mathbf{M}}. \tag{11}$$

The actual matrices will depend on n. For $n = 2$,

$$\mathbf{Q} = \mathbf{M}/2 \pm \left(\left(\sqrt{3}+i\right)\mathbf{1} - 2\,i\,\mathbf{M}\right)/4, \tag{12}$$

where $\mathbf{1}$ is the all-1s matrix. These two matrices have eigenvalues $1+\sqrt{3}, 0, 0, -1+i$ and $1-\sqrt{3}, 0, 0, -1-i$.

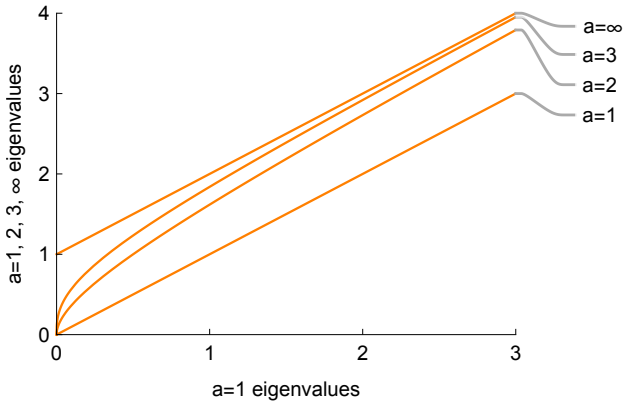

FIG. 3: The eigenvalues for maximum age $a = 1$ (of the matrix $\mathbf{Q}$) are mapped to higher eigenvalues for $a = 2$ and $a = 3$, according to the curves shown here. The $a = 2$ curve is the solution to $x^2 = \lambda(1+x)$ and the $a = 3$ curve is the solution to $x^3 = \lambda(1+x+x^2)$. Curves for higher maximum age a approach the top line $x+1$, which corresponds to $a = \infty$ (immortality). These new eigenvalues correspond to the transition matrix $\mathbf{Q}$, described by our mortality equation in eq. (9). This mapping of the eigenvalues is independent of the fitness function $\mathbf{F}$, which only determines the initial eigenvalues, but not how they get shifted.

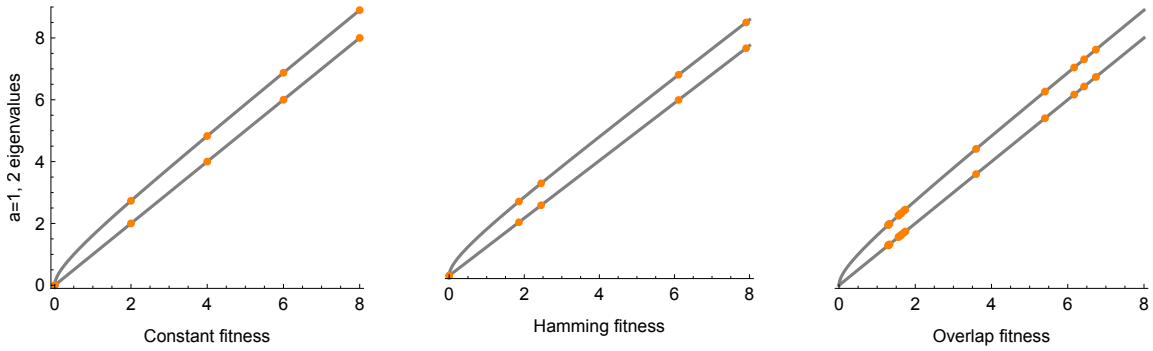

FIG. 4: We tested our theory for the shift in eigenvalues by simulating populations evolving according to three different fitness functions: constant fitness, Hamming fitness and overlap fitness. In particular, for each fitness we compare the eigenvalues for populations with maximum age $a = 1$ and $a = 2$ (points), and find complete agreement with our prediction (lines).

For $n = 3$,

$$\mathbf{Q_3} = \mathbf{M_3}/2 \pm \Big(\big(\sqrt{21} - \sqrt{5}\big)\mathbf{1} + 4\big(\sqrt{5} + \sqrt{3}i\big)\mathbf{I} - 2\Big(\sqrt{5} - \sqrt{3}i\Big)(\mathbf{J} - \mathbf{M_3})\Big)/16, \quad (13)$$

where $\mathbf{I}$ and $\mathbf{J}$ are the identity and antidiagonal identity matrices. These two matrices have eight eigenvalues each. Notice that even though $\mathbf{M}$ and $\mathbf{F}$ are real, $\mathbf{Q}$ can be complex.

We can in principle extend this process to solving for higher n, and indeed for larger maximum age a. However, this requires painstaking matrix diagonalization and inversion and is not analytically tractable. The good news it that there is a shortcut for computing the eigenvalues of $\mathbf{Q}$ from those of $\mathbf{MF}$.

When it comes to the eigenvalues of $\mathbf{Q}$, what's good for the goose is good for the gander. The eigenvalues of $\mathbf{Q}$ have the same functional relation to those of $\mathbf{MF}$ as the matrix $\mathbf{Q}$ does to the matrix $\mathbf{MF}$. Let $\mu_1, \ldots, \mu_{2^n}$ be the eigenvalues of $\mathbf{Q}$. We can express them in terms of $\lambda_1, \ldots, \lambda_{2^n}$, the eigenvalues of $\mathbf{MF}$, using the analogue to our mortality equation in eq. (9), but for numbers rather than matrices:

$$\mu_i^a(1 + \lambda_i - \mu_i) = \lambda_i, \quad (14)$$

where λ_i is constant and we are solving for μ_i. For $a = 2$ we can solve this explicitly:

$$\mu_i = \lambda_i/2 \pm \sqrt{(\lambda_i/2)^2 + \lambda_i}. \quad (15)$$

As a test, we know from eq. (2) that the eigenvalues of M for $n = 2$ are $-2, 0, 0, 2$. Plugging these into eq. (15) gives the same eigenvalues as we obtained for the matrices in eq. (12).

Discussion

The growth rate of our population is set by μ_1, the largest eigenvalue of $\mathbf{Q}$. As Fig. 1 illustrates, μ_1 increases with the maximum age a, from λ_1 at $a = 1$ to $\lambda_1 + 1$ at $a = \infty$. Therefore the fastest growing population is the immortal one. In a constant environment, there is no growth rate benefit afforded by aging; mortality is a losing strategy.

However, in reality populations do not reach steady state, where they are optimally suited to their environment. Rather, they are in continually out of equilibrium, adapting to a continuously shifting environment. The rate of convergence is determined by the so-called spectral gap, which is the difference between the first and second eigenvalues. As we shall see, the spectral gap always decreases with the maximum age a, regardless of the choice of fitness function $\mathbf{F}$. First let's prove that the gap decreases from $a = 1$ to $a = 2$. The condition is

$$(\lambda_1 - \lambda_2)/\lambda_1 > (\mu_1 - \mu_2)/\mu_1.$$

Substituting eq. (15) in for μ_1 and μ_2 and rearranging, this becomes

$$\lambda_1 + \sqrt{\lambda_2(\lambda_2 + 4)} < \lambda_2 + \sqrt{\lambda_1(\lambda_1 + 4)}.$$

Since $\lambda_1 > \lambda_2$, this is always true.

We can show that the spectral gap continues to diminish as a increases using the convexity of the eigenvalue curves in Fig. 1.

To test our approach for predictions, we simulated a population evolving according to three fitness functions: constant fitness, Hamming fitness and overlap fitness.

Constant fitness, which we looked at earlier, is our control. In this case the eigenvalues for $a = 1$ (and thus of the matrix $\mathbf{MF}$) occur at regular intervals, and their aging equivalents for $a = 2$ lie on the line $x/2 + \sqrt{((x/2)^+ x)}$. We plot both sets of eigenvalues in Fig. 2C, and see that they lie on the two predicted lines.

Second, we considered the Hamming fitness h, described earlier. Third, we considered the overlap fitness v, which is the length of the longest initial sub-string over which g and the $\tilde{g}$ match. For example, if $\tilde{g}$ is 11111 and g is 11001, then $v = 2$, since the first two bits in g and $\tilde{g}$ match, but not the first three.

That programmed death, which we also refer to as mortality, can have an evolutionary advantage is further evidence that it is not fundamental to life itself.

This work motivates a number of open questions, and we hope that others will pick up where we have left off. First, we conjecture that the general transition matrix **MF** has only real eigenvalues and that they come in plus-minus pairs. Second, a complete solution for a non-trivial fitness function **F**, such as our Hamming or overlap fitness or something else, would shed light on any connections between maximum age a and genome size n. Third, even if aging does confer an evolutionary benefit, it remains to show how this might be bootstrapped into existence.

This research was supported by a grant from Bit.Bio. We acknowledge Robert Farr for helpful discussions and for introducing the concept of the overlap fitness. T. F. acknowledges hospitality from Oneiro in Corfu, Greece.

[1] M. Rozing, T. Kirkwood, R. Westendorp, Is there evidence for a limit to human lifespan?, *Nature* **546**, E11 (2017).
[2] T. B. Kirkwood, *Aging Cell* **15**, 986 (2016).
[3] T C Goldsmith, Aging is programmed!(A response to Kowald-Kirkwood), *Aging Cell* (2016).
[4] TB Kirkwood, S Melov, On the programmed/non-programmed nature of ageing within the life history, *Curr Biol* **21**, R701 (2011).
[5] T. C. Goldsmith Aging, evolvability, and the individual benefit requirement, *J. Theor. Biol.* **252**, 764 (2008).
[6] T. C. Goldsmith On the programmed/non-programmed aging controversy. *Biochemistry* **77**, 729 (2012).
[7] T. C. Goldsmith Arguments against non-programmed aging theories. *Biochemistry* **78**, 971 (2013).
[8] J. Werfel, D.E. Ingber, Y. Bar-Yam, Programed death is favored by natural selection in spatial systems, *Phys Rev Lett*, **114**, 238103 (2015).
[9] J. Werfel, Ingber DE, Bar-Yam Y, Theory and associated phenomenology for intrinsic mortality arising from natural selection, *PLoS ONE*, **12**, e0173677 (2017).
[10] J. Travis, The evolution of programmed death in a spatially structured population, *J Gerontology*, **59A**, 301 (2004).
[11] J Mitteldorf, A C Martins Programmed life span in the context of evolvability, *Am. Nat.* **184**, 289 (2014)