Modeling and optimization of populations subject to time-dependent mutation

(affinity maturation/somatic hypermutation/optimal control/stochastic model)

THOMAS B. KEPLER* AND ALAN S. PERELSON†

*Biomathematics Program, Department of Statistics, North Carolina State University, Raleigh, NC 27695-8203; and †Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545

Communicated by Manfred Eigen, Max-Planck-Institut für Biophysikalische Chemie, Göttingen, Germany, May 15, 1995

ABSTRACT It has become clear that many organisms possess the ability to regulate their mutation rate in response to environmental conditions. So the question of finding an optimal mutation rate must be replaced by that of finding an optimal mutation schedule. We show that this task cannot be accomplished with standard population-dynamic models. We then develop a “hybrid” model for populations experiencing time-dependent mutation that treats population growth as deterministic but the time of first appearance of new variants as stochastic. We show that the hybrid model agrees well with a Monte Carlo simulation. From this model, we derive a deterministic approximation, a “threshold” model, that is similar to standard population dynamic models but differs in the initial rate of generation of new mutants. We use these techniques to model antibody affinity maturation by somatic hypermutation. We had previously shown that the optimal mutation schedule for the deterministic threshold model is phasic, with periods of mutation between intervals of mutation-free growth. To establish the validity of this schedule, we now show that the phasic schedule that optimizes the deterministic threshold model significantly improves upon the best constant-rate schedule for the hybrid and Monte Carlo models.

The optimal replication fidelity for a population of organisms represents a compromise between the need to preserve hard-won gains in fitness and the need to maintain adaptability in an unpredictably changing environment. If replication fidelity is too low, deleterious mutations will quickly accumulate, and the population will “diffuse” away from its adaptive fitness peak (1). If, on the other hand, replication fidelity is too high, a population that is presently well-adapted may not be able to adjust to subsequent changes in the fitness landscape. Several investigators have examined this question in several contexts and found optimal mutation rates (2–4).

It is now apparent that some organisms partially resolve these conflicting needs by regulating their mutation rates in response to environmental circumstances. When the population is thriving, their genomes are maintained with high fidelity, but, when stressed, a specific mechanism is invoked to up-regulate the mutation rate. The best studied among these systems is the response to DNA lesions in *Escherichia coli* known as the SOS response (5). The SOS response increases mutation rates throughout the genome as well as in the specific area of the lesions. Genes that play a central role in the regulation of the SOS response have recently been implicated (6) in the phenomenon of “directed mutation” (7, 8). Induced genomic alteration in response to environmental conditions has been studied in eukaryotes as well (9, 10). An inducible down-regulation of replication fidelity mediated by biasing the nucleotide pool has been noted in association with human

immunodeficiency virus (11), although the implications of this finding are unclear at present.

A remarkable system for the study of mutation and selection is the affinity maturation of the humoral immune response by somatic hypermutation (12–14). Upon exposure to a foreign antigen, B cells whose antigen receptors are specific for that antigen proliferate and differentiate. A subset of these colonizes germinal centers in the lymph nodes and spleen (15, 16). Here they divide rapidly and the genes that encode the variable regions of their immunoglobulins (heavy and light chains) experience a 10^4 - to 10^5 -fold increase in their mutation rate (17). If a higher-affinity variant immunoglobulin arises by chance in this process, the B cell that produces it binds antigen more easily and is preferentially stimulated. When antigen becomes limiting, these high-affinity variants grow to dominate the population. When the germinal center reaction dwindles, after 14 days or so, the mutation rate appears to return to normal. This extraordinary miniature Darwinian evolution typically produces increases in affinity of from 10- to 100-fold (12).

For organisms that change replication fidelity as a result of conditions in the environment, the issue of optimizing their mutation rate is more complicated. Optimization must be performed over the mutation rate considered as a function of time. One must find the best mutation *schedule*. This problem is considerably harder and presents some interesting mathematical and computational challenges. One approach that has been successfully employed to find optimal *constant* mutation rates is to apply deterministic ordinary differential equation (ODE) population-dynamic modeling and include appropriate terms for mutation (18–20). This approach can be shown to give accurate results when the population growth is not density-dependent or when all populations are always large. Of course, when novelty arises, the population representing it is very small; uncertainty as to the time of arrival of the first mutants implies a large variance in the size of the mutant subpopulation, which, in the presence of nonlinearities in the growth equations, feeds back to influence the mean. Thus, the deterministic model is no longer accurate. The problem is greatly exacerbated by time-dependent mutation rates, which can act to maintain a large coefficient of variation.

In an ordinary continuum description using concentrations or real-valued population sizes, mutations do not generate integer-valued jumps in population size. Thus, it is not possible, within the context of these models, to ask *when* a mutation takes place. This situation is similar to one faced in modeling diffusion, for example, where the solution of Fick's equation gives no sense of when the first particle in a finite population will cross a given boundary; the concentration is non-zero at all points immediately and there is no sense of “one particle.” But in a population model, these “phantom” concentrations of instantly produced mutants experience exponential growth and can produce wildly mistaken results.

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. §1734 solely to indicate this fact.

Abbreviation: ODE, ordinary differential equation.

A single short burst of mutation seeds all mutant classes connected to the initial population. The size of these mutant "populations" might be minuscule, but those that are advantageous will quickly grow to dominate the population, even though the time that mutation was active can be made arbitrarily short, so that the probability that a single advantageous mutant would have been created is vanishingly small.

Rather than using ODEs, one could model mutation and growth with a Markov model. Agur *et al.* (21) used a Markov model with linear growth rates to determine the optimal start time for a mutation process when seeking a single-step mutant. This approach is very difficult to implement when generalized to multiple mutations and to more realistic density-dependent growth rates. Furthermore, in spite of their relative complexity, single age-class Markov models are themselves inaccurate, greatly overestimating the population variance by overestimating the variance in the length of the generation times.

We have developed an alternative approach to modeling populations with time-dependent mutation rates. Our strategy is to neglect fluctuations due to random variations in the length of generation times but to account fully for the randomness in the times at which mutations arise. These mutation arrival times will be introduced as random variables, upon which the vector fields of the population ODEs will then depend. In turn, the differential equations describing the behavior of the probability density functions for the mutation arrival times will depend on the state of the population. We call this model a *stochastic hybrid model*. The introduction of mutants into a population at random times has been a strategy used in ODE-based simulation models of AIDS (22).

From this hybrid model we derive an alternative deterministic model by replacing the mutation arrival times by dynamically determined characteristic values for those times. This new ODE model, which we call a *threshold model*, contains discrete thresholds that need to be surpassed before mutant populations are introduced and is appropriate for investigating the behavior of systems experiencing time-dependent mutation and, in particular, for optimizing over mutation schedules.

We will illustrate the process by modeling the somatic mutation of B cells. We have previously shown that, for a B-cell population model of the "threshold" type, the schedule that maximizes a measure of "fitness" is one in which periods of high mutation rates alternate with periods of mutation-free growth (23, 24). Having obtained this optimum schedule using an algorithm based on the Pontryagin maximum principle, we now test this schedule using the hybrid stochastic model as well as Monte Carlo simulations of the population dynamics regarded as an age-structured Markov model. We cannot easily say whether or not this schedule remains optimal for the Markov model, but we do show that it is a significant improvement over the best constant-rate schedule.

Models

We write a general population-dynamic equation with mutation, under the assumption that mutations are expressed only during replication, as follows:

$$\frac{dx_i}{dt} = g_i(x, \mu)x_i + \sum_{j \neq i} m_{ji}(\mu)b_j(x)x_j, \quad [1]$$

where x_i is the population size of the i th strain or affinity class (24), g_i is the net growth rate of this class, which owing to competition generally will depend on the total population composition as well as on the mutation rate μ . The mutation matrix m_{ji} gives the probability that a given daughter of a cell of class j will be of class i . The birth rate is b_i . We consider the case where the fitnesses of the various classes differ and ask what mutation schedule most quickly facilitates the dominance

of the population by fitter classes. This is an appropriate question when, for example, the environment is changing in such a way that the survival of the population depends on the rapid establishment of fitter classes. In the case of antibody affinity maturation, we assume fitness is equivalent to the binding affinity of antibody for the eliciting antigen; increasing the affinity enhances the efficacy of the immune response.

We have formulated a model for germinal center B-cell dynamics (23, 24) that is a special case of Eq. 1, with

$$g_i = r[-1 + 2e^{-\mu}h_i(\alpha)]. \quad [2]$$

Here r is the cell division rate [cell death is assumed to also occur at rate r (24)], α is the concentration of free antigen, which we may think of as a resource necessary for growth, and h_i is the probability that a cell reaching the end of its cycle divides rather than undergo programmed cell death, and is given by

$$h_i(\alpha) = \frac{\alpha K_i}{1 + \alpha K_i}, \quad [3]$$

where K_i is the affinity of the antigen receptor of class i for the antigen. Furthermore,

$$b_i = 2rh_i(\alpha), \quad [4]$$

and

$$m_{ij} = \frac{\{\mu(1 - p_L)\}^{|i-j|} e^{-\mu}}{|i-j|!(1 + \Lambda^{-i})}, \quad [5]$$

where μ is the effective genomic mutation rate (gen^{-1}), i.e., the mutation rate (ignoring silent and functionally silent mutations) per base times the number of base pairs targeted for mutation, p_L is the probability that a non-silent mutation is lethal (or nonbinding), and Λ is a constant that determines the fraction of expressed, nonlethal mutations of a cell that are advantageous or disadvantageous (see ref. 24 for details). The free antigen obeys the following conservation equation,

$$\alpha + \sigma \sum_i x_i h_i(\alpha) = \alpha_0, \quad [6]$$

where σ is the number of antigen receptors per B cell and α_0 is the (constant) total amount of antigen. In writing Eq. 6 we assume that h_i is equal to the proportion of antigen receptors that are bound by antigen.

We will maximize the total affinity $A(T)$, defined as the product of the average affinity and the total population size, or $A(t) \equiv \sum_i K_i x_i(t)$ evaluated at some specified time, T . This maximization is over all mutation schedules and is performed using an algorithm based on the Pontryagin maximum principle (23, 25). A direct maximization of the system specified by Eqs. 1–6 with a single class initially occupied and one lower- and two higher-affinity classes available yields the schedule of Fig. 1A, showing an initial sharp, brief peak of mutation followed by mutation-free growth for the remaining time.

With this "optimal" schedule, the highest-affinity class completely dominates the population by generation 40 (Fig. 1B). But when this schedule is used in a Monte Carlo simulation of the population, we find that <0.2% of replicate runs produce affinity class 2 cells and that the unmutated cells, $i = 0$, dominate the population (Fig. 1C). In the ODE model, the explosive process of cell proliferation has taken a tiny concentration of mutants, corresponding to much less than one cell, and blown it up into the dominant clone. To avoid this problem, we must reinterpret concentrations that correspond to less than one individual and adjust our models accordingly.

FIG. 1. (A) Optimal mutation schedule (see text for optimality criterion) for the B-cell hypermutation model given by Eqs. 1–6, with affinity classes $i = -1, 0, 1, 2$, and initial condition $x_0(0) = 2$ cells, $x_i(0) = 0$ for $i \neq 0$. The parameters $r = \log 2 \text{ gen}^{-1}$, $\alpha_0 = 3.3 \text{ } \mu\text{M}$, $\sigma = 3.3 \times 10^{-4} \text{ } \mu\text{M} \cdot \text{cell}^{-1}$ (corresponding to 10^5 antigen receptors per B cell and a system of volume $5 \times 10^5 \text{ } \mu\text{m}^3$), $K_i = 5^i \times 10^6 \text{ M}^{-1}$, $p_L = 0.5$, and $\Lambda = 100$, corresponding to a frequency of advantageous mutations among all replacement mutations of about $1/100$. (B) Population dynamics of conventional ODE model. The seed class is $i = 0$ and i designates the number of advantageous mutations acquired. Note that by generation 40, the highest affinity class, $i = 2$, is at maximum size, and all other classes are negligible. (C) Mean population sizes over 10,000 replicate runs of the Monte Carlo simulation of the age-structured Markov model described in the text (dashed lines) and the mean population sizes for the hybrid model of Eqs. 7 and 8 and 2–6 (solid lines), both using the schedule shown in A.

A Stochastic Hybrid Model

The failure of this and other standard ODE models stems from their lack of a size scale corresponding to a single individual. This is manifest in the absence of any sense of the time it takes for the first mutant to arise (in fact, all mutants arise immediately). Therefore, we extend the state space for our model, by including discrete random variables that indicate whether or not a given class has become occupied—i.e., whether a mutation to that class has yet occurred. Since these are stochastic variables, we keep track of their probability distributions. Let $P_i(t)$ be the probability that the i th mutant class has *not* become occupied by time t . A reasonable description, based on a Markov assumption, is given by

$$\frac{dP_i}{dt} = -P_i \sum_{j \neq i} m_{ji}(\mu) b_j(x) x_j. \quad [7]$$

The coefficient on the right is the intensity, or the rate at which individuals of class i are produced by mutation from individuals of all other occupied classes.

Now Eq. 1 is modified to read

$$\frac{dx_i}{dt} = [g_i(x, \mu) x_i + \sum_{j \neq i} m_{ji}(\mu) b_j(x) x_j] \theta(t - t_i) + \delta(t - t_i), \quad [8]$$

where $\theta(t)$ is the Heavyside function, which is 1 for $t > 0$ and is 0 otherwise. The Dirac delta is the derivative of the Heavyside function, and t_i is the stochastic variable giving the time at which the first mutant of class i appears. The term $\delta(t - t_i)$ bumps the population in class i from zero to one at time t_i , after which the familiar dynamics (identical to those prescribed by Eq. 1) take over.

Notice that once the t_i s are prescribed, Eq. 8 is deterministic, hence the term “hybrid” model. This system may be solved as a “path integral.” One fixes the t_i and integrates Eq. 8. Then Eq. 7 is used to find the probability density at this t_i . This model has some similarity to the Luria–Delbrück model (26, 27) for mutation in bacterial populations, in which an ODE was used to model for the growth of the unmutated population and a stochastic model for the generation of variants.

A Deterministic Threshold Model

Eqs. 7 and 8 prescribe a stochastic dynamical system that more accurately models populations experiencing mutations and especially temporally variable mutations. We can derive a simpler deterministic model, useful for qualitative analysis, by replacing the stochastic variables t_i by real variables τ_i defined by

$$P_i(\tau_i) = 1/e. \quad [9]$$

The τ_i are similar to medians [which are defined by $P_i(\tau_i) = 1/2$]. Note that the mean is an inappropriate characteristic value because it depends on the *future* course of the dynamical system. The mode is also poor because it is very sensitive to local fluctuations in the intensity, and hence, to spikes in the mutation rate.

The deterministic system we study is given by replacing t_i with τ_i in Eq. 8, retaining Eq. 7, and adding Eq. 9. Note that the equations for P_i and x_i are never used simultaneously. P_i is calculated until it reaches the value $1/e$, after which it is no longer needed. On the other hand, x_i is inert until $P_i = 1/e$, at which time it jumps to the value 1 and then changes according to the modified Eq. 8. Under these circumstances, we can define a new variable ξ_i that equals x_i when $t > \tau_i$ and $-\log P_i$ when $t \leq \tau_i$. This new variable is continuous and obeys

$$\frac{d\xi_i}{dt} = g_i(\xi, \mu) \xi_i \theta(\xi_i - 1) + \sum_{j \neq i} m_{ji}(\mu) b_j(\xi) \xi_j \theta(\xi_j - 1), \quad [10]$$

which is quite similar to the “traditional” ODE, but contains Heavyside functions to switch from passive to active growth.

Example: Optimization

We have previously taken Eq. 10 as our starting point in the optimization of the mutation schedule (23, 24). Here, we show that the schedule obtained using the threshold model represents a significant improvement over the best constant-rate schedule for corresponding stochastic models realized by Monte Carlo methods.

In the hybrid model cell division is deterministic. To generate a comparable stochastic model we want the variance in the time to division to be a parameter. The Markov model we use for this purpose is related to the Smith and Martin (28) model of the cell cycle, and has N “immature” stages or “ages” and a single mature state. Upon passing through the mature state, the cell either divides or dies. The rate of passage from age a to age $a + 1$, or from age N to either death or division, depends on the affinity class and is taken for convenience to be $(2N - a)h_i$. The probability of dividing at the mature state is h_i as in the hybrid model, and the mutation matrix is also carried over. The standard deviation of the intermitotic time for rapidly dividing eukaryotic cells (in this case, neuroblas-

tomas) has been carefully measured (29) and amounts to 10% of the mean. We have used $N = 25$ age classes, giving a cycle-length standard deviation about 20% of the mean. It is computationally unfeasible to use significantly larger N to lower the standard deviation further.

We have found, using techniques described earlier (23, 24), the optimal schedule for the threshold ODE system corresponding to the above Markov model (Fig. 2A). The interpretation of this schedule is straightforward. The initial period when the mutation rate is essentially zero allows the seed population to grow as rapidly as possible, without loss to deleterious mutations. When this population has reached a critical size, inversely proportional to the probability of a given single mutation being advantageous, the mutation rate is quickly raised, so that an advantageous mutant will quickly arrive. Once this advantageous mutation has occurred, the mutation rate is again lowered to let this new subpopulation grow at the maximum rate.

This simple description works well for the deterministic model that generated this schedule, but there remains some question about how well this strategy will work in the more realistic situation when the mutation arrival time is random. In Fig. 2C we show the population dynamics predicted using the hybrid and Monte Carlo models. These agree well with each other and with the predictions of the threshold model (Fig. 2B). In Fig. 3 we compared the phasic schedule to the optimum constant rate schedule (which we determined occurs with $\mu = 0.075 \text{ gen}^{-1}$) in the hybrid and Monte Carlo models. Fig. 3 shows the density of first mutation times for the passage from the seed class to class one and from class one to class two. The histograms represent 10,000 replicates of the Monte Carlo simulation and the solid curves show the results for our hybrid model. Differences between the Monte Carlo and hybrid models are due primarily to the large variance of cell cycle length in the former. Even with 25 age classes, the standard deviation of the cell division time is 20% of the mean division

FIG. 2. (A) Optimal mutation schedule using the threshold ODE model given by Eq. 10 and Eqs. 2–6, with the same parameters as in Fig. 1. (B) Population dynamics for this threshold model. (C) Analogous to Fig. 1C, but computed using the schedule in A.

FIG. 3. Histograms (normalized to probability density functions) for the mutation first-arrival times in 10,000 replicate runs of the Monte Carlo simulation. The solid lines are the density functions calculated using the hybrid model. (A and C) Arrival times of the first cells in class $i = 1$. (B and D) Arrival times of the first cells in class $i = 2$. A and B were produced using a constant mutation rate of $\mu = 0.075 \text{ gen}^{-1}$, which we have found, within the limits imposed by fluctuations, to be optimal among constant schedules. C and D were produced using the schedule shown in Fig. 2A. During the first mutation phase 88.9% of mutations to class $i = 1$ occur, 11.0% occur during the second phase, and 0.1% do not occur at all. The analogous percentages for mutations to class $i = 2$ are 0.6%, 78.5%, and 20.9%.

time. The hybrid model neglects this variance; all the variance in this model is due to the uncertainty in the mutation arrival time. The agreement is worse when the variance is still higher, but lowering the variance further by increasing the number of age classes is computationally prohibitive. The agreement at this level is quite good nevertheless. The mutation arrival times are moved forward in the optimal schedule with respect to the constant rate schedule.

We define an advantage index by

$$n_a \equiv \frac{\sigma}{2\alpha_0} \sum_i i x_i, \quad [11]$$

where the factor in front of the sum is the reciprocal of the maximum possible population size. Then n_a is the product of the average class index in the population and the size of the population relative to its "carrying capacity." The average $\langle n_a \rangle$ of the advantage index over 10,000 Monte Carlo replicates is, for the best constant schedule ($\mu = 0.075 \text{ gen}^{-1}$), $\langle n_a \rangle = 0.965$, and for the optimal schedule, $\langle n_a \rangle = 1.300$. Fig. 4 shows the empirical density functions for n_a for both schedules, as well as the density functions calculated using the hybrid model, and shows that our optimal schedule significantly enhances the density over large values of n_a .

Conclusions

We have developed a population-dynamic model that treats population growth as deterministic but the arrival times of mutations as stochastic. This hybrid model can then be further simplified to yield a deterministic approximation appropriate for finite populations experiencing temporally varying mutation rates. Using this threshold model, we have shown that the mutation schedule that maximizes the efficiency of the mutation and selection process is phasic, with periods of increased mutation rate punctuating intervals of mutation-free growth. This schedule maintains an advantage over the best constant-

FIG. 4. Histograms (normalized to probability density functions) for the advantage index (described in the text) based on 10,000 runs of the Monte Carlo simulation of the age-structured Markov model. The panel labeled "optimum" was made using the schedule of Fig. 2A; that labeled "constant" was made using a constant schedule with $\mu = 0.075 \text{ gen}^{-1}$. The solid lines were calculated using the hybrid model.

rate schedule when applied to a Monte Carlo simulation of an age-structured Markov model. For the constant and phasic mutation schedules, the hybrid model fits the Monte Carlo data adequately.

While the phasic mutation schedule was derived from an affinity maturation model it may have advantages over constant-rate mutation schedules in many other contexts. When this phasic schedule is viewed in the context of species evolution, it is suggestive of punctuated equilibrium (30). The significance of this relationship is unclear at present.

Self-regulation of mutation rate does occur among prokaryotes and eukaryotes and somatically among cellular constituents of higher organisms. The manipulation of these mutation rates may enhance fitness in a significant way.

We thank George I. Bell for useful comments. This work was partially performed under the auspices of the U.S. Department of Energy and supported by National Institutes of Health Grant AI28433 (A.S.P.), National Science Foundation Grant MCB-9357637 (T.B.K.),

and the Santa Fe Institute Theoretical Immunology Program through a grant from the Joseph P. and Jeanne M. Sullivan Foundation.

1. Muller, H. J. (1964) *Mutat. Res.* **1**, 1–9.
2. Leigh, E. G. (1973) *Genetics*, Suppl., **73**, 1–18.
3. Sasaki, A. & Isawa, Y. (1989) *Theor. Popul. Biol.* **39**, 201–239.
4. Kaneko, K. & Ikegami, T. (1992) *Physica D* **56**, 406–429.
5. Echols, H. & Goodman, M. F. (1991) *Mutat. Res.* **236**, 301–311.
6. Harris, R. S., Longerich, S. & Rosenberg, S. M. (1994) *Science* **264**, 258–260.
7. Cairns, J. J., Overbaugh, J. & Miller, J. (1988) *Nature (London)* **283**, 26–30.
8. Hall, B. G. (1988) *Genetics* **120**, 887–897.
9. Harrison, B. J. & Fincham, J. R. (1964) *Heredity* **19**, 237–258.
10. Hall, B. G. (1992) *Proc. Natl. Acad. Sci. USA* **89**, 4300–4303.
11. Vartanian, J.-P., Meyerhans, A., Sala, M. & Wain-Hobson, S. (1994) *Proc. Natl. Acad. Sci. USA* **91**, 3092–3096.
12. Berek, C. & Milstein, C. (1987) *Immunol. Rev.* **96**, 23–41.
13. Kocks, C. & Rajewsky, K. (1989) *Annu. Rev. Immunol.* **7**, 537–559.
14. French, D. L., Laskov, R. & Scharff, M. D. (1989) *Science* **244**, 1152–1157.
15. Jacob, J., Kelsoe, G., Rajewsky, K. & Weiss, U. (1990) *Science* **354**, 389–392.
16. Jacob, J., Kassir, R. & Kelsoe, G. (1991) *J. Exp. Med.* **173**, 1165–1175.
17. Rada, C., Gupta, S. K., Gheradi, E. & Milstein, C. (1991) *Proc. Natl. Acad. Sci. USA* **88**, 5508–5512.
18. Eigen, M., McCaskill, J. & Schuster, P. (1989) *Adv. Chem. Phys.* **75**, 149–263.
19. Nowak, M. A. (1992) *J. Theor. Biol.* **155**, 1–20.
20. Sasaki, A. (1994) *J. Theor. Biol.* **168**, 291–308.
21. Agur, Z., Mazon, G. & Meilijson, I. (1991) *Proc. R. Soc. London B* **245**, 147–150.
22. Nowak, M. A., Anderson, R. M., McLean, A. R., Wolfs, T. F. W., Goudsmit, J. & May, R. M. (1991) *Science* **254**, 963–969.
23. Kepler, T. B. & Perelson, A. S. (1993) *Immunol. Today* **14**, 412–415.
24. Kepler, T. B. & Perelson, A. S. (1993) *J. Theor. Biol.* **164**, 37–64.
25. Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V. & Mischenko, E. F. (1962) *The Mathematical Theory of Optimal Processes* (Wiley, New York).
26. Luria, S. & Delbrück, M. (1943) *Genetics* **28**, 491–511.
27. Lea, D. E. & Coulson, C. A. (1949) *J. Genet.* **49**, 264–285.
28. Smith, J. A. & Martin, L. (1973) *Proc. Natl. Acad. Sci. USA* **70**, 1263–1267.
29. Van Zoelen, E. J. J., Van der Saag, P. T. & De Laat, S. W. (1981) *Exp. Cell Res.* **131**, 395–406.
30. Eldredge, N. & Gould, S. J. (1972) in *Models in Paleobiology*, ed. Schopf, T. J. M. (Freeman, San Francisco), pp. 82–115.