Mutation Rates as Adaptations

Carlo Maley*
545 Technology Sq., NE43-803
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
email: cmaley@ai.mit.edu

May 2, 1996

Abstract

In order to better understand life, it is helpful to look beyond the envelop of life as we know it. A simple model of coevolution was implemented with the addition of a gene for the mutation rate of the individual. This allowed the mutation rate itself to evolve in a lineage. The model shows that when the individuals interact in a sort of zero-sum game, the lineages maintain relatively high mutation rates. However, when individuals engage in interactions that have greater consequences for one individual in the interaction than the other, lineages tend to evolve relatively low mutation rates. This model suggests that different genes may have evolved different mutation rates as adaptations to the varying pressures of interactions with other genes.

1 The Possibilities of Life

Lewontin suggested that when we think about what life could be, rather than what we see before us, there are two dimensions that might be considered: mortality and heredity.[27] The life with which we are most familiar is mortal and reproduces with high fidelity. But why not mortal organisms that pass on little information to their offspring? For that matter, why not immortal organisms with varying degrees of hereditary transmission of information? This paper is an attempt to address one aspect of these questions, the evolution of mutation rates. Neither mortality nor heredity need be dichotomous qualities. Mortality might range over varying degrees of longevity. Heredity can be addressed

*This work was supported in part by the generosity of the MIT AI Lab. I am greatful to R. Brooks, H. Westerhoff, M. Donoghue, P. Goss, K. Rice and L. King for their comments and support. M. Maley deserves special recognition for his time and efforts to untangle the mysteries of this model. I could not have completed the project without him.

through varying rates of mutation. At one extreme, the lack of mutation is like Charybdis. It is an absorbing boundary such that if a lineage ever evolves a complete lack of mutation, it will never escape its current state. Unfortunately for a lineage stuck in such a stasis, both the abiotic and biotic environment change over time. It is unlikely that a static lineage would be able to survive long in the real world, even if it were possible to evolve a complete lack of mutation. At the other extreme, a high mutation rate is a Scylla which will tear apart any genetic adaptations of the parent to produce an almost totally random offspring. In order to evolve adaptations, a lineage must happen to avoid both extremes.

1.1 Modifier Genes and the Reduction Principle

Any gene that affects the mutation rate of other genes is a modifier gene. That is, the mutation rate gene effects transmission and is not directly under selective pressure except to the extent that the genes it modifies are under selective pressure. According to the "evolutionary reduction principle" [12, 28, 1] an invading allele for the modifier gene should only succeed if it increases the fidelity of transmission of the selected gene. In other words, only lower mutation rates should invade.

The modifier gene itself is assumed to have perfect transmission and to produce linear variation in the transmission of the selected loci [1]. Furthermore, selection does not fluctuate in these models. If selection is allowed to fluxtuate or the population experiences random genetic drift, optimal mutation rates are non-zero [16, 15, 21]. The specific optimal mutation rate depends at least on parameters for the degree of dominance in the selected gene, the variance and autocorrelation in the environment, and the mean selective differences between the homozygotes [15]. These models generally do not focus on coevolution. While the environment may fluctuate, the selective effects of the environment are not determined by the population of organisms.

1.2 Genetic Algorithms

The evolution of mutation rates has also been studied from an engineering point of view, using genetic algorithms. The challenge is to find a mutation rate that will allow the population of solutions to converge on the best solution in the shortest amount of time [5, 4, 33, 9, 17, 14]. However, others point out that the optimal mutation rate depends on how a solution is encoded as a genome[39] and the choice of fitness function to evaluate a solution.[4] Bäck goes as far as to include a gene for mutation rate in each solution in the population.[4] This allows a "schedule" of mutation rates over time. Early on, mutation rates tend to be high. As the population approaches the optimal solution, mutation rates descend. This reduction in mutation rates over time is a common improvement technique in genetic algorithms.[5, 9, 14]

Traditional genetic algorithms employ an unrealistically static concept of fitness. The "problem" that the lineages are trying to solve in a genetic algorithm does not change over time. That is, the environment of the genetic algorithm is completely static. If we allow the organisms to interact with each other, then the organisms are suddenly immersed in a dynamic, biotic environment. With interactions, there is no longer any single, optimal solution.[1] Would the evolution of mutation rates change under coevolutionary pressures? The following model provides a tool for addressing that question.

2 The Model

Each organism in a population is explicitly represented in this configuration model.[10, 31, 20] Time is discrete and is organized by "generations." During each generation all of the organisms are allowed to interact in parallel with the other near by organisms in the environment. The form of these interactions is highly abstract, though not as simplified as other models of ecosystems[24, 23, 6], and have been inspired by the ECHO model.[18, 22, 19]

2.1 An Organism

An organism consists of three genes: *offense, defense,* and *mutation.* Both the *offense* and *defense* genes consist of eight bits. These are treated solely as bit patterns. They determine the results of an interaction with another organism through matching of the bit patterns as detailed below in Section 2.2.

The gene for *mutation* is an integer from 0 to 100. The *mutation* gene codes for the mutation rate of the *offense* and *defense* genes. Initially the model was implemented with the *mutation* gene coding for the mutation rates of all the genes, including itself [30]. This meant that if the *mutation* gene ever reached zero, all mutation and all evolution stopped. However, having this absorbing boundary at zero adds a bias to the average mutation rates over time. To avoid this bias, the *mutation* gene was altered at a fixed rate independent its integer value. The qualitative results were not affected by this change.

Mutation only occurs during reproduction when the parental genes are passed on and then modified in the offspring. Mutation is different for the integer and bit pattern genes. In the case of the integer gene, *mutation,* the gene of the offspring differs randomly from the parental gene along a Poisson distribution centered on the parent's gene with a λ of 5, as defined in (1).

$$\frac{(\lambda^{x+\lambda-parental\,gene})e^{-\lambda}}{(x+\lambda-parental\,gene)!} \tag{1}$$

[1] For a discussion of optimality see [26, 38]. For examples of coevolutionary models see [36, 18, 29].

Here λ was fixed at 5 when the *mutation* gene itself was being mutated, x is a potential offspring's gene, and *parental gene* is the gene of the parent that is being transmitted to its offspring. This gives a Poisson distribution around the parental gene. The offspring's gene is truncated at 0 and 100.

Mutation of a bit pattern gene depended on the value of the *mutation* gene of the parent, scaled down via equation (2).

$$\lambda = (mutation/50)^3 \tag{2}$$

The λ in equation (2), ranging from 0 to 8, was again used in a Poisson distribution. Here λ is the average number of bits that are flipped in an 8-bit gene during reproduction. This scaling was chosen to give both the full range of possibilities from 0 to 8 and to allow evolution to act on a fine grain of resolution of different mutation rates when the rates were low. When the parent's *mutation* gene is low, the offspring's gene is likely to be very close to the parent's gene. However, when the parent's *mutation* gene is high, the offspring's gene may vary widely from the parent's gene.

2.2 An Interaction

To model an ecosystem, interactions must at the least allow for the full scope of coevolutionary relationships, from mutualism to competition. More rigorously, an interaction can have a positive, negative, or neutral effect on either of the participants in the interaction. A model ecosystem should minimally be able to simulate all eight of these possible interactions. Comparison of bit patterns is one realization of this minimal requirement [18]. In this model, an interaction between two organisms, X and Y, consists of a comparison of X's *offense* gene bit pattern versus Y's *defense* gene bit pattern, and X's *defense* gene bit pattern versus Y's *offense* gene bit pattern.

A comparison is scored by the number of bits in the two genes that match. This is the reverse of the hamming distance. An organism with an *offense* gene of 11111111 would match 7 bits against an organism with a *defense* gene of 01111111 and so score 7 points. A full interaction is scored by subtracting the *defense-punishment* parameter times the comparison of X's *defense* gene and Y's *offense* gene from the *offense-reward* parameter times the comparison of X's *offense* gene and Y's *defense* gene, as per equation (3). An organism's *fitness* score is calculated as the sum total of all the interaction scores that organism collected during one generation.

$$\begin{aligned} \text{Score} \quad = \quad & \textit{offense-reward} \times (\text{match}(\text{offense}(\textit{me}),\ \text{defense}(\textit{opponent}))) - \\ & \textit{defense-punishment} \times (\text{match}(\text{offense}(\textit{opponent}),\ \text{defense}(\textit{me}))) \ . \end{aligned}$$

2.3 A Generation

The population of organisms is organized along a ring[6], so that every organism has exactly one neighbor on either side. A single generation consists of three steps:

1. All the organisms interact with each neighbor on either side. An organism's *fitness* score is the sum of two interaction scores.

2. The organisms with the lowest fitness scores are then killed off until only half of the population remains. This simulates the organisms producing twice the carrying capacity of the environment each generation.

3. The remaining half of the population is allowed to reproduce, according to their *mutation* genes, to fill the empty positions on the ring.

Reproduction is asexual. The newborn offspring receives the genes of a single parent, modified by that parent's *mutation* rate gene. Furthermore, reproduction is fair. Each surviving organism gets a turn to reproduce, by order in the ring. The position of the last organism to reproduce in a generation is saved so that reproduction in the next generation starts with the following organism on the ring. Only living organisms are allowed to reproduce.

2.4 Missing Complexities

The model includes many simplifying assumptions. Interactions are highly stylized. The fitness payoffs for the interactions are simple, discrete, linear functions. The environment has a fixed carrying capacity and the organism's reproduction rates are assumed to be both identical and high enough to fill the environment beyond the carrying capacity. There are no fatal mutations. A mutation merely changes an organisms interactions with its neighbors or changes its mutation rate. Organisms reproduce asexually. Finally, the organisms lack life histories. An organism that has survived many generations has no *a priori* competitive advantage over a new born organism.

2.5 An Elaboration and A Null Model

The model has a single gene, *mutation,* that governs the mutation rates of both the *offense* and the *defense* genes. An interesting elaboration involves splitting the single *mutation* gene into two mutation genes, *offense-mutation* and *defense-mutation.* The *offense-mutation* gene codes for the mutation rate of the *offense* gene, and the *defense-mutation* gene codes for the mutation rate of the *defense* gene. Both of the mutation gene are themselves mutated at a fixed rate ($\lambda = 5$).

A null model was also implemented to illuminate any hidden dynamics in the experimental model. The null model was identical to the experimental model except that when an organism interacted with its neighbors, the payoff was a

random number between 0 and 8, rather than a number generated through the matching of the *offense* and *defense* genes.

2.6 Running the Model

When the model is started, all of the genes in the initial population are generated randomly.[2] During the course of the simulation statistics are gathered each generation on the average fitness, minimum fitness, maximum fitness, average mutation rate, average age, and average number of offspring. Finally, at the end of a simulation, the resulting population is saved in a file for later examination.

3 The Results

In all of the following results the model was run with a population of 200 organisms for 2000 generations. Organisms interacted with their immediate neighbors on either side. The independent variables in the experiment were the *offense-reward* and *defense-punishment* parameters.

The initial model, with the single mutation rate gene, was run 50 times under each set of parameter conditions. The elaborated model with two mutation rate genes was run 16 time under each condition. Since the parameters had no affect on the dynamics of the null model, it was only run for one set of 16 trials.

3.1 Mutation Rates

The evolution of mutation rates differed dramatically depending on the settings of the *offense-reward* and *defense-punishment* parameters. The central results for the basic model are depicted in Fig. 1.

The mutation rates drop dramatically with the divergence of *offense-reward* and *defense-punishment*. When the average mutation rates are compared for the condition where *offense-reward* = *defense-punishment* = 1 against the condition where the *offense-reward* is twice the *defense-punishment*, the later condition results in significantly lower mutation rates (t=6.113, d.f.=76.708, $p < 0.001$).[3] Simlarly, *offense-reward* = 2, *defense-punishment* = 1 condition resulted in significantly higher mutation rates than when the *offense-reward* = 5 and the *defense-punishment* = 2 (t=6.0262, d.f.=74.2, $p < 0.001$). Furthermore, the *offense-reward* = *defense-punishment* = 1 populations have significantly larger mutation rates as compared to the populations in the null model (t=3.7452, d.f.=19.785, $p < 0.001$). However, once *offense-reward* grew to be 2.5 times

[2] Knuth's pseudo-random number generator was used, with a seed taken from the computer's clock.

[3] A Welch modified two-sample t-test was used with the alternative hypothesis that the second condition resulted in lower mutation rates. The variances for the two sets of populations were not assumed to be equal.

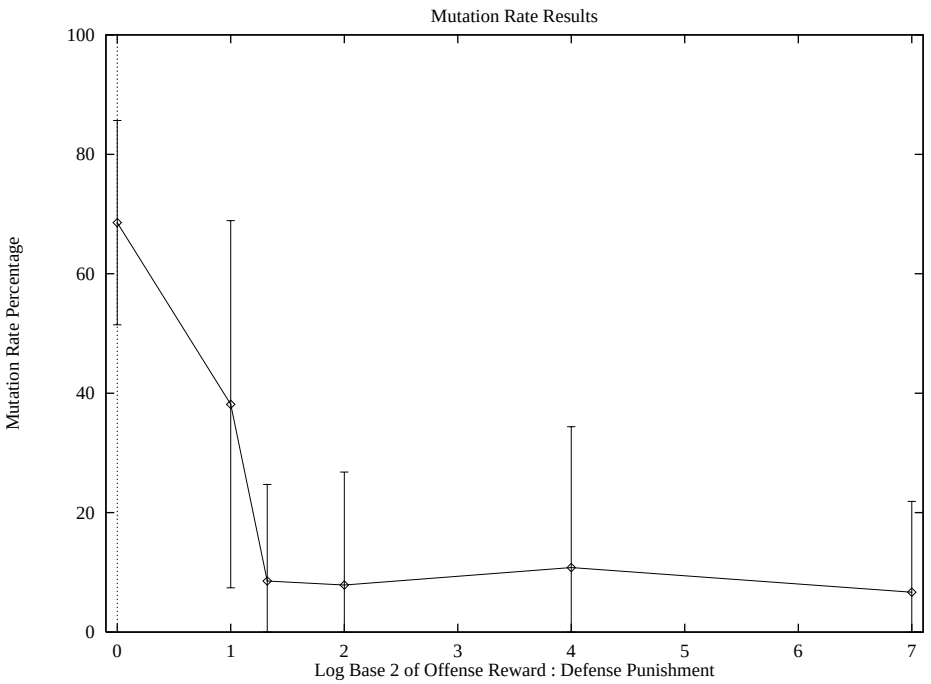

Figure 1: This is a $\log_2$-normal plot of the average mutation rates after 2000 times steps as a function of $(\frac{\textit{offense-reward}}{\textit{defense-punishment}})$. Each data point represents the average of 50 different populations with an error bar at one standard deviation.

defense-punishment, the experimental populations had significantly lower mutation rates than the null model populations (t=-5.3156, d.f.=19.258, $p < 0.001$).

The average mutation rates for the 50 trials in the *offense-reward* = *defense-punishment* = 1 condition cluster in an approximately normal curve. However, the relatively large variances observed in the data for the conditions where *offense-reward* $\gg$ *defense-punishment* is due to a small number of outliers with extremely high average mutation rates.

The model is symmetric around the two parameters. That is, if *defense-punishment* is twice the *offense-reward* then the average mutation rates drop by approximately the same amount. The only difference is that when *defense-punishment* is greater than *offense-reward*, there is pressure on the lineages to evolve bit pattern genes that mismatch the bit pattern genes of their neighbors. In any case, the average mutation rates decreases as *offense-reward* and *defense-punishment* diverge.

The elaborated model produced the same pattern of results. Fig. 2 shows

the average mutation rates for the *offense* and *defense* genes when their rates of mutation are independent.

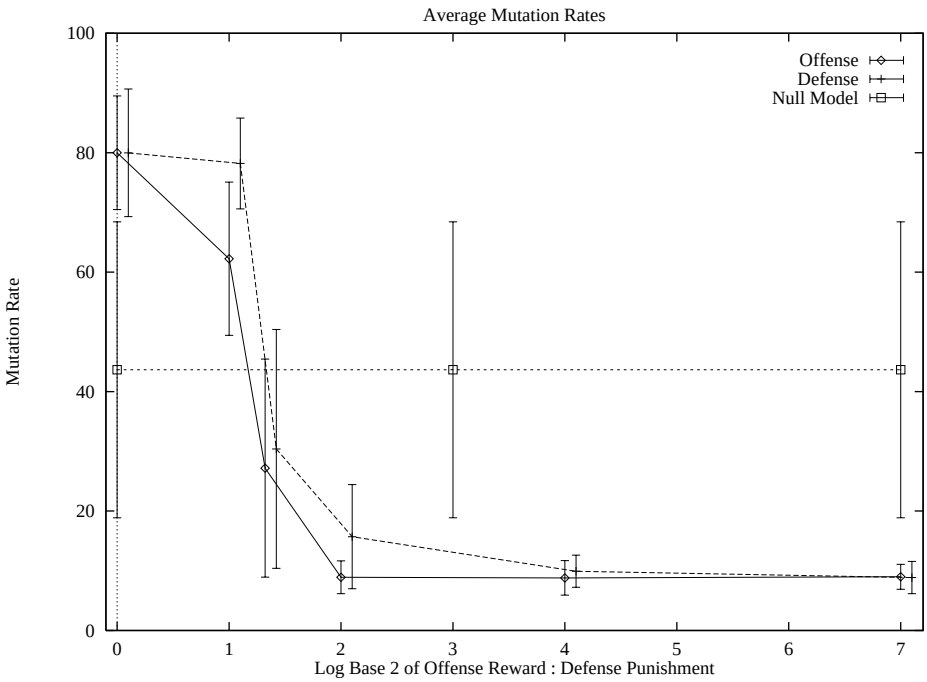

Figure 2: This is a $\log_2$-normal plot of the average mutation rates for the two mutation rate gene model and the null model. Except for the "Null Model" curve, each data point represents the average of 16 different populations with an error bar at one standard deviation. The "Offense" and "Defense" curves graph the average mutation rates of the two genes that independantly controlled the *offense* and *defense* genes respectively. The "Null Model" curve represents the average mutation rates when fitness scores for interactions were completely random. Although the null model is not affected by the *offense-reward* and *defense-punishment* parameters, the data for the 16 populations has been graphed at three different points for ease of comparison. The "Defense" mutation curve has been slightly offset to the right for readability.

4 Discussion

Game theory[37, 11] provides a useful tool for analyzing the evolution of mutation rates in the model. Why should an inequality in the *offense-reward* and *defense-punishment* parameters lead to the evolution of low mutation rates?

The players in these games are lineages. Each organism in the environment interacts with its two neighbors and the fitness value of those interactions form an explicit payoff matrix. A "move" in the game happens when an organism reproduces. If the offspring carries a mutation in the *offense* or *defense* genes, the new organism's interactions with the other lineages will be different from its parent's interactions. That is, a "move" in the game is a reproduction event with or without a mutation. The lineages play mixed strategies, sometimes mutating at reproduction, and sometimes not, depending on their mutation rates.

The structure of the payoff matrix in this coevolutionary game is determined by the *offense-reward* and *defense-punishment* parameters. To simplify the analysis it helps to assume that the lineages are coevolving over time. Of course, there is no guarantee that the next generation in a lineage will interact with the same "player." However, as a small number of lineages come to dominate the environment, the participants in the games will stabilize. The explanatory value of a game theoretic analysis appears to be robust to this simplification.

When *offense-reward* > *defense-punishment* and two lineages have tag genes that match well, the lineages engage in an iterated prisoner's dilemna (IPD) regarding the mutation of their *defense* genes.[3, 2, 34, 32, 7, 8, 35, 40] That is, mutual mutation will result in lower interaction payoffs than mutual lack of mutation. Yet, if one lineage mutates and the other doesn't, then the mutant lineage will gain a better payoff than mutual lack of mutation while the non-mutant lineage will score worse than mutual mutation. Table 1 illustrates one example of an IPD payoff matrix in the model.

Table 1: *Offense-reward* = 2 and *defense-punishment* = 1. When lineages 1 and 2 perfectly match, a mutation in the *defense* gene produces these scores. This is a Prisoner's Dilemma game.

		Lineage 2
Lineage 1	No Mutation	Mutation
No Mutation	8, 8	6, 9
Mutation	9, 6	7, 7

Under the same conditions, a mutation in the *offense* genes will always lower a lineage's score. The combination of this pressure on the *offense* genes and the IPD for the *defense* genes leads to pressure towards lowering the mutation rates of the lineages. The difference between mutual mutation and mutual lack of mutation increases as *offense-reward* and *defense-punishment* diverge. Thus, as *offense-reward* grows large relative to *defense-punishment* the pressure to lower mutation rates increases and the lineages can be seen to respond accordingly.

The fact that the *mutation* gene evolves under the pressures on both the *offense* and *defense* gene suggested an elaboration to the model. Instead of one gene regulating mutation in both the bit pattern genes, a mutation gene

was created for each bit pattern gene. The fact that the the *defense* gene's mutation rate appears higher than the mutation rate for the *offense* gene would be expected under conditions where the lineages match each other well, so that a mutation in the *defense* gene would tend to lower a lineage's score. In contrast, when lineages are well matched, a mutation in a lineage's *offense* gene would tend to lower the lineage's score.

4.1 High Mutation Rates

The game theoretic analysis fails to illuminate why the symmetric case, where *offense-reward* = *defense-punishment*, results in high mutation rates. In the symmetric case, when organisms match their neighbors, the chance of a beneficial mutation to the *defense* gene is matched by an equal chance for a deleterious mutation to the *offense* gene. This balance does not change when we relax the assumptions that organisms match their neighbors. So it would seem the mutation rates of the populations should do a random walk. They should turn out exactly like the null model populations. So why are the mutation rates in the experimental populations significantly higher than the mutation rates in the null model populations?

We hypothesized that in the symmetric case, a high mutation rate is advantageous for a lineage in an environment where the average mutation rate of the population is low. Thus, when the random walk of the population wanders into the region of low mutation rates, high mutation rates become selectively advantageous and so the population average bounces back up. Thus the population should do a random walk in the region of high mutation rates with a reflecting boundary at 100 and some sort of soft boundary at low mutation rates. We then designed a series of ecological experiments to test this hypothesis.

4.1.1 Ecological Experiments

To test whether some mutation rates are selected over others, we altered the model so that the mutation rates no longer evolve themselves, but remain fixed. We then inoculated the model with two interleaved populations of organisms. One population had a "high" mutation rate and the other had a "low" mutation rate. The *offense* and *defense* genes were set randomly for each individual in both populations. We then ran the model until one of the two populations reached fixation. If there was no selective advantage for one of the mutation rates, then each sub-population should reach fixation in 50% of the trails.

4.1.2 Ecological Results

When we ran a population of organisms with a mutation rate of 30 with a population of organisms with a mutation rate of 10, 199 out of 200 trails resulted in the mutation = 30 population growing to fixation. In contrast, when a

population of mutation = 80 organisms was mixed with a population of mutation = 60 organisms, the mutation = 80 population grew to fixation in only 96 of 200 trials. This was not a statistically significant deviation from the expected 50% of the trials.

4.1.3 Ecological Conclusions

The ecological experiments exactly matched the predictions of our hypothesis. The population seems to be following a random walk between 100 and some elevated lower bound. It remains to be elucidated why a high mutation rate is selectively advantageous for a lineage in a low mutation rate biotic environment. However, we believe the advantage lies in the interaction of a high variance in a lineage with non-linearities in the interactions and spatial structure of the environment. A lineage with a high mutation rate will produce offspring with widely divergent *offense* and *defense* genes. Whereas a lineage with a low mutation rate will produce offspring similar to the parent. This is the equivalent to "putting all of one's eggs in one basket." When the general population wanders into a low average mutation rate, regularities quickly spread through the population. Due to these regularities, it is likely that if any of the low mutators score poorly in their interactions, all of the low mutating offspring will be culled from the population and the lineage will go extinct. In contrast, a lineage with a higher mutation rate is more likely to produce a collection of offspring spanning the range from low to high scoring individuals. Thus, the high mutating lineage is liable to lose a few individuals every generation, but less likely to lose all of its constituent individuals. In this way, the lineage with the higher mutation rate survives to pass on its gene for a high mutation rate. This hypothesis may be supported by Kephart's finding that the addition of spatial structure to a predator-prey model can create large undampted oscillations in the population dynamics [25].

4.2 Implications for Biology

The results of the model follow the evolutionary pressures that are applied to the lineages of organisms in the model. But what has this got to do with biology? The dramatic distinction in the model is between the case where *offense-reward* = *defense-punishment* and the cases where *offense-reward* $\neq$ *defense-punishment*. This distinction roughly corresponds to zero-sum games versus non-zero-sum, or mixed-motive games. In more biological terms, the first case corresponds to interactions where the consequences are roughly the same for each participant. Intraspecific competition might be a common example. Under this form of interaction, the model predicts that the participants will evolve relatively high mutation rates. The second case corresponds to an interaction where the consequences are significantly different for the participants. Here the model predicts that the participant lineages will settle down into a mutualis-

tic association with relatively low mutation rates. The dinner principle may be a familiar example. The cost/benefit analysis is dramatically different for predator and prey, as long as the predator is not on the brink of starvation. The preceding results suggest that such an imbalance in consequences should dampen the pressures in a predator-prey genetic arms races.

Mutation should be understood under this theory as phenotypic variation. Any heritable, causal factor that affects the similarity of a phenotypic trait between the parent and offspring should be considered as a possible adaptation to the coevolutionary pressures on that trait. In a zero-sum game, the evolution of a relatively high "mutation rate" might be realized through the evolution of a greater sensitivity of the relevant trait to environmental conditions during development. This would produce a lower correlation between the trait of the parent and the trait of the offspring. On the other hand, the evolution of a high phenotypic mutation rate might be realized as a "hot spot" in the genes that influence the development of the relevant traits.

There is evidence for differences amongst mutation rates in active genes of mammalian genomes.[41] However, the causal mechanism for these differences is not yet clear. Filipski suggests that some regions of the genome are replicated by the relatively error-prone β DNA polymerase while most others are replicated by the relatively high fidelity α DNA polymerase [13]. In contrast, Wolfe *et al.* observe that genes are replicated at different times during the S phase of the cell cycle. They suggest that the differences in mutation rate are caused by differences in dNTP pool conditions at the time that the genes are replicated [41]. These are by no means the only possible explanations for differences in mutation rates.

The results suggest that different forms of interactions put different evolutionary pressures on the mutation rates of the participants. Expressed genes should have evolved different intrinsic mutation rates as adaptations to their environment of interactions. We should expect to find structures in the cell and genome for selectively allowing some genes to mutate faster than others. This may mean actively boosting the mutation rates. Conversely, it may simply mean that some genes are the subject of better replication and error correction mechanisms than other genes.

References

[1] L. Altenberg and M. W. Feldman. Selection, generalized transmission and the evolution of modifier genes. i. the reduction principle. *Genetics*, 117:559–572, 1987.

[2] R. Axelrod. *The Evolution of Cooperation.* Basic Books, New York, NY, 1984.

[3] R. Axelrod and W. D. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.

[4] T. Bäck. The interaction of mutation rate, selection, and self-adaptation within a genetic algorithm. In R. Männer and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 85–94. Elsevier Science Pubishers, Amsterdam, 1992.

[5] T. Bäck. Optimal mutation rates in genetic search. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 2–8, San Mateo, CA, 1993. Morgan Kaufmann Publishers.

[6] P. Bak and K. Sneppen. Punctuated equilibrium and criticality in a simple model of evolution. *Physical Review Letters*, 71(24):4083–4086, 1993.

[7] R. Boyd. Mistakes allow evolutionary stability in the repeated prisoner's dilemma game. *Journal of Theoretical Biology*, 136:47–56, 1989.

[8] R. Boyd and J. P. Lorberbaum. No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game. *Nature*, 327:58–59, 1987.

[9] M. F. Bramlette. Initialization, mutation and selection methods in genetic algorithms for function optimization. In K. Belew and B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, page 100=107, San Mateo, CA, 1991. Morgan Kaufmann Publishers.

[10] H. Caswell and A. M. John. From the individual to the population in demographic models. In D. Deangelis and L. Gross, editors, *Individual-based Models and Approaches in Ecology*, pages 36–61, New York, 1992. Chapman and Hill.

[11] A. M. Colman. *Game Theory and Experimental Games.* Pergamon Press Inc., Elmsford, NY, 1982.

[12] M. W. Feldman and U. Liberman. An evolutionary reduction principle for genetic modifiers. *Proceedings of the National Academy of Science, USA*, 83:4824–4827, 1986.

[13] J. Filipski. Correlation between molecular clock ticking, codon usage, fidelity of dna repair, chromosome banding and chromatin compactness in germline cells. *FEBS Letters*, 217:184–186, 1987.

[14] Terence C. Fogarty. Varying the probability of mutation in the genetic algorithm. In *Proceedings of the Third International Conference on Genetic Algorithms*, pages 104–109, Palo Alto, CA, 1989. Morgan Kaufmann.

[15] J. H. Gillespie. Evolution of a mutation rate at a heterotic locus. *Proceedings of the National Academy of Science, USA*, 78:2452–2454, 1981.

[16] J. H. Gillespie. Mutation modification in a random environment. *Evolution*, 35:468–476, 1981.

[17] J. Hesser and R. Männer. Towards an optimal mutation probability for genetic algorithms. In H.-P. Schwefel and R. Männer, editors, *Parallel Problem Solving from Nature*, pages 23–32. Springer-Verlag, New York, 1990.

[18] J. H. Holland. *Adaptation in Natural and Artificial Systems.* MIT Press, Cambridge, MA, 1992.

[19] J. H. Holland. Echoing emergence: Objectives, rough definitions, and speculations for echo-class models. Technical Report 93-04-023, Santa Fe Institute, 1993.

[20] M. Huston, D. DeAngelis, and W. Post. New computer models unify ecological theory. *Bioscience*, 38(10):682–691, 1988.

[21] K. Ishii, H. Matsuda, Y. Iwasa, and A. Sasaki. Evolutionary stable mutation rate in a periodically changing environment. *Genetics*, 121:163–174, 1989.

[22] T. Jones and S. Forrest. An introduction to sfi echo. Technical Report 93-12-074, The Santa Fe Institute, 1993.

[23] S. A. Kauffman. *The Origins of Order.* Oxford University Press, Oxford, UK, 1993.

[24] S.A. Kauffman and S. Johnsen. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of Theoretical Biology*, 149:467–505, 1991.

[25] J. O. Kephart. How topology affects population dynamics. In C. G. Langton, editor, *Artificial Life III*, Reading, MA, 1994. Addison-Wesley.

[26] R. C. Lewontin. Fitness, survival, and optimality. In D. J. Horn, G. R. Stairs, and R. D. Mitchell, editors, *Analysis of Ecological Systems*, pages 3–22. Ohio State University Press, Columbus, OH, 1979.

[27] R. C. Lewontin, October 1993. Personal communication.

[28] U. Liberman and M. W. Feldman. Modifiers of mutation rate: A general reduction principle. *Theoretical Population Biology*, 30:125–142, 1986.

[29] K. Lindgren and M. G. Nordahl. Cooperation and community structure in artificial ecosystems. *Artificial Life*, 1:15–38, 1994.

[30] C. C. Maley. The coevolution of mutation rates. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón, editors, *Advances in Artificial Life*, pages 234–245, Berlin, 1995. Springer-Verlag.

[31] C. C. Maley and H. Caswell. Implementing i-state configuration models for population dynamics: An object-oriented programming approach. *Ecological Modelling*, 68:75–89, 1993.

[32] M. Milinski. Cooperation wins and stays. *Nature*, 364:12–13, 1993.

[33] H. Mühlenbein. How genetic algorithms really work i: Mutation and hill-climbing. In R. Männer and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 15–26. Elsevier Science Pubishers, Amsterdam, 1992.

[34] M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, 364:56–58, 1993.

[35] M. A. Nowak and K. Sigmund. Tit for tat in heterogeneous populations. *Nature*, 355:250–253, 1992.

[36] T. S. Ray. An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 371–408, Reading, MA, 1992. Addison-Wesley.

[37] J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK, 1982.

[38] J. Maynard Smith. Optimization theory in evolution. In E. Sober, editor, *Conceptual Issues in Evolutionary Biology*, pages 91–118. MIT Press, Cambridge, MA, 1994.

[39] D. M. Tate and A. E. Smith. Expected allele coverage and the role of mutation in genetic algorithms. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 31–37, San Mateo, CA, 1993. Morgan Kaufmann Publishers.

[40] G. S. Wilkinson. Reciprocal food sharing in the vampire bat. *Nature*, 308:181–184, 1984.

[41] K. H. Wolfe, P. M. Sharp, and W.-H. Li. Mutation rates differ among regions of the mammalian genome. *Nature*, 337:283–285, 1989.