The Coevolution of Mutation Rates

C.C. Maley*
NE43-803
Massachusetts Institute of Technology
Cambridge, MA 02139
617-253-6567
email: cmaley@ai.mit.edu

30 August 1994

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 genes for longevity and mutation rate in the individuals. This made it possible for a lineage to evolve to be immortal. It also allowed the evolution of no mutation or extremely high mutation rates. 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

> ...we badly need a comparative biology. So far, we have been able to study only one evolving system and we cannot wait for interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we will have to look at artificial ones.
>
> - J. Maynard Smith[Smith, 1992]

We only have a single example of the evolution of life. With a sample of one, it is difficult, if not impossible, to distinguish between historical and systematic constraints on the evolution of life. Artificial Life represents an emerging complementary approach to understanding biology.[Taylor and Jefferson, 1994, Langton, 1989, Langton *et al.*, 1992,

*This work was supported in part by the generosity of the MIT AI Lab. I am greatful to M. Donoghue, P. Goss, K. Rice and L. King for their comments and support.

Langton, 1994] It is an attempt to expand the size of the sample of living systems we have to study. Artificial Life also commonly includes the synthesis of artificial systems designed to address specific questions in biology that are intractable under more traditional laboratory techniques. The principle weakness in the Artificial Life approach is the sometimes tenuous connection between the artificial systems and biological systems. This weakness stands in contrast to the high degree of experimental control over the artificial systems. An artificial evolving system can be restarted under identical initial conditions, excepting the random number generator seed, slightly different initial conditions, or different system parameters. Furthermore, the experimenter has access to all of the information in the system, and so can avoid many of the difficulties inherent in ecological research. The fact that Artificial Life research depends on models of life means that the conclusions must be predicated on the acceptance of the model as a reasonable description of life. Even given such acceptance of the realism of the model, the results generally can only be qualitative rather than quantitative. Given these limitations, Artificial Life helps us to understand both the constraints and possibilities of life.[Maley, 1994, Valentine and Walker, 1986, Kauffman, 1993, Kauffman and Johnsen, 1991, Boerlijst and Hogeweg, 1991, Collins and Jefferson, 1992]

Lewontin suggested that when we think about what life might be, rather than what we see before us, there are two dimensions that can be considered: mortality and heredity.[Lewontin, 1993] 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 those questions. Neither mortality nor heredity need be dichotomous qualities. Mortality might range over varying degrees of longevity. Heredity can be addressed through varying rates of mutation. At the extremes there is either no mutation and so perfect replication, or the mutation rate is so high that the offspring's genetic information is near random.

1.1 A Naive Model

Organisms described only by genes for longevity and mutation rates would be difficult to model on a computer. Assuming a constant rate of reproduction, the computer's memory would fill up with near immortal creatures. Once a lineage struck upon a low mutation rate combined with a high longevity, this growth would become exponential. Of course, in the real world, such growth would be limited by the available resources. The computer's memory in this case is an appropriate analogy for just such a limited resource. So if the model contained a fixed maximum carrying capacity, then an organism could only successfully reproduce if another died. Death, in this naive model, only comes from the longevity gene. Once an immortal organism evolved it would be there for good. Given non-zero mutation rates and a reasonable sized population, immortal organisms would continue to arise. In fact, the only selective pressure in this naive model is a pressure against mortality. Eventually, the entire carrying capacity of the environment would be filled with immortal organisms and evolution would come to a halt. An implementation of this naive model has confirmed this prediction.

1.2 Genetic Algorithms

Research into genetic algorithms has addressed the issue of mutation rates. Genetic algorithms are generally used in problems of optimization. See [Goldberg, 1989] for a good introduction to genetic algorithms. A genetic algorithm employs an evolving population of potential solutions to an artificial problem in order to find a good solution to the problem. Each individual in the population is evaluated as a solution to the problem via the application of a "fitness function" to the individual. Then some scheme is used to remove the poor solutions from the population and replace them with variants on the better solutions in the population. These new variants are usually generated by minor mutations in the successful solutions and through recombination of pieces of the successful solutions. Thus a highly stylized form of evolution is employed to find a good, and sometimes optimal, solution to the problem at hand. The individual solutions in the population do not directly interact with each other. The only form of interaction is indirect, through the fitness function. In general, a solution must be better than other solutions in the population in order to survive and reproduce.

A significant amount of work on genetic algorithms has focused on finding a mutation rate that will allow the population of solutions to converge on the best solution in the shortest amount of time.[Bäck, 1993, Bäck, 1992, Mühlenbein, 1992, Bramlette, 1991, Hesser and Männer, 1990, Fogarty, 1989] However, others point out that the optimal mutation rate depends on how a solution is encoded as a genome[Tate and Smith, 1993] and the choice of fitness function to evaluate a solution.[Bäck, 1992] Bäck goes as far as to include a gene for mutation rate in each solution in the population.[Bäck, 1992] 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. Scaling down mutation rates over time is a common improvement technique in genetic algorithms.[Bäck, 1993, Bramlette, 1991, Fogarty, 1989]

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. On the other hand, the naive model misses the point of environmental carrying capacity in biology. Rather than reproduction filling a viable hole in the population, the biological world follows the more Malthusian dynamic of explosive reproduction curtailed by starvation, disease, and conflict. If we allow the organisms to interact with each other, then we open up the possibility for competition, exploitation, and even cooperation. With interactions, there is no longer any single, optimal solution.[1] Would the evolution of mutation rates and longevity 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.[Caswell and John, 1992, Maley and Caswell, 1993, Huston *et al.*, 1988] Time is discrete and is orga-

[1] For a discussion of optimality see [Lewontin, 1979, Smith, 1994]. For examples of coevolutionary models see [Ray, 1992, Holland, 1992, Lindgren and Nordahl, 1994].

nized by "generations." During each generation all of the organisms are allowed to interact in parallel with the other organisms near by in the environment. The form of these interactions is highly abstract, though not as simplified as other models of ecosystems[Kauffman and Johnsen, 1991, Kauffman, 1993, Bak and Sneppen, 1993], and have been inspired by the ECHO model.[Holland, 1992, Jones and Forrest, 1993, Holland, 1993]

2.1 An Organism

An organism consists of four genes: *offense, defense, longevity* 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 genes for *longevity* and *mutation* are integers over the range 0 to 100, inclusive. The gene for *longevity* codes the percent likelihood that the organism will die in a generation. Thus an organism with a 0 for a *longevity* gene would not survive past one generation. An organism with 100 for a *longevity* gene would be immortal.[2]

The *mutation* gene codes for the mutation rate of all the genes in the organism. 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 genes, *longevity* and *mutation* itself, the genes of the offspring differ randomly from the parental genes along a Poisson distribution centered on the parent's gene, as defined in Equation 1.

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

In Equation 1, λ is the mutation gene, 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. When the mutation rate is low, the offspring's gene is likely to be very close to the parent's gene. However, when mutation rate is high, the offspring's gene may vary widely from the parent's gene. Figure 1 depicts three different distributions. When the mutation rate gene is zero, the offspring's genes are identical to the parent's genes.

Mutation of a bit pattern gene is scaled down via Equation 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. Figure 2 shows the scaling curve.

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

[2]Such an "immortal" organism could still fall prey to competition, described in Sections 2.2 and 2.3.

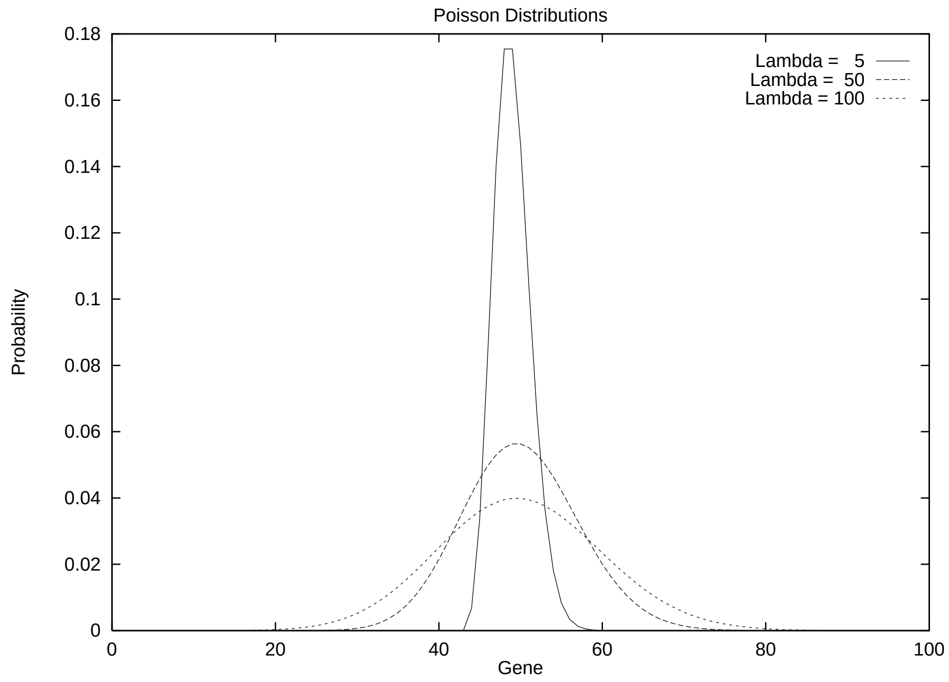

Figure 1: For a gene with a value of 50, the genes of its offspring are distributed around it depending on the λ of the poisson distribution. This figure shows poisson distributions for $\lambda = 5$, 50 and 100.

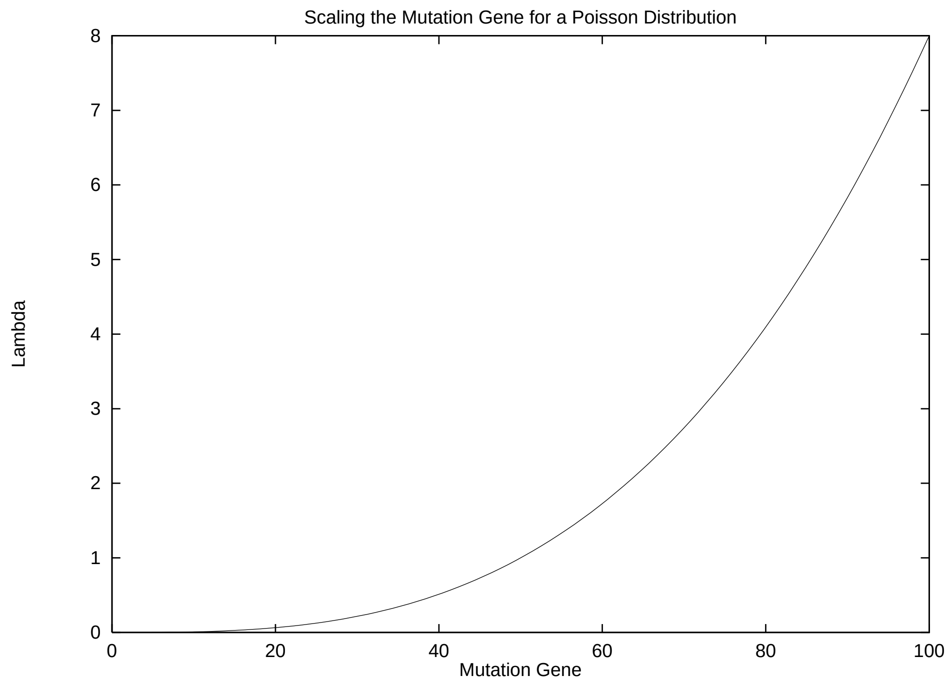

Figure 2: The mutation gene is scaled from a range of 0 to 100 to a range of 0 to 8 by this curve. This scaling has the advantage of a fine grain resolution of different mutation rates when the mutation gene is low.

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.[Holland, 1992] 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 subtracting four from the number of bits in the two genes that match. This is the reverse of the hamming distance. Subtracting 4 gives a score spread from -4 to 4. An organism with an *offense* gene of 11111111 would match 7 bits against an organism with a *defense* gene of 01111111 and so score 3 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. An organism's *fitness* score is calculated as the sum total of all the interaction scores that organism collected during one generation.

2.3 A Generation

The population of organisms is organized along a ring[Bak and Sneppen, 1993], so that every organism has exactly one neighbor on either side. A single generation consists of four 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) A random number from 0 to 99 is generated for each organism in the population. If that number exceeds the *longevity* gene number for an organism, that organism is considered dead. (3) 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. (4) 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 new born offspring receive 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. Organisms reproduce asexually. Perhaps most limiting is the absence of any life histories for the organisms. This means that an organism that has survived many generations has no *a priori* competitive advantage over a new born organism. It also means that longevity, and its complement, senescence, cannot be examined in the context of age dependent genetic effects.

2.5 Running the Model

When the model is started, parameters are input for population size, the number of generations to simulate, the number of neighbors on each side with which an organism will interact, the *offense-reward* and *defense-punishment* multipliers. All of the genes in the initial population are generated randomly.[3] During the course of the simulation statistics are gathered each generation on the average fitness, minimum fitness, maximum fitness, average longevity, 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 and an interaction radius of one. The independent variables in the experiment were the *offense-reward* and *defense-punishment* parameters.

3.1 Longevity

There are two separate selective pressures in the model, senescence and competition. The first depends solely on the longevity genes. It is perhaps not surprising then that the longevity genes of the organisms approach 100, or immortality. However, the population generally does not reach fixation on the maximum. To test if these results were significant, a null model was implemented. All interactions were removed from the null model and instead of killing off the organisms with the lowest fitness scores, half of the population was killed at random. The results showed that longevity of the population does not fixate at 100 because mutation continually introduces organisms with longevities below 100. There was no statistically significant difference between the evolution of longevity in the coevolutionary model and the null model.

3.2 Mutation Rates

The evolution of mutation rates differed dramatically depending on the settings of the *offense-reward* and *defense-punishment* parameters. The central results are detailed in Table 1. 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 = -4.0175, df = 9.075, p-value = 0.0015).[4] As can be seen, the mutation rates often dropped

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

[4] A Welch modified two-sample t-test was used with the alternative hypothesis that the second condition resulted in lower mutation rates.

trial	*O=1, D=1*	*O=2, D=1*	*O=1, D=2*	*O=100, D=1*	*O=1, D=0*	*O=0, D=1*
1	53.525	61.170	0	0	0	0
2	47.575	0	0	0	0	0
3	51.310	0	50.275	0	0	0
4	67.455	31.980	0	0	0	0
5	66.485	0	22.805	0	0	0
6	71.195	33.480	55.465	0	0	0
7	67.470	34.890	24.380	0	0	0
8	66.725	45.505	0	0	0	0

Table 1: These are the average mutation rates for a population of 200 after 2000 generations. In the headings of the columns, "O" and "D" refer to the *offense-reward* and *defense-punishment* multiplier that were used for that trial.

to fixation at a zero mutation rate. The same is true if the *defense-punishment* is twice the *offense-reward*. The effect becomes even more dramatic as one of the two parameters is weighted more than the other. In all the extreme cases, where the *offense-reward* was weighted 100 times as much as the *defense-punishment* or when one of the two parameters was 0, mutation rates always dropped to 0.[5]

4 Discussion

There seems to be nothing of interest in the evolution of longevity in this model. The results of the null model clearly argue that the only reason that longevity does not reach its maximum is that the mutation rates of the organisms constantly introduce a certain amount of genetic load. The central issue in the evolution of longevity is the attempt to distinguish between two theories for the phenomenon of senescence.[Partridge and Barton, 1993, Cristofalo, 1988, Bell, 1984, Sonneborn, 1978, Cutler, 1978] Senescence may be caused by a life history trade off. Genes might persist that enhanced survival and reproduction in an organism's early life at the expense of impairment in later life. The competing theory asserts that senescence is caused by deleterious mutations that are only expressed in old age, and so are not easily weeded out of the population. That is, senescence is just a form of genetic load. Clearly, without an implementation of life histories in the organisms, it is impossible to distinguish between these theories.

4.1 Game Theory Analysis of Mutation Rates

Game theory[Smith, 1982, Colman, 1982] 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. The outcome of this interaction is fixed over the life time of the pair. As long as they

[5] In the two trials when one of the two parameters was zero, the population reached fixation at a zero mutation rate in less than 150 generations.

	Lineage 2	
Lineage 1	No Mutation	Mutation
No Mutation	0, 0	-1, 1
Mutation	1, -1	0, 0

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

both survive, the interaction will not change. There is no selection of alternate strategies within the life time of an organism. However, when an organism reproduces, its *offense* and *defense* genes may mutate. A mutation in either of these genes changes the way in which the new organism interacts with other organisms. So a lineage may fix on a certain form of interaction by evolving a zero mutation rate, or it may continually change its interactions by maintaining a non-zero mutation rate.

We can analyze the payoff matrices for a lineage mutating or not mutating in interaction with other lineages. To simplify the analysis, it helps to consider the payoffs around the extremes of interactions, with either a perfect match or mismatch between the *offense* and *defense* genes. This is not an unreasonable approach, because these extremes are often observed to be equilibria in the model. It also 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 predictive value of a game theoretic analysis appears to be robust to these simplifications.

In all of the following matrices, it is helpful to remember that a perfect match scores positive (e.g., 4) fitness points and a perfect mismatch scores negative (e.g., -4) points. The maximum payoff for an interaction, if *offense-reward* = *defense-punishment* = 1, would be $4-(-4) = 8$, when the organism's *offense* gene perfectly matches its neighbor's *defense* gene, and the organism's *defense* gene perfectly mismatches the neighbor's *offense* gene. Let us begin with those parameter settings. If we consider the case where the two players perfectly match each other, in Table 2, then a mutation in an *offense* gene will always lower that player's payoff and a mutation in a *defense* gene will always raise that player's payoff. The interactions are completely symmetric. The first number in each paired entry of a matrix is the score for player 1 and the second number is the score for player 2.

At the opposite extreme, the two players perfectly mismatch. Then a mutation in a player's *defense* gene will always lower its score. If we consider the mutations in the *offense* genes, then the exact same payoff matrix applies. A mutation in the player's *offense* gene will always raise its score. Under this condition, the benefits of mutation in one gene is always matched by a penalty for mutation in the other gene. Thus, when *offense-reward* = *defense-punishment* = 1 there is no selective pressure on mutation rates. The result is a population of lineages doing a sort of slow random walk in their mutation genes. The larger the population, the slower this walk.

Now consider what happens if *offense-reward* = 2 and *defense-punishment* = 1. If the players initially have a perfect match then the payoff matrix for mutating the *defense* gene is

Lineage 1	No Mutation	Mutation
	Lineage 2	
No Mutation	4, 4	2, 5
Mutation	5, 2	3, 3

Table 3: *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 1	No Mutation	Mutation
	Lineage 2	
No Mutation	-4, -4	-5, -2
Mutation	-2, -5	-3, -3

Table 4: *Offense-reward* = 2 and *defense-punishment* = 1. When lineages 1 and 2 perfectly mismatch, a mutation in the *offense* gene produces these scores.

shown in Table 3. This is a Prisoner's Dilemma payoff matrix.[Axelrod and Hamilton, 1981, Axelrod, 1984, Nowak and Sigmund, 1993, Milinski, 1993, Boyd, 1989, Boyd and Lorberbaum, 1987, Nowak and Sigmund, 1992, Wilkinson, 1984] A lineage scores better by mutating, or "defecting" in the Prisoner's Dilemma terminology, regardless of what its opponent does. However, mutual cooperation scores slightly better for both players than mutual defection. The table for a mutation in the *offense* gene reverses each pair of scores in Table 3. Thus there is pressure against mutation in the *offense* gene and mutual cooperation still scores better than mutual defection. Thus, on balance, there is pressure towards cooperation, or zero mutation rates. A pair of well matched organisms with low mutation rates will thrive in a sea of organisms with high mutation rates. An historical accident could lead to the spread of cooperation and even fixation on a zero mutation rate.

If we consider the opposite extreme where the players are perfectly mismatched, as in Table 4, mutual defection scores better than mutual cooperation. As long as the lineages generally mismatch they naturally evolve towards relatively high mutation rates. Once the lineages tend to match each other, the prisoner's dilemma payoff matrix emerges with its reward for mutual cooperation. However, the difference in payoff between mutual cooperation and mutual defection is minimal. It still requires an adjacent cooperating pair of lineages to let cooperation spread. This explains why some of the trials fixate at zero mutation rates and others do not. Those that do not fixate at a zero mutation rate still have lower average mutation rates than the symmetrically weighted trials discussed above.

Reversing the parameters does not change the analysis significantly. If *offense-reward* = 1 and *defense-punishment* = 2, mutual benefit can be gained when the players mismatch. The payoff matrix in this case is identical to Table 3. Similarly, in the case that the genes match, the payoff matrix is identical to Table 4. Reversing the two parameters results in an opposite but equivalent set of pressures. So these trials ought to behave similarly to the trials

	Lineage 2	
Lineage 1	No Mutation	Mutation
No Mutation	396, 396	296, 397
Mutation	397, 296	297, 297

Table 5: *Offense-reward* = 100 and *defense-punishment* = 1. When lineages 1 and 2 perfectly match, a mutation in the *defense* gene produces these scores.

	Lineage 2	
Lineage 1	No Mutation	Mutation
No Mutation	-396, -396	-397, -296
Mutation	-296, -397	-297, -297

Table 6: *Offense-reward* = 100 and *defense-punishment* = 1. When lineages 1 and 2 perfectly mismatch, a mutation in the *offense* gene produces these scores.

where *offense-reward* = 2 and *defense-punishment* = 1 except that cooperating organisms must mismatch rather than match. This is, in fact, what is observed in the populations that fixated on zero mutation rates.

When we increase the differences between the two parameters, the interactions become more lopsided, but the prisoner's dilemma remains. Table 5 shows the payoff matrix for *offense-reward* = 100 and *defense-punishment* = 1. It still benefits a lineage to mutate its *defense* gene. However, the advantage of such a mutation is only a single point, whereas the disadvantage if both lineages mutate is 99 points. The analysis of the mismatching case, in Table 6, is also just a caricature of the Table 4.

When the differences in the parameters are exaggerated like this, there is even greater evolutionary pressure for the lineages to reduce their mismatches as well as to maintain cooperation. Correspondingly, all of the trials run under these conditions fixated on zero mutation rates.

The limit of this trend towards asymmetric interactions is to set one of the parameters to zero. When *offense-reward* = 1 and *defense-punishment* = 0, there is an overall bias towards cooperation. The payoff matrix of Table 7 makes it clear that there is no benefit to mutating the *defense* gene. When the pairs of payoffs are reversed in the table, it becomes clear that there is a pressure to prevent mutations in the *offense* gene. In total, mutual cooperation scores higher than mutual defection when the lineages perfectly match. All of the experimental runs under these conditions fixated at zero mutation rates.

When the lineages mismatch, as in Table 8, there is pressure on the lineages to mutate to reduce mismatches and then to stop mutating when a match is achieved. Under this analysis, it is not surprising that occasionally a pair of species fixated at a zero mutation rate when they had a good but not perfect match. As above, the case where *offense-reward* = 0 and *defense-punishment* = 1 follows an exactly equivalent analysis.

Lineage 1	*Lineage 2* No Mutation	Mutation
No Mutation	4, 4	3, 4
Mutation	4, 3	3, 3

Table 7: *Offense-reward* = 1 and *defense-punishment* = 0. When lineages 1 and 2 perfectly match, a mutation in the *defense* gene produces these scores.

Lineage 1	*Lineage 2* No Mutation	Mutation
No Mutation	-4, -4	-4, -3
Mutation	-3, -4	-3, -3

Table 8: *Offense-reward* = 1 and *defense-punishment* = 0. When lineages 1 and 2 perfectly mismatch, a mutation in the *offense* gene produces these scores.

4.2 Implications for Biology

It is now clear that 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* ≠ *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 mutualistic 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 put the breaks on predator-prey genetic arms races.

Mutation should be understood under this theory as phenotypic mutation. Any heritable, causal factor that affects the rate of change in a phenotypic trait 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.

These results suggest that different forms of interactions put different evolutionary pressures on the mutation rates of the participants. This leads to the prediction that different 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. A possible example might involve the selective methylization of cytosine into 5-methylcytosine in a gene. Conversely, it may simply mean that some genes are the subject of better error correction than others.

References

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

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

[Bäck, 1992] 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.

[Bäck, 1993] 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.

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

[Bell, 1984] G. Bell. Evolutionary and nonevolutionary theories of senescence. *American Naturalist*, 124:600–603, 1984.

[Boerlijst and Hogeweg, 1991] M. Boerlijst and P. Hogeweg. Spiral wave structure in prebiotic evolution: Hypercycles stable against parasites. *Physica*, 48D:17–28, 1991.

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

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

[Bramlette, 1991] 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.

[Caswell and John, 1992] 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.

[Collins and Jefferson, 1992] R. Collins and D. Jefferson. The evolution of sexual selection and female choice. In F. J. Varela and P. Bourgine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 327–336. MIT Press, 1992.

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

[Cristofalo, 1988] V. J. Cristofalo. An overview of the theories of biological aging. In J.E. Birren and V.L. Bengtson, editors, *Emergent Theories of Aging*, pages 118–126. Springer Publishing Company, Inc., New York, NY, 1988.

[Cutler, 1978] R. G. Cutler. Evolutionary biology of senescence. In J.A. Behnke, C.E. Ellicott, and G. Moment, editors, *The Biology of Aging.*, pages 311–359. Plenum Press, New York, NY, 1978.

[Fogarty, 1989] 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, 1989.

[Goldberg, 1989] D. E. Goldberg. *Genetic Algorithms: In Search, Optimization and Machine Learning.* Addison-Wesley, Reading, MA, 1989.

[Hesser and Männer, 1990] 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.

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

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

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

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

[Kauffman and Johnsen, 1991] 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.

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

[Langton *et al.*, 1992] C. G. Langton, C. E. Taylor, J. D. Farmer, and S. Rasmussen, editors. *Artificial Life II*, Reading, MA, 1992. Addison-Wesley.

[Langton, 1989] C. G. Langton, editor. *Artificial Life*, Reading, MA, 1989. Addison-Wesley.

[Langton, 1994] C. G. Langton, editor. *Artificial Life III*, Reading, MA, 1994. Addison-Wesley.

[Lewontin, 1979] 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.

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

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

[Maley and Caswell, 1993] 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.

[Maley, 1994] C. C. Maley. A model of the effects of dispersal distance on the evolution of virulence in parasites. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 152–159, Cambridge, MA, 1994. MIT Press.

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

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

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

[Nowak and Sigmund, 1993] 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.

[Partridge and Barton, 1993] L. Partridge and N.H. Barton. Optimality, mutation and the evolution of aging. *Nature*, 362:305–311, 1993.

[Ray, 1992] 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.

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

[Smith, 1992] J. Maynard Smith. Byte-sized evolution. *Nature*, 355:772–773, 1992.

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

[Sonneborn, 1978] T. M. Sonneborn. The origin, evolution, nature, and causes of aging. In J.A. Behnke, C.E. Ellicott, and G. Moment, editors, *The Biology of Aging.*, pages 361–374. Plenum Press, New York, NY, 1978.

[Tate and Smith, 1993] 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.

[Taylor and Jefferson, 1994] C. Taylor and D. Jefferson. Artificial life as a tool for biological inquiry. *Artificial Life*, 1:1–14, 1994.

[Valentine and Walker, 1986] J. W. Valentine and T. D. Walker. Diversity trends within a model taxonomic hierarchy. *Physica*, 22D:31–42, 1986.

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