

Experimental Lamarckism

Brian Hayes

A reprint from

American Scientist

the magazine of Sigma Xi, the Scientific Research Society

Volume 87, Number 6
November–December, 1999
pages 494–498

This reprint is provided for personal and noncommercial use. For any other use, please send a request to Permissions, *American Scientist*, P.O. Box 13975, Research Triangle Park, NC, 27709, U.S.A., or by electronic mail to perms@amsci.org. Entire contents © 1999 Brian Hayes.

Experimental Lamarckism

Brian Hayes

A curious young elephant got its nose stretched by a crocodile, with the result that elephants everywhere now carry a trunk. What this story tells us is that Rudyard Kipling was a thoroughgoing Lamarckian—a believer in the inheritance of acquired characteristics. As it happens, Lamarckian ideas were already in disrepute when Kipling wrote his “Just So” stories. The German biologist August Weismann, in a remarkably Kiplingesque experiment, had shown that chopping off a rat’s tail did not lead to the birth of tailless ratlets. Experimental protocols have gotten more sophisticated since then, but the verdict is the same: There’s no sign of Lamarckian inheritance anywhere in the kingdoms of life.

But why not? A few years ago Colin McGinn wrote (in a review of a book by Daniel Dennett): “Why have Lamarckian organisms never evolved? Surely a mutation which made the genes responsive to changes of phenotype (‘learning’) would have selectional advantage, and there seems to be no *physical* impossibility in such a set-up. Wouldn’t natural selection favour a physiological mechanism that allowed learned characteristics to be passed genetically to offspring?”

These are good questions. One way of answering them is to note that the molecular pathways needed for Lamarckian inheritance just don’t exist. Within the context of life-as-we-know-it, there’s no way for the elephant’s nose to talk to the elephant’s genes—especially the germ-line genes. The “central dogma” of molecular biology says that information flows from DNA to RNA to protein, not the other way around. A Lamarckian feedback loop would seem to require some mechanism by which the proteins of the phenotype could alter the DNA of the genotype.

The trouble with such an answer is that it invites a further annoying question: Why is it that such feedback loops have never evolved? Given all that *has* evolved in the way of genetic detours and shortcuts—plasmids, transposons, retroviruses, prions—it seems a bit arbitrary to declare this one pathway out of bounds. The case of retroviruses is particularly provocative,

since they produce an enzyme (reverse transcriptase) that violates the central dogma, copying information from RNA back into DNA.

Here’s another possible reason for the absence of Lamarckian inheritance in nature: Maybe it’s just not worth the bother. Many authors seem to take for granted that a genetic means of passing on learned traits would be beneficial if it could exist. They assume Lamarckism would make for a smoother and quicker kind of evolution than Darwin’s blindfolded selection of random variations. But what are the true costs and benefits of Lamarckism? Perhaps the reason we see no Lamarckian organisms is not that nature cannot invent the necessary apparatus but rather that the result is maladaptive. Lamarckism could be a trick that nature has tried and discarded.

I have attempted to investigate this issue through some simple computer simulations. Specifically, I’ve addressed the following question: If you were offered a Lamarckian capability, how much should you be willing to pay for it, when the price is exacted in the form of some compensating detriment to fitness? My experiments in free-market genetics are too crude to yield a definitive answer, but I can report that within the rather narrow bounds of this one model, I’ve been unable to find any situation where the benefits of Lamarckism would justify paying more than a small price.

The Evolution of Evolution

Jean Baptiste Pierre Antoine de Monet, Chevalier de Lamarck, was treated badly by his contemporaries and worse by history. At the Muséum d’Histoire Naturelle he held the lowliest professorship, namely Professor of Insects and Worms, but he turned this academic insult to good advantage, establishing the division between vertebrate and invertebrate animals. And he devised a theory of speciation through gradual evolution 60 years before Darwin published his *Origin*. Today, however, Lamarck is remembered only for his great error—his thesis that evolution works by the transmission of traits acquired through habits of use or disuse.

The idea must have seemed irresistible. If you play a lot of basketball, Lamarck says, you’ll have taller children. And he appears to be right: The

Brian Hayes is a former editor of *American Scientist*. Address: 211 Dacian Avenue, Durham, NC 27701. Internet: bhayes@amsci.org.




Figure 1. Three hypothetical organisms adapt to a changing environment in different ways. The vertical position of each dot represents the camouflage color of a moth; the diagonal line indicates the shifting background color the moths must match to survive. Darwinian moths cannot change as individuals, but the distribution of colors shifts in the population. Learning allows a moth to change more quickly but also reduces selection pressure, so that moths with maladaptive genotypes continue to appear. Lamarckian moths—which learn and then pass on their learned color to their progeny—track the changing environment closely. The question remains whether the benefits of learning and Lamarckism would outweigh the costs. Note that the graph for Darwinian evolution covers a longer interval.

children of basketball players surely *are* taller than average. Likewise, if you want your children to get into Harvard, go to Harvard yourself; the high rate of acceptance for children of alumni argues that education too is heritable. The fallacy in this reasoning is now plain, and no one would propose a Lamarckian mechanism to explain such correlations. Nevertheless, the suspicion lingers that if only the world *did* work Lamarck's way, it would work a little better.

The Harvard basketball team is not the most convenient context for a computer model of Lamarckian evolution. In searching for a simpler system, I have been inspired by the famous case of the melanic moths in industrial Britain. Dark-pigmented forms of the peppered moth *Biston betularia* were first noticed in the 19th century; they grew in abundance for several decades and then receded again after the 1950s. The cause of the original color shift was apparently the darkening of tree trunks by coal soot, which impaired the camouflage of lighter moths and left them exposed to predators. The later reversal of the trend coincided with measures to reduce air pollution.

My model of these events is highly abstract, with all the naturalistic details stripped away. It is not meant to reveal anything new about melanic moths but merely uses the idea of selection based on camouflage to explore some mechanisms of adaptation. The computer model is written in the programming language StarLogo, created by Mitchel Resnick of the Massachusetts Institute of Technology. (I discussed StarLogo in the January-February, 1999, "Computing Science" column.) For this project I employed StarLogoT, a variant developed by Uri Wilensky of Tufts University. The model and additional technical details are available on the *American Scientist* Web site.

Costs and Benefits

The rise and fall of melanism in the peppered moth was unquestionably a Darwinian event, brought about by natural selection acting on random mutations. Inheritance of acquired characteristics was impossible simply because there were no acquired characteristics to inherit. A moth has no way to change its color over the course of its lifetime, even if it could somehow figure out that making the change would be advantageous. And if the moth cannot adjust its color, it obviously cannot transmit any adjustments to its descendants.

But in imagination—or in the computer—we can rerun the experiment without the constraints of insect physiology. We can create chameleon moths that sense the color of their environment and adjust their own color to match. I shall refer to this adjustment process as *learning*, although it needn't imply any kind of cognitive capacity; the term is meant to encompass any adaptation within the lifetime of an organism.

Would moths that learn have an advantage over those that don't? It seems like a sure bet—and yet if adjustable camouflage is such an obvious asset, why don't all prey species have it? A likely answer is the no-free-lunch theorem. Learning has a cost, which in some cases may outweigh the benefits. At a minimum there is a complexity cost: Sensing the state of the environment and responding to it requires metabolic machinery that a simpler organism could do without. Building and maintaining that machinery incurs an energy cost; resources that might have gone into growth and reproduction have to be diverted into learning. Thus a creature that does a lot of learning could be expected to have a slower reproductive cycle than one with more

hard-wired traits. (*H. sapiens* takes 20 or 30 years to accomplish what *E. coli* can do in 20 or 30 minutes.) Thus adjustable camouflage might reduce mortality, but the price would be reduced fertility.

The cost-benefit analysis for Lamarckism is similar. In a Darwinian world, any acquired improvements cannot cross the generation gap. A smart moth born with white wings might darken gradually to match a sooty environment, but the moth's offspring would be white again (barring mutations). The moth's acquired pigment is no more heritable than a suntan. Lamarckism creates a link between learning and genetics. A moth that adjusts its color during its lifespan will give birth to offspring that share at least some of this adjustment. Is this shortcut advantageous? Again it would seem so. The young moths are hatched with protective coloration already in place. But, as with learning, maintaining the Lamarckian mechanism imposes a metabolic cost, so that lowering the death rate limits the birth rate. The balance between these two effects determines whether Lamarckian inheritance pays off. Finding the point of balance is the aim of the computer simulation.

Modeling Melanism

The main actors in a StarLogo program are mobile, animal-like objects. For historical reasons they are known as turtles, but they can just as well represent moths. Each moth has its own internal state, which includes a genome, a camouflage color and a level of energy reserves. The moths move over a background of "patches," which represent the color of the environment.

Melanism is often a polygenic trait, producing a more-or-less continuous range of hues. For simplicity I encode the camouflage color in a single gene with a continuous range of alleles; color can take on any value from 0 (darkest) to 1,000 (lightest). Separate variables represent the color genotype and the color phenotype; at birth the two variables have identical values, but in animals

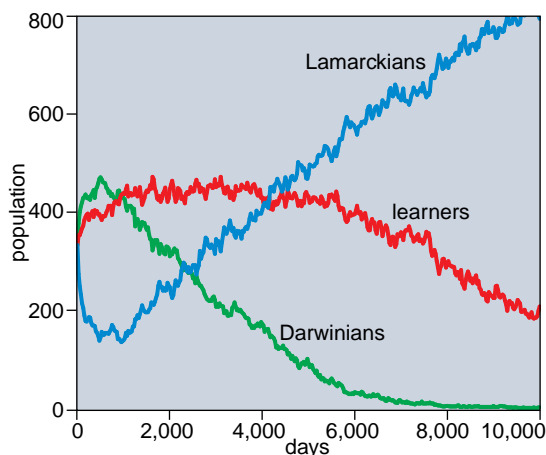


Figure 2. Competition of three species was initially dominated by Darwinians and then by learners but finally by Lamarckians. Paradoxically, though, the gene for Lamarckian inheritance was driven to its minimal value.

that learn they can later diverge. Lamarckian inheritance is implemented as feedback from the color phenotype of the parent to the color genotype of the offspring. In the extreme case of perfect Lamarckism, the color gene of the offspring is set equal to the parent's color phenotype; lesser degrees of Lamarckism interpolate between the original genotype and the acquired phenotype.

The moths have three other genes, which also range in value between 0 and 1,000. Following the custom of geneticists, I give the genes names: *kudzu*, *harvard* and *vanderbilt*. *Kudzu* is a growth gene: It determines the rate at which the moth absorbs resources, gains weight and grows toward reproductive maturity. Other things being equal, natural selection would drive this gene toward its maximum value, and that's what happens in the purely Darwinian case. For learning and Lamarckian moths, however, *kudzu* is linked to the *harvard* gene, which governs the rate of learning, and the *vanderbilt* gene, which controls the inheritance of acquisitions. The linkage is negative, so that faster learning or more complete inheritance of acquired traits entails a more severe penalty in growth rate. The constants that determine the degree of linkage—the tuition charged for learning and the inheritance tax imposed on Lamarckism—are the main parameters under investigation in the model.

A simulation begins with randomly assigned genotypes. Each moth grows at a rate determined by the value of the *kudzu* gene. On reaching a threshold weight (which takes 100 days at the highest possible growth rate), the moth produces two offspring and immediately dies. Meanwhile, each day a fraction of the moths are killed by predators. The probability of being killed increases in proportion to the difference between the moth's color and the background color. The overall death rate is adjusted to match the birth rate, keeping the population constant.

In those moths that learn, the color phenotype is adjusted every day, bringing it into closer correspondence with the background color at a rate determined by the value of the *harvard* gene. In the same way, Lamarckian inheritance adjusts the color genotype of the offspring toward the parental phenotype by an amount proportional to the value of the *vanderbilt* gene. All of the genes are also subject to random mutation and natural selection. (Even in Lamarckian moths, only the camouflage gene evolves by Lamarckian methods; all the other genes are purely Darwinian.)

The Backdrop to Evolution

When I first started up the model, I got an immediate reminder of a fundamental principle of evolutionary biology: No organism evolves in isolation. Evolution only makes sense as an interaction between the organism and its environment. I should not have needed a reminder—after all, the driving force in the peppered moth

story was environmental change—but in fact I had given too little attention to the backdrop against which the moths play out their drama.

What matters most about the background is not its specific color but the rate at which the color changes. In a static environment, learning is useless; there's nothing to learn. Darwinian mutation and selection can match an unchanging background just as closely as learning can, and so organisms unburdened by the overhead of learning will be favored. Running the model with an unchanging environment illustrates this effect clearly. Starting with random values of the *harvard* gene, the distribution shifts within a few dozen generations to favor the lowest values—those that produce the least learning but also incur the least penalty. In a mixed population of nonlearning Darwinian moths and learners, the outcome is even more dramatic. The Darwinians take over the population and drive the learners to extinction. And if learning is disadvantageous in these circumstances, then Lamarckism must also be unfavorable, since organisms that don't learn acquire nothing to bequeath their offspring.

The rate of environmental change does not have to be exactly zero to favor Darwinians. The rate merely has to be low enough to ensure that change is insignificant within the lifetime of an individual. Even a world with large and abrupt environmental transitions can penalize learners if the upheavals are separated by long interludes of stasis. Learners are better equipped to deal with the upsets, but they are wiped out by the faster-breeding Darwinians during the periods of calm.

Watching the simulations in action gives a new perspective on the relation between learning and life cycle. The idea that learning takes so much time and energy that it delays reproductive maturity is only half the story. The other half is that only a long-lived organism has any use for learning. Bacteria can rely on Darwinian evolution to fine-tune their metabolism to seasonal changes in temperature; as individuals they don't need to learn about hot and cold. Large mammals, on the other hand, would get no benefit from winter genes and summer genes, because they must cope with both seasons.

The More Things Change...

If slow change favors Darwinian selection, a rapidly fluctuating environment is where learning proves its worth. What is a little less obvious is that evolution becomes irrelevant here.

Even in the absence of learning, natural selection is helpless when change is faster than a generation time. If bark color can go from black to white in a week, and moths live several months, then the genotype can't keep up. Light-colored specimens might be favored one week, and produce more offspring than dark-colored moths, but the genes for paleness would be maladaptive by the time these descendants reproduced.

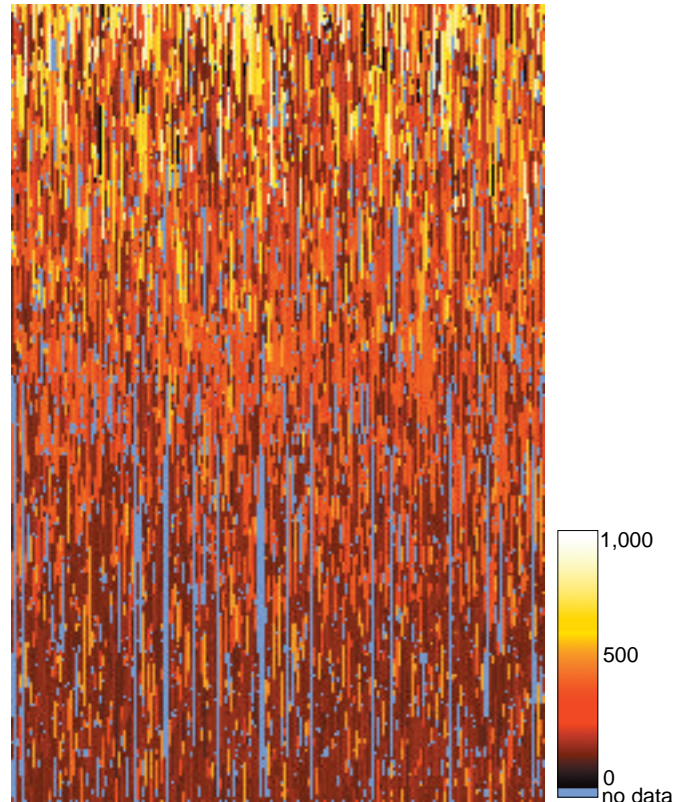


Figure 3. Gene for Lamarckism evolves toward its minimal value in almost all simulations, suggesting the trait is unfavorable. The gene values range from 0 to 1,000. The illustration tracks the genes of 200 individuals over a period of 10,000 days. Time progresses from top to bottom.

Turning on the *harvard* gene does nothing to restore the efficacy of natural selection; on the contrary, learning further decouples the genotype from the phenotype. When learning is rewarded, it becomes so efficient that there is little selection pressure on the genotype. Learning provides *any* newborn moth with excellent camouflage in a day or two, and so survival is essentially independent of the color gene.

What happens if Lamarckian inheritance is turned on in this rapid-change regime? Not much. If Lamarckism is assessed the same penalty as learning, the *vanderbilt* gene is disadvantageous, and the distribution of values slides toward the low end of the scale. This result is not hard to fathom. In the model the sole benefit of Lamarckian inheritance is being born pre-adapted to the color of the environment. But if that environment is changing rapidly, the benefit won't last long. Furthermore, in a population dominated by fast learners, most of the newborn moths would come to match their background in a few days anyway, even without the Lamarckian head start.

To put it another way: Learning is a valuable survival skill every day of your life, whereas Lamarckism helps only on the first day. This formulation suggests a way to quantify the worth of the *vanderbilt* gene. If your expected lifespan is L , then you should be willing to pay about $1/L$ as much for a Lamarckian legacy as you would pay

for learning. In the model, learners live more than 100 days, so Lamarckism should be worth less than 1 percent of the price of learning. If the penalties are adjusted accordingly—making the *vanderbilt* inheritance tax less than 1 percent of the *harvard* tuition—Lamarckism ought to spread through the population. In my experiments I could not see this effect clearly. Even at a cost of zero I couldn't be sure whether the *vanderbilt* gene was growing in frequency or merely drifting neutrally, but it may be that I wasn't patient enough to wait for the trend to become apparent.

Looking for Lamarck in All the Wrong Places

If Lamarckism has no value in a static environment and only the slimmest of marginal benefits in a fast-changing environment, it's natural to wonder if there might be some intermediate condition where the utility of Lamarckian inheritance is maximized. This would be a condition where change is quick enough to make learning worthwhile, but not so rapid as to make genetics irrelevant. I have surveyed a broad range of rates of environmental variation looking for this point of optimality, without reaching any firm conclusion. The level of statistical fluctuations in the output of the model suggests that much longer runs and larger populations would be needed to settle the question. I do feel confident in saying there is no level of variation where Lamarckian inheritance is worth as much as learning is, or even half as much. But there may be a range of variation rates where a species could benefit from a Lamarckian mechanism if it cost no more than a tenth of what learning costs.

Even where the model's answers are clear, they are at best preliminary and provisional. The model is too simple to capture much detail about the lives of real organisms. On the other hand, it's not simple enough to explore the entire space of parameter values. Another reason for caution is that the model sometimes behaves in ways I don't understand. For example, in certain mixed-population experiments the Lamarckian gene is driven toward the lowest possible values, indicating it is unfavorable and "wants" to be zero, yet at the same time the Lamarckian breed drives its competitors to extinction. What does *that* mean?

It would be interesting to test the model on problems other than moth camouflage, especially problems where the environment is not a passive background but can react and evolve on its own. One realm where Lamarckian mechanisms look particularly attractive is the immune system. Every child must reinvent immunity to measles and chickenpox and other diseases. It's done through a miniature Darwinian process of generating many random antibodies and selecting those that recognize a pathogen. Wouldn't it be better to pass on the selected antibody genes to later generations, so that babies would be born pre-immunized? Edward J. Steele of the

University of Wollongong in Australia argues that something like this does go on in the immune system, through the agency of reverse transcriptase. But Steele has won few converts. It's hard to be a believer in genetically transmitted immunity when your parents had chickenpox and you had chickenpox and your children get chickenpox.

If Steele's mechanism could exist, would it prevent disease? The obvious drawback is resistance: A generation born with a high titer of antibody would exert powerful selective pressure on the pathogen, so that mutants with slightly different surface markers would proliferate. Thus every generation would have to come up with a new defense anyway, and Lamarckian inheritance would be rendered superfluous. But this is speculation; a model might well reveal subtler effects.

In closing, I want to say a word about cultural evolution, which is often described as a Lamarckian process. Suppose it were *truly* Lamarckian: Suppose some neurogenetic innovation allowed your children to be born already knowing everything you know. What a boon to humanity! What a head start! No one would ever again have to spend all those years learning the alphabet and the multiplication tables and the conjugation of Latin verbs and the law of cosines and the preamble to *Evangeline* and the date of the Battle of Hastings and how to ride a bicycle. But the more items I add to this list—let's not forget the state capitals or the nine orders of angels or the 20 amino acids or the recipe for mom's meatloaf—the more I'm struck by the fundamental problem of Lamarckism. Which acquired traits do you choose to pass on?

Bibliography

- Ackley, David H., and Michael L. Littman. 1994. A case for Lamarckian evolution. In Christopher G. Langton, ed.: *Artificial Life III*. Reading, Mass.: Addison-Wesley, pp. 3–10.
- Belew, Richard K., and Melanie Mitchell (eds.). 1996. *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Reading, Mass.: Addison-Wesley.
- Hinton, Geoffrey E., and Steven J. Nowlan. 1987. How learning can guide evolution. *Complex Systems* 1:495–502. (Reprinted in Belew and Mitchell 1996, pp. 447–454.)
- Johnston, Timothy D. 1982. Selective costs and benefits in the evolution of learning. *Advances in the Study of Behavior* 12:65–106. (Reprinted in Belew and Mitchell 1996, pp. 315–358.)
- Lamarck, Jean Baptiste Pierre Antoine de Monet de. 1963. Original publication 1914. *Zoological Philosophy: An Exposition with Regard to the Natural History of Animals*. Translated, with an introduction, by Hugh Elliot. New York: Hafner Publishing Company.
- Majerus, Michael E.N. 1998. *Melanism: Evolution in Action*. Oxford, New York: Oxford University Press.
- McGinn, Colin. 1995. Left-over life to live. (A review of *Darwin's Dangerous Idea*, by Daniel C. Dennett.) *Times Literary Supplement*, November 24, pp. 3–4.
- Steele, Edward J., Robyn A. Lindley and Robert V. Blanden. 1998. *Lamarck's Signature: How Retrogenes Are Changing Darwin's Natural Selection Paradigm*. Reading, Mass.: Perseus Books.