Whole Earth Review, Spring 1989 n62 p90(10)

Universal parasitism and the co-evolution of extended phenotypes. (genetic influences reach outside the body) Richard Dawkins.

To understand Dawkins' thesis you'll need to keep the following distinctions in mind: Phenotype is the observable appearance of an organism, while genotype is the hidden governing constitution. The genotype manifests itself through the physical attributes of the phenotype. An organism that is of a particular genotype is called a genome.

IN MANY RELIGIOUS CULTS AROUND THE world, ancestors are worshipped. And well they may be, for ancestors, not gods, hold the key to understanding why living things are the way that they are. Of all organisms born, the majority die before they come of age. Of the minority that become parents, an even smaller minority will have descendants alive 1,000 years hence. A tiny minority are the only ones that future generations will be able to call ancestors. This minority had what it takes to be successful. Every organism alive can look back at its ancestors and say the following: Not a single one of my ancestors was killed by a predator, or by a virus, or by a misjudged footstep on a precipice, or a mis-timed handhold on a high tree branch, before begetting or bearing at least one child. Not a single one of my ancestors was too unattractive to find at least one copulation partner, or too selfish a parent to nurture at least one child through to adulthood. Thousands of my ancestors' contemporaries failed in all these respects, but not a single, solitary one of my ancestors failed.

Since all organisms alive inherit their genes from their ancestors, rather than from their ancestors' unsuccessful contemporaries, all organisms alive tend to possess successful genes. This is why organisms tend to inherit genes that build a well-designed machine, a machine that behaves as if it is striving to become an ancestor.

The rationale for this view of life can be seen only if we focus attention on the genes themselves (Williams, 1966; Dawkins, 1976). Genes are documentary information handed down, in the form of copies, from generation to generation. But genes are not only archival documents, passed like a family Bible from ancestor to descendant. They also exert a causal influence on each of the bodies in which they successively reside. They influence the development of arms and legs, of eyes and skins, brains and behavior patterns. Those genes that just happen to cause successive bodies to be more likely to die young, or to be unattractive to the opposite sex, or to fail in caring for children, are not the genes that pass through the net of natural selection into future generations of bodies. It follows that the animals that we see tend to be built by good genes: genes that are good at making bodies that, in turn, are good at passing those same genes on to future generations. It further follows that we can regard an individual animal as a machine for passing on the genes that it contains, a "survival machine" as I have put it.

The way that behavioral ecologists normally express this is to say that individual animals behave in such a way as to maximize their reproductive success. More precisely, it is referred to as their inclusive fitness (Hamilton, 1964). This doctrine has become orthodoxy. When a modern behavioral ecologist sees an animal doing behavior pattern A in situation P, his immediate reaction is to ask: "In what way is behavior pattern A good for the animal in situation P?" His colleagues may disagree with the answer he comes up with. Some of them may dispute the premise of the question, accusing him of being too "adaptationist, " perhaps of neglecting a "developmental constraint, or of neglecting the power of neutral drift. But, following my book The Extended Phenotype (1982), I want to raise a very different kind of problem. I suspect that the animal we are watching may be being manipulated by some other animal or plant, perhaps behind the scenes.

The animal we are watching is moving under the power of its own muscles, of course, and its own brain is giving the orders. Since the brain and muscles grew under the influence of the animal's own

genes we assume, as good neo-Darwinians, that the brain and muscles are working for the benefit of the animal's own genes. But what if there is some other animal lurking behind the scenes, pulling the puppet strings? Then, instead of asking "In what way is this animal benefiting from its behavior?" we should ask: "Which animal is this behavior benefiting?"

Parasites provide most of the examples we know about so far. Many flukes have a complicated life cycle, involving one or more intermediate hosts, before they finally infect their definitive host. For instance, flukes of the genus Leucochloridium have a snail as their intermediate host. From this they have to pass to a bird, and, in order for this to happen, their snail must be eaten by a bird, or at least the part of the snail containing the fluke. They could just sit back and wait for this to happen, but in fact they take active steps to make it happen. They burrow up into the tentacles of the snail, where they can be seen through the snail's skin, conspicuously pulsating. This makes the tentacles look to a bird like tempting morsels in their own right. Wickler (1985) suggests that they look like insects. Anyway, birds peck them off, and the fluke achieves the next stage in its life cycle.

What is more interesting from our point of view is that the flukes even manage to change the snails' behavior. The snails are normally negatively phototactic: they tend to avoid light, and therefore do not approach the tops of plants on which they feed. Infected snails change their behavior. They become positively phototactic, actively seeking light. This carries them up to the open tops of the plants, and makes them more likely to be seen by birds. Perhaps the fluke achieves this by interfering with the optic nerves of the snail: the eyes are, after all, in the tips of the tentacles into which the flukes have burrowed. From our point of view, it is sufficient that the parasites do change the behavior of the host, in such a way as to benefit the parasite, but not the host. If a behavioral ecologist watched the behavior of the snail, and asked: "In what way does its light-seeking behavior benefit the snail?" he would seek in vain for an answer. The truth is that some other animal, in this case a fluke, is manipulating the snail from behind the scenes. The behavioral ecologist would have done better to ask: "Which animal is this behavior benefiting?"

It is not just behavior that parasites manipulate. There is a protozoan parasite, Nosema, that infects beetle larvae. As far as the beetle larva is concerned, the purpose of its existence is to feed and grow until it is big enough to metamorphose into an adult beetle and reproduce. But the parasite has no interest in its host's reproducing. The parasite simply "wants" its host to go on growing and providing food for more and more of the parasite's descendants. It achieves this by a remarkable feat of biochemical manipulation. The parasites together (presumably they are a clone) succeed in synthesizing the juvenile hormone, or a close chemical analog of it. juvenile hormone is the substance that insects normally synthesize to maintain larval growth and inhibit metamorphosis. Human experimenters have shown that, if you inject an insect larva with juvenile hormone, you can stop it metamorphosing. These Nosema parasites have "discovered" the same thing! They synthesize the juvenile hormone and secrete it into the beetle larva's body. Instead of metamorphosing, the larva continues to grow through as many as six extra larval moults, end- ing up as a giant larva more than twice the normal size.

In the case of the snail's phototaxis, it might have been possible to regard the change as an accidental byproduct, not as a true adaptation by the parasite. In the case of Nosema, it is hardly possible to maintain this. juvenile hormone is not something that protozoa ordinarily have anything to do with. Achieving the feat of synthesizing a specific molecule like a hormone indicates true adaptation by natural selection over many generations.

Once again, the conclusion I want to draw concerns the kind of question that behavioral ecologists should ask. We are tempted to look at a giant beetle larva and ask: "How does this giantism benefit the insect?" Instead, we should ask: "Who is benefiting from the giantism?" The answer, once again, is not the animal itself, but a manipulator hidden behind the scenes.

These examples are all from the point of view of individual organisms. But, as stated at the outset, all adaptation should fundamentally be seen at the genetic level. If the animal we are watching is behaving for the benefit of a manipulator behind the scenes, we must express this at the genetic

level. Just as, in normal adaptation, we say that an animal behaves so as to benefit the genes that it contains, so, in the case of these parasites, we must say that the host behaves in such a way as to benefit the parasite's genes. And the reason is the same. just as, normally, an animal's development is influenced by the genes that it contains, so a parasitized beetle larva's development is influenced by the genes of the parasite. The conclusion of the doctrine of the extended phenotype" is that a gene in one animal may have phenotypic expression in the body of another animal. It is this doctrine that I want to persuade you of, and I am doing so largely by talking about parasites.

The snail can be regarded as a vehicle exploited by a fluke. A beetle larva can be regarded as a vehicle exploited by a protozoan parasite. But the selfish gene view of life sees this as just a larger version of the normal relationship of a gene to the body in which it sits. A body is just a gene's vehicle for getting into the next generation, and hence into an indefinite series of future generations. A snail is just a fluke's way of getting into a sheep, and hence of getting its genes into the future.

But why do we assume that the fluke genes work with a kind of group loyalty to one another, while the snail genes oppose them and work with a group loyalty to one another? Many people do not see this as a question that needs an answer at all. They see it as the starting assumption, that the whole of a body works together for the entire reproductive success of all of that body, in other words, for the propagation of all its genes.

But it is more fundamental for genes to work in their own interests. Under what circumstances might we expect genes within one genome to rebel, and not to pull together with one another for the common good? We would expect this if some genes had found a way of breaking out of the ordinary meiotic lottery involved in making gametes [the random division of chromosomes], and succeeded in manipulating their bodies into spreading them some other way. Suppose, for instance, that a gene succeeded in making its bodies sneeze them out, so that they could be breathed in by another body. Such a gene might well share with ordinary genes the same interest in preserving the individual body alive. But it would not share with ordinary genes the same interest in making that body have offspring, via sperm or eggs. This partial divergence of interests will tend to make the sneezed genes behave in a more detrimental, parasitic" manner. Are there any examples of such genes? Well, if there were, by definition we would not call them members of the body's own genome. We might call them virus genes.

The only reason all genes are not rebels like this is that all the genes in one individual organism normally stand to gain from the propagation of the gametes of that organism. Rebelling is difficult, for reasons that in themselves require an explanation, and which have to do with the disciplined fairness of the meiotic lottery. Given that rebelling is difficult because of the way meiosis works, selfish genes can normally actually benefit themselves best by cooperating with others in the same body, in order to promote the reproduction of that body, as a coherent entity.

Briefly, I believe that this amicable state of affairs comes about in the following general way. Genes that can make use of one another's products tend to prosper in one another's presence. This sets up a climate in which genes that cooperate are favored. "Climate" means a climate provided by other genes. From any one gene's point of view, other genes can be regarded as part of the environment, in much the same way as the external temperature and humidity can be regarded as part of the environment. "Cooperate" just means work together, especially work together to make the whole genome behave as a single coherently purposeful unit. This in turn increases the unitariness and coherence of the body, which in turn increases the pressure for the genes to be even more cooperative, and specifically increases the pressure for all the genes to converge upon the same method of leaving the body. So we have a self-sustaining, self-reinforcing evolutionary trend towards large units of phenotypic power. To go back to the example of snails and flukes, we normally think of parasites as weakening their hosts. But there are some cases where, at least at first sight, they strengthen their hosts. Cases have been reported of snails parasitized by flukes having thicker and stronger shells than unparasitized snails. Does this mean that the snails actually derive some benefit from the flukes? In the sense of being better protected, the answer may well be

yes, but it will not be a net benefit. When we consider benefits, we must not forget economic costs. It costs calcium and perhaps other resources to make a thick shell. We may be sure that the snail, and not the fluke, is bearing these costs. From the snail's point of view, a shell that is too thin is bad, for the obvious reason that it provides inadequate protection. But a shell that is too thick is also bad, because it consumes resources that could have been spent more profitably elsewhere in the economy of the snail. for instance, in making more eggs. Admittedly a super-thick shell presumably provides even better protection than a normal shell, but if, so to speak, the snails thought it worthwhile for this reason, they would have invested in it anyway! By making them have a thicker shell than they "want, " the flukes are not doing the snails a favor, unless the flukes are, in some way, shouldering the economic cost of the extra thickness. We may be pretty sure that they are not.

Is there any reason for the flukes to "prefer" a thicker shell than the snail does? Yes, I think a plausible case can be made, precisely because the flukes are not shouldering the economic burden. From the snail's point of view, the weighing up of costs and benefits can be thought of as a trade-off between survival and reproduction. A thicker shell means that the snail's own life expectancy is increased, but the economic costs of the thicker shell are felt as reduced reproductive success. Natural selection presumably arrives at an optimum balance.

But from the fluke's point of view the optimum balance looks different. The fluke is also inter- ested in the snail's survival, since its own survival is intimately bound up with the survival of its host (at least for a while). But the fluke has no specific interest in the reproductive success of its host. To be sure, it has a vague interest in the entire species of snails having reproductive success, so that there will be a new generation of snails to parsitize. But it has no specific interest in the reproductive success of its particular host, since the benefits of this to the next generation of flukes would be shared by all its rival flukes. As far as its particular host is concerned, it would be quite happy if that host were castrated. Indeed some parasites, as we know, do castrate their hosts, probably gaining benefits in the increased bodily growth of the host (Baudoin, 1975).

So, as far as snail shell thickness is concerned, there are two optima. The snail's optimum shell is thinner than the fluke's optimum. Switching, now, to gene language and the language of the extended phenotype, the snail phenotype is influenced not only by snail genes but also by fluke genes. These influences, to some extent, tug in opposite directions. The phenotype that we actually observe is probably a compromise between the two influences.

This is a slightly unfamiliar way of looking at life, so I will explain it in another way. Imagine three geneticists all doing research on the genetics of snail shell thickness. All three geneticists, in other words, are studying the same set of varying phenotypes. They differ with respect to the genes that they consider. One of the the three geneticists is a snail scientist. He studies the inheritance of shell thickness in pedigrees of snails. To him, the contribution of flukes to variations in the phenotype is strictly an environmental contribution to the variance. The second geneticist is a fluke geneticist. He studies the inheritance of host shell thickness in pedigrees of flukes. To him, the contribution of snail genes to variation in shell thickness is strictly an environmental contribution! I hope it is clear that both geneticists are practicing perfectly respectable genetics, albeit the fluke geneticist is a little unconventional. Yet each of them is relegating the genes studied by his colleague to the environmental category.

As you may have guessed, the resolution of this apparent paradox is achieved by the third geneticist. The third geneticist is an extended geneticist. He treats the variation in the shell phenotype as being under the joint influence of both snail genes and fluke genes. When you think about it, this is just what geneticists do all the time anyway, when they are studying genes within one genome. Geneticists are entirely accustomed to the idea that several genes influence the same phenotype. They normally think in terms of several genes of the "same" genome, but the whole point I am making is that there is nothing particularly special about the "same" genome. Fluke genes and snail genes can jointly influence the same phenotype, in just the same kind of way as snail genes and snail genes ordinarily interact with one another.

We have again reached our puzzle. Why do we assume that all the snail genes pull together as a team, while all the fluke genes pull together as a different team? The answer is not that there is anything qualitatively different about fluke genes and snail genes, some essence of snailiness or flukiness that pervades the substance of the genes. What, then, is the answer? The answer lies in the fact that the snail genes all share the same method of leaving the present snail body, and the fluke genes do not. The fluke genes in their turn all share the same method of leaving the present snail body, and the snail genes do not.

Why does the method of leaving the body matter so much? It matters because on it depends the series of events, in the future, from which the two sets of genes stand to gain. There is a partial overlap of interests. Both fluke genes and snail genes stand to gain from the snail's succeeding in finding food of the kind that best suits the snail's health. Both stand to gain from the snail's finding shelter from cold and other climatic hazards. Both, to a large extent at least, stand to gain from the snail's continuing to survive, But the two do not overlap in benefiting from the snail's reproducing. Snail genes that make the snail successful in finding a mate will be favored in the snail gene pool. Fluke genes that have the same effect on the snail will not be favored in the fluke gene pool.

In general, parasitologists should pay attention, above all other things, to the extent of overlap between methods of leaving the shared (host) body. Those parasites that put their gametes inside host gametes stand to gain from an almost identical set of future events to their host genes. They can therefore be expected to cooperate with their host as benign parasites or symbionts.

Some bacterial parasites of beetles not only live in the beetle's body. They also use the beetle's eggs as their transport into a new beetle. The genes of such a parasite therefore stand to gain from almost exactly the same set of future circumstances as the genes of their host. The two sets of genes, therefore, would be expected to pull together, for exactly the same reasons as all the genes of one organism pull together. It is irrelevant that some of them happen to be beetle genes while others happen to be bacterial genes. Both sets of genes are interested in the propagation of beetle eggs. Both sets of genes, therefore, are interested in making the beetle bodies successful in all departments of life, in both survival and reproduction. This is not true of the fluke genes and snail genes. The fluke genes care about snail survival, but not about snail reproduction. Therefore the cost/benefit calculations of snail genes and fluke genes come out differently. In the case of transovarial parasites like these bacteria, the cost/benefit calculations of host genes and parasite genes come out the same in all departments of life.

We now can take a radically unfamiliar view of any animal's "own" genes, and why they pull together for the good of all. The reason, quite simply, is that all expect to leave the present body by the same route as each other, by the same sperm or eggs. To be sure, in sexually reproducing organisms, not all genes get into all gametes. Indeed, each gene has only a 50-percent chance of getting into any given gamete. But all have the same statistical chance of getting into each gamete. As long as rogue genes do not cheat, and increase these odds - which some genes, the so-called segregation distorters, actually do (Crow, 1979) - all the genes stand to gain from the same set of events in the future. Fundamentally the reason is that meiosis is largely a fair, unbiased lottery.

This opens the new question of why meiosis is largely a fair, unbiased lottery. This is not a question I will tackle here. For now, I shall just accept that it is, and note what follows from it. The conclusion is that the genes of any one organism pull together for just the same reason as the genes of a transovarially transmitted bacterium pull together with the genes of its host. just as transovarially transmitted parasites are exceedingly "gentle" parasites - indeed not true parasites at all but mutualistic symbionts - so all the genes of a body can be regarded as gentle parasites of that body. The gentler the parasite, the more intimate the mutualism of a symbiotic relationship, and the less obvious it will be to us that it is a parasite at all. The parts will come to merge, until we cease to call the relationship parasitic or symbiotic, and think of the entire partnership as a single body. This is what has happened to mitochondria and other cell organelles, if Lynn Margulis's (1970) symbiotic theory is right. I want to go even further than Margulis, and regard all "normal" nuclear genes as symbiotic in the same kind of way as mitochondrial genes.

Parasites do not have to live inside their hosts. Cuckoos are perfectly good parasites, but they do not live inside their host's body, merely in its nest. They do not exploit the host's physiology directly, but indirectly via its behavior. But the principle is exactly the same, and the doctrine of the extended phenotype applies in the same kind of way.

It is easy to sympathize with the host foster parent when the cuckoo is at the egg stage. The eggs laid by a female of any one race closely resemble the eggs of the host species. The foster parent is fooled, in the same way as any victim of mimicry. We, can sympathize because human egg collectors - for such disreputable creatures were once, I regret to say, common - have frequently been fooled. We find it much harder to sympathize with the foster parent when the cuckoo youngster has grown near to the point of fledging. It seems to us the height of absurdity when we see a picture of a tiny reed warbler, standing on the back of its monstrous foster child in order to reach its huge open gape and drop food into it (Hamilton and Orians, 1965). Surely any fool could see that the nestling cuckoo is not a reed warbler. It is one thing to be fooled by subtle egg mimicry, but who could be fooled by a fake child seven times the size of the real thing? Putting the problem in a less subjective and more Darwinian way, how can natural selection be so efficient in perfecting the egg mimicry of the cuckoo, yet so inefficient in allowing grossly oversized nestlings to survive their foster parents' discrimination?

The problem is lessened by the following consideration. The cost of failure, from the point of view of the foster parent, is less at the egg stage of the cuckoo than at the nestling stage. A reed warbler who succeeds in detecting a cuckoo egg gains an entire breeding season. A reed warbler who succeeds in detecting a nearly fledged cuckoo has little to gain, since the season is nearly over anyway. But, even so, it seems hard to believe that a visual system sharp enough to detect the mimicry of cuckoo eggs could be "stupid" enough to be fooled by a cuckoo fledgling.

Perhaps "fooled" is the wrong word. A human male may be sexually aroused, even physiologically aroused, by a photograph or drawing of a female. Suppose a Martian ethologist observed this phenomenon. Would he say: "How silly to be fooled by this fake woman. Surely anyone can see that she is only a pattern of printing ink on paper, and only about a tenth of natural size. " Men of course are not actually "fooled" by the picture. They do not really think it is a woman. They simply find themselves aroused by it in the same kind of way as they might be by a real woman. Perhaps something like this is true of the cuckoo's foster parent. There are many well-documented observations of adult birds, of many species, flying home with food for their own young, and being diverted by the chance sighting of a gaping cuckoo nestling in another bird's nest. They then feed the cuckoo in the other bird's nest, in apparent preference to their own young in their own nest. Perhaps the cuckoo nestling is, as Oskar Heinroth is reported to have said, a "vice" of its foster parents. He said that the parents behave like "addicts. " Is the colored gape of the young cuckoo like an irresistible drug? Following Dawkins and Krebs (1978) and Krebs and Dawkins (1984), I want to make the general case that animals may manipulate other animals with weapons that we can best understand if we think of metaphors like "drugs" and "hypnosis. Keith Nelson once gave a talk about bird song entitled: "Is bird song music? Well, then, is it language? Well, then, what is it?" I want to make the case that, at least in some cases, it may be akin to hypnotic persuasion, spellbinding oratory, hauntingly irresistible music. The poet Keats wrote, in his Ode to a Nightingale,

My heart aches, and a drowsy numbness pains My sense, as though of hemlock I

had drunk, Or emptied some dull opiate to the drains One minute past, and

Lathe-wards had sunk. What I am suggesting is that nightingale song, cuckoo gapes, and many pheromones perhaps are exerting an influence on their receivers' ner- vous systems which is irresistible in the same kind of way as a drug may be irresistible. Or as the electric currents of a neurophysiologist may be irresistible. A neurophysiologist can implant electrodes in carefully chosen parts of the brain of a cat or a chicken and, by passing current down them, manipulate the

behavior of the animal like a puppeteer pulling strings. If the brain is vulnerable to such manipulation, should not natural selection, working on other animals, have perfected the power to manipulate? To be sure, animals cannot literally bore holes in one another's brains, cannot literally pass electric current in. But there are convenient holes already bored: eyes, ears, and noses. They provide ready-made channels into the deep parts of the brain and they are, in some senses, predisposed to be manipulated. A reed warbler's brain already has the predisposition to be attracted to the open gapes of it,,; own young. The young cuckoo has only to tap into this ready-made channel into the brain, and it apparently is not all that difficult to go one better and evolve a supernormal stimulus. Natural selection would surely favor animals that succeed in manipulating the nervous systems of other animals in this kind of way.

The obvious question now stands out. Why do victims of manipulation stand for it? Just as natural selection would favor manipulators who discover and exploit portholes into the brains of their victims, so natural selection will favor those would-be victims who close off those very portholes. How can there be any long-term future in manipulation as a way of life? One possible answer is that there is not any long-term future. It could be that cuckoos can survive only by exploiting evolutionary time lags. Perhaps cuckoos can exploit any one host species for only a few centuries, before the host gene pool accumulates enough genes for resisting manipulation. Then selection in the cuckoo gene pool favors those who start exploiting a new species which is still, evolutionarily speaking, naive about the dangers of being manipulated, There is some direct evidence that this may be at least a part of the truth (N. B. Davies and M. de L. Brooke, in preparation). But I doubt if it is the whole truth. I think we also need to consider the theory of evolutionary arms races, and how they may end (Dawkins and Krebs, 1979).

An evolutionary arms race is a process of co-evolution in which advances on one side are matched by counter-advances on the other, which in turn provoke further advances on the first side, and so on. Arms races are common between predators and prey, and parasites and hosts, and are one of the principal forces driving towards progressive evolution of ever more complex and sophisticated biological armament and instrumentation (Dawkins, 1986). As so far described, there seems no obvious way for an arms race to end. But this is too simple. We have left economics out of the discussion. Arms races do not, in any case, make sense without economic considerations.

There are economic and other costs to each side in each advance in the arms race. For a deer to evolve faster running, for example, it must develop bigger muscles. This means spending more resources on muscle tissue, resources which could have been spent on, say, reproduction. There will be some optimum compromise between amount spent on leg muscles and amount spent on reproduction. Any individual deer that spends less than the optimum will be vulnerable to being eaten. But also, any individual deer that spends more than the optimum will be less reproductively successful than an individual spending the optimum amount. The overspender, to be sure, may live longer as an individual. But it will not pass so many genes on to future generations, so genes for overspending will not increase in the gene pool. If it were not for such economic considerations, all animals would run as fast as cheetahs and would be as clever as humans.

Now, what happens to this optimum if there is an arms race going on? If the predators increase their running speed, there will be a shift in the timum balance within the deer gene pool. Individuals that previously would have been classed as overspenders now propagate more genes than individuals that previously would have been classified as optimal. So the deer population takes a step in the direction of greater average running speed. This in turn changes the optimum in the predator population, and so on.

But now, what if there are asymmetries in the economic calculations on the two sides of the arms race? Two thousand years ago, Aesop noted that the rabbit runs faster than the fox, because the rabbit is running for his life, while the fox is only running for his dinner. The cost of failure in running speed, for the fox, is merely a lost dinner. The cost of failure in running speed, for the rabbit, is a lost life. In the trade-off between spending resources on leg muscles and on reproduction, therefore, the optimum for the fox population could well come out very different from the optimum for the

rabbit population.

We can apply this kind of economic thinking to the case of cuckoo nestlings manipulating their foster parents. The cost of failure to a young cuckoo is death. The cost of failure to a foster parent is the loss of part of one breeding season. To put it another way, the cuckoo is descended from a long line of ancestors, every single one of whom has succeeded in manipulating a foster parent. The foster parent is descended from a long line of ancestors, only a proportion of which ever met a cuckoo in their lives, and even that proportion had another chance to reproduce after failing in that one year. Maybe the arms race between cuckoos and reed warblers has ended in a kind of stable compromise.

If there are economic costs to a reed warbler in resisting manipulation by cuckoos, it is even possible that natural selection among reed warblers favors complete capitulation. if cuckoos, for instance, were rare, then any individual reed warbler that was prepared, genetically speaking, to pay the cost of resistance, might actually be less successful than a rival individual that made no attempt whatever to resist cuckoos. Total nondiscrimination could be, for economic reasons, a better policy than costly discrimination, even though nondiscrimination carries the risk of parasitization.

If animals can manipulate other animals, and if the economics of arms races leads to stable equilibria in which the victims of manipulation acquiesce in being manipulated, we once again arrive at the same conclusion as before. When a behavioral ecologist looks at some feature of an animal's behavior, or anatomy, he should not necessarily ask, "How does this feature benefit the animal?" Instead, he should ask, "Which animal is this feature benefiting?" Whereas, before, the hidden manipulator behind the scenes was assumed to be a parasite inside the host's body, with direct access to the host's physiology and biochemistry, we have now extended our view to include manipulators outside the victim's body. The manipulator can even be a long way away, manipulating its victim by sound, or by chemical means.

I can summarize the extended phenotype view of life by contrasting it with two others in the form of diagrams. The two others can conveniently be labeled with the names of the great biologists who advocated them, Lamarck and Weismann. In the Lamarckian view of life (actually Lamark simply adopted a prevailing view of his contemporaries and predecessors, but his name is conveniently used as a label), bodies pass on their attributes to descendant bodies fig. 1). Hence new characteristics acquired during the body's life can be passed on. The Lamarckian view was replaced by the Weismannian view, according to which the germ-lines (we should now say the genes) are passed down the generations, influencing bodies as a side issue. A very important side issue, it has to be hastily said, since the survival or nonsurvival of the genes largely depends upon their effects upon bodies. The extended phenotype view of life (fig. 3) is an extension of the Weismannian view. Indeed, I would maintain that it takes Weismannism to its logical conclusion. There is still an immortal germ-line, and genes still survive or perish by virtue of their phenotypic consequences. But those phenotypic consequences are no longer limited to the body in which the genes are sitting. Genetic influences reach out beyond the body of the individual organism and affect the world outside, both the inanimate world and other living organisms. Coevolution, and the interaction between organisms, is best seen as an interlocking web of extended phenotypes.

Literature CITED

Baudoin, M. (1975). Host castration as a parasitic strategy. Evolution., 29: 335-352.

Crow, 1. F. (1979). Genes that violate Mendel's rules. Scientif. Am., 240 (2): 104-113.

Dawkins, R. (1976). The Selfish Gene. Oxford University Press, Oxford.

Dawkins, R. (1982). The Extended Phenotype. W. H. Freeman, San Francisco.

Dawkins, R. (1986). The Blind Watchmaker. Longman, London.

Dawkins, R. and J. R. Krebs (1978). Animal signals: information or manipulation? In Krebs. J. R. and N. B. Davies (eds.), Behavioural Ecology: An evolutionary approach, pp. 282-309. Blackwell Scientific Publications, Oxford.

Dawkins, R. and J. R. Krebs (1979). Arms races between and within species. Pmc. Roy. Soc. Lond. B., 205: 489-511.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. 1. J. Theor. Biol., 7: 1-16.

Hamilton, W 1. and G. H. Orians (1965). Evolution of brood parasitism in altricial birds. Condor, 67: 361-382.

Krebs, J. R. and R. Dawkins (1984). Animal signals: mind-reading and manipulation. In Krebs, J. R. and N. B. Davies (eds.), Behavioural Ecology: An evolutionary approach, pp. 380-402. Second Edition. Blackwell Scientific Publications, Oxford.

Margulis, L. (1970). Origin of Eukaryotic Cells. Yale University Press, New Haven.

Wickler, W (1968). Mimicry. Weidenfeld & Nicolson, London.

Williams, G. C. (1966). Adaptation and Natural Selection. Princeton University Press, Princeton.

Article A7422481