GRESHAM

COLLEGE

WHY HAVE SEX?

A Lecture by

PROFESSOR IAN STEWART MA PhD FIMA CMath
Gresham Professor of Geometry

11 February 1998

Gresham Lecture

Why Have Sex?

Ian Stewart 11 February 1998

From the standpoint of classical genetics, one of the biggest problems about sex is that it appears not to convey any evolutionary advantage. So there must be more going on than just classical genetics. Mathematical models help biologists to work out why — and when — sex can be a good thing.

Mendelian Heredity

We now know the physical — in fact, chemical — basis of of heredity. Organisms pass on characters through their DNA, and errors occur when the DNA copying process makes a mistake. This discovery has come only recently, but a thriving theory of genetics, founded upon clear mathematical principles, has been around for much longer. The theory is called Mendelism, after its discoverer, the monk Gregor Mendel. Mendel was a mathematics student at the University of Vienna, and he was required to take courses in other subjects too. Ironically, he failed his botany course and was therefore denied his teaching diploma, so he never rose above the ranks of temporary teacher. In order to study, he became a monk; he then succeeded *too* well, becoming an abbot: he was then overwhelmed with administrative duties and had to give up his scientific research. But in between, he made one of the key discoveries of his century.

Mendel studied heredity in peas. He kept lots of pea plants, and he cross-bred them, pollinating plants with various characteristics, using the pollen from other plants. For example, some of the plants produced green peas, some yellow; some produced wrinkled peas, some round. A typical experiment involved breeding a population of 'green' pea plants (by which terminology we mean that they produce green peas, not that the leaves are green) that would 'breed true', in the sense that all offspring were also green. At the same time he would breed a population of 'yellow' pea plants (producing yellow peas) that would also breed true. Now: what happens if you cross-breed green and yellow? Well, one plausible possibility is that you get plants that produce greenish-yellow peas, as if the colours get blended together. But that doesn't happen. Another possibility is that you get a mixture of plants, about half green and half yellow. That doesn't happen either. What *does* happen is that you get a mixture of plants, about three quarters yellow and one quarter green — a neat, tidy, and puzzling mathematical ratio of 3:1.

Mendel also looked at combinations of characters — say smooth/wrinkly peas as well as yellow/green. Now there are four possibilities: smooth yellow, smooth green, wrinkly yellow, wrinkly green; and here he found equally simple proportions, 9:3:3:1. Mendel could *explain* these proportions in a simple mathematical recipe. Give 'yellow' and 'smooth' the value 3, and 'green' and 'wrinkly' the value 1 — the proportions that occur for those characters *alone*. Then multiply the values together when characters occur in combinations — for instance, 'yellow wrinkly' gets the value 'yellow' × 'wrinkly' = 3 × 1= 3. The numbers you get for the pairs are 9, 3, 3, and 1; so the magic 9:3:3:1 proportions are governed by a kind of hereditary arithmetic. Might there be laws of heredity, as there are laws of motion and laws of chemistry? Mendel thought there were, and he found them.

Later generations, armed with the tools of statistics, have decided that Mendel's numerical results were a little *too* tidy: the good monk had perhaps unconsciously interpreted borderline cases in favour of the results he expected to obtain. But his central idea was spot on: there must be factors that determine the colour of the peas, and they must be inherited from *both* parents. Nowadays these factors are called *alleles*: these are related to, but distinct from, *genes*, a word that has made its way into everyday language. Alleles are *differences* between genes — the different forms that a gene can take. For instance, in any pea plant there is a gene for the colour of its peas, which may be either yellow or green, so the gene 'pea colour' has at least two alleles, yellow and green.

The simplest theory would be that each plant possesses a single allele that determines its pea colour. If so, then the offspring of a yellow-pea parent and a green-pea one would inherit two alleles — but it couldn't keep them both, because, like its parents, it must possess a single allele for colour. Would it select at random between its parents' alleles? If so, then on average, half of its offspring would be yellow-pea, half green-pea.

But the ratios are 3:1, not 1:1.

What actually happens is that the parents each have have *two* alleles, and the offspring inherits one from each of them, chosen at random. Call the alleles Y and G, so that the possible pairs are YY, YG, GY, and GG. If a pea plant has alleles YY or GG, then it is clear what colour its peas should be; but what about YG? Mendel's answer is that in such cases one particular allele always wins. That allele is said to be *dominant*, the other *recessive*. In peas, Y is dominant, G recessive. There is a familiar example in humans: black hair is dominant, red recessive. Two red-headed parents (both RR) cannot produce black-haired children; but two black-haired parents *can* produce red-haired children, provided each is of type RB, with a recessive allele for red hair and a dominant one for black.

At any rate, for pea plants each of the pairs YY, YG, and GY leads to offspring with yellow peas; and only GG to green. Note the numbers: three pairs yellow, one green — the magic 3:1 ratio.

Reproduction

In 1953 the great mathematician John von Neumann set himself a puzzle. Is it possible, in principle, to build a self-replicating machine? In other words, is the replicative ability of life (from which its flexible cousin, reproduction, stems) just a matter of getting the right kind of organisation? He wanted an answer that was not just a verbal description or a machine schematic: he wanted a mathematical theory of replication.

The technology of the time was not up to actually building such a machine, otherwise he would have wanted that too. Von Neumann, one of the fathers of the computer and a founder of some areas of mathematical economics, often strayed from his ivory tower. Today we probably could build a von Neumann machine if we really wanted to, but there's no very good reason to do so at the moment — and SF writers have argued in their stories that this might not be a terribly wise move anyway.

Von Neumann's answer not only showed that reproducing machines are indeed possible: it also led him to the same abstract scheme that living creatures — on this planet at least — actually use. It would have been a wonderful prediction, a major contribution of mathematics to biology, but for one coincidence: 1953 was the year in which Crick and Watson announced the structure of DNA. Biologists didn't have to *deduce* how living creatures reproduced: they could go and see what really happened.

Although von Neumann made a real mathematical breakthrough, the follow-up has been slow. We could do with a well-developed theory of replication in the abstract: what we actually have is a few scattered ideas and one golden, glowing theorem.

One of the things that that such a theory would have to do is distinguish between trivial and genuine replication. Every replicative system operates in some environment, and if that environment is too rich, then the system has easy ways to cheat. For example a letter is a self-replicating machine in an environment of photocopiers. As Karl Sigmund says in *Games of Life*, a robot that walked through a factory full of deactivated robots and switched them on would be a kind of self-replicating machine — but not a terribly interesting one. Closer to genuine replication are things like flames. A flame will spread, given combustible material and oxygen. Moreover, a flame has complex structure (flames really are much more complicated than we tend to think, as all combustion theorists know to their cost). What flames lack is heredity: their form is determined by their environment, not by 'what they are'.

So what we want is a machine that replicates itself stably and accurately, which employs only 'low level' materials from its environment. A robot that can dig up ores, smelt them into metal, dig sand and grow silicon crystals, etch its own chips, manipulate circuits and components, and eventually build another robot just like itself. Such a 'von Neumann machine' would be a pretty complicated object — unless it was designed a lot more cleverly than my description suggests. One potential route is nanotechnology, with microminiaturised machines in a microscopic environment. Biology went several levels of scale smaller, and used molecules. Von Neumann headed into the abstract territory of the mind, and came up with a mathematical scheme that replicated. And he invented a new kind of mathematical system, now known as a *cellular automaton*, to implement his scheme.

Von Neumann's machine has (at least) two parts. One is a manufacturing unit which, given suitable instructions and enough raw materials, can build just about anything. The other is the list of instructions needed to build *itself.* At first sight these requirements are paradoxical, for the following reason. If the instructions describe the complete Von Neumann Machine, then somewhere inside the Von Neumann Machine we find its instructions. So inside the machine's instructions there has to be a description of those instructions, which means that inside that description, there is a description of the description of the instructions... and so on in an apparent infinite regress. But Von Neumann realised that this is not necessary. Why should the instructions contain a description of themselves, when you've *already got* the full instructions anyway? So he added a third component, an instruction-copier. It didn't 'understand' what was written on the instructions, and it didn't interpret or obey them: it just copied them.

Now we start with a machine consisting of a builder, a copier, and an instruction list. When the builder is given the instruction list it makes a new builder and a new copier, and bolts them together. Then it puts the instruction list into the copier and makes a new, identical instruction list. Finally, still obeying instructions, it places the instruction list in its proper place inside the new machine.

The trick for avoiding the infinite regress is simple but subtle: it is to treat the information in the instruction list in two different ways at two stages of the replicative process. First, the information is *interpreted* and obeyed. At this stage, what is written on the list is crucial: mistakes don't just get copied, they make the builder go wrong. Then the information is *copied.* At this stage its meaning is irrelevant. Mistakes do not affect the copying process itself — though they do lie dormant, ready to wreak havoc when the new machine starts to build its own self-copy. One wonders if subconsciously von Neumann got this idea from the work of Kurt Gödel and Alan Turing, who proved that certain mathematical questions exist to which neither 'yes' or 'no' is a provable answer. Because their work rests on exactly the same dual interpretation of information: first as a sequence of instructions that has meaning within some chosen system, then as a string of symbols with no meaning whatsoever. The genome has exactly this feature: during the development of an organism it is interpreted, during cellular replication it is also simply copied.

Von Neumann could have stopped at that point, but there was still a residual question 'can a single system really do that?' So, at the suggestion of another mathematician called Stanislaw Ulam, he invented a gadget called a cellular automaton. This was an infinite square grid of 'cells', each of which could exist in a variety of internal states. The cells were also equipped with rules telling them how to change their states in response to those of their neighbours....

This may sound like a rather limited system, but it can carry out any process that a computer can. The difference is not in the abilities, but in how laborious they are. Von Neumann invented a cellular automaton with twenty-nine states, laid down in an inital pattern of around 200,000 cells. One region of those cells is a builder-cum-copier, the rest is an instruction list. The builder-copier region, following the rules of the automaton, extends a long tentacle out into empty space. Then it builds a copy of itself, follwing the instructions. Then it copies the instructions, adds them to the new region, and withdraws its tentacle.

Von neumann's ideas were simplified by E.F.Codd. Very recently they were taken up by Chris Langton, a complexity theorist, who devised an even simpler replicating automaton. Unlike those of von Neumann and Codd it did not contain a universal computer — but it seems unlikely that a bacterium contains a universal computer either. Its is overkill to make the builder capable of building anything: all it really needs to be able to do is build a copy of *itself*, copy unit included, and transfer the instructions around. Langton's automaton is a P-shaped region that fits inside a 10 × 15 rectangle. After a mere 151 time steps it spawns a copy if itself, nearby but turned through a right angle. Now both P's repolaicate again, yielding four copies. At generation four there ought to be eight copies; but because the P's are immobile, two of them compete for the same space and end up creating a 'dying' unit. By generation five this has died, and is surrounded by four other dying regions — but outside those there are ten living P's. And so it continues, with an ever-expanding ring of P's surrounding an area of dead and dying ones.

Sexual Balance

Reproductive techniques in the plant and animal kingdom can be divided into two drastically different types: asexual and sexual. In asexual reproduction, the only genetic changes from parent to offspring are those induced by accidental mutations. But in sexual reproduction, genes are 'mixed and matched' from those of *two* parents. The parents themselves are of two types: one male, one female. The most familiar example of asexual reproduction is probably *Amoeba*, a unicellular organism that reproduces by division; a more complex one is *Hydra*, which reproduces by budding — tiny Hydras grow on the surface of the parent, and eventually detach. Among the 'higher' animals, and most plants, sexual reproduction is the norm. Some creatures, for instance aphids, can alternate between sexual and asexual phases.

One of the earliest puzzles about sex to be solved by mathematical means was: how do the two sexes remain in balance? First you must understand that it is not obvious that they should. There is a reproductive asymmetry between males and females. For simplicity let's just think about mammals, though much the same story applies far more widely. In order to produce offspring, the female supplies a *small* number of eggs, primed and ready for fertilisation. The male, in contrast, supplies a *large* number of sperm. The imbalance is huge: in humans it is one egg versus about a hundred billion sperm. Even allowing for wastage, it is clear that the female numbers are what limits the growth of the population. Would it not be a better strategy to have just a few males, but an enormous number of females? Of course the ratio would have to be maintained by evolving some mechanism to produce many more daughters than sons, but that could surely have happened if there were good enough evolutionary reasons.

Or could it? The sex of a mammal, say, is determined by its two sex chromosomes, and these are of two types, called X and Y. Females have the pair XX, males the pair XY. This particular system has the mathematical advantage of 'closure': if you select one sex chromosome from each parent then again you get either XX or YY. In contrast, if females were XX and males YY, say, then there would be an awful lot of unaccounted XYs: either they would have to be eliminated, or there would be three sexes. An alternative in which there was only one sex chromosome would be X for females, Y for males. Why did evolution not choose this simpler mechanism? Let's leave that question on the back burner for a moment. Assuming that the choice of X or Y from the male parent is equally probable, it looks as if the 1:1 sex ratio is a direct consequence of the biochemical mechanism of heredity. However, the genetic dice can easily be loaded by selection processes. In humans, for instance, the sex ratio at conception is biased in favour of males. At three months after conception there are 20% more males than females, a 12:10 ratio. At birth the excess of males is 6%. By age 20, the sexes balance. After that there is an excess of females. Throughout life, males die off faster than females.

The 1:1 ratio at reproductive age therefore requires a different explanation. The answer to this particular puzzle was supplied by the geneticist Ronald A Fisher. He argued on genetric grounds that if there were more females than males, it would be in the females' genetic interests to produce more sons; conversely if there were fewer females than males, it would be in the females' genetic interests to produce more daughters. The reason is that she would then have more grandchildren.

This takes a little effort to see. Consider three successive generations, and suppose, for the sake of argument, that the overall sex ratio is always 3:1 in favour of females. In particular, suppose that there are 300 females and 100 males in generation 2. Consider a mother-to-be M in generation 1 and suppose that, contrary to all other members of her generation, she gets to choose the sex of her forthcoming child C. Now consider a randomly selected member G of generation 3, not necessarily related to M. What is the probability that G carries M's genes, that is, G is the child of C? Suppose C is male. Then he was competing with 99 others for the privilege of fathering G, giving a chance of 1/100. But if C was female, then he chance would be only 1/300, which is smaller. In other words, it would be in M's interests to have a male child rather than a female one. The same argument holds whenever there is an excess of females over males. And when there is an excess of males over females, then the same kind of argument shows that the best strategy is to 'choose' female offspring. Now an organism cannot actually choose the sex of its offspring, but evolution could develop a bias towards one sex or the other — so any organism whose sex ratio was biased in the strategically correct direction would gain an advantage over its fellows. And the strategically correct direction is to move the sex ratio back towards 1:1. In other words, if the sex ratio were not 1:1 then evolution, following the process just sketched, would move it to that value anyway.

Fisher's argument was couched in more formal mathenatical terms, but that's what it boiled down to. The argument makes a number of assumptions — for example that only one organism at a time gets to change its sex ratio — but more elaborate models yield the same conclusion. The theory has been tested experimentally by breeding fish in tanks. It is also confirmed by the remarkable behaviour of mealy bugs. Here the sex ratio depends upon when the female emerged from her pupa. If she mates immediately, her male offspring have an excess of 2% over females. If she mates after six weeks, the excess is 80%. After ten weeks, it is 890%! Evolution has 'deduced' that since she has waited such a long time to find a mate, there are precious few males around — and therefore the wise bet it to put your money on male offspring. More accuraetely, since the strategy *works*, evolution has blindly but inexorably settled upon it

Fisher's argument can be generalised, to yield the important concept of an evolutionarily stable strategy. Any strategy, be it for reproduction, securing food, acquiring habitat, or whatever, is said to be *evolutionarily stable* if on average it does not pay any one organism to deviate from the strategy. Fisher's argument shows that a 1:1 sex ratio is an evolutionarily stable strategy. I hasten to add that there are creatures in which the sex ratio is not evenly balanced, but in those cases other special factors come into play, and the sex ratios adopted turn out to be those of an evolutionarily stable strategy suited to the new context.

Why Two Sexes?

All this establishes that when there are two sexes then their numbers balance out. A similar agument explains why the males and females evolve 'in synch' so that they remain sexually compatible. It does not pay any individual to evolve in a direction that reduces its range of possible mates, and evolving away from sexual compatibility is one of the easiest ways to do just that. Less evident are the answers to two other questions: why have more than one sex, and why not have three or more?

There are several apparent genetic advantages in having (at least) two sexes. The most obvious is that change can occur not just via random mutation of individual DNA bases, or random insertion/deletion of subsequences, but by mix-and-match permutations of entire genes. This introduces a much greater element of variability. Moreover, because each gene comes in two versions, one from each parent, the diversity of genes in the population remains potentially high even when a great deal of selection has gone on as regards the genes actually being expressed. 'One to wash, one to wear'. And since many genes occur in multiple copies in the genome anyway, 'one to keep in the drawer in case'.

This may sound like a disadvantage, because it makes it a lot harder to weed out 'bad' genes. However, genes operate in a context: they may be bad in one context but good in another. For example imagine a gene that helped an animal survive on less water, but which also slowed its metabolism down. In dry spells such a gene might be 'good', giving it a competitive edge: although it is slow, the rest are dead. But in wet periods, its slowness works against it. Now think of a climate that runs to long relatively dry spells, over hundreds or thousands of years, intermingled with comparably long wet spells. It would not be good for the survival of the species if the gene for less water consumption got bred out of the population during a wet spell. With sexual reproduction and its spare copies, that is far less likely to happen. But a *really* bad gene — one that was fatal, for instance — would still slowly be bred out of the population.

Fisher had his own explanation of sexual reproduction, a rather simpler one. The ability to mix and match genes implies that useful genes can occur more easily in combination with each other — say a gene for better eyesight *plus* one for faster movement. If each has to evolve slowly, mutation by mutation, then it takes long enough to get one of these genes into the gene pool: the chance of getting both in one individual is very small indeed. Unlike his explanation of sex ratios, this theory has not survived. It runs into trouble with timescales: on the whole, useful characteristics tend to be introduced one at a time anyway. And even of that were not the case, selection can act on many characters simultaneously: in a sense it is the whole genome that evolves, not individual genes.

Whatever the details, until recently the accepted theory for the existence of sexual reproduction was that it increased potential diversity. In the mid-70s, however, the view changed dramatically, until the big problem seemed to be why sex had not disappeared altogether. The difference in attitude stemmed from a better understanding of selection. Selection, it was believed, does not act on species, but on individuals. If you are a cat then your main competition is not birds: it is other cats. Because the other cats are after the same resources as you, whereas birds are mainly looking elsewhere. In 1966 George Williams explained how beneficial effects on entire species could stem from selection at the level of individuals.

The one glaring exception was sex. Sex seemed to be good for the spceies, yes — but bad for the individual. Calculations made by John Maynard Smith indicated that the beneficial results of sexual reproduction had to lead to the survival of at least twice as many offspring before sex was worth bothering with. The reason: in asexual reproduction you pass on *all* our genes to your offsrping: in sexual reproduction, you only pass on half of them.

So why sex? Obviously there *is* an advantage, otherwise it would have evolved away. But what is that advantage? According to the lottery theory, sex lets you spread your genetic bets over a lot of different tickets. This idea did not survive the research if Graham Bell. If it's a matter of winning genetic lotteries, then sex should be associated with ecological uncertainty. It should be most common among creatures that live in rapidly changing environments and produce large quantities of offspring. But in fact the reverse is true. It is the large, slow-breeding animals in stable environments that are sexual.

Muller's Ratchet and Red Queen

By 1985 only two theories still survived. In 1982 Alexey Kondrashov argued that sex could reverse an effect known as 'Muller's ratchet' after Hermann Muller, who pointed out in 1964 that random mutations are far more likely to be bad than good, and therefore should accumulate in asexual populations. The main objection, again, is timescale. Mathematical models show that Kondrashov's mechanism is too slow. So since 1980 another theory has been gaining adherents. It is named after the Red Queen, Lewis Carroll's character in *Through the Looking Glass* who had to keep running to stay in the same place. The Red Queen effect was so named by Leigh Van Valen, who encountered it when studying marine fossils. Typically, the leading journals — or more accurately their referees — refused to allow such a revolutionary hypothesis to be published, but it gained ground anyway. The Red Queen theory leads to the conclusion that sex exists in order to combat disease. It is my argument about a gene for surviving in dry weather, but with genes that protect against parasitic diseases. The extra twist, which adds more weight to the argument, is that the scenario no longer depends on accidental environmental changes. Instead, the entire dynamic is unavoidable. If many organisms evolve a gene that protects them against a parasite, then the parasite will die down to a low level. It is then in the parasite's interests to evolve a new attack strategy, so the organisms no longer need to protect themselves against it. Now, if the original anti-parasite gene died out altogether because of this, as it would in asexual species, then the few parasites that were still operating the old strategy could make a dramatic come-back. However, sexual species, keeping a copy of that obsolescent gene 'in the drawer in case', can select from a far larger genetic wardrobe. Genes that helped fight parasites many generations back can rapidly come back into general use if the need arises. In confirmation of this idea, it is striking that many of the most polymorphic genes — those with the greatest variety of alleles — are the genes involved in disease resistance.

The idea that sex helps species resist disease has a lengthy heritage: a version goes back to J B S Haldane in 1949 and it was made more explicit by Suresh Jayakar a couple of years later. In the late 1970s five people independently developed similar ideas, among the Bill Hamilton. Hamilton tested his theory with a computer model of sex and disease: 'a slice of artifical life,' as Matt Ridley has described it. First, sexual and asexual versions of an artificial organism were pitted against each other, and against Death, who struck at random. The asexual organisms won hands down. But then Hamilton introduced parasites into the simulations, with virulence genes that affected how rapidly they spread and how severe their effects were, and corresponding resistance genes in the hosts. Now the sexual creatures often won.

Sexual Selection

Not only does sex in a species lead to improved resistance to parasites: it opens up a whole new range of phenomena. One, known to Darwin, is sexual selection. Why , Darwin asked, do peacocks have such massive tails? On the face of it, this obvious suvival handicap makes no sense. But if females just happen to prefer males with big tails, everything changes. Darwin did not explain why such a preference might arise, but in 1930 Fisher pointed out that it could be the result of a runaway feedback loop. The ingredient that makes the process possible is a link between the females' perceptions and their choice of mates. Suppose that a few of the females happen, possibly by chance, to prefer males with slightly bigger tails. If so, then this trait will be passed on preferentially to their offspring, so after a time most males will have slightly bigger tails. But now it is in the interests of the females to prefer larger-tailed mates, because those are what is available. So male genes for larger tails and female genes for preferring larger tails chase each other round and round a mathematical feedback loop, which only stops when it is brought to a halt by something with an even stronger effect on reproductive success. It all depends on that accidental slight intiial preference, and the whole story could have gone differently, with a preference for small tails, or curly beaks, or big feet instead.

Apart from its feedback stucture, this effect is not especially mathematical, but I mention it because a closely related one — symmetry — most definitely is. The prevalence of symmetry in organisms is striking: why does it happen? One possible explanation is simply that the form of an organism arises by a series of breakages of symmetry, occurring in a world that is governed by the everywhere symmetric laws of physics. Organisms may not have to work hard to be symmetrical: that could be the default.

Provided, that is, the symmetric state is stable. However, it might not be difficult for small perturbations to drive the developing creature away from symmetry. Indeed most apparent symmetries in organisms are imperfect. If loss of symmetry is easy, then biological development has to work rather hard to maintain it.

Why should it bother? Why is it bad for organisms to be asymmetric?

For some, there are mechanical reasons. If a horse had shorter legs on one side, it wouldn't be able to move so effectively, implying that the tendency towards asymmetry would be snuffed out by evolution. But why should it matter if, say, a peacock's tail becomes lopsided?

There are at least three theories of this. Two involve sexual selection; the third is a strange 'accidental' property of neural nets. Sex first. Fisher's idea is that female preference is dictated by an accident. Short tails could equally well have become the *sine qua non.* Amotz Zahavi begged to differ, arguing that the very thing that makes a big tail so bad for the male will cause females to prefer it. The reason, he said, is that only males with very good genes can survive the handicap of carrying around an extravagant tail. Females who mate with such males will pass on those genes to the next generation; those that do not will not, so their offspring will be at a disadvantage. Small tails never get into the picture, because *any* bird, good genes or not, can manage a small tail. Zahavi's suggestion found little favour until mathematicians such as Alan Grafen modelled the system, and found that it would work provided the handicap is a real one that really is indicative of 'good genes'. If it's possible to cheat, then the handicap theory bites the dust.

A variant of this theory holds that being able to maintain symmetry during development is itself an indicator of good genes. Animals with defective genetic machinery are liable to exhibit asymmetries — such as a lopsided tail. Evolution will therefore select for females that mate with symmetric males, because that's one way to give your offspring a better chance of inheriting those good genes. I think that the emphasis on genetics here may be overdone, and offer one further contextual observation: the most common way for an adult animal to be badly asymmetric is not to have poor genetics — which would make it less likely to survive into adulthood anyway — but to have been injured. In any species in which males and females cooperate to raise the children, an injured mate does not improve your offspring's survival chances. I'm not sure that a leopardess, say, is up to counting how many legs her intended mate possesses, but she may well be able to spot that he is asymmetric.

Indeed the final rival theory explains how she might do this, while rejecting the notion that sexual selection is the primary factor in the story. Before entering into the details, let me first remark that many experiments and observations support the notion that females often have a preference for symmetric males. For example in 1992 Anders Møller was trying to test the theory that female swallows prefer mates with long tails. He found that they did; but he also found that they preferred symmetrical tails too. This was unexpected. Even more unexpectedly, Randy Thornhill found that female Japanese scorpion flies prefer the smell of males with symmetrical bodies, compared to the smell of those with asymmetrical ones — even when they could not see the males concerned. (Good smell goes with good genes? I'm still trying to puzzle that one out.) Bees, it seems, prefer symmetrical flowers with nice, even petals. And in humans there is evidence that symmetry and beauty are allied, and that both males and females prefer beautiful mates.

Back to that final theory. In 1994 Magnus Enquist and Anthony Arak showed that asymmetry preferences might arise because of a side-effect in the neural networks used for perception. Such networks 'learn' to respond to certain stimuli but not others. Imagine a network that is learning to *recognise* a tail. You have to do this before you can evolve a preference for certain kinds of tail, right? So on an evolutionary timescale the perceptual net is subjected to various stimuli, some of them being tail-like, some not. Now a tail presents itself to the eye from a variety of angles — related by the physical symmetries of the universe. If tails are lopsided, then the recognition system is likely to be presented with both left- and right-handed tails. So it will learn to respond strongly to lopsided tails in either orientation. Now think what happens when it is presented with a symmetrical tail. This bears a strong resemblance to both left- and right-handed tails, and because of the way neural nets respond to stimuli, the net will respond even more strongly when resemblances to *two* stimuli are present. Computer experiments verify this contention. At virtually the same time, Rufus Johnstone carried out similar experiments, reaching the identical conclusion. So a preference for symmetric mates could well be a by-product of the need to *recognise* which things are mates. This theory does not conflict with sexual selection: it just indicates that a preference for symmetry is not just an accident of fashion, but is pretty much unavoidable.

Further Reading

Jack Cohen and Ian Stewart, *The Collapse of Chaos*, Viking, New York 1994.

David Concar, Sex and the symmetrical body, *New Scientist* (22 April 1995) 40-44;

Charles Darwin, *The Origin of Species*, Penguin Books, Harmondsworth 1985.

Richard Dawkins, *The Blind Watchmaker*, Longman, London 1986;

Richard Dawkins, *The Selfish Gene*, Oxford University Press, Oxford 1989;

Richard Dawkins, *The Extended Phenotype*, Oxford University Press, Oxford 1982;

Richard Dawkins, *River out of Eden*, Weidenfeld and Nicolson, London 1995;

Richard Dawkins, *Climbing Mount Improbable*, Weidenfeld and Nicolson, Viking 1996.

Magnus Enquist and Anthony Arak, Symmetry, beauty and evolution, *Nature* **372** (1994) 169-172

Magnus Enquist and Anthony Arak, The illusion of symmetry?, *The Journal of NIH Research* **7** (July 1995) 54-55.

Martin Giurfa, Birgit Eichmann, and Randolf Menzel, Symmetry perception in an insect, *Nature* **382** (1996) 458-461.

Rufus A. Johnstone, Female preference for symmetrical males as a by-product of selection for mate recognition, *Nature* **372** (1994) 172-175.

Mark Kirkpatrick and Gil G. Rosenthal, Symmetry without fear, *Nature* **372** (1994) 134-135;

David Palliser, Symmetry in the human body is sexier and healthier as well as aesthetically pleasing, says scientist, *The Guardian*, 10 August 1996 page 4.

Mark Ridley, *Evolution*, Blackwell Science, Oxford 1996.

Matt Ridley, Swallows and scorpionflies find symmetry is beautiful, *Science* **257** (1992) 327-328;

C. David Rollo, *Phenotypes*, Chapman and Hall, London 1994.

Karl Sigmund, *Games of Life*, Oxford University Press, Oxford 1993.

John Maynard Smith, *Evolution and the Theory of Games*, Cambridge University Press 1978.

John Maynard Smith, *The Evolution of Sex*, Cambridge University Press, Cambridge 1982.

Ian Stewart and Jack Cohen, *Figments of Reality*, Cambridge University Press, Cambridge 1997. Ian Stewart and Martin Golubitsky, *Fearful Symmetry*, Blackwell, Oxford 1992; Penguin, Harmondsworth 1993.

R.Thornhill, S.W. Gangstead, and R.Comer, Human female orgasm and mate fluctuating asymmetry, *Animal Behaviour* **50** (1995) 1601-1615;

GRESHAM COLLEGE

Policy & Objectives

An independently funded educational institution, Gresham College exists

- to continue the free public lectures which have been given for 400 years, and to reinterpret the 'new learning' of Sir Thomas Gresham's day in contemporary terms;
- to engage in study, teaching and research, particularly in those disciplines represented by the Gresham Professors;
- to foster academic consideration of contemporary problems;
- to challenge those who live or work in the City of London to engage in intellectual debate on those subjects in which the City has a proper concern; and to provide a window on the City for learned societies, both national and international.

Gresham College, Barnard's Inn Hall, Holborn, London EC1N 2HH
Tel: 020 7831 0575 Fax: 020 7831 5208
e-mail: enquiries@gresham.ac.uk