CHAPTER 7

THE GENE'S EYE VIEW

7.1. Genes and Darwinian Populations

[Genes as parts of organisms; genetic accounting; selection of scaffolded reproducers; transposons, homing endonucleases, meiotic drive.]

For many biologists and philosophers there has been an elephant in the room throughout the last few chapters. A very tiny elephant, but a watchful and omnipresent one. The elephant is the gene, and the possibility of dealing with many or all of these problems from a "gene's eye view." This is not just a matter of making more use of genetics; it involves taking the perspective of individual genes as evolutionary units.

The "gene's eye view" is associated with a family of ideas, some more empirical and some less so. First, there is the idea that in *some* cases, perhaps very special ones, genes are the units of selection rather than organisms or other entities (Burt and Trivers 2006). A second, and quite different claim, is that *all* cases of biological evolution (or almost all) can be *represented* that way. The gene's eye view gives us one available description that coexists with other descriptions of the same cases. To this may be added the claim that a few phenomena *only* have a gene-level description (Sterelny and Kitcher 1988). Then there is the strongest version of the view, seen in Dawkins (1976), which holds that it is a mistake to describe most cases as selection on anything else. Natural selection is a contest between replicators, and genes are almost the only replicators—certainly organisms and groups do not usually qualify.

There has been a great variety of philosophical criticism of these ideas, mostly directed at the second two views listed above. (See Okasha 2006 and Lloyd 2006 for reviews.) But against this there can be found not just philosophical defenses but a growing list of empirical cases where the gene's eye view seems to help—a list of phenomena that we have got a *handle* on via this way of thinking. Some of these cases may be accepted by critics of the gene's eye view, as a short list of special phenomena that can be accommodated in a multi-level picture of evolution. But then we have to work out why the gene-level description is appropriate in those cases and not in others. For defenders of a gene's eye view,

the "special" cases are just those where the general nature of gene action is seen in a pure form.

The treatment here proceeds, again, by application of the framework defended in earlier chapters. This makes sense of the appropriateness of a gene-level description in some special cases, explains the artificiality of descriptions in purely genic terms in other cases, and explains how the two are related. Here is an outline of the main points. To treat genes as "units of selection" in some situation is to treat them as making up a Darwinian population. Roughly speaking, genes are scaffolded reproducers; the copying of DNA is a kind of reproduction. In some respects, genes are like cells: they are small parts of organisms which, because they can reproduce, make up lower-level Darwinian populations. Genes are also central to inheritance in both cells and whole organisms. As a consequence, a lot of evolutionary change at those other levels can be *tracked* in genetic terms. Much of that description has a special status that is easy to misread. It often appears to be a description in which a Darwinian pattern of explanation is applied to genes, when in fact the Darwinian description is being applied to *organisms* that are described and categorized in terms of their *genetic properties*. In addition, however, there are some phenomena in which genes enter into distinctive Darwinian processes of their own. Most of these cases are the products of the details of eukaryotic sexual machinery. From an evolutionary point of view, in fact, the very existence of genes as units is dependent on this modern sexual machinery. The question of why that machinery exists then looms large.

The status of genes as Darwinian individuals is shown to be quite different from the other cases discussed in this book. By some strict but reasonable standards, genes as evolutionary units do not exist at all. By more relaxed standards, they do. Evolutionary explanations at the genic level require the relaxed standards. And "selfish gene" cases do not function as models or exemplars, in which the overall nature of evolution is seen in a pure form.

To begin, let us look at the place of genes in the biological hierarchy. Cells are parts of organisms. Chromosomes are parts of cells. And genes are parts of chromosomes. This last claim might be met with some unease, for reasons discussed below. But initially at least, I will treat genes as small parts of living things.

In the preceding chapters, when low-level parts of organisms were discussed I usually chose cells. Cells give us a partial model of how to treat genes, but the two are different in several ways. First, genes are scaffolded reproducers, unlike cells which are simple reproducers. DNA is replicated as part of the process of cell division, via the larger machinery of the cell. In a sense, cells too cannot reproduce "on their own," as environmental conditions need to be suitable. But in the case of genes, the required environmental conditions are *very* specific—they need to contain almost all of the machinery of genetic reproduction. In addition, the

role of genes is made complex by sex and meiosis, which create cells with new combinations of genes derived from separate sources. In the previous chapter I noted that in the case of cells, biologists usually focus on evolution in the separate populations *within* each organism, rather than the total population of (say) human cells. In the case of genes, in contrast, people usually think in terms of a population extending *across* organisms—evolution in the total pool of human genes. But in principle, both kinds of analysis can be done in both cases.[1]

This is also an area where the literature has evolved special "mixed" ways of talking that do not apply Darwinian concepts in a straightforward manner. I will give two illustrations. First, within a simple and direct treatment of genes as potential Darwinian individuals, *all* the gene-sized bits of DNA are potentially competing with each other. It is often said, in contrast, that a gene competes only with its alleles at a given locus. This organizes much talk of genetic competition, but it is also known not to be true. Various phenomena associated with "transposons"—genes which move around within the genome—show the falsity of the claim. Those phenomena are, and are accepted as, Darwinian. The "only compete within a locus" rule is not one that genes themselves respect. Genetic competition occurs across loci as well as within them; these are just different environments which a genetic element can occupy.

Second, a lot of talk about the evolutionary role of genes is really talk about organisms, genetically characterized. This can be seen in the accounting. In standard models, a diploid organism will be described as *AA*, *Aa*, or *aa* at a particular locus. It might be said that the *A* allele is selected over the *a* allele, for example, when the fitness of the *AA* combination is higher than that of *Aa*, and that of *Aa* is higher than that of *aa*. This will lead to an increase in frequency of *A* over *a*. This familiar description is more unusual than it looks. If there are more *AA* organisms than *aa* organisms, in the sense above, that does not imply there are more physical copies of the *A* allele than of the *a* allele in the population. The *aa* organisms might contain many more cells than the *AA* organisms, and hence there may be more physical copies of *a*. In the standard accounting, each diploid organism is counted as equivalent, and each contributes two units to the calculation of the genetic composition of the population at that locus.

So there are different possible ways of counting genes in a population. Call the "standard count" the one that counts each organism as equivalent regardless of the number of cells and hence gene copies. The alternative one, which treats all gene copies on a par, can be called the "simple count." The special features of the standard ways of accounting show up also in the treatment of loci. When a

[1] A cellular population within an organism like us is asexual. The "total" population of human cells, in contrast, shows asexual cell division, reductive cell division producing gametes, and fusion of gametes. So it is a complicated system—but so are many protist populations.

gene jumps to another locus, the standard count does not treat this as an extra copy of the old gene. It is a new allele at that locus. Within mainstream models, talk about how selection leads to the proliferation of a gene (and so on) is usually a mixture. Some of it is directly aimed at tracking the spread of copies of a gene. But much of it is really talk of the natural selection of organisms, genetically characterized.[2]

It is no accident that organisms are being genetically described and categorized. Genetic properties do not merely give a handy label. Gene action is often causally responsible for one organism reproducing more than another, and genes are also central to the patterns of inheritance in organism-level populations. So genetic properties are pivotal to the evolutionary role of organisms. But this is a situation in which a Darwinian population is being recognized at the level of organisms—they are the things whose reproduction we are tracking and explaining—and that Darwinian population is being affected by the genetic properties of its individuals. This is one reason why genes, as discussed in evolutionary biology, sometimes seem partially abstract, less than fully material: talk of genes is not being used to refer to physical particulars made of DNA, but as a way of talking about sharable properties of organisms.

Let's think about what taking a *purely* gene-level view would involve. Shifts of this general kind were discussed in Chapter 6, for the case of organisms and their constituent cells. In the cell case, the shift in perspective is mainly a matter of "zoom," or how closely we look. Organism-level reproduction mostly *is* cell-level reproduction, suitably organized. The total goings-on at the cell level comprise most of what goes on in organisms. There we have a relationship between simple reproducers and collectives. That is not true in the case of genes. In the genetic case, an analogy that can be used is one of "staining" the organism-level population, as when using a microscope. Suppose we could stain all the DNA in the world, in a way that makes the rest of each organism invisible. We then also zoom in. We will see a great collection of facts about genetic reproduction, variation, and inheritance. We will also see packets of genetic material grouped into various kinds of interacting clumps. Suppose we are looking at humans. Then there will be many small packets containing a fixed number of tangled strands. These packets are internally diverse except that particular pairs of strands within them are similar. These (cell-level) packets are collected into (organism-level) clumps that are mostly very internally similar across packets. Variation exists mostly across, not within, the organism-sized clumps. We also find that most

[2] Arguments that the gene's eye view is merely a matter of "bookkeeping" were originally developed by Wimsatt (1980) and Gould (2002). I add that the bookkeeping being done is not a simple and direct form of genetic accounting, but more a counting of organisms guided by their genetic properties.

of the copies of any stretch of DNA are, because somatic, dead-ends and the producers only of short lineages. A few give rise to long lineages, and they are distinguished from others in their clump largely by location.

When we look at a case of evolution this way, most of the machinery by which genes reproduce, proliferate, and have their effects on the world is missing—is not visible under the stain. A proponent of the gene's eye view does not deny the importance of that machinery, but adds it back into the picture as *background*, as a context to gene action, and an arena in which genes compete. Why might we have reason to look at a case like this? There is good reason to do so in some cases. These are the ones classified as "selfish genetic elements" (Burt and Trivers 2006). Talk of "selfishness" will be discussed below, but the key feature of these cases is that the crucial step in their working is that there is a gene-level reproductive difference that does not go *via* a difference made to organism-level reproduction.

This list of such cases is long and fascinating. I will discuss three examples (drawing on Burt and Trivers's review) which illustrate different forms of the phenomenon. Transposons, mentioned earlier, are genetic elements that move to different places within the genome. If the old copy is retained as new ones are added (which happens in some but not all cases), then the genetic element increases in frequency over others. The description of transposons puts pressure on familiar ways of talking about genetic competition, as it does not involve competition "for representation at a locus." But if mitochondria can compete reproductively within a cell, so can stretches of nuclear DNA. There are various mechanisms by which this reproductive advantage can be gained. One example was given in Chapter 4, in my discussion of "formal reproduction." A LINE transposon codes for an mRNA molecule which is translated to produce a couple of proteins that bind to the mRNA and reverse-transcribe the RNA back into the cell's genome in a new location. So there are now two copies of that element in the genome where before there was one.

This a case where a genetic element proliferates *within* a cell and *across* loci. There are also cases of proliferation within a cell and within a locus. "Homing endonuclease" genes exploit the cell's machinery for DNA repair. When a chromosome breaks in a diploid organism, the cell uses the other matching or "homologous" chromosome as a template to repair it. This is because as well as joining the break, some DNA often must be replaced around the two sides of the gap. If the homologous intact chromosome differs from the broken one (if the cell is a heterozygote at that locus), then this process of repair creates a new copy of the DNA sequence of the unbroken chromosome. Homing endonuclease genes take advantage of this fact. They code for an enzyme that cuts DNA at a specific site, called a "recognition sequence." The DNA that codes for the cutter is also inserted into the middle of the recognition sequence itself. This disrupts the recognition sequence, so the cutter does not cut itself. But in a heterozygote

cell where one chromosome contains the cutter gene and the other does not, the cutter will break the other chromosome and thereby induce the cell to copy it into the other chromosome in the process of repair.

A third example is "meiotic drive." These genetic elements are diverse, but the general pattern is like this. A chromosome that can "drive" contains a "killer" element and a "resistant" element. In a heterozygote cell that is about to form haploid sex cells, there will be one chromosome with the "driving" complex and a homologous chromosome without it. Each sex cell produced will contain just one of the pair. The killer acts at some point during meiosis to sabotage the newly produced sex cell that contains the other chromosome. The driving chromosome has a "resistant" element at the place in the genome that the killer targets. This prevents the driving chromosome from destroying the sex cell that contains itself. This case is different from the other two because now the gene-level advantage does often go via a contribution to a difference in cell-level fitness. Whole cells are being sabotaged by the driving complex, those with particular genetic properties. In the cases of transposons and homing endonucleases, the process that generates a reproductive difference between genes takes place within a single cell.

So in the cases of transposons and homing endonucleases there is a gene-level reproductive difference that does not go via a contribution to cell-level or organism-level reproductive differences. In order for the genetic element to spread, the normal machinery of cell and organism reproduction must then enter the picture. But it makes sense to treat that machinery as mere background, because the crucial gene-level advantage was gained by processes within a cell. In the meiotic drive case, the gene-level advantage goes via a contribution to cell-level reproductive differences, but not via a contribution to organism-level reproduction. At least, that is true when the driving mechanism is acting alone, and often it does not. In many cases, an individual with two copies of a driving gene dies or is sterile. Even regardless of this, because of the role of cell-level fitness meiotic drive is a less pure case of gene-level selection than the others.

A biologist says: "A gene appears, which does X ... and it will proliferate." What this usually means is that an organism appears, with a new genetic property. The organism will reproduce successfully as a consequence. (The gene, given its context, makes for better camouflage, better disease resistance, a more impressive song.) The result is more organisms with that genetic feature. What the biologist sometimes means, instead, is that an organism with that genetic property will help *other* organisms with the same genetic property to reproduce. This is a case of the kind discussed in Chapter 6. If you help a brother, or a worker bee helps a queen reproduce, one organism contributes indirectly to the proliferation of organisms with its genetic features. Sexual reproduction leads to complicated cross-cutting patterns of genetic similarity in populations. The result is often a need for head-spinningly fine-grained genetic accounting (Queller and Strassman 2002).

Regardless, the Darwinian population in question is made up of reproducing individuals, with cells and genes as some of the individuals' parts. And what the biologist above *very* occasionally means is that the gene does something within individual cells to give it an advantage over other genetic material *in that cell*, or gives that cell an advantage over other cells within or produced by the same organism.

So in some cases it does make sense to focus on the activities of DNA per se, backgrounding most of the machinery of cell-level and/or organism-level reproduction. But those cases are unusual ones; to the extent that we organisms are products of gene action, we are not complex joint products of *that kind* of gene action. The picture, instead, is like this. To have evolution at all there must be some kind of reproduction, and evolution often gives rise to very sophisticated kinds. Once certain kinds of reproductive machinery are in place there is scope and space created for various additional Darwinian possibilities, via highly scaffolded, sometimes parasitic, reproduction. Given the presence of elaborate sexual reproductive machinery, we would *expect* some of this to arise. But these phenomena do not represent the general pattern of Darwinian evolution any more than parasites represent the general pattern of living activity.

Some of this picture can be seen in the Burt and Trivers survey I have drawn on in this section (2006: 25). Selfish genetic elements reliably arise, and tend to sweep through populations. But they often create conditions that undermine themselves; "selfish genetic elements almost invariably set in place forces that cause their own deterioration."

7.2. The Evolution of Genes

[*Genes as evolutionary units; dependence on crossing-over; team-shuffling analogies; evolution of recombination and the origin of genes.*]

This chapter has worked so far within a particular picture of what genes are like. Genes are treated as small stretches of DNA that are mostly preserved intact across generations while constantly entering into new combinations through sexual reproduction. This view is often described using analogies: genes are like cards that are repeatedly shuffled; genes are like rowers who are mixed into new teams (Dawkins 1976). Much of the time, that picture is accurate enough. But it involves an idealization, an imposed simplifying picture. In some contexts the idealization becomes misleading. A close look at where and how this picture breaks down leads to more conclusions regarding the status of genes as Darwinian individuals and units of selection.

I will introduce the main point immediately, and then approach it from several angles. When genes are recognized as units in an evolutionary context, a stretch of

DNA is said to count as a gene not only because of its effects on the organism, but because of how it is passed on. Genes are taken to have a degree of independence and persistence in this process. Thus a particular gene copy can, it is said, give rise to a definite lineage of descendants even though a chromosome cannot.

The fact that chromosomes cannot do this while genes can is a consequence of *crossing-over*: the exchange of genetic material between homologous chromosomes during meiosis. As advocates of the gene's eye view note, the "size" of a gene for the purposes of evolutionary explanation depends on the rate of crossing-over in the population. The way people commonly talk has it that the facts about crossing-over make chromosomes temporary, but leave genes as persisting and definite units. I will argue that this is not so.

This has consequences for evolutionary questions. In the most straightforward cases, a Darwinian population is made up of a set of definite countable things. Talk of genes as entities subject to natural selection relies on different and looser standards. My argument will not be that talk of gene-level reproduction and fitness makes no sense at all—the argument is not intended to contradict the previous section. But genes are not nearly as straightforward examples of Darwinian individuals as they look. In some ways they are marginal cases. It is not that there are no natural units at all in the genetic domain. Chromosomes and nucleotides are bounded natural units; we know where one ends and another ends. But the unit between these, the gene, is more dubious. In an evolutionary context it is more accurate to talk of *genetic material*, which comes in smaller and larger chunks, all of which may be passed on and which have various causal roles.

I will now go through these ideas in more detail. We can start with bacteria. How many genes are there in a typical bacterium? The standard answer is a few thousand (for example, four thousand in *E. coli*). This figure is basically a count of *cistrons*, genetic elements responsible for production of a single protein. The reality of these units gives us the length of each unit (a thousand nucleotides or so) and the rough location of boundaries between them. Then it seems that in a local population of a million bacteria there may be a few billion gene copies present (American billions, that is). We can count the bacteria, and we said there were a few thousand genes in each one.

But as many evolutionists will be quick to interject, this is not really the right count in an evolutionary context. Assume, for simplicity, that these bacteria do not engage in plasmid exchange, and the reproduction and spread of genetic material occurs only by simple cell division. Then the entire bacterial genome functions as an evolutionary unit; there is no basis for seeing it as a collection of distinct replicating things.

We now move to the case of humans, who are diploid and sexual. How many genes are there per individual? A standard figure is about 25,000. We could then go through, as above, a calculation for the number of gene copies in a local

human population. Each cell contains two sets of the 25,000. We multiply by a few trillion for the number of cells in a human, and then multiply by the number of humans in the population. But that 25,000 figure is, again, not a figure that is directly applicable in the evolutionary context. It is again a count of cistron-like units. This count is more complicated and problematic than it was in the bacterial case, due to the elaborate organization of eukaryotic genomes, but let us accept it as close-enough to right—certainly it is not a completely arbitrary number.[3]

In the bacterial case, several thousand of these things were combined in a single genetic reproducing entity. How many such things are there in the human case? Chromosomes, at least, are easy to count. Suppose for a moment that no crossing-over existed in the human population. Then each human diploid cell would contain 46 units with a genuine evolutionary role; chromosomes would be high-fidelity scaffolded reproducers, reshuffled into new combinations during sex and gradually diverging along their own asexual lineages due to mutation. But in humans there is crossing-over, and chromosomes are not passed on intact. Advocates of the gene's eye view argue that this forces us to recognize smaller genetic units as replicating entities.

Suppose we tried to follow this logic in a pure way, recognizing genetic units simply on the basis of the facts about crossing-over. We then run into the problem that crossing-over does not respect the boundaries between cistrons, or any similar boundaries. Crossing-over breaks and recombines genetic material, and the only boundary marking discrete units that cannot be broken is the individual nucleotide. Crossing-over does not break chromosomes entirely at random; some regions do not break, and other regions are "hot spots" where it occurs more often. Most crossing-over events will hit non-coding regions, just because they take up most of the genome. But crossing-over does not reshuffle discrete genetic units in a way that respects functional boundaries between them. So the facts about crossing-over determine the *length* of a stretch of DNA that is likely to persist for a given period of time, but they do not determine the division of a chromosome into definite segments of that length. The lengths can start and finish anywhere. They are like stretches of time, not card-like units that are picked up and rearranged. Crossing-over may give us "units" in a units-of-measurement sense (the *centimorgan*, in fact), but not in the building-blocks sense.

[3] In the arguments in this section I will, for simplicity, often not make use of the fact that many "genes" with a known evolutionary role are not cistron-like at all, but are regulatory elements which have, given a specific context, definite phenotypic effects (Moss 2003). These phenomena strengthen the argument. For the ever-increasing complexities involved in counting genes as our knowledge of the organization of genomes grows, see Griffiths and Neumann-Held (1999), and Griffiths and Stotz (2006).

Summing up this set of arguments: the analogy with teams of rowers is a misleading one. Rowers are organized countable units that remain individually intact as they are shuffled into new combinations, and they have reasonably consistent causal roles within their teams. No genetic element in a population like ours has that combination of properties.[4]

I will also approach the point from a slightly different direction. Advocates of the gene's eye view say, as noted above, that the size of a gene in an evolutionary context depends on the particular rate of crossing-over (Williams 1966, Dawkins 1982a). As the rate of crossing-over gets higher, the number of genes presumably gets higher—or at least, it gets higher until there is a kind of collapse. Suppose crossing-over occurred every couple of nucleotides in each meiotic event; there is an almost complete re-shuffling of the sequences on the two chromosomes. Then, I take it, genes as evolutionary units would not exist at all; there would be no genetic element between nucleotide and chromosome that was copied as a unit.

The average rate of crossing-over in humans is roughly two times per pair of homologous chromosomes, per meiotic event. (The number varies across chromosomes of different sizes.) How do we feed that number back in to calculate the number of evolutionary genes in humans? At this point, defenders of the gene's eye view say that the size of an evolutionary gene depends not only on the crossing-over rate, but also on the strength of selection. Williams said in his classic 1966 book that a gene is any stretch of DNA subject to a "selection bias equal to several or many times its rate of endogenous change" (p. 25). Endogenous change includes both mutation and crossing-over. But we then encounter the fact that "selection differentials" change as the environment, the overall genetic composition, and the behaviors found in the population change. Genes will fade in and fade out of existence—though without changing their intrinsic physical features—as various selection differentials get larger and smaller.

The question "how many evolutionary gene tokens in humans?" is turning out to be unanswerable—not because we do not know the facts well enough, but because there is no definite number to learn. There is a fairly definite number of human beings, human cells, human chromosomes, and human DNA nucleotides. There is also a rougher number of human cistron tokens. But there is not a definite number of evolutionary genes, and the only sketch of a calculation that has been offered would yield a number that has an obvious element of arbitrariness and would also change as selection pressures change.

On one hand, there is an agreed-on set of facts about meiosis, selection, and the structure of genomes. On the other hand, there is a standard way of talking about genes, as units that are passed on intact while doing things that affect their

[4] Except perhaps for Y chromosomes, or at least the main non-recombining part of them.

rate of replication. I am arguing that the relation between the descriptive habits and the facts that provide the grounding for the habits is less straightforward than often supposed. When looking at genes in a close-up and empirical way, people discussing gene-level selection often use language that acknowledges these facts. Burt and Trivers, in their extensive 2006 review, do not talk of "selfish genes" but instead of "selfish genetic elements," a phrasing that steers the reader away from the idea that these are discrete units of the sort posited in classical genetics. I hear their term "element" not as suggesting something elemental, but as referring to any *piece of genetic material* that has, given the local context, an ability to causally affect "its" reproductive rate in ways other than by fostering organism-level reproduction.

Dawkins himself argues that the sorts of facts raised here simply do not matter, as there is a harmless "elasticity" in the concepts of a replicator and evolutionary gene (1982a: 90). I agree that these facts do not matter much if one's point of view is sufficiently pragmatic. We can pick any piece of DNA found in some organisms in a population and note, given that specific context, its ability to causally affect its reproductive rate. When the context or our interests change, that bit of DNA will no longer be a salient-looking unit. But these considerations do matter if the aim is to give an account of the real entities that undergo the kind of change Darwin described.[5]

Within a more general discussion of genetics, Sterelny and Griffiths (1999) say that the word "gene" has become a "floating label" for any reasonably small stretch of DNA whose role is significant in the circumstances of a particular discussion. Whether this is generally true or not, it does apply to the case of evolutionary genetics, where the stretch of DNA that makes an evolutionary difference can be very unlike a classical gene (Moss 2003). At this point a defender of genic selection might argue that *all* the terms used to pick out alleged Darwinian individuals, including "organism," "cell," and "group," are floaters in this sense. Chapter 4 showed some buoyancy in the case of "organism," I accept, and looser collectives raise problems. But the problems with genes are more acute than the uncertainties found in the cases of organisms and cells; no one has argued that organisms fade into and out of existence as selection pressures change.

At this point it is useful to make a comparison with the status of populations themselves. There is some freedom, I argued, in recognizing the boundaries of

[5] Dawkins also addresses one of these issues directly. He says that crossing-over within a cistron will not usually break up a gene; only if the break is between two polymorphic sites will the old structure be lost. The rest of the time, it will be broken and then put back together (1982a: 90). This reply conflates *copying* with the mere fact of *reappearence* of a structure. A replicator is not faithfully copied if it is broken in half and then, owing to the luck of the gene pool, the same sequence is restored.

a population. Two evolutionary factors discussed earlier (sex and competition) affect how well "glued" a collection of individuals is into a natural unit. But as discussed in the previous chapter, it would be possible to recognize several small evolving populations and "stitch them together" to get a picture of evolutionary change in the larger collection. The situation with genes is different. It is not one where we first recognize a set of definite individual entities and then encounter some flexibility in how we collect them into groups for analysis. Here there is no clear inventory of the entities themselves.

So far the arguments in this section have been based on present-day facts about genes and chromosomes. To finish the section I will look at these issues from a more historical point of view. One part of the argument above can be summarized by saying that genes are only evolutionary units—to the extent that they are at all—as a consequence of crossing-over. That makes vivid the question: what is the evolutionary origin of crossing-over?

Crossing-over is one of three main kinds of exchange of genetic material. Bacteria engage in a kind of sex, "conjugation", which is severed from reproduction, by giving and receiving small packets of genetic material that may be integrated into a bacterial chromosome or carried around separately. (They also pick up stray bits of DNA that may be floating around, or gain them from viruses.) We can imagine a schematic history in which bacterial sex is initially the only kind of genetic exchange between organisms, until the appearance of the eukaryotic cell. In eukaryotes, the circular bacterial chromosome is replaced by some larger number of linear ones. And at some point, eukaryotes began to engage in cycles of haploidy and diploidy, requiring the successive doubling and halving of genetic material. At this point I will move the story forward within the framework of one particular hypothesis about the selection pressures that then took hold. This hypothesis is chosen provisionally, and because it throws some relationships into particularly sharp relief, not because I have special reasons to think it is superior to others. The hypothesis, developed by Haig and Grafen (1991), focuses on the role of intra-cell conflict.

Imagine a situation where an organism is cycling between haploid and diploid stages. A meiosis-like process forms haploid cells from diploid ones at a particular point, and it distributes half the diploid set of chromosomes into each haploid cell in the way familiar from ordinary meiosis, but with no crossing-over. This, Haig and Grafen argue, creates rich opportunities for destructive conflict. If a "killer" chromosome arose that could sabotage the haploid cell that the killer did not end up in, it would spread through the population. However, typical "driving" chromosomes, of the kind described in the previous section of this chapter, have two components, a destructive element and also an element that prevents the killer from destroying itself. In the driving complexes that are found these are tightly linked, and that is a clue to the hypothesis. If there was no crossing-over at

all in a situation with a haploid–diploid cycle of this kind, the resources of entire chromosomes could be turned to the evolution of devious driving mechanisms. A killer and its protection could be far apart on the chromosome but still reliably passed on as a unit, and they could have subsidiary devices working with them on the chromosome as well. Crossing-over prevents such large weaponed complexes, such "terminator chromosomes," from being possible, as it breaks up genetic associations on chromosomes except for those that are very physically close (or protected by an inversion). The complex and technical part of the Haig and Grafen hypothesis is showing that these facts will actually select for genetic elements that foster crossing-over. Here I will assume their hypothesis gives a *bona fide* possible mechanism for the evolution of crossing-over, whether or not it is the actual mechanism.

If this was how evolution went, the resulting picture would be like this. Genes, roughly speaking, are late-comers. They are products of complex evolutionary measures taken by cells to suppress what would otherwise be carnage at the chromosomal level. Before the advent of haploid–diploid cycles in eukaryotes, genes as evolutionary units did not exist. I said "roughly speaking" because of the complicating role of genetic exchange at earlier stages in evolution. A more exact way to put it is like this. Gene-like units only have an evolutionary role as a consequence of *some* process of shuffling, so that small genetic elements can be passed on independently from others. If all we have is the reproduction of whole bacterial genomes, or whole eukaryotic genomes for that matter, then we may have identifiable cistrons (and various regulatory elements) but we do not have genes as evolutionary units. Bacterial conjugation is one kind of shuffling; segregation of chromosomes is another; crossing-over is another again. And the evolution of crossing-over is what set small genetic elements free as evolutionary players. This is the "evolutionary transition" that gave us genes. Not all transitions make big things out of small ones; the evolution of crossing-over created new small things out of bigger ones.

People often think of gene-like things as early arrivals in the history of life. In some "RNA world" scenarios, they are the earliest of *all* arrivals. If so, they were then largely *lost* as evolutionary players, for at least a billion years. Here I have in mind the long interval between the origin of the bacterial cell and the origin of something like eukaryotic sex. That event gave small genetic elements a role on the evolutionary stage once again. The trouble they periodically cause for the organisms that contain them is part of the "cost of sex" for organisms (though a very different kind of cost from the twofold cost that arises from the presence of males).

Even within the strong empirical assumptions made to tell this story, there are qualifications. Bacterial conjugation and related phenomena involve pre-eukaryotic shuffling, and may have had an extensive evolutionary role

(Woese 2002). In the previous section I also allowed that a simple increase in copy number of a genetic element within a cell, as with transposons, does have a Darwinian character. That could, in principle, happen without shuffling. But sex, which makes the fragmentation of genomes routine, set small genetic elements free in a new way, and the in-principle possibility here is also one that is important: an evolutionary transition that unleashes small things which were previously only aspects or variables characterizing a whole.

7.3. Agents, Interests, and Darwinian Paranoia

> There are a great many important characters ... that are in the nature of collective attributes, all possessing the common quality of contributing to the welfare and survival of the group as such, and when necessary subordinating the interests of the individual. One of these is the reproductive rate. (Wynne Edwards 1962: 19).

> Four thousand million years on, what was to be the fate of the ancient replicators? They did not die out, for they are past masters of the survival arts. But do not look for them floating loose in the sea; they gave up that cavalier freedom long ago. Now they swarm in huge colonies, safe inside gigantic lumbering robots, sealed off from the outside world, communicating with it by tortuous indirect routes, manipulating it by remote control. They are in you and in me; they created us, body and mind; and their preservation is the ultimate rationale for our existence. They have come a long way, those replicators. Now they go by the name of genes, and we are their survival machines. (Dawkins 1976: 19–20).

Treating genes as evolutionary units is often allied with a version of the "agential" approach to evolution, the approach that understands evolutionary processes as contests between agents with goals and strategies. The idea that genes are units of selection is often expressed by saying that genes are the units of evolutionary *self-interest*—the true contestants battling on evolutionary timescales, the agents who organic adaptations are "for."

Overtly agential description of evolution is part of a larger family, or a graded series, of metaphorically loaded usages. At the extreme end we have talk of strategies and cabals. These shade into less tendentious talk of welfare and goals, and those shade into talk of costs and benefits understood directly in terms of components of fitness—chance of survival, number of matings, and so on. It can be unclear where metaphor ends and literal usage begins. Talk of this kind can also have several different intended roles. It may be seen as a metaphorical expression of a deep truth (as in Dawkins 1976), or merely as a practical tool for thinking about some complex matters in a simple way (Haig 1997).

A good way to approach the status of this talk is to put it into a context provided by parts of recent psychology and anthropology, work which also casts new light on the history of ideas (Medin and Atran 1999, Griffiths 2002). Some of this work was outlined in Chapter 1. It argues that when dealing with the living world, people naturally make use of a particular package of conceptual tools.

These include an essentialist causal model of organisms, habits of teleological thinking, and a willingness to explain events in terms of agents and agendas. These habits do not operate the same way in all domains. The biological world triggers it more than some others, though Aristotle's science shows the potential for very systematic application of some elements of the package, and in the face of complexity and frustration it can be applied more broadly.

Descriptions of evolutionary processes often feature a mix of populational concepts with teleological and agential forms of description.[6] Shifts in views about how evolution works will then be accompanied by shifts in how agential talk is applied. This phenomenon can be illustrated by looking at some famous debates about evolutionary processes in the middle to late twentieth century, summarized in the quotes at the start of this section.

A number of evolutionary writers in the middle of the last century freely invoked selection at higher levels than the individual organism, especially in explaining cooperation and various kinds of restraint. Eventually there was a reaction against this thinking, spearheaded by Hamilton, Williams, and Maynard Smith. They argued that commonly invoked mechanisms of high-level selection would not in fact be evolutionarily efficacious, as lower-level evolution would lead to subversion of cooperative groups even if such groups were, in some sense, better adapted than non-cooperative ones. The attention to genetic models of such processes led to the development of the gene's eye view itself. Towards the end of the century there was a revival of explanations in terms of multi-level selection, but in a more rigorous form, as in the models discussed in Section 6.2.

This process was accompanied by successive shifts in use of the language of agency and benefit. Such language often has a significant communicative role. When a student is told that the gene is the ultimate unit of evolutionary self-interest, for example, he or she is supposed to hear that as gesturing towards one family of evolutionary mechanisms—which can be more precisely described in other terms—and away from another. In the case of descriptions from a genic point of view, however, these formulations developed an unusual power and role. They became more than a shorthand, being used not just to summarize complicated ideas but to shape foundational descriptions of evolution. An example of this was discussed in Chapter 2. There I discussed a passage by Dawkins in which an agential picture of evolution was used to argue for a requirement that any process of natural selection contain long-term persisting entities of some sort. (You cannot get evolution by selecting between things

[6] Influence of the older habits shows up in psychological work on the understanding of evolutionary ideas in students, even those who have had extensive instruction (Lombrozo et al. 2006, Shtulman 2006).

when there is only one copy of each.) I argued that this observation shows the limitations of agential descriptions of evolution, not a requirement of long-term persisting entities. This is truly a case of the metaphorical tail wagging the scientific dog.[7]

How did the gene's eye view acquire such apparent power as a foundational description? I conjecture that this is because the gene's eye view of evolution is a special kind of agential narrative.

Two explanatory schemata can be distinguished, within the general agent-positing category, which have a special psychological potency. The first is a *paternalist* schema. Here we posit a large, benevolent agent, who intends that all is ultimately for the best. This category includes various gods, includes the Hegelian "World Spirit" in philosophy, and includes stronger forms of the "Gaia" hypothesis, according to which the whole earth is a living organism. The second schema is a *paranoid* one. Now we posit a hidden collection of agents pursuing agendas that cross-cut or oppose our interests. Examples include demonic possession narratives, the sub-personal creatures of Freud's psychology (superego, ego, id), and selfish genes and memes.

As the examples suggest, I think it is common for paranoid explanatory projects to posit small agents and paternalist projects to posit large ones. The tendency is not invariable, of course, as Satan and the angels attest. And while sometimes there *are* large and kind agents or small and vicious ones at work, the list of examples is intended to suggest that the psychological appeal of such hypotheses often far outruns their empirical warrant.

The transition between styles of explanation in biology was accompanied, I said above, by the exchange of one set of beneficiaries for another. Harmonious groups were replaced by selfish genes. But the new set of beneficiaries acquired too powerful a role, and one tradition of foundational description of evolution devolved into Darwinian paranoia.

My talk of "paranoia" in this context draws on the work of Richard Francis (2004). Francis argues that parts of contemporary biology have come to prize, above all others, explanations in terms of hidden *rationales* for biological characteristics. The biologist is induced to expect that there is some such rationale for nearly everything, and if we cannot find one that is a kind of scientific failure. Francis uses the phrase "Darwinian paranoia" more broadly and less psychologistically than I do. For Francis, pure adaptive thinking in biology itself tends towards paranoia, even without posits of hidden plotting agents. The concept of adaptation has a special intermediate status here, being

[7] Sometimes the direction of wagging is explicit: "The whole purpose of our search for a 'unit of selection' is to discover a suitable actor to play a leading role in our metaphors of purpose" (Dawkins 1982a: 91).

useable in both thick quasi-teleological senses and much more minimal ones (Lewontin 1985, Burian 1992). But there is a style of selectionist thinking in biology that I think does not involve any form of paranoia. This is the kind of investigation where someone asks: suppose a population was like this, and such-and-such a mutation appeared, what would happen to it? Thinking this way does not require the idea that genes are "ultimate beneficiaries" of anything.

The reproduction of genetic material is part of the reproduction of cells and organisms. The scrambling of genetic material is one of the consequences of sex. These facts support two kinds of Darwinian description of genes, one weaker and one stronger. The weak kind is supported by the fact that any collection of stretches of DNA located in the right places within organisms can be described as varying, passing on their differences in reproduction, and influencing their chances of being copied. The stronger kind is seen in the special cases where a gene proliferates through a process whose crucial steps do not involve a contribution to organism-level reproduction. Those processes are largely dependent on the machinery of sex. With the scrambling of genetic material comes the possibility of independent action, and selfish genetic elements are part of the cost.

During early and classical genetics, a "particulate" way of thinking about genes was undoubtedly progressive. The positing of Mendelian "factors," and then genes, which remain intact and pure across generations despite combining with other factors inside different organisms, was a huge advance. But as our knowledge gets finer-grained, talk of genes as units is slowly being replaced by talk of genetic "material"—a stuff, not a discrete unit—and by flexible talk of genetic "elements" when the causal roles of particular pieces of this material are under investigation. Known paradigm cases of evolution by natural selection depend on the high-fidelity copying of genetic material, but rather than being the clearest and most fundamental units of selection, genes themselves in most cases are marginal Darwinian individuals.