Chapter 10

Social Evolution

> Although *The Origin of Species* is developed in the crude English style, this is the book which contains the basis in natural history for our view.
> — Marx, *Letter to Engels*

10.1 Natural and Social History

Why does Darwin matter to anyone other than a professional biologist? The reason is that Darwin proved that we were not created in the image of God (which was the previous best theory) but are products of natural selection. He proved that we are not miraculous, in a religious sense, but that we are miraculous in a much more interesting, scientific, way.

However we must be careful to distinguish between what Darwin can explain and what he cannot. In general natural selection can only explain those traits that all members of a population have in common, not the differences between them. The commonalities will be due to selective pressures acting on populations over many generations, and the differences will be due to variations in the genetic make-up and developmental histories of individuals. For example, Darwin can explain why all humans have eyes but he cannot explain why *my* eyes are brown, which is due to my own particular genetic inheritance. Darwin can explain why all humans produce scars when their skin is cut but he cannot explain the scar on my chin, which was due to a surfing accident.

When we look at other animals we find that, in general, each species has a characteristic way of living. They form social groups, find their food and shelter, and raise their offspring in ways that are common throughout the species; and these patterns of behaviour are relatively static over the generations. But when we look at the thousands of distinct societies that humans have lived in over the last million years we see a vast variety. Some societies are vegan and others are cannibalistic. Some are violent and others are peaceful. Some are patriarchal and others are matriarchal. Some build palaces and cities while others do not build any permanent shelters at all. Some are rigidly divided into classes or castes, and others have no strict social hierarchies. Some force their children into labour or education outside the home from the age of six, while parents in others support and educate their offspring well into adulthood. Some are democratic, some are dictatorships, and some have virtually no social organisation at all. Some subsist by hunting and gathering in small bands, and others are organised into technologically-advanced empires that exploit the resource of

many continents.

Where did these differences come from? Are they dueto the effects of natural selection or to socialhistory? In chapter 8 I argued that only changes that are due tochanges in gene frequencies can be explained by the actions of naturalselection, and since humans form a distinct biological species (or, atleast, as distinct as any species ever is), with individuals more geneticallysimilar than they are distinct, then those traits that are due to naturalselection will *tend* to be those that are universal. But what on earthhas, for example, a stone-age hunter-gatherer got in common with a 21st Century computer tycoon? Well they have certain obvious biological similarities (two arms and legs, can suckle their young, and so on), they both use some form of language, are very good at learning and reasoning compared to other animals, and they are both (presumably) conscious. And that is about it. Darwin can explain the *natural* history of these human universals. All the rest, all the interesting fast-changed variety,must be explained by a theory of *social* history.

However there are two caveats that should be mentioned here. The first is that we have relatively little idea which psychological traits, if any, are universal and innate products of natural selection. This is the question that Evolutionary Psychology is attempting to answer. But the problem is that to determine whether a trait is universal and innate it is necessary to test a large enough sample of human subjects raised in a sufficiently wide range of social environments, from neolithic tribes to Western industrialised cities. Done properly this can yield fascinating insights into the common cognitive mechanisms that underlie our ability to form a wide variety of stable social groups — such as Cosmides and Tooby's investigation into our ability to detect social 'cheats' (Barkow, Cosmides, & Tooby, 1992). But even when a psychological trait is found to be universal this does not yet prove that it is innate. It is equally possible that it is what Dennett describes as a 'forced move in design space': 'so far as I know, in every culture known to anthropologists, the hunters throw their spears pointy-end-first, but this obviously doesn't establish that there is a pointy-end-first gene that approaches fixation in our species' (1995, p486). As I argued in section 8.3, in order to prove that a trait is amenable to adaptation by natural selection it is necessary to identify the brain mechanisms that are responsible for those traits and then discover whether development of those mechanisms is canalised and robust against environmental effects. Unfortunately we have barely begun to scratch the surface of these processes.

The second caveat is this. Natural selection does not usually have a direct impact on social history for the simple reason that the rate of reproduction of individuals is rarely correlated with genetic factors. The rate of reproduction of a Kenyan peasant is currently roughly twice that of a Western urban professional, for example, but this difference is due to social pressures not biological 'fitness'. Nonetheless, in some circumstances, genetic factors *can* play a significant role in history through their effects on rates of reproduction. The most important of these has been resistance to infectious disease. For example, the largest single factors in the defeat of the mighty Aztec and Inca empires by the Spanish Conquistadors were smallpox, tuberculosis, and measles. The pathogens for these diseases originally evolved from infections of Eurasian cattle, and over the millennia natural selection produced genetic resistance in Eurasian natives. Amerindians were not exposed to these pathogens until the Spanish invasions, and so had no chance to build up resistance. The result was fatal, both for individual Amerindians and their societies. The native population fell by 80% within 100 years, mostly due to Eurasian diseases, and Latin America

became a colony of Europe as a result.

But apart from such particular examples, social changes are rarely driven by changes in gene frequencies. They are not the result of natural selection, but social forces. In this chapter I argue that these social forces can be understood using an analogue of Darwin's theory of natural selection applied to the social domain.

Now on the face of it there is no *a priori* reason why a theory developed to explain biological evolution should have anything useful to tell us about the evolution of culture. The two domains are different and so surely different rules will apply? On the other hand Darwin's theory is completely general: as long as a system meets that conditions described in chapters 8 and 9, then adaptation will be the necessary consequence. Moreover natural selection is quite simply the best theory we have ever had at explaining the history of complex, highly structured, and reproducing systems; therefore if any aspect of cultural change fits a Darwinian framework then we may expect that the analytical rewards would be great.

The central issue in the philosophy of history is the relationship betweenideas and the material social environment. There are two extreme positions we can take. The first is a pluralist attitude usually attributed to Hegel, and the second is a reductionist attitude that is usually, though wrongly, attributed to Marx[1]. The Hegelian position is *ideological determinism*, according to which ideas are independent of the material world, and whose evolution is driven by their own internal dynamic. (Hegel also insisted on their steady progression towards some final goal, but you can be an ideological determinist without holding to this teleology). The not-Marxist position is *economic determinism*, according to which ideas are determined by the current state of the material social world. (The Orthodox Marxists, such as Stalin, also insisted that the materialstructures of society also undergo their own inexorable progress, but again this teleology is optional). In this chapter I defend a synthesis of these two caricatures — which I believe is closer to what Marx was 'really' getting at — and argue that ideas evolve through an an analogue of artificial selection such that they become adapted to a changing material social environment, even though they are not determined by it.

In the next two sections I argue that the evolution of ideas cannot beunderstood in isolation from the material social context. In section 10.4 Iargue that Marx's theory of history provides a way of understanding thiscontext that is analogous to Darwin's theory of natural selection. And in theremainder of the chapter I use this combination of Darwin and Marx to examineparticular aspects of the evolution of cultural traits.

10.2 Memes and Vehicles

One's theory of cultural evolution will depend on one's theory of biological evolution. For example Dawkins argues that natural selection is the result of the differential replication of genes. Some genes are better able to replicate in a particular phenotypic and environmental context — they will be fitter — and so they will spread more than others. Individuals or groups of organisms (vehicles) have no privileged conceptual or causal role in this picture. They are just the genes way

[1] It may plausibly be attributed to the orthodox Marxists of the Second International such as Kautsky and Plekhanov, Stalin, and Lenin (at least until the outbreak of the first world war forced him to revisit Hegel). But this is another question.

of making another gene. Dawkins also proposed an analogy with the social realm in which cultural entities (memes) play the role of genes. Some memes are better are reproducing in a particular social environment, and so will spread more than others. Persons, or larger social units, play no privileged causal role in this picture: they are just the meme's way of creating more memes.

In chapter 9 I argued that the underlying conceptual glitch with Dawkins' picture was what Sober called the counting problem: genetic replication, in itself, does not constitute successful reproduction. Fatter is not fitter. The same problem applies to the differential reproduction of memes. Suppose that an artist becomes obsessed with a particular idea and produces canvass after canvass of the same painting in his lonely garret, or that a monk becomes obsessed with repeating the same mantra in the seclusion of his cell. The meme is being replicated with every repetition but this process will not lead to it spreading: it will die with the artist or monk.

Dawkins solved the counting problem by measuring the spread of genes in terms of *germ line replication*, which he defines as those replicators that are 'potentially the ancestor of an indefinitely long line of descendants'. However the meaning of 'indefinite potential' is as undefined in cultural evolution as it is in biology. Through what process does a meme or gene gain the potential for immortality? The artist's canvasses may lay preserved and undiscovered in his garret long after his lonely death — like the cave paintings of Lascaux — but this does not, in itself, constitute success. There is more to becoming part of cultural evolution than just existing. Mere immortality is not enough for a meme (or for an artist): recognition is key. An unread book is just so much wood-pulp. An unheard speech is just hot vibrating air. In chapter 9 argued that a gene is only successfully (germ-line) replicated when it is carried by a new vehicle. And, in the same way, a meme is only successfully replicated when it lives on in the minds of others[2]. However it is not just individual minds that can carry memes, so too can larger social organisations such as religious organisations, media conglomerates, nations, classes, state structures, and so on. Each of these are capable of possessing memes, ideologies, or policies which they transmit to others, possibly even despite the particular views of their individual members. Thus the memes of the artist may impress their personal friends and colleagues, but they would be more effectively spread, and stand a better chance of surviving into posterity, if they were taken up by an influential gallery or magazine. In this case it would not matter if the individuals within those organisations did not think much of the artist's work, as long as the official policy of the organisation was to support them and so spread their memes.

The rate at which a meme spreads is its *fitness*, and when discussing the concept of fitness we have to distinguish two questions. The first is how to define and measure it: what does it mean to say that a gene (or meme) is fit? The second is to ask how particular genes (or memes) fulfill those criteria: what properties make a particular gene (or meme) fit? The answer to the first question is that the fitness of a gene is the ability of the class of vehicles that carry it (its *genotype*) to transmit to others. Genes are normally passed on by the reproduction of their vehicles, therefore fitness is usually operationally defined as the expected number of offspring of a individual member of a genotype. However, as we saw in section 9.4, what matters for the genes of a parasite is not the rate of reproduction of the vector *per se*, but the rate of transmission to a new host. Memes are more like infectious parasites than germ-line genes in this respect, since they too can spread

[2] Dennett also discusses this issue (1995, p349) but does not take it to its logical conclusion.

'horizontally' through a population of vehicles rather than just vertically through a reproductive lineage: the only way to give someone my brown hair is by giving birth to them, but I can give them my ideas and my bacteria by talking to them. Therefore if the definition of fitness is translated from genes to memes we find that it should be understood as the ability of the class of vehicles that carry it to transmit their memes to others.

Memes spread because their carriers are able to convince or persuade others to hold them. But some carriers are more persuasive than others. We are more likely to believe — i.e. accept the memes of — someone who we trust and respect. We are more likely to take notice of an artist whose previous work we enjoyed. Social organisations too can be more or less persuasive. An influential gallery is more likely to spread the work of an artist than an obscure one. Sociologists define this ability to spread memes as *ideological power*, and the concept of power is as central to sociology as force is to physics and natural selection is to population biology. In each discipline it is the concept that we use to explain *change*; the change in motion of a particle, the change in traits of organisms and, in the case of the social sciences, how we explain changes in people's social behaviour.

It may be useful at this point to develop an analogy that Dawkins (1993) draws between the spread of memes and the spread of computer viruses. One of the key problems that the designer of a virus faces is how to achieve effective transmission. The basic method is a form of memetic hitchhiking in which the virus is embedded in a set of files that contain useful data; when a user copies those files into their computer then the virus hitches a ride. However this is a rather haphazard method of infection that depends on the rate of individual transfers between computers, so hackers soon hit upon the idea of embedding the viruses in central resource locations that many users would access. In Europe until the late 1980s these central resources would usually be bulletin-board services (BBS) accessed *via* modems. The success of virus transmission would then depend on the popularity of the BBS: if it contained a lot of useful information that encouraged many downloads then the virus would spread rapidly. The most potent place to embed a virus today is an internet portal or search engine, or an application produced by a major software house. Even quite a 'dumb' virus placed in these locations can be very successful; consider, for example, the 'features' placed in Microsoft's code by disgruntled employees that produce derogatory messages about Bill Gates in response to certain obscure commands and key presses. 80% of PCs use Microsoft software and so carry these viruses. But the reason why they are so successful is not because they are intrinsically 'fit' — they are completely incapable of replicating themselves — but because they are carried by a vehicle that has the infosphere analogue of great ideological power. This is why advertisers buy banner space at popular web-sites: a meme carried by these vehicles will have a greater chance of spreading than one placed at a more obscure location.

The fitness of a meme is the ability of the class of vehicles that carry it to transmit their memes to others. This ability, as in the case of genetic fitness, is measured as the expected number of individual vehicles that they succeed in passing their memes/genes on to. This seems to directly contradict one the central tenets of Dawkins' account of memes:

> The most important point that Dawkins makes is that there is no necessary connection between a meme's replicative power, its *fitness* from *its* point of view, and its contribution to *our* fitness (by whatever standard we judge that). (Dennett, 1995, p363)

In other words a meme may spread simply because it is good at spreading — quite independently of the 'fitness' of those vehicles that carry them. The issue here depends on what we mean by 'our fitness'. If by 'our fitness' we mean our *biological* fitness then this is certainly true[3]. But if by 'our fitness' we mean our ideological power — i.e. our ability to transmit our ideas — then there is a necessary, even *analytic*, connection between memetic replication and the fitness of the social vehicles that carry them. (Though the question of precisely who benefits from memes, and in what sense, will be discussed below.)

10.3 Memes and Power

Fitness is the ability of a vehicle to pass its genes on to others, but it does not in itself *explain* that ability. In order to this we have to understand how those genes make the vehicles that carry them stronger or faster or more efficient, and so on. Similarly, the concept of ideological power describes the ability of a social vehicle to influence the memes of others, but in order to *explain* that ability we have to answer the second question of fitness, namely what properties makes a meme fit? Why do certain memes spread? In other words how do memes contribute to the ideological power of their carriers?

Ideas can contribute to ideological power in many ways. The most basic ideological role of a meme is the way in which it enables an individual to make sense of a common experience and communicate it clearly and convincingly to others. If an idea 'makes sense' then the person that carries it is 'sensible' and we take notice of what they say. Or ideologically influential memes may take the form of novel solutions to common problems, including advances in science and technology — the memes of flint tools, fire, agriculture, and industry spread through their evident potential to satisfy our material needs. Or memes can spread because of their ability to fulfill spiritual, aesthetic, or psychological needs. Thus the exercise of ideological power does not always involve foisting unwanted ideas on other people; it only requires that memes are transferred to heads that would not otherwise carry them. It is quite irrelevant whether or not those ideas are welcome.

Ideological power may come from the popularity of the transmitted memes themselves, but it can also stem from other sources of social power. Social power is defined generally as the ability of one social entity to affect the behaviour of another and, in addition to ideological power, Runciman (1998, ch4)[4] identifies two other forms of social power: *economic* power is the ability of one social unit to endow or deprive another of wealth, income, or resources; and *coercive* power is the ability of one social unit to exercise physical force over another. All these forms are interconnected and so, like physical energy, social power may be converted from one to another: economic power may enable one to buy arms, and so yield coercive power; coercive power may enable a nation to invade the oil fields of another, and so yield economic power; and both coercive and economic power enables one to take over television stations and so yield ideological power. All forms of power have to be taken into account when we are analysing the spread of memes in practice. For

[3] Though there are some exceptions. The reason why the meme of Catholicism spread to South America, for example, while the meme of sun-worship failed to infect Spain, is in large part due to the difference in mortality rates of their carriers.

[4] Runciman, following Weber, defines ideological power as that social power which stems from the ideological status of the social unit; but the current issue is the *effects* of social power. Its causes are discussed below.

example, in order to understand the evolution and spread of pop songs we must take into account the social dynamics of record companies, how they use their money and influence to get air-play, use advertising, associate their product with other ideological forces including fashion, sport, film, and politics, and so on. These tactics increase sales and profits, which can then be used to push more product in the future. Such explanations do not deny the role of the individual meme. Some pop songs are just plain catchy and so confer ideological influence on the artist; but the ideological power derived from the catchiness of an individual pop song is only one of the forces that play a role in shaping the evolution of popular culture. And unless we understand the role of all these social forces then the movements of the pop-charts remains a mystery.

Consider another example. Chong and Zanfarlin (1999) compare the spread of new technologies to the spread of infectious diseases. In particular they demonstrate, and explain, the existence of 'threshold' effects in which technologies only start to spread rapidly once a certain critical mass has been reached. This is a common characteristic of biological epidemics. For example, it is likely that the ancestors of HIV have been transmitted from monkeys to humans on many occasions. But it was not until the mid-20th Century, and the growth of a migrant labour force in African urban population centers, that a large and dense enough pool of infection was created sufficient to trigger an epidemic (Hooper & Hamilton, 1999). The disease had not necessarily become any more virulent but the change in conditions increased the ability of each host to infect others.

Chong and Zanfarlin argue that there are similar threshold effects in the case of new technologies. They find that the rate at which a superior new technology is adopted is not constant. Rather it depends on whether that superiority has been demonstrated by increasing the social (in this case, *economic*) power of those who currently adopt it; and this, in turn, may require an 'incubation' period in which the work-force is retrained, the work process is re-organised, and so on. Only once this stage has been passed will the new technology spread rapidly. Some technologies (such as non-qwerty keyboards or Betamax video) never reach this stage. They may be technologically superior but that superiority is never translated into social power, and so they fail to overcome the threshold.

In biological evolution there is only one force that matters, namely the force of selection acting on inherited characteristics. Of course other factors play a role, such as the impact of comets on the fate of dinosaurs or the impact of humans on natural habitats, but these exogenous forces only have an effect by defining the environment in which selection acts. Therefore *given* an environment then natural selection is an autonomous process. However one does not have to be a Marxist to notice that the evolution of culture is not autonomous from other, non-ideological, social forces. Payola has effects on the pop-charts and dictatorships can enforce ideologies. Maynard Smith (1961, p90) uses this fact to argue that historical processes cannot be analogous to natural selection:

> If it is true that any adequate theory of history must take into account all the causal connections [from material social environment to culture] then it follows that no formal analogy between historical and evolutionary processes as a whole is possible. In other words, I do not think that a helpful theory of history can be derived by starting from evolution theory and attempting to find historical analogues for the various entities and processes involved.

The effects of the material social environment on culture do not have strict analogies in natural selection, but they do have analogies in *artificial* selection. When a record company uses their economic power to push a particular record, or when a dictatorship uses its coercive power to enforce an ideology, they are acting like animal or plant breeders using their power of intervention to choose which traits will multiply and which will not. But, as Dennett argues (1997), artificial selection should be seen as a subcategory of natural selection rather than an alternative to it:

> The short legs of daschunds and the huge udders of Holsteins are just as much products of natural selection as the wings of the eagle; they just evolved in an environment that included a particular well-focussed selective pressure consisting of human agents. These phenotypes fall under the same laws of transmission genetics, the same replicator dynamics, as any others — as special and extreme cases in which the default 'randomness' or noisiness of selective pressure has been greatly reduced.

In the appropriate artificial environment even sterile organisms — such as mules, wheat hybrids, or seedless fruits — can flourish. If we apply this argument back to the case of cultural evolution then exogenous forces — such as payola and police states — define the social environment in which certain memes prove more popular than others. The question of where these forces come from and how they change is therefore outside the remit of the theory of memetic evolution. But this is a limitation on the theory, not a counter-argument. We should not expect a theory of narrowly cultural evolution to explain the origin of police states any more than we would expect a theory of natural selection to explain the origins of comets or plant breeding stations. Memetics can only explain the evolution of ideas *given* a particular social environment. Memetics is a micro-theory of historical change that makes sense only in the context of processes described using a wider macro-theory. Marx supplied one such theory, and in the next section I place memetic evolution within the context of Marx's theory of history.

10.4 Marx's Theory of History

Two kinds of things are needed for humans to survive and reproduce. The first are material things: tools, fields, hand axes, offices, bricks, computers, and so on. These are what Marx called the *forces of production*. Second, they need the appropriate social organisation: there must be someone to reap and someone to sow, these people must communicate and discuss and organise and control the process of production; and they need ideas and inspiration and reasons for acting. These are what Marx called the *social relations of production*, of which memes are one part. The forces of production are the hardware of society and the relations of production are the software; and Marx's theory of history describes the dynamic of the relationship between the two.

The basic idea is very simple. If a social unit — an individual, or nation, tribe, corporation, or kin-group etc — develops new, more efficient, forces of production then they will gain more social power, their relations of production will prevail, and may be propagated to others. As Torrance puts it (1985, p392):

> It is in the interest of some producer units to make experimental mutations of existing production relations to accommodate new productive forces. Some of these succeed better in enabling reproductions of these forces. ... Those units with new production relations that most successfully reproduce the new productive forces will be most

> successful also in reproducing their new production relations. A probable result would be that since those producer units survive while others fail, the production relations they embody tend to supplant both older ones and new but less successful ones.

For example, consider the politics of Europe after the first world war. In 1914 most of Europe was ruled by various forms of monarchy, from the semi-feudal aristocracy of Russia to the constitutional monarchy of Britain. And each of these regimes were supported by variations of the ideological theme of 'God, King, and Country'. However after the war those countries that had seen their productive resources decimated — most notably Germany, Russia, and the Austro-Hungarian empire — were thrown into political upheaval, the monarchies were overthrown, and other social structures grew in their place: the Weimar republic in Germany, Kerensky's provisional government and the Soviets in Russia, and Bela Kun's republic in Budapest, respectively. And each of these regimes were supported by variations on the ideological themes of democracy, socialism, and peace. However none of these regimes were effectively able to develop the productive forces necessary to rebuild their countries. These societies collapsed, though the reasons were different in each case: the reparations written into the Treaty of Versailles, the continuing civil war in Russia, and the lack of political experience and organisation in the Hungarian communist party. As a result these social structures gave way in turn to totalitarian regimes that put the entire economy under centralised state control, each supported by the memes of extreme nationalism: Nazism in Germany, Stalin's doctrine of socialism in one country, and Horthy's White Terror in Hungary[5]. It was the fact that these ideologies supported regimes that could develop the productive forces of their respective countries that ensured their success.

None of these social changes were inevitable. The particular histories of each country were the result of social conflicts involving political parties, classes, armies, and millions of people acting of their own volition — i.e. conflicts and changes within the relations of production. However, the wishes of an individual (or larger organisation) can only effect social change to the extent that they can wield social power, and the material basis of social power is the forces of production. Relations of production are not determined by the forces of production, but nor are they autonomous from them. The relations of production are *adapted* to the task of developing the forces of production, just as biological traits are adapted to the task of the survival and reproduction of the organism. This is summed up in what Cohen (1978, p160) calls Marx's *Primacy Thesis*:

> We hold that the character of the forces *functionally* explains the character of the relations ... The favoured relations take this form: *the production relations are of kind R at time t because relations of kind R are suitable to the use and development of the productive forces given the development of the latter at t*.

Now the idea that societies are functionally organised in a way similar to living systems is hardly novel. It is a truism that runs from Plato to Radcliffe-Brown (1952) and Malinowski (1944). But such theories have often been viewed with suspicion by historians since simply noting that a social trait serves a social function is not yet an *explanation* of why that trait exists. To say that religion sustains a certain social order, for example, may be true but does not imply that religion

[5] Italy showed the same pattern of social development, from the *biennio rosso* of 1919–21 to the rise of the *Fascisti*. But the reason for the initial social turmoil was more the weakness of the Italian state, rather than the economic effects of the war *per se*.

exists *because* it sustains a certain social order. The existence of an entity (such as some element of the relations of production) can only be explained in terms of its effects (such as the development of the forces of production) if those effects have contributed to the persistence of that in entity in the past. Marx's theory of history fills this explanatory gap — just as Darwin did for biological functionalism — by arguing that those social relations that are best adapted to the development of the productive forces are more likely to prevail over those that are less well adapted (Carling, 1993). Functionalism as a way of making sense of social systems only works if we can use those functions to explain the existence — i.e. the historical origins — of the traits we are trying to understand[6]. The fact that functional organisation is a product of historical processes — rather than simply an empirical fact, as the functionalists claimed — has three important implications.

The first is that Marx provides the framework for functional explanations in sociology, just as Darwin provides the framework for functional explanation in biology. But is still up the empirical historian (or biologist) to investigate the details of each particular case. Social relations are adaptations for harnessing and developing the forces of production, not determined by them. Relations can not be simply 'read off' from forces, therefore it is necessary to understand the particular history of adaptation that produced the current situation. For example, Darwin showed that those Galapagos finches that were better able to exploit the food sources of the islands would survive, but it took empirical investigation to find out which shape of beak would achieve this. Similarly, Marx showed how those social relations of production that were best able to exploit the productive resources of German capitalism would prevail, but it is up to the historian to investigate the particular historical reasons — and historical accidents — that meant Weimar failed and Hitler succeeded.

Second, a false dichotomy is sometimes presented between explaining history in terms of the intentional acts of individuals and in functional terms. For example, suppose we wanted to claim that the function, Y, of the doctrine of Divine Right, X, is to strengthen the social position of the King, Z. Elster insists we can only use this function to explain the existence of the doctrine, if it is *unintended*:

> the fact that some part of the relations of production X has the function Y is only explanatory iff
>
> 1. Y is an effect of X
> 2. Y is beneficial for Z
> 3. *Y is unintended by the actors producing X*
> 4. Y — or at least the causal relation between X and Y — is unrecognised by the actors in Z
> 5. Y maintains X by a causal feedback loop passing through Z
>
> (1983, p57, emphasis added)

However, as Dennett points out (1990)(1998b), intentional explanations should be seen as a subset of functional explanations, rather than as a distinct category. It does not matter if the

[6]Though there is some debate as to the strength of this requirement — see (Elster, 1980) (Cohen, 1980) (Wright, Levine, & Sober, 1992, ch2) and (Callinicos, 1987, sec2.3).

doctrine of Divine Right was originally a devious plan concocted by an ambitious chief, or was the product of dumb generate-and-test; the effects on history will be the same in either case.

The third point is this. Darwin identified the conditions, the *mechanism*, that causes species to evolve. When those conditions are not met, i.e. when the mechanism fails, then evolution does not occur. Similarly, Marx identified the mechanism underlying social evolution, but sometimes that mechanism can fail. Thus neither Darwin's theory of evolution nor Marx's theory of history are *teleological* in Aristotle's sense. Neither natural nor social history are pulled along towards some determinate final goal *a la* Hegel or Stalin, rather they are pushed along by the mechanism of descent with modification. When that mechanism fails then history stalls. So, although it is true that social relations of production prevail because they are capable of harnessing and developing new forces of production, this does *not* imply that there will be a monotonic increase in the development of the productive forces, as Cohen and many other 'orthodox' historical materialists argue (Callinicos, 1987, p23).

One way in which natural or social evolution can stall is by interrupting the generation of new mutations. For example, selection for sex ratios in cattle has no effect because variations in this trait are not compatible with the rest of the developmental biology of the cow: i.e. sex ratio is not a *dissociable* trait (section 8.4). The underlying point here is that biological traits do not enter into competition as atomistic parts, rather they come grouped in whole organisms with the development of each trait dependent on many others. Particular traits may be very useful, but unless they are compatible with the rest of the organism then they cannot be selected for. The same process happens in social evolution. Social relations of production come grouped together in what Marx called *modes of production*. For example, it may be very useful for a society to have scientists, craftsmen, other cultural specialists. But you cannot just have scientists on their own. They also need other people to produce food for them, they need training, a wider culture that values learning, and so on. If a scientist were plonked down in, say, a hunter-gatherer society then they would not contribute anything useful; they would just starve.

In some circumstances modes of production can become what Marx described as 'fetters' on the further development of productive forces. For example, Brenner (1986) has argued that the social structures of feudal societies prevented economic competition between peasant producers: it was not worth their while developing new productive techniques since any extra surplus would simply be stolen by the lord. In such situations 'economic growth in agriculture will, by and large, take the form of the multiplication of units on already existing lines', while 'the long-term developmental trend will be toward stagnation, if not crisis' [p28]. This produced the recession of the feudal economy in Europe, and the stagnation of China from 1300 to 1800. A similar process also led to Japan's abandonment of guns:

> Firearms reached Japan in 1543, when Portuguese adventurers armed with harquebuses arrived on a Chinese cargo ship. The Japanese were so impressed by the new weapon that they commenced indigenous gun production, greatly improved gun technology, and by 1600 owned more and better guns than any other country in the world.
>
> But there were also factors working against the acceptance of firearms in Japan. The country has a numerous warrior class, the *samurai*, for whom swords rated as class symbols and works of art (and as means for subjugating the lower classes).

> Japanese warfare had previously involved single combats between *samurai* swordsmen, who stood in the open, made ritual speeches, and then took pride in fighting gracefully. Such behaviour became lethal in the presence of peasant soldiers ungracefully blasting away with guns. In addition, guns were a foreign invention and grew to be despised, as did other things foreign in Japan after 1600. The *samurai*-controlled government began by restricting gun production to a few cities, then introduced a requirement of a government license for producing a gun, then issued licenses only for guns produced for the government, and finally reduced government orders for guns, until Japan was almost without functional guns again. (Diamond, 1997, p257)

The manufacture of guns was not compatible with the samurai shogunate, and so they were repressed; they were not a dissociable element of the overall mode of production. But this entire mode of production could not persist in the face of foreign competition. In 1853 Japan's safety in isolation came to an end when Commander Perry's U.S. fleet arrived bristling with cannons. This threw Japanese society into a crisis. It obviously could not compete without adopting Western military technology, but this very technology was incompatible with the structures of traditional Japanese society. The result was that within 20 years the Japanese feudal order was comprehensively re-ordered by the Meiji Restoration, with new forces (and relations) of production imported from the most advanced models of the West: the legal system was re-organised along French lines, and the army according to the Prussian model; the primary education system was copied from Germany; agriculture, banking, universities, and the postal service were modelled on the systems of the U.S.; and railways, telegraphs, the textile industry, business methods, and the Navy were based on the British. Western costumes, diets, architectural styles, and hair fashions soon followed (Hobsbawm, 1962, ch8). The social relations of production — and memes — that were most compatible with the development of the productive forces spread, while those that were not perished.

10.5 Memes and Symbols

Where do we get our ideas? Where do our ideas come from? There are two obvious sources. We can get ideas from other people or we can think them up for ourselves. But where is the line between the two? In the *Meno*, for example, Plato attempts to demonstrate that a young slave boy can be induced to see a truth of geometry on the basis of his innate grasp of certain mathematical ideas. The teacher, like a midwife, merely has to 'draw the knowledge out'. However in some cases the birth can be so laboured as to make the process more like implantation. Where is the line between simply telling the student the answer and creating the environment in which they can work it out for themselves? All teachers will be familiar with this dilemma. Now it is not obvious that there is a principled distinction to be made here; however it is one that is vital for a robust notion of memetic inheritance.

The problem is how to distinguish the effects of evolution from the effects of a changing environment. For example, suppose we progressively add fertiliser to a field over a number of years and watch the yield of the crop steadily increase. This is just a common response to a changing environment, not evolution. A meme can spread for the same reasons. For example there is no *a priori* reason why the memes necessary for making stone hand-axes should have had a single origin and then spread through communication. In an environment that includes hard fractured rocks

and food-sources that need killing and butchering, then it would be no surprise if many different groups of *Homo sapiens* thought of stone axes for themselves. Thus the meme may have spread because of a common response to a common situation, rather than through communication. In biological terminology this is the problem of distinguishing between analogous and homologous traits, and it is solved by examining the genetic basis of those traits. For example, all insects and mammals have heads at one end, and it turns out that this is due to a common set of *homeobox* genes. Therefore this common trait must have been inherited from a common origin. The eyes of insects and mammals, on the other hand, evolved separately, as revealed by their distinct genetic basis. The same applies to the social domain: in order to be able to claim that the spread of an idea is due to memetic evolution rather than the effects of a common environment we need to identify the units of memetic inheritance. Those units are *symbols*. Symbols are the genes of cultural evolution.

Genes share two properties with symbols that enable them to play the role of units of inheritance. The first is that they are discrete and atomistic. As Maynard Smith (1989a, p108) observes

> Dawkin's meme concept has been criticised on the grounds that an 'atomic' theory of culture is necessarily wrong. This may well prove correct, although I am astonished at the confidence with which it is sometimes asserted. Animal bodies show a far higher degree of coherence and functional interrelationship than do human societies, and yet an essentially atomic theory of genetics has had a lot to say about the evolution of animal bodies.

The reason why an atomistic theory of genetic inheritance is so successful is that *genes* are atomistic: they are replicated base-pair by base-pair. Similarly, in order to claim that one token meme is inherited from another then they must, in some way, be 'the same'. And in order that two meme tokens can be 'the same' even when they are in two different contexts, then the form of those tokens must be individuated independently of those contexts: i.e. they must be realised in symbols.

The second property that makes genes suitable as units of inheritance is what Monod (1997) calls 'gratuity', i.e. the fact that there is no chemical necessity why particular DNA sequences should fulfill particular biological functions. Similarly, Saussure argued that what distinguishes symbolic language from other representations (like pictures, or 'natural icons' such as footprints in the snow) is that the elements of language are arbitrary. There is nothing inherent in the string of letters W-A-T-E-R (or the equivalent phonemes) that relate it to H_2O. If we find that two cultures represent water using similar pictograms, such as parallel wavy lines, then this may simply be a common response to an environmental stimulus. It is coincidence, not inheritance. But when we find the words *agua* and *aqua* in different European languages then this is good evidence that they share a common origin — i.e. that these are inherited memes. Consider another example. A common defence against an accusation of plagiarism is to insist that the similarity between the original work and the alleged copy is due to common inspiration, rather than memetic theft: it is no surprise that artists in similar conditions will come up with similar ideas. The usual way of separating between the two hypotheses is to match symbol strings. Two musicians may come up with similar ideas, but if they came up with identical sequences of notes then this would be more than coincidence. Teachers expect to find similar ideas in the essays of students who have

attended the same lecture course, but suspicions would be aroused if they contained the exact same sentences. The similarity of content may be explained by a common response to a common environment, but the similarity of *form* may not.

Symbols allow us to distinguish between memetic inheritance and the effects of a common environment. But is this is purely an epistemological distinction? Surely an idea could be transmitted from one head to another even though they were expressed in different symbols, and thus showed no definite evidence of transmission? This is true, but epistemological considerations may be as important as ontological ones in social systems. In other words, the fact that a meme is *seen* to be inherited may be more important than whether or not it is.

The evolution of culture depends on the ideological power of social vehicles, but that power is measured in the spread of symbols; thus the power of the vehicle becomes invested in the symbol. This is the importance of symbols in the broad sense, including trademarks, icons, uniforms, brand names, and slogans. Symbols do not just express or represent a commonly held idea, but they also make their inheritance explicit. It is for this reason that ritual is so important for organised, as opposed to personal, religion. One can show allegiance to God without ritual, but it is through ritual that one expresses allegiance to His representatives on earth. When you make the sign of the cross you are not just assenting to the general idea of the holy trinity, you are also making explicit your allegiance to the church of which the cross is a symbol. In many cases the ritual can become more important than the ideas of which they are supposed to be an expression. Even the Gideon Society makes sure its name is prominent in every bible it leaves in a hotel room. Their aim is not simply to spread the meme of the Word of God *per se*, but to make sure that they are acknowledged as the organisation that is doing the spreading.

The importance of symbols in differentiating memetic inheritance from other environmental effects can also be seen at work in copyright and patent laws. The memes that such coercive systems are designed to cover have direct economic benefits, and so their carriers have an interest in either preventing their replication or ensuring that replication results in some payment to the originator. This is why there has been such strong pressure from powerful individuals, corporations and nations to extend intellectual property rights across the globe.

The relationship between the carriers and originators of such controlled memes is akin to that of the workers and queens in colonies of eusocial insects. Queens pass on their genes to the workers of the colony but the latter are a genetic dead-end: only the queen can pass on her genes to a new generation. Therefore the function of the colony system is to ensure that the benefit derived from those genes accrues in the fitness of the queen, measured by her ability to pass them on to new colonies. The workers in the colony thus act ‘altruistically’ in order to increase the fitness of the queen. Similarly the owner of a copyrighted meme will pass it on to many users, but those users will be prevented by the terms of the agreement from passing it on to others in turn: they are a memetic dead-end. Therefore the function of the copyright system is to ensure that at least a portion of the benefit of the meme accrues to the originator, either in the form of direct economic payment or through an increase in their ideological power — hence copyright statements of the form ‘*X* asserts the right to be acknowledged as the author of this work’.

10.6 Lamarckian Inheritance and Signalling

It is often said that a fundamental difference between memetic and genetic inheritance is that the former is Lamarckian whilst the latter is not; i.e. we can pass on acquired memes, but acquired traits are not passed on through germ-line genes. However the situation is not as simple as all that.

The possibility of Lamarckian inheritance is often regarded as an anathema to the modern synthesis. Dawkins, for example, admits that he is *scared* of the possibility:

> I use the word 'scare' because, to be painfully honest, I can think of few things that would more devastate my world view that a demonstrated need to return to the theory of evolution that is traditionally attributed to Lamarck. It is one of the few contingencies for which I might offer to eat my hat. ...Naturally any scientist hopes that the truth, whatever it is, will out. But a scientist is also entitled to his innermost hopes as to what that truth will turn out to be — a revolution in one's own head is bound to be a painful experience. (1982, p164)

Why does Dawkins regard the possibility of Lamarckian inheritance as such a threat to his understanding of biological evolution? The reason is that Lamarckianism threatens the Central Dogma on which the Weismann diagram, and hence Dawkins' gene-centrism, rests (sections 9.1 and 9.3). If characteristics acquired by the vehicle through development were inherited then vehicles would play an ineliminable causal role in selection: they too would count as replicators, not just genes.

However it is interesting to note that Darwin himself — because he took organisms to be the units of selection — saw no fundamental contradiction between Lamarckianism and natural selection. Indeed he saw in Lamarckianism a possible solution to the problem of blending inheritance in which, in the absence of a discrete genetic mechanism, phenotypic variation in the population would tend to be blended out by sexual reproduction; but if acquired characteristics *were* inherited then variations in the environment would produce corresponding phenotypic variations in the population which would then be inherited. Darwin included this very possibility in the later editions of *The Origin of Species* (Desmond & Moore, 1992, p567). Thus Lamarckianism is a threat to gene-centrism, not the theory of natural selection *per se*[7]. Natural selection only requires 'descent with modification', and there is no *a priori* reason why Lamarck could not provide the mechanism of that descent.

There are plenty of biological processes which are pseudo-Lamarckian; i.e. in which acquired characteristics are inherited, though not through germ-line genes. Obvious examples include the viral and bacterial flora that we inherit from our mothers. Moreover these pseudo-Lamarckian processes have a significant effect on natural and artificial selection. Young calves, for example, are often treated with large doses of broad-spectrum antibiotics, which boost yield but often leave them with very weak immune systems. The inheritance of memes from our parents is strictly analogous to such biological pseudo-Lamarckian processes.

However there are also ways in which acquired characteristics can affect the germ-line, and so be Lamarckian in a stronger sense. Consider sexual selection. When an animal acquires a mate it is thereby acquiring a new part of its extended phenotype. Moreover in proactively *selecting* a mate it is thereby choosing a new pack of genes to shuffle into its germ-line. And this selection is based on the phenotypic effects of those genes: if their genes seem to improve the fitness of

[7] Thanks to Dan Dennett for this point.

the candidate mate, then it is likely to confer a similar benefit on future offspring. Thus sexual selection involves the directed inheritance of acquired characteristics.

Zahavi (1975)(1977) has argued that sexual selection is a special case of the more general phenomenon of the evolution of signalling. Animals cannot always judge the fitness of possible mates directly, therefore they select on the basis of other observable traits which thereby act as signals which advertise the availability of a fitness resource. However in all such signalling situations there is a pressure to cheat. If an animal invested in the signalling trait even though it was not actually fit then it would gain a reproductive benefit. If the density of cheaters reaches a certain point then the signalling system becomes effectively useless. Zahavi's solution is that if signalling traits are expensive to produce then only the fit can carry them. It takes a lot of energy to produce and carry a peacock's tail, therefore the males with the most impressive ones must be well-equipped in other, less directly observable, ways. Waste can be a signal of quality: if the production of signals involve a handicap, then honesty is enforced. Grafen (1990) has proven the stability of such handicap-signalling systems using game-theoretic techniques (Maynard Smith, 1982)(1991), and Bullock has generalised these results to less analytically-amenable situations using evolutionary simulations (1998).

Sexually-reproducing animals may choose who to acquire their germ-line genes from, and this leads to the evolution of expensive signalling systems. There is an exact analogy here with memetic evolution. We are selective about who we accept our memes from. We do not believe just anybody. Therefore there is a pressure on social vehicles to convince others of the fitness of their memes — i.e. their ability to confer social power on their possessor — which leads to the evolution of signalling systems. The function of these signals is to convince others of the value of the memes (or relations of production more generally) of those that produce them. The most obvious example of this is conspicuous consumption (Veblen, 1899). Social-evolutionarily speaking, the point of being rich is to gain social influence. Therefore there is no point in being rich if no one knows that you are. (Reclusive millionaires are an historical dead-end). One way to gain social influence from riches is to *demonstrate* that you are rich; and conspicuous consumption is a way of doing this that is hard to fake. The cold-war space race can be seen in the same light. The ability to send rockets into space was a way of honestly signalling the technological muscle of the superpowers in a situation in which their destructive muscle — i.e. their nuclear weapons — had to be kept secret. You can't fake a Sputnik or a moon landing. The purpose of this phenomenal waste of resources was not simply to let the world know of the power of the USA and USSR, but was to thereby to convince other nations to accept their memes and relations of production (i.e. to join the western free-market economies or to become part of the Soviet bloc, respectively).

The problem of honest signalling is acute for the advertising industry. It is not enough to tell the audience that 'Omo washes whiter' since they have no reason to believe that the advertiser is telling the truth. The obvious solution is to act as free-loaders on other signalling systems which the audience trust. Science, for example, has a reputation for honest signalling which is ultimately due to the peer-review process. It is hard to cheat in science (i.e. spread memes which do not accurately reflect empirical results) because the fitness of those memes may be independently tested. Therefore if a scientist (or at least someone in a white coat who looks like a scientist) tells the audience that Omo washes whiter then they are more likely to be believed.

Another form of honest signalling is style. It is hard to fake style since, almost by definition, to look stylish is to be stylish. Thus corporations who produce intellectual property (from computer games to fashion) will invest heavily in associating their products with style icons, either through product placement or through incorporating those styles into their advertisements, in an attempt to convince the audience of the ideological benefits of the meme-product. The advertised product hitchhikes on the success of its advertisement.

Sexual selection enables suitors to test the suitability of their prospective mates according to some visible criterion. However all teachers know the problem of setting tests: they end up testing the ability to pass tests, rather than the ability that the test was supposed to be a measure of. The same occurs in sexual selection. The peacock's tail originally evolved as an accurate measure of fitness. But once the peahen evolved to select on the basis of that criterion then it is in the interest of the peacock's genes to generate bigger tails, rather than increase 'underlying' fitness — i.e. its efficiency at finding resources, and so on. (Underlying fitness may be thought of as the fitness a genotype would have if the effects of sexual selection were factored out.) Similarly the space technology of the USSR ended up as the most advanced in the world, even while the rest of their military aerospace industry lagged behind. The signal was still honest — the USSR really did have the best space industry in the world — but it became misleading.

10.7 Memes and Adaptation

If memes are adaptations, then who are they adaptations for? Who do they benefit? Dawkins warns us that it is a mistake to assume that biological traits must benefit the organism that carry them, and argues instead that a trait must benefit the germ-line genes through which it is inherited. This is true but does not yet answer the question; it just generalises it. Germ-line genes are always germ-line genes *of* a lineage of reproducing vehicles, and unless we identify those vehicles then the question remains unanswered. Who benefits from the diarrhoea induced by cholera, for example? It is the parasite, not the host. They are the originator and propagator of the trait, and so they must be the beneficiary if anything is. The same problem occurs when determining who benefits from memes (i.e. who gains social power from them). Just because a person carries a meme does not in itself prove that is an adaptation *for* them. Again, the key is to discover the social *origins* of the trait. As Marx put it:

> The ideas of the ruling class are in every epoch the ruling ideas: i.e. the class which is the ruling *material* force of society is at the same time its ruling *intellectual* force. The class which has the means of material production at its disposal consequently also controls the means of mental production, so that the ideas of those who lack the means of mental production are on the whole subject to it. ... The individuals composing the ruling class possess among other things consciousness, and therefore think. Insofar, therefore, as they rule as a class and determine the extent and compass of an epoch, it is self-evident that they ... rule as thinkers, as producers of ideas, and regulate the production and distribution of ideas of their age. (1846, p64)

Now the power of this argument has often been overstated: the evolution of culture is not one big conspiracy by an omniscient and omnipotent ruling class. You can't fool all of the people all of the time. Nonetheless the dominant memes in society will *tend* to be those that maintain the

ideological power of those that most successfully utilise and control the productive forces, rather than those that increase the social power of all those that carry them. Who gains from the meme of patriotism, for example? *Dulce et decorum est pro patria mori*[8], but it was *dulce et decorum* for the rulers of our country, not for those individuals that died in the trenches.

The question of who benefits from a meme is an empirical question that can only be answered by examining who gains social power from its propagation. This may be the person who carries it. It may be the ruling class of a society. Or, in some cases, it may benefit nobody. For example, an office I worked in once became horribly obsessed with humming the theme from *The Dambusters*; a meme that I managed to infect a new office with when I changed jobs. No social vehicle benefited from this meme, it was just catchy; a meme that spread solely because it was good at spreading, like so much selfish DNA. It is a neutral mutation, as far as social fitness is concerned. Dennett (1999), following Dawkins, uses such examples to argue that:

> in the domain of memes, the ultimate beneficiary, the beneficiary in terms of which the final cost-benefit calculations must apply is: the meme itself, not its carriers. This is not to be heard as a bold empirical claim, ruling out (for instance) the role of individual human agents in devising, appreciating and securing the spread and prolongation of cultural items. As I have already noted, the traditional perspective on cultural evolution handsomely explains many of the patterns to be observed. My proposal is rather that we adopt a perspective or point of view from which a wide variety of different empirical claims can be compared, *including the traditional claims*, and the evidence for them considered in a neutral setting, a setting that does not prejudge these hot-button questions.

The question that the memes-eye perspective raises is which social vehicles, if any, benefit from a meme; and the yard-stick that it uses for measuring that benefit is the rate at which those vehicles are able to spread their memes, i.e. their social ideological power. But this is *not* equivalent to claiming that the ultimate beneficiary is the meme itself. The reason for introducing the Darwinian terminology of adaptive benefit is to ground functional explanations of social history. But a meme, on its own, plays no role in history. Memes only have effects through the activity of social agents. This activity may result in the further spread of those memes (i.e. the exercise of ideological power) but it may also have effects through the exercise of *other* forms of social power, including economic and coercive. Therefore the patterns of cultural evolution can, in general, only be explained by extending our Darwinian analysis to include those wider processes.

10.8 Conclusion

In chapters 8 and 9 I argued that Darwinian evolution occurs whenever genes are differentially replicated according to the fitness that they confer on the vehicles that carry them. The analogical process in the social domain occurs when cultural entities are differentially reproduced according to the ideological power of the social structures that carry them. However this ideological power may *also* stem from other forms of social power, including the economic and coercive. Therefore a Darwinian analysis of social history must go deeper than the level of memes, and should include

[8] 'It is sweet and honourable to die for one's country'. This was a British patriotic motto of the first world war, brilliantly satirised by Wilfred Owen in his poem of the same title.

the evolution of the relations and forces of production that define the environment in which memes themselves evolve.

Marx argued that relations of production survive and spread to the extent that they enable social vehicles utilise and develop the forces of production. This is the basis of a Darwinian mechanism in social evolution, and it makes functional explanations of society possible. Memes are one part of the relations of production and so they too will become adapted to the forces of production.