

Is morality a unique, humanising
quality, or is it just another product of
evolution?

To discuss the evolution of morality, and whether
humans are a special, or indeed unique, case.

Nicholas Giles
Single Honours Philosophy - Long Dissertation
Van Mildert College

Contents

1	Introduction	1
2	The evolution of morality in all things	3
2.1	What is morality?	3
2.2	How can evolution produce altruism?	4
2.3	Why would “selfish genes” be altruistic?	5
2.4	How can altruism to unrelated individuals be fitness increasing?	7
2.5	How do unrelated individuals manage cooperation?	8
2.6	Do we see altruism in the natural world?	9
2.7	Is there therefore morality in the natural world?	11
3	How do humans implement morality?	13
3.1	What are the human characteristics that put them in a different moral realm?	13
3.2	How does humanity’s rational nature affect their moral nature?	13
3.3	Are humans motivated by moral sentiments, and if so, what is their source?	15
3.4	Does the development of culture bring anything to human morals?	16
3.5	Could memes play a part in human moral behaviour?	17
3.6	Why is morality so prevalent in humans?	19
4	What are the differences between human and animal implementations of morality, and what are their implications?	21
4.1	What features of human consciousness distinguish us morally from animals?	21
4.2	What meta-ethics supports this view of morals?	23
4.3	Why do agents act morally or immorally?	24
4.4	How have humans as social animals exceeded the original designs of morality?	26
4.5	What level of control do we have over our morals?	27

4.6	What are the key human modifications of morality?	28
5	Conclusions	29
5.1	Is morality the result of evolution?	29
5.2	Are humans a special case of morality?	30
5.3	Final statements	30
6	Bibilography	31

Chapter 1

Introduction

From the very beginning, people have been unwilling to apply evolutionary theory to humans. Darwin himself kept mention of humans out of his *Origin Of The Species*, devoting the much later book *The Descent Of Man* to it when he could no longer avoid the implications. The recent movement known as sociobiology seeks to treat humans as just another animal to be studied by biology, with their own special attributes and qualities, but no ‘better’ or ‘higher’ than any other. Of particular interest to philosophers is the study of philosophy of mind in these terms, and the suggestion that we can gain significant insight into our own behaviours by examining relatives, such as the primates. The opposing camp seeks to separate humans from the rest of nature, classifying us as special. Their reasons may be that they do not in fact consider us to be products of evolution, as some theological doctrines still maintain, or they may simply consider that we have progressed too far since our ancestors to learn anything about ourselves from them. I disagree with both ideas, and maintain that we are definitely the result of evolutionary processes, as biology seems to have proven, and also that the differences between ourselves and our ancestors are significant, but not so significant that we have nothing to learn from them.

The particular focus of this project will be an evolutionary study of morality. I will examine how morality could have evolved, and whether there is in fact any evidence that it has in the animal kingdom. I will try to combine the concept of morality, as exemplified by altruistic behaviour, with evolutionary theory, and see whether the two are compatible. I will then look at how humans are affected by morality, and how they have tailored it to their own needs. Particular human traits will be examined to see how they have led us to have a concept of a binding morality. Finally I will examine the differences between the behaviours of animals and humans, and see whether there is any significant reason to claim that only humans can act morally,

and whether in fact humans do act morally.

Before beginning the study itself, one point in particular needs to be explained. Throughout this project, I am using altruism, morality, and cooperation as roughly the same concept. Whilst I find the conflation of cooperation and altruism to be relatively unproblematic, each cooperator giving a little of themselves for the benefit of others, morality is much more than simple altruism. However, in order to simplify the comparative study of humans and other organisms, I decided to choose the paradigm example of moral action, and in fact the one that requires the most explanation. Action that is both morally good and to the direct benefit of the individual concerned is not only rare, but unproblematic to explain in terms of motivation. On the other hand, altruism is very difficult to explain if one considers self-interest to be a prime motivator, as one must in evolutionary theory. The vast majority of morally good actions are altruistic ones, ones that are not done for the benefit of the individual. If these can be explained with the help of evolutionary theory, then that is the biggest hurdle cleared. The issue of whether altruism can be classified as morality is addressed towards the end of the project, and the answer found to be in the positive.

Chapter 2

The evolution of morality in all things

2.1 What is morality?

Morality is a way of living, a set of rules for life. Whether you consider it to be the dictates of a deity, a part of the structure of the universe, or something that your parents made you think, it is still a set of rules for living. Morality tells you what is acceptable, and what is bad. It encourages certain actions, discourages others, it seems to be the central set of rules of our life. Whilst we may have others, such as what it is that we want, morality always seems to be the limiting factor. We can't take the money at gunpoint, because that would be wrong. We must be nice to our friends, because that is good. Morality is the set of rules, dictated or created, by which we decide what to do. And it allegedly makes us nice. Anyone who consistently refuses to follow the rules of society is deemed a threat, and removed. The immoral are dangerous, only the moral can enjoy society. And this is because morality is a control mechanism, it is the means by which we can even have a society. It promotes altruism and tries to limit selfishness so that we are able to cooperate to produce the fruits of civilisation. Most commentators agree that if we were without morality, civilisation would not exist. Even those who see it as God's divine law say that without it we would fall apart. The core of morality is the promotion of altruism, the giving of the self to benefit others. It would seem to make no sense that a person should benefit others at cost to themselves, yet that is what morality encourages. Therefore, it is altruism that is the primary focus of this investigation, the aspect that needs to be explained.

In order to spread the scope of this investigation as wide as possible, I

will use a definition of altruism that doesn't talk of motives, only of observable costs and benefits. Whilst there are many who will object that this is not altruism, I will address such issues later. The important distinction is between psychological altruism and biological altruism. The common sense of the term altruism is psychological, where someone intends to do something that benefits another individual, but this is only applicable to creatures with intentionality, so the definition I will use is:

“An entity is said to be altruistic if it behaves in such a way as to increase another entity's welfare at the expense of its own.” [Daw99, pg 4]

This is biological altruism, and “welfare” is usually defined in terms of reproductive success. Thus, an organism which increases another's chances of successfully reproducing, at cost to its own chances of doing so, is being biologically altruistic.

2.2 How can evolution produce altruism?

Evolution is, at its most basic, a synthesis of two concepts - natural selection and population variance. If one considers there to be different characteristics within a population, and competition for some resources, then those set of characteristics that are most successful at getting the resources will thrive at the expense of less successful characteristics. The standard formulation is that individuals who are more successful will have greater reproductive success than those who are unsuccessful, and thus more of the successful characteristics will be passed on to the next generation. If we wish to explain how altruism could have evolved, this seems to be a very large obstacle. Organisms that reduce their own chances of reproductive success are going to lose out to those that do not, yet that is exactly the quality that we want to see. However, there are alternative formulations of evolution, based on different units of selection, that claim to be able to explain altruism.

The first of these is group selection. This is a theory that agrees with the idea of competition for resources as in standard evolutionary theory, but asserts that the competing units are groups, and that the individuals within these groups do not compete, but rather cooperate, thus benefiting the group. Under this system, there is no conflict with individuals being altruistic, since that is what they should do. It would explain our tendency to live in groups, and why we are so keen to be altruistic. However, there are two large problems with this theory. The first is that there is no evidence that populations work in this manner, or that individuals do. Whilst the theory is coherent, given certain assumptions, those assumptions are unfounded, and most evolutionary theorists discard the idea of group selection. The

other large problem is of particular note to the question of altruism. Given that you had a population of altruists cooperating with each other against other groups, if just one individual had a mutation that caused it to be selfish, then it would benefit from the altruism of the others, and from its own selfishness. It would thus prosper, and produce more offspring than its altruistic neighbours. This would continue until the whole population was selfish, and that is on the assumption that a wholly altruistic population existed to begin with. Thus, “selection is ... for the group only if it is for the individual.” [Rus79, pg 14]. Group selection is a wishful alternative that doesn’t work.

The other alternative is to see the competing units as the genes of organisms. The idea, often termed ‘selfish gene’ theory, is that since it is the genes which persist through replication, it is in fact the genes that are competing. Richard Dawkins, one of the original advocates of selfish gene theory, calls organisms “gene survival vehicles” to emphasise that it is not the interests of the organism that are to be served, but those of its genes. After all, it is the genes that are the more direct descendants of those initial successful replicators. We are just a more versatile cell membrane. Given that there could be a split between what benefits us and what benefits our genes, it would seem possible for altruism to have crept in. However, there would still need to be a great deal of explanation of how it can be beneficial for the genes.

2.3 Why would “selfish genes” be altruistic?

Having determined that we can no longer assume that it is the interests of the individual that must be served in order to be successful, we need to examine what it is about altruism that could be of benefit to genes, when we must now consider that it is they who need to maximise their fitness. The first point to note is that genes are not unique, as we suppose individuals to be. They make copies of themselves all the time, whether in standard cell mitosis, producing more cells for a body or as asexual reproduction, or in meiosis, selecting half the alleles for inclusion in germ line cells, for sexual reproduction. This means that there are lots of copies of the genes in existence, and reproduction by any of them is reproduction for all of them. This realisation led sociobiologist E.O. Wilson to coin the term “inclusive fitness” to describe the fitness of all copies of a gene, and by association, an individual and all their relatives. Inclusive fitness measures the probability of successful reproduction of a certain gene line, so can encompass an entire family group’s reproductive chances. Once one realises that genes ‘stretch’

between bodies, altruism at the level of the individual can be selfishness at the level of the genes. A parent who looks after their child is being biologically altruistic at the level of the individual, since they could be spending that time and effort producing more offspring, but are clearly being biologically selfish at the level of the genes, since if they don't look after the child, it has less chance of surviving and reproducing the genes it carries, which is the reason that the parent produced it in the first place. Therefore, the strategy of helping those individuals who are related to you is a very good one from the point of view of the genes, if not for the individual concerned. A precise definition of how altruism between relatives can be profitable is provided by Michael Ruse:

“altruism towards relatives is worthwhile if the ratio of gain to loss in fitness exceeds the reciprocal of the average co-efficient of relationship of benefiting relatives.” [Rus79, pg 44]

Thus, the more closely related, the more aid an individual is likely to give. As an Arab proverb states, “My brother and I against our cousin. My cousin and I against the stranger.” [Rus95, pg 246] The benefits that the offspring, and other relatives, receive from the strategy allow them to be more successful, and thus reproduce more successfully, and therefore, the genes of the individual, including those that incline their carrier to be altruistic towards their relatives, will be more successful than the genes of the other individuals who refuse to help their family. One could challenge how such a strategy could begin, since it would seem to need to spring fully formed from nowhere, but not only can one simply point out how ubiquitous it is among animals now, but it can be seen how it could start slowly with small benefits, which natural selection has the ability to take advantage of. A challenge to this is much like the challenge that the eye cannot be a product of evolution since it would not work except as a whole, and it can be defeated in much the same way.

Thus, the ability of individuals to discriminate their relatives and be altruistic towards them is an advantageous trait, and one which does not require a leap of faith to begin it. Therefore, from this explanation of the origins of the process, and simple observation of parental care and relation care in the natural world, we can see that there is an initial group of cooperators. However, one must remember that there are still conflicts of interest in the family, we are not talking about a version of group selection. Cooperation only occurs because it is in the interests of the genes of the individuals involved, not because the group demands it. It may well be that groups are another level of advantage for genes, but that does not lead us to group selection. Additionally, altruism does not exist solely between related individuals, so there must be some other mechanism at work.

2.4 How can altruism to unrelated individuals be fitness increasing?

When dealing with cooperation between non-relatives, there is no way that the genes can simply be helping out other instances of themselves, so there has to be a wider-reaching explanation. It is not, as one could argue from group selection, that individuals feel a bond toward their species, and will do what they can to benefit them, since as we have seen that is easily subverted. Rather, there must be some reason why helping others can be profitable. The answer lies within the field of maths known as game theory.

Game theory is based around the prisoner's dilemma, a logic puzzle that has been in existence for a long time, and is a good simplified version of any interaction where two individuals have the option to cooperate or not. The prisoner's dilemma is a model of two individuals (A and B) who each have the option to cooperate or defect, producing a different result depending upon a function of their action and the action of their counterpart. Their goal is to derive the most benefit for themselves. The four possible outcomes (from the point of view of A) are:

- 1) A cooperate and B defect=A gets 0 (sucker's payoff)
- 2) A cooperate and B cooperate=A gets 3 (reward)
- 3) A defect and B defect=A gets 1 (punishment)
- 4) A defect and B cooperate=A gets 5 (temptation)

B gets the appropriate points for their actions in relation to A's. The precise values of each outcome are unimportant, so long as they obey the rule

'Sucker's Payoff ; Punishment ; Reward ; Temptation'.

If we examine the possible outcomes, it seems that defection is the best policy, since in any case where A cooperated, they could have done better by defecting. If B cooperates, then A gets 3 points from cooperation, but 5 from defection, and if B defects, A gets 0 points from cooperation and 1 point from defection. There seems to be no reason for either party to cooperate, and thus "the only strategy that can be called a solution to the game is to defect." [AH] However, this is only the case in single interaction models, as if one extends the number of interactions, one can see how cooperation may be a more effective method than defection. In the iterated prisoner's dilemma, each party is presented with the same options, but also has information on the past performance of their opponent, if there is any. Under

such a system, defection is a viable strategy, but cooperation can be as well. The system known as Retaliator, or Tit-For-Tat, is one that suggests beginning by cooperating, then copying your opponent's last move. Thus, when someone is defecting on you, you defect on them, getting the most out of the situation. However, if a cooperator is your opponent, you cooperate, getting more. When Robert Axelrod ran a computer simulation of the iterated prisoners dilemma and challenged people to supply the strategies, Tit-For-Tat did indeed come out with the best score. Since then, there have been several refinements of the model of the prisoners dilemma, including the option not to participate with a given opponent, and each player's turn being taken consecutively, not concurrently. Each time Tit-For-Tat came out well, and with minor revision, came out on top again.

Thus, cooperation can be a viable and beneficial strategy, under certain circumstances. The circumstances seem to involve some sort of implicit agreement to keep cooperating, an unlikely state of affairs. In fact, the situation is that in any finite series of interactions, defection is the best policy, it is only when the number of interactions is unknown that cooperation can be feasible. This is because of a logical regress when there is a finite series of interactions. If you know that there are n interactions, then it makes sense to defect on interaction n , since your opponent will have no means of retaliation. However, they will realise this, and thus defect on $n-1$. You realise this, and so on, and thus, defection from the start is how you should play. However, since in the natural world, nothing is absolutely fixed, the number of interactions is uncertain, and therefore fits the criteria. The exception that proves the rule is that you feel less need to be polite to someone you are never likely to see again or who will not recognise you if you encounter them again, leading for instance to the phenomenon of 'road rage', where individuals feel less identifiable, and thus are less inclined to be altruistic. So, cooperation can work, but how it can be implemented in real situations is still an issue.

2.5 How do unrelated individuals manage cooperation?

In order for individuals who have no inherent interest in each other to cooperate, there must be some mechanism that ensures it is worth their while to be altruistic in the first place, to ensure that they are not just going to get cheated upon the whole time. Obviously, those individuals who are bad at spotting those who will defect on them are selected against, and

therefore, some system should have resulted from the evolutionary pressures. That system is reciprocal altruism, first proposed by Robert Trivers. Under this system, individuals are altruistic precisely to receive altruism from their counterpart at some stage. Individuals will be altruistic so long as the benefit they receive is greater than the cost they incur. This makes perfect evolutionary sense, since they gain fitness overall, and any difference in fitness expresses itself significantly given time. There is of course the question of how this can be altruism, but we still retain the biological / psychological split, and this counts as biological altruism, if only in the short term, and says nothing about psychological attitudes. Trivers suggests that the basic principles are much like the game theory strategy of Tit-For-Tat, in that individuals will cautiously cooperate, unless given a reason not to. However, the punishments for defection are more varied than simple defection back. The most common way to play the strategy is to only play with those who are inclined to be altruistic back, thus any individual who cheats tends to be isolated, and receive no benefit. Thus, defection can in fact be made less profitable than cooperation, since a defector who never gets to play receives much less than cooperators who may occasionally be cheated on.

Trivers' favourite examples of reciprocity are so called 'cleaning stations' on reefs, where certain types of fish will gather and clean other fish who arrive, removing parasites and dead tissue. They clean all fish who come to them, including predatory fish, and their 'customers' in turn fail to eat them. According to Trivers, this arrangement works because both sides get a benefit, a meal and a clean, and it is in neither's interests to defect, since the predator will start to suffer from parasites, and the cleaners will lose a meal, and probably their lives. Trivers also points out that cleaning stations only exist on reefs, not in the open ocean, where there is much less repeat custom. There is no reason for a stranger not to have a clean, then eat the cleaners if it will never pass through the area again.

A combination of reciprocity and inclusive fitness can therefore have been produced by evolution, and would seem to approach what we term morality, but we must consider whether such behaviour actually exists in the natural world before examining humans for its existence.

2.6 Do we see altruism in the natural world?

If these altruistic tactics are profitable, we would expect evolution to have produced organisms that take advantage of them. Whilst inclusive fitness seems relatively uncontroversial, do we see other examples of cooperation and altruism amongst animals? In fact, cooperation is rife throughout all

organisms, from a certain point of view. All multi-cellular organisms have cells that are eukaryotic, a combination of simple prokaryotes and what were once parasitic invaders. Eukaryotes are now an inextricable synthesis of the two, with mitochondria being the best example of the parasites. Mitochondria are cell organelles that have their own genes, distinct from the rest of a cell's DNA, and they are responsible for producing the energy that the cell uses in all its activities. The common theory is that they are in fact resident parasites, that gave up their independence in return for greater protection and proliferation. Even our genetic code is a combination of elements, each capable of replicating without the others. Bacteria have only one chromosome, whilst we have twenty three. The cooperation and altruism of these tiny elements are not what we normally call morality, but under the biological definition of altruism, combined with the notion of reciprocity, they are very much cooperators and altruists.

On a more macroscopic scale, one of the classic examples of cooperation in the animal world is the case of the social insects, the hymenoptera. Ants, termites, bees and their like live in cooperating colonies, all working for the benefit of the colony, and were a classic case used in the promotion of group selection theory. However, after the discovery of their unusual haplo-diploidal genetic nature, a better understanding of their motives became possible. There are three varieties of these social insects, the male drones, the female workers and the queens. The drones carry two sets of genes, making them diploid, whilst the females, both workers and queens, carry only one set, making them haploid. Because of this, each worker is more closely related to their sisters than they would be to their own offspring, and it is therefore in their interests to help the queen, their mother, produce more daughters. Therefore, whilst it is unfamiliar to us, worker ants are willing to give their lives for the colony without producing offspring, since their sisters are more like their children than their children would be.

An excellent example of reciprocity in the animal world is that of vampire bats. These bats feed on the blood of other mammals, but they often fail to feed successfully. It only takes a few days without feeding for a bat to die, so it is imperative that they feed. Coupled with this is the fact that when a bat feeds successfully, it takes much more blood than it needs. Therefore, there is the opportunity for unsuccessful bats to receive surplus from successful individuals, thus ensuring their survival. This is in fact what happens, and the roosts in which the bats live contain mostly unrelated individuals, so inclusive fitness is not the reason for the altruism. However, should an individual refuse to give up some of their feed to a hungry roost mate, that individual is unlikely to receive any from the bat that it turned away. And given that the bats groom each other, they can tell who has a bloated stomach, and who

does not, defection is easily detected. Therefore, they bats tend to share their meals, thereby ensuring a more stable food supply for all.

The closest example to humans in the animal world is of course chimpanzees, with their fabled 2% genetic divergence from us. In chimpanzees we see much social behaviour involving reciprocity, including grooming, nursing other individuals offspring, and most significantly, fighting. Chimps appear to form alliances within their troops, and fight against each other for leadership. However, if a leader fails to reward those loyal chimps who put him in charge, they are quite likely to topple him from power themselves. This is a clear example of retaliation at the contravening of a reciprocal arrangement, a defection deterrent mechanism.

2.7 Is there therefore morality in the natural world?

It seems clear to me that there is a large amount of cooperation in the wild, from the shoaling movements of fish, to the collaboration of genes in all multicellular organisms. The strongest motivation to be altruistic is that of genetic relatedness, thus loyalty is to the family first. However, reciprocal altruism allows unrelated individuals to enter into mutually beneficial relationships. Whilst this always entails an overall benefit for the individual or their genes, that does not mean that altruism is a front for egoism. Perhaps we simply set the criterion for what counts as altruism too high. Is it unreasonable to suppose that one can be altruistic and derive benefit from it when we spout aphorisms such as “One good turn deserves another”?

Chapter 3

How do humans implement morality?

3.1 What are the human characteristics that put them in a different moral realm?

Evolution has given humans several abilities that set them apart from the rest of the inhabitants of the planet, to a greater or lesser extent. Our conscious nature is considerably above that of any other animal, if indeed you attribute any mental states to non-human animals at all. With our increased brain power and special social adaptations come several key factors that could affect our implementation of morality. Our rationality and ability to go far beyond what is given to us in our genetic code leads to the possibility of forming our own systems of behaviour, and the medium of culture has the potential to pass on many ideas through a non-genetic medium, and indeed could create ideas of its own. These factors undoubtedly separate humans from other animals, but whether they affect our morality is under question. Potentially they could mean that we are in fact the only creatures capable of being moral.

3.2 How does humanity's rational nature affect their moral nature?

There are two possible ways in which rationality could affect our moral nature. The first is that we may be able to determine what is the best course of action through application of our rational powers, and the second is that we may do what is good because we rationally choose to. These two positions

are very close, but importantly distinct. The idea that we can, or possibly need to rationally determine the moral course of action is best typified by the ideas of Immanuel Kant, who asserted that actions could not be moral unless done deliberately in view of the moral law. Under his view, rationality was the only way to be moral, and moral sentiments should essentially be ignored, correct or not. If we acted simply in accordance with our moral sentiments, then we were simply looking for the gratification that came in fulfilling them, not seeking to do something right. This is a fairly unpopular view, a fact that is not surprising if one holds that we have in fact inherited our moral intuitions from our non-human, non-rational ancestors. There is of course no guarantee that reason would deliver only one acceptable moral system, or indeed that it would deliver a moral system at all. If we examine game theory again, it is clear that rationally, defection is always a good choice, especially if one can get away with it. Kant certainly would not approve of a moral system where one could murder a stranger if no-one would ever find out, yet that is what game theory would suggest we do if rationality and self interest are the two motivating factors. Kant does of course add several rules of his own that should be followed to produce appropriate answers, but one then gets into problems about the justification for those rules, which is well outside the scope of this project. A slight variant on this idea is that we choose our moral system, that whilst we might not decide each time what the right course is, we choose our principles, then act in accordance with them, treating them almost like moral sentiments. Kant would not have approved, but it is perhaps a more reasonable view. However, there are still requirements for a basis upon which to make those first judgments about what to select as a moral system, and, as products of evolution, we may well exhibit the same altruistic tendencies as animals before our specifically human attributes affect our response.

The second way that rationality can impinge upon morality, the idea that we choose to be moral, is also somewhat Kantian, but has significant differences. The idea is that we would already know what the right thing to do was, but would then decide whether or not to do it, using rationality as a guide. We would then only be good if it made rational sense to do so. This is similar to the doctrine of externalism, which says that simply accepting that something is morally the right thing does not entail being motivated to do it, as opposed to internalism which equates recognition of goodness with motivation to act that way. This view almost seems to remove morality from the equation, as you would be choosing the rational thing to do anyway, and if it happens to coincide with what is moral, then that is a benefit, but not a necessary element. If humans operated like this, it would indeed make them a special case of morality, since they would essentially have none. However,

if one ascribes to an evolutionary view, which would be similar to emotivism in this area, then recognising an action as good inherently requires being motivated to do it. This idea touches on the question of what makes an action a moral action, whether intentionality is required, but that will be discussed fully later. At the very least, it would seem that for an action to be moral, some morally significant causes should be a necessary, even if not sufficient part of its explanation.

3.3 Are humans motivated by moral sentiments, and if so, what is their source?

If we accept that morals exist in the animal kingdom, in one form or another, then there must be some mechanism causing those animals to act that way. If we choose not to attribute mental states to animals, as most would not, certainly not at levels below primates, then we seem to be left with instincts determined by their genetic codes. If that is the case, then as humans are also the result of evolution, we should expect there to be at the very least remnants of these moral instincts in our behaviour. Whether or not our additional capacities add or detract anything from them is another matter, but it seems that these moral sentiments would have to have some effects. Few claim to be able to completely escape their biological origins, unless one subscribes to a form of dualism, a view somewhat incompatible with an evolutionary explanation. Hume for one considered humans to have strong moral sentiments, and that they were in fact pointing the correct way. His idea of the moral sense is the kind of way we would expect moral instincts to manifest.

Most people do not have a developed moral system, in the sense of a full set of justifications for their moral beliefs. They do however quickly form judgments on what they consider to be right and wrong, and this would suggest that there is no requirement for consideration in moral judgments. The idea of a moral sense, either in touch with objective morals, or simply ingrained into all of us by nature or nurture, is a common one, and certainly gains weight when one considers the degree of similarity of morals between cultures. I shall not be considering the possibility of objective morals for the moment, rather I shall focus on the idea that humans have a certain set of moral precepts that result from their path of evolution. Certainly there are many examples of moral disagreements between cultures over ethical choices, but more interesting are the underlying commonalities. If one considers the Islamic belief that it is right to cut the hands off thieves, and compare that

with a more western belief in humane treatment of criminals, they seem incompatible. However, if you then realise that both views imply punishment of people who break the rules of society, and do not allow for unjustified acts of violence, it seems that there are common themes. I would argue that the common themes are in fact the necessary elements of a social system, and that the superficial variations are merely local flavours, in much the same way as languages may differ, but the universal presence of a spoken language is significant.

The source of these moral sentiments is somewhat debateable. Many have postulated some moral area of the brain where a form of ethical calculus is carried out, but it seems more reasonable to suppose that our brain structure, as laid down according to our genetic code, involves predispositions to act in certain ways. Certainly, the whole concept of moral sentiments implies that they are not something under conscious control. One may be able to change an ethical belief by adding more information, but one cannot simply decide to consider something else moral.

3.4 Does the development of culture bring anything to human morals?

The interplay between culture and morality is a delicate one. For the most part, people consider morality to be a prerequisite for culture. Whether you follow Rousseau or Hobbes, morals are inherent in society. Without some system for encouraging social behaviour, living in groups would not be an advantage. However, there is also the possibility that culture brings extra qualities to morality. For example, culture can act as a way of passing on information, so that we even now know what Plato thought was the right thing to do. It is therefore possible that, thanks to culture, we could have escaped our genetic programming entirely. Since the development of the phenotype is a result of the interaction of the genes with the environment, and the culture shapes the environment, one can see that it is perhaps possible that non-genetic moral beliefs could perpetuate. Indeed, the moral divergence between cultures seems to require some mechanism such as this. Since culture to any appropriate level is indeed unique to humans, this would most definitely set us above the rest of the planet's creatures, although one could look to the tool-using of certain bands of chimps as examples of cultural heredity in non-humans. Perhaps the morals that come from the genes are very basic, and to compare it to a moral system enriched by culture shows up its primitiveness.

The idea that we can supervene our genetic dispositions, either deliberately or not, is clearly shown in cases where people are asked to act in ways which are not normal for them. Brainwashing someone is perhaps difficult, but is entirely possible, demonstrating that beliefs, including moral ones, can be altered. If we were genetically hard-wired with certain unchangeable morals, this would not be possible. As with most behavioural adaptations, our innate morals are best conceived of as predispositions, not imperatives. We are inclined to help our offspring, but that doesn't stop some people from killing them. A good example of how our natural tendencies can be overridden is provided by Matt Ridley in his description of the Hutterites, a religious group who model themselves on insect colonies. They rigidly enforce a policy of pro-social behaviour and working for the good of the community, with strict punishments for infractions. This supposedly very moral behaviour is successful, but the need for the punishments highlights that it is not natural behaviour to all. The natural tendencies of a human is to combine altruistic actions with egoistic actions, to the best degree for themselves, and cultural conditioning hasn't removed that, simply discouraged it. Whilst this example is of encouraging moral behaviour, there is no practical difference between this and encouraging less moral behaviour, it is all behavioural conditioning, and demonstrates the flexibility in our phenotype.

These examples show both that our moral predispositions are not absolute, but that at the same time they are strong tendencies. Culture is capable of modifying them, but is unlikely to remove them entirely, and cannot cease enforcing them if the offspring of the society are to continue to follow the same rules.

3.5 Could memes play a part in human moral behaviour?

First of all, a description of what memes are. A 'meme' is a "distinct memorable unit" [Den95, pg 344], roughly analogous to a gene. Some regard memes as metaphors for more complex but less well understood concepts, whilst others regard them as genuine entities. Richard Dawkins first coined the term to help explain some ideas about the evolution of cultural ideas, but their main champion of late has been Daniel Dennett. Roughly, memes are ideas, such as the idea of the wheel or a poem, and can be transmitted between individuals in a process analogous to reproduction. The analogy with genes is somewhat shaky, but Dennett uses it as a starting point, and develops a more realistic model specifically for memes, rather than copying the genetic

model wholesale. If one considers that memes exist, and inhabit the human mind, then one can draw many interesting conclusions from this. Given that there are a limited number of hosts (humans), there would be competition for resources, and this entails selection between memes. The idea of cultural evolution is not new, but the meme metaphor is perhaps the most promising avenue of inquiry into it. As with genes, memes that are more successful prosper at the expense of the less fit, and the fitness of the memes is not necessarily equated with that of the individual host. Dennett's example is a meme that is a compelling reason for suicide, but also causes the host to leave a note containing the reason for suicide. Whilst the host would then kill themselves, and that instance of the meme would be destroyed, another host could well read the note and receive the meme themselves. In this way, the meme would prosper, if only until a meme for not reading suicide notes developed.

The interesting implications of memes for human morality are that memes could in fact carry some morality altering effects, such as particular moral beliefs. This would allow an additional line of development for morals in humans, one that could be viewed as corrupting, or simply enhancing our 'true' morality. One somewhat worrying concept is that we do not choose the memes that inhabit our mind, and indeed our mind may be constituted by the memes we carry. This leads to unpleasant, although not new, conclusions about human freedom. If we are constituted by the memes we carry, does that leave anything which is 'us'. This is relevant not only in personal identity, but could also be relevant in terms of moral responsibility. If memes can change our morals, then what happens if we act in a way coherent with our morals at that moment, then receive a morality altering meme, which tells us that the act we committed is wrong. Which perspective should we take on the action, and are we morally culpable if our society deems the action to be wrong?

If memes are a valid idea, then one phenomenon they could explain is that of cultural diversity. If certain memes prosper in one group, but encounter 'hostile' memes when attempting to enter another group, this would lead to groups with different ideas, and Dennett postulates, memes can in fact set up screening processes for other memes, so that only memes that adhere to the current network of memes will be examined, let alone allowed to inhabit the host. Memes could create genuine differences in morals between cultures, perhaps even to the level of changing the basic tenets of morality.

Personally, I do not find the meme metaphor coherent, despite Dennett's improvements to it. I think that there are simply too many differences for the analogy to hold, and that whilst ideas are transmissible, the concept of memes is too strong an emphasis on the ideas themselves. There is a good

degree of selection of the incoming ideas by the ‘host’, both on the conscious and unconscious level. We do indeed rationally examine most new ideas, rather than them simply vying for space, and there are many genetically based behavioural predispositions to protect us from such mental viruses. Additionally, as Dennett himself admits, memes mutate and hybridise more frequently and severely than natural selection can cope with, so they should descend into chaos very rapidly. In short, whilst culture has its part to play in human morality, I do not consider memes to be a relevant way of describing that part.

3.6 Why is morality so prevalent in humans?

Whilst, as we have seen, morality may vary somewhat between one culture and another, there is no human culture that exists without a moral system. It may in fact be one of the prerequisites for a society emerging. Certainly, there seem to be large fitness benefits for those individuals that live in societies, and evolution seems to have left behind those individuals who do not live socially. Hermits are not known for their reproductive capabilities. However, there is some question, as this chapter has attempted to show, as to whether morality comes intuitively from within individuals, whether it is developed progressively by a society, or whether it is devised rationally by individuals. It seems that a combination of all three is most likely, with perhaps less emphasis on the rational development, perhaps mediating it through the social development. The meme idea is untenable, and should be discarded, but some form of cultural transmission is essential, and observable. The simple fact that societies persist requires the transmission of the societies precepts to the next generation, and whilst the mechanism may be disputed, the brute fact of cultural inheritance cannot be.

One question of morality that is likely to be unique to humans is how morality can be binding. If we consider that morality is part of the behavioural disposition that our genetic code lays down, then there is no reason per se for us to adhere to it. It is true that our intuitive reactions would involve those moral concepts, but humans are eminently capable of exceeding their genetic behavioural constraints, as examples such as peaceful protest show. Our natural instinct is to fight back or flee when attacked, but if one considers the case of Gandhi and his followers letting themselves be clubbed down, one can see that it is possible to overcome these behavioural dispositions. What is it about morality that makes it so resistant to contravening, if not quite immune to it? In attempting to answer this, Michael Ruse suggests what would be one uniquely human aspect to morality, a tendency to

‘objectify’ morality. He asserts that whilst there is no real justification for what we consider to be moral, we inevitably tend to consider there to be an objective justification. The implication is that our rational nature needed a counterbalance so that we would continue to act morally. If we did not tend to objectify, then we would in fact lose the tendency to be moral, and thus lose the advantages that come with morality.

Chapter 4

What are the differences between human and animal implementations of morality, and what are their implications?

4.1 What features of human consciousness distinguish us morally from animals?

Our conscious nature is usually the feature that is picked out as the one that makes us so special. Our ability to have abstract thoughts as opposed to simple perceptions has allowed us to achieve all the technologies we have, and must have strongly influenced our formation of societies. In addition to this, many moral philosophies require consciousness behind an action for it to be considered as moral. Kant's is the paradigm case, where action that is done with any element that is not rational is not moral, but most philosophers distinguish between motives in action. Thus the common consequentialist counter-example of helping an old lady across the road. If she is then mown down by a bus halfway across, that doesn't make the action a bad one, since the intention was good, despite the unfortunate outcome. At worst an action such as this would be considered morally neutral. The suggestion that moral action requires an appropriate intention would tend to exclude animals from being capable of moral action, but is it in fact necessary to have intent to act morally? Consequentialists would obviously contend that motives are not the relevant part of the action, but all too often our intuition runs counter to that.

Certainly in law it is the intention that is most relevant, and most consider law to be at least partly enshrined morality. Thus we have manslaughter as opposed to murder when the killing was not deliberate, and one can be convicted of attempted murder, or conspiracy to commit murder. These and many other crimes, which we evidently consider to be wrongdoings, are based on intention rather than action. However, counter to this, when one considers a thought experiment where individuals are assessed for likelihood of criminal behaviour, and those who are likely to do wrong are imprisoned, it seems that we do also require some part action for moral culpability. The error here is a conflation of actual intent with possible intent. Had the individuals been found to actually be harbouring intent to commit crimes, we would be less worried about punishing them, it is because they have not yet formulated such an intent that we have qualms, as well we should.

Given then that intent is a requirement for moral action, we must examine whether animals are capable of intent. Most people's instant reaction is that they are of course not capable of it, since they do not have the level of consciousness required. However, as Bruce Waller points out, we are anthropomorphising intention. Given our point of view as humans, rationality is one of our key factors in forming intent, but it is all too often considered to be the only one. In fact, rational deliberation is only one way of forming intent, and is perhaps not the most morally relevant. He agrees that "proper intent is essential for moral behaviour [but] reason and deliberation is not." [Wal] He points out that we would not claim that someone reaching out to catch a falling child 'without thinking about it' was not acting morally. The intention was to catch the child, thus protecting it from harm, and that is a morally good act, despite the lack of rational deliberation on the part of the catcher. Indeed, he goes further, pointing out that we may in fact think less of someone who has to take a moment to decide whether catching the child is the right thing to do before acting. Given that humans do not need more than instinct to form intentions, it seems feasible that animals may also be able to form intent. Indeed Waller cites examples of hyenas "intending to feed upon an animal carcass" [Wal] before being chased away by lions, or chimpanzees intending to mate with a female, but then being approached by a dominant chimp. These intentions are not the product of rational thought, but that does not make them any less intentional. If we require more than instinct, then I doubt that most human action would count as moral, barring the difficult cases that give people moral dilemmas and can only be resolved by rational inquiry.

4.2 What meta-ethics supports this view of morals?

Evolutionary theory has a bad name in terms of the ethics it justifies. In its first incarnation, as evolutionary ethics, many ridiculous, or even abhorrent, ideas were promoted in the name of evolution. Thomas Huxley was adamant that evolution was a progression towards perfection, and that therefore that which promoted evolution was good, and that which hindered it bad. This led to him calling for the cessation of support to the poor, much as Malthus, one source of Darwin's inspiration, had done. Evolutionary ethics often called for harsher conditions for people, following from the belief that selection would weed out the weak, leaving the strong to survive, thus improving the human race. Whilst there were some intuitively good ideas that came under the name evolutionary ethics, such as Carnegie's idea that all individuals should be given the best chance to develop their full potential, most of the ideas involved the removal of what were deemed the 'unfit' elements of the society. However, there are two large errors inherent in this view of evolutionary ethics. The first is that, as Anthony Flew points out, there is no way that the unfit can be surviving to the detriment of the fitter members, since evolution calls for fitness to be defined in terms of survival and reproduction. The very fact that these supposedly unfit classes continue to exist denotes the fact that they are fit enough to survive. Evolutionary ethics is making unfounded value judgments in order to determine what it will deem morally acceptable. The second major error that was part of evolutionary ethics is the notion of progress that it uses. As Michael Ruse, amongst many others, points out, "palaeontology does not, in itself, thrust some notion of an upward progressive climb of life upon you." [Rus86, pg 13] Again, the evolutionary ethicists were making assumptions.

However, since then, there have been other syntheses of evolutionary theory and ethics, of which I will describe the meta-ethical systems of two. The first is that of Michael Ruse, who asserts that evolution is an explanation of ethics, and other areas, rather than a justification for anything. He lays out a system essentially without justification for normative ethical beliefs, so that "the claims of normative ethics are like the rules of a game." [Rus95, pg 249] That is to say that they are essentially arbitrary rather than anything real. Part of the reason that Ruse avoids making meta-ethical claims is that evolutionary ethics is all too often attacked with claims that it is committing the naturalistic fallacy. What this means is that one is allegedly determining 'ought' from 'is', something that is supposedly logically impossible. Both Hume and G.E. Moore pointed out the gap involved here, although it is

more Hume's ideas and Moore's term that have stuck. There have been repeated attempts to demonstrate that there is no gap, but none have been particularly successful. Thus, Ruse's sidestep of the whole notion of 'ought' places him well outside the attack of the naturalistic fallacy. However, what it does do is leave him at the mercy of those who demand that there should be a justification for ethics, since he goes even beyond subjectivism in terms of a lack of support for moral differences, whilst asserting that morality is still coherent and forceful.

One such view is espoused by Rottschaefer and Martinsen in their response to Ruse[RMb]. They claim that Ruse is too afraid of the naturalistic fallacy to formulate a reasonable system of justification. They seek to expound a theory that evolution has given us the ability to identify objective morals, but that unlike many moral realist positions, these morals are relations, not things. This idea of supervenient relational properties is able to avoid most of the standard objections to moral realism, such as Mackie's argument from queerness. What Rottschaefer and Martinsen suggest is that there are real objective morals, but that they are the relation of the person to the environment, equivalent to the person's fitness. This means that what is morally good is what is fitness enhancing, an idea that would fit well with the game theoretic analysis of the profitability of morality. If they are right about these relational properties, that would give a real justification for the ethical systems that exist, and would be an objective measure of the 'rightness' of actions and choices. However, it seems to me that they are making something of a category mistake by terming relational properties objective. Whilst one could reasonably argue for terming fitness a quantitative value, and therefore assigning some scientific impartiality to its measurement, there are moral assumptions inherent in deciding that fitness is what is in fact the morally right thing, and in these assumptions the naturalistic fallacy has been committed, where our tendency to see mutually fitness enhancing behaviour as morally good is believed to entail that fitness is what is good. The only response Rottschaefer and Martinsen could supply would be that they are in fact defining moral goodness as fitness, but that leads directly to a response that they are making arbitrary decisions about morality, and creating new moral systems, not explaining the extant one.

4.3 Why do agents act morally or immorally?

This may seem like a restatement of the question that this whole project is set to answer, but what I am in fact asking here is what causes individuals to act morally in a certain case, rather than in general. I am examining the

causes of moral action rather than the explanation of it. The question to examine is whether we are in fact moral to benefit ourselves. Given that the examination of game theory has shown us that altruism can in fact lead to profit, does that make altruism egoism, as many egoists argue. In fact this is the point where the biological / psychological motivational split is relevant. In order for evolution to have produced us, we cannot be truly biologically altruistic, and that is what game theory shows. By being altruistic in the short term, we can in fact be biologically selfish in the longer run. However, if we want to retain the sense of moral action that we believe we have, that we are in fact capable of altruistic action, then we need to show that our selfish genes can produce psychological altruism. The way to do this is to point out that our genes do not in fact have motives, only effects. Thus, by saying that the explanation for a behaviour is that it is in the interests of the genes is not equivalent to saying we are doing it for the genes. In fact, so long as our behaviour matches the required behaviour for the genes, the mechanism to get us there is irrelevant, except perhaps morally. What has in fact happened is that “in order to make us ‘altruists’ in the metaphorical biological sense, biology has made us altruists in the literal, moral sense.” [Rus95, pg 241] As with all behavioural dispositions, evolution has not produced something that is deliberately acting with a strategy in mind, we are not playing Tit-For-Tat, but has produced within us inclinations to act in certain ways, which may well happen to produce effects similar to such a strategy. Thus, we are capable of wholly altruistic moral action, and it can produce biologically selfish results in the long run, and there is no contradiction. Our psychology is such that we want to look after our children, and we feel guilt for having injured someone without cause, and we do this without the intention to benefit ourselves. This is an important point, since there are very few people who wish to remove the altruism from moral action. This emphasis on psychological altruism does not exclude non-humans from real altruism, since as we have already asserted, intentionality is the key aspect of moral responsibility, and animals are capable of forming intentions.

If we are therefore motivated to act morally in a very real sense, not out of selfish desires, then why is it that we do in fact act immorally as well? If cooperation is the best policy, what is the benefit of defecting. In fact, if we look again at game theory, cooperation is not the best policy, defection whilst the other party cooperates is. What this tells us is that, if evolution finds the most beneficial paths, then we should be inclined to cheat when we think we can get away with it. Robert Axelrod points out that “while an individual can benefit from mutual cooperation, each one can also do even better by exploiting the cooperative efforts of others.” [AH] And that is exactly what happens. All too often, people are willing to help while they must, but cheat

while they can. Benefit fraud has to be the paradigm example of this, where an individual damages their whole community for their own gain. The reason that we are not immoral more is that there has been an evolutionary arms race of better cheats and better cheating detectors. Thus, it is very difficult for people to cheat without getting caught, and it is therefore more normally in their interests to simply cooperate, and therefore avoid the possibility of sanction. Excellent evidence of this is provided by Matt Ridley [Rid97] when he describes the results of a Wason selection task based on moral choices. People are considerably better at getting the correct answers when asked to discover those individuals who are cheating, rather than those who are cooperating properly. We seem to have a greater mental faculty for spotting cheats, who are a threat, than cooperators, who are not.

4.4 How have humans as social animals exceeded the original designs of morality?

The bulk of human evolution occurred well before we started living in villages, and so we should perhaps not expect all the behavioural attitudes that formed there to be ideal for modern life, in particular the large number of individuals that we interact with. However, morality seems to fit surprisingly well with our new lifestyle. Why is this when we have had no time to change our genetic rules enough to compensate? The two possible answers are that we did not in fact need to alter the rules, that they were already suitable for the new lifestyle, or that the rules have been altered by some non-genetic factors, perhaps culture or rational direction as discussed earlier. One key point that suggests that some alteration must have occurred is that we have nowhere near enough brainpower to keep track of the reciprocity involved in the hundreds of people we meet all the time. Our capabilities are sufficient for a tribe, not a city. By rights, we should be aggressive towards other ‘tribes’, which basically means individuals we do not recognise, yet politeness is a virtue, and a common one at that.

The cultural explanation for this would be that as society began to be implemented, one of the earliest rules it would have enforced was that of cooperation, and it has continued to reinforce that behaviour, leading to general acceptance that cooperation with all is the appropriate form of behaviour. However, a more biologically based explanation, in line with the biological / psychological distinction discussed earlier, is that our genes in fact give us a tendency to be cooperative to those individuals in our group, where that was originally our tribe, but the same psychological inclination applies now,

since we have simply expanded our concept of ‘group’. This is possibly also an explanation of how individuals first began to cooperate with unrelated individuals. Genes simply provided the instruction to be cooperative in most circumstances, in expectation that the individual would be in a family group, but when outsiders joined the group, the tendency continued to be that of cooperation, which proved to be beneficial.

4.5 What level of control do we have over our morals?

One factor that our conscious nature could bring to morality is the ability to change it. Whether it is through the cultural modifiers discussed earlier, or some other method, the ability to extend and modify our morality would be a definite difference from animals. However, whether we have any active control over our own morals is very much debateable. Most people would be inclined to say that they do not choose what they find to be moral, that whilst they can have their mind changed about whether a particular action is right or wrong, that is more a result of realising that certain moral standards apply differently to how they initially thought. There is no change of morals involved, only a change of perception. Whilst some philosophers believe that we shape our own morality, often in our childhood development, many more consider our moral phenotype to be a result of our genotype’s interaction with the external environment, rather than the internal one.

Linked with the idea of changing one’s morals is the idea of opting out of mainstream morality. Many people, including some philosophers, have considered this to be a real possibility. Nietzsche thought that through force of will we could throw off the restrictions of slave morality, and become stronger, better individuals with a more natural morality of the strong ruling the weak. However, not only does it seem from biology that much of our cooperative, herd morality is innate and not forced upon us, it also seems unfeasible that we could abscond from morality by force of will. Our morality may develop in such a way as to lead to us being suitably ruthless for Nietzsche, but we do not have the power to abandon it. This can be seen when one looks at the people who are described as ‘amoral’, or ‘psychotic’. This is usually either an innate tendency, i.e. there from birth, or the result of some massive psychological trauma. Certainly such cases are considered to be aberrations rather than deliberate acts.

4.6 What are the key human modifications of morality?

Obviously, being non-identical to other species, we will have a slightly different version of morality, but there should not be too many changes. However, we possess some particularly remarkable attributes that have the potential to massively alter our behavioural patterns. In the case of morality, they appear not to have made significant alterations, presumably because the benefits derived from morality and cooperation are so significant that any individual or group altering them significantly was at a disadvantage when compared with the moral individuals or societies out there. Thus, we have not altered the morality that evolution has provided us, and our ancestors, with.

Another way in which humans have the potential to be significantly different to their animal ancestors is that they can conceptualise morality, and can realise things about it. There have been attempts to require this to be involved in moral action, but I think that intention is the relevant part of a mental state when determining if the morality of an action should be judged, and I also consider that animals are capable of forming intentions, and can thus be considered moral agents. There is the need for a decision on the level at which intention ends, but the relevant result is that some animals are capable of moral action, and therefore morality is not unique to humans. My own view is that one can argue successfully that even unicellular organisms can have intentions, but that is not in the scope of this discussion.

Humans may therefore bring additional qualities to morality, such as a reflective capacity, and a better capability for working their way out of moral grey areas through reasoning, but they are not the only residents of the moral domain. They also do not leave behind their genetic heritage, including the moral tendencies that it imparts.

Chapter 5

Conclusions

5.1 Is morality the result of evolution?

From an understanding of game theory combined with the relatively new concept of selfish genes, one can see that the evolution of morality is most certainly not an impossibility. Limited altruism between individuals can lead to mutual increases in fitness, a trend that will lead to such tendencies developing throughout populations of individuals. And whilst I have not included empirical data in this project, there is plenty of evidence that animals certainly do act in a biologically altruistic manner in order to be biologically selfish in the longer term. There is some disagreement about whether such benefit seeking action can really be classed as altruism, and thus moral action, but I think that I have shown that an action having a beneficial result does not entail a selfish motive or intention on the part of the agent. On the question of when an action counts as a moral action, I think that animals are entirely capable of forming the intentions to act that moral action requires, and that any attempt to keep such cases out of moral action is highly anthropocentric and biased. Certain acts of biological altruism may fall short of being classified as moral, such as the cooperation of cells in a body or genes in a cell, but certainly the cases in the higher animals seem to qualify as genuine altruism. The individual acting may not recognise it as such, that is our luxury, but that does not remove the intent, or the moral quality. Animals are capable of acting as morally as humans, and often with the same moral system, as one would expect given their common ancestry.

5.2 Are humans a special case of morality?

The nature of humans as rational animals tends to lead people to classify them separately from the rest of nature. However, the project of sociobiology is validly applied to us, as we are as much the result of evolution as any other organism on the planet. That much is inescapable. We started out from a common ancestor much like a chimpanzee, and with many of the behavioural tendencies of the higher primates, including those tendencies that lead to moral action. Whether we have improved upon that is the relevant question. It seems to me that culture and rationality may have made some improvements, but they are in the manner of additions, not changes to basic principles. The base principles of morality are the result of the drive to succeed in competitive environments. The level of complexity of our morality is a result of more precise and abstract renditions of these principles in response to new problems being raised both by intellectual inquiry and changing cultural environment. The common themes in all extant moral systems are the beneficial result of blind evolution. They may be able to be removed from individuals, such as in the cases of psychopaths, but only at the cost of those individual's humanity.

5.3 Final statements

One reaction to the idea that morals are the result of evolution is that they are suddenly devalued, that because they are beneficial to us in some ways, they are not in fact moral. On the contrary, we are moral very much in spite of our selfish tendencies. In order that we can take full advantage of the benefits cooperation can bring, we have become moral beings. However, that 'we' does not extend only as far as humans. Animals too are capable of moral action to the same extent as us, only they are much less capable of expressing it. Additionally, the lack of justification that an evolutionary explanation brings to morality can be dismaying, but that is no reason to dismiss the idea. Our genes have given us the tendency to consider morality to be absolutely true, and that has been highly beneficial to us. However, it is only an adaptation, it need not be absolute truth. To expect that from evolution is to make assumptions about progress, and thus alter the true sense of evolution. If it had been more beneficial, evolution would have made us all wholly immoral, and pleased about it.

Chapter 6

Bibliography

Bibliography

- [AH] R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396.
- [Ard70] R. Ardrey. *The Social Contract*. Garden City Press, 1970.
- [Aya94] F. Ayala. *Biology, Ethics, And The Origins Of Life*, chapter The Differences Of Being Human. Jones And Bartlett, 1994.
- [Bla98] S. Blackburn. *Ruling Passions*. OUP, 1998.
- [Cal97] W.H. Calvin. *How Brains Think*. Science Masters, 1997.
- [CS] J. Collier and M. Stingl. Evolutionary naturalism and the objectivity of morality. *Biology And Philosophy*, 8:47–60.
- [Dar85] C. Darwin. *The Origin Of The Species*. Penguin, 1985.
- [Daw99] R. Dawkins. *The Selfish Gene*. OUP, 1999.
- [Den95] D. Dennett. *Darwin's Dangerous Idea*. Penguin, 1995.
- [Dia97] J. Diamond. *Why Is Sex Fun?* Science Masters, 1997.
- [Fle67] A. Flew. *Evolutionary Ethics*. MacMillan Publishing, 1967.
- [Gei] G. Geiger. Why there are no objective values. *Biology And Philosophy*, 7:315–330.
- [Gri] P.E. Griffiths. A sober view of life. *Biology And Philosophy*, 12:427–431.
- [Haj] I. Haji. Evolution, altruism, and the prisoner's dilemma. *Biology And Philosophy*, 7:161–175.
- [Har] W.F. Harms. Adaptation and moral realism. *Biology And Philosophy*, 15:699–712.
- [Joy] R. Joyce. Darwinian ethics and error. *Biology And Philosophy*, 15:713–732.
- [Lob] A. Lobo. Man: The rational or irrational animal. Thesis, Murdoch University.

- [Mac77] J.L. Mackie. *Ethics: Inventing Right And Wrong*. Penguin, 1977.
- [Mid94] M. Midgley. *The Origin Of Ethics, A Companion To Ethics*, pages 3–13. Blackwell, 1994.
- [Mid95] M. Midgley. *The Ethical Primate*. Routledge, 1995.
- [mon] monkey. Flarn.
- [Rap47] D.D. Raphael. *The Moral Sense*. OUP, 1947.
- [Rid97] M. Ridley. *The Origins Of Virtue*. Penguin, 1997.
- [Rid00] M. Ridley. *Genome*. Fourth Estate, 2000.
- [RMa] W. Rottschaefer and D. Martinsen. The insufficiency of supervenient explanations of moral actions. *Biology and Philosophy*, 6:439–445.
- [RMb] W. Rottschaefer and D. Martinsen. Really taking darwin seriously. *Biology And Philosophy*, 5:149–173.
- [Rot98] W. Rottschaefer. *The Biology And Psychology Of Moral Agency*. CUP, 1998.
- [Rus79] M. Ruse. *Sociobiology: Sense Or Nonsense?* D. Reidel, 1979.
- [Rus86] M. Ruse. *Taking Darwin Seriously*. Basil Blackwell, 1986.
- [Rus95] M. Ruse. *Evolutionary Naturalism*. Routledge, 1995.
- [Ryaa] J.A. Ryan. Taking the ‘error’ out of ruse’s error theory. *Biology And Philosophy*, 12:385–397.
- [Ryab] J.A. Ryan. Woolcock, ruse, again. *Biology And Philosophy*, 15:733–735.
- [Sob] E. Sober. The evolution of altruism. *Biology And Philosophy*, 7:177–187.
- [Sob94] E. Sober. *From A Biological Point Of View*. CUP, 1994.
- [Sob00] E. Sober. *Philosophy Of Biology*. Westview Press, 2000.
- [SW98] E. Sober and D.S. Wilson. *Unto Others*. Harvard Press, 1998.
- [Tri] R. Trivers. The evolution of reciprocal altruism. *Quarterly Review Of Biology*, 46:35–57.
- [vST84] F. von Schilcher and N. Tennant. *Philosophy, Evolution and Human Nature*. Routledge and Kegan Paul, 1984.
- [Wal] B. Waller. What rationality adds to animal morality. *Biology And Philosophy*, 12:341–356.
- [War00] M. Ward. *Virtual Organisms*. Pan, 2000.

- [Wila] D.S. Wilson. A critique of r.d. alexander's view on group selection. *Biology And Philosophy*, 14:431–449.
- [Wilb] D.S. Wilson. On the relationship between evolutionary and psychological definitions of altruism and selfishness. *Biology And Philosophy*, 7:61–68.
- [Wim] W.C. Wimsatt. Genes, memes and cultural heredity. *Biology And Philosophy*, 14:279–310.
- [Woo] P. Woolcock. Ruse's darwinian meta-ethics: A critique. *Biology And Philosophy*, 8:423–439.