

Group Selection in the Evolution of Religion: Genetic Evolution or Cultural Evolution?

Taylor Davis

Department of Philosophy, Purdue University, 100 N. University Street,
W. Lafayette, IN 47907-2098, USA
taylordavis@purdue.edu

Abstract

In the scientific literature on religious evolution, two competing theories appeal to group selection to explain the relationship between religious belief and altruism, or costly, prosocial behavior. Both theories agree that group selection plays an important role in cultural evolution, affecting psychological traits that individuals acquire through social learning. They disagree, however, about whether group selection has also played a role in genetic evolution, affecting traits that are inherited genetically. Recently, Jonathan Haidt has defended the most fully developed account based on genetic group selection, and I argue here that problems with this account reveal good reasons to doubt that genetic group selection has played any important role in human evolution at all. Thus, considering the role of group selection in religious evolution is important not just because of what it reveals about religious psychology and religious evolution, but also because of what it reveals about the role of group selection in human evolution more generally.

Keywords

altruism – cooperation – cultural evolution – group selection – prosocial behavior – religion

1 Introduction

In the blooming scientific literature on the evolution of religion, all theories explain the traits of religion by appealing to natural selection. Theorists differ, however, with regard to the kinds of selection processes to which they appeal.

Some adhere closely to the orthodoxy of inclusive fitness theory, appealing only to genetic selection, and only to selection among individuals (e.g., Boyer, 2001; Barrett, 2004; Bering, 2006; Bloom, 2007; Baumard and Boyer, 2013; Baumard et al., 2014). Others adopt a broader adaptationist perspective, appealing to cultural selection in addition to genetic selection, and to multilevel selection instead of selection at the individual level alone (e.g., D.S. Wilson, 2002; Bulbulia and Mahoney, 2008; Norenzayan and Shariff, 2008; Atran and Henrich, 2010; Haidt, 2012; Norenzayan, 2013). And since both genetic evolution and cultural evolution may occur at either the individual level or the group level, this broader approach recognizes no less than four distinct forms of selection that could potentially explain the evolution of any given trait. I thus distinguish between the theoretical frameworks of broad adaptationism, in which four distinct forms of selection process might explain a given trait, and narrow adaptationism, in which only individual-level genetic selection is recognized.

Here I consider a debate that occurs within the broad adaptationist camp, regarding one particularly important feature of religious psychology: the relationship between religious belief and altruism, or prosocial behavior. In Section 2 I discuss in detail what “altruism” and “prosociality” mean in this context. The upshot is that I take them to be synonymous, so for simplicity’s sake I will use the term “altruism” throughout.

Both of the theories involved in this debate explain this relationship by appealing to group selection, but in the broad adaptationist framework there are two different ways of doing this. One is to appeal to genetic group selection, or group selection acting on genetically inherited traits. The other is to appeal to cultural group selection, or group selection acting on culturally inherited, socially learned traits. Developing a theory first put forward by biologist David Sloan Wilson (2002), psychologist Jonathan Haidt (2012) defends a theory of religious altruism according to which both genetic group selection and cultural group selection have played important roles. In contrast, other broad adaptationists have argued that group selection played a significant role only in cultural evolution (Atran and Henrich, 2010; Norenzayan and Gervais, 2012; Norenzayan, 2013). I argue here against Haidt and Wilson’s appeals to genetic group selection. Even in its most sophisticated form, the account of religious altruism based on genetic group selection is conceptually flawed and empirically implausible.

Moreover, diagnosing the flaws in this account is important not just because of what it tells us about religion, but also because of what it tells us about human evolution in general. The problems underlying this account of religious altruism reveal good reasons to doubt that genetic group selection has played any significant role at all in human evolution, but they do not pose a corresponding threat to theories based on cultural group selection. I conclude

that if group selection has played any role at all in our evolutionary history – whether for religion or anything else – then it has done so only in cultural evolution, shaping only psychological traits that individuals acquire through learning, not through genes.

Of course, to show that some trait did not evolve through genetic group selection is not to show that it did evolve through cultural group selection instead. Narrow adaptationists continue to deny any role at all for either group selection or cultural selection (Pinker, 2010, 2012; Krasnow et al., 2012; Baumard and Boyer, 2013; Baumard et al., 2013), whether for religion or anything else, so while they will agree with my treatment of genetic group selection, they will claim that appeals to cultural group selection should be rejected as well. I disagree, but this is not an issue I take up here. It is beyond the scope of my argument to provide a general defense of the broad adaptationist framework, or to rehearse the positive case for the theory of cultural group selection and its application to religion. Fortunately, all of this has been done elsewhere, and a growing number of social scientists from across the disciplines of psychology, anthropology, biology, economics and philosophy continues to marshal evidence in favor of this approach (for cultural group selection in general: Boyd and Richerson, 2002; Henrich, 2004; Richerson and Boyd, 2005: Chapter 6; Bell et al., 2009; Turchin, 2010; Baldini et al., 2013; Sterelny, 2014; for theories of religion based on cultural group selection: Bulbulia and Mahoney, 2008; Atran and Henrich, 2010; Henrich et al., 2010; Norenzayan and Gervais, 2012; Norenzayan, 2013; Norenzayan et al., 2014). I make the assumptions of broad adaptationism here for the sake of argument, in order to address a conflict between two broad adaptationist theories. Accordingly, while I assume that cultural evolution and group selection have both occurred, I do not defend that assumption here.

This article is structured as follows. After identifying what the relevant target traits are in Section 2, Section 3 describes Haidt's account of how these traits evolved. In Section 4, I identify conceptual problems in Haidt's evolutionary reasoning, and show that they lead to a crucial empirical claim that Haidt does not address, and that related evidence suggests is false. I conclude by discussing some important implications for religion, human altruism, and human evolution in general.

2 What is Religious Altruism?

By “religious altruism” I mean behavior that is (1) altruistic and (2) caused specifically by religious belief and practice. Except for the notoriously difficult problem of defining the concept of religion itself, another issue I will not attempt to address here, I assume that what it means for a behavior to be

caused specifically by religious beliefs and practices is sufficiently clear. I will add only that I generally take religion to constitute a partial cause of the behaviors in question, one factor among others.

I have much more to say about what it means for a behavior to be altruistic, because this issue has become quite complicated in recent years. I am not referring here to already well-known complications arising from the ambiguity between psychological altruism, which is defined in terms of the mental states that cause altruistic behavior, and evolutionary altruism, which concerns the consequences of altruistic behavior for genetic fitness (Sober, 1994). In case it does not go without saying, the focus here will be exclusively on evolutionary altruism.

Evolutionary theories of altruism began in biology, and were advanced by narrow adaptationists who were not particularly concerned with human cognition and behavior. But the evolution of altruism has now become a prominent issue in the social sciences as well, among broad adaptationists who are concerned specifically with human cognition and behavior. And in the social scientific context, the term “altruism” is often treated as equivalent in meaning to two other key terms: “prosociality” and “cooperation”. In the original theories of altruism from biology, however, the concept of altruism itself is not defined in a way that is equivalent to these other two terms. As a result, broad adaptationists have ushered into the literature (unwittingly, as far as I can tell) a new, competing definition of altruism. So in order to properly identify the target traits whose evolutionary history is under debate, some terminological ground clearing is needed.

In an article from 2008, psychologists Ara Norenzayan and Azim Shariff review behavioral and experimental evidence for “religious prosociality, or the idea that religions facilitate acts that benefit others at a personal cost” (Norenzayan and Shariff, 2008: 58). They then suggest that cultural group selection explains (in part) why the most successful and widespread religious systems emphasize prosocial attitudes, and succeed in producing prosocial behavior. More recently, Norenzayan (2013) has published a book-length defense of this account, which contains a more detailed discussion of the proximate psychological mechanisms connecting religious belief and practice to prosocial forms of behavior. Haidt disagrees about the nature of the evolutionary processes involved, but he does not disagree about what the relevant target traits are. The studies he cites (Haidt, 2012: 297) are the same studies that Norenzayan and Shariff discuss in their 2008 article, which include several studies of their own. Indeed, Haidt simply adopts the supernatural monitoring hypothesis that forms the core of Norenzayan’s account. In Haidt’s terms,

“Creating gods who can see everything, and who hate cheaters and oath-breakers, turns out to be a good way to reduce cheating and oath-breaking” (Haidt, 2012: 297). Both theorists routinely use the terms “altruism”, “prosociality” and “cooperation” interchangeably when referring to the suppression of behaviors such as cheating and oath-breaking.

In the literature on cooperation in humans, “cooperation” is typically used in a technical sense that is derived from a standard experimental methodology, in which economic game theory is used to measure behavior. Though particular experiments differ in all kinds of ways, what these experimental designs have in common is that they require participants to choose between two moves in a game. One move, called “defection,” represents the pursuit of self-interest. The other move is to “cooperate,” or to choose to pay some cost that will benefit others. Thus, as long as it is understood that “cooperation” is being used in this technical sense, the term does indeed have the same meaning as “prosociality”, as Norenzayan and Shariff (2008) define it above. Both refer to behavior in which one individual makes a sacrifice for the sake of at least one other individual, where the costs and benefits involved are measured in the currency of genetic fitness.

Altruism is a more problematic concept, because it now appears in the evolutionary literature in both a weak and a strong sense. In the stronger sense, the costs of altruism are not just costs, but net costs. That is, in order for a behavior to count as altruistic, the costs involved must first be weighed against any benefits the behavior might also confer, and it is only if the overall impact on fitness is negative that the behavior counts as altruism. This focus on the overall balance of costs and benefits is not shared by the weaker definition. Indeed, the weak definition makes no mention of any benefits at all. On the weak definition, the costs of altruism are just costs – not net costs. These costs may of course be offset or outweighed by various fitness benefits, but that has nothing to do with whether the behavior counts as a form of altruism. The mere existence of any costs at all is sufficient.

As a result, on the weak definition, altruistic behavior may either increase fitness or decrease fitness overall. It is deemed costly either way, but this means that it is costly only in the same mundane sense in which all behavior and cognition are costly: at the very least, there are metabolic costs to pay. Thus, like any other behavior, the weak definition of altruism simply leaves it an open question whether any particular form of altruistic behavior will ultimately be selected for or selected against. By contrast, on the stronger definition this question is not left open. For a behavior to count as altruism it must be costly overall, and this entails that selection can only work against it. Thus, altruistic

behavior is, by definition, costly in a special way that is different from the costliness of other kinds of behavior.

The strong definition came into prominent use along with the first major breakthroughs in the study of altruism, which were developed by biologists in response to a puzzle that seemed to challenge the theory of selection itself. The puzzle is this: If altruistic behavior is costly by definition, and if the relevant costs and benefits are measured in the currency of genetic fitness, then altruistic behavior must decrease fitness by definition. If so, however, it follows that altruistic behavior must have been selected against from its very inception, and so it should never have evolved. Yet it seems that altruistic behavior has evolved, both in humans and many other animal species. Biologists responded to this puzzle by showing that, contrary to appearances, purportedly altruistic behavior actually has not evolved. Their claim is not that the behaviors in question do not occur, but rather that they do not count as forms of altruism. Instead, apparently altruistic behavior actually reduces fitness only in the short term, or only in some specific way, while it also increases fitness in the long run, or in some other way such that the benefits outweigh the costs. This explains how the behavior evolved, but in a way that is said to strip it of its status as a form of altruism. For if the behavior is ultimately beneficial, then it does not decrease fitness after all. And if it does not decrease fitness, then it is not actually costly in the relevant biological sense. But if it is not costly, then it does not count as altruism. So the solution to the puzzle is not to explain how altruism evolved, but to explain altruism itself away. As Trivers famously put it, "Models that attempt to explain altruistic behavior in terms of natural selection are models designed to take the altruism out of altruism" (Trivers, 1971: 35). And as Ghiselin famously put it, "Scratch an 'altruist' and watch a 'hypocrite' bleed" (Ghiselin, 1974: 247).

Both the puzzle of altruism and the theories devised to solve it depend on the strong definition of altruism. Only if altruism reduces fitness overall does the puzzle arise in the first place, because if the definition of altruism leaves open whether a behavior is ultimately selected for or against, then the theory of selection does not yield the puzzling, problematic prediction that altruism should not have evolved. And it is only if altruism reduces fitness overall that the puzzle can be solved by explaining altruism away, since it is explained away precisely by showing that behaviors that are costly are nevertheless not costly overall.

Accordingly, when social scientists working from the broad adaptationist framework treat the concept of altruism as equivalent in meaning to the concepts of prosociality and cooperation, they are employing a definition of altruism according to which the puzzle that originally motivated biologists does not

even arise. If altruism is defined in the weak sense, then the theory of selection does not yield the prediction that altruism should never have evolved, and there is no longer anything puzzling about the fact that it did evolve. This is important, because it shows that social scientists working from the broad adaptationist framework are interested in the phenomenon of altruism for reasons very different from those that motivated early theories of altruism by biologists working from the narrow adaptationist framework.

While David Sloan Wilson is a biologist, and while it is his theory of religion that Haidt adopts and develops, in this context Wilson is acting as a social scientist, in the sense that he is focusing specifically on human social behavior and cognition. Biologists responding to the puzzle of altruism were not setting out specifically to explain human behavior. They were interested in saving the theory of genetic selection in general from what appeared to be a false prediction about the behavior of many different species. As a result, they had no reason to be interested in behavior that increased fitness overall. Given their theoretical concerns, it was only because altruism was special – different from other forms of behavior – that it presented an interesting issue, and what was special about it was specifically that it decreased fitness overall, by definition. So to remove from the definition of altruism the requirement that it reduce fitness overall would be, from their point of view, to remove what it is about altruism that makes it a phenomenon worthy of study. If altruistic behavior is just costly in the same weak sense in which all behavior is costly, then there is nothing special about altruism per se that stands in need of explanation.

This is not the case for social scientists working from the broad framework. They are interested, in the first instance, in explaining human social behavior and cognition. And for these purposes, behavior that increases fitness is every bit as important as behavior that decreases fitness. In other words, even if it is true that all social behavior is costly somehow, it is also true that social behavior that ultimately increases fitness stands just as much in need of explanation as social behavior that ultimately decreases fitness. What is special about altruistic behavior is not that it is costly in some special, overall sense, but that its costs go to benefit others. And what is often special about religious behavior is that it is especially costly. But if what one sets out to explain is human altruism, the role of religious psychology in producing it, then there is little point in adopting a concept of altruism that carves up the domain of target behaviors specifically on the basis of their net impacts on fitness. All of the relevant behaviors are costly, and while some of them increase fitness, some of them decrease fitness instead. All of this behavior needs to be explained.

In light of all this, I must state explicitly that, because I am addressing a debate occurring within the broad adaptationist framework, I will be adhering

strictly to the weak definition of altruism. “Religious altruism”, as I define it, is costly by definition. But while its costs are measured in the currency of the genetic fitness, to say it is costly is not to say that it is costly overall. It may well be costly only in the short term, or only in some specific way, and so it may well increase fitness overall.

3 Haidt’s Appeals to Genetic Group Selection

The basic, general form of the group-selectionist theory of altruism holds that altruistic traits spread by increasing the fitness of the altruist’s group relative to other groups, even if those traits simultaneously reduce the fitness of the altruist as an individual, relative to other individuals. If group-level benefits outweigh individual-level costs, then altruism can evolve in spite of its costs (for a detailed survey of the general debate over levels of selection, see Okasha, 2006). In a collaborative paper written with biologist Edward O. Wilson (2007), David Sloan Wilson offers a pithy summary of what he hopes will become “sociobiology’s new theoretical foundation”: “Selfishness beats altruism within groups. Altruistic groups beat selfish groups. Everything else is commentary.” And in other work, both of these Wilsons take this principle to explain altruism in humans in the same way that it explains altruism in other species, most notably ants, bees and wasps (E.O. Wilson 1978, 2012; D.S. Wilson, 2002).

David Sloan Wilson (2002) exploits this reasoning to develop a theory of religion. He gathers a wealth of evidence supporting the claim that successful religions increase altruistic behavior. For example, he describes the difference between the responses of Christians and pagans to a plague in Rome around 260 CE. While pagans fled, and pushed the sick out onto the streets before they were dead, many Christians fulfilled their religious duty to care for the sick, making martyrs of themselves in the process (D.S. Wilson, 2002: 153). This example does not show that these beliefs actually increased the fitness of the Christian group, because it does not show that Christians only provided nursing care to other Christians. But it does show that Christian beliefs succeeded in motivating an exceedingly costly form of altruism. The theory of group selection is then offered as an explanation for why such costly traits have not only persisted in spite of their costs, but have actually become far more frequent in the overall population over time.

Haidt follows D.S. Wilson in appealing to religion, but he also follows E.O. Wilson in placing heavy emphasis on the similarity between humans and eusocial insects. Haidt (2012) adopts the social behavior of bees as an emblem of the power of group selection, and he offers a slogan describing the role of

this process in human evolution: “We are 90% chimp and 10% bee”. This claim is puzzling, however, because of a crucial dissimilarity between humans and bees. What group selection would explain in bees would be the evolution of altruistic genotypes; bees inherit altruistic traits through their genes. But even when humans inherit the specific religious beliefs and practices of their own parents, they do not do so by inheriting genes for those beliefs and practices. In the same way, genetic inheritance does not explain why people speak the same language as their parents. Moreover, Haidt’s appeals to cultural evolution show that he takes such traits to be transmitted culturally, through social learning, rather than genetically. This difference between humans and bees thus generates important questions about what it means for humans to be 10% bee.

The fundamental difference between cultural selection and genetic selection lies in the mechanism through which traits are inherited. The key insight of cultural evolutionary theory, and of the methodology of cultural evolutionary modeling, is that social learning may be treated as a non-genetic form of inheritance – an alternative mechanism by which psychological traits may spread through populations. Mathematical models based on this notion of inheritance rigorously describe a form of natural selection that is wholly distinct from genetic selection (cf., Richerson and Boyd, 2005). While the term “natural selection” is sometimes reserved specifically for the process of genetic selection, I see no good reason for this. Perhaps the most important thing about marking off a specifically natural form of selection process is that it removes all teleology from the evolutionary process; there are no goal states toward which evolution itself is progressing, because there is no designing agent involved in doing the selecting. In this important sense, cultural selection is just as natural as genetic selection.

It is thus crucial to recognize that the basic, general form of the group-selectionist theory of altruism remains neutral with regard to the mechanism of inheritance for altruistic traits. Altruistic behavior may be selected for at the group level whether individuals acquire altruistic traits genetically or through learning.

Accordingly, when Haidt claims that human evolution is similar to bee evolution, one might hope what he means is just that the logic of group-level selection is the same regardless of how the relevant traits are inherited. The claim that group selection explains altruism in both bees and humans is entirely compatible with the claim that bees inherit altruistic traits genetically, while humans have to learn to be altruistic. But Haidt thinks the similarity runs deeper, encompassing not just the general logic of group selection, but the mechanism of inheritance as well. That is, like D.S. Wilson (2002: 155, 175), he appeals not just to group selection, but to genetic group selection.

Haidt describes two general features of moral psychology that he thinks have been influenced by genetic group selection. First, he posits a mechanism called the hive switch, a “group-related adaptation” (Haidt, 2012: 258) to which he devotes a full chapter of his book. The hive switch is meant to be a capacity to “transcend self-interest and lose ourselves (temporarily and ecstatically) in something larger than ourselves” (Haidt, 2012: 283). And according to Haidt, it “cannot be explained by selection at the individual level,” because it is “for making groups more cohesive, and therefore more successful in competition with other groups” (Haidt, 2012: 259). Clearly, the hive switch is meant to be part of the 10% of our psychology that is bee-like, having evolved by the same process of genetic group selection that explains altruism in bees.

The other important effect of genetic group selection on moral psychology, according to Haidt, is quantitative rather than qualitative. At the core of his theory of moral psychology are six basic capacities that he calls moral foundations. Concerns about fairness, for example, are explained by appeal to a fairness foundation: a package of cognitive and motivational capacities that allows people to track the social distribution of goods and respond in functionally appropriate ways. The other foundations he posits are care, loyalty, authority, sanctity and liberty. And because Haidt is laudably concerned that this model should explain more than just the judgments of American university students, the moral foundations are posited as universal traits that originally evolved by individual-level genetic selection (the 90% that is part chimp). So while Haidt supposes that these traits were already in existence before group selection began to take effect, he also argues that, in recent evolutionary history, group selection has modified the relative levels of activity exhibited by various foundations (Haidt, 2012: 216).

But while he makes these claims about genetic group selection, Haidt also assumes that the moral foundations of individuals may be influenced in substantial ways by social learning, making them subject to cultural evolution. He does not view genetic group selection as a process in which group selection occurs in genetic evolution instead of cultural evolution. Rather, he sees it as a process in which genetic evolution occurs through cultural evolution. Here he takes himself to be appealing to the theory of gene-culture coevolution developed by Robert Boyd and Peter Richerson, according to which cultural evolution can exert a causal influence upon genetic evolution. However, as I explain in what follows, Haidt misinterprets the theory of gene-culture coevolution in an understandable but fundamental way. For cultural evolution to causally influence genetic evolution is not for genetic evolution to occur through cultural evolution, in the way he supposes. As a result, he mistakenly sees gene-culture coevolution as a way of integrating appeals to cultural group selection

with appeals to genetic group selection. In fact, however, gene-culture coevolution provides reasons to reject appeals to genetic group selection altogether, as Boyd and Richerson themselves argue.

4 Group Selection in Humans

Whereas Haidt follows Wilson and Wilson in treating group selection in humans as an instance of a more general biological pattern, Boyd and Richerson take precisely the opposite approach. They too appeal to group selection to explain altruism, but they motivate this appeal specifically by pointing out that group selection explains what is unique about the altruistic behavior of humans. They first note that altruism in other species occurs only in close-knit groups composed of genetically related individuals, although a small number of trusted partners in reciprocal exchange are sometimes included as well. They then point out (Boyd and Richerson, 2006: 108) that altruism in human cooperation stands out starkly against this background:

In contrast to the societies of other animals, virtually all human societies are based on the cooperation of large numbers of unrelated people. This is obviously true of modern societies in which complex tasks are managed by enormous bureaucracies like the military, political parties, churches, and corporations . . . Even contemporary hunter-gathers who are limited to the least productive parts of the globe have extensive exchange networks and regularly share food and other important goods outside the family. Other animals do none of these things.

Not all of this cooperation is altruistic, of course, but much of it is. For example, much to the chagrin of social conservatives, governmental welfare systems often benefit “have-nots” at the expense of “haves” who are not their genetic relatives, and whom they will never even meet. Further, people often donate anonymously to charities, tip at restaurants they do not expect to return to, and give up seats to pregnant women on buses. None of these forms of altruism can be explained by genetic relatedness or reciprocity, and none occur in other species.

The specialness or uniqueness of human altruism constitutes what Boyd and Richerson (2006) call “the puzzle of human cooperation,” and they offer a solution to the puzzle that they call the tribal instincts hypothesis. Haidt takes himself to be adopting this hypothesis, but the view he actually adopts is distorted by his misinterpretation of the theory of gene-culture coevolution.

4.1 *Misinterpreting the Theory of Gene-Culture Coevolution*

That genetic evolution causally influences cultural evolution is a mundane point. It is no surprise that genetic selection acting on our frugivorous primate ancestors has had some effect upon cultural traditions for preparing sweet foods, since a diet based on fruit largely explains why humans like sugar in the first place. What is striking about gene-culture coevolution, however, is that this causal relation can also run in the other direction: prior events in cultural evolution can also affect the subsequent course of genetic evolution. Since culturally inherited traditions can sustain stable forms of behavior that persist for centuries and even millennia, and since some of these forms of behavior have serious consequences for individuals' genetic fitness, cultural evolution can create social environments in which genetic selection pressures are strong enough, and last long enough, to influence genetic evolution.

Perhaps the best-documented example concerns the evolution of lactose tolerance among adults (Richerson and Boyd, 2005: 192; Holden and Mace, 2009; Laland et al., 2010). Before humans began keeping livestock, they were unable to digest milk beyond the first few years of age. But when people in some groups established cultural traditions for domesticating livestock that produced milk, an environment evolved in which a new source of nutrition was readily available – if only it could be digested. In other words, the cultural evolution of domestication practices created a selection pressure for genes enabling the digestion of lactose. Eventually, the genes responded, and today adults in these cultures (but not all cultures) have genes for digesting dairy products.

Since the distinction between cultural evolution and genetic evolution is based upon a prior distinction between genetic inheritance and cultural inheritance, it follows by definition that any genetically inherited trait that evolves by group selection is a product of genetic group selection, not cultural group selection. Haidt, however, interprets the thesis of gene-culture coevolution in such a way that cultural group selection can produce genetic adaptations (Haidt, 2012: 210):

If cultural innovations (such as keeping cattle) can lead to genetic responses (such as adult lactose tolerance), then might cultural innovations related to morality have led to genetic responses as well? Yes. Richerson and Boyd argue that gene-culture coevolution helped to move humanity up from the small-group sociability of other primates to the tribal ultrasociality that is found in all human societies. According to their “tribal instincts hypothesis,” human groups have always been in competition to some degree with neighboring groups. The groups that figured out

(or stumbled upon) cultural innovations that helped them cooperate and cohere in groups larger than the family tended to win these competitions (just as Darwin said).

Haidt is correct to say that, according to the tribal instincts hypothesis, cultural traits (“innovations”) can be selected for at the group level. But this would be cultural group selection, not genetic group selection, and Haidt does not recognize this. Instead, he focuses on the “genetic responses” to group-selected cultural traits, drawing the conclusion that gene-culture coevolution makes it possible for genetically inherited traits (e.g., the hive switch and the innate settings of moral foundations) to be selected for at the group level.

The problem here is that while both genetic selection and group selection are occurring in this coevolutionary process, the traits selected for at the group level are not the traits that are inherited genetically. Like any other coevolutionary process, gene-culture coevolution involves two distinct processes of selection, each of which explains the evolution of different traits, which are inherited in different ways. The tribal instincts hypothesis holds that cultural evolution for one trait (a culturally inherited one) may subsequently influence genetic evolution for a different trait (a genetically inherited one). But even if selection for the relevant cultural trait takes place at the group level, this would imply nothing at all about which level of selection is responsible for the relevant genetic trait.

In the same way, many traits of flowers have coevolved with traits of the bees that pollinate them. But if the relevant bee traits were to evolve by group selection, this would obviously be no reason to suppose that the traits of the flowers also evolved by group selection. Whether they did or not would be an entirely separate empirical question. In such a case, all of the relevant traits are inherited genetically, yet distinct processes of selection explain bee traits and flower traits, and these distinct processes may have occurred at different levels. In gene-culture coevolution, the difference between the two co-evolving traits is in one sense even more pronounced, since they are not even inherited in the same way. Traits acquired by learning are distinct from traits acquired genetically, so distinct processes of selection are needed to explain the evolution of each. Just as in the coevolution of bee traits and flower traits, the fact that one trait evolves by group selection implies nothing at all about the level of selection responsible for the other trait.

According to the tribal instincts hypothesis, group selection occurs only for traits inherited by cultural learning. And while cultural group selection may then subsequently influence the course of genetic evolution, the resultant process of genetic evolution is said to take place at the individual-level,

not at the group level. Indeed, Richerson and Boyd explicitly argue against appeals to genetic group selection in humans, as I explain in the next section. Haidt clearly disagrees with the conclusion of this argument, but he offers no response to it.

4.2 *Against Genetic Group Selection in Humans*

Richerson and Boyd's argument against genetic group selection in humans is based on the role played by variation in selectionist explanations, and it is a fairly technical argument. Accordingly, a brief digression on variation is necessary.

In order for selection to be what explains the frequency of a trait in a population, there must have been distinct variants of that trait to select from. Put metaphorically, if Mother Nature does not have distinct options to select from, then her role as Selector does not do any explanatory work. This is a matter of degree; the greater the variation, the more significant the explanatory role of selection. If the giraffes in a population have wildly different neck lengths, then even weak selection can be an important part of the explanation for changes in average neck length. But if all the giraffes have necks of almost exactly the same length, then only very strong selection pressures will be capable of making such minor variations significant in an explanation for changes in average neck length.

Moreover, group-level selection requires group-level variation. If variation exists within a group – such that members of the same group possess different variants – then selection at the group level will not be selection for some variants over others; when the whole group gets selected, all of these variants are selected together, in a package deal. Variation within the group is variation among individuals, and it can only support individual-level selection. For selection to take place at the group level, members of the same group must possess the same variant, while members of different groups possess different variants. As it is often put, variation must be “concentrated” at the group level.

For the same reasons, genetic selection requires genetic variation. If the traits that vary are inherited culturally, rather than genetically, then by definition group selection acting on those traits would be cultural group selection. Of course, for any given psychological trait it is an open empirical question how individuals inherit it, and for some traits this question is quite difficult. But to explain a trait by appeal to selection is to assume that this empirical question has already been closed. To say a trait is a product of genetic group selection already entails that variants of this trait are inherited genetically, while to say it is a product of cultural group selection already entails that it is inherited by social learning. Since gene-culture coevolution necessarily involves two

distinct traits, appeals to this process require identifying two distinct forms of phenotypic variation: one due to genetic differences, another due to differences among individuals' histories of learning.

Given all of this, one major reason why Boyd and Richerson appeal to cultural group selection is that cultural evolution tends to concentrate variation at the group level. Cultural inheritance causes individuals to become more similar to one another with regard to a particular trait, thereby reducing the variation between those two individuals. When the individuals in question happen to be members of the same group, this reduces variation within the group. And when the trait in question is not also present in other groups, this decrease in variation at the individual level is not accompanied by a corresponding decrease in variation between groups. That is what it means to concentrate variation at the group level. And since individuals often do, in fact, acquire from fellow group members beliefs and practices that do not appear in competing groups, cultural evolution does tend to concentrate variation at the group level, at least with regard to certain traits.

This shows that group selection can occur in cultural evolution, but it also shows that the same question arises regarding genetic evolution as well. Since group-level genetic selection would require group-level genetic variation, how much group-level variation do we actually observe among human genes? By appealing to genetic group selection, Haidt is making a substantial empirical commitment. However, he never addresses this question, and he provides no evidence that the required form of variation exists.

Boyd and Richerson do consider the question, and they argue that group-level genetic variation is virtually non-existent in humans. For any kind of genetic selection to occur, the relevant genetic variation must be continuously maintained throughout the course of the selection process. But genetic selection takes a very long time, and during that time some individuals will migrate from one group to another, bringing their genes with them. This reduces genetic variation at the group level, concentrating it at the individual level instead. The question then becomes a matter of degree: if rates of migration are relatively low, or if group-level selection is relatively strong, then group-level genetic variation can be maintained, and genetic group selection may occur, in principle. In fact, however, Richerson and Boyd cite modeling work showing that this is not what happens: "Even very small amounts of migration are sufficient to reduce genetic variation between groups to such a low level that group selection is not important" (Richerson and Boyd, 2005: 203).

In addition, a recent study found that group-level genetic variation in humans is no higher than group-level genetic variation in chimpanzees (Langergraber et al., 2011). Since chimpanzees are not especially altruistic, this

bodes ill for the hypothesis that genetic group selection explains why humans are especially altruistic. Chimpanzee groups also face conditions of significant intergroup competition, but chimpanzees do not engage in widespread altruism toward strangers. If these psychological differences between humans and chimps were the result of genetic group selection in humans, as Haidt (2012) suggests, then we should observe higher levels of group-level genetic variation among humans than we do among chimps. According to this study, we do not.

Once the errors in Haidt's evolutionary reasoning are identified, we are left with no good reason to suppose that genetic group selection explains any form of altruism in humans. Indeed, what we are left with are positive reasons to think otherwise. And because this argument against genetic group selection is based on general concerns about the amount of group-level genetic variation in humans, the problem with Haidt's appeal to genetic group selection is not specific to religiously motivated forms of altruism, or even to altruistic traits in general. Any theory that purports to explain any trait in humans via genetic group selection will have to overcome this obstacle. Appeals to cultural group selection, however, do not incur this burden.

5 Conclusion

All of this substantially deflates the content of Haidt's claim that we are 10% bee. If group selection does explain altruism in bees, then humans may indeed have something important in common with them: group selection may explain why both species are more altruistic than most other species. If this were all it meant for humans to be 10% bee, then the metaphor would be apt. But the purpose of identifying the role of group selection in human evolution is not merely to explain the fact that humans are altruistic to a high degree, so to explain what humans have in common with bees is not to explain most of the facts about human altruism that most need to be explained.

Most importantly for present purposes, we learn nothing from bee evolution about the role of religion in producing altruism. Since religious beliefs and practices play no role in producing the altruistic behavior of bees, similarities between bee evolution and human evolution are beside the point. Yet the phenomena that Haidt first set out to explain were facts about religious psychology. These distinctly human forms of cognition and behavior were the target traits that led him to appeal to group selection in the first place. I fully endorse the general approach of explaining human psychology in terms of the same biological principles that explain cognition and behavior in other species, but this is not the right tool for the job of explaining forms of cognition and behavior that do not occur in other species.

Indeed, precisely because religion is a uniquely human form of thought and behavior, explaining its evolution requires emphasizing the differences between human evolution and evolution in other species. Boyd and Richerson's tribal instincts hypothesis is just the right kind of tool for this job. But while Haidt recognizes the importance of this theory, his emphasis on similarities between humans and other species leads him to misinterpret it in a way that erases the crucial differences. An evolutionary process in which cultural group selection leads to gene-culture coevolution would be very different from the process of genetic group selection, even if both types of process would ultimately lead to high levels of altruism.

From the broad adaptationist perspective, it is only because religious altruism is a cultural phenomenon that group selection may affect its evolution. Human individuals inherit their particular religious beliefs and practices non-genetically, by learning from parents, siblings, friends, religious leaders and many others. Unlike bees, we must learn to behave altruistically, even if the relevant capacities for learning are themselves inherited genetically. Accordingly, as the debate over group selection continues between broad and narrow adaptationists, it should focus specifically on the theory of cultural group selection.

References

- Atran, S. and Henrich, J. (2010). The evolution of religion: How cognitive by-products, adaptive learning heuristics, ritual displays, and group competition generate deep commitments to prosocial religions. *Biological Theory* 5, 18-30.
- Barrett, J.L. (2004). *Why would anyone believe in god?* AltaMira Press, Plymouth.
- Baumard, N., Hyafil, A., Morris, I. and Boyer, P. (2014). Increased affluence explains the emergence of ascetic wisdoms and moralizing religions. *Current Biology* 25, 1-6.
- Baumard, N. and Boyer, P. (2013). Explaining moral religions. *Trends in Cognitive Sciences* 17, 272-280.
- Baumard, N., Andre, J.-B. and Sperber, D. (2013). A mutualistic approach to morality: The evolution of fairness by partner choice. *Behavioral and Brain Sciences* 36, 59-78.
- Bell, A., Richerson, P. and McElreath, R. (2009). Culture rather than genes provides greater scope for the evolution of large-scale human prosociality. *Proceedings of the National Academy of Sciences of the United States of America* 106, 17671-17674.
- Bering, J.M. (2006). The folk psychology of souls. *Behavioral and Brain Sciences* 29, 453-461.
- Bloom, P. (2007). Religion is natural. *Developmental Science* 10, 147-151.
- Boyd, R. and Richerson, P. (2002). Group beneficial norms can spread rapidly in a structured population. *Journal of Theoretical Biology* 215, 287-296.

- and ———. (2006). Solving the puzzle of human cooperation. In Levinson, S. and Jaisson, P. (Eds), *Evolution and culture*. MIT Press, Cambridge, MA, pp. 105-132.
- Boyer, P. (2001). *Religion explained: The evolutionary origins of religious thought*. Basic Books, New York, NY.
- Bulbulia, J. and Mahoney, A. (2008). Religious solidarity: The hand grenade experiment. *Journal of Cognition and Culture* 8, 295-320.
- Ghiselin, M.T. (1974). *The economy of nature and the evolution of sex*. University of California Press, Berkeley, CA.
- Haidt, J. (2012). *The righteous mind*. Pantheon Books, New York, NY.
- Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization* 53, 3-35.
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A. et al. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science* 327, 1480-1484.
- Holden, C. and Mace, R. (2009). Phylogenetic analysis of the evolution of lactose digestion in adults. *Human Biology* 81, 597-619.
- Krasnow, M., Cosmides, L., Pedersen, E.J. and Tooby, J. (2012). What are punishment and reputation for? *PLoS One* 7, e45662.
- Laland, K.N., Odling-Smee, J. and Myles, S. (2010). How culture shaped the human genome: Bringing genetics and the human sciences together. *Nature Reviews Genetics* 11, 137-148.
- Langergraber, K., Schubert, G., Rowney, C., Wrangham, R., Zommers, Z. and Vigilant, L. (2011). Genetic differentiation and the evolution of cooperation in chimpanzees and humans. *Proceedings of the Royal Society London Series B: Biological Sciences* 278, 2546-2552.
- Norenzayan, A. (2013). *Big gods: How religion transformed cooperation and conflict*. Princeton University Press, Princeton, NJ.
- Norenzayan, A. and Gervais, W.M. (2012). The cultural evolution of religion. In Slingerland, E. and Collard, M. (Eds), *Creating consilience: Integrating the sciences and the humanities*. Oxford University Press, New York, NY, pp. 243-265.
- Norenzayan, A. and Shariff, A.F. (2008). The origin and evolution of religious prosociality. *Science* 322, 58-62.
- Norenzayan, A., Shariff, A.F., Gervais, W.M., Willard, A.K., McNamara, R.A., Slingerland, E. and Henrich, J. (2014). The cultural evolution of prosocial religions. *Behavioral and Brain Sciences*, 1-86.
- Okasha, S. (2006). *Evolution and the levels of selection*. Oxford University Press, New York, NY.
- Pinker, S. (2012). The false allure of group selection. Available online at <http://edge.org/conversation/the-false-allure-of-group-selection> (retrieved 7 November 2013).

- . (2010). The cognitive niche: Coevolution of intelligence, sociality, and language. *Proceedings of the National Academy of Sciences of the United States of America* 107 (Suppl. 2), 8993-8999.
- Richerson, P., Baldini, R., Bell, A., Demps, K., Frost, K., Hillis, V., et al. (2014). Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, 1-71.
- Richerson, P.J. and Boyd, R. (2005). *Not by genes alone*. The University of Chicago Press, Chicago, IL.
- Sober, E. (1994). *From a biological point of view*. Cambridge University Press, Cambridge.
- Sterelny, K. (2014). Cooperation, culture, and conflict. *The British Journal for the Philosophy of Science*, doi:10.1093/bjps/axu024.
- Trivers, R.L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* 46, 35-57.
- Turchin, P. (2010). Warfare and the evolution of social complexity: A multilevel-selection approach. *Structure and Dynamics* 4, Article 2.
- Wilson, D.S. (2002). *Darwin's cathedral: Evolution, religion and the nature of society*. University of Chicago Press, Chicago, IL.
- Wilson, D.S. and Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and Brain Sciences* 17, 585-608.
- Wilson, D.S. and Wilson, E.O. (2007). Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology* 82, 327-348.
- Wilson, E.O. (1978). *On human nature*. Harvard University Press, Cambridge, MA.
- . (2012). *The social conquest of earth*. Liveright, New York, NY.