

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

Taylor Davis
Department of Philosophy
The University of British Columbia

Abstract

Religious altruism is altruistic behavior caused specifically by religious belief. In the scientific literature on religion, two different accounts of religious altruism appeal to group selection. The difference between them is the type of evolution said to be involved—genetic evolution or cultural evolution. Recently, Jonathan Haidt has defended the most fully developed account appealing to genetic group selection. I argue that problems with this account expose general reasons for skepticism about genetic group selection in humans. An important upshot is that controversy over group selection in biology need not threaten appeals to group selection in social science.

Keywords

Group Selection; Cultural Evolution; Altruism; Religion; Sociobiology

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 types of selection to which they appeal. Some adhere closely to the biological orthodoxy of inclusive fitness theory, appealing only to genetic selection occurring at (or below) the level of the individual (e.g., Boyer 2001; Barrett 2004; Bering 2006; Bloom 2007). Others adopt a broader adaptationist perspective, appealing to cultural evolution in addition to genetic evolution, and to multilevel selection instead of selection at only the individual level (e.g., Wilson 2002; Haidt 2012; Atran and Henrich 2010; Norenzayan and Shariff 2008; Bulbulia and Mahoney 2008). 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 natural selection that could potentially explain the evolution of a given trait. I thus distinguish between the theoretical framework of *broad adaptationism*, which recognizes four distinct forms of selection, and *narrow adaptationism*, which recognizes only one.

Naturally enough, narrow adaptationists tend to focus specifically on religious traits that are amenable to explanation by genetic selection at the individual level, while broad adaptationists focus on other kinds of traits. Here I consider a religious trait that narrow adaptationists have not attempted to pursue (perhaps for principled reasons, as I explain below), but for which two competing accounts have been offered by broad adaptationists. The trait in question is *religious altruism*—altruistic behavior that is caused specifically by religious beliefs and practices. Both theories attempt to explain this trait by appealing to group selection, but in the broad adaptationist framework there are two different ways of doing this. One is to claim that

group selection explains the *genetic* evolution of religious altruism. The other is to claim, instead, that group selection explains the *cultural* evolution of this trait.

I argue here against the gene-based approach: even in its most sophisticated form, the account of religious altruism based on genetic group selection is conceptually flawed. Moreover, diagnosing this flaw is important, because its implications extend beyond the study of religion to the study of human evolution more generally.

Biologist David Sloan Wilson (2002) was the first to appeal to genetic group selection in a theory of religion, but his account is no longer the most fully developed version of this approach. Psychologist Jonathan Haidt (2012) has recently adopted Wilson's account, and integrated it into his own theory of moral psychology. This brings to Wilson's approach a rich, detailed description of the psychological traits being posited as targets of group selection, allowing Haidt to say much more about precisely what it is that is said to be selected at the group level. More importantly, however, Haidt also says more about the details, as he sees them, of the relationship between genetic evolution and cultural evolution. It's in these details that the confusion arises. Haidt makes a serious attempt to integrate his appeals to genetic group selection with his appeals to cultural evolution, but he nevertheless becomes tangled in the complex web of evolutionary dynamics that arises within the broad adaptationist framework. And when the source of the confusion is identified, it reveals general reasons to remain skeptical about genetic group selection in humans—reasons that do not apply to cultural group selection as well. As a result, the failure of Haidt's account is significant for two reasons.

First, with regard to religious altruism, it renders the account based on cultural group selection the only viable theory available. Since there are only two accounts on offer, the failure of one of them leaves only the other one to explain religious altruism. Of course, to show that a

theory is “the last man standing” is not exactly to show that it is true, and the task of marshalling positive evidence in its favor is not one I take up here. Fortunately, however, this evidence exists, and a growing number of scientists has been marshalling it elsewhere (Norenzayan and Shariff 2008; Norenzayan 2013; Atran and Henrich 2010; Henrich and Gervais 2010; Bulbulia and Mahoney 2008; Atkinson and Whitehouse 2011).

Second, the failure of Haidt’s appeals to genetic group selection also carries important implications for the wider application of the theory of group selection in social science. For while it shows that appeals to genetic group selection should be rejected, it also shows that this is no reason to reject appeals to group selection altogether. This is important, because it illustrates how social scientists may safely appeal to group selection without making questionable biological commitments.

Though he adopts D. S. Wilson’s theory of religion as the basis of his account, Haidt also follows Edward O. Wilson, founder of the field of sociobiology, in emphasizing evolutionary similarities between humans and other species, especially eusocial insects (E. O. Wilson 1978, 2012). All three of these theorists take group selection in humans to be an instance of a more general biological pattern, and all three take group selection to explain altruism in humans in the same way that it explains altruism in other species. But in biology appeals to group selection are extremely controversial, even for eusocial insects (cf. Wilson and Sober 1994; Okasha 2006, ch. 6). When E. O. Wilson and his collaborators published their group-selectionist account of eusociality (Nowak, Tarnita, and Wilson 2010), the number of authors included on a response paper objecting to it (Abbott et al. 2011) made the response look as much like a petition as a theoretical objection. So to explain human altruism by appeal to this more general biological

pattern is to make the account of altruism in humans dependent upon a much larger, more general issue, and one plagued by long-standing and deeply entrenched debate.

Haidt accepts this burden, devoting a full chapter of his book to a “retrial” of the theory of group selection, which he claims has been “falsely convicted and unfairly banished” in biology (p. 222). But if it is true that genetic group selection does not occur in humans, then it isn’t true that human altruism is explained by this more general biological pattern, and appeals to group selection in humans do not depend on the outcome of this more general biological debate. Analyzing the problems in Haidt’s account thus illustrate how psychologists, anthropologists, sociologists and economists may employ the theory of group selection without making biological claims that are highly controversial.

2. What is Religious Altruism?

In an article from 2008, psychologists Ara Norenzayan and Azim Shariff review the behavioral and experimental evidence for religious altruism. And in his recent book-length defense of the cultural group selection account of religion, Norenzayan (2013) provides a detailed discussion of the proximate psychological mechanisms that connect beliefs about morally concerned supernatural agents with forms of behavior that he refers to interchangeably as “altruism,” “prosociality” and “cooperation.” Haidt disagrees about how such psychological traits evolved, but he does not disagree about what the relevant target traits are. The studies he cites (2012, p. 297) are the same ones Norenzayan and Shariff mention in the 2008 article, which include some of their own studies. Indeed, Haidt simply adopts the hypothesis of supernatural monitoring 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” (2012, p. 297).

Here it becomes significant that both of these theories employ the broad adaptationist approach, because the meaning of the concept of altruism changes in the broad framework in a subtle but crucial way. To appeal to cultural evolution and group selection in addition to individual-level genetic selection is to weaken the conditions necessary for a behavior to count as an instance of altruism. This may partially explain both why narrow adaptationists have not attempted to explain religious altruism at all, and why broad adaptationists are often tempted to use other terms, such as ‘prosociality’ and ‘cooperation,’ to stand in for the term ‘altruism.’

Narrow adaptationists have famously insisted that to explain how altruism evolved is actually to explain altruism *away*. As Trivers 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). And as Ghiselin put it, “Scratch an ‘altruist’ and watch a ‘hypocrite’ bleed” (1974, p. 247). The theory of inclusive fitness has succeeded in explaining many forms of behavior that are intuitively considered altruistic according to the ordinary, common-sense concept of altruism, but narrow adaptationists have taken this explanation to show that, in fact, the relevant behaviors don’t count as forms of altruism after all. The narrow framework explains so-called ‘altruism’ by showing that the behaviors in question aren’t actually costly, and thus don’t count as altruistic, because their overall contributions to fitness are positive, not negative. This shows that the definition of altruism being employed is based on overall costs; a behavior only counts as altruistic if it yields net costs in the currency of genetic fitness.

Broad adaptationists agree in defining altruism as behavior that is costly for the agent and beneficial for some other individual, and they agree that the relevant costs and benefits are to be

measured in the currency of genetic fitness. What they do not agree about is what it means for altruism to be costly. According to broad adaptationists, to say that altruism is costly is *not* to imply that it reduces genetic fitness *overall*. In the broad framework, altruism implies costs, but it does not imply net costs. This difference is crucial, because if ‘religious altruism’ is taken to refer to behavior that confers *net* costs, as narrow adaptationists assume, then it is no longer clear that religious altruism even exists. Accordingly, it is easy to see why narrow adaptationists have not bothered to offer a theory of religious altruism.

Altruism once posed an important challenge to adaptationist explanations of behavior in general, and narrow adaptationists met this challenge precisely by drawing a distinction between costs and net costs. At that time, narrow adaptationists had the choice of adopting either the weaker definition of altruism, based on simple costs, or the stronger definition, based on *net* costs. They chose the stronger definition, apparently for rhetorical reasons: saying that altruism doesn’t even exist is a striking way of emphasizing the difference between costs and net costs that forms the basis of their account. However, it is important to see that there is no principled theoretical reason why narrow adaptationists could not have chosen the weaker definition instead. In the terms of inclusive fitness theory, it makes perfect sense to say that foraging behavior, for example, is costly. In addition to metabolic costs for locomotion and cognition (if food is hard to find), foraging exposes many animals to increased risk of predation. These are all costs, and there is no principled reason why altruism could not be defined as costly in exactly the same sense in which foraging is costly. The difference is only that altruism possesses these costs by definition, while it is no part of the definition of foraging that it must be costly. Adopting this weaker definition of altruism would rob narrow adaptationists of the rhetorical effect of claiming that the behaviors they explain don’t actually count as altruistic at all, but it wouldn’t change

anything about the substance of their explanation: just like foraging, these costly behaviors evolved because the costs involved were offset by countervailing benefits.

However, since narrow adaptationists did choose the stronger definition, they are not likely to regard as altruistic the behaviors that are explained by broad adaptationist theories of religious altruism. In the studies that Norenzayan and Haidt both appeal to, participants do such things as refraining from cheating even when their cheating can't be detected, or giving money to charity or to strangers in economic games. These behaviors are obviously costly, and this suggests a clear sense in which *prosocial* behavior is something over and above ordinary social behavior. But nothing in the design of these studies shows that these forms of behavior are fitness-reducing *overall*. The relevant costs might easily be outweighed by other fitness benefits. For example, perhaps individuals who give money to strangers acquire good reputations, and perhaps this ultimately increases their fitness by increasing their access to the benefits of cooperation.

Norenzayan and Haidt are both fully aware of this, as are broad adaptationists in general. They do not assume that the behaviors they explain reduce fitness overall. Rather, when referring to such behaviors as 'altruism,' they take advantage of the fact that these behaviors are obviously costly, and thus obviously altruistic, but in a weaker sense. The fact that they are costly at all is already sufficient to justify use of the term 'altruism' according to the common-sense concept of altruism. However, it does not justify use of this term according to the technical definition adopted by narrow adaptationists a few decades ago. This may be one reason why broad adaptationists are sometimes tempted to use terms other than 'altruism,' but also remain tempted to use 'altruism' at times as well.

More importantly, unlike narrow adaptationists, broad adaptationists have no motivation to adopt the stronger definition of altruism, because they do not presuppose, a priori, that costs must

be balanced against corresponding benefits. In the narrow framework, once some prevalent form of costly behavior has been observed, the question of overall fitness arises automatically, because observed costs must be explained by corresponding benefits. But cultural evolution and group selection each provide separate ways of explaining how a costly form of behavior can evolve even if it *is* costly, overall, from the point of view of individual genetic fitness (cf. Wilson and Wilson 2007, Richerson and Boyd 2005, ch. 5). As a result, the existence of benefits to individual genetic fitness does not simply follow from the existence of costs to individual genetic fitness. It is an empirical question whether there are any such benefits, a question that will depend upon which form of selection explains how the trait evolved. In the broad framework, the question of a balance between costs and benefits in the currency of individual fitness arises only after it has been determined empirically (1) how individuals inherit the trait in question, and (2) what level of selection is responsible. For many of the traits of religion, and for many other forms of cooperative or prosocial behavior pursued by broad adaptationists, these questions do not have obvious answers.

In light of all this, it is necessary to say explicitly that in the broad adaptationist theories I will be discussing, ‘altruism’ refers to traits that are costly only in the weaker, short-term sense, not in the overall sense. This definition preserves the meaning of ‘altruism’ in ordinary, common usage, according to which giving money to a stranger is already considered altruistic regardless of its overall effect on the giver’s fitness. But this definition also conflicts with the stronger technical definition adopted by narrow adaptationists, according to which it is not clear whether giving money to strangers counts as altruism until the overall fitness value of the act has been determined. Thus, according to the definition of ‘religious altruism’ that I employ here, to say

that some behavior is altruistic is to say that it is costly only in the same sense in which foraging is costly.

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 herself, relative to other individuals. If group-level benefits outweigh individual-level costs, then altruism can evolve in spite of its costs.¹ In a collaborative paper, David Sloan Wilson and Edward O. Wilson (2007) offer the following summary of what they hope 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 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 1998, 2012; D. S. Wilson 2002).

D. S. 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 (2002, p. 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

¹ For a detailed survey of the general debate over levels of selection, see Okasha (2006).

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.

It is D. S. Wilson that Haidt follows in appealing to religion, but it is E. O. Wilson that Haidt follows in placing heavy emphasis on the similarity between humans and eusocial insects. Haidt 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 one 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 specific beliefs and practices. In the same way, genetic inheritance does not explain why people speak the same language as their parents. Haidt’s appeals to cultural evolution show that he takes traits like these to be transmitted culturally, rather than genetically, but this difference between humans and bees introduces important questions about what it means for us to be 10% bee.

In the theory of cultural evolution subscribed to by those who study religion, cultural transmission is defined as a specific form of social learning known as *cultural learning*. Cultural learning is high-fidelity copying, or imitation, and the key insight of cultural evolutionary theory is that imitation may be treated as a non-genetic form of inheritance—an alternative mechanism by which psychological traits may spread through human 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). It is thus crucial to recognize that the basic, general form of the group-selectionist account 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 their altruistic traits genetically or through cultural 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, p. 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” (2012, p. 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” (p. 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” (p. 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 their innate settings, or their relative levels of activity in various contexts (p. 216).

But while he makes these claims about genetic group selection, Haidt also assumes that the moral foundations of individuals may be affected by cultural learning, and are thus 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 co-evolution* 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 misunderstands the theory of gene-culture co-evolution. For cultural evolution to causally influence genetic evolution is not for genetic evolution to occur *through* cultural evolution, in the way Haidt imagines. As a result, he mistakenly sees gene-culture co-evolution as a way of integrating appeals to cultural evolution

with appeals to genetic group selection. In fact, however, the theory of gene-culture co-evolution provides general reasons to remain skeptical of appeals to genetic group selection in humans, as Boyd and Richerson themselves argue.

4. Why Group Selection in Humans is Unique

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 explain human altruism by appealing to group selection, but they motivate this appeal specifically by claiming that it explains what is *unique* about altruism in 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 (2006) 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 systems of social welfare 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 don't expect to return to, and give up seats to pregnant women on buses.² None of these forms of altruism can be explained by genetic relatedness or reciprocity, and none of them occurs in other species.

Such uniquely human forms of altruism constitute what Boyd and Richerson call “the puzzle of human cooperation,” and they offer as a solution to the puzzle what 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 co-evolution.

4.1 Misinterpreting Gene-culture Co-evolution

That genetic evolution causally influences cultural evolution is a mundane point. It's no surprise, for example, that cultural traditions for preparing sweet foods have been impacted causally by the events of genetic selection that explain why humans like sugar in the first place. What is striking about gene-culture co-evolution, 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 a social environment containing selection pressures that are strong enough, and that last long enough, to influence genetic evolution.

Perhaps the most well documented example concerns the evolution of lactose tolerance among adults (Richerson and Boyd 2005, p.192; Holden and Mace 2009; Laland, Odling-Smee,

² In case it isn't obvious where the fitness costs lie in giving up one's seat on a bus, note that standing is more metabolically expensive than sitting down. This is why we become “tired” —motivated to *stop* standing—after standing for long periods, and why pregnant women become tired sooner than others.

and Myles 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 milk-producing livestock, this created an environment in which a new source of nutrition was readily available—if only it could be digested. In other words, the cultural evolution of livestock domestication 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 co-evolution in such a way that *cultural* group selection produces *genetic* adaptations (2012, p. 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 co-evolution 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 acknowledge this. Instead, he focuses on the

“genetic responses” to group-selected cultural traits, drawing the conclusion that gene-culture co-evolution 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 co-evolutionary process, the traits selected for at the group level are not the traits that are inherited genetically. Like any other *co*-evolutionary process, gene-culture co-evolution involves two distinct processes of selection, each of which explains the evolution of different traits 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 co-evolved 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 co-evolution the difference between the two co-evolving traits is in one sense even more pronounced, since they aren't 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 co-evolution of bee traits and flower traits, the fact that one

trait evolves by group selection implies nothing at 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. While cultural group selection is said to influence the subsequent 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 makes no attempt to respond 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 of *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 doesn't have distinct options to select from, then her role as Selector doesn't do any explanatory work. This is a matter of degree; the greater the variation, the more significant the 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 cultural learning. Since appeals to gene-culture co-evolution always refer to two different forms of selection acting on two distinct traits, they also require reference to two distinct forms of variation: one due to genetic differences, another due to differences in learning.

Given all of this, a major reason why Boyd and Richerson appeal to cultural group selection is that cultural evolution concentrates variation at the group level. Cultural inheritance causes individuals to become more similar to each other with regard to a given psychological trait, thereby reducing the variation between those two individuals. When the individuals

involved are members of the same group, this reduces variation within the group. And when, as is often the case, the relevant trait is not also present in competing groups, this decrease in variation within the group is not accompanied by a corresponding decrease in variation between groups. That is what it means to concentrate variation at the group level.

By itself, of course, this is no evidence against genetic group selection. It shows, however, that by appealing to genetic group selection Haidt is making a substantial empirical commitment. Since genetic group selection presupposes the existence of group-level genetic variation, a crucial empirical question arises: how much group-level genetic variation is there among humans? Haidt never considers this question, and he provides no evidence for the required form of variation.

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 long time, and during that time individuals tend to migrate, bringing their genes with them across group boundaries. 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 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” (2005, p. 203).

In addition, a recent empirical 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 aren't 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 don't 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 suggests, then we should observe higher levels of group-level genetic variation among humans than we do among chimps. According to this study, we don't.

Once these errors in Haidt's evolutionary reasoning are identified, we are left with no good reason to suppose that genetic group selection explains human altruism. 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 by appealing to genetic group selection will have to overcome this obstacle. Appeals to cultural group selection 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 would appear to have *something* important in common with bees: group selection would explain why both species display higher levels of altruism 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 explaining how altruism evolved in humans is not

merely to explain why humans are more altruistic than other species, and to explain similarities between the altruism of humans and the altruism of other species is not to explain many of the features of human altruism that we most want to understand.

For example, we learn nothing from bee evolution about the role of moral education in producing altruistic behavior. And, more important for present purposes, we learn nothing from bees about the role that religion plays in producing altruistic behavior. Yet the psychological facts that Haidt originally set out to explain were facts about moral and religious psychology. These were the target phenomena that led him to appeal to group selection in the first place. So, since altruism in bees has nothing to do with morality or religion, similarities between human evolution and bee evolution are simply beside the point.

Fortunately, an account of uniquely human forms of altruism is still available, even if the evolution of altruism in humans is *not* like the evolution of altruism in bees. The reasons I've given for remaining skeptical of genetic group selection in humans do not apply to cultural group selection as well, so the role of religion and morality in producing altruistic behavior may yet be explained by group selection. Among all other species in which high levels of altruism are observed, group-level variation in traits is due to an underlying pattern of group-level variation in genes. That does not appear to be the case for humans. Yet it may still be that humans inherit their religious beliefs and moral practices non-genetically, by imitating the beliefs and practices of their parents, siblings, friends, religious leaders and others. If so, then cultural group selection may explain the role that religious and moral forms of cognition play in causing people to do costly things for the sake of others, even when those others are people they will never meet.

If cultural group selection does explain these distinctly human traits, then it might also affect the subsequent course of genetic evolution. This would not be genetic *group* selection,

however, because even if it were selection for genetically inherited traits, it would not be selection for genes that vary at the group level. A co-evolutionary process of this kind would be nothing at all like the process of genetic group selection that Haidt, D. S. Wilson and E. O. Wilson claim is the explanation for altruism in bees and other eusocial species. As a result, appeals to group selection in humans may stand (or fall) independently of the debate over group selection in biology. The distinction between cultural group selection and genetic group selection thus makes it possible for broad adaptationists in the social sciences to appeal to group selection without thereby committing themselves to more sweeping—and more tenuous—claims about the role of group selection in the rest of the biological world.

References

- Abbot, Patrick, Jun Abe, John Alcock, Samuel Alizon, Joao A. C. Alpedrinha, Malte Andersson, and Jean-Baptiste Andre, et al. (2011). “Inclusive Fitness Theory and Eusociality.” *Nature* 471 (7339): E1–E1. doi:10.1038/nature09831.
- Atkinson, Quentin D., and Harvey Whitehouse. (2011). “The Cultural Morphospace of Ritual Form.” *Evolution and Human Behavior* 32 (1): 50–62. doi:10.1016/j.evolhumbehav.2010.09.002.
- Atran, Scott, and Joseph Henrich. (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 (1): 18–30.
- Barrett, Justin L. (2004). *Why Would Anyone Believe in God?* Plymouth: AltaMira Press.
- Bering, Jesse M. (2006). “The Folk Psychology of Souls.” *Behavioral and Brain Sciences* 29 (5): 453–61.
- Bloom, Paul. (2007). “Religion Is Natural.” *Developmental Science* 10 (1): 147–51. doi:10.1111/j.1467-7687.2007.00577.x.
- Boyd, Robert, and Peter J. Richerson. (2006). “Solving the Puzzle of Human Cooperation.” In *Evolution and Culture*, ed. Stephen Levinson and Pierre Jaisson, 105–32. Cambridge: MIT Press.
- Boyer, Pascal. (2001). *Religion Explained: The Evolutionary Origins of Religious Thought*. New York: Basic Books.
- Bulbulia, Joseph, and Andrew Mahoney. (2008). “Religious Solidarity: the Hand Grenade Experiment.” *Journal of Cognition and Culture* 8 (3): 295–320. doi:10.1163/156853708X358191.
- Ghiselin, Michael. T. (1974). *The Economy of Nature and the Evolution of Sex*. Berkeley, CA: University of California Press.
- Haidt, Jonathan. (2012). *The Righteous Mind*. New York: Pantheon Books.
- Holden, Clare, and Ruth Mace. (2009). “Phylogenetic Analysis of the Evolution of Lactose Digestion in Adults.” *Human Biology* 81 (5/6): 597–619.

- Laland, Kevin N., John Odling-Smee, and Sean Myles. (2010). "How Culture Shaped the Human Genome: Bringing Genetics and the Human Sciences Together." *Nature Reviews Genetics* 11 (2): 137–48. doi:10.1038/nrg2734.
- Langergraber, Kevin, Grit Schubert, Carolyn Rowney, Richard Wrangham, Zinta Zommers, and Linda Vigilant. (2011). "Genetic Differentiation and the Evolution of Cooperation in Chimpanzees and Humans." *Proceedings of the Royal Society B: Biological Sciences* 278 (1717): 2546–52. doi:10.1098/rspb.2010.2592.
- Norenzayan, Ara, and Azim F. Shariff. (2008). "The Origin and Evolution of Religious Prosociality." *Science* 322 (5898): 58-62.
- Norenzayan, Ara. (2013). *Big Gods: How Religion Transformed Cooperation and Conflict*. Princeton: Princeton University Press.
- Nowak, Martin A., Corina E. Tarnita, and Edward O. Wilson. (2010). "The Evolution of Eusociality." *Nature* 466 (7310): 1057–1062. doi:10.1038/nature09205.
- Okasha, Samir (2006). *Evolution and the Levels of Selection*. New York: Oxford University Press.
- Richerson, Peter J., and Robert Boyd. (2005). *Not By Genes Alone*. Chicago: The University of Chicago Press.
- Trivers, Robert L. (1971). "The Evolution of Reciprocal Altruism." *Quarterly Review of Biology* 46 (1): 35–57.
- Wilson, David Sloan. (2002). *Darwin's Cathedral: Evolution, Religion and the Nature of Society*. Chicago: University of Chicago Press.
- Wilson, David Sloan, and Elliott Sober (1994). "Reintroducing Group Selection to the Human Behavioral Sciences." *Behavioral and Brain Sciences* 17 (4): 585-608.
- Wilson, David Sloan, and Edward O. Wilson. (2007). "Rethinking the Theoretical Foundation of Sociobiology." *The Quarterly Review of Biology* 82 (4): 327–48.
- Wilson, Edward O. (1978). *On Human Nature*. Cambridge: Harvard University Press.
- Wilson, Edward O. (2012). *The Social Conquest of Earth*. New York: Liveright Publishing Corp.