

# Cultural Group Selection and the Evolution of Religion

Taylor Davis

## Abstract

The role of group selection in the evolution of religion cannot be properly understood without carefully distinguishing cultural group selection from genetic group selection. Moreover, the concept of cultural group selection itself is often misunderstood, because a proper understanding of it requires a correct and fairly detailed understanding of the dual-inheritance theory of cultural evolution in general. Accordingly, my aim in this chapter is to explain the role of group selection in religion by situating it within the more general commitments of dual-inheritance theory.

## 1. Introduction

While debates about group selection in biology may not have led to any actual warfare, they sometimes seem as heated and as polarized as debates about religion. Biologists Edward O. Wilson and David Sloan Wilson note that, “many evolutionists have felt that their very careers would be jeopardized if they invoked group selection. In some cases, their fears were well founded; we could provide numerous examples of colleagues whose articles and grant proposals were rejected when stated in terms of multilevel selection theory, and then accepted when restated using other terms” (2007, 329). When Edward O. Wilson teamed up with modeling experts Martin Nowak and Corina Tarnita (2010) to publish a high-profile critique of attempts to explain eusociality, or ultrasocial behavior, *without* appealing to group selection, their article garnered an extremely critical response letter (Abbot et al. 2011) coauthored by 136 researchers from 103 different institutions. The list of authors and affiliations for this letter was actually longer than the letter itself, giving the impression that it is as much a petition against group selection as a theoretical argument against it.

My goal in the following chapter does not require me to wade into such treacherous waters. This is because my focus is not on the role of group selection across the whole of the living world, but rather only on its role in a single, unusual species—our own—and only in a specific domain of social life in that species—religion. And in this special context, the form of group selection at issue is quite different from the one that is so controversial in biology. The general debate in biology concerns group selection acting on *genetically* inherited traits, but it is assumed here that the role of group selection in religion involves only *culturally* inherited traits (Davis, 2015). The most prominent and fully developed account of religion that appeals to group selection is the one developed and defended by Norenzayan, Henrich and their

colleagues (Atran & Henrich, 2010; Norenzayan, 2013; Norenzayan et al., 2016), which is based on the *dual-inheritance theory* of cultural evolution (see Chapter 7 of this volume; Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 2005; Richerson and Boyd, 2005; Henrich, 2016). The primary tenet of dual-inheritance theory is that humans inherit phenotypic traits in two distinct and independent ways: genetically, through reproduction, and culturally, through social learning. But it follows from the basic logic of Darwinian models that two separate channels of inheritance entail two separate processes of selection. Thus, *cultural group selection* is conceptually and empirically distinct from genetic group selection, and the dual-inheritance theory of religion appeals only to cultural group selection. Indeed, some dual-inheritance theorists have repeatedly argued *against* appeals to genetic group selection in humans (Henrich 2004, 15-16; Richerson and Boyd 2005, 203; Richerson et al. 2016, 9). As a result, reasons for skepticism about group selection in biology are separate from questions about group selection in religion.

Instead, reasons for skepticism about cultural group selection depend on the viability of the dual-inheritance framework itself. Here, there continue to be detractors, who defend what they consider to be competing accounts. These issues have recently been aired and debated in the literature (Richerson et al., 2016; Norenzayan et al., 2016). However, a general defense of the dual-inheritance framework and its application to religion would extend far beyond the scope of this chapter. Generally, many objections to this approach are due to misunderstandings and misinterpretations of its claims. Dual-inheritance theory is based on highly technical mathematical models, so communicating its commitments to wider audiences is a difficult challenge. Accordingly, in the relevant sense, the best defense is a good offense: one good way to defend the dual-inheritance theory of religion is simply to provide a clear, straightforward articulation of its claims and commitments. This is the purpose of the current chapter.

While the focus of this chapter is on the role of group selection in the evolution of religion, my aim is to address this topic in a way that brings out the more basic assumptions of dual-inheritance theory in general. Accordingly, I ask readers to bear with me: Because a proper understanding of the role group selection plays in religion depends on a proper understanding of group selection acting on cultural traits in general (i.e., cultural group selection), religion will enter the picture only near the end of the chapter, after the concept of cultural group selection itself has been explained.

## **2. What Is Cultural Group Selection?**

The standard view of natural selection in population genetics rests on three basic assumptions: (1) that traits vary, (2) that these variants are heritable, and (3) that variants compete, with some increasing in frequency in the population at the expense of others. In dual-inheritance models of cultural evolution, all of these assumptions are preserved (Richerson and Boyd 2005, 76). What changes is the specific mechanism of inheritance. Instead of modeling traits transmitted via genes, the focus is on traits transmitted through social learning, or cultural inheritance. As a result, the cultural fitness of a trait has nothing to do with the genetic fitness of the organism who bears it, because a trait's propensity to be transmitted through learning is

independent from the contribution it makes to its bearer's ability to have offspring. A lot of people can learn a lot of things from a person who never has any children at all. Thus, cultural fitness is *model* fitness: it depends on the contribution a trait makes to its bearer's potential to function as a model for others to learn from (Ramsey & De Block, 2015). And while most culturally inherited beliefs, desires, skills and practices do enhance genetic fitness, others actually reduce genetic fitness (Richerson and Boyd 2005, Ch. 4-5). Accordingly, it remains an open empirical question whether the genetic fitness consequences of any given culturally inherited trait are positive, negative, or neutral.

All of this applies to selection acting on cultural traits in general, whether at the level of individuals or of groups. The difference between group selection and individual selection depends on where the relevant variation occurs. Variation is often described as the "raw material" of selection, because one phenotype can only be selected for, over a competing phenotype, if there is some difference between them that causes one to increase in frequency relative to the other. This means that for selection to occur at the individual level, variation must occur at the individual level, rather than the group level. Put differently, wherever two individuals share the same trait, there can be no selection between them. Likewise, group selection can only occur when variation exists at the group level, such that the traits that vary belong to groups, not to individuals.

While the properties of groups can always be reduced to corresponding properties possessed by individual group members, the emergent properties of groups are nevertheless not identical to the corresponding traits of individuals from which they emerge (Smaldino, 2013). No individual can possess the trait of being a democracy, for example. Rather, democracy is a group-level property that *emerges* from interactions among individuals who play different roles—voters, presidents, senators, judges, etc. While these interactions are reducible to the behavior of individuals, interactions are not themselves properties, or traits, of the individuals who engage in them. Similarly, statistical properties such as frequencies and averages are emergent properties of groups, not properties of individuals. It is possible for the average family in America to have 1.89 children, but it is not possible for any individual family to have 1.89 children. Even if a group is made up of just two individuals, one who is literate and one who is illiterate, the literacy *rate* in that group is not a property of either individual; there is no individual who is 50% literate. Whenever the variation that selection acts upon involves emergent traits of groups, selection acting on this variation is group selection.

*Cultural group selection*, then, requires two things. First, the traits that vary must be traits of *groups*, not individuals. Second, this group-level variation must be transmitted culturally, through social learning. This does not mean that transmission itself occurs at the group level; facts about social learning are facts about how *individuals* acquire their traits (Davis and Margolis, 2014). Thus, in order to explain how group-level variation is inherited, it is necessary to reduce the relevant emergent traits to the corresponding traits of individuals. The spread of democracy does not involve whole countries learning things from other whole countries, nor does literacy spread by whole groups inheriting literacy rates from other whole groups. Rather, democracy spreads from one country to another when individuals learn to perform the roles

from which democracy emerges. High literacy rates spread when more individuals learn how to read. To say that selection favors an emergent, group-level trait is to say that this trait, which does not itself belong to any individual, makes some causal contribution to the subsequent spread of the relevant individual-level traits—the traits from which it emerges.

Suppose, for example, that a passionate, charismatic idealist convinces a number of other members of her own group that democracy is better than monarchy, because it is more just. In such a case, the individual-level traits of democracy spread, but not because of any emergent, group-level trait. Thus, this would be selection at the individual level, not at the group level. By contrast, suppose that a number of individuals become convinced that democracy is better than monarchy for a completely different reason: they observe the economic success of other groups that possess the trait of democracy (Miller, 2016). Whether this economic success is enjoyed by individuals or by the group as a whole, the democracy that caused it would be an emergent, group-level trait, not a trait possessed by any individual. Thus, to the extent that democracy contributes to economic success, which in turn contributes to the spread of the individual-level traits from which democracy emerges in the new group, selection for democracy would occur through group selection. This is so even though the transmission of democracy must be explained at the individual level, by appeal to facts about how individuals acquire the various roles from which democracy emerges.

Indeed, there are a number of different ways in which the emergent trait of democracy might cause the individual beliefs and practices of democracy to be adopted by more people. Individuals in non-democratic nations might witness the economic success of democratic nations, form preferences for democracy themselves, and then advocate for change in their own governments. Alternatively, witnessing the economic success of other democracies might cause individuals in non-democratic nations to migrate into democratic nations, exchanging the political practices of their old home for those of their new home. This would entail that the frequency of the relevant individual-level traits has increased in the overall human population, while the frequency of some competing cultural variant has decreased. A third potential mechanism by which democracy could spread through cultural group selection involves reproduction without genetic inheritance. If democracy were to lead to increased birth rates, and if the growing population were to *learn* the beliefs and practices of democracy after they were born, then again the overall frequency of democratic traits would have increased in the human population. A fourth potential mechanism is somewhat darker: Democratic nations might use their economic success to support military conquest, waging successful wars against non-democratic nations, and then forcing the beliefs and practices of democracy onto individual members of the conquered groups. In all of these hypothetical scenarios, an emergent, group-level trait makes a causal contribution to the spread of the individual-level traits from which it emerges.

Likewise, compare two ways in which literacy might spread from a country with a high literacy rate to one with a low literacy rate. In one scenario, literate individuals from the high-literacy nation voluntarily travel to the low-literacy nation and teach some of its illiterate members to read and write. Here, cultural selection is occurring at the individual level, because the group-

level trait of high literacy *rate* plays no role in causing the relevant individual-level traits to increase in frequency. By contrast, suppose instead that the president of the low-literacy country learned that literacy is positively correlated with per capita GDP (Coulombe & Tremblay, 2006), causing her to initiate a government-funded literacy program in her country's schools. In this case, the group-level property of high literacy rate *does* play a causal role in increasing the frequency of literate individuals in the overall population. The cultural *inheritance* of literacy would still be explained solely in terms of facts about individuals learning how to read and write. But the cultural *fitness* of literacy would be explained, at least in part, by an emergent, group-level trait. This is cultural group selection.

### 3. Cultural Group Selection in Humans

In light of all this, it is an important fact that humans are by far the most cultural species, exhibiting far more group-level variation among cultural traits than any other species. Since variation provides the raw material for selection to work upon, comparing humans to other animals helps to illustrate how much raw material cultural group selection has to work upon.

Social learning occurs in a wide range of vertebrates, with most species possessing only rudimentary capacities, such as an ability to attend to stimuli with which other individuals are observed to interact. On other end of the spectrum, however, other species possess advanced capacities for high-fidelity imitation of complex action sequences (Galef & Laland, 2005). For example, chimpanzees learn from others how to use tools to crack nuts (Boesch & Tomasello, 1998), while whales and dolphins learn from others to execute complex calls, mating songs and foraging techniques (Cantor & Whitehead, 2013). These sophisticated social learning capacities produce distinct cultural traditions in these species. Specific skills and practices are preserved across many generations within distinct social groups, such that all members of one group crack nuts or catch fish in one way, while all members of a neighboring group crack nuts or catch fish in some other way, even when the same materials and techniques are equally available to both groups. In other words, group-level cultural variation does occur in a few other species that possess exceptionally sophisticated social learning capacities.

Nevertheless, the *extent* of group-level cultural variation in these other species still pales in comparison with humans, where striking differences between groups arise across such varied domains as language, clothing, tool use, hunting, agriculture, political organization, art, music, myths and legends, supernatural beliefs, etiquette, moral values, and more. Of course, not all group-level variation in human behavior is cultural. We need not appeal to social learning to explain why people in Bermuda wear shorts while people in Siberia wear fur. Yet in all of the domains just mentioned, a large quotient of group-level variation is indeed due to social learning. And this extreme degree of group-level cultural variation is the result of a correspondingly extreme degree of sophistication in social learning capacities at the individual level.

Chimpanzees are capable of imitating a human model to learn a complex sequence of actions necessary for opening a "puzzle box" to obtain a food treat. But chimpanzees are

discriminating. They only imitate actions that have a clear instrumental function. Human children, by contrast, indiscriminately copy almost everything they see the cultural model do, including not just actions relevant to opening the puzzle box, but a wide range of nonfunctional, irrelevant actions as well (Whiten, McGuigan, Marshall-Pescini, & Hopper, 2009). Indeed, children do this even after they have been explicitly told that many of the model's actions will be pointless, and not need not be copied (Lyons, Young, & Keil, 2007). This phenomenon, called *overimitation*, shows that humans are voracious and promiscuous social learners, even by comparison with the most cultural of other species. Yet overimitation is just one of many ways in which human social learning capacities are unlike those found in any other species. Dual-inheritance theorists have identified a wide range other complex biases and heuristics for social learning as well (reviewed in Chudek, Brosseau-Laird, Birch, & Henrich, 2013; Chudek, Heller, Birch, & Henrich, 2012).

Dual-inheritance theory offers a novel explanation for these unique features of human psychology, and for the unique degree of group-level cultural variation they yield. It's fairly obvious that genetic evolution influences the course of cultural evolution, since what people learn from each other depends in many ways on their genetically inherited cognitive, motivational, and learning capacities. Less obviously, however, dual-inheritance theory predicts that this causal relationship should also run in the opposite direction; selection acting on cultural traits should also influence the course of genetic evolution. The most carefully documented case involves the evolution of genes for digesting lactose in adulthood (Holden & Mace, 2009). In societies that had cultural traditions of pastoralism, the presence of livestock created a readily available source of nutrition: milk. Originally, however, humans couldn't digest milk beyond age four or five. Thus, selection acting on culturally inherited practices of pastoralism created a strong and stable genetic selection pressure for the ability to digest milk, and over several thousand years human genes eventually responded. This bidirectional evolutionary dynamic is called "gene-culture coevolution" (Richerson and Boyd 2005, p. 192), and it also explains many other important features of human life (Henrich, 2016; Laland, Odling-Smee, & Myles, 2010), including our hypertrophied social learning capacities themselves.

In an article entitled "The Evolution of Cultural Evolution," Henrich and McElreath (2003) describe a positive feedback loop in which genetic selection first favored sophisticated capacities for cultural learning, which led to a social environment in which cultural learning was increasingly adaptive, resulting in even stronger genetic selection for cultural learning, and a social environment in which cultural learning was even more adaptive, and so on. In a chimpanzee social environment, where one can learn from cultural ancestors how to crack nuts or extract bone marrow, social learning enhances genetic fitness. But in a human social environment, where the accumulation of cultural knowledge carries vast amounts of adaptive information across a wide range of functional domains, social learning enhances genetic fitness to a far greater extent.

Indeed, overimitation can be explained by the fact that in a human social environment, it is often adaptive to imitate even beliefs and behaviors that have no apparent instrumental value. Henrich (2016, 97-100) offers the example of a small-scale society in the Amazon that has a

long-standing traditional process for removing cyanide from manioc, an important root vegetable. However, individuals don't know about cyanide poisoning, and they don't know that their traditional process removes cyanide. They just know that there is a *right* way to process manioc—the way their cultural ancestors did. Thus, indiscriminate, high-fidelity social learning provides an abundant, healthy food source. In this manner, gene-culture coevolution explains not just why imitation is adaptive for humans, but even why *overimitation* is adaptive. Chimpanzee culture is not like this.

Another unique feature of human psychology explained by gene-culture coevolution is norm psychology. Unlike most other species, human social groups contain cultural values and standards—rights and wrongs—which are enforced by reward and punishment. This is underwritten by a suite of genetic adaptations we can refer to as “the norm system” (Davis, Hennes & Raymond 2018). Once our general social learning capacities became sufficiently sophisticated, one thing people could learn from one another is how to control each other's behavior using punishments and rewards. But such enforcement brings with it strong genetic selection pressures. During the Spanish Inquisition, it was wrong to question the existence of God or the authority of the Pope, and violating these norms could bring imprisonment, torture or execution—all of which are very bad for genetic fitness. Conversely, the norms of nearly all societies deem it right to treat people fairly, and Baumard, Andre and Sperber (2013) show in detail how compliance with norms of fairness is rewarded with fitness-enhancing opportunities for increased cooperation. Of course, many norms are not enforced very strongly, and many norms come and go too quickly to affect genetic evolution. But many other norms remain stable across millennia, and are enforced by rewards and punishments that constitute powerful genetic selection pressures. Norms against theft, rape and adultery, for example, are stable features of most societies today, and are enforced by significant punishments. It is hard to know just how old they are, but they were in place by the time of the Code of Hammurabi approximately 4,000 years ago (Harper 1904), and there is little reason to think they were new at that time. Thus, gene-culture coevolution predicts the genetic evolution of psychological capacities for negotiating an environment filled with normative enforcement.

Among the features of the norm system are cognitive capacities for identifying norms, as well as motivational capacities for following norms and enforcing them on others (Chudek & Henrich, 2011; Sripada & Stich, 2007). These motives can then be divided into *instrumental* desires to follow and enforce norms, in order to avoid punishment and gain rewards, and *intrinsic* desires to follow and enforce norms, or desires to “do the right thing” just because it is right (Sripada & Stich, 2007, p. 285). While people clearly follow many norms just because they want others to like them, or because they don't want to go to jail or be ostracized, they follow other norms because the idea of violating them is already abhorrent in itself, regardless of what anyone else thinks. And such intrinsic motivations to follow norms can lead people to pay very high costs, or take very big risks, for the sake of what they believe is right. “Martyr” is the common term for people who give their very lives for the sake of such intrinsic desires. Moreover, while motivations to follow norms are functionally distinct from motivations to enforce norms, when people *internalize* norms, they acquire both of these forms of intrinsic motivation at once, in a package deal (Sripada & Stich, 2007, p. 289).

Interestingly, the norm system may play an important role in overimitation. Developmental psychologists have found that children are “promiscuous normativists,” as Schmidt et al. (2016) put it, “who overattribute objective social norms even when there actually are none.” This already suggests that overimitation occurs for norms in the same way that it occurs for other cultural traits, but Kenward (2012) argues, further, that promiscuous normativity may be an underlying *cause* of overimitation. Children, he suggests, follow unnecessary steps in action sequences because of an innate bias to assume that whatever a good cultural model does is probably the *right* thing to do, normatively. Children may then be motivated to follow this inferred norm either by instrumental desires to avoid punishment and gain approval, or by intrinsic desires to do the right thing.

One important consequence of overimitation and promiscuous normativity is that they help to concentrate cultural variation at the group level. The more individuals in the same group imitate one another, the less their phenotypes vary, and, thus, the less raw material there is for individual-level selection to work with. And since people imitate fellow group members far more often than foreigners, this reduced variation does not occur at the group level to the same degree. Further, when the traits in question are norms, individuals in the same group are literally forced—by one another—to possess the same phenotypes. Thus, to the extent that different groups have different cultural values, norms reinforce cultural variation at the group level, rather than at the individual level (Richerson et al., 2016, p. 4).

Thus, as impressive as social learning is in chimpanzees, whales and dolphins, social learning in humans is a new thing under the sun. In these other species, group-level cultural variation exists only in small quantities, so group-level selection has little raw material to work with. It is far from clear that cultural group selection occurs at all in these species, and even if it does, its effects are minimal. In humans, by contrast, hypertrophied social learning capacities caused by gene-culture coevolution create and maintain large amounts of group-level cultural variation, providing plenty of raw material for selection to act upon. And many of the beliefs, norms and practices of religion are among the cultural traits that vary, and are selected for, at the group level.

#### **4. Religion, Cooperation and Cultural Group Selection**

Cultural group selection came to be applied to religion through a prior account of the evolution of cooperation in humans. Dual-inheritance theorists have repeatedly argued, over many years, that human cooperation is unique in the animal kingdom in a number of ways, and that this uniqueness is best explained by the unique degree to which we engage in, and rely upon, cultural learning (Boyd & Richerson 2005; 2009; Chudek & Henrich, 2011; Henrich, 2004; Richerson et al., 2016; Richerson & Boyd, 1999). Central to this account is the claim that groups whose emergent cultural traits are more effective at promoting cooperation are more likely to grow in size, and compete successfully for resources, than groups whose cultural traits are less effective at promoting cooperation. As a result, more cooperative cultures are more successful at spreading their more-cooperative cultural traits. In other words, cultures of cooperation are

often selected for at the group level. But since other species lack the social learning capacities necessary to support group-level selection, cooperation in humans has evolved along a trajectory all its own—especially in very recent evolutionary history, since the beginning of settled agriculture (Sterelny, 2016). Religion enters the picture, then, by virtue of its role in promoting cooperation.

First, some background about the evolution of human cooperation: In this literature, cooperation is operationalized in terms of game theory, most often according to the Prisoner's Dilemma, where agents must choose between "defecting," or maximizing their own payoffs by free riding on other players, or "cooperating," or paying some cost that makes them vulnerable to free riding themselves. Accordingly, cooperation in this context is defined more specifically as *prosocial* cooperation, or behavior that is costly to the agent, but beneficial to others (see also Chapter 9, this volume). These costs and benefits are typically measured in the currency of genetic fitness, but it is important to recognize that they are only *short-term*, rather than *overall*, costs and benefits. That is, "prosocial" implies costs, but not *net* costs, so the costs paid in a particular cooperative interaction might (or might not) result in long-term fitness benefits overall. As a result, "prosocial," as it is used in dual-inheritance theory, is different from "altruistic," as it is often used in biology. In biology, altruism is often defined as behavior that exacts *net* fitness costs, which makes altruistic behavior a theoretical puzzle for genetic selection. In dual-inheritance theory, by contrast, there is no reason to assume, a priori, that prosocial behavior should evolve specifically through genetic selection. It is equally possible that it should evolve through cultural selection instead, and, as mentioned above, it is always an open empirical question whether the genetic fitness consequences of any given cultural traits are positive, negative or neutral. Consequently, it is necessary to avoid *defining* "prosocial" in terms of net costs and benefits, because it would then no longer be an empirical question what the genetic fitness consequences of prosocial behavior are. All prosocial behavior would have negative fitness consequences by definition, leaving out important cases where short-term costs are outweighed by long-term benefits.

Note that prosocial behavior does not imply that the benefits accrue indiscriminately to all human beings. In the vast majority of cases, the beneficiaries are members of one's own social group. Accordingly, prosocial behavior is often "parochial" (Bernhard, Fischbacher, & Fehr, 2006), such that solidarity within the group is directed *against* out-group members. Thus, prosocial behavior can be quite *antisocial* from the point of view of those who are not the relevant beneficiaries.

Within a given group, however, the benefits of prosocial cooperation often take the form of non-excludable public goods, or goods that benefit all members equally. For example, when a soldier risks his life defending his group, everyone in the group benefits in the same way from the costs he pays. Likewise, everyone benefits alike from roads, navies, aqueducts, parks, hospitals, resource management systems, social welfare programs, and many other public goods. Of course, individual group members benefit significantly from these goods. But it is essential to keep in mind that because all group members enjoy the *same* benefits, these benefits do not exhibit variation at the individual level, and thus cannot play a role in selection

at the individual level. By contrast, public goods do produce variation at the group level. Groups that produce more public goods possess more resources for growing in size and competing for resources than do groups that produce fewer public goods. And larger groups with better resources are more capable of increasing the frequency of their cultural traits in the overall human population—including the prosocial cultural traits that enable them to generate public goods.

Generating public goods is not easy, of course, because of the central problem of free riding. Precisely because everyone gets the same benefits from cooperation, everyone also shares the same temptation to avoid paying the costs, and to leave it to others to do the work, pay the money, or take the risks. Consequently, one especially important form of group-level variation is variation in cultural norms that suppress free riding. In many different ways, across a vast range of distinct cultures, prosocial behavior is considered right, and free riding is considered wrong. And by enforcing such prosocial norms, cultural groups use punishments and rewards to ensure that individual group members cooperate instead of free ride. In effect, the enforcement of prosocial norms changes the payoff structure of cooperative interactions, such that, whatever the costs of prosocial cooperation may be, the costs of free riding are even greater. Thus, when enforcement is effective, norms prevent cooperative interactions from becoming Prisoner's Dilemmas in the first place, transforming them into other games, in which cooperation is less difficult to achieve.

Moreover, norms can sometimes achieve this effect even when enforcement is *not* effective. This is due to the above-mentioned phenomenon of internalization. When people are intrinsically motivated to follow and enforce norms, they are sometimes willing to pay the costs of prosocial behavior even in the absence of any enforcement. Almost any movie, novel or television drama will illustrate what this looks like, because most such narratives contain at least one hero or heroine who pays significant costs, or takes serious risks, for the sake of some prosocial norm to which he or she is intrinsically committed.

Such norms may or may not be formalized into explicit laws or policies. Laws and policies also specify standards of behavior, or right and wrong ways of acting, along with the punishments and rewards involved in enforcing these standards. In this sense, laws and policies are just a special kind of norms: those which are codified in some formal way. For example, corresponding to ordinary, informal norms against theft, rape and adultery are the formal norms laid down in the Code of Hammurabi—as well as in the laws of modern nations—which prohibit the same actions. Other informal norms, however, have no formal counterparts. In addition, norms, whether formal or informal, may or may not come bundled together in the form of institutions, or organized systems of norms, which (together with the relevant beliefs) specify coordinated roles for individuals to perform in the service of some larger goal or mission.

The government of the United States, for instance, is an institution based originally on formal norms—laws explicitly codified in the Constitution—which specify the right ways of performing the roles of voter, senator, president, supreme court justice, etc. In addition, many informal

norms have evolved around this original structure of formal norms, specifying in further detail how these roles should be performed in a way that integrates the formal structure into the culture of American politics. The result is the emergent, group-level trait of American democracy. And this institution contributes in many ways to the generation of public goods within the group. Through the enforcement of the formal and informal norms in this institution, free riding is suppressed well enough to produce armed forces, an interstate system, the Federal Reserve Bank, funding for research, national parks, and much, much more. Not all of these public goods contribute much to the spread of American cultural traits in the overall human population, and many contribute only in ways that are quite indirect (e.g., by fostering economic success in general, as described above). But to the extent that a democratic central government contributes to an increase in frequency of American cultural traits in the overall population, it can be used to illustrate the role that prosocial norms play in cultural group selection.

Here, finally, is where religion comes in. “Religion” typically refers to internally coherent belief systems that include supernatural agents and forces, and bundle them together with systems norms, including norms that specify roles for individuals to play in coordinated interactions. That is, religions are institutions made up of culturally inherited beliefs, norms and practices. They may be entirely informal, but are usually formalized to a significant degree. And the religions that exhibit the greatest cultural fitness, relative to other religions, tend to be those that heavily promote prosocial behavior.

In making this point, Norenzayan and his collaborators (Norenzayan, 2013; Norenzayan et al., 2016; Shariff & Norenzayan, 2007) have emphasized the role of *supernatural monitoring* in world religions, such as Judaism, Christianity, Islam, Buddhism and Hinduism. These religious systems include many norms asserting the rightness of prosocial behavior, and the wrongness of pursuing self-interest at the expense of other group members. They also include beliefs about supernatural agents who are preternaturally well-suited to the task of enforcing such prosocial norms. What Norenzayan calls “big gods” are supernatural agents who know about violations of prosocial norms even when no mortal beings do. Big gods are preoccupied with human norm-following behaviors, and they have the power to bring misfortune to free riders and norm violators, as well as fortune to those who follow and enforce the right norms. In groups in which the individual-level traits from which these institutions emerge occur in high frequencies, prosocial behavior is thus enforced not just by real people, but, in addition, by supremely powerful divine agents. Not all religions involve big gods, of course. Local ancestor spirits may not count as *big* gods, since they have limited knowledge and power. Nevertheless, they may engage in supernatural monitoring and supernatural enforcement to some degree. Further, group selection may have played little role in the evolution of many religions, especially prior to the advent of settled agriculture (Sterelny 2016). Nevertheless, big-god religions exemplify the special relationship between supernatural beliefs and prosocial cooperation that makes cultural group selection a potent force in the evolution of religion.

That said, it is important to separate the specific function of supernatural monitoring in this account from the more general relationship between supernatural beliefs and prosocial

behavior. Supernatural monitoring by big gods is just one mechanism by which religion promotes prosocial cooperation. In Hindu and Buddhist systems, for example, karma is not a supernatural *agent*, but is rather a supernatural *force*. It is simply built into the metaphysical fabric of the universe that souls strive toward nirvana, and that compassionate, prosocial behavior brings positive karma, moving one closer to this goal, while selfishness and free riding bring negative karma, moving one further from this goal. Karma is not a god at all, much less a big god, so there is no agent who monitors human behavior in this belief system, and doles out rewards and punishments. Nevertheless, supernatural beliefs about the relationship between norm-following behavior and future rebirth function to promote prosocial behavior, by positing a mechanism through which those who violate prosocial norms are punished, while those who follow them are rewarded.

In addition, religions can facilitate prosocial behavior in ways that don't depend on genuine belief. For example, non-believers can still grasp the prosocial themes of religious narratives—the moral of the story—and they may be influenced by such normative messaging in the same way that children are influenced by narratives about fictional heroes. No one believes that Superman exists, but everyone understands that he's a good guy, who takes a lot of risks for the benefit of others. By populating culture with such models, religious systems can contribute to the spread of prosocial behavior even among those who do not adopt the relevant supernatural beliefs. *Any* manner in which supernatural beliefs function to reinforce prosocial norms will be favored by cultural group selection.

Groups in which prosocial religions are common and influential possess an effective mechanism for generating public goods, in addition to any secular institutions they may also possess. This gives them an edge in material competition with other groups, and, consequently, an edge in the competition for the spread of cultural traits. And, of course, religious beliefs and norms are themselves among the cultural traits that spread in this manner. Thus, prosocial religions are emergent, group-level traits, which are transmitted through cultural inheritance, that contribute positively to the spread of the individual-level traits from which they emerge.

A positive feedback loop results. Cultural group selection favors the spread of more prosocial religions at the expense of less prosocial religions, causing religious institutions in general to become more oriented toward prosociality. But then, as religions become more prosocial, religious beliefs and practices are themselves selected for at the group level. Due to selection at the group level, the spread of religion increases prosocial behavior, while increases in prosocial behavior increase the spread of religion.

On this view, this feedback loop depends on a set of genetically inherited social learning capacities that were already in place before religion became a significant factor in human evolution. This is due to the role gene-culture coevolution has played in the evolution of human cooperation. However, dual-inheritance theory predicts that over the last 12,000 years, since the beginning of settled agriculture, the cultural evolution of religion has increasingly been driven by an evolutionary dynamic establishing an intimate connection between religion and prosocial cooperation. And group selection is at the heart of this dynamic.

Note, finally, that for all I have said about the role of prosocial behavior here, religious beliefs may foster group-beneficial behaviors in other ways as well. Norms of self-discipline or industry, such as the “Protestant work ethic,” for example, need not address cooperative interactions at all. Nevertheless, groups in which such norms occur in high frequencies might compete successfully against other groups, in ways that enable them to spread their culture more effectively—including norms about hard work and discipline themselves. So too might norms favoring large families. All else being equal, large groups compete well against small groups. Thus, simply by contributing to the group’s overall size, the emergent trait of having a high frequency of norms favoring high reproductive rates might cause the cultural traits of that group to spread more widely than those of its neighbors. Yet there is nothing especially prosocial about having a large family. Thus, while prosocial norms provide a systematic link between cultural group selection and religion, which is manifest in a global evolutionary trend, this does not mean that religion can *only* spread through cultural group selection in virtue of promoting prosocial behavior. In various ways, other religious beliefs and norms may create other sorts of group-level selection pressures.

## 5. Conclusion

The dual-inheritance theory of religion does not claim that prosocial behavior *depends* upon religion, or that cultural group selection is *necessary* for the existence of religion. It does claim, however, that group selection is an essential part of the cultural evolutionary process through which religious beliefs became intertwined with prosocial norms. It also claims that, via this link with prosocial cooperation, cultural group selection is one important factor (among others) explaining why religion in general is such a universal and fundamental feature of human social life.

## References

- Abbot, P., Abe, J., Alcock, J., Alizon, S., Alpedrinha, J. A. C., Andersson, M., et al. (2011). Inclusive fitness theory and eusociality. *Nature*, 471(7339), E1. <http://doi.org/10.1038/nature09831>
- Atran, S., & 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(1), 18–30.
- Baumard, N., Andre, J.-B., & Sperber, D. (2013). A mutualistic approach to morality: The evolution of fairness by partner choice. *Behavioral and Brain Sciences*, 36(1), 59–78. <http://doi.org/10.1017/S0140525X11002202>
- Bernhard, H., Fischbacher, U., & Fehr, E. (2006). Parochial altruism in humans. *Nature*, 442(7105), 912–915. <http://doi.org/10.1038/nature04981>
- Boesch, C., & Tomasello, M. (1998). Chimpanzee and human cultures. *Current Anthropology*, 39(5), 591–614. <http://doi.org/10.1086/204785>
- Boyd, R., & Richerson, P. J. (2005). Solving the puzzle of human cooperation. In P. Jaisson & S. Levinson (Eds.), *Evolution and culture: A fyssen foundation symposium* (pp. 105-132). Cambridge, MA: MIT Press.

- Boyd, R., & Richerson, P. J. (2009). Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1533), 3281–3288. <http://doi.org/10.1098/rstb.2009.0134>
- Cantor, M., & Whitehead, H. (2013). The interplay between social networks and culture: Theoretically and among whales and dolphins. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1618), 20120340–20120340. <http://doi.org/10.1098/rstb.2012.0340>
- Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.
- Chudek, M., & Henrich, J. (2011). Culture-gene coevolution, norm-psychology and the emergence of human prosociality. *Trends in Cognitive Sciences*, 15(5), 218–226. <http://doi.org/10.1016/j.tics.2011.03.003>
- Chudek, M., Brosseau-Laird, P., Birch, S., & Henrich, J. (2013). Culture-gene coevolutionary theory and children’s selective social learning. In M. Banaji & S. Gelman (Eds.), *Navigating the Social World: What Infants, Children and Other Species Can Teach Us* (pp. 181-185). New York: Oxford University Press.
- Chudek, M., Heller, S., Birch, S., & Henrich, J. (2012). Prestige-biased cultural learning: Bystander’s differential attention to potential models influences children’s learning. *Evolution and Human Behavior*, 33(1), 46–56. <http://doi.org/10.1016/j.evolhumbehav.2011.05.005>
- Coulombe, S., & Tremblay, J.-F. (2006). Literacy and growth. *Topics in Macroeconomics*, 6(2). <http://doi.org/10.2202/1534-5998.1404>
- Davis, T. (2015). Group selection in the evolution of religion: Genetic evolution or cultural evolution? *Journal of Cognition and Culture*, 15(3-4), 235–253. <http://doi.org/10.1163/15685373-12342149>
- Davis, T., & Margolis, E. (2014). The priority of the individual in cultural inheritance. *Behavioral and Brain Sciences* 37(3), 257-258.
- Davis, T., Hennes, E. P., & Raymond, L. (2018). Cultural evolution of normative motivations for sustainable behaviour. *Nature Sustainability*, 1, 218-224. <http://doi.org/10.1038/s41893-018-0061-9>
- Galef, B. G., & Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *Bioscience*, 55(6), 489–499. [http://doi.org/10.1641/0006-3568\(2005\)055\[0489:SLIAES\]2.0.CO;2](http://doi.org/10.1641/0006-3568(2005)055[0489:SLIAES]2.0.CO;2)
- Harper, R. (1904). *Code of Hammurabi, King of Babylon*. Chicago: The University of Chicago Press, Callaghan Co.
- Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior & Organization*, 53(1), 3–35. [http://doi.org/10.1016/S0167-2681\(03\)00094-5](http://doi.org/10.1016/S0167-2681(03)00094-5)
- Henrich, J. (2016). *The secret of our success: How culture is driving our evolution, domesticating our species, and making us smarter*. Princeton: Princeton University Press.
- Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews* 12(3), 123-135.
- Holden, C., & Mace, R. (2009). Phylogenetic analysis of the evolution of lactose digestion in adults. *Human Biology*, 81(5/6), 597–619.

- Laland, K. N., Odling-Smee, J., & Myles, S. (2010). How culture shaped the human genome: Bringing genetics and the human sciences together. *Nature Reviews Genetics*, *11*(2), 137–148. <http://doi.org/10.1038/nrg2734>
- Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences*, *104*(50), 19751–19756.
- Miller, M. K. (2016). Democracy by example? Why democracy spreads when the world's democracies prosper. *Comparative Politics*, *49*(1), 83–116. <http://doi.org/10.5129/001041516819582946>
- Norenzayan, A. (2013). *Big gods: How religion transformed cooperation and conflict*. Princeton: Princeton University Press.
- Norenzayan, A., Shariff, A. F., Gervais, W. M., Willard, A. K., McNamara, R. A., Slingerland, E., & Henrich, J. (2016). The cultural evolution of prosocial religions. *Behavioral and Brain Sciences*, *39*(1), 1–65. <http://doi.org/10.1017/S0140525X14001356>
- Nowak, M. A., Tarnita, C. E., & Wilson, E. O. (2010). The evolution of eusociality. *Nature*, *466*(7310), 1057–1062. <http://doi.org/10.1038/nature09205>
- Ramsey, G., & De Block, A. (2017). Is cultural fitness hopelessly confused? *The British Journal for the Philosophy of Science*, *68*, 305–328. <http://doi.org/10.1093/bjps/axv047>
- Richerson, P. J., & Boyd, R. (1999). Complex societies. *Human Nature*, *10*(3), 253–289.
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. Chicago: University of Chicago Press.
- Richerson, P., Baldini, R., Bell, A. V., Demps, K., Frost, K., Hillis, V., et al. (2016). Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, *39*, 55–68. <http://doi.org/10.1017/S0140525X1400106X>
- Schmidt, M. F. H., Butler, L. P., Heinz, J., & Tomasello, M. (2016). Young children see a single action and infer a social norm. *Psychological Science*, *27*(10), 1360–1370. <http://doi.org/10.1177/09567976166661182>
- Shariff, A. F., & Norenzayan, A. (2007). God is watching you: Priming god concepts increases prosocial behavior in an anonymous economic game. *Psychological Science*, *18*(9), 803–809. <http://doi.org/10.1111/j.1467-9280.2007.01983.x>
- Smaldino, P. (2013). The cultural evolution of emergent group-level traits. *Behavioral and Brain Sciences* *37*(3), 243–254. <https://doi.org/10.1017/S0140525X13001544>
- Sripada, C., & Stich, S. (2007). A framework for the psychology of norms. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind, vol 2: Culture and cognition* (pp. 280–301). New York: Oxford University Press.
- Sterelny, K. (2016). Cooperation, culture, and conflict. *The British Journal for the Philosophy of Science*, *67*(1), 31–58. <http://doi.org/10.1093/bjps/axu024>
- Whiten, A., McGuigan, N., Marshall-Pescini, S., & Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1528), 2417–2428. <http://doi.org/10.1098/rstb.2009.0069>
- Wilson, D. S., & Wilson, E. O. (2007). Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology*, *82*(4), 327–348. <http://doi.org/10.1086/522809>