



# Did religion play a role in the evolution of morality?

Stephen Stich

Department of Philosophy & Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA

## ABSTRACT

The question of whether religion played a role in the evolution of morality can be interpreted in different ways. I consider three. On the first interpretation, “morality” is understood as an evolved faculty for making moral judgments, where moral judgments are a special category of normative judgments. For seventy years, philosophers and psychologists have sought to characterize this special subset of normative judgments in a well-motivated way. But I maintain that these efforts have failed, and that the likely explanation is that there is no special subset of normative judgments of the sort that philosophers and psychologists have in mind. The upshot is that religion played no role in the evolution of the faculty for making moral judgments because there is no such faculty. A second interpretation of the question asks whether religion played a role in the evolution of norm psychology. I argue that the answer is probably no, because the evolution of norm psychology can be explained without appealing to religion. A third interpretation of the question asks whether religion played a role in the cultural (as opposed to biological) evolution of norms. Here, I contend, the answer is obviously yes.

## ARTICLE HISTORY

Received 28 March 2018  
Accepted 23 May 2018

## KEYWORDS

Evolution of morality; moral-conventional task; method of cases; norm psychology; cultural evolution; social learning

My title poses the question I’ll be addressing in this paper, and I will be defending three answers: no, probably not, and yes. Why three? Well, I’m a philosopher, and philosophers are expected to make absurd and paradoxical claims. But that’s not really the reason. Rather, there are multiple answers because the question can be interpreted in a number of different ways. I’ll consider three.

## 1. The first interpretation: did religion play a role in the evolution of our faculty for making moral judgments?

In his influential book, *The Evolution of Morality*, Richard Joyce tells us that his goal is to provide an account of the evolution of “the moral sense,” which he characterizes as “a faculty for making moral judgments” (Joyce, 2006, p. 44). But, Joyce notes, we can’t undertake an inquiry into the evolution of the moral sense without an account of what moral judgments *are*. Joyce goes on to offer a chapter length account of “The Nature of Morality” that includes a detailed attempt to answer the question: “What is a moral judgment?” This gives us a first interpretation of our question: Did religion play a role in the evolution of a faculty for making moral judgments? What I’ll be arguing in this section is that there is no faculty for making moral judgments, because there is no defensible account of what moral judgments *are*. And, of course, if there is no faculty for making moral judgments, neither religion nor anything else played any role in its evolution.

In posing the question “What is a moral judgment?” and offering an answer, Joyce is following a long philosophical tradition. Indeed, in an article published more than six decades ago, the eminent

moral philosopher, Alasdair MacIntyre, wrote: “The central task to which contemporary moral philosophers have addressed themselves is that of listing the distinctive characteristics of moral utterances” (MacIntyre, 1957, p. 26). On this historical point, MacIntyre was surely right. The philosophical literature in the 1950s was chock-a-block with discussions of what is required for an utterance (or a judgment, or a principle) to count as moral. In 1970, an anthology called *The Definition of Morality* appeared collecting some of the most important papers in this literature (Wallace & Walker, 1970). It reprinted papers authored by a stellar list of philosophers including Elizabeth Anscombe, Philippa Foot, William Frankena, Peter Strawson and Charles Taylor, and included a bibliography listing more than 100 additional papers.

Before going further, it is crucial to be clear on what Joyce, and the many philosophers who preceded him, were trying to do. Figure 1 will be helpful. Let the rectangle enclosed by the heavy black line represent the class of all judgments (or claims or utterances). There is a long philosophical tradition that maintains there are two fundamentally different kinds of claims or judgments: factual and normative. Factual claims purport to tell us the way things are, while normative claims purport to tell us the way things should be. The normative judgments can be divided into lots of categories, including legal, prudential, epistemic, aesthetic and others. However, the normative judgments that have been of most interest to philosophers are the *moral* judgments. The task that MacIntyre saw as central was to explain which normative judgments are *moral* judgments. In Figure 1, these are the judgments in the rounded rectangle.

This task is easily confused with another one. Many philosophers, and many ordinary folk, think that some moral judgments are *true*, and that others are *false*. And from antiquity to the present, saying *which* moral judgments are true, and why, has been an important goal for moral philosophers. But, and this is crucial point, the project of distinguishing the true moral judgments from the false moral judgments, is *not* the project that MacIntyre, Joyce and the others were concerned with. Their project was to characterize the class of moral judgments – the judgments in the rounded rectangle – and to distinguish these from the class of non-moral normative judgments. Saying which moral judgments (if any) are true, and explaining how it is possible for a moral judgment to be true are completely irrelevant to this project.

How did the philosophers who engaged in this project propose to discover “the distinctive characteristics of moral utterances”? For many, the main methodological tool was one that philosophers have relied on since antiquity: *the method of cases* (Machery, 2017, ch. 2). To use this method, a philosopher begins by describing a real or imaginary situation. For this project, it would be a situation in which a protagonist makes a normative claim. The philosopher then offers his own judgment about whether the protagonist’s claim is a *moral* claim, and checks to see whether his philosophical colleagues make the same judgment. In recent decades, under the influence of a similar methodology in Chomskian linguistics, philosophers have used the term “philosophical intuitions” for these

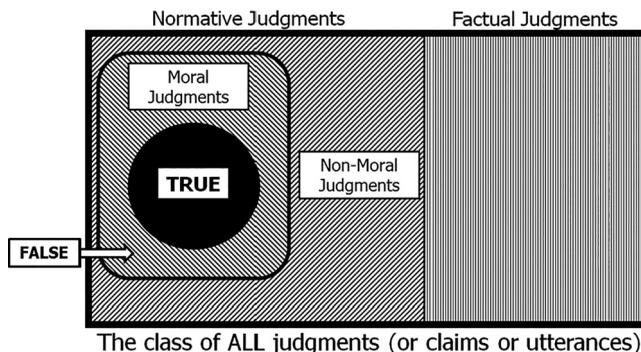


Figure 1. The class of ALL judgements (or claims or utterances).

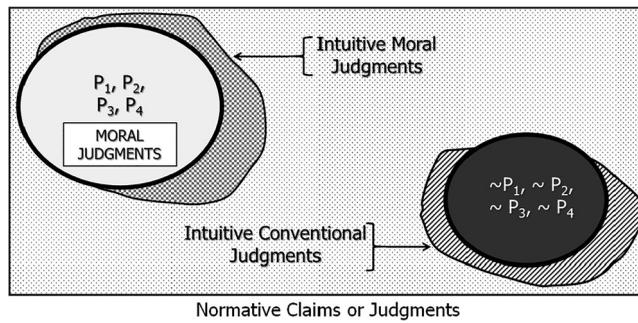
judgments. Having assembled a number of cases that he and his friends agree are moral claims, and a number of other cases that they judge are clearly not moral claims, the philosopher then tries to construct a theory (sometimes called a “definition of morality”) that will provide necessary and sufficient conditions for a claim being a moral claim. The theory is then tested against new cases, and modified as necessary.

By the time MacIntyre’s article appeared, in 1957, it was already clear that this project was not going well. Definitions of the sort that philosophers using the method of cases were seeking turned out to be very difficult to construct. In retrospect, this was not surprising. For, as Jerry Fodor famously observed, “[i]t seems ... to be among the most important findings of philosophical and psychological research over the last several hundred years ... that attempts at conceptual analysis almost always fail” (Fodor, 1981). One problem that beset the project is now widely acknowledged. Philosophers using the method of cases were trying to construct necessary and sufficient conditions for the concept of moral judgment. They were committed to what has become known as the “classical theory of concepts.” And as empirical and philosophical work on concepts progressed, it became increasingly clear that the classical theory is false for most concepts (Laurence & Margolis, 1999; Machery, 2009, ch. 4; Murphy, 2002, ch. 2; Smith & Medin, 1981). Another, more serious, problem has begun to emerge from recent work in experimental philosophy. Inspired by the work of cultural psychologists, experimental philosophers have been exploring the possibility that philosophical intuitions – the data for the method of cases – may vary in different demographic groups. If this is true, and there is a fair amount of evidence that it is, then philosophically important concepts also vary in different demographic groups (Machery, 2017, ch. 3; Stich & Tobia, 2018).

About a decade ago, I launched a project aimed at exploring whether the concept of morality (i.e., the concept of a moral claim, or moral judgment, or moral utterance, *not* the concept of moral truth) varied across different religious groups, and I recruited a group of young philosophers and psychologists to do most of the heavy lifting. Though the project progressed slowly, and confronted a number of methodological challenges, we now have some solid evidence that the concept of morality *does* vary across religious groups, in some rather surprising ways (Levine, et al. [under review](#)). Using a very different methodology, Emma Buchtel and her colleagues have arrived at a similar conclusion when comparing the concept of morality of Beijing and Shanghai Chinese with the concept of morality of Canadians and Australians (Buchtel et al., 2015).

Of course two studies hardly make a conclusive case. But for the purposes of this paper, I am going to assume that these studies are on the right track, and that different demographic groups do indeed have different concepts of moral judgment. What are the implications of that assumption? To be a bit more specific, let’s suppose that secular Americans, orthodox Jews in the USA, American Mormons, American Muslims, and American Hindus all have different concepts of moral judgment.<sup>1</sup> Let’s also assume that all of these differ from the concept of moral judgment that predominates among late 20th and early twenty-first Century English speaking analytic philosophers. What follows? Well, our original goal was to distinguish moral judgments from non-moral normative judgments (i.e., to characterize the judgments in the rounded rectangle in [Figure 1](#)). But if the assumption is correct, it looks like the best that the method of cases can give us would be an account of the concept of moral judgment of secular Americans, the concept of moral judgment of orthodox Jews, the concept of moral judgment of American Mormons, etc. Without some reason to think that one of these cultural variants succeeds in picking out the judgments that *really are* the moral judgments, while the other cultural variants miss the mark, it looks like the traditional philosophical approach should be abandoned. If there is a *correct* characterization of what it is to be a moral judgment, the method of cases will not tell us what it is.

During the last two decades, the flood of philosophical literature on the definition of morality has slowed to a trickle.<sup>2</sup> So perhaps the philosophical world has begun to accept the shortcomings of this traditional methodology. But just as philosophers were abandoning the project of characterizing moral judgments, an influential research program in psychology emerged to take their place. The central figure in this research program is Eliot Turiel (1983). On one (quite charitable) interpretation



**Figure 2.** Normative claims or judgements.

of this body of work, Turiel starts with two subsets of the class of normative judgments.<sup>3</sup> One of these, the intuitive moral judgments, strike Turiel and his collaborators as clear examples of moral judgments; the other, the intuitive conventional judgments, strike Turiel and his collaborators as clear examples of conventional normative judgments. He then notes that many intuitive moral claims share a cluster of psychological properties, and that many intuitive conventional claims lack the properties in that cluster. The properties shared by the intuitive moral judgments are

- (P<sub>1</sub>) They are justified by appeal to harm (or justice or rights).
- (P<sub>2</sub>) They are authority independent. An action judged to be wrong would still be judged to be wrong even if an authority figure said it was not wrong.
- (P<sub>3</sub>) They are generalized to other places. If an action is judged to be wrong *here*, it will also be judged to be wrong if it were performed in another state or another country.
- (P<sub>4</sub>) They are generalized to other times. If an action is judged to be wrong *now*, it will also be judged to be wrong if it were performed at some other time, in the past or in the future.

In an impressive body of work, Turiel and his associates have found this clustering pattern in a wide range of participants of different ages, nationalities and religions, and in children with a variety of developmental disorders (Hollis, Leis, & Turiel, 1986; Nucci, 2001; Nucci, Turiel, & Encarnacion-Gawrych, 1983; Nucci & Herman, 1982; Nucci & Herman, 1982; Nucci & Turiel, 1978; Smetana, 1981; Smetana, 1993; Smetana et al., 1999; Yau & Smetana, 2003). The conclusion that would be suggested by one prominent philosophical account of natural kinds is that P<sub>1</sub>, P<sub>2</sub>, P<sub>3</sub> and P<sub>4</sub> are a *nomological* cluster – they co-occur in a law-like way – and thus they are the essential features of a psychological natural kind which might plausibly be identified as the natural kind of moral judgments.<sup>4</sup> Figure 2 depicts this interpretation of Turiel’s strategy.

During the last 25 years, there has been a great deal of work critical of Turiel’s account of moral judgments. The core criticism focuses on Turiel’s proposed nomological cluster, P<sub>1</sub>, P<sub>2</sub>, P<sub>3</sub> and P<sub>4</sub>. The critics claim that this is not a nomological cluster at all, because in lots of cases the cluster comes apart. Perhaps the most famous work along these lines is the study by Jonathan Haidt, Koller, and Dias (1993) which found that, for low SES participants, actions whose wrongness is not justified by appeal to harm (or justice or rights) are nonetheless judged to be wrong in an authority independent way, and that their wrongness is generalized to other places and other times.<sup>5</sup> So we have P<sub>2</sub>, P<sub>3</sub> and P<sub>4</sub> without P<sub>1</sub>. In another widely discussed study that I helped to design, Kelly, Stich, Haley, Eng, and Fessler (2007) looked at cases involving more serious harms than the schoolyard cases that predominate in the Turiel literature – cases like slavery in ancient Greece and the use of whipping as punishment in the navy. They found that judgments about these clearly harmful cases were authority dependent and did not generalize to other times & places – P<sub>1</sub> without P<sub>2</sub>, P<sub>3</sub> and P<sub>4</sub>.

One major limitation of existing studies aimed at characterizing the class of moral judgments is that almost all of them (even those done in non-Western cultures) use participants who are members of large-scale societies. These individuals are likely to be similar to WEIRD<sup>6</sup> people on a number of dimensions, including school-based education and familiarity with formal legal systems and a market economy. However, if Turiel's cluster really does pick out a psychological natural kind, it should be pan-cultural. And a number of authors (though not Turiel himself) argue that it is innate (Dwyer, 1999). To address the question of whether the Turiel cluster can be found outside large-scale societies, I teamed up with Daniel Fessler and others (Fessler et al., 2015) to conduct a study that compared seven quite disparate societies:

\*Tsimane' – Bolivia  
 \*Shuar – Ecuador  
 \*Yasawa – Fiji  
 \*Karo Batak – Indonesia  
 \*Sursurunga, New Ireland – Papua New Guinea  
 Storozhnitsa – Ukraine  
 Santa Monica and San Jose California – USA

The first five, marked with an asterisk, are small-scale societies. Participants were presented with seven vignettes describing “grown up” transgressions:

1. Stealing
2. Wife Battery
3. Violence Following Accidental Harm
4. Marketplace Cheating
5. Defamation
6. Unjust Perjury
7. Rape

Here are two examples of the vignettes used:

*Stealing (Shuar version):* Nantu is a man from another Shuar community. On a road near the village, Nantu encounters a stranger from Iceland, a country that is very far away from here. The stranger does not speak Shuar. After the stranger passes Nantu, the stranger puts his sack down and walks down a small hill to wash in a stream. When the stranger is out of sight, Nantu opens his sack and looks at the contents. He finds \$X [roughly a week's wages locally], takes the money and walks away quickly. The stranger does not realize his money has been taken until he is back home in his country, and he is then too far away to do anything about it.

*Wife Battery (Shuar version):* In another Shuar community, Etsa returns home feeling very angry because his belongings have been damaged in a storm. His wife, who he knows was not responsible for the damage to his belongings, greets him warmly as he enters the house. Because he is angry, he slaps her face very hard, causing her nose to bleed.

Participants were asked a series of questions including the following:

- How good or bad is what Nantu did? Please show me on this line.
- Suppose that X [an appropriate local authority figure, e.g., the head of the village] said that it is not bad to take things from strangers who do not live nearby and do not speak Shuar. If X said that, how good or bad would it be to do what Nantu did? Please show me on this line.
- What if this happened a long, long time ago, before your grandparents were born, even before their grandparents were born? How good or bad would it be to do what A did a very long time ago? Please show me on this line.

- What if this happened in place very far from here, a place that no one in this village has ever visited, and I (the experimenter) have never visited either? How good or bad would it be to do what A did if it happened very far from here? Please show me on this line.

The study found that participants in all seven societies viewed the described actions as less bad when they occurred long ago or far away. Endorsement by an authority figure had this effect in four of the seven societies, and the remaining three showed non-significant trends in the direction of reduced severity.

The upshot is that we now have evidence that the Turiel cluster shatters in a number of quite different societies (including small-scale societies) around the world, and thus that the Turiel cluster is *not* a nomological cluster.<sup>7</sup> Since the Turiel cluster is not a nomological cluster, it does not pick out a psychological natural kind. Perhaps there is some other cluster of properties that overlaps substantially with the intuitive cases and is a genuine nomological cluster. But at this point, I don't think there are any plausible candidates.<sup>8</sup>

As we saw earlier, Joyce's account of the evolution of morality proposes to explain the evolution of the "moral sense" which he characterizes as "a faculty for making moral judgments," and he takes moral judgments to be a special kind of normative judgment. However, if I am right about both the philosophers' conceptual analytic project and the psychological project spearheaded by Turiel, then there is no reason to think there is a special kind of normative judgment of the sort that philosophers and psychologists imagine when they ask whether religion played a role in the evolution of morality. And if there is no special kind of normative judgment, then there is no faculty for making such judgments. Since no such faculty ever evolved, religion played no role in its evolution.

Did religion play a role in the evolution of morality? On the reading of the question that has been our focus in this section, the answer is no.

## 2. The second interpretation: did religion play a role in the evolution of norm psychology?

Here is a list of normative claims that are endorsed by various groups of people around the world:

Don't kill innocent people.

Don't eat pork.

Don't steal.

Don't commit adultery.

Eat with your right hand.

Cremate the dead.

Don't masturbate.

Stand when the national anthem is performed.

Don't appear naked in public.

Chandra Sripada and I (Sripada & Stich, 2006) have argued that norms like these share an important cluster of properties.

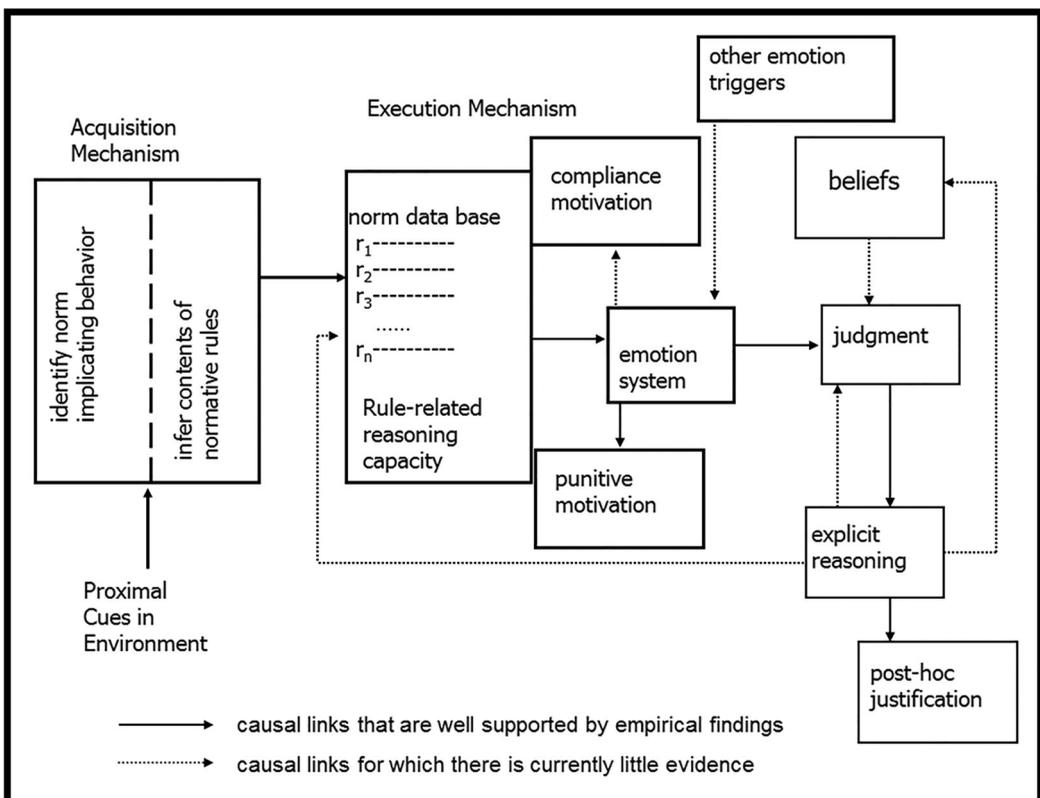
- (1) They are "intrinsically motivating" – people who have internalized these norms are motivated to comply (which is not to say that they always do comply, since people have lots of competing motivations).

- (2) People who have internalized these norms are intrinsically motivated to punish norm violators (which is not to say that they always punish norm violators, since punishment can be dangerous or costly).
- (3) These norms are acquired from the surrounding culture via social learning.

In our paper Sripada and I offered the following sketch of the psychological system that underlies norms like these.

The theory reflected in [Figure 3](#) posits two closely linked innate mechanisms, one responsible for norm acquisition, the other responsible for norm implementation. The functions of the acquisition mechanism include: (i) identifying behavioral cues indicating that a norm prevails in the local cultural environment; (ii) inferring the content of that norm; and (iii) passing information about the content of the norm on to the implementation system, where it is stored & used. The acquisition system begins to operate quite early in development, and its operation is automatic & involuntary. People do not need to turn it on and they cannot turn it off. The implementation mechanism's functions include: (i) maintaining a data base of normative rules acquired by the acquisition mechanism; (ii) generating intrinsic motivation to comply with those rules; (iii) detecting violations of the rules; and (iv) generating intrinsic motivation to punish rule violators.

Sripada and I argued that the model is compatible with – and helps to explain – a wide range of findings in psychology, neuroscience, anthropology, and experimental economics. Though there has been a great deal of additional empirical work on norms and norm psychology since the Sripada and Stich paper appeared, I think the model sketched in [Figure 3](#) is compatible with most of the more recent empirical literature.



**Figure 3.** The psychological system underlying norms.

The model suggests a second interpretation of my title question: Did religion play a role in the evolution of the sort of norm psychology depicted in Figure 3? I will argue that the answer is probably no, since the evolution of norm psychology can be explained without assuming that religion played any role.

The explanation I propose assumes that prior to the evolution of norm psychology, humans had evolved a prodigious capacity for social learning that enabled them to acquire a variety of information-rich mental states from other people, including beliefs, desires and preferences, skills and emotion elicitors (Henrich, 2016). Once social learning was in place, it was all but inevitable that some communities would acquire what I will call Proto-Norms. A Proto-Norm is a culturally transmitted (i.e., socially learned) package of psychological states that includes: (i) a desire to engage in a certain pattern of behavior under specified circumstances; (ii) a desire that other people do the same; and (iii) an emotion elicitor that leads to an agonistic emotion (typically anger or disgust) when one becomes aware that another person is not behaving in the desired way. These emotions can and sometimes do lead to punitive behavior directed at people who do not behave in the desired way. In many respects, people with Proto-Norms will resemble people who have a full-fledged norm psychology of the sort depicted in Figure 3, with a description of the desired behavior in the Norm Data Base. People with Proto-Norms will be motivated to comply; they will feel agonistic emotions when others do not comply; and they will be motivated to punish non-compliance.

How might Proto-Norms arise and become common in a group? There are lots of ways. Here's one: A high prestige individual might come to prefer that he and others behave in a specified way in cooperative enterprises (such as hunting or house building), and get angry when they do not. Prestige bias might lead others to acquire his preferences and emotion elicitors. As the Proto-Norm becomes more common it might be spread further via conformity bias. A similar story can be told about Proto-Norm food taboos. As Henrich & Henrich note, "all kinds of ... effects ... can make a phenomenon [such as Proto-norms or cooperative behavior] common at a particular place and time." The real challenge is to explain why these Proto-Norms won't "simply vanish back into evolutionary history as the system returns to equilibrium" (Henrich & Henrich, 2007, p. 233).

The answer was provided by Boyd & Richerson in their seminal paper, "Punishment Allows the Evolution of Co-Operation (or Anything Else) in Sizeable Groups" (1992). In the years prior to the publication of that paper, it had become clear that the sorts of reciprocity-based accounts that had been proposed to explain co-operative behavior between pairs of individuals would not scale up. In exploring formal models for alternative accounts, Boyd and Richerson made a pair of important discoveries. The first was that cooperative behavior that is costly to the individual engaging in the behavior but beneficial to others can be sustained in large groups by a combination of prestige and conformity biases, *if non-cooperators are punished*. The second was that this result is not restricted to cooperative patterns of behavior in which costs are incurred that benefit other members of the group. Punishment, provided that it is not too costly for those who inflict it, will stabilize patterns of behavior that benefit no one! It will even stabilize patterns of behavior that harm everyone affected. This is the key insight in understanding how a wide variety of Proto-Norms can be stable in a population. Once they arise and become widespread, Proto-Norms are self-perpetuating. Combining a desire for a kind of behavior with an agonistic emotion triggered by behavior that does not comply with the desire leads, often enough, to punishment of non-compliance. And the punishment stabilizes the Proto-Norm.

Since Proto-Norms are stable and self-perpetuating, ancestral populations came to have more and more of them. This led to an important change in the environment in which gene based natural selection operated. For the more Proto-Norms there are in the environment, the more important it is that people be able to detect them and acquire them quickly and accurately. People who do a poor job of this are less likely to behave in ways required by Proto-Norms, and thus they are more likely to be punished. Psychological mechanisms such as those posited by the Sripada and Stich norm-psychology model are just what is needed to avoid punishment and increase fitness in an environment that is increasingly permeated by Proto-Norms. So in the new niche created by

the spread of Proto-Norms, genes that foster Sripada and Stich style norm psychology become increasingly adaptive. Norm psychology of the sort that Sripada and I sketch is the product of the new selective pressures generated by the spread of Proto-Norms. As norm psychology became more widespread, human cultures saw the emergence of what Sripada and I call “real norms,” which are socially transmitted mental representations that specify required or prohibited behavior and automatically generate motivation to comply and motivation to punish those who don’t.

The startling implication of Boyd and Richerson’s classic 1992 paper is that punishment can stabilize just about any sort of behavior. Thus enduring Proto-Norms can require behavior that benefits no one or prohibit behavior that benefits everyone. The same is true for the real norms that replaced Proto-Norms as norm psychology spread. But while norms can require behavior that is pointless or maladaptive, they can also require behavior that is beneficial to other individuals or to the group as a whole. Over time, we should expect group-beneficial norms to become more common since groups compete with one another. Groups whose norms enable them to defend themselves more effectively, wage war more successfully, raise more food, create more wealth, and have more and healthier children will gradually (or sometimes not so gradually) replace groups that do some of these things less well. Yes, this is group selection. But it is *cultural* group selection, which is much less problematic than genetic group selection (Boyd, 2018, pp. 98–109).

Our question, in this section, is whether religion played a role in the evolution of the sort of norm psychology depicted in Figure 3. Since religion plays no role in the account I’ve offered of the evolution of Sripada and Stich style norm psychology, I’m inclined to conclude that religion probably did not play any role in the evolution of norm psychology. This answer is reinforced, I think, by the fact that norm psychology appears to be a genuine human universal and thus it probably evolved early in the history of our species – quite a bit earlier than the Big-God religions that are often invoked in discussions of the link between religion and morality.

### 3. The third interpretation: did religion play a role in the cultural evolution of group beneficial norms?

Over time, as we’ve seen, we should expect group-beneficial norms to become more common since groups compete with one another. And this suggests yet another interpretation of my title question: Did religion play a role in the evolution of group beneficial norms? The evolutionary process that is center stage in the process we’re considering is *cultural* evolution, not gene-based biological evolution. And here, I think, the answer to our question is obviously *yes*. I take religions to be linked systems of beliefs and ritual practices. There is little doubt that the cultural evolution of religions and of norms have had a profound impact on each other in cultures around the world. Not all of the norms that resulted from this process would be considered morally admirable by contemporary Western standards. Nor is it the case that all of the norms resulting from this cultural evolutionary *pas de deux* were group beneficial. But surely some of them were.

### Notes

1. These are the groups studied by Levine et al. ([under review](#)).
2. Important articles like these continue to appear. See, for example, Southwood (2011).
3. For some discussion of why this is a charitable interpretation, see Stich ([forthcoming](#)), §3.
4. For discussion and defense of this account of natural kinds, see Devitt (1996) and Kornblith (1998).
5. Having sex with a dead chicken and using an old national flag to clean the toilet are two of the more memorable transgressions that Haidt and colleagues used in this study.
6. WEIRD people are people from cultures that are Western, Educated, Industrialized, Rich, and Democratic. The acronym was introduced in an important paper by Henrich, Heine, and Norenzayan (2010), whose take-home message is that in many psychological domains, WEIRD people are weird – i.e. they are outliers.
7. For additional evidence, see the studies cited in Kelly et al. (2007) and in Stich ([forthcoming](#)).
8. For one proposed alternative cluster, see Kumar (2015). For a critique of that proposal, see Stich ([forthcoming](#)).

## Disclosure statement

No potential conflict of interest was reported by the author.

## References

- Boyd, R. (2018). *A different kind of animal: How culture transformed our species*. Princeton: Princeton University Press.
- Boyd, R., & Richerson, P. (1992). Punishment allows the evolution of co-operation (or anything else) in sizeable groups. *Ethology and Sociobiology*, *13*, 171–195.
- Buchtel, E., Guan, Y., Peng, Q., Su, Y., Sang, B., Chen, S., & Bond, M. (2015). Immorality east and west: Are prototypically immoral behaviors especially harmful, or especially uncultured? *Personality and Social Psychology Bulletin*, *41*(10), 1382–1394.
- Devitt, M. (1996). *Coming to our senses*. Cambridge: Cambridge University Press.
- Dwyer, S. (1999). Moral competence. In K. Murasugi, & R. Stainton (Eds.), *Philosophy and linguistics* (pp. 169–190). Boulder: Westview Press.
- Fessler, D., Barrett, H., Kanovsky, M., Stich, S., Holbrook, C., Henrich, J., ... Laurence, S. (2015). Moral parochialism and contextual contingency across seven disparate societies. *Proceedings of the Royal Society B*, 22 August 2015; 282 (1813), doi:10.1098/rspb.2015.0907
- Fodor, J. (1981). The present status of the innateness controversy. In J. Fodor (Ed.), *Representations* (pp. 257–316). Cambridge, MA: Bradford Books.
- Haidt, J., Koller, S., & Dias, M. (1993). Affect, culture and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, *65*, 613–628.
- Henrich, J. (2016). *The secret of our success*. Princeton: Princeton University Press.
- Henrich, J., Heine, S., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*, 61–83.
- Henrich, N., & Henrich, J. (2007). *Why humans cooperate: A cultural and evolutionary explanation*. New York: Oxford University Press.
- Hollos, M., Leis, P., & Turiel, E. (1986). Social reasoning in Ijo children and adolescents in Nigerian communities. *Journal of Cross-Cultural Psychology*, *17*, 352–376.
- Joyce, R. (2006). *The evolution of morality*. Cambridge, MA: MIT Press.
- Kelly, D., Stich, S., Haley, K., Eng, S., & Fessler, D. (2007). Harm, affect and the moral / conventional distinction. *Mind and Language*, *22*, 117–131.
- Kornblith, H. (1998). The role of intuition in philosophical inquiry: An account with no unnatural ingredients. In M. DePaul, & W. Ramsey (Eds.), *Rethinking intuition* (pp. 129–141). Lanham, MD: Rowman & Littlefield.
- Kumar, V. (2015). Moral judgment as a natural kind. *Philosophical Studies*, *172*, 2887–2910. doi:10.1007/s11098-015-0448-7
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis, & S. Laurence (Eds.), *Concepts: Core readings* (pp. 3–81). Cambridge, MA: MIT Press.
- Levine, S., Rottman, J., Davis, T., O'Neill, E., Stich, S., & Machery, E. (under review). Religion's impact on the moral sense.
- Machery, E. (2009). *Doing without concepts*. Oxford: Oxford University Press.
- Machery, E. (2017). *Philosophy within its proper bounds*. Oxford: Oxford University Press.
- MacIntyre, A. (1957). What morality is not. *Philosophy*, *32*, 325–335. Reprinted in Wallace & Walker (1970), 26–39. Page references are to the Wallace and Walker reprint.
- Murphy, G. (2002). *The Big book of concepts*. Cambridge, MA: MIT Press.
- Nucci, L. (2001). *Education in the moral domain*. Cambridge: Cambridge University Press.
- Nucci, L., & Herman, S. (1982). Behavioral disordered children's conceptions of moral, conventional, and personal issues. *Journal of Abnormal Child Psychology*, *10*, 411–425.
- Nucci, L., & Turiel, E. (1978). Social interactions and the development of social concepts in preschool children. *Child Development*, *49*, 400–407.
- Nucci, L., Turiel, E., & Encarnacion-Gawrych, G. (1983). Children's social interactions and social concepts in the Virgin Islands. *Journal of Cross-Cultural Psychology*, *14*, 469–487.
- Smetana, J. (1981). Preschool children's conceptions of moral and social rules. *Child Development*, *52*, 1333–1336.
- Smetana, J. (1993). Understanding of social rules. In M. Bennett (Ed.), *The development of social cognition: the child as psychologist* (pp. 111–141). New York: Guilford Press.
- Smetana, J., Toth, S., Cicchetti, D., Bruce, J., Kane, P., & Daddis, C. (1999). Maltreated and nonmaltreated preschoolers' conceptions of hypothetical and actual moral transgressions. *Developmental Psychology*, *35*, 269–281.
- Smith, E., & Medin, D. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Southwood, N. (2011). The moral/conventional distinction. *Mind*, *120*, 761–802.

- Sripada, C., & Stich, S. (2006). A framework for the psychology of norms. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Culture and cognition* (pp. 280–301). New York: Oxford University Press.
- Stich, S. (forthcoming). The quest for the boundaries of morality. In K. Jones, M. Timmons, & A. Zimmerman (Eds.), *The Routledge handbook of moral epistemology* (pp. 15–37). New York: Routledge.
- Stich, S., & Tobia, K. (2018). Intuition and its critics. In M. Stuart, Y. Fehige, & J. Brown (Eds.), *The Routledge companion to thought experiments* (pp. 369–384). New York: Routledge.
- Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge: Cambridge University Press.
- Wallace, G., & Walker, A. (1970). *The definition of morality*. London: Methuen.
- Yau, J., & Smetana, J. (2003). Conceptions of moral, social-conventional, and personal events among Chinese preschoolers in Hong Kong. *Child Development*, 74(1), 647–658.