VIII.10
Cultural Evolution

Elizabeth Hannon and Tim Lewens

Outline

1. What cultural evolution is not
2. Memetics
3. Cultural evolution
4. Nonhuman animal cultural evolution
5. Defining culture

For most of the twentieth century, evolutionary theory focused on phenotypic variation underpinned by inherited genetic variation. Any comprehensive account of the evolution of the human species, and some animal species, must acknowledge that this is at best a simplification of the forces affecting change and stasis in these lineages. Habits, know-how, and technology—what we might consider cultural traits—can also contribute to survival and reproduction. Moreover, these traits are often maintained, in our own species at the very least, by learning from others—that is, they are inherited nongenetically. Further, these traits often show patterns of cumulative improvement as discoveries made in one generation are built on and modified. There also exist subgroups with distinct traits, again often generated and maintained through learning. Theories of cultural evolution start from the observation that humans, and possibly other species, display these important prerequisites for evolution—variation and inheritance of cultural traits—and attempt to build rigorous accounts of cultural change based on this. How exactly these accounts should be fashioned, what relationship cultural evolution should have with organic evolution, and how culture itself should be

conceivedremain open questions. This chapter describes some recent attempts to address these questions.

Glossary

cultural evolution. A process of change in the traits manifested within a population that is explained by various forms of social learning among species members.

horizontal transmission. Transmission within a generation, sometimes also used to refer to transmission from any nonparent.

meme. A cultural entity, intended to be analogous to a gene, capable of being replicated and transmitted between individuals.

replicator. An entity capable of being replicated and capable of influencing its own chances of being replicated through its effects on the world.

vertical transmission. Transmission from parent to offspring, usually of genetic material.

1. What Cultural Evolution Is Not

Not all evolutionary approaches that seek to account for cultural phenomena are theories of cultural evolution. Evolutionary psychologists, for example, tend to regard human behavior and culture as the output of cognitive adaptations, and they assume that the most important mechanism producing such cognitive adaptations is natural selection acting on genetically inherited variation (see chapter VII.12). Evolutionary psychologists acknowledge that changes in cultural environments can affect the behavioral outputs of cognitive adaptations, but they tend to downplay the role of nongenetic inheritance. As a consequence, they also tend to be skeptical of the thought that cultural change may affect and underpin the generation of cognitive adaptations themselves.

Nonetheless, cultural inheritance and genetic inheritance are processes that may affect

each other in important ways. Cultural changes bring about alterations to the environment, which in turn affect how genes act in development and what selection pressures act on genes. For example, dairy farming, thought to have developed somewhere between 6000 and 8000 years ago, appears to have created a selective environment that facilitated the proliferation of lactose tolerance in those populations where it was practiced. Far from removing humans from the evolutionary fray, our cultural environments may exert selective pressures and thus may be implicated in the ongoing evolution of our physiological natures. These sorts of processes are captured by models of gene-culture coevolution that explore the ways in which cultural changes, adaptive or otherwise, can affect genetic evolution, and vice versa.

The targets of explanation for evolutionary psychology and gene-culture coevolutionary accounts differ from those of theories of cultural evolution. Theories of cultural evolution address cultural trends in their own right, rather than as outputs of cognitive adaptations or as selective environments for the natural selection of genes. Further, unlike evolutionary psychology, though in common with gene-culture coevolutionary accounts, they allow a significant amount of nongenetic, learning-based inheritance. The remainder of this chapter describes two distinct attempts to build a theory of cultural evolution from this starting point.

2. Memetics

A theory of cultural evolution needs some systematic way of modeling the effects of cultural inheritance. One such approach is memetics, which has attracted considerable attention in the popular scientific literature. Initially developed by Richard Dawkins, and following the "gene's-eye view" of natural selection he popularized, memetics takes the view that to explain the sort of transgenerational resemblance needed for cumulative evolutionary change, entities that have the ability to make faithful copies of themselves—so-called replicators—are

required. In standard biological models of evolution it is assumed that genes are the relevant replicators; genes make copies of themselves, and (so the story goes) this explains why offspring resemble their parents. For culture to evolve, memeticists argue that replicators—called *memes*—are once again required. Dawkins lists some exemplary memes: "tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches." Note that while it is sometimes assumed that all memes are ideas (and vice versa), Dawkins's list includes other types of things, such as ways of making pots, which are techniques.

Memetics proposes that ideas, skills, practices, and so on, are entities that can be understood to hop from mind to mind, making copies of themselves as they go. For example, you hear a song on the radio as you leave your house in the morning and you sing it at work that day. Your colleague later whistles it as she prepares dinner that evening. Her child hums the song in school the next day and passes it on to his classmates. The meme—in this case, the tune—spreads through its being "catchy." Like genes, memes have differential success in replication: some songs are catchier than others. The rate at which a meme may replicate itself is thought to be dependent on the same factors that determine the rate at which a gene may replicate itself—namely, its effects on the organism it inhabits and on the local environment (partly constituted by the downstream effects of other memes) in which the organism finds itself.

Critics of memetics put pressure on the claim that cultural inheritance is analogous to genetic inheritance. Such criticisms are as likely to originate from those sympathetic to the broad project of developing evolutionary approaches to culture as they are from those outside this project and doubtful of its merits. The remainder of this section outlines some of the key criticisms leveled against memetics.

First, cultural items rarely behave like replicators, and imitation is often very error prone. If you see us dance the tango, and this inspires in you the desire to also dance the

tango, we will almost certainly not dance exactly the same steps. Our dances will have been influenced by a wide variety of factors such as who taught us the dance and our own particular physiologies. If this is copying, it is very bad copying indeed.

A second, and closely related, criticism of memetics draws on the fact that while genetic replication allows us to trace a token copy of a gene back to a single parent, ideas are rarely copied from a single source in a way that allows us to trace clear lineages. Perhaps you learned the tango from several teachers, and your style has been influenced by watching expert dancers. There is no clear single origin for your "tango" meme. Within the realm of biological evolution, an understanding of Mendel's laws has been important in explaining some aspects of evolutionary dynamics. Mendel's laws rely on an understanding of genes as discrete, transmitted units. But if token ideas can appear in an individual by virtue of that individual's exposure to several sources, then it is unlikely that anything close to Mendel's laws will be discovered within cultural evolution. Such an objection need not be fatal for theories of cultural evolution in general, as we shall see, but it does threaten the tight analogy memetics draws between ideas and genes.

Third, memetics seems to demand that we be able to divide culture into discrete units. But it is not clear how this should be done. Ideas stand in logical relations to one another. It is impossible to believe in the theory of relativity without understanding it, and one cannot understand it without holding many additional beliefs relating to physics. It is not at all clear that it makes sense to think that the theory of relativity can be isolated from the rest of physics as an individual meme. One might respond that we have delineated the meme incorrectly here, and those ideas to which the theory of relativity stands in logical relation all form a part of a single more inclusive meme. The worry, now, is that even if we "step back" and consider some broader group of theories, we cannot understand even them without further basic mathematical training, understanding of the operation of measuring apparatus, and so forth. A

form of holism looms, according to which single memes will correspond to massive complexes of belief.

These criticisms focus on whether memes exist. Another matter of concern for memetics in particular, and theories of cultural evolution in general, stems from the fact that cultural transmission occurs both vertically and horizontally. Inheritance as understood in mainstream evolutionary biology involves *vertical transmission*, whereby genetic inheritance is taken to occur between parents and their offspring. But in the case of cultural inheritance, we acquire traits or memes from a wide variety of sources. This is *horizontal transmission*, and it can occur more rapidly than vertical transmission: a significant proportion of a population might come to possess a meme within a single generation. The same degree of saturation (assuming fairly low fitness differences) can take many generations via vertical transmission. The speed with which memes can spread in a population, potentially ousting other memes along the way, means that they may not have the longevity required for cumulative selection processes. Even if cumulative selection processes are occasionally established, horizontal transmission leaves them extremely vulnerable. So, the criticism runs, if memes or cultural traits cannot, or only very rarely, attain any degree of longevity, we are unlikely to see complex cultural adaptations; and this limits the explanatory power of memetics, or any other theory of cultural evolution that permits horizontal transmission. Thus, since horizontal transmission undermines cumulative natural selection, we are left with no reason to suppose that horizontally transmitted traits will be fitness enhancing. This is a consequence well known to memeticists: Dawkins cites the celibacy of priests as an example of a meme that decreases reproductive potential.

The potential for horizontal transmission in cultural evolution is held to mark a significant break from mainstream evolutionary theory, and a radical overhaul of the latter would be required if theories of cultural evolution were to be brought under its umbrella. It is

worth noting, however, that horizontal transmission is not unique to cultural evolution, and it seems an especially important phenomenon among bacteria: horizontal or lateral gene transfer (see chapter II.11) also puts pressure on the strictly vertical notions of inheritance assumed by much mainstream evolutionary theory.

A final worry stems from asking whether, even if memes do exist, the meme concept is of any use. The charge here is that memetics is not particularly enlightening: it only dresses up familiar explanations in a slightly different guise. So, for instance, we might allow that clothes fashions are memes, but even if that is the case, memetics does not explain why fashion memes differentially replicate. To explain why one such meme propagates throughout the population while another one perishes in obscurity, we still require reference to local conditions, consumer psychology, and so on. Any value memetics can bring to the explanation of why one meme is fitter than another is parasitic on conventional work done in psychology . And if individual preferences are subject to change over time, then there may be no general and informative theory of cultural evolution to be had; rather, we will have to settle for local explanations that look to shifting preferences. The upshot of this, the argument goes, is that memetics never gets beyond conventional narrative cultural history and cannot provide us with a new scientific framework for understanding culture.

3. Cultural Evolution

Another line of investigation, pioneered by Luigi Luca Cavalli-Sforza, Marcus Feldman, Peter Richerson, and Robert Boyd, concerns the ways in which cultural inheritance can affect evolutionary processes. These models do not assume that cultural inheritance works in the same way as genetic inheritance: in these ways they differ significantly from memetic approaches. Indeed, they model cultural inheritance in ways that depart quite markedly from genetic inheritance. So it is that many of Robert Boyd and Peter Richerson’s models

explicitly assume that an individual's cultural makeup is an error-laden blend—synthesized from, and influenced by, many cultural "parents"—rather than a collection of discretely transmitted, self-replicating, gene-like particles. Their work then focuses on the population-level evolutionary consequences of such an inheritance system; moreover, they tend to concentrate on this form of modeling while remaining noncommitted regarding the precise way in which cultural variants are physically realized.

Such a move can be defended by an appeal to history. Darwin's theory of evolution by natural selection lacked a plausible material theory of inheritance for some time, but this did not prevent Darwin's theory from being useful in the interim. So even without an account of what exactly is inherited in cultural inheritance, work can be done to explain the changes in (cultural) trait frequencies in a population by focusing on the population-level consequences of (cultural) inheritance, selection, mutation, and other forces. So, although cultural evolutionary theorists may deny that cultural change should be understood in just the same way as we understand biological change, their approach remains recognizably evolutionary in style.

All the same, one might think that even if cultural change does not require cultural replicators such as memes, cultural replicators are necessary if cultural change is to be adaptive. As discussed in section 1, longevity is required for cumulative selection processes to operate and complex adaptations to arise. One obvious concern in this context stems from the fact that learning is often very error prone. If an individual hits on a fitness-enhancing behavior, that trait may be lost to future generations either because it is miscopied, or because it is combined with other, less adaptive traits to produce an averaged, "blended" behavior (recall that a particular version of the tango may be an amalgamation of the influences of several teachers and famous dancers). Again, we might fear that cultural traits will not persist long enough for selection to act on them.

Richerson and Boyd argue that these problems are not fatal, and they have developed a number of valuable models to demonstrate why this is the case. These models assume that individuals will pick up cultural traits from a variety of sources and will frequently make mistakes. They also assume that we possess certain kinds of cognitive biases, and they show how these biases can dampen the spread of error in the population. So, even if errors are occasionally made, these isolated errors will tend not to be imitated by others if we possess a *conformist bias*, such that we are more likely to imitate or learn the traits we most commonly encounter. A *prestige bias*, whereby individuals are more likely to imitate a trait possessed by those members of the population who are deemed to be successful, is also thought to keep error in check. It is likely that at least some of the traits possessed by a successful individual will be instrumental in that individual's success. A bias toward imitating successful individuals increases the chances that it is those success-generating traits that are imitated. But if the correct, success-generating trait is not identified, the result is that the individual fails to become successful and so does not become a target for future imitators. Cognitive biases such as these, it is argued, allow for mistakes to be made at an individual level but protect against those mistakes being repeated so widely that they undermine the distribution of cultural traits in a population.

While the existence of such biases can dampen the spread of error, it is far from obvious that they will keep error in check to the extent that cultural inheritance will be robust enough for selection processes to operate. To show that these properties of individual psychology combine to yield population-level inheritance requires some abstract mathematical modeling; much of the novel explanatory payoff of recent work in cultural evolutionary theory comes from the insights gained from this sort of modeling. And the establishment of such population-level consequences is important, for it enables the investigator to revise the constraints one might naively think must bear on cultural inheritance

if cumulative cultural evolution is to occur.

This approach allows cultural evolutionists to agree that cumulative evolution requires that fitness-enhancing cultural traits are preserved in the offspring generation as a whole while denying that this requires faithful transmission between individuals. This move also answers one of our earlier criticisms of meme theory: in taking a population-level perspective, cultural evolution offers genuinely novel explanatory resources that go beyond cosmetic redescriptions of what we already know

One may ask why it should be the case that we are able to learn from nonparents at all, given that horizontal transmission enables the spread of maladaptive traits. Cultural evolutionists have defended the thought that the overall adaptive benefits of learning from nonparents outweigh the overall adaptive costs. Determining how best to live in an environment can be difficult, even dangerous, if one attempts to do so without guidance; one may not be able to tell until it is too late which foodstuffs are nutritious and which are poisonous, for example. Similarly, a prestige bias is an achievable solution to a tricky problem: "determining who is a success is much easier than determining how to be a success" (Richerson and Boyd 2005, 124). The contention here is that cognitive biases that incline us toward imitating or learning from certain individuals may not rule out all maladaptive traits spreading through the population, but these biases are nonetheless more adaptive, overall, than available alternatives.

4. Nonhuman Animal Cultural Evolution

There is widespread acceptance of the existence of some degree of culture (sometimes referred to as "tradition") in nonhuman species. For example, distinct dialects exist within the songs of certain species of birds, and tool use in some primates can vary from group to group within a single species. Japanese macaques are a particularly well studied population, and

differences in everything from their grooming behavior to diet have been documented. The macaques of Koshima Island have developed some remarkable behaviors: to attract the macaques to open land so they could be observed, primatologists left sweet potatoes on a beach. This technique was effective but it did leave the potatoes covered in sand. One member of the troop solved this problem by washing the potatoes in a nearby stream. Soon, her peers were following suit. After a while, the group began to use the sea instead of the stream for washing, preferring the taste the saltwater imparted, and their young took to playing in the sea for the first time. The group also discovered fish discarded by local fishermen and added this to their diet. These changes in their behavioral repertoires have taken place over the course of 50 years and identify them as distinct from other troops of macaques.

This example appears to have some key ingredients for cultural evolution: an individual hits on an innovation, and the innovation spreads throughout the population, which creates behaviorally distinct populations within the species. However, critics argue that this process alone is unlikely to secure cumulative evolution. Not all cultural inheritance involves "observational learning" or imitation, and the worry is that only these forms of social learning will allow the appearance of complex adaptations. We can see why observational learning might be considered crucial with the following example: Certain populations of blue tits learned to remove the foil tops from milk bottles to gain access to the milk inside. The birds' attention was drawn to the milk bottles by the activity of their conspecifics. But it was only through its own trial-and-error exploration of the milk bottles that each tit worked out how to get to the milk inside. If an individual bird happened on a particularly efficient means of removing the foil top, this technique could not be transmitted to any other bird. Social learning of this sort, which does not rely on observational learning or imitation, means that innovations cannot be combined and built on, and cumulative evolution is unlikely to get off

the ground.

Although these sorts of considerations have left some pessimistic about the possibility of significant cultural evolution in nonhuman animals, at worst they merely make complex adaptations as the result of cumulative evolution unlikely. Cultural variations may still play an important role in any evolutionary story of a given species. We mentioned earlier that dairy farming led to selection pressure for the ability to digest lactose. Cultural changes in nonhuman species will alter the selective environment of those species and may instigate a sequence of evolutionary changes. Further, as the macaques of Koshima Island demonstrated, although novel behaviors may not be built on and made more complex, one new innovation can open up previously unexplored parts of the environment and inspire further innovation. At the very least, this sort of example leaves room for the kind of gene-culture coevolutionary models briefly discussed in section 2.

Cultural evolution in nonhumans is an underresearched area, and it remains unclear how widespread observational learning is. Further work will help us establish the significance of culture on the overall evolutionary trajectory of a species, as well as the extent to which we may speak of distinct cultural evolutionary processes in nonhuman species.

5. Defining Culture

According to Richerson and Boyd, "culture is (mostly) information in brains" (2005, 81). Cultural inheritance is then understood as the transmission of this information from one person to the next. So even though Richerson and Boyd deny any strong similarity between genes and cultural variants, they maintain that cultural variants "must be genelike to the extent that they carry cultural information" (2005, 81). If we are to understand what is meant by culture and cultural evolution, then we must understand what is meant by *information*.There is no consensus on the meaning of this term, and the definitions that are

offered—where they are offered at all—are problematic. For example, in their earlier work, Boyd and Richerson offered a definition of information as "something which has the property that energetically minor causes have energetically major effects" (1985, 35). This is a curious definition: presumably it is meant to evoke intuitive examples whereby small informational "switches" (whether they are literally switches in a designed control system or metaphorical "genetic switches" in developmental pathways) have magnified downstream effects on the systems they influence. However, there are plenty of cases of information-bearing relations in which the energetic inequality is reversed. An instrument's display screen can carry information about solar flares: here, an energetically major cause has an energetically minor effect. Perhaps because of these oddities, their more recent work describes information as "any kind of mental state, conscious or not, that is acquired or modified by social learning and affects behavior" (Richerson and Boyd 2005, 5). They later qualify this description with the concession that, in some cases, "cultural information may be stored in artefacts" (2005, 61). Alex Mesoudi, on the other hand, does not offer a definition of information in his work but tells us it is "intended as a broad term to refer to what social scientists and lay people might call knowledge, beliefs, attitudes, norms, preferences, and skills" while also insisting that culture is "information rather than behaviour" (2011, 2–3). However, because skills involve practiced, embodied behaviors, it is unclear how they can count as a form of cultural information.

Eva Jablonka and Marion Lamb defend the use of "information" on the grounds that it provides us with a term that can free us from worrying about the specifics of modes of transmission. It is taken to cover what is transmitted in genetic material, epigenetic material, environmental structures such as nests, behaviors learned from conspecifics, and the kind of knowledge stored in books. Although this kind of abstraction allows us to formulate hypotheses and theories that bear on all these cases of transmission, grouping them together

will highlight what differences exist among them, too, which may encourage us to attend to features of certain types of information and its transmission that we might otherwise overlook. For example, repositories of symbols, most obviously in the form of libraries and computer databases, are vital inheritance systems for humans, allowing the preservation and accumulation of knowledge across generations. Nonsymbolic transmission occurs when some birds inherit their song from adult birds around them. Jablonka and Lamb use the characteristic differences among typical modes of social inheritance in animals and humans to illuminate the impact our own symbolic transmission systems have on human cultural evolution.

In sum, there is some confusion here over what is meant by information and thus how we define culture. The worry is that the term *information* masks some serious issues that any theory of cultural evolution ought to be addressing; we really ought to be clear about what it is we are trying to explain with our theories.

Developing a more fine grained analysis of cultural inheritance, as Jablonka and Lamb suggest the concept of "information" may allow, can only add to the explanatory power of theories of cultural inheritance, but more work is needed first to clarify some conceptual confusion. While more research exists on human cultural evolution than on nonhuman cultural evolution, both areas are in their infancy. Thus, we should not be surprised to find that we are faced with a paucity of data and concepts in need of some untangling. But although the precise details have yet to be ironed out, the research so far has at least demonstrated that cultural evolution is both possible and plausible.

Acknowledgments

The research leading to this work received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant

agreement no 284123.

Further Reading

Boyd, R,. and P. J. Richerson. 1996. Why culture is common but cultural evolution is rare. Proceedings of the British Academy 88: 73–93. *An overview of nonhuman culture and difficulties for theories of cultural evolution in nonhuman animals.*

Cavalli-Sforza, L., and M. Feldman. 1981. Cultural Transmission and Evolution: A Quantitative Approach. Princeton, NJ: Princeton University Press. *One of the first attempts to construct a theory of cultural evolution.*

Dawkins, R. 1976. The Selfish Gene. Oxford: Oxford University Press. *Chapter 11 contains Dawkins's first discussion of his meme concept.*

Lewens. T. 2007. Darwin. London: Routledge. *Chapter 7 contains a detailed critique of memetics.*

Jablonka, E., and M. Lamb. 1995. Epigenetic Inheritance and Evolution: The Lamarckian Dimension. Oxford: Oxford University Press. *Contains a summary of research into animal culture and defends a coevolutionary model for nonhuman animals. Also includes a discussion of the concept of "information" and its place in evolutionary theory.*

Odling-Smee, F. J., K. Laland, and M. W. Feldman. 2003. Niche Construction: The Neglected Process in Evolution. Princeton, NJ: Princeton University Press. *One of the most developed accounts of gene-culture coevolution.*

Richerson, P., and R. Boyd. 2005. Not by Genes Alone: How Culture Transformed Human Evolution. Chicago: University of Chicago Press. *An accessible introduction to cultural evolution by two of the founders of the field.*