7

The Experimental Study of Cultural Transmission and Its Potential for Explaining Archaeological Data

Alex Mesoudi

My objective in this chapter is to review experimental studies of cultural transmission within the field of psychology and to explore how the findings and methods of this work can potentially be used to help explain some of the patterns seen in the archaeological record. Two general experimental methods for studying cultural transmission are identified: Bartlett's (1932) transmission chain method and Jacobs and Campbell's (1961) replacement method. For various historical reasons, neither method became established within psychology, and both remained neglected for a large part of the twentieth century. Recently, however, there has been a renewal of interest in such methods among certain evolutionarily minded psychologists and anthropologists (e.g., Baum et al. 2004; Mesoudi and Whiten 2004; Mesoudi, Whiten, and Dunbar 2006). This renewed interest coincides with a similar revival of cultural transmission concepts within evolutionary archaeology (e.g., Lipo et al. 1997; O'Brien and Lyman 2000a, 2003d).

Here I argue that both fields can benefit from a cross-fertilization of ideas and methods. Specifically, I argue that the experimental methods used by psychologists might be used to help explain archaeological data by simulating in the laboratory the microscale transmission mechanisms that potentially underlie the macroscale changes observed in the archaeological record. This proposal draws inspiration from similar work in evolutionary biology in which patterns in the fossil record are investigated using laboratory-based breeding experiments. Indeed, the theory of Darwinian cultural evolution (Mesoudi et al. 2004) adopted by many of these researchers can facilitate such cross-fertilization of ideas and methods and help generate a unified and productive science of culture (Mesoudi, Whiten, and Laland 2006).

I define cultural transmission as the process by which information is passed by means of social learning from individual to individual. This contrasts with the transmission of genetic information from parents to offspring at conception and with the acquisition of information nonsocially through individual learning from the nonsocial environment. Culturally transmitted information may be stored neurally in the brain or in artifacts. Cultural transmission may occur by means of any one of a number of social learning mechanisms, including imitation, emulation, stimulus enhancement (Whiten and Ham 1992; Whiten et al. 2004), teaching (Csibra and Gergely 2005), and spoken or written language (Smith 2002a, 2002b).

THE EXPERIMENTAL STUDY OF CULTURAL TRANSMISSION IN PSYCHOLOGY

For reasons of space and to maintain focus, this review covers only experimental studies of human cultural transmission. Observational studies of cultural transmission, such as those that examine the diffusion of innovations (Rogers 2003) or rumor propagation (Rosnow 1980, 1991), are not included. Also not considered are experimental studies of social learning in which transmission occurs only from a demonstrator to a single learner (e.g., Bandura et al. 1961; Horowitz 2003; Want and Harris 2002). To count in a study of cultural transmission, information must be passed along a chain of more than two individuals, for only then can conclusions be drawn about persistence and long-term cumulative cultural change. Finally, experimental studies of nonhuman cultural transmission (e.g., Curio et al. 1978; Galef and Allen 1995; Laland and Plotkin 1990) are also not included.

The Transmission Chain Method

The transmission chain method represents the

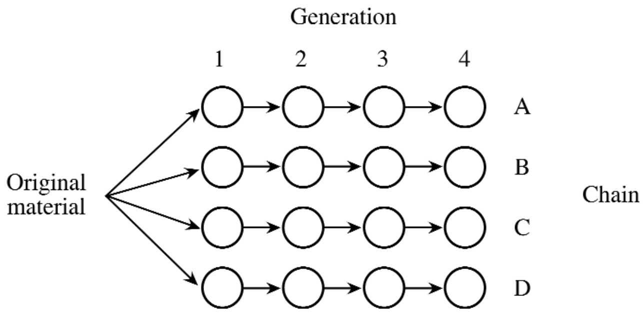

Figure 7.1. The general design of an experiment using the transmission chain method.

simplest experimental procedure within psychology for studying cultural transmission. The method is similar to the children's games "Chinese Whispers" or "Broken Telephone." The first participant in the chain reads or hears some verbal material and then attempts to recall it. This recall is then given to the second participant, who does the same, and so on along the chain. The general design of the interaction is illustrated in Figure 7.1. Changes that occur to the material as it is transmitted along the chain, or the different degradation rates of different types of material, can then reveal systematic biases in transmission. Although highly simplified compared with the transmission of real human culture, the transmission chain method affords a high degree of experimental control and, as will be seen below, has the potential to provide important contributions to the study of human cultural transmission. Indeed, it has been described by Plotkin as "close to an experiment tailor-made for those interested in culture" (1995:219).

This method has its origin in the work of Sir Frederic Bartlett, who in his classic book *Remembering* (1932) developed the transmission chain method (he termed it "method of serial reproduction") and described a series of studies that used the method. One of Bartlett's main objectives was to move away from abstract, asocial methods of studying memory, such as having participants memorize sequences of numbers or nonsense letter strings, and toward a more ecologically valid method. He therefore used more meaningful stimulus material such as stories and devised the transmission chain method to study the social aspects of memory. Unlike many of his contemporaries, however, Bartlett was also aware of the wider implications of his transmission chain method for studying human culture: "Elements of culture, or cultural complexes, pass from person to person within a group, or from group to group, and, eventually reaching a thoroughly conventionalised form, may take an established place in the general mass of culture possessed by a specific group" (1932:118).

Bartlett's precise method consisted of a participant reading some written material twice at normal reading speed and then performing a distractor task for 15–30 minutes before recalling the material. That recall was then taken by Bartlett (or another experimenter) and given to the next participant in the chain, who went through the same procedure, as did each subsequent participant. Bartlett reported results for various types of material: two folktales, "The War of the Ghosts" (from Native American culture) and "The Son Who Tried to Outwit His Father" (from the Congo); passages describing a cricket match, an air raid, and how to play tennis; a joke; two arguments; and a series of pictures. The participants were predominantly Cambridge undergraduates, with some replications using undergraduates from India.

A general finding for all of the studies and participants was that the material rapidly became considerably shorter in length and lost much of its detail, with only the overall gist being preserved. A second general finding was that participants tended to distort the material to make it more coherent and consistent with their own preexisting knowledge. "The War of the Ghosts," for example, contained many supernatural elements that were nonsensical to the English participants and were subsequently removed or replaced with more familiar events. These two processes, loss of detail and rationalization, led Bartlett to propose that remembering is primarily a reconstructive process and hardly ever a process of exact replication. Only the gist of the material is preserved and is rebuilt around preexisting knowledge structures, or schemas. Bartlett also found that the folk stories were transmitted with greater accuracy than any of the other material, which he explained by arguing that people already possess implicit story schemas that contain the structure of a typical folktale, thus aiding recall.

During the 30 years following the publication of Bartlett's *Remembering*, a series of transmission chain studies was published, primarily in the *British Journal of Psychology* (probably attributable in part to Bartlett's position as editor of that journal). These studies all shared Bartlett's general method but varied in the material used and participants tested. The transmission of stories, for example, was studied using chains of children (Northway 1936) and adults of different professions (Maxwell 1936) and nationalities (Talland 1956). Allport and Postman (1947), Ward (1949), and Hall (1951), meanwhile, all studied the transmission of pictorial rather than written stimuli. The results of these studies largely confirmed Bartlett's original findings that there is a general reduction in the length or complexity of the material, that much of the detail is lost, and that only the overall impression of the material is preserved.

Although Bartlett's work was groundbreaking, it had a number of shortcomings. First, his analyses were entirely subjective and qualitative, with no attempt at quantitative or statistical analyses. Second, details of the precise method that he used are sketchy at best. No detail of the distractor task was given, and there was no discussion of why it varied in length from 15 to 30 minutes. Neither was any mention made of the instructions given to the participants or whether they were standardized, despite the use of several different experimenters. Indeed, Gauld and Stephenson (1967) found that instructing participants to reproduce only information that they were absolutely certain was in the original led to significantly fewer errors in the transmission of "The War of the Ghosts." Third, Bartlett's conclusions were all drawn on single chains. He himself noted that "the main turning points [in the reproductions] are the work of individual interests or idiosyncrasy, and in the total social product the outstanding individual can be shown to have played a crucial part" (1932:126). Surely it is improper to draw general conclusions about human cultural transmission if the data are so susceptible to individual idiosyncrasies, which also make independent replication of experimental findings almost impossible. Most of the early studies that followed Bartlett inherited not only his methods but also the flaws, although there were occasional attempts at quantitative analysis (Talland 1956) and the use of multiple chains (Brissey 1961).

These studies should not be criticized too harshly for failing to adopt standards of scientific practice that seem obvious today but had yet to be established when the studies were carried out. Indeed, as time passed, transmission chain studies benefited from modern scientific standards, such as the reporting of standardized experimental procedures, the use of statistics, and sample sizes large enough to yield significant results. However, there was also a marked decline in the popularity of the transmission chain method during the latter part of the twentieth century, perhaps as a result of the rise of cognitive psychology, which has tended to ignore social processes. Only in the last few years has the transmission chain method been revived by a handful of researchers interested in the experimental study of social transmission and culture. Whiten and I (Mesoudi and Whiten 2004) used the method to test a hypothesis derived from the cognitive psychological work of Schank and Abelson (1977), who have argued that people's knowledge of routine events, such as going to a restaurant, is represented in hierarchically organized "action scripts." The overall goal—for example, to go to a restaurant—constitutes the highest level of this hierarchy and can be subdivided into a series of constituent subroutines—enter the restaurant, sit down, order food, eat the food, pay, and leave—each of which can in turn be subdivided into a series of low-level actions such as pick up the menu, read it, decide on the food, and tell the waiter.

Whiten and I (Mesoudi and Whiten 2004) tested this hypothesis in the context of cultural transmission by passing descriptions of three everyday events (getting up, going shopping, going to a restaurant) expressed entirely at the lowest hierarchical level along 10 chains of four participants each. Consistent with Schank and Abelson's (1977) script hypothesis, the low-level actions were gradually subsumed into their higher-level parent goals as they were transmitted along the chains, demonstrating that the participants were spontaneously imposing a hierarchical structure onto the event knowledge and transforming the events into higher levels of that structure. So, rather than the participants just simplifying the material (Bartlett 1932), Whiten and I have demonstrated that this simplification occurs in a systematic manner according to a hierarchical bias in cultural transmission.

A version of the transmission chain method has also been used within the emerging field of experimental economics, although it does not directly derive from Bartlett's work. "Intergenerational games" have been set up to study the transmission of behavioral traditions along chains of participants. For example, Schotter and Sopher (2003) had successive pairs of participants play the "Battle of the Sexes" game, in which the two players must choose one of two possible options without communication. (The name "Battle of the Sexes" comes from the original framing of the game, in which a husband and wife must decide how to spend an evening. The husband prefers a baseball game, whereas the wife prefers the theater, although both prefer to be together than alone.) The payoffs are determined by two rules. First, if the players choose different options, then neither player gets any payoff, whereas if both players choose the same option, then they both get a payoff. Second, the two options differ in their payoffs to the two players. If both players choose the first option, then Player 1 gets a larger payoff than Player 2, whereas if both players choose the second option, then Player 2 gets the larger payoff. Hence, the first rule encourages the players to cooperate by coordinating their responses, whereas the second rule creates a conflict in that one player will always get a larger payoff than the other. Transmission between successive generations was effected by giving each pair two sources of information from the previous pair(s) in the transmission chain. This was either a behavioral history of every previous generation—the options chosen by previous players and the outcomes they received—or explicit advice given by the previous generation as to which option they should choose and why. After approximately 50 generations, one of these two sources of information was removed to assess the effect of each independently.

Schotter and Sopher (2003) found that the repeated games exhibited long periods during which both players chose one of the options interspersed with brief periods of instability and rapid change. This tendency to coordinate over many generations was attributed to the role of advice: Removing the option to view the behavioral history of previous generations had no significant effect, whereas removing the explicit advice from the previous generation significantly disrupted the stable conventions. This suggests that conformity to explicitly expressed social norms played a powerful role in maintaining these cultural traditions.

In conclusion, Bartlett's (1932) transmission chain method constitutes a simple yet effective means of testing hypotheses regarding cultural transmission under controlled experimental conditions. Although early studies featured a number of methodological flaws, later studies demonstrate that the transmission chain method can be updated to meet modern standards of scientific practice, specifically by the use of multiple parallel chains, quantitative and statistical analysis, and properly standardized and controlled methods. Whiten and I (Mesoudi and Whiten 2004), for example, used Kintsch's (1974) propositional analysis to divide each resulting recall into measurable units of meaning, allowing quantitative statistical analyses to be performed. One general limitation of the transmission chain method, however, lies in the linearity of the one-to-one chains. Actual cultural transmission may frequently involve more than one model and more than one receiver, and to study these more group-based aspects of transmission, a slightly different method is needed.

The Replacement Method

In the replacement method, originally proposed by Gerard et al. (1956), a norm or bias is established in a group of participants, and one by one the participants are replaced by new, untrained participants (Figure 7.2). Each replacement represents one cultural generation. The degree to which the norm remains in the population during successive replacements/generations represents a measure of its transmission to the new members. Jacobs and Campbell (1961) used the replacement method to study the persistence of an artificially exaggerated perceptual judgment of the "autokinetic effect," a perceptual illusion in which a stationary point of light is perceived as constantly moving slightly when viewed in an otherwise pitch-black room. In earlier work by Sherif (1936), a group of participants was shown this illusion and asked individually to estimate how much they thought the light was moving. The group was, in fact, composed of only one genuine participant, the rest being confederates of the experimenter who had been instructed to give unrealistically exaggerated estimates of the light's movement. Sherif's (1936) now-classic finding was that the majority of the participants he tested gave similar estimates to the

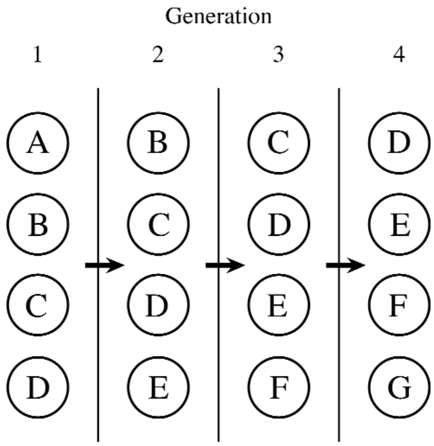

Figure 7.2. The general design of an experiment using the replacement method.

confederates despite that estimate being obviously false, illustrating the powerful effect of conformity in group settings. Jacobs and Campbell (1961) repeated Sherif's (1936) experiment with the addition that, after each round of estimating, one group member was replaced with a new, naive participant. Significant evidence of the inculcated norm was found for about four or five generations after the replacement of all of the confederates, after which the perceptual judgment returned to that exhibited by naive control groups. This demonstrates some degree of transmission but no long-term persistence.

Insko et al. (1980) used the replacement method to simulate the trading of goods. Three groups of four participants each were taught to produce different products (paper models). Earnings were maximized if these different products were combined, encouraging trade. One group was placed in a position of economic and communicative superiority: Its products were more valuable than those of the other two groups, and all trade had to be conducted through it. After each round of trading, one member of each group was replaced with a new participant, with a total of nine replacements (generations). The results showed that although the dominant group earned more than the other two groups, all groups increased their earnings over the successive generations. This increased productivity was attributed to increasingly efficient trading and division of labor, rules concerning which were being transmitted to each new generation.

More recently, Baum et al. (2004) used the replacement method to study the transmission of traditions in an anagram-solving task. Groups of participants could choose to solve an anagram printed on a red or blue card. The red anagrams gave a small but immediate payoff, whereas the blue anagrams gave a larger payoff but were followed by a "time-out," during which no anagrams could be solved. By manipulating the length of this time-out, the experimenters were able to determine which of the two anagrams gave the highest overall payoff—that is, where the blue time-out was short, blue was optimal; where the blue time-out was long, red was optimal. Every 12 minutes one member of the group was replaced by a new participant. It was found that traditions emerged, defined by whichever choice gave the highest payoff under each experimental condition (the optimal choice), with existing group members instructing new members in this optimal tradition by transmitting either accurate or inaccurate information about payoffs or (less frequently) through coercion.

In conclusion, the replacement method offers a useful complement to the transmission chain method. Whereas the transmission chain method has been mainly used to study the transmission of verbal material along one-to-one chains, the replacement method has been used to study the emergence and persistence of groupwide behavioral traditions. The replacement method is therefore more suited to investigating the effects of social/interpersonal factors on cultural transmission, such as conformity (Jacobs and Campbell 1961) and power (Insko et al. 1980).

THE HISTORICAL STUDY OF CULTURAL TRANSMISSION IN ARCHAEOLOGY

Long before experimental psychologists started running their own simulated transmission chains, archaeologists were collecting evidence of actual cultural transmission from historical artifacts. This was achieved through the method of seriation, in which artifacts or collections of artifacts are ordered according to their similarity: The more features they share, the closer they are in the order; the fewer they share, the farther apart they are placed. As noted by O'Brien and Lyman (2000a), the use of seriation rests on two assumptions. First, the similarity of features corresponds to closeness in historical time, so that a sequence of gradually changing artifacts is also a chronological sequence of artifacts. This is the assumption of *historical continuity*. Second, the reason why similarity of features corresponds with closeness in time is that the artifacts in the sequence are part of a lineage of inherited information, causally connected by cultural transmission. O'Brien and Lyman (2000a) name this process *heritable continuity*.

Although the existence of historical continuity often denotes heritable continuity, it is possible that the former could occur without the latter. An example of historical continuity without heritable continuity is the anti-Darwinian notion that cultures evolve through a sequence of fixed stages such as band, tribe, chiefdom, and state. Another example is the case of convergence, in which the same trait emerges more than once independently, with no connection through transmission. O'Brien and Lyman (2000a) argue that heritable continuity can be detected by showing that the different types of artifacts in the seriation overlap in time. For example, type A might occur during periods 1–3; type B, during periods 2–4; and type C, during periods 4–6. Types A, B, and C therefore show heritable continuity: Their similarities are the result of inheritance through cultural transmission. In general, however, if the change is continuous and gradual, then heritable continuity can be assumed.

The earliest recorded use of seriation is probably Evans's (1850) sequence of gold coins found in Britain, dating back to before the Roman invasion of 54 B.C. These seriations are shown in Figure 7.3. Evans's sequence of artifacts begins with coins featuring the head of Phillip II of Macedon on one side and a horse-drawn chariot on the other. As the seriation progresses, these images become less lifelike and more schematized, with the head of Phillip II eventually being lost altogether and replaced with an abstract pattern. The horse-drawn chariot, meanwhile, loses the chariot and becomes either a horse or a horse with a rider. This simplification resembles the changes that Bartlett (1932) observed in his experimental chains of stories and pictures (although in other lineages of Evans's [1850] coins the simplification is reversed and the design again becomes naturalistic).

Seriation was later used by Pitt Rivers (1875a) to reconstruct lineages of stone tools, copper and bronze axes, and muskets, all demonstrating gradual change in which each specimen can be seen as a slight modification on the one before. Petrie (1899) used seriation to order approximately 4,000 ceramic vessels excavated along the Nile. There are several noteworthy aspects to his seriations. First, there is considerable overlapping of features,

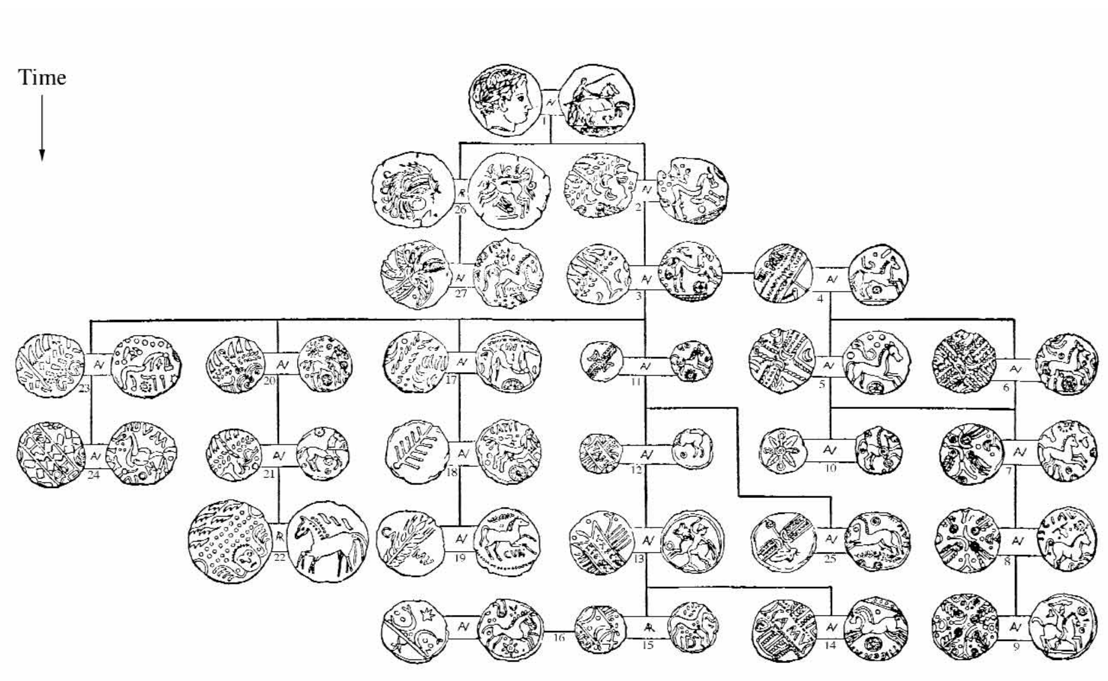

Figure 7.3. Evans's (1850) sequence of British coin designs from around the time of the Roman invasion of 54 B.C. (from O'Brien and Lyman 2000).

confirming the assumption of heritable continuity. Second, various lineages of designs can be observed in the seriation, with some lineages becoming extinct and others merging to form new lineages. Third, the seriation allowed Petrie to speculate that the handles on the vessels, which originally served a practical purpose, gradually became less functional and more decorative, eventually forming a kind of vestigial feature. Kidder (1915) performed a similar seriation using pottery from New Mexico, finding that the decorative patterns on the pottery became gradually less intricate through time. O'Brien and Lyman (1999, 2000a) label this kind of seriation *phyletic seriation*.

Despite their use of the method, archaeologists rarely were explicit in attributing the historical similarity seen in their seriations to cultural transmission (O'Brien and Lyman 1999, 2000a), although it sometimes was implicitly stated (e.g., Kroeber 1916). By about 1920, however, even this implicitness was gone, and seriation was viewed strictly in terms of its ability to measure the passage of time (O'Brien and Lyman 1999, 2000a). Several recent studies have shifted the focus of seriation back to heritable continuity. Lipo et al. (1997) have used a form of seriation to examine prehistoric cultural transmission in the lower Mississippi Valley (see ch. 10), and O'Brien and his colleagues (O'Brien and Lyman 2002b, 2003a; O'Brien et al. 2001; O'Brien et al. 2002) use it to help construct phylogenetic orderings of projectile points from the southeastern United States.

INTEGRATING THE EXPERIMENTAL AND HISTORICAL STUDY OF CULTURAL TRANSMISSSION (I)

The central argument of this chapter is that both the psychological experiments and the archaeological studies discussed in the previous two sections are studying the same phenomenon at different time scales and in different contexts. Both study the cultural transmission of information by means of social learning from individual to individual along a chain or lineage. The experimental studies look at simulated small-scale transmission along chains of a few individuals or within small groups under controlled laboratory conditions. The archaeological studies look at actual, unsimulated, large-scale transmission through entire populations involving countless unidentified individuals from successive biological generations.

Despite these differences, the same process—cultural transmission—lies at the heart of both. Consequently, it may be both possible and mutually beneficial to integrate the findings, methods, and theories of the two fields. The historical lineages observed in the archaeological record should be amenable to investigation in the laboratory using the experimental methods described earlier. Such an approach could provide valuable insights into the underlying transmission mechanisms that were originally responsible for the actual cultural patterns. Real history, of course, is vastly more complex than anything that could ever be re-created in the laboratory. The value of experimental simulation, however, is not that reality can be exactly reproduced. Rather, experimental and historical methods should be seen as complementary, each with its own advantages. What experiments may lack in ecological validity, for example, they gain in the following:

1. Experiments allow history to be rerun in the laboratory, where changes can be directly observed, recorded, and measured. Barring the invention of a time machine, real history can never be directly observed.
2. This rerunning of history can be repeated in multiple chains. Replication allows the investigation of such factors as contingency and random drift.
3. Experimental methods allow single variables to be isolated and manipulated, whereas it is difficult to retrospectively infer the effects of single variables on historical data.
4. Experiments generate complete data, unlike historical data, which often are fragmentary.
5. Different steps in the transmission chains can be directly compared, for example, in terms of relative efficiency or fitness.

Although I cannot yet give any actual empirical examples of these advantages as applied to cultural transmission, I can take advantage of the fact that biological evolution constitutes an inheritance system that resembles human culture in important ways (Mesoudi et al. 2004) and provide an example of the experimental simulation of historical (paleontological) data within evolutionary biology.

The Experimental Simulation of Historical Data in Evolutionary Biology

I believe it is more than a coincidence that many of the researchers who have sought to reintroduce cultural transmission into their fields (e.g., Baum et al. 2004; Lipo et al. 1997; Mesoudi and Whiten 2004; O'Brien and Lyman 2000a, 2002b) share an evolutionary view of culture. That is, these researchers are committed to a theory of cultural evolution (e.g., Mesoudi et al. 2004; Richerson and Boyd 2005) in which human culture—the entire body of socially transmitted knowledge that is passed on from individual to individual, generation to generation, by means of cultural transmission—is governed by the same general Darwinian evolutionary principles as govern changes in biological species.

Transmission (inheritance) plays a central role in any theory of evolution, be it biological or cultural. Darwin discussed at great length the inheritance of biological traits and the importance of inheritance to his theory of biological evolution, stating that "any variation which is not inherited is unimportant for us" (Darwin 1859:75). A focus on transmission in culture naturally leads experimentalists to adopt methods that permit the study of the social aspects of cognition, memory, and decision making and leads archaeologists to recognize explicitly that seriations may be causally connected by cultural transmission (heritable continuity) and not simply by time (historical continuity). Perhaps more importantly, Darwin's theory of biological evolution served to unify and integrate the diverse and hitherto separate branches of biology—the "Modern Synthesis" of the late 1930s and 1940s (Huxley 1942). Indeed, a key aspect of the synthesis was the realization that the macroevolutionary patterns studied by paleobiologists are the result of the microevolutionary mechanisms studied by population geneticists (Mayr 1982).

Viewing human culture within a similar evolutionary framework allows a similar integration between the micro- and macrolevels of cultural change, leading to a recognition that the macroevolutionary patterns studied by archaeologists are the product of the microevolutionary transmission mechanisms studied by psychologists (Mesoudi, Whiten, and Dunbar 2006). We can also take advantage of the similarities between biological and cultural change and look at evolutionary biologists' attempts to experimentally simulate historical (paleontological) data in order to illustrate the advantages listed in the previous section and inform the proposals outlined in the following section.

For example, Lenski and colleagues (e.g., Cooper et al. 2003; Lenski and Travisano 1994; Lenski et al. 1991) have conducted a series of experiments in which laboratory populations of the bacterium *Escherichia coli* are exposed to controlled environmental conditions and allowed to breed for approximately 10,000 generations. Lenski and Travisano (1994) describe one such study in which an initially genetically identical population of *E. coli* was allowed to breed in an unchanging environment. Mean cell size was found to increase rapidly during the first 2,000 generations before entering a period of stasis during the next 8,000 generations. Lenski and Travisano argue that this pattern of rapid change followed by lengthy stasis resembles the punctuated equilibria (Eldredge and Gould 1972) seen in the fossil record. (Note that the data obtained by Schotter and Sopher [2003] described above also showed long periods of stasis punctuated by rapid change.) Although this experiment involved only a single species under specific environmental conditions, the fact that punctuated equilibria were observed in the continuously measured laboratory data suggests that such instances may not simply be by-products of gaps in the fossil record. They might instead constitute a genuine pattern of change.

Lenski and Travisano (1994) repeated the experiment using 12 replicate populations, again all initially genetically identical (both within and between populations) and breeding under identical environmental conditions. A similar pattern of change in cell size was found in all populations—a rapid increase, followed by stasis—ruling out random drift as an explanation for the observed change in cell size. Lenski and Travisano then placed populations that had evolved for a number of generations in direct competition with their ancestors (which in the case of *E. coli* can be stored in refrigeration) in order to measure changes in relative fitness. The finding that relative fitness showed similar change over time as cell size, with a rapid increase and then stasis, ruled out the possibility that adaptation had continued after cell size stabilized.

Lenski and Travisano's (1994) study demonstrates a number of the advantages of experimental simulations listed above: History can be rerun in replicate chains, generating complete data and allowing different steps in the chain to be directly compared. However, studies that attempt to integrate experimental population genetics and paleobiology are relatively rare in the biological literature. Indeed, some paleobiologists argue that their studies of macroevolution cannot be integrated with studies of microevolution such as the experiments and models of population biologists. Kemp (1999), for example, argues that a chronological gap exists between the two, implying that experiments can never capture the vast time scales studied by paleobiologists.

Might the same chronological gap prevent a similar integration within the social sciences? There are two reasons why I think not. First, as noted by Bell, the chronological gap between paleobiology and population biology "does not exist because it is impossible to collect data at the appropriate time scale; it exists because neither paleobiologists nor population biologists have invested enough effort in collecting these data" (2000:1459). Pioneering studies such as those of Lenski and colleagues accord with this view. Second, and more fundamental, the chronological gap as it applies to human culture is simply much smaller. Whereas meaningful biological change frequently occurs only over millions of years, the cultural patterns observed by archaeologists occurred instantaneously in comparison. Moreover, biological transmission experiments are not feasible using species with generations of more than a few days, whereas in the case of culture, generations may be extremely short. Furthermore, there is a much greater incidence of horizontal transmission (transmission within a single biological generation [Cavalli-Sforza and Feldman 1981]) in cultural inheritance than in biological inheritance, potentially allowing much greater change in culture than in species during the same period of time. An originally vertically transmitted culturally acquired behavior, or the knowledge required to construct a tool, may be learned, mastered, and transmitted horizontally in a matter of minutes, hours, or days. Finally, a species found in the fossil record has almost certainly gone extinct or else undergone significant change, whereas even relatively old culturally acquired behaviors, such as stone knapping, are still practiced today by traditional societies or amateur enthusiasts and are thus more amenable to experimentation.

In conclusion, we have seen how some biologists have attempted to bridge the gap between macroevolutionary studies of paleontology, on the one hand, and experiments simulating the processes of microevolution, on the other. Such studies are rare because of the perceived chronological gap between the two approaches. These limitations are not as restrictive, however, in the realm of culture, and the next section explores this integration further.

INTEGRATING THE EXPERIMENTAL AND HISTORICAL STUDY OF CULTURAL TRANSMISSSION (II)

What would an experimental simulation of archaeological data look like, and what specifically would it achieve in terms of the advancement of archaeological theory? It is important to recognize that the experimental simulation of historical lineages is unlikely to be able to investigate the effects of large-scale environmental factors such as temperature or subsistence. Its use lies instead in studying social factors that affect transmission. Such factors may include the potential social learning mechanisms involved—for example, imitation, teaching, or language; different kinds of transmission biases—for example, Boyd and Richerson's (1985) indirect or conformist biases; or factors such as intergroup competition (e.g., Insko et al. 1980).

Surprisingly, one early and largely forgotten transmission chain study attempted to do just what is proposed here—integrate experimental and historical work. Ward (1949) used Bartlett's (1932) transmission chain method to simulate the transmission of the same coin designs that were studied by Evans (1850) and are shown in Figure 7.3. The original and earliest Macedonian design was used as the starting point in a series of chains, the results of which were compared with the artifact chain obtained by Evans. Although Ward claimed that changes in his experimental reproductions did indeed mirror actual changes in the historical artifacts, there unfortunately was no quantitative analysis, and the few reproductions that were published in his article are fairly unconvincing. Nevertheless, the idea of simulating the actual transmission of an artifact was there.

Ward's study was exploratory in nature rather than being theoretically driven. With the benefit of over 50 years of additional theoretical work in archaeology, we can today perhaps provide a more rigorous rationale and precise predictions regarding Ward's intuitions. Specifically, we can draw on Dunnell's (1978a) distinction between style and function, in which stylistic traits are selectively neutral and spread through drift, and functional traits have a detectable effect on fitness and spread through selection. The changes observed by Evans are most likely stylistic, as he acknowledged (see O'Brien and Lyman 1999, 2000a). Hence if history were to be rerun, then the designs on the coins would have followed an entirely different course (within certain limits imposed by the original design, the size of the coins, and so on).

Experimental simulations allow us to in a sense rerun history to see whether this claim is upheld. Although Ward (1949) claimed that key changes in the artifacts had been reproduced in his laboratory study, it is hard to tell from his figure, and he offered no quantitative or systematic statistical analyses. A modern transmission chain study might run multiple linear chains, all starting with the earliest coin design, and compare the changes with Evans's (1850) seriation and then compare the changes in the multiple chains using quantitative statistics. If the coin designs are stylistic and subject to drift, then the chains should show no significant similarities and thus fluctuate randomly over transmission generation. If, on the other hand, each chain shows a similar pattern of change consistently in the same direction, then there may be a functional explanation. Again, it should be emphasized that such a simulation could never provide a definitive answer to the question, "Are the changes in coin designs observed by Evans the result of stylistic or functional change?" There were undoubtedly many more factors affecting the actual transmission of coin designs across medieval Europe than could ever be re-created in a laboratory. Rather, experimental simulation is one potentially useful means of investigating historical change, to be used in conjunction with historical data and computer models.

Another area of theory that might be explored using simulation is Boyd and Richerson's (1985) cultural transmission biases. Boyd and Richerson propose that human cultural evolution features a number of transmission biases with no obvious analogy in biological transmission but which have potentially significant effects on large-scale cultural dynamics. Some of these biases have recently been invoked by archaeologists to help explain historical data. For example, Bettinger and Eerkens (1997, 1999; see also ch. 3) have identified regional differences in the spread of bow and arrow technology in the western United States around A.D. 300–600. Specifically, the weight and the width of projectile points from central Nevada were found to be correlated, and there is little variation in point form. In contrast, contemporary projectile points from eastern California show independent variation in both weight and width. Bettinger and Eerkens argue that these differences reflect the operation of different underlying transmission mechanisms. The central Nevada points vary little because they were transmitted by *indirect bias* (Boyd and Richerson 1985), in which point designs were copied wholesale from prestigious individuals. If everyone is copying the same high-status individual, then all the point designs will end up similar. With respect to the eastern California points, Bettinger and Eerkens argue that they were transmitted according to a process of *guided variation* (Boyd and Richerson 1985), in which a trait is copied from a model but then modified by individual trial and error according to the needs of the copier. This latter component of individual modification resulted in the independently varying width and weight seen in the Californian points.

Of course, in this case the operation of indirect bias or guided variation cannot be directly observed and can only be inferred from the archaeological evidence. We can, however, be more confident in this inference if experimental methods generate similar findings. Given that we are interested here in group-level factors, the replacement method might be more suitable than the transmission chain method. Successive groups of participants might be asked to design computer-generated "virtual projectile points," with the possible means of transmission manipulated in a manner similar to that of Schotter and Sopher (2003) discussed earlier. One condition could simulate indirect bias, where each successive generation can view the point designs of each member of the previous generation together with their associated "virtual hunting success," thus allowing them potentially to copy the most successful. Another condition could simulate guided variation, where each successive generation is allowed to modify the virtual point of a participant from the previous generation. Bettinger and Eerkens's (1997, 1999) proposal would be supported if the experimental data mirror their archaeological data—that is, if the indirect bias condition generates traditions in which the point designs vary less than those in the guided variation condition. This experimental paradigm mirrors recent work by Eerkens et al. (2006) in which phylogenetic analyses were performed on data generated by a simple computer model rather than on data from actual participants, as proposed here.[1]

Another of Boyd and Richerson's (1985) cultural transmission biases is frequency-dependent bias, in which cultural traits are copied on the basis of their popularity in the population. This bias has been used by Kohler et al. (2004) to explain village formation in New Mexico around A.D. 1325–1525. They found that ceramic patterns from larger settlements show significantly less diversity than would be expected from a neutral model of cultural change. They therefore infer that individuals were exhibiting a positive frequency-dependent, or conformist, bias when designing vessels. Kohler et al. argue that this within-group conformity—the result of increased within-group cooperation—was one factor responsible for maintaining large settlements (see Henrich and Boyd 1998; Richerson and Boyd 2005).

Conformity might also be simulated in an experimental setting. The virtual projectile point paradigm could be adapted so that each generation has information about only the frequency of different point designs in the previous generation rather than their associated payoffs, allowing them potentially to select the most popular design. Alternatively, a similar paradigm to that in Jacobs and Campbell's (1961) study of conformity could be adopted, in which confederates are used to establish a predetermined majority. Kohler et al.'s (2004) contention that conformity is characteristic of large cooperating groups might also be tested by dividing participants into groups and allowing them to interact, as in Insko et al.'s (1980) study. Finally, it might be instructive to see which of the three proposed transmission biases—indirect bias, conformity, or guided variation—is preferred by participants and under what environmental conditions.

SUMMARY

In the previous section I sketched some ideas for using the experimental methods described at the beginning of this chapter to simulate and help explain some of the archaeological patterns outlined subsequently. The transmission chain method might be used to produce multiple replicate chains of historical lineages in order to explore the style–function dichotomy, whereas the replacement method could be used to explore the role of cultural transmission biases in generating regional differences in trait distributions. I am not claiming that such experimental methods can provide a definitive resolution to these issues but, rather, that they can be used in conjunction with historical data and computer models to provide a more comprehensive answer than any one of these methods alone. This cross-disciplinary borrowing of methods, tools, and hypotheses follows naturally from viewing human culture in the context of a larger unifying evolutionary framework.

ACKNOWLEDGMENTS

I am grateful to Andrew Whiten for providing useful comments on this chapter and to Mike O'Brien for inviting me to contribute to this volume.

NOTE

1. Recent work conducted after this chapter was written used actual participants, who were given the task of designing "optimal" projectile points (Mesoudi and O'Brien 2007). The results confirm that periods of indirectly biased cultural transmission are associated with low variation and highly correlated virtual point attributes, whereas periods of isolated individual learning (simulating guided variation) are associated with high variation and poorly correlated virtual point attributes, as predicted.