

Cultural Transmission, Phylogenetics, and the Archaeological Record

Jelmer W. Eerkens, Robert L. Bettinger, and Richard McElreath

Phylogenetic studies of material culture (e.g., Jordan and Shennan 2003; Lyman and O'Brien 2000; O'Brien et al. 2001; O'Brien and Lyman 2002a, 2003a; Tehrani and Collard 2002) are gaining popularity in the anthropological sciences and are giving us new insights into the transmission of human culture in the past and present. But in their focus on phylogeny such studies tend to bypass extended discussion of the processes of transmission and modification that produce variability in the items of material culture. Yet as numerous anthropologists and others have discussed (e.g., Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981; Henrich 2001), cultural information can be passed on and modified in a multitude of ways. Surely, these different modes of transmission affect our ability to reconstruct and interpret phylogenetic trees.

Here we attempt to fill this vacuum in applications of phylogenetics to material culture by examining how well such analyses are able to reconstruct the descent of traits inherited through different cultural-transmission systems. In particular, we are interested in transmission systems of the kind discussed and defined by Boyd and Richerson (1985). We use computer simulation to assess the effects of different forms of cultural transmission on cultural phylogenies and evaluate these concepts to a dataset comprising projectile points from the Great Basin of North America.

Transmission Processes

Before describing our simulations, we briefly define the transmission systems modeled here. "Guided variation" entails transmission of information from a social model to a recipient, who then experiments with that information

in search of an optimal or better character state. Traits are acquired and modified individually in piecemeal fashion. For example, an individual might initially learn how to make a projectile point from his father, but later improve techniques of manufacture or selection of raw materials based on his own experience. "Conformist transmission" involves cases in which individuals survey the social-model pool and find the most common attribute or attribute state and adopt it. Again, this is done piecemeal for different attributes. In this form of transmission, an individual making an arrow point would seek to identify the most common raw material, the most common flaking tool, the most common point length, and so on. All of these traits would be combined during process of point construction. "Indirectly biased transmission" involves cases in which individuals survey the social-model pool and select a single social model on the basis of an indicator trait and then copy the entire cultural repertoire of this model. Here traits are not adopted one at a time but together as a package. For example, success in some communities might be judged by the number of surviving children. An individual might preferentially imitate the beliefs and behaviors of members of the community with the most children over the age of, say, five.

There are good reasons to suspect that all these potential transmission mechanisms exist in the real world. First, theoretical analyses suggest that no general cultural-learning strategy is optimal in all circumstances. Instead, social and ecological contexts strongly influence the efficacy of particular models and the payoffs that result from integrating their behaviors and beliefs into one's cultural repertoire (Boyd and Richerson 1985; Henrich and Boyd 1998). While social learning may have a low cost, it sometimes pays to ignore the behavior of others and rely on individual learning, or guided variation, to acquire behavior (Boyd and Richerson 1985; Rogers 1988). Second, evidence from psychology and other social sciences suggests that people employ a diversity of learning strategies, including the three outlined above. Bandura (1973, 1977), for example, found that people select models based on cues such as observed payoffs, prestige, and specific similarity between themselves and the model.

Psychologists (e.g., Asch 1952, 1955; Milgram 1974; Moscovici 1985; Rock 1990; Wilder 1977) have also observed a tendency for individuals to conform to the beliefs or behaviors of a majority. In addition, psychologists and economists working on individual choice and judgment have studied how people apply experience to modify existing beliefs and preferences, providing evidence for a mechanism that will lead to guided variation in natural circumstances. Lastly, evidence from the literature on the diffusion of innovations, such as hybrid corn or computers, suggests that people employ different imitation and learning strategies (Henrich 2001; Henrich and Boyd 2001; Rogers 1988), including conformity, indirect biases, and guided variation.

Simulation

We used a Microsoft Excel spreadsheet to model these transmission systems. Our simulations deliberately minimize the number of rules and initial conditions that govern transmission to expose the general patterns that emerge when a few simple traits are transmitted by simple rules in a relatively small population over a small number of generations. Our assumptions are highly unrealistic; they are certainly not intended to capture any specific real-world situation. However, the simulations define a simple transmission system with enough detail to directly observe how the mechanisms of transmission affect phylogenetic reconstruction. We think that more complicated simulations would likely arrive at the same conclusions, but with much less clarity.

Each of our simulations models a population of individuals. The first generation of each population consists of a single individual, whom we refer to as the “founding father.” This founding father creates two offspring (the second generation), each of whom creates two offspring of his own (the third generation), and so on until we reach the sixth generation, in which there are thirty-two individuals. In its basic form this simulation is akin to asexual reproduction. Each individual is characterized by four traits, so in each simulation there are 124 trait transmissions: $1 + 2 + 4 + 8 + 16 = 31$ parent-to-offspring transfers down the tree, each entailing four traits ($4 \times 31 = 124$).

Each of the four traits can take one of ten different states, coded in our simulations with numbers between 0 and 9. One way to envision this is as people who wear four articles of clothing (e.g., shirt, shoes, trousers, socks), each of which comes in ten possible colors. We follow that analogy in the discussions below. Two different types of models, directed and undirected, are presented. In the undirected models, each of the ten color states is equal in value, and there is no inherent ordering to the character states (*i.e.*, the states are nominal variables). In the directed models, the ten states are ordered from 0 to 9, with 0 being the lowest and least desirable color and 9 the highest and most desirable state (*i.e.*, the traits are ordered, integer-scale variables).

During transmission, one of three things will happen. First, with a predefined frequency, referred to as the “error rate,” individuals mis-copy the character state of their parents and randomly generate a new character state. This introduces heritable variation that is tracked during sequences of transmission events through the simulation run. For consistency, the copy error rate was set at 5 percent for all runs. Second, individuals employ an inheritance system to acquire character states from other individuals in the population. Inheritance between individuals occurs at a set frequency that is defined as the “cultural transmission rate.” Third, if there is neither copy error nor cultural transmission, individuals simply inherit and keep the character state of their parents. Only one inheritance system was simulated at a time. That is, within a simulation run, individuals use one and only one mode of transmission—they

cannot use guided variation for one trait and conformist transmission for another. In reality, individuals probably employ many mechanisms at once, even for the same traits, but our aim here is to illustrate the consequences of transmission mechanism on phylogenetic reconstruction, so we keep the simulations simple.

In each simulation we know exactly how traits are transmitted and how populations are related to one another. In this study, our interest is in determining how well phylogenetic analyses are able to reconstruct an inheritance tree when traits are transmitted under the three systems of inheritance. To make this determination we used trait data from the sixth generation of each simulation run in a common cladistics package, PAUP* 4.0 (Swofford 1998). All thirty-two individuals in this generation were treated as independent terminal taxa even if they exhibited the same set of traits. We examined general trends by slowly increasing the amount of transmission that was allowed to take place under a particular inheritance system—guided variation, conformist transmission, and indirect bias—from 0 percent to 100 percent, recording the resulting inheritance tree (i.e., the tree directly created from the simulation data from every generation) and reconstructed tree (i.e., the tree produced using only data from the final generation).

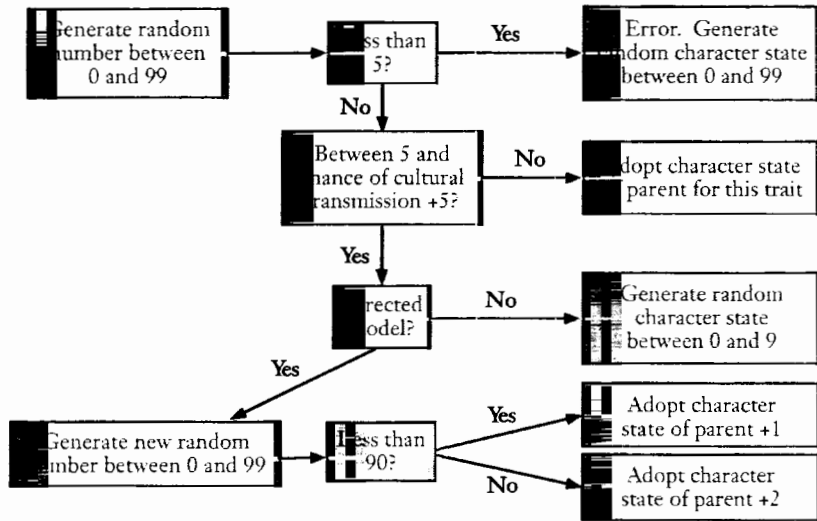
Guided variation—inheritor followed by experimentation and learning—was simulated as follows. For each of the four traits, a random number between 0 and 99 was generated. If this random number fell within the range of the error rate (5 percent), a color between 0 and 9 was randomly assigned to that trait. If the random number fell within the range of cultural transmission, the trait was transmitted by undirected or directed guided variation. In the undirected model, where no trait value was recognizably better than any other, this was simulated by generating a random number between 0 and 9 corresponding to trait color. In the directed model, the parent trait value was raised by 1 with 90 percent chance and raised by 2 with 10 percent chance, up to the maximum value of 9. Thus, the undirected model simulates the outcome of experiments where all outcomes (trait values) are equally valuable or desirable by randomly generating a new trait value, whereas the directed model simulates experiments that lead to trait improvement by slowly pushing trait states toward the maximum value of 9. Finally, if the random number fell outside the combined ranges of error and cultural transmission, the value of the parent was adopted without change (neither copy error nor guided variation occurred). Figure 11.1 presents a flowchart depicting the simulation.

The simulation of conformist transmission operates in the same way. Again, for each trait a random number was generated between 0 and 99. If this number fell within the range of copy error, a new color was randomly generated. If it fell within the range of conformist transmission, individuals surveyed the previous generation (the generation of their parent) and adopted the modal color displayed by individuals in that generation. If neither random copy error

Figure 11.1
Flowchart for Guided-Variation Model

Guided-Variation Model

For each of 4 characters (traits):



nor conformist transmission occurred, the color of the parent was adopted. Because conformist transmission is directed by definition toward modal behaviors, we did not model an undirected version. Figure 11.2 presents a flowchart for our simulation of conformist transmission.

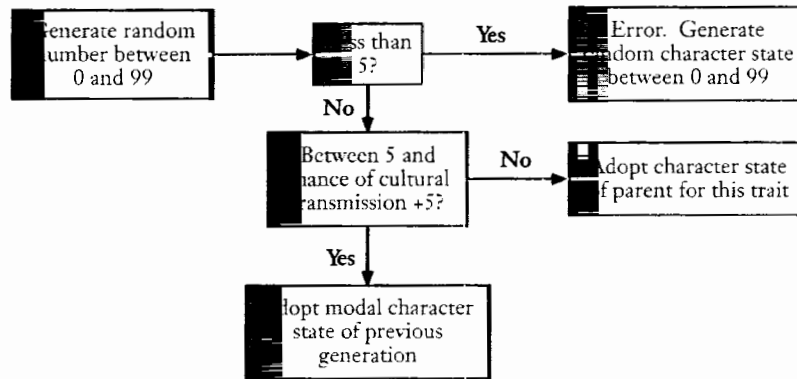
Indirectly biased transmission was modeled slightly differently. In addition to the four character states, each individual also possessed an indicator trait (varying between 1 and 4) that informed the individual which locus (trait) to inspect when selecting a social model. In the undirected model, each individual also possessed a preference trait (varying between 0 and 9) that informed him of the value of the indicator trait that was to be sought in a social model. The indicator and preference traits were transmitted vertically from parent to offspring with a 5 percent copy error. Transmission copy error was modeled by generating a random number between 1 and 4 for a miscopied indicator trait and between 0 and 9 for a miscopied preference trait.

A random number was generated for each individual to determine whether all four traits would be obtained as a package by indirect bias. If so, in the undirected model individuals consulted their indicator trait, surveyed the previous generation for that specific trait, and if any individual in that generation displayed a value within 1 of their own preference trait, they copied all four traits of that individual perfectly, that is, copied their entire color scheme.

Figure 11.2
Flowchart for Conformist-Transmission Model

Conformist-Transmission Model

For each of 4 characters (traits):

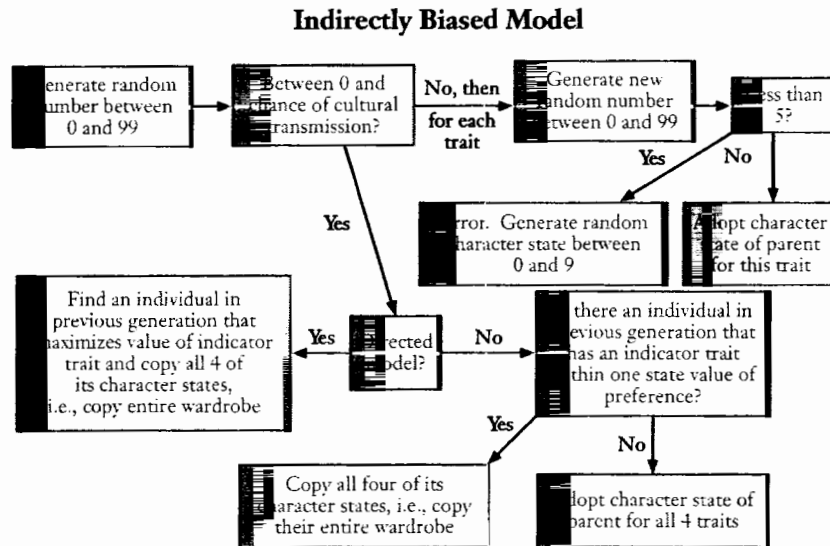


If no individual in the preferred range was present, individuals simply adopted the full set of parental values. In the directed model, on the other hand, individuals consulted their indicator trait, surveyed the previous generation for that trait, selected an individual with the maximum value of that trait (at random, when two or more individuals displayed the maximum value), and copied all four traits of that individual perfectly. Figure 11.3 presents a flowchart for how we simulated indirectly biased transmission.

Results

We were interested in what happens to the pattern of inheritance when the amounts of guided variation, conformist transmission, and indirectly biased transmission are increased. More specifically, we wanted to know how much phylogenetic information would be retained in the terminal taxa as this occurred. Note that the data structure was constant across all simulations: thirty-two individuals (terminal taxa), each displaying four traits that can take any one of ten different states. The similarity in data structure allows us to compare results from the three transmission systems. The stochasticity built into all our simulations by the use of random numbers in decision making will generate variability. This obviously obscures patterns of central tendency but is highly informative about the process of transmission and the feasibility of constructing phylogenetic trees under the simulated conditions. Accordingly, in the discussion below we focus on both central tendency and variation. We report and compare two statistics reported by PAUP* across the different simulations—tree length (TL) and consistency index (CI)—which are strongly and

Figure 11.3
Flowchart for Indirectly Biased-Transmission Model




inversely correlated. Smaller TLs correspond to higher CIs. Both statistics help describe the amount or strength of the phylogenetic signal within a particular simulation run.

Guided Variation

The amount of guided variation (t) was increased from 0 percent ($y = 0$) to 95 percent ($y = 0.95$) in small increments (recall that the remaining 5 percent corresponds to copying error). Figure 11.4 shows the CI for the simulations for the undirected model. As the figure shows, when individuals never use guided variation ($t = 0$) they produce very consistent and simple trees with strong phylogenetic signals. Only random errors contribute heritable variability to the tree. At the other extreme, when individuals always use guided variation ($t = 0.95$), which means they always experiment with new colors (or make copy errors, which amounts to the same thing), they produce completely random trait sets with no phylogenetic signal. At this extreme ($t = 0.95$), CI values hover near 0.30, and TL approaches 120, indicating very long trees. However, experimentation removes phylogenetic signal long before guided variation becomes this common.

This result was expected. As people experiment more and more without direction (randomly), they produce more and more variation, which becomes harder and harder to capture in a short, simple tree. This means that as undi-

Figure 11.4
Consistency-Index Values with Increasing Strength of Undirected Guided Variation

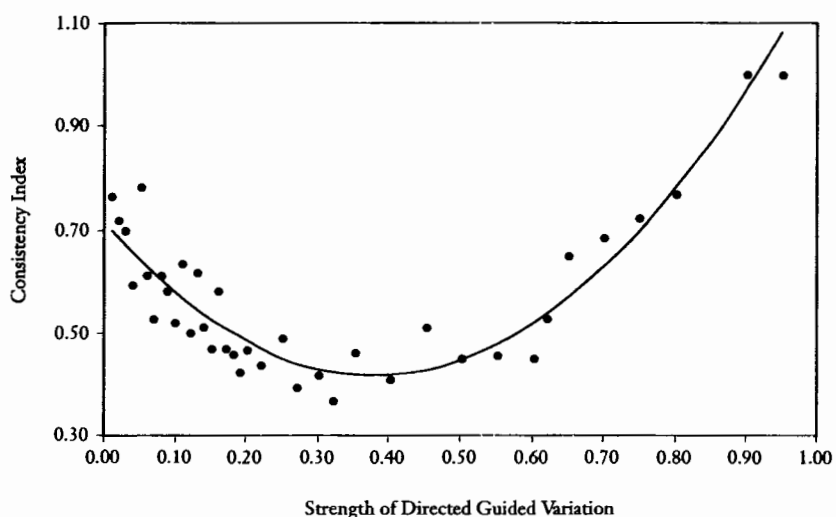


rected guided variation becomes more important, cladistic analysis becomes less and less useful in reconstructing inheritance systems, simply because less and less information is being transmitted across generations. The inheritance “signal” is quickly lost due to variability from experimentation.

A different result obtains when we make individual experimentation directional. Figure 11.5 shows that as in the undirected version, when the strength of directed guided variation is low, CI values are initially high, and as the strength of guided variation (and experimentation) increases, CI values fall, indicating weakening of the phylogenetic signal. However, at the point where the undirected CI values hit their minimum (roughly at 35 percent guided variation), the directed CI values begin to rise again. Although more and more people are experimenting, they increasingly arrive at the same conclusion, thereby decreasing overall variation. In other words, all individuals begin to adopt the same character traits, much as they did when there was little experimentation and mostly vertical transmission.

Using cladistics, we would interpret this as a better phylogenetic signal because of the higher CI and lower TL. Note, however, that this “better” signal is not really the result of vertical transmission of heritable information. It is really individuals separately arriving at the same optimal trait combination through experimentation—the convergent evolution of analogous traits as a result of the environment. In a variable environment in which different traits are optimal in different places, strong guided variation will destroy phylogenetic signals, just as strong natural selection destroys phylogenetic signals in

Figure 11.5
Consistency-Index Values with Increasing Strength of Directed
Guided Variation



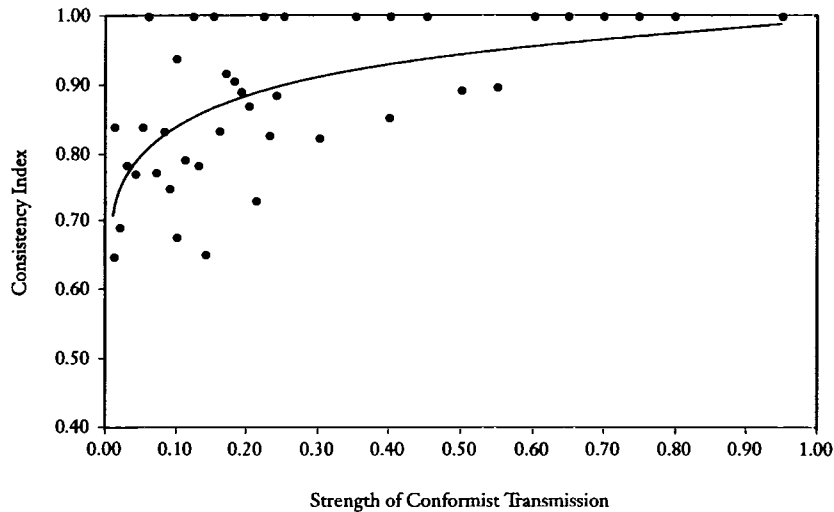
genetic data. When a strong attractor erases variation, phylogenetic information is destroyed in any system. This observation points out the weakness of using CI (and TL) to estimate the “strength” of the phylogenetic signal in a data set. Other statistics are much better in this regard (see McElreath 1997), which we aim to employ in future work.

Conformist Transmission

We obtain still a different result when individuals acquire traits by conformist transmission. In contrast to guided variation, where CI values decrease (at least initially) as the force of transmission increases, with conformist transmission CI values increase as the force of transmission increases. CI values also vary much more around this basic trajectory. When conformist transmission is low, CI values are intermediate, between 0.65 and 0.75. When conformist transmission exceeds 50 percent, CI values consistently exceed 0.85. These trends are shown in figure 11.6.

This result was not unexpected, in that conformist transmission tends to decrease variation, at least within communities (Boyd and Richerson 1985; Henrich and Boyd 1998). As more and more individuals use conformist transmission to obtain traits, these individuals look more and more alike, which greatly increases the overall phylogenetic signal displayed by a population. The increased variability around this mean tendency is the result of the potential for novel recombination of traits when some are inherited vertically and

Figure 11.6
Consistency-Index Values with Increasing Strength
of Conformist Transmission.

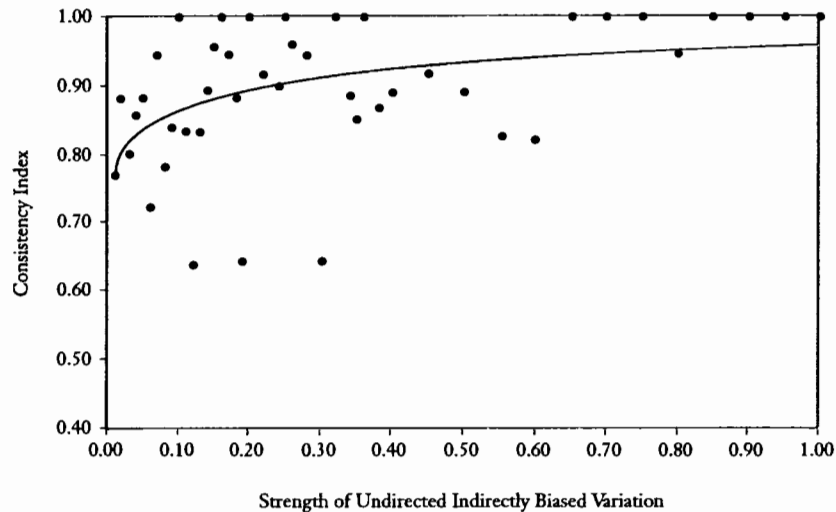


others are acquired by conformist transmission, which results in the creation of new taxa (or color schemes in our simulations). As the strength of conformist transmission increases, the potential for novel recombination decreases, and the variability around the regression line decreases (right side of figure 11.6). Overall, in comparison to guided variation, conformist transmission tends to preserve a phylogenetic signal within terminal taxa. It follows, then, that traits acquired by conformist transmission are more suited to cladistic analysis than traits acquired through guided variation.

Indirect Bias

The undirected and directed simulations of indirectly biased transmission behave much like those of conformist transmission, but there is more variation around central tendencies. In both models as the strength of indirectly biased transmission increases, CI values increase, from around 0.7, when there is no indirect bias, to 1.0 as the frequency of indirect bias approaches 95 percent. This result is shown in figures 11.7 and 11.8. As with conformist transmission, CI values increase directly with the amount of indirect bias because indirect bias is in general a variation-reducing process. Because indirect bias does not generate any novel trait-value combinations, unless copy error is built into the system, TL will decrease, and CI values will increase monotonically as indirect bias increases.

Figure 11.7
Consistency-Index Values with Increasing Strength of Undirected Indirect Bias



The relatively high variation around the regression lines in figures 11.7 and 11.8 is a result of these indicator- and preference-trait copy errors. The errors generate subsets of individuals whose color schemes are identical but whose indicator and preference traits are different because some were copied incorrectly. These random copy errors in the indicator and preference traits force sets of individuals who should have chosen the same social models to copy different social models. Whenever random copy errors generate significant within-generation variation in the indicator and preference traits, phylogenetic trees will become longer and more complicated. On the other hand, if we remove the 5 percent error rate built into the transmission of the indicator and preference traits, so that all individuals have the same value for these traits, indirectly biased transmission forces everyone to choose the same social model. If we use PAUP* to generate trees using data from these simulations, the amount of variation around the regression line is much lower, as expected. Overall, our studies suggest that indirectly biased transmission tends to preserve a phylogenetic signal as well or better than conformist transmission and better than guided variation.

Summary of Results

These analyses lead us to several observations about the possibilities for cladistic analysis in archaeology. First and most importantly, some transmission systems preserve phylogenetic signatures better than others. Phylogenetic signals are most likely to be retained in collections of archaeological

traits that were culturally transmitted in packages by indirect bias, for example. Conversely, phylogenetic signatures are the least likely to be retained in trait complexes whose individual components were acquired one at a time (not as packages) by guided variation, particularly in instances characterized by experimentation over many generations. The phylogenetic signal will be stronger, and retained longer, if the individually transmitted components of such complexes were acquired by conformist transmission, but the signal will be weaker and shorter-lived when trait packages are obtained by indirectly biased transmission. Apart from this, the utility of cladistic analysis will decrease for traits that are subject to external (noncultural) evolutionary forces that lead to convergence, especially where guided variation is strong and to a lesser degree where indirectly biased transmission is strong. In the following section, we compare these generalizations against the results of phylogenetic analyses of projectile points from two areas—Owens Valley and Monitor Valley—in the western Great Basin.

Application to Great Basin Projectile Points

Approximately half of the described projectile points in Owens Valley and Monitor Valley belong to either the Rosegate or Elko series. Rosegate represents the introduction and first phase of bow and arrow use, from roughly A.D. 600 to A.D. 1300. Rosegate replaced dart points belonging to the Elko series, which are similar in shape but larger and older than Rosegate; in Owens and Monitor Valleys Elko-series points date from roughly 1500 B.C. to A.D. 600.

Previous analyses (Bettinger and Eerkens 1997, 1999) of these two kinds of points led us to argue that Owens Valley Rosegate points were the product of a transmission system that depended heavily on experimentation—guided variation or something similar—that governed the introduction of bow-and-arrow technology in that area. Conversely, we argued that Monitor Valley Rosegate points were part of a transmission system characterized by conformist or indirectly biased transmission. Our simulations imply that if these things are true, Monitor Valley points should display a stronger phylogenetic signal than Owens Valley points.

To facilitate comparison, we quantified the point assemblages in much the same way as our simulated data. We sampled thirty-two points more or less at random from each area, giving preference to specimens that were independently dated by obsidian hydration in Owens Valley (where obsidian is common) and by stratigraphy in Monitor Valley (where obsidian is rare). For each point, we recorded four attributes that previous research suggested were governed by more or less neutral selective forces: length/width ratio, basal width, proximal shoulder angle, and thickness. Each attribute was rescaled into ten states (0 to 9) by subtracting the minimum value across the entire sample for each attribute from its observed value, dividing by the range for that type in that region, multiplying by 9, and dropping the values to the right of the

decimal point. This mathematical manipulation transformed each projectile-point measurement from a continuous number (in centimeters or degrees) to an integer that varies between 0 for the lowest observed value and 9 for the largest.

The PAUP* results are presented in table 11.1 (rows 1 and 2). The results of our simulations of completely random and perfectly transmitted data are listed for comparison (rows 3 and 4). As can be seen, in both regions CI values are only slightly higher, and TL values only slightly lower, than the results obtained from random data, suggesting a negligibly small phylogenetic signal. Because our previous work suggested that guided variation dominated the transmission of projectile-point technology in Owens Valley, we expected this type of result for the Owens Valley points. That work, however, suggested that point transmission in Monitor Valley was dominated by conformist or indirectly biased transmission, which should have produced a stronger phylogenetic signal, but did not.

Further, although there was strict consensus of trees in all cases, these reconstructed trees do not predict point chronology. Hydration readings (Owens Valley) and stratigraphic positions (Monitor Valley) do not segregate by tree branch. Within these two regions at least, point morphology does not appear to vary over time in a way that it should under cultural transmission. The lack of tree-time correspondence is not an artifact of our continuous-scale to ordinal-scale (between 0 and 9) data transformation. For example, rescaling the data in units of standard deviation (so that each character state expresses the number of standard deviations between an individual observation and its mean) produces a narrower range of character states, between 0 and 5 in most cases. This results in a predictable increase in the CI (table 11.1, rows 5 and 6) as well as in the production of nearly identical trees. In short, changing the way data are quantified does not change how PAUP* reconstructs trees for these data sets.

Table 11.1
PAUP* Results on Rosegate Points from Owens and Monitor Valleys, California

Data Set	CI	TL
1 Owens Valley Rosegate	0.33	109
2 Monitor Valley Rosegate	0.35	104
3 Random Data	0.27	132
4 Perfect Transmission	1.00	1
5 Owens Valley Rosegate – (rescaled)	0.59	34
6 Monitor Valley Rosegate – (rescaled)	0.63	35
7 Owens Valley Elko and Rosegate	0.36	99
8 Monitor Valley Elko	0.38	97

Attempts to detect a phylogenetic signal in Elko points were similarly unproductive (see table 11.1, rows 7 and 8). For the Owens Valley analysis, we were able to assemble only twenty-two points that were sufficiently complete for phylogenetic analyses. To obtain the requisite thirty-two terminal taxa, we augmented the sample with 10 Owens Valley Rosegate points. With one exception, PAUP* did place the Elko and Rosegate points on different branches, but the PAUP* subgroupings of Elko points again failed to predict hydration readings, which were inconsistently distributed across the subgroups. This suggests that PAUP* is not detecting a temporal component in these data.

For Monitor Valley, Elko points from Gatecliff Shelter were also subjected to phylogenetic analysis in PAUP*. Low CI and high TL suggest a low phylogenetic signal here as well. PAUP* does divide these points into two major branches, however, which visual inspection suggests is primarily a function of thickness, (thinner vs. thicker points). Once again, this split fails to predict the stratigraphic position of Elko points within the Gatecliff deposit. Overall, points from different levels are indiscriminately mixed between all branches and subbranches, including the major thin-thick division.

Does this mean that information about projectile points was not transmitted and/or subject to evolutionary forces in prehistory in these two areas? No. One obvious possibility is that Rosegate and Elko points assumed their optimal forms almost immediately following their appearance and then drifted very tightly and nondirectionally around these optimal forms for the balance of the time they remained in use. In such cases PAUP* will probably not highlight changes that are related to time.

Conclusions

Our simulations demonstrate that all transmission systems do not act on variation in ways that are amenable to analysis using cladistic methods. Some systems, such as indirectly biased transmission, are variance reducing and lead to detectable signal preservation over time; others, such as guided variation, increase variation and cause loss of a detectable signal relatively quickly. The system in place prehistorically dramatically affects our ability to tease apart cultural inheritance with phylogenetic methods.

We did not detect phylogenetic signals in our projectile-point data, despite previous work suggesting we should have been able to (Bettinger and Eerkens 1997, 1999). Our lack of success may be a result of the scale of analysis. Most previous studies of projectile-point phylogeny have dealt with point types (Lyman and O'Brien 2000; O'Brien et al. 2001; O'Brien and Lyman 2002a), with distinct types being treated as terminal taxa that are assembled into a phylogenetic tree. In our analyses individual projectile points, not the types, are the terminal taxa, and it may be that the amount of inter-individual variability at any given time overwhelms what is being transmitted during the time a particular point style is in use. Thus, in practice, cladistic methods may

be unable to highlight and segregate heritable variation at this more fine-scale level. This is unfortunate because on a theoretical level the technique is well suited to track the evolution of small heritable changes at such a scale.

This is not a failing of cladistic analysis *per se*. PAUP* detected heritable variability in our simulations, which was at the scale of the individual, though unrealistically simplified. The devil may be in the details of the real-life example. It may simply be that the majority of archaeological data at the individual level are too subject to idiosyncratic variation to be suitable for cladistic analysis. In other words, cladistic analyses in archaeology might be better suited to broad-scale changes rather than microscale ones. Unfortunately, while the modes of transmission at the individual level are relatively straightforward and the subject of substantial theoretical and empirical work (e.g., Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981; Henrich 2001; Henrich and Boyd 2001; McElreath 1997), an understanding of the processes that produce broad-scale changes are still a mystery. It seems to us, then, that there is still a substantial gap in our ability to link the trees produced by cladistic analyses of such broad-scale data (types) to prehistoric behavior.

Part of the problem may relate to the assumption of parsimony—that the best trees are those that minimize the number of evolutionary changes. In the biological sciences parsimony-based methods work well because heritable genetic changes are rare. However, as previously discussed by one of us (McElreath 1997), there is no good reason to assume this is true in the case of the transmission of human material culture. Indeed, the rate of change witnessed in the archaeological record over the last 10,000 years suggests that human material culture is characterized by an extremely high rate of change. The penchant for human learning suggests that sudden convergence may be much more common than in biological evolution. Perhaps part of our failing to find a phylogenetic signal in our projectile-point sample stems from the fact that specimens were subject to a high rate of evolutionary change, which included instances of convergence. Thus, the phylogenetic methods we employed may have failed because they naively assumed evolution to be rare.

