Journal of Archaeological Science 37 (2010) 1348–1358

Contents lists available at ScienceDirect

Journal of Archaeological Science

journal homepage: http://www.elsevier.com/locate/jas

Ceramic diversity, random copying, and tests for selectivity in ceramic production

James Steele*, Claudia Glatz, Anne Kandler

AHRC Centre for the Evolution of Cultural Diversity, Institute of Archaeology, University College London, 31–34 Gordon Square, London WC1H 0PY, UK

ARTICLE INFO

Article history:
Received 27 September 2009
Received in revised form
28 December 2009
Accepted 28 December 2009

Keywords:
Cultural transmission
Neutral model
Drift
Selection
Archaeological ceramics

ABSTRACT

This paper examines the evidence for selective reproduction of ceramic types based on their evolving frequency distribution, and on correlations of the abundance of types with underlying functional characteristics. How can we tell if a variant has been preferentially selected or discriminated against, and become more or less frequent than would be expected under pure drift (i.e. random copying, with the compounding over time of proportionate sampling error)? We examined a database of Hittite ceramic bowl types from two successive Phases of occupation of the Upper City of Boğazköy–Hattusa, capital of the Hittite empire and the largest Bronze Age settlement in Turkey. We applied two tests used in genetics to assess departures from neutrality, to assess selectivity in rates of reproduction of ceramic bowl types (the Ewens-Watterson and Slatkin's Exact tests). We also examined the effects of ceramic fabric and vessel dimensions on changes in vessel abundance between the two Phases, using regression analysis. We found that while the frequency distribution of rim sherds did not in itself enable us to reject the null hypothesis of random copying, closer examination of the characteristics of these types enabled us to recognize latent dimensions of functional variability (including ware type and bowl diameter) that had demonstrably been the subject of selective decision-making by the potters. The present case study suggests that we should be wary of applying the neutral model from genetics uncritically in archaeology, because it is much harder to prove that the cultural traits whose frequencies are being modelled are genuinely functionally equivalent (as that model requires).

1. Introduction

This paper examines the evidence for selective reproduction of ceramic types based on their evolving frequency distribution, and on correlations of the abundance of types with underlying functional characteristics. It is intended as a contribution to evolutionary archaeology.

It is a truism that cultural traditions and successful innovations are socially transmitted both between and within generations (respectively, by vertical or oblique and by horizontal transmission routes; Cavalli-Sforza and Feldman, 1981), with learners applying heuristics or rules of thumb in choosing when to engage in independent trial-and-error learning, and in selecting which models to copy when this is the preferred strategy (transmission biases; Boyd and Richerson, 1985). Boyd and Richerson (1985) distinguish several different decision mechanisms affecting the adoption of new cultural traits, each of which is determined by an underlying expectation that people seek to copy the most advantageous variants.

For quantitative hypothesis testing, we also need a random copying model as a null hypothesis, if we want to recognize the outcomes of non-random (selective) copying. The neutral model of cultural diversity in finite populations in the absence of selection is one in which new traits appear by a process akin to random mutation, and existing traits are randomly chosen to be copied so that their frequencies change only because of chance sampling effects. Numerous recent archaeological studies have examined the stylistic choices made by potters and their clients from such a theoretical perspective. Following Neiman's (1995) work on cultural transmission of selectively neutral stylistic elements in Woodland ceramics, which introduced neutral theory from genetics and explored evidence for changes in the scale of interaction and in the number of potters whose vessels were being circulated in a particular study region, more recent archaeological applications of cultural transmission theory have concentrated on the influence of social norms on the rates of transmission of particular cultural variants (Shennan and Wilkinson, 2001; Kohler et al., 2004; Eerkens and Lipo, 2005; Schauer, 2008).

We will not present a detailed analysis of the neutral model here, since this model has already been expounded for archaeologists at greater length by Neiman (1995) and by Shennan and Wilkinson (2001). Briefly, Eq. (1) summarizes the expected cultural

---

* Corresponding author.
E-mail address: j.steele@ucl.ac.uk (J. Steele).

0305-4403/$ – see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jas.2009.12.039

homogeneity F within a population at a time t under this model, as a function of the homogeneity in the previous time step (F_{t-1}), the number of individuals producing successive generations of cultural variants (N_e, with a generation meaning an independent episode of cultural production at which variation may be introduced), and the mutation rate μ (Neiman 1995, p. 14, Eq. (2)):

$$F_t = \left[\frac{1}{N_e} + \left(1 - \frac{1}{N_e}\right)F_{t-1}\right](1-\mu)^2 \tag{1}$$

The rate of mutation μ (a constant between 0 and 1) is the probability that an individual copying episode introduces a new variant rather than faithfully copying an existing one. The likelihood of an individual making a faithful copy of an existing model is $(1-\mu)$. The probability that any two randomly sampled individual copying episodes at time t have faithfully copied the same exemplar from the previous copying round t_{-1} is $1/N_e$, while $1 - 1/N_e$ represents the probability that they have not. Equilibrium homogeneity $\widehat{F}_{eq}$ is reached when the rate of mutation corresponds to that of random variant loss ($F_t = F_{t-1}$), since loss and mutation are opposing evolutionary forces, the former reducing diversity, the latter increasing it (Shennan and Wilkinson, 2001, p. 583):

$$\widehat{F}_{eq} = \frac{(1-\mu)^2}{\left[N_e - (N_e - 1)(1-\mu)^2\right]} \tag{2}$$

For transmission systems in which the rate of mutation μ is very low, this can be approximated as:

$$\widehat{F}_{eq} \approx \frac{1}{2N_e\mu + 1} \tag{3}$$

which indicates that the equilibrium homogeneity of neutral variants within a population is inversely proportional to $2N_e\mu$.

We emphasize that to apply this model to cultural production we must assume one of two simplifications of reality. On the one hand, we might assume that each human producer always produces identical artifacts (in which case a generation is a human biological generation, a copying episode is an episode of vertical or oblique transmission of craft skills between a skilled and an unskilled individual, and N_e is the number of active teachers). On the other hand, we might assume that a single human producer can produce many successive cultural generations of variants, and that at each successive copying episode an individual is no more likely to copy a variant in his or her own repertoire than an extant variant found elsewhere in the population (in which case a generation is an independent bout of craft production, a copying episode may involve inter-individual transmission or may simply involve the same individual reproducing artifacts in a new bout of production, and N_e is the number of active producers). In this paper, we make the second of these assumptions; however, we note that future developments of this model could usefully address the contrasting empirical rates of mutation and of sampling of alternative variants in within-individual versus between-individual copying episodes.

Fig. 1 shows typical output from a simulation of evolving ceramic assemblage diversity under drift. In this case, we start with a uniform distribution of 10 variants, each making up 10% of the assemblage, and then propagate them over a hundred copying cycles using the random copying rule, with new variants introduced with a likelihood $\mu = 0.01$. As this example shows, it is obvious that under drift most variants dwindle towards extinction while a few others increase in abundance, and that the diversity profile evolves towards one in which the empirical homogeneity statistic $F = \sum_{i=1}^{k} p_i^2$ fluctuates around the value expected at mutation–drift equilibrium ($\widehat{F}_{eq} \approx 1/2N_e\mu + 1$).

Archaeological use of this model has focused on two issues: the inference of changes over time in the values of N_e or μ (respectively the size of the population reproducing the variants, and the mutation rate); and inference of non-neutrality (i.e., the existence of other processes than drift influencing variant frequencies) at a given location and time step. Where there is assumed to have been no departure from neutrality, variation in the size of the population reproducing the variants and/or in the mutation rate has been inferred from the frequency distribution of variants. Neiman (1995) found a trend across time for an increase in the value of an index of diversity, θ, in inventories of Woodland ceramic stylistic attributes, which he interpreted as reflecting an increase in the scale of the social network and in the number of pottery producers. This index of diversity can be calculated empirically as the coefficient $\theta = (1/F) - 1$, and can also be estimated as a maximum likelihood for an observed sample size and number of variants using Ewens' sampling formula (Ewens, 1972). Instead of calculating θ, Bentley et al. (2004) have used the magnitude of the power law slope coefficient α to make similar inferences about varying innovation rates in other cultural contexts where the population size is known. They draw on another technique for estimating how well an empirical assemblage matches the predictions of the neutral model by plotting the frequency distribution of variants, with the expectation under neutrality of a linear fit on a log–log plot of trait frequency against the likelihood of a trait appearing at that frequency (which means that the underlying distribution can be fitted as a first approximation by a power-law curve, cf. Hahn and Bentley, 2003; Bentley et al., 2004; Herzog et al., 2004). This derives from Kimura and Crow's (1964) demonstration that the equilibrium homogeneity $\widehat{F}_{eq} = 1/2N_e\mu + 1$ corresponds to a frequency distribution of neutral variants with expected number of variants at frequency ν, $P(\nu) = \theta\nu^{-1}(1-\nu)^{\theta-1}$, which approximately corresponds to a power law distribution $P(\nu) \sim \nu^{-\alpha}$ for small mutation rates μ (with α varying as a function of θ; Hahn and Bentley, 2003; Bentley et al., 2004).

How can we tell if a variant has been preferentially selected or discriminated against, and become more or less frequent than would be expected under pure drift? One approach, popular in both population genetics and in ecology during the early phases of development of neutral theory, has been to test for departures from the theoretical frequency distribution of variants that is predicted by the neutral model. This is the approach explicitly used by Lipo (2001), and which is implicit in the studies by Shennan and Wlikinson (2001) and by Kohler et al. (2004). In genetics, two such tests use Ewens' (1972) sampling formula. This formula predicts the properties of the underlying distribution for a neutral process, where an empirical sample of n cases has been observed with k variants. The Ewens–Watterson test of homozygosity (Watterson, 1977, 1978) evaluates the empirical homogeneity statistic $F = \sum_{i=1}^{k} p_i^2$ against the value expected under neutrality, and Watterson (1978) provides a table of values for F for different values of n (up to 500 observations) and k (up to 10 variants) and for different significance thresholds. This is similar to the approach of Shennan and Wilkinson (2001) except that the latter did not estimate statistical significance. The Slatkin Exact Test evaluates the empirical frequency distribution against all possible configurations for a given n and k under the Ewens sampling distribution, and is a more general test which makes fewer assumptions about which aspects of distribution shape might indicate selection (Slatkin, 1994, 1996). Slatkin (1997) has helpfully archived online his source code for a Monte Carlo simulation program that conducts both tests (Ewens–Watterson and Slatkin Exact). The outputs are tail probabilities, which indicate the position of an empirical set of data in a probability distribution of possible sets derived from the Ewens