Author for correspondence:
Anne Kandler
e-mail: anne.kandler.1@city.ac.uk

THE ROYAL SOCIETY
PUBLISHING

A generative inference framework for analysing patterns of cultural change in sparse population data with evidence for fashion trends in LBK culture

Anne Kandler[1,2] and Stephen Shennan[3]

[1]Department of Mathematics, City University London, Northampton Square, London EC1V 0HB, UK
[2]Santa Fe Institute, 1599 Hyde Park Road, Santa Fe, NM 87501, USA
[3]Institute of Archaeology, University College London, 31–34 Gordon Square, London WC1H 0PY, UK

Cultural change can be quantified by temporal changes in frequency of different cultural artefacts and it is a central question to identify what underlying cultural transmission processes could have caused the observed frequency changes. Observed changes, however, often describe the dynamics in samples of the population of artefacts, whereas transmission processes act on the whole population. Here we develop a modelling framework aimed at addressing this inference problem. To do so, we firstly generate population structures from which the observed sample could have been drawn randomly and then determine theoretical samples at a later time t_2 produced under the assumption that changes in frequencies are caused by a specific transmission process. Thereby we also account for the potential effect of time-averaging processes in the generation of the observed sample. Subsequent statistical comparisons (e.g. using Bayesian inference) of the theoretical and observed samples at t_2 can establish which processes could have produced the observed frequency data. In this way, we infer underlying transmission processes directly from available data without any equilibrium assumption. We apply this framework to a dataset describing pottery from settlements of some of the first farmers in Europe (the LBK culture) and conclude that the observed frequency dynamic of different types of decorated pottery is consistent with age-dependent selection, a preference for 'young' pottery types which is potentially indicative of fashion trends.

1. Introduction

How does cultural change happen? What are the cultural transmission processes that underlie observed episodes of cultural change? Those and similar questions have been the subject of intense research in the last few decades. Naturally, fine-grained individual-level data detailing who learns from whom would be most suited to answer these questions empirically. However, outside experimental conditions, large individual datasets of this kind are difficult to obtain, especially in pre-modern contexts. Studies such as [1] that used historical records of the board game Go, or [2] that used contemporary data from Fiji to determine the relative importance of different cultural transmission processes are remarkable exceptions. Much available data in archaeology, biological anthropology, animal behaviour or psychology, however, describe the population-level outcomes of the transmission processes. The data are often in the form of frequency distributions of a number of different variants of a cultural trait in the population at a certain point in time, or in a time-series describing the dynamics of the frequency change of cultural variants over time. Furthermore, the data usually comprises sparse samples from the whole population.

But how well can population-level data answer questions about the underlying cultural transmission processes that produced observed episodes of cultural change? It has been suggested that the shape of the adoption curve

(detailing the fraction of the population that has adopted a specific cultural variant) might be indicative of the underlying transmission processes (e.g. [3–6]). However, the diagnostic power of these approaches, especially when applied to sparse data, has been proved to be limited (e.g. [7–10]).

In archaeological applications, a lot of emphasis has been put on testing the consistency between observed empirical patterns and the hypothesis of neutral evolution. It has been argued that neutral evolution provides the most parsimonious mechanism to account for the common observation that new artefact attributes, such as decoration patterns, appear, rise in popularity then decline in frequency and go extinct [11]. To determine whether the composition of an observed archaeological assemblage is consistent, or not, with unbiased social learning (the mechanism of neutral cultural evolution), empirical diversity statistics of the observed assemblage have been compared with their theoretical expectations at innovation–drift equilibrium (e.g. [11–14]). Further, Bentley et al. (e.g. [15–17]) have used the full variant frequency distribution of a dataset and compared it with the right-skewed power-law distribution expected to arise under unbiased social learning at innovation–drift equilibrium [18]. However, it has been pointed out that the power law as limiting distribution can be generated by a large number of underlying mechanisms [19,20]. Nevertheless, Premo [21] showed that the power-law approach produces more robust results when time-averaging processes are considered than the diversity approach. To examine the temporal dynamic of neutral evolution, Bentley et al. [22] and Acerbi & Bentley [23] used the rate of turnover within the n most popular variants in a population to test for departures from the neutral evolution hypothesis within a model comparison framework. Potentially, it may be difficult to apply this approach to sparse archaeological data where the total number of observed cultural variants is low (such that a 'top list' of size sufficient to accurately calculate turnover rates cannot be defined). Crucially, the approaches described above are based on the assumption that the cultural system considered is at innovation–drift equilibrium. This means the system is in a state where the loss of cultural diversity through drift and the introduction of diversity through innovations are balanced. For a more detailed review of this literature, see [24].

In this paper, we contribute to this discussion by developing a generative inference framework, going beyond our previous work on this topic [25] to address sampling and other issues. Generative inference frameworks have already been successfully applied to genetic applications (e.g. [26,27]). Importantly, the framework will not rely on the assumption that the cultural system is at equilibrium. It has been shown that the inference of underlying transmission processes should, if possible, be based on the temporal dynamic of cultural change and not on the cultural composition of the population at a single point in time [28,29]. Therefore, the aim of our framework is to evaluate whether frequency changes of different types of cultural variants observed in two archaeological assemblages between two time points t_1 and t_2 (with $t_2 > t_1$) are consistent with a specific hypothesis about the underlying mechanisms of cultural transmission [25]. The approach is similar in spirit to the model selection framework developed by McElreath et al. [30], but we do not restrict the choice of the modelling framework by the availability of an analytical representation of the

corresponding likelihood function. In particular, we use the frequency distribution of the observed sample, denoted by $\bar{S}(t_1)$ (throughout the paper, the 'overline' symbol $\bar{}$ stands for observed quantities), of cultural variants at time t_1 as the starting point of our analysis. But importantly a sample describes only a fraction of the whole population of cultural variants. Cultural transmission processes, however, act on the whole population and therefore any inference framework must be based on the cultural composition of the population rather than the observable sample. To incorporate this fact, we firstly generate population structures from which the observed sample could have been drawn at random. Then we determine the statistical properties of theoretical samples, denoted by $S(t_2)$, produced under the assumption that changes in frequencies seen at time t_2 are caused by a specific transmission mechanism. By this means, we also account for potential time-averaging processes, where time-averaging refers to the phenomenon where cultural remains deposited at different times are preserved together and cannot be distinguished from one another, the rule rather than the exception in archaeology (e.g. [21]). Subsequent comparisons of the properties of the theoretical sample with properties of the observed sample $\bar{S}(t_2)$ will allow for conclusions about the consistency between the transmission hypothesis and the observed changes in frequency of the cultural variants in the time period $[t_1, t_2]$. In particular, we are interested in exploring the consistency with neutral evolution, frequency- and age-dependent selection processes. For that we use approximate Bayesian computation (ABC) to establish which transmission processes could have produced the observed frequency data and therefore infer the underlying processes directly from the available data without any equilibrium assumption. The outcomes of the framework developed are posterior distributions for the model parameters (characterizing the specific transmission process), indicating the range of the transmission processes that could have produced the observed changes in frequency.

We stress that it is not known a priori whether observable population-level patterns, such as observed changes in frequency of cultural artefacts, carry a strong signature of the underlying transmission process and therefore whether reliable inference is possible on the base of the observed data [3]. However, we envision that a generative inference framework might shed light on this problem. Knowing when inference should not be based on population-level data is equally important, as in these situations other lines of evidence have to be considered.

2. The inference framework

The inference framework is designed to investigate whether observed changes in frequencies of cultural variants between two time points t_1 and t_2 are consistent with a specific hypothesis of cultural transmission. It consists of two main parts. Firstly, we develop a generative model which, based on the observed sample $\bar{S}(t_1)$ at time t_1, generates theoretical samples, $S(t_2)$, at time t_2 conditioned on a specific cultural transmission process. Secondly, we use Bayesian inference techniques to establish whether the cultural transmission process considered can produce samples $S(t_2)$ that are similar to the observed sample $\bar{S}(t_2)$ at time t_2. In this section, we discuss both parts of the framework in detail.

Throughout the paper we use the following notation. A sample of cultural variants is characterized by the number k of different types of variants, the total number n_i of variants of type i, $i = 1, \ldots, k$, the sample size $n = \sum_{i=1}^{k} n_i$ and the relative frequencies $r_i = n_i/n$, $i = 1, \ldots, k$. Similarly, the corresponding population, denoted by P, is described by the following usually unknown quantities: the number K of different types of variants ($k \leq K$), the total number N_i of variants of type i, $i = 1, \ldots, K$, the population size $N = \sum_{i=1}^{K} N_i$ and the relative frequencies $R_i = N_i/N$, $i = 1, \ldots, K$. Importantly, we summarize the population structure by $[R_1, R_2, \ldots, R_k, R_{k+1}]$ with $R_{k+1} = 1 - \sum_{i=1}^{k} R_i$ or $[N_1, N_2, \ldots, N_k, N_{k+1}]$ with $N_{k+1} = N - \sum_{i=1}^{k} N_i$ and therefore consider all variant types which are not observed in the original sample as type $k + 1$. This is a simplifying assumption as variant type $k + 1$ most probably consists of several unobserved (and potentially rare) variant types. The determination of the number of unobserved variant types is not the subject of this inference approach as we aim to establish the statistical power of inference methods based on observable population-level quantities; it will be left for future research.

2.1. The generative model

As already mentioned, the observed sample $\bar{S}(t_1)$ of size $n(t_1)$ only describes a fraction of the population of cultural variants present at t_1. Cultural transmission processes, however, act on the whole population. Consequently, the first step of our model must include the determination of a population structure $P(t_1)$ that is consistent with the observed sample $\bar{S}(t_1)$ or in other words, of a population structure from which the observed sample could have been drawn by a process of random sampling. Grounding the inference framework on sample frequencies would inevitably overestimate the role of cultural drift. Based on a consistent population $P(t_1)$, we generate the population-level frequencies of the different cultural variant types conditioned on the assumption that frequency changes between the time points t_1 and t_2 are caused by the assumed evolutionary process. In this way, we generate theoretical populations $P(t_1 + \tau)$, $\tau = 1, \ldots, t_2 - t_1$ and the last modelling step consists of sampling from those populations to generate a sample $S(t_2)$ of size $n(t_2)$ ($n(t_2)$ denotes the size of the observed sample $\bar{S}(t_2)$) at time t_2.

2.1.1. Generating population structures based on a given sample

Given the observed sample $\bar{S}$, we determine consistent population structures $P = [R_1, \ldots, R_k, R_{k+1}]$, meaning population structures from which the observed sample $\bar{S}$ could have been drawn by a process of random sampling, by using a Dirichlet distribution approach. This approach has already been successfully applied in population genetics (e.g. [26,27]). It results in a distribution of consistent population P of the form

$$P(P|\bar{S}) \propto P(P)P(\bar{S}|P)$$
$$\propto R_1^{\alpha_1 - 1 + n_1} R_2^{\alpha_2 - 1 + n_2} \ldots R_k^{\alpha_k - 1 + n_k} R_{k+1}^{\alpha_{k+1} - 1}. \quad (2.1)$$

Consequently, sampling from distribution (2.1) generates population structures $P = [R_1, \ldots, R_k, R_{k+1}]$ from which the observed sample could have been drawn at random. As an example, we assume a sample consisting of seven variant
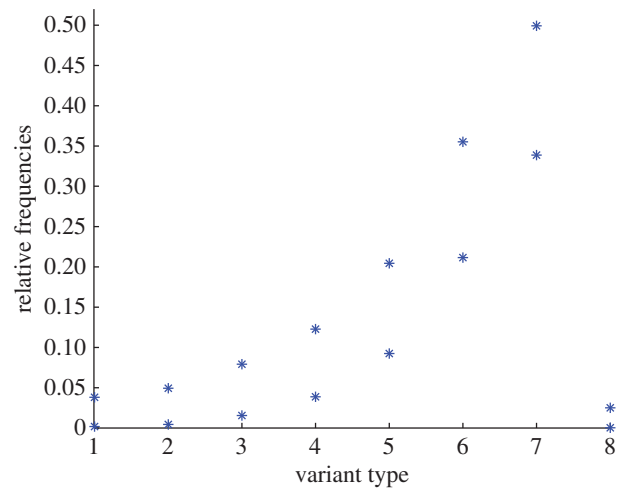

Figure 1. 95% intervals of the marginal (population-level) frequency distributions of the seven variant types which are present in the sample with the frequencies $\bar{S} = [1, 2, 5, 10, 20, 40, 60]$ and variant type eight summarizing all variants not observed in the sample with $\alpha_i = \alpha = 1$. (Online version in colour.)

types with the frequencies $\bar{S} = [1, 2, 5, 10, 20, 40, 60]$. Figure 1 shows the 95% intervals of the possible population-level frequency ranges of the seven variant types present in the sample and of variant type eight summarizing all variants not observed in the sample. As expected, those 95% intervals are relatively wide, demonstrating the importance of this modelling step.

In more detail, we assume that the total population size N is much larger than the sample size n. Now, given a population structure $P = [R_1, \ldots, R_k, R_{k+1}]$ the probability of a specific sample $\bar{S}$ with the absolute frequencies $[n_1, \ldots, n_k, 0]$ is determined by the multinomial distribution

$$P(\bar{S}|P) = \frac{n!}{\prod_{i=1}^{k} y_i!} R_1^{n_1} \ldots R_k^{n_k} R_{k+1}^{0}.$$

In this case, the conjugate prior is given by the Dirichlet distribution of the form [31]

$$P(P) \propto R_1^{\alpha_1 - 1} \ldots R_k^{\alpha_k - 1} R_{k+1}^{\alpha_{k+1} - 1},$$

with exponents $\alpha_i > 0$ and we can derive the unnormalized posterior distribution of the population P given the sample $\bar{S}$ (see equation (2.1)) [31]

$$P(P|\bar{S}) \propto P(P)P(\bar{S}|P) \propto R_1^{\alpha_1 - 1 + n_1} R_2^{\alpha_2 - 1 + n_2} \ldots R_k^{\alpha_k - 1 + n_k} R_{k+1}^{\alpha_{k+1} - 1}.$$

We note that this procedure does not need an estimate of the actual population size. The parameter setting $\alpha_1 = \cdots = \alpha_{k+1} = 1$ assigns equal density to any population P satisfying $\sum_{i=1}^{k+1} R_i = 1$. However, one needs to be cautious when there are many variant types with very low frequencies in the sample. In this situation, a uniform Dirichlet distribution (i.e. $\alpha_1 = \cdots = \alpha_{k+1} = 1$) tends to underestimate the fraction of unobserved variant types and the parameters α_i might have to be chosen appropriately.

2.1.2. Modelling the evolutionary process

In the following, we develop a framework that models the change in frequency of the different variant types in

the population at each time point $t_1 + \tau$, $\tau = 1, \ldots, t_2 - t_1$ (where t_1 and t_2 are measured in years and a time step has the duration of a year) conditioned on a specific underlying transmission mechanism. To do so, we determine the total population sizes $N(t_1 + \tau)$, $\tau = 1, \ldots, t_2 - t_1$ for each time point between t_1 and t_2, sample a population structure $P(t_1)$ at t_1 from equation (2.1) and assume that in each time point $t_1 + \tau$, $\tau = 1, \ldots, t_2 - t_1$, there is a fraction of the population of cultural variants that is removed from the population and subsequently replaced. While the removal process is modelled in a random fashion, the replacement process occurs according to the cultural transmission process considered. In this way, we generate theoretical population structures $P(t_1 + \tau)$ at times $t_1 + \tau$ with $\tau = 1, \ldots, t_2 - t_1$ conditioned on a specific underlying transmission mechanism, and additionally, allow for the inclusion of potential time-averaging effects (see [21] or [32] for detailed investigations of these effects). The assumption of a random removal process is a modelling choice. This is motivated by the fact that we later apply the inference framework to a dataset describing temporal frequency changes of different types of decorated ceramic vessels. The removal of certain vessels from the population might coincide with being physically broken. It seems reasonable to assume that this occurs in a random rather than in a selective manner.

More specifically, we determine the population sizes $N(t_1)$ and $N(t_2)$ at times t_1 and t_2 by

$$N(t_1) = \frac{1}{s} n(t_1) \quad \text{and} \quad N(t_2) = \frac{1}{s} n(t_2),$$

where s ($0 < s \leq 1$) describes the ratio between population size and sample size. To account for uncertainties in the estimate of this ratio, we define s as normally distributed with mean s and variance σ_s^2. This implies that on average the sample represents $\langle s \rangle \cdot 100\%$ of the population; however, the concrete realization varies around this mean ratio. The population sizes $N(t_1 + \tau)$ for time points between t_1 and t_2 are obtained by linear interpolation; in other words, they are inferred on the basis that the difference between the sample sizes at t_1 and t_2 is evenly distributed across the time steps between them.

Next, we sample a population structure $P(t_1) = [R_1(t_1), \ldots, R_k(t_1), R_{k+1}(t_1)]$ from equation (2.1) and determine the absolute frequencies $P_N(t_1) = [N_1(t_1), \ldots, N_k(t_1), N_{k+1}(t_1)]$ by multiplying R_i by the estimate of the total population size $N(t_1)$ at t_1.

To describe the evolutionary process between t_1 and t_2, we model the frequencies of the $k + 1$ variant types as a Markov process. In each time step $\tau = 1, \ldots, t_2 - t_1$, we firstly randomly choose u variants to be removed from the population. It holds

$$u = \begin{cases} \lfloor rN(t_1 + \tau) \rfloor \\ \qquad \text{if } N(t_1 + \tau) \leq N(t_1 + \tau + 1) \\ \lfloor rN(t_1 + \tau + 1) \rfloor + (N(t_1 + \tau) - N(t_1 + \tau + 1)) \\ \qquad \text{if } N(t_1 + \tau) > N(t_1 + \tau + 1) \end{cases} \quad (2.2)$$

with $\lfloor \cdot \rfloor$ denoting the integer part. The variable $0 < r \leq 1$ describes the fraction of the population being replaced based on the population size $N = \min\{N(t_1 + \tau), N(t_1 + \tau + 1)\}$ and

therefore $1/r$ defines the average lifespan of each variant. We assume that r is modelled as a normally distributed random variable with mean $\langle r \rangle$ and variance σ_r^2. Again the use of a random variable aims to capture the uncertainty and temporal variation in the fraction of the population that is replaced in each time step and therefore in the average lifespan of a cultural variant. Equation (2.2) ensures that the number of variants subsequently added to the population is not influenced by the direction of the change in population size between two time steps. Further, u_i records the total number of removed variants of type i in this time step; it holds $\sum_{i=1}^{k+1} u_i = u$. Secondly, we add $v = N(t_1 + \tau + 1) - (N(t_1 + \tau) - u)$ cultural variants to the population and ensure that the total population size $N(t_1 + \tau + 1)$ at $t_1 + \tau + 1$ is maintained. Importantly, the choice of which variant type to add is guided by the assumed cultural transmission process and the variables v_i denote the total number of added variants of type i so that it holds $\sum_{i=1}^{k+1} v_i = v$.

Summarizing, the two random processes, the removal and subsequent replacement of cultural variants, result in a one-step transition probability for transforming the population $P_N(t_1 + \tau) = [N_1(t_1 + \tau), \ldots, N_{k+1}(t_1 + \tau)]$ of cultural variants at time $t_1 + \tau$ into $P_N(t_1 + \tau + 1) = [N_1(t_1 + \tau + 1), \ldots, N_{k+1}(t_1 + \tau + 1)]$ at time $t_1 + \tau + 1$ of the form

$$P(P_N(t_1 + \tau + 1) | P_N(t_1 + \tau))$$
$$= \sum_{\forall (\vartheta, \nu) \in G} P(\text{removing}[\, u_1, \ldots, u_{k+1}]\text{ variants})$$
$$\cdot P(\text{adding }[v_1, \ldots, v_{k+1}]\text{ variants}). \quad (2.3)$$

The set G contains all pairs of vectors $\vartheta = [u_1, \ldots, u_{k+1}]$ and $\nu = [v_1, \ldots, v_{k+1}]$ satisfying $0 \leq u_i \leq N_i(t_1 + \tau)$, $\sum_{i=1}^{k+1} u_i = u$ and $N(t_1 + \tau) - u + \sum_{i=1}^{k+1} v_i = N(t_1 + \tau + 1)$. (see the electronic supplementary material, appendix A for more mathematical details). In the following, we explore the temporal dynamic of the frequency change under the assumption that the replacement process is guided by neutral evolution, frequency- and age-dependent selections.

Neutral cultural evolution means that in finite populations cultural variant types are chosen to be added to the population according to their relative frequency and new variants are introduced by a process of random mutation at rate μ (see [11,12,17]). This implies that the probability for choosing a variant of type i to be added to the population is given by

$$p_i(t) = \left[\frac{N_i(t) - u_i}{N(t) - u} \right] (1 - \mu) \quad \text{for} \quad i = 1, \ldots, k$$

and $\quad p_{k+1}(t) = \frac{N_{k+1}(t) - u_{k+1}}{N(t) - u} + \mu \sum_{i=1}^{k} \frac{N_i(t) - u_i}{N(t) - u}.$

Variant type $k + 1$ contains all variant types not observed in the sample $\bar{S}(t_1)$ and also by definition all mutations.

Frequency-dependent selection is defined as the disproportional preference for either low-frequency variant types (anti-conformist bias) or high-frequency variant types (conformist bias) (e.g. [4]). In order to incorporate those biases in our framework, we define the probability of choosing a variant of type i, $i = 1, \ldots, k + 1$ to be added to the population by $p_i(t) = \tilde{p}_i(t)1/w$ with

$$\tilde{p}_i(t) = \begin{cases} 0, & \text{if} \quad \frac{N_i(t)-u_i}{N(t)-u} + b_{\text{freq}}\left(\hat{k}\frac{N_i(t)-u_i}{N(t)-u}-1\right) < 0 \\ \left[\frac{N_i(t)-u_i}{N(t)-u} + b_{\text{freq}}\left(\hat{k}\frac{N_i(t)-u_i}{N(t)-u}-1\right)\right](1-\mu), & \text{else} \end{cases}$$

and

$$\tilde{p}_{k+1}(t) = \begin{cases} \dfrac{\mu}{1-\mu}\sum_{i=1}^{k}\tilde{p}_i, & \text{if} \quad \frac{N_{k+1}(t)-u_{k+1}}{N(t)-u} + b_{\text{freq}}\left(\hat{k}\frac{N_{k+1}(t)-u_{k+1}}{N(t)-u}-1\right) < 0 \\ \left[\dfrac{N_{k+1}(t)-u_{k+1}}{N(t)-u} + b_{\text{freq}}\left(\hat{k}\frac{N_{k+1}(t)-u_{k+1}}{N(t)-u}-1\right) + \dfrac{\mu}{1-\mu}\sum_{i=1}^{k}\tilde{p}_i\right], & \text{else} \end{cases}$$

(2.4)

and $w = \sum_{i=1}^{k+1}\tilde{p}_i(t)$. The parameter b_{freq} defines the strength and direction of the selection. For $b_{\text{freq}} > 0$, selection favours variant types with frequencies higher than the relative majority $1/\hat{k}$ (where $\hat{k}$ denotes the number of variant types with non-zero frequency at t) and for $b_{\text{freq}} < 0$, selection favours variant types with frequencies lower than $1/\hat{k}$. In other words, b_{freq} controls the importance of selection compared with neutral evolution (equation (2.4) coincides with neutral evolution for $b_{\text{freq}} = 0$). The simplifying assumption that variant type $k + 1$ contains all unobserved and newly invented variants (i.e. variant type $k + 1$ has a higher frequency than any of the individual types) will slightly overestimate the selection pressure on this variant type in conformist situations and underestimate the selection pressure in anti-conformist situations. This naturally influences the temporal frequency change dynamic but simulation studies indicate that those influences are very small for sample sizes of the magnitude considered in §3.

Age-dependent selection (or pro-novelty selection) is here defined as a disproportional preference for 'young' variant types (e.g. [33]). Naturally, there can exist different preferences; however, we restrict ourselves to the situation where young variant types possess a selective advantage just because they are novel. We note that age-dependent selection and anti-conformity result in different cultural dynamics. While anti-conformity favours variant types below a certain threshold regardless of their age, age-dependent selection only supports 'young' variants. Consequently, the effects of both transmission processes might be similar on newly invented and therefore low-frequency variant types; however, they are very different on 'older' variant types returning to a low frequency.

Besides its frequency, each variant type i is now also characterized by its age, denoted by a_i, defined as the time it has been present in the population. We assume: the younger the age of type i the stronger it is favoured by selection. Then the probability of choosing a variant of type i, $i = 1, \ldots, k + 1$ to be added to the population is given by $p_i(t) = \tilde{p}_i(t)1/w$ with

$$\tilde{p}_i(t) = \left[\frac{N_i(t)-u_i}{N(t)-u}\mathrm{e}^{-b_{\text{age}}a_i(t)}\right](1-\mu)$$

$$\text{and } \tilde{p}_{k+1}(t) = \frac{N_{k+1}(t)-u_{k+1}}{N(t)-u}\mathrm{e}^{-b_{\text{age}}a_{k+1}(t)} + \mu\sum_{i=1}^{k}\frac{N_i(t)-u_i}{N(t)-u}\mathrm{e}^{-b_{\text{age}}a_i(t)}$$

(2.5)

and $w = \sum_{i=1}^{k+1}\tilde{p}_i(t)$. The parameter b_{age} defines how quickly the selective advantage of young variant types fades with increasing age. Similar to frequency-dependent selection, b_{age} controls the importance of selection compared with neutral evolution (equation (2.5) coincides with neutral evolution for $b_{\text{age}} = 0$). As variant type $k + 1$ contains all unobserved and newly invented variants the age-dependent selection pressure on this variant type can only be an assumed average value. Again this influences the temporal frequency change dynamic. But as it is almost impossible to obtain information about the age of unobserved variant types from observable data, we believe this simplification is justified.

2.1.3. Generating theoretical samples
Lastly, we generate a theoretical sample $S(t_2)$ at time t_2 of size $n(t_2)$. Instead of sampling from the theoretical population $P(t_2)$ at time t_2 only, we randomly draw $n(t_2)$ variants evenly from the populations $\{P(t_1 + \tau), \tau = 1, \ldots, t_2 - t_1\}$ that have been determined as described in the previous section. In this way, we account for possible time-averaging effects in the considered interval $[t_1, t_2]$ by, for example, allowing variant types to be present in the sample $S(t_2)$ at t_2 which are already extinct in the population $P(t_2)$.

2.2. Inference of underlying transmission process
The second part of the inference framework includes the statistical comparison between observed frequency changes in the interval $[t_1, t_2]$ and frequency changes expected under a specific hypothesis $\mathcal{H}$ about the underlying mechanisms of cultural transmission. As described above, we consider in the following the hypotheses $\mathcal{H}$: 'neutral evolution', $\mathcal{H}$: 'frequency-dependent selection with strength b_{freq}' and $\mathcal{H}$: 'age-dependent selection with strength b_{age}'.

To do so, we generate 100 000 theoretical samples $S(t_2)$ conditioned on a specific hypothesis $\mathcal{H}$ according to the procedure described in §2.1 and derive their statistical properties. In particular, we determine (i) the marginal frequency distributions $P(S_i(t_2)|\bar{S}(t_1), \mathcal{H})$ of all $k + 1$ variant types at t_2 and (ii) the probability distribution $P(\kappa(t_2)|\bar{S}(t_1), \mathcal{H})$ of the number of variant types, denoted by $\kappa(t_2)$, that are present in the sample at t_1 and are still present in the sample $S(t_2)$ at t_2 conditioned on the assumed transmission hypothesis $\mathcal{H}$. Comparisons of these distributions with the observed frequencies of the k variant types and the observed number of remaining initially present variants at time t_2 allow for conclusions on whether the assumed transmission mechanism $\mathcal{H}$ could have produced the frequency changes in the samples $\bar{S}(t_1)$ and $\bar{S}(t_2)$.

While the procedure described above is appropriate for neutral evolution, our definition of frequency- and age-dependent selections (equations (2.4) and (2.5)) allows for

varying selection strengths b_{freq} and b_{age}, respectively. Naturally, different selection strengths affect the distributions $P(S_i(t_2)|\bar{S}(t_1), \mathcal{H})$ and $P(\kappa(t_2)|\bar{S}(t_1), \mathcal{H})$. This leads to the question of which parameter values of b_{freq} and/or b_{age} describe the observed sample $\bar{S}(t_2)$ best. As the complete exploration of the relevant parameter spaces is too costly, we use the statistical inference method ABC.

2.2.1. Approximate Bayesian computation

The general idea of this method is to approximate the (joint) posterior distributions of the model parameters, summarized in the vector θ, given the observed sample $\bar{S}(t_2)$ in situations where the likelihood function of the considered model is either impossible or computationally prohibitive to obtain (e.g. [34]). Broadly speaking, the posterior distributions describe the parameter ranges which could have produced the observed sample under a given tolerance level ε. These distributions are obtained by generating samples $S(t_2)$ with parameter values θ sampled from appropriately chosen prior distributions. Parameter combinations resulting in samples 'close enough' to the observed sample $\bar{S}(t_2)$ are retained and provide a sample of the distribution $P(\theta|d(S(t_2), \bar{S}(t_2)) \leq \varepsilon)$. The function d describes the distance between the theoretical and observed sample and ε, the tolerance level that determines the level of approximation. As $\varepsilon \to 0$, the retained sample asymptotically approaches the true posterior distribution $P(\theta|\bar{S}(t_2))$ (e.g. [34]). In other words, the distributions $P(\theta|d(S(t_2), \bar{S}(t_2)) \leq \varepsilon)$ are only a meaningful approximation of the likelihood function of the generative model if the tolerance level ε is sufficiently small. In the following, we use a sequential Monte Carlo approach (e.g. [35,36]). The idea here is to decompose the problem of sampling from the distribution $P(\theta|d(S(t_2), \bar{S}(t_2)) \leq \varepsilon)$ into a series of simpler sub-problems. For that we choose a sequence of tolerance levels $\{\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_n\}$ with $\varepsilon_0 > \varepsilon_1 > \cdots > \varepsilon_n \geq 0$ and draw in the first step a sample from $P(\theta|d(S(t_2), \bar{S}(t_2)) \leq \varepsilon_0)$ with ε_0 large and subsequently from an increasingly constrained sequence of distributions

$$P(\theta|d(S(t_2), \bar{S}(t_2)) \leq \varepsilon_i), \qquad i = 1, \ldots, n$$

until the desired tolerance level $\varepsilon = \varepsilon_n$ is reached. The crucial question, however, is how we define the distance between the theoretical samples $S(t_2)$ and the observation $\bar{S}(t_2)$. In practice, data are often reduced to a set of summary statistics, denoted by $\mathcal{S}$, producing samples of the distribution $P(\theta|d(\mathcal{S}(S(t_2)), \mathcal{S}(\bar{S}(t_2))) \leq \varepsilon)$. This approach only generates exact inference results if the summary statistic is sufficient. Here we determine the distance between theoretical and observed sample by calculating

— the absolute distance between observed and theoretical frequencies

$$d_d = \frac{1}{k} \sum_{i=1}^{k} |\bar{S}_i(t_2) - S_i(t_2)|, \tag{2.6}$$

— the absolute difference between the median m_i of the distributions $P(S_i(t_2)|\bar{S}(t_1), \mathcal{H})$ and the observed frequencies

$$d_m = \frac{1}{k} \sum_{i=1}^{k} |\bar{S}_i(t_2) - m_i|. \tag{2.7}$$

Additionally, we require that the number of remaining initially present variant types coincides with the observed

number. We note that the statistics d_d and d_m are based directly on the actual frequencies. While d_d compares the observed and theoretical frequencies directly d_m compares the observed frequencies to the median frequencies (a measure of the average value of the possible frequency range) under the considered transmission hypothesis $\mathcal{H}$. Consequently, using d_m as the statistic infers likely parameter constellations whereas using d_d infers possible parameter constellations which, however, might only produce the observed sample $\bar{S}(t_2)$ with a small likelihood. In the electronic supplementary material, appendix B (cf. figure S1), we show that inference based on the statistic d_m provides the tightest posterior distribution of the model parameters and therefore we build the analysis in the next section on this statistic. In the electronic supplementary material, appendix F, we show the corresponding results for the analysis based on d_d.

2.3. Theoretical limitations of the inference procedure

To explore the theoretical limits of inference based on population-level data, we assume the presence of seven cultural variant types with the frequencies $\bar{S}(t_1) = [1, 2, 5, 10, 20, 40, 60]$ at $t_1 = 0$. We use the modelling framework to generate samples $\bar{S}(t_2)$ under the hypothesis of neutral evolution (i.e. $b_{\mathrm{freq}} = 0$), negative frequency-dependent selection (i.e. $b_{\mathrm{freq}} = -0.01$), positive frequency-dependent selection (i.e. $b_{\mathrm{freq}} = 0.0.1$) and the parameter constellation $r = 0.1$ and $s = 0.3$ at $t_2 = 50$. Application of the ABC procedure identifies the cultural transmission processes (parametrized by the selection strength b_{freq}) that are consistent with the observed data. Figure 2 shows the posterior distributions of the parameter b_{freq} if the 'true' selection strength is $b_{\mathrm{freq}} = -0.01$ (figure 2a), $b_{\mathrm{freq}} = 0$ (figure 2b) and $b_{\mathrm{freq}} = 0.01$ (figure 2c). Those posterior distributions describe the parameter range of b_{freq} for which the distance between theoretical and observed samples after 50 time steps as measured by equation (2.7) is almost zero. The widths of the distribution show that the information that can be inferred from sparse frequency data is limited. For example, in figure 2b, the data were produced under neutral evolution but frequency-dependent selection can generate almost identical frequency patterns resulting in a relatively wide posterior distribution for b_{freq}. It is theoretically impossible to distinguish between neutral evolution and weak frequency-dependent selection on the basis of such sparse data. This implies that there exist theoretical limits on the amount of information that can be extracted from population-level frequency data.

3. Case study

In the following, we apply the framework developed in §2 to a dataset describing the culture of the earliest farmers in Central Europe, the so-called Linear Pottery Culture (or LBK after its German name), *ca* 7500–7000 years ago. The dataset (given in [25,37]) records the frequency of different types of decorated vessels at seven different points in time, denoted with t_j, $j = 1, \ldots, 7$, defining six phases of cultural change which vary in duration. Our aim is to explore whether observed frequency changes in different types of pottery between the beginning and the end of each of the six phases are consistent with a specific hypothesis about the underlying mechanism of cultural transmission. The
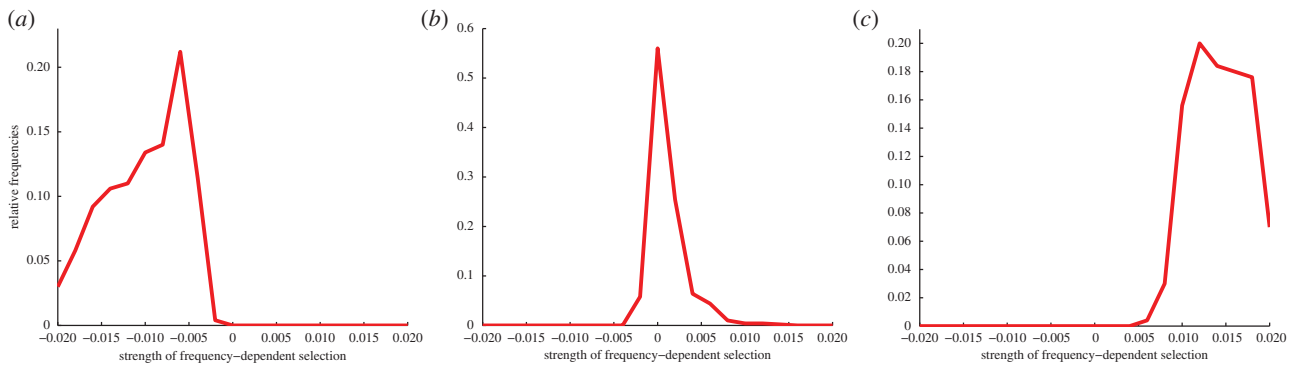
Figure 2. Posterior distributions of the model parameter b_{freq} for (a) anti-conformist selection with $b_{freq} = -0.01$, (b) neutral evolution and (c) conformist selection with $b_{freq} = 0.01$ based on the statistic d_m. (Online version in colour.)

Table 1. Number of variants outside the 95% intervals of the marginal frequency distribution for $r = 0.1$; 0.3; 0.5.

replacement	phase I out of 10	phase II out of 10	phase III out of 19	phase IV out of 26	phase V out of 35	phase VI out of 20
$\langle r \rangle = 0.5$	2 (=20%)	1 (=10%)	3 (≈16%)	6 (≈23%)	7 (=20%)	3 (=15%)
$\langle r \rangle = 0.3$	2 (=20%)	5 (=50%)	6 (≈32%)	9 (≈35%)	8 (≈23%)	4 (=20%)
$\langle r \rangle = 0.1$	2 (=20%)	5 (=50%)	7 (≈37%)	15 (≈58%)	8 (≈23%)	7 (=35%)

assumption that cultural transmission is relevant in this case is based on ethnographic studies showing that knowledge of pottery-making techniques is transmitted between potters. Previous studies of the transmission of LBK decoration patterns [12,25,38,39] have been inconclusive, some consistent with neutral evolution and others rejecting it in favour of anti-conformist transmission.

For the analysis, we use the observed sample sizes $n(t_j)$ at times t_j and assume that the corresponding population sizes $N(t_j)$ are determined by $N(t_j) = (1/s)n(t_j)$, $j = 1, \ldots, 7$ with $s \sim \mathcal{N}(\langle s \rangle, \sigma_N^2)$. Estimating even the average sampling fraction s is, of course, very problematical for archaeologists because the currently available record is the result of a large number of local factors relating to site survival and recovery. However, in the case of the LBK in Germany the work of Zimmermann *et al.* [40,41] provides a basis. Using data from small areas that have been very thoroughly excavated they calculated the estimated number of households for the different LBK settlement areas of Germany [40, fig. 15]. For the areas included in Strien's study this is *ca* 2300 at the peak density. The total number of distinct finds in Strien's dataset [37], which should roughly correspond to households, is 270. While this is based on all the phases, it suggests that an average sampling fraction of 0.1 is the right order of magnitude and we assume $s \sim \mathcal{N}(0.1, 0.01^2)$. The samples represent on average 10% of the population at t_j but we allow for sufficient variation to account for uncertainties in those estimates. Population sizes for intermediate time steps are obtained by linear interpolation between $N(t_j)$ and $N(t_{j+1})$. Further, in temperate climates such as Central Europe pottery making is largely seasonal and generally takes place in one bout of activity per year [42]. Accordingly, the lengths of the phases correspond to the number of potential opportunities for copying within a given phase and thus the number of opportunities for the

population of decorative motifs to change as a result of the operation of drift, innovation and other potential forces.

3.1. Analysis of neutral evolution

We start our analysis by investigating whether observed frequency changes in different types of pottery between different phases are consistent with the hypothesis of neutral evolution. For that we determine the marginal frequency distributions $P(S_i(t_{j+1})|\bar{S}(t_j), \mathcal{H})$ of all initially present variant types and the probability distribution $P(\kappa(t_{j+1})|\bar{S}(t_j), \mathcal{H})$ of the remaining initially present variant types for all phases under neutral evolution. We assume that the fraction r in equation (2.2) is given by $r \sim \mathcal{N}(\langle r \rangle, 0.05^2)$ with $\langle r \rangle = 0.1$; 0.3; 0.5. Varying the mean value $\langle r \rangle$ allows us to explore the effects of different replacement fractions and consequently different individual lifespans of the cultural variants on the frequency change dynamic. Shott [43] reviews a range of ethnoarchaeological studies of pottery use life and finds a general correlation with vessel size: smaller vessels tend to have short use lives with larger vessels surviving longer. The data show that the replacement fractions used are plausible ones. It should also be noted that in the case of the LBK the vessels are very uniform in size and shape.

The 95% intervals of the marginal frequency distributions $P(S_i(t_{j+1})|\bar{S}(t_j))$ of all initially present variant types at the end of each phase represent the frequency ranges of each type of pottery that are possible under neutral evolution (see the electronic supplementary material, appendix C, figure S3 for the concrete intervals). Table 1 shows that in each phase a number of observations lie outside these intervals for $\langle r \rangle = 0.1$; 0.3 and therefore we argue that the observed frequencies are not likely to be produced by neutral evolution with an average lifespan of an individual variant of 10 and 3.33 years, respectively. For $\langle r \rangle = 0.5$ the neutral hypothesis
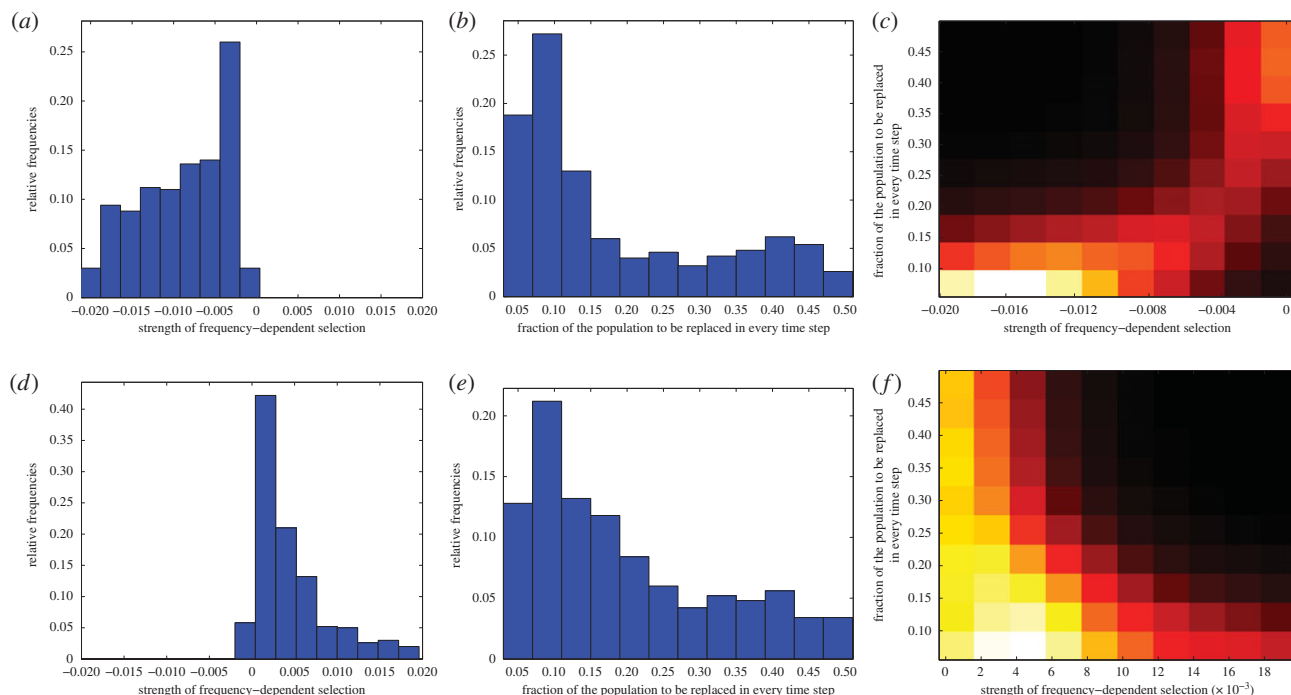
Figure 3. Posterior distribution of b_{freq} (a,d), r (b,e) and joint posterior distribution (c,f: lighter areas coincide with regions of the parameter space that possess a higher probability) at the end of phase I (a–c) and V (d–f). We assume uniform priors over the intervals $[-0.02, 0.02]$ and $[0.05, 0.5]$. (Online version in colour.)

gains slightly more support; fewer observations lie outside the intervals; however, with the possible exception of the phases II, III and VI, we also conclude that the observed frequencies are not likely to be produced by neutral evolution with an average lifespan of an individual variant of 2 years.

In the next step, we analyse the distributions of the number of remaining initially present variant types for each phase. The relatively long lifespan of the individual cultural variants for $\langle r \rangle = 0.1$ predicts that under neutral evolution more variant types than observed should survive until the end of phases IV and V. For $\langle r \rangle = 0.3$; 0.5, the number of remaining initially present variants in phase V contradicts the hypothesis of neutral evolution (see the electronic supplementary material, appendix C, figure S4 for details).

Summarizing, based on the analysis of the two statistics we found no consistent evidence that neutral evolution with an average sampling fraction $\langle s \rangle = 0.1$ was the dominant transmission mechanism over the 500 years the LPK lasted.

3.2. Analysis for frequency-dependent selection

In this section, we ask whether frequency-dependent selection can describe the observed changes better than neutral evolution and assume that the probability of choosing a particular variant type is given by equations (2.4). To do so, we apply the ABC inference procedure based on the statistic d_m. The results of this procedure are posterior distributions of the model parameters r and b_{freq} describing the parameter ranges for which the distance between theoretical and observed samples as measured by equation (2.7) is smaller than or equal to some tolerance level ε. Figure 3 shows those distributions for the phases I and V (figure 3a,d: posterior distribution of b_{freq}; figure 3b,e: posterior distribution of r; figure 3c,f: joint posterior distributions of $\theta = (r, b_{freq})$); the complete analysis can be found in the electronic supplementary material, appendix D. It is apparent that $b_{freq} = 0$ is contained in the ranges of the posterior distribution of all

phases meaning that frequency-dependent selection does not explain the observed samples $\bar{S}(t_j)$ better (but also not worse) than neutral evolution. In detail, based on the posterior distribution for b_{freq}, we exclude the presence of conformist selection in the population for phases I–III. By contrast, we exclude the presence of anti-conformist selection in phase V. The posterior distributions in the phases IV and VI indicate no strong frequency-dependent selection pressure. We conclude that a relatively large parameter range of b_{freq} could produce the observed frequency changes between the different phases (given a certain tolerance level) which demonstrates the uncertainties associated with inferences based on sparse data. But importantly we note that the achievable tolerance level ε in the ABC analysis is relatively large in the different phases, indicating that frequency-dependent selection (and neutral evolution) might not be the best model to explain the observed data. To explore how well the identified parameter ranges are able to replicate the observed changes in frequency in more detail, we determine the marginal frequency distributions $P(S_i(t_{j+1})|\bar{S}(t_j), \mathcal{H})$ of all k variant types present at t_j based on the joint distributions of the model parameters $\theta = (r, b_{freq})$ shown in figure 3. Electronic supplementary material, figure S6 in appendix D shows that a similar number of variant types as in the case of neutral evolution with $\langle r \rangle = 0.5$ (cf. Table 1) lie outside the 95% intervals of $P(S_j(t_{j+1})|\bar{S}(t_j), \mathcal{H})$ pointing again to the fact that the hypotheses of neutral evolution and frequency-dependent selection have similar explanatory power.

3.3. Analysis of age-dependent selection

Next, we explore whether age-dependent selection can describe the observed changes better than neutral evolution and assume that the probability p_i of choosing a particular variant type is given by equations (2.5). This means that p_i not only depends on the frequency but also on the age of
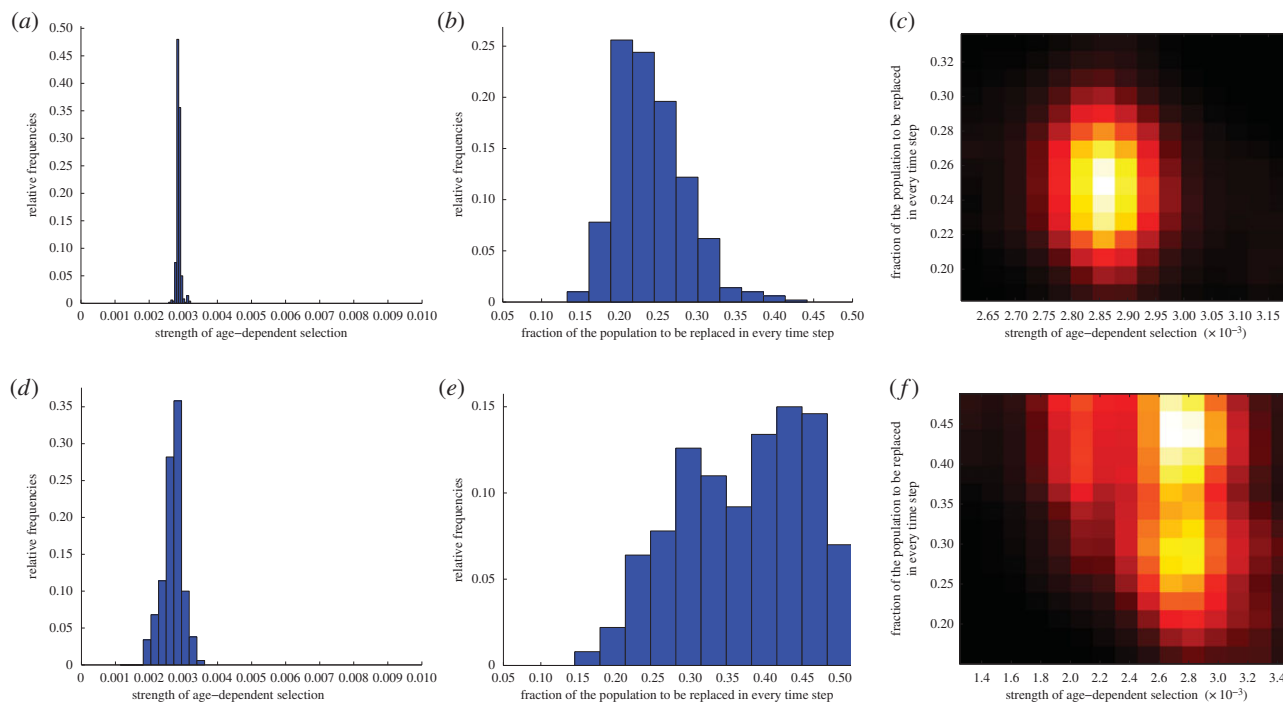
Figure 4. Posterior distribution of b_{age} (a,d), r (b,e) and joint posterior distribution (c,f: lighter areas coincide with regions of the parameter space that possess a higher probability) at the end of phase I (a–c) and V (d–f). We assume uniform priors over the intervals [0, 0.01] and [0.05, 0.5]. (Online version in colour.)

the variant type, and consequently, we need to estimate the time of innovation for each type. This is a challenging task as archaeological insights can help constructing an interval estimate, but an accurate innovation time cannot be obtained with any certainty. Therefore, we apply the ABC inference procedure not only to determine the posterior distribution of the fraction r of the population that is replaced in every time step and the strength of age-dependent selection b_{age} but also to determine the posterior distributions of the innovation times of all variant types and the average age of the variant types not observed in the sample at the beginning of each phase. Consequently, age-dependent selection represents a more complex hypothesis than frequency-dependent selection. Archaeological interval estimates allow us to define relatively tight prior distributions for the innovation times (at least for the observed variant types) which restrict the variability of those parameters. Figure 4 shows the posterior distributions of b_{age} (figure 4a,d), r (figure 4b,e) and the joint posterior distribution of $\theta = (r, b_{age})$ (figure 4c,f) for the phases I and V; a complete analysis can be found in the electronic supplementary material, appendix E. We note that in contrast with frequency-dependent selection, the achievable distance ε between theoretical and observed samples as measured by equation (2.7) is relatively small for all phases. We further observe that $b_{age} = 0$ is not contained in the ranges of the posterior distributions and therefore conclude that the hypothesis of age-dependent selection describes the observed samples $\bar{S}(t_j)$ better than neutral evolution. Further, the marginal frequency distributions $P(S_i(t_{j+1})|\bar{S}(t_j), \mathcal{H})$ of all k variant types present at t_j produced by the joint distributions of the model parameters $\theta = (r, b_{age})$ shown in figure 4 indicate that, apart from phases II and V, almost no variant type lies outside the 95% intervals of $P(S_i(t_{j+1})|\bar{S}(t_j), \mathcal{H})$ (see the electronic supplementary material, appendix E, figure S8). Therefore, we conclude that the observed frequency changes especially between the beginning and the end of the phases I, III, IV

and VI are consistent with the hypothesis of age-dependent selection. The Linear Pottery early farmer population might have had a mild preference for novel as opposed to longer established types of decorated pottery.

4. Qualitative differences between neutral evolution and age-dependent selection

We have shown that age-dependent selection, i.e. the preference for novel cultural variants, is able to explain the LBK dataset better than neutral evolution. But what are the qualitative differences between the two cultural transmission processes? To explore these differences, we use our simulation framework and record the cumulative frequencies of each variant type that has been introduced in the population at some point over a period of 20 000 time steps under the assumption of $r = 0.3$ and neutral evolution and age-dependent selection, respectively. To study the effects of the transmission processes only we keep the population size constant over time ($N = 500$) and allow for an initial burn-in time. Figure 5 shows that age-dependent selection produces variant frequency distributions which differ from neutral evolution (see [15–17] for results for neutral evolution with $r = 1$). Age-dependent selection results in more variant types with medium-high frequencies and fewer variant types with high frequencies.

In the electronic supplementary material, appendix G, we show that under neutral evolution variant types with a high frequency are usually 'old'. Consequently, the lack of high-frequency variant types in situations of age-dependent selection is explained by the fact that 'old' variant types are selected against and therefore variant types do not have sufficient time to accumulate a high frequency. Further, we demonstrate in the electronic supplementary material, appendix G, that age-dependent selection leads to a higher turnover rate of top lists compared with neutral evolution (see [22,45]
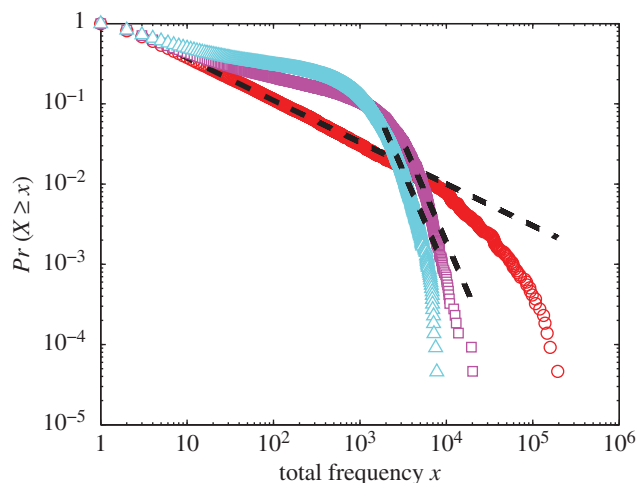
Figure 5. Distribution of the cumulative frequencies of all variant types that have been present in the population at some point over a period of 20 000 time steps for $r = 0.3$ and neutral evolution (red circles), age-dependent selection with $b_{age} = 1/300$ (magenta squares) and $b_{age} = 1/100$ (cyan triangles). The dashed lines illustrate the fitted power laws (see [44] for details on the fitting procedure). (Online version in colour.)

for results under neutral evolution). Additionally, under age-dependent selection, variant types show a more monotonic decrease behaviour after they have reached their frequency maximum (see electronic supplementary material, table S1). Summarizing, age-dependent selection accelerates the speed of cultural change (see also [33]) by selecting against 'old' and therefore cumulatively high-frequency variant types.

5. Conclusion

In this paper, we developed an inference framework which evaluates whether observed frequency changes of cultural variant types between two time points t_1 and t_2 are consistent with a specific hypothesis about the underlying process of cultural transmission. In particular, we were interested in exploring the consistency with neutral evolution, frequency- and age-dependent selections. Naturally, this is not an all-encompassing analysis and alternative transmission hypotheses could also be consistent with the observed frequency changes. The framework developed takes into account that observed frequency data often describe the composition of samples and *not* of whole populations of cultural variants. Using the Dirichlet distribution approach, we generated population structures from which the observed sample at time point t_1, denoted by $\bar{S}(t_1)$, could have been drawn randomly. Based on those inferred population structures, we determined the temporal evolution of the frequencies of the cultural variant types conditioned on a specific transmission hypothesis. In doing so, we assumed that only a fraction of the population, denoted by r, is renewed in every time step, which implies that individual variants have an average life-span of $1/r$. In this way, we generated populations of cultural variants at each time point in the considered interval $[t_1, t_2]$. To further account for potential time-averaging processes, we produced samples with a size equal to the observed sample at t_2, denoted by $\bar{S}(t_2)$, by randomly sampling from all generated populations in $[t_1, t_2]$. Consequently, the outputs of the developed framework are theoretical samples at t_2 conditioned on the cultural

transmission hypothesis considered. Statistical analysis of the samples then informs about, for example, the frequency ranges of the cultural variants that are expected under the considered hypothesis.

Here we use ABC to determine which transmission processes are consistent with the observed changes in frequency. To test whether the hypothesis of frequency- or age-dependent selection explains the data better than neutral evolution, we approximated the posterior distribution of the parameters $\theta = (r, b_{age})$ (with b_{age} describing the strength of frequency-dependent selection) and $\theta = (r, b_{age})$ (with b_{age} describing the strength of age-dependent selection), respectively. Those posterior distributions indicate the parameter values and therefore the transmission hypotheses which could have produced sample frequencies similar to the observed ones. To this end, we evaluated the similarities between the theoretical samples obtained by frequency- or age-dependent selection and the observed sample at t_2 in two different ways. Firstly, we compared the observed frequencies $\bar{S}(t_2)$ to the median frequencies (a measure of the average value of the possible frequency range) of the considered transmission process with the chosen parametrization θ (cf. statistic d_m described in equation (2.7)). We retained the parameter constellations which produced frequency data sufficiently close to the observed data and subsequently generated posterior distributions for the model parameters describing probable parameter constellations. Secondly, we compared the obtained theoretical frequencies directly to the observed frequencies $\bar{S}(t_2)$ (cf. statistic d_d described in equation (2.6)). Again we retained the parameter constellations which produced frequency data sufficiently close to the observed data and subsequently generated posterior distributions for the model parameters describing possible parameter constellations. Possible means that the transmission hypothesis with the parametrization θ could have produced the observed sample frequency but sometimes only with a very small probability. We found, unsurprisingly, that applications of the ABC algorithm based on d_d resulted in substantially wider posterior distributions indicating that many transmission processes could theoretically produce the observed frequencies $\bar{S}(t_2)$ although this might only happen with a very small probability (see the electronic supplementary material, appendices D and F). This strongly points to the uncertainties associated with inferences based on sparse population-level data.

The observed widths of the posterior distributions led to the question: How well should we expect to be able to distinguish between different transmission processes on the base of the observed changes in frequency? We used the developed model to generate frequency data under a known underlying transmission hypothesis (in detail, we assumed frequency-dependent selection with $b_{freq} = -0.01$; 0; 0.01 and $r = 0.3$, see the electronic supplementary material, appendix B) and applied the inference framework to these data. The resulting posterior distributions of the model parameters r and b_{freq} indicate the ranges of the parameter values that could have produced the observed data (figure 2) and consequently the widths of those distributions describe the information about the underlying transmission process contained in the population-level data. Our example showed clearly that the information is limited. For example, it is theoretically impossible to distinguish between neutral evolution and weak frequency-dependent selection on

the base of sparse population-level frequency data (cf. figure 2b). Even though the data were produced under the assumption of neutral evolution, i.e. $b_{freq} = 0$, values of b_{freq} within the interval $[-0.005; 0.01]$ corresponding to weak/medium anti-conformity and conformity can generate almost identical frequency patterns. The main insight from the described inference analysis is the exclusion of cultural transmission processes (e.g. conformity with $b_{freq} > 0.01$) that could *not* have produced the observed sample based on the inference procedure used and therefore the reduction of the pool of potential hypotheses [46].

Further, our results show that the cultural system considered is more sensitive to changes in the selection strengths, b_{freq} or b_{age}, than to the fraction r of the population that is renewed in every time step (and therefore the average lifespan of an individual variant) (cf. the widths of the posterior distributions for r in figures 3 and 4). Considering the difficulties of estimating r from archaeological data, this points to the fact that moderate estimation errors do not affect the temporal dynamic of cultural change substantially.

We applied this framework to a dataset describing the culture of the earliest farmers in Central Europe, the so-called LBK, *ca* 7500–7000 years ago. Our aim was to investigate whether the observed frequency changes between the beginning and the end of a phase of the different types of pottery decoration are consistent with the hypothesis of neutral evolution. To do so, we determined the marginal frequency distributions $P(S_i(t_{j+1})|\bar{S}(t_j), \text{neutral evolution})$ of all variant types initially present at t_i and the probability distribution $P(\kappa(t_{j+1})|\bar{S}(t_j), \text{neutral evolution})$ of the initially present variant types remaining at t_{j+1} under the assumption that cultural transmission occurs in an unbiased way. We concluded that neutral evolution is not a likely hypothesis for the underlying transmission process as in almost all phases a relatively large proportion of the observed frequencies of the different variant types were outside the 95% interval of the ranges expected under neutrality. Additionally we found that in phase V the observed number of surviving variant types between the beginning and the end of this phase was too low to have been produced by neutral evolution.

To test whether frequency-dependent selection describes the observed frequency changes better than neutral evolution, we employed ABC. Based on the resulting posterior distributions of the model parameters r and b_{freq}, we were able to exclude conformity in phases I–III and anti-conformity in phase V as probable causes of the observed frequency changes. Overall, we found that frequency-dependent transmission did not explain the data better than neutral evolution. The parameter value $b_{freq} = 0$ (indicating neutral evolution) was contained in the range of the posterior distribution of b_{freq} in all phases and a similar number of variant types lay outside the 95% intervals of marginal frequency distributions

$P(S_i(t_{j+1})|\bar{S}(t_j), \mathcal{H})$. Comparing these results with [25] points to the crucial importance of rooting any inference framework in population frequencies and not in sample frequencies.

We applied the same statistical procedure to test whether age-dependent selection (defined as a disproportional preference for 'young' variant types) explains the data better than neutral evolution. But besides the parameters r and b_{age}, describing the strength of the age-dependent selection process, we also estimated the innovation times of the variant types present in the dataset (as generally this information cannot be obtained precisely from the archaeological record), resulting in a more complex model than considered for frequency-dependent selection. We concluded that with the possible exception of phases II and VI, the observed frequency changes between the beginning and the end of a phase are consistent with the hypothesis of weak age-dependent selection. In particular, the hypothesis of age-dependent selection is able to explain the LBK data better than neutral evolution.

In this context, it has been demonstrated that neutral evolution is able to replicate the population-level pattern of fashion trends (e.g. popularity of baby names [15,17], dog breeds [16], popular music [22]). In particular, it has been argued that neutral evolution may explain why a few cultural variant types become highly popular but then fade away, producing the well-observed temporal dynamic of fashion trends [22]. Age-dependent selection favours 'young' recently invented variant types but disfavours 'old' variant types which have been in the population for some time. Consequently, age-dependent selection equally produces the characteristic frequency behaviour of fashion trends (see §4 and electronic supplementary material, figure S11) and we hypothesize that it could be an alternative explanation for the dynamic of fashion trends (see also [47] for a different hypothesis), or rather that the concept of 'fashion trends' may actually encompass dynamics of different types as age-dependent selection does not produce a power law. Summarizing, the observed frequency dynamic of the different types of decorated pottery in the LPK is consistent with age-dependent selection potentially indicating that the early farmers treated decorated pottery as fashion items, with more recent innovations being preferred because of their novelty. This observation might present the earliest instance of such fashion trends yet recorded.

References

1. Beheim AB, Thigpen C, McElreath R. 2014 Strategic social learning and the population dynamics of human behavior: the game of Go. *Evol. Hum. Behav.* **35**, 351–357. (doi:10.1016/j.evolhumbehav.2014.04.001)

2. Henrich J, Broesch J. 2011 On the nature of cultural transmission networks: evidence from Fijian villages for adaptive learning biases. *Phil. Trans. R. Soc. B* **366**, 139–148. (doi:10.1098/rstb.2010.0323)

3. Cavalli-Sforza LL, Feldman MW. 1981 *Cultural transmission and evolution*. Princeton, NJ: Princeton University Press.

4. Boyd R, Richerson PJ. 1985 *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.

5. Laland KN, Richerson PJ, Boyd R. 1996 Developing a theory of animal social learning. In *Social learning in animals: the roots of culture* (eds CM Heyes, BG Galef Jr), pp. 129–154. London, UK: Academic Press.

6. Henrich J. 2001 Cultural transmission and the diffusion of innovations: adoption dynamic indicate that biased cultural transmission is the predominant force in behavioural change. *Am. Anthropol.* **103**, 992–1013. (doi:10.1525/aa.2001.103.4.992)

7. Reader SM. 2004 Distinguishing social and asocial learning using diffusion dynamics. *Anim. Learn. Behav.* **32**, 90–104. (doi:10.3758/BF03196010)

8. Kendal JR, Kendal RL, Laland KN. 2007 Quantifying and modelling social learning processes in monkey populations. *Int. J. Psychol. Psychol. Ther.* **7**, 123–138.

9. Franz M, Nunn CL. 2009 Network-based diffusion analysis: a new method for detecting social learning. *Proc. R. Soc. B* **276**, 1829–1836. (doi:10.1098/rspb.2008.1824)

10. Hoppitt WJE, Kandler A, Kendal JR, Laland KN. 2010 The effect of task structure on diffusion dynamics: implications for diffusion curve and network-based analyses. *Learn. Behav.* **38**, 243–251. (doi:10.3758/LB.38.3.243)

11. Neiman FD. 1995 Stylistic variation in evolutionary perspective. *Am. Antiq.* **60**, 7–36. (doi:10.2307/282074)

12. Shennan SJ, Wilkinson JR. 2001 Ceramic style change and neutral evolution: a case study from neolithic Europe. *Am. Antiq.* **66**, 577–593. (doi:10.2307/2694174)

13. Kohler TA, Van Buskirk S, Ruscavage-Barz S. 2004 Vessels and villages: evidence for conformist transmission in early village aggregations on the Pajarito Plateau, New Mexico. *J. Anthropol. Archaeol.* **23**, 100–118. (doi:10.1016/j.jaa.2003.12.003)

14. Steele J, Glatz C, Kandler A. 2010 Ceramic diversity, random copying, and tests for selectivity in ceramic production. *J. Archaeol. Sci.* **37**, 1348–1358. (doi:10.1016/j.jas.2009.12.039)

15. Hahn MW, Bentley RA. 2003 Drift as a mechanism for cultural change: an example from baby names. *Proc. R. Soc. Lond. B* **270**, 120–123. (doi:10.1098/rsbl.2003.0045)

16. Herzog HA, Bentley RA, Hahn W. 2004 Random drift and large shifts in popularity of dog breeds. *Proc. R. Soc. Lond. B* **271**, 353–356. (doi:10.1098/rsbl.2004.0185)

17. Bentley RA, Hahn MW, Shennan SJ. 2004 Random drift and culture change. *Proc. R. Soc. Lond. B* **271**, 1443–1450. (doi:10.1098/rspb.2004.2746)

18. Kimura M, Crow JF. 1964 The number of alleles that can be maintained in a finite population. *Genetics* **49**, 725–738.

19. Mesoudi A, Lycett SJ. 2009 Random copying, frequency-dependent copying and culture change. *Evol. Hum. Behav.* **30**, 41–48. (doi:10.1016/j.evolhumbehav.2008.07.005)

20. Powell A, Shennan SJ, Thomas MG. Submitted. The power of power-laws in cultural evolution.

21. Premo LS. 2014 Cultural transmission and diversity in time-averaged assemblages. *Curr. Anthropol.* **55**, 105–114. (doi:10.1086/674873)

22. Bentley RA, Lipo CP, Herzog HA, Hahn MW. 2007 Regular rates of popular culture change reflect random copying. *Evol. Hum. Behav.* **28**, 151–158. (doi:10.1016/j.evolhumbehav.2006.10.002)

23. Acerbi A, Bentley RA. 2014 Biases in cultural transmission shape the turnover of popular traits. *Evol. Hum. Behav.* **35**, 228–236. (doi:10.1016/j.evolhumbehav.2014.02.003)

24. Kandler A, Powell A. 2015 Inferring learning strategies from cultural frequency data. In *Learning strategies and cultural evolution during the palaeolithic* (eds A Mesoudi, K Aoki), pp. 85–101. Tokyo, Japan: Springer.

25. Kandler A, Shennan S. 2013 A non-equilibrium neutral model for analysing cultural change. *J. Theor. Biol.* **330**, 18–25. (doi:10.1016/j.jtbi.2013.03.006)

26. Pritchard JK, Stephens M, Donelly P. 2000 Inference of population structure using multilocus genotype data. *Genetics* **155**, 945–959.

27. Wilde S *et al.* 2014 Direct evidence for positive selection of skin, hair, and eye pigmentation in Europeans during the last 5,000y. *Proc. Natl Acad. Sci. USA* **111**, 4832–4837. (doi:10.1073/pnas.1316513111)

28. Kandler A, Wilder B, Fortunato L. Submitted. Inferring the process of cultural evolution from population-level patterns.

29. Wilder B, Kandler A. In press. Inference of cultural transmission modes based on incomplete information. *Hum. Biol.*

30. McElreath R, Bell AV, Efferson C, Lubell M, Richerson PJ, Waring T. 2008 Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Phil. Trans. R. Soc. B* **363**, 3515–3528. (doi:10.1098/rstb.2008.0131)

31. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. 2014 *Bayesian data analysis*, 3rd edn. Boca Raton, FL: CRC Press.

32. Madsen ME. 2012 Unbiased cultural transmission in time-averaged archaeological assemblages. (http://arxiv.org/abs/1204.2043)

33. Aoki K, Lehmann L, Feldman MW. 2011 Rates of cultural change and patterns of cultural accumulation in stochastic models of social transmission. *Theor. Popul. Biol.* **79**, 192–202. (doi:10.1016/j.tpb.2011.02.001)

34. Beaumont MA, Zhang W, Balding DJ. 2002 Approximate Bayesian computation in population genetics. *Genetics* **162**, 2025–2035.

35. Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MPH. 2009 Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. R. Soc. Interface* **6**, 187–202. (doi:10.1098/rsif.2008.0172)

36. Del Moral P, Doucet A, Jasra A. 2011 An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Stat. Comput.* **22**, 1009–1020. (doi:10.1007/s11222-011-9271-y)

37. Strien H-C. 2000 *Untersuchungen zur Bandkeramik in Württemberg, Universitätsforschungen zur Prähistorischen Archäologie*, 69. Bonn, Germany: Habelt.

38. Bentley RA, Shennan SJ. 2003 Cultural transmission and stochastic network growth. *Am. Antiq.* **68**, 459–486. (doi:10.2307/3557104)

39. Shennan SJ, Bentley RA. 2008 Style, interaction, and demography among the earliest farmers of central Europe. In *Cultural transmission and archaeology: issues and case studies* (ed. MJ O'Brien), pp. 164–177. Washington, DC: SAA Press.

40. Zimmermann A, Richter J, Frank T, Wendt KP. 2004 Landschaftsarchäologie II. Überlegungen zu Prinzipien einer Landschaftsarchäologie. *Ber. Römisch-Germanischen Kommission* **85**, 7–66.

41. Zimmermann A, Hilpert J, Wendt KP. 2009 Estimations of population density for selected periods between the neolithic and AD 1800. *Hum. Biol.* **81**, 357–380. (doi:10.3378/027.081.0313)

42. Arnold DE. 1985 *Ceramic theory and cultural process*. Cambridge, UK: Cambridge University Press.

43. Shott MJ. 1996 Mortal pots: on use life and vessel size in the formation of ceramic assemblages. *Am. Antiq.* **61**, 463–482. (doi:10.2307/281835)

44. Clauset A, Shalizi CR, Newman MEJ. 2009 Power-law distributions in empirical data. *SIAM Rev.* **51**, 661–703. (doi:10.1137/070710111)

45. Eriksson K, Jansson F, Sjöstrand T. 2010 Bentley's conjecture on popularity toplist turnover under random copying. *Ramanujan J.* **23**, 371–396. (doi:10.1007/s11139-010-9280-z)

46. Kandler A, Laland K. 2013 Tradeoffs between the strength of conformity and number of conformists in variable environments. *J. Theor. Biol.* **332**, 191–202. (doi:10.1016/j.jtbi.2013.04.023)

47. Acerbi A, Ghirlanda S, Enquist M. 2012 The logic of fashion cycles. *PLoS ONE* **7**, e32541. (idoi:10.1371/journal.pone.0032541)