Evolution and Human Behavior 30 (2009) 41–48

Random copying, frequency-dependent copying and culture change☆

Alex Mesoudi[a,b,*], Stephen J. Lycett[c,d]

[a]*Department of Social and Developmental Psychology, University of Cambridge, Free School Lane, CB2 3RQ Cambridge, UK*
[b]*School of Biological and Chemical Sciences, Queen Mary, University of London, Mile End Road, E1 4NS London, UK*
[c]*Department of Anthropology, University of Kent, Marlowe Building, Canterbury, CT2 7NR Kent, UK*
[d]*British Academy Centenary Research Project, SACE, University of Liverpool, Hartley Building, Brownlow Street, L69 3BX Liverpool, UK*

Initial receipt 3 June 2008; final revision received 30 July 2008

Abstract

Previous evolutionary analyses of human culture have found that a simple model of random copying, analogous to neutral genetic drift, can generate the distinct power-law frequency distribution of cultural traits that is typical of various real-world cultural datasets, such as first names, patent citations and prehistoric pottery types. Here, we use agent-based simulations to explore the effects of frequency-dependent copying (e.g., conformity and anti-conformity) on this power-law distribution. We find that when traits are actively selected on the basis of their frequency, then the power-law distribution is severely disrupted. Conformity generates a “winner-takes-all” distribution in which popular traits dominate, while anti-conformity generates a “humped” distribution in which traits of intermediate frequency are favoured. However, a more passive frequency-dependent “trimming”, in which traits are selectively ignored on the basis of their frequency, generates reasonable approximations to the power-law distribution. This frequency-dependent trimming may therefore be difficult to distinguish from genuine random copying using population-level data alone. Implications for the study of both human and nonhuman culture are discussed.

Keywords: Conformity; Cultural evolution; Cultural transmission; Frequency-dependent copying; Neutral drift; Power law; Random copying

1. Introduction

In the past few decades, researchers within the cultural evolution tradition (incorporating gene-culture coevolution/dual inheritance theory) have used concepts, tools and methods from evolutionary biology to explain various aspects of human cultural change (Boyd & Richerson, 1985, 2005; Cavalli-Sforza & Feldman, 1981; Durham, 1992; Henrich & McElreath, 2003; Laland, Kumm, & Feldman, 1995; Mesoudi, Whiten, & Laland, 2006; Shennan, 2002). This cross-disciplinary borrowing from the biological to the cultural sciences is justified by the observation that human culture constitutes a Darwinian evolutionary system of inheritance that acts in parallel to genetic evolution (Campbell, 1965; Mesoudi, Whiten, & Laland, 2004; Plotkin, 1994).

A significant advance in evolutionary biology came with the formal realisation that biological change can occur through the nonselective process of neutral genetic drift (Crow & Kimura, 1970; Kimura, 1983), in which changes in gene frequencies occur through chance alone. Since the very earliest cultural evolution models (Cavalli-Sforza & Feldman, 1973, 1981), drift has similarly been recognised as an important factor in explaining patterns of human cultural evolution. That is, changes in the frequencies of cultural traits that are selectively neutral will be determined by random events such as sampling error, just as for selectively neutral alleles.

In recent years, several studies have elaborated on this earlier work and proposed that the frequency distributions of various cultural traits — first names (Hahn & Bentley, 2003), dog breeds (Herzog, Bentley, & Hahn, 2004), pottery decorations (Bentley, Hahn, & Shennan, 2004; Neiman,

---

☆ AM was supported by a Mellon Foundation Fellowship. SJL was supported by the British Academy Centenary Research Project, *Lucy to Language*.

* Corresponding author. School of Biological and Chemical Sciences, Queen Mary, University of London, Mile End Road, London E1 2NS, UK.
E-mail address: am786@cam.ac.uk (A. Mesoudi).
URL: http://amesoudi.googlepages.com.

1090-5138/$ – see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.evolhumbehav.2008.07.005

1995), scientific paper citations (Simkin & Roychowdhury, 2003) and patent citations (Bentley et al., 2004) — can be explained using a simple model of *random copying* (Table 1), the cultural analogue of genetic drift. The random copying model assumes that while a small number of individuals engage in innovation (the generation of novel cultural traits, analogous to genetic mutation), the majority of individuals select another member of the population entirely at random and copy their cultural trait, with no intrinsic fitness differences between traits and no bias in the selection of cultural models. Through simulation, Bentley et al. (2004) showed that this random copying model generates a distinct "power-law" distribution, characterised as:

$$y = Cx^{-\alpha} \tag{1}$$

where y is the proportion of cultural traits that occur with frequency x in the population, and C and α are constants. This power-law distribution appears as a straight line with a slope of $-\alpha$ on a graph of proportion (y) vs. frequency (x) with both axes on logarithmic scales. This distribution indicates a small number of very popular traits and a large number of very rare traits in the population, such that the power-law distribution is described as "scale free", i.e., the distribution has the same shape for all ranges of frequencies. The observation that the aforementioned cultural datasets (first names, dog breeds, pottery decorations, scientific paper citations and patent citations) all exhibit this power-law distribution suggests that these particular cultural traits are all selectively neutral and copied at random.

Random copying can be contrasted with *independent decisions* (Bentley & Shennan, 2003), in which cultural traits are chosen at random from other members of the population but without reference to their frequency, such that each extant cultural trait has an equal probability of being adopted (Table 1). Bentley & Shennan (2003) showed that independent decisions are described not by a scale free power law but by an exponential decay function of the form:

$$y = Ce^{-\alpha x} \tag{2}$$

Under independent decisions, each trait has an equal probability of going extinct in every generation, resulting in a curved (i.e., non–scale-free) distribution on a log–log plot with fewer traits achieving high frequencies.

Just as in population genetics (e.g., Kreitman, 1996), the (neutral) model of random copying provides a powerful null hypothesis for cultural change (Bentley et al., 2004). However, we note that the assumption of random copying appears to be at odds with much evidence that actual human cultural transmission is often far from random. Cognitive psychologists have long known that people have a limited capacity for randomness (Baddeley, 1966; Baddelely, Emslie, Kolodny, & Duncan, 1998), while social psychologists (e.g., Asch, 1951; Bond, 1995) have repeatedly demonstrated that many people exhibit strong tendencies for conformity, i.e., adopting the most common behaviour/trait in the population. This latter finding is consistent with gene-culture coevolution models which suggest that conformity readily evolves whenever social learning itself evolves (Henrich & Boyd, 1998), as well as the observation that other cultural patterns show signatures of conformity, such as the diffusion rates of technological innovations (Henrich, 2001). In general, these empirical and theoretical studies suggest that cultural transmission is often *frequency-dependent*, i.e., traits are preferentially adopted on the basis of their frequency in the population beyond that expected by random copying.

These findings do not necessarily conflict with the random copying models cited previously, which make no specific predictions regarding the proximate, individual-level biases that govern cultural change. Perhaps some mix of conformity, anti-conformity and innovation combine to produce aggregate, population-level data that are indistinguishable from random copying. However, to our

Table 1
Description and consequences of different cultural learning rules simulated in the model

Learning rule		Description	Resulting frequency distribution
Random copying		Choose another member of the previous generation at random and copy their cultural trait	A scale-free power-law distribution, with a small number of very popular traits and a large number of very rare traits (Fig. 1)
Independent decisions		Choose a cultural trait from the previous generation at random and copy it, ignoring trait frequency	A non–scale-free exponential decay distribution in which fewer traits reach high frequencies (Fig. 1)
Frequency-dependent copying	Conformity	Adopt the most common trait in the previous generation with a probability c_p	A "winner-takes-all" distribution dominated by a very small number of very popular traits (Fig. 2A)
	Anti-conformity	Adopt the least common trait in the previous generation with a probability c_n	An irregular "humped" distribution, favouring traits of intermediate frequency (Fig. 2B)
Frequency-dependent trimming	Common-trait trimming	Copy randomly but ignore traits that have frequencies in the top T fraction of the frequency distribution	A power-law distribution with fewer traits with high frequencies (Fig. 3A)
	Rare-trait trimming	Copy randomly but ignore traits that have frequencies in the bottom B fraction of the frequency distribution	A power-law distribution with more very low frequency traits and more very high frequency traits (Fig. 3B)

knowledge, this claim has not yet been tested explicitly. Moreover, without a clear understanding of how different forms of frequency-dependent copying affect the power-law distribution, any observation that a particular cultural dataset deviates from the power-law distribution can only be explained in a post hoc fashion. Our aim here was to use agent-based simulations to explore more formally the effect on the power-law distribution of non–random-copying biases, in particular frequency-dependent biases such as conformity, given the evidence cited above that such biases appear often to be prevalent in actual instances of cultural transmission.

We define two kinds of frequency-dependent biases: frequency-dependent *copying* and frequency-dependent *trimming* (Table 1). Frequency-dependent copying describes the case where agents actively choose to copy a trait on the basis of its frequency in the population. Positive frequency-dependent copying, or *conformity*, is where agents adopt the most common trait in the population with a certain probability. Negative frequency-dependent copying, or *anti-conformity*, is where agents adopt the least common trait in the population with a certain probability. An alternative way of implementing frequency-dependent bias is by "trimming" the distribution of traits from which agents select models, excluding traits on the basis of their frequency in the population. For example, parents selecting a name for their child might selectively ignore very popular and/or very unpopular names in the population and choose at random from the remaining names. *Common-trait trimming* agents selectively ignore traits that have frequencies in some upper portion of the frequency distribution, while *rare-trait trimming* agents selectively ignore traits that have frequencies in some lower portion of the frequency distribution.

2. Methods

We constructed an agent-based model in C++ that replicates and extends the model presented by Bentley et al. (2004). The model is available for download from the Electronic Supplementary Materials in the form of an executable Windows-compatible (.exe) program, with which readers can replicate all of the results presented below. This program is freely distributable for educational purposes, and we particularly encourage its use as a teaching aid. Source code is also available on request from the first author. The following sections describe the model assumptions for each set of transmission biases.

2.1. Random copying/independent decisions

In the basic model, a population of N agents each exhibit a single cultural trait, represented by a single integer. Initially, each agent possesses a unique trait. In every subsequent generation, each agent either innovates with a probability μ ($0 \leq \mu \leq 1$), i.e., adopts a novel trait not possessed by any previous agent, or copies with a probability $1-\mu$. Copying is by default random: the agent selects another agent from the previous generation at random and adopts their trait. We additionally simulated independent decisions (Bentley & Shennan 2003), in which agents adopt a trait from the previous generation at random irrespective of the trait's frequency. Following Bentley et al. (2004), the basic model was run with parameters $N=250$ and $\mu=0.008$ for 1000 generations. Results shown below are the average data from five separate runs of the model and are displayed as log–log plots of the proportion of traits (y) falling into a particular frequency bin (x). Bin sizes increased logarithmically (e.g., 0–1, 1–2, 2–4, 4–8, 8–16, …) and frequencies were normalised according to bin size (e.g., total frequencies in the aforementioned bins were divided by 1, 1, 2, 4, 8, …), as described in Bentley & Shennan (2003:467).

2.2. Frequency-dependent copying

The basic model was extended by adding frequency-dependent copying, where agents preferentially adopt either the most common or the least common trait in the population. As before, agents innovate with probability μ and copy with probability $1-\mu$. Of the $1-\mu$ agents (on average) who copy, two additional parameters operate: c_p, which is the probability that an agent engages in positive frequency-dependent copying ("conformity"), i.e., adopts the most common trait from the previous generation; and c_n, which is the probability that an agent engages in negative frequency-dependent copying ("anti-conformity"), i.e., adopts the least common trait from the previous generation. When two or more traits have the joint highest or lowest frequency, one of these tied traits is chosen at random. The remaining copiers ($1-\mu-c_p-c_n$ on average) copy at random as described in the previous section (independent decisions were not implemented here). The restriction $(\mu+c_p+c_n)\leq 1$ ensured probabilities did not exceed 1.

2.3. Frequency-dependent trimming

Frequency-dependent trimming is where agents selectively ignore traits that are either very common (common-trait trimming) or very rare (rare-trait trimming), such that the frequency distribution of potential traits is trimmed. In our model, the $1-\mu$ (on average) copying agents copy at random as described above, except that they selectively ignore any trait that has a frequency in the top T fraction (common-trait trimming) or in the bottom B fraction (rare-trait trimming) of the total number of trait copies, including multiple copies of the same trait (where $T+B\leq 1$). Because each agent possesses a single copy of a trait, the total number of trait copies will always be equal to the total number of agents, N. Hence trimming agents copy the trait of a randomly selected member of the previous generation as long as $NB<F<N(1-T)$, where F is the frequency of that trait in the previous generation. For example, if $T=0.2$ and $B=0.1$ and there are 250 agents ($N=250$), there will be 250 trait

copies (one per agent) and agents will ignore all traits that have a frequency greater than $N(1-T)=200$ and less than $NB=25$, i.e., they only copy traits where $25<F<200$.

3. Results

3.1. Random copying/independent decisions

As expected, and replicating the results of Bentley et al. (2004), random copying generated a straight line on the log–log plot (Fig. 1) with a good fit to a power-law function ($r^2=0.996$). The value of $\alpha=1.572$ was similar to that found by Bentley et al. (2004) of $\alpha=1.52$, and additional analyses (not shown) confirmed that varying μ and N had the same effects as found by Bentley et al. (2004), i.e., α increases with $N\mu$, and the power-law relationship breaks down for large values of μ. As predicted by Bentley and Shennan (2003), the frequency distribution resulting from independent decisions fit an exponential decay function ($r^2=0.960$). Compared to random copying, independent decisions resulted in a narrower range of trait frequencies, with fewer traits having very low or very high frequencies at extinction. This is as expected, given that each trait has a constant and equal probability of going extinct during each generation.

3.2. Frequency-dependent copying

Both conformity (Fig. 2A) and anti-conformity (Fig. 2B) disrupt the power-law distribution quite markedly. Conformity ($c_p>0$) makes already-popular traits even more popular, thus increasing the absolute frequency of popular traits as indicated in the far rightmost points in Fig. 2A. This resembles a "winner-take-all" distribution (Albert & Barabasi, 2002:77; Frank & Cook, 1995), in which a single trait dominates the population. The remaining data points under conformity show a good fit to an exponential decay function (e.g., $r^2=0.985$ for $c_p=0.6$), similar to independent decisions (Fig. 1). Here, because high-frequency traits dominate the frequency distribution, the range of possible frequencies that other traits can take is much reduced (in the extreme "winner-take-all" situation, minority traits would all have frequencies of 1). Consequently, a trait's frequency is less likely to affect its chances of being copied, hence the resemblance to independent decisions in which frequency information is ignored.

Anti-conformity ($c_n>0$) results in an irregular "humped" distribution (Fig. 2B), such that it cannot be described by either a power or exponential function. Anti-conformity simultaneously reduces the frequency of very popular traits (because common traits are less likely to be copied) and increases the probability that a rare trait becomes popular (because rare traits are more likely to be copied). Hence popular traits in the upper end of the frequency distribution suffer from anti-conformity, and rare traits in the lower end of the frequency distribution benefit from anti-conformity; added to their existing frequencies, this produces the distinctive "humped" distribution in which traits of intermediate frequency are favoured.

3.3. Frequency-dependent trimming

Common-trait trimming ($T>0$; Fig. 3A) results in a straight line power-law distribution ($r^2=0.989$, $\alpha=1.601$) that

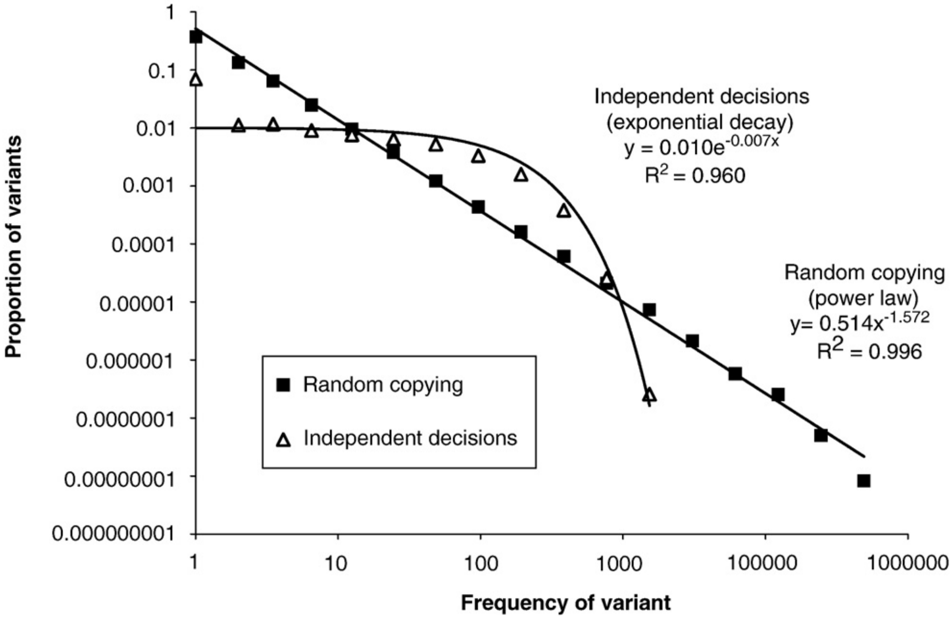

Fig. 1. A log–log plot showing the resulting frequency distribution of random copying, where agents choose a model from the previous generation at random and adopt their trait, and independent decisions, where agents choose a trait from the previous generation at random ignoring trait frequency. Parameters: $\mu=0.008$, $N=250$, 1000 generations, five runs. Random copying generates a straight line on the log–log plot, indicating a scale-free power-law distribution ($r^2=0.996$), while independent decisions are better described by an exponential decay function ($r^2=0.960$).

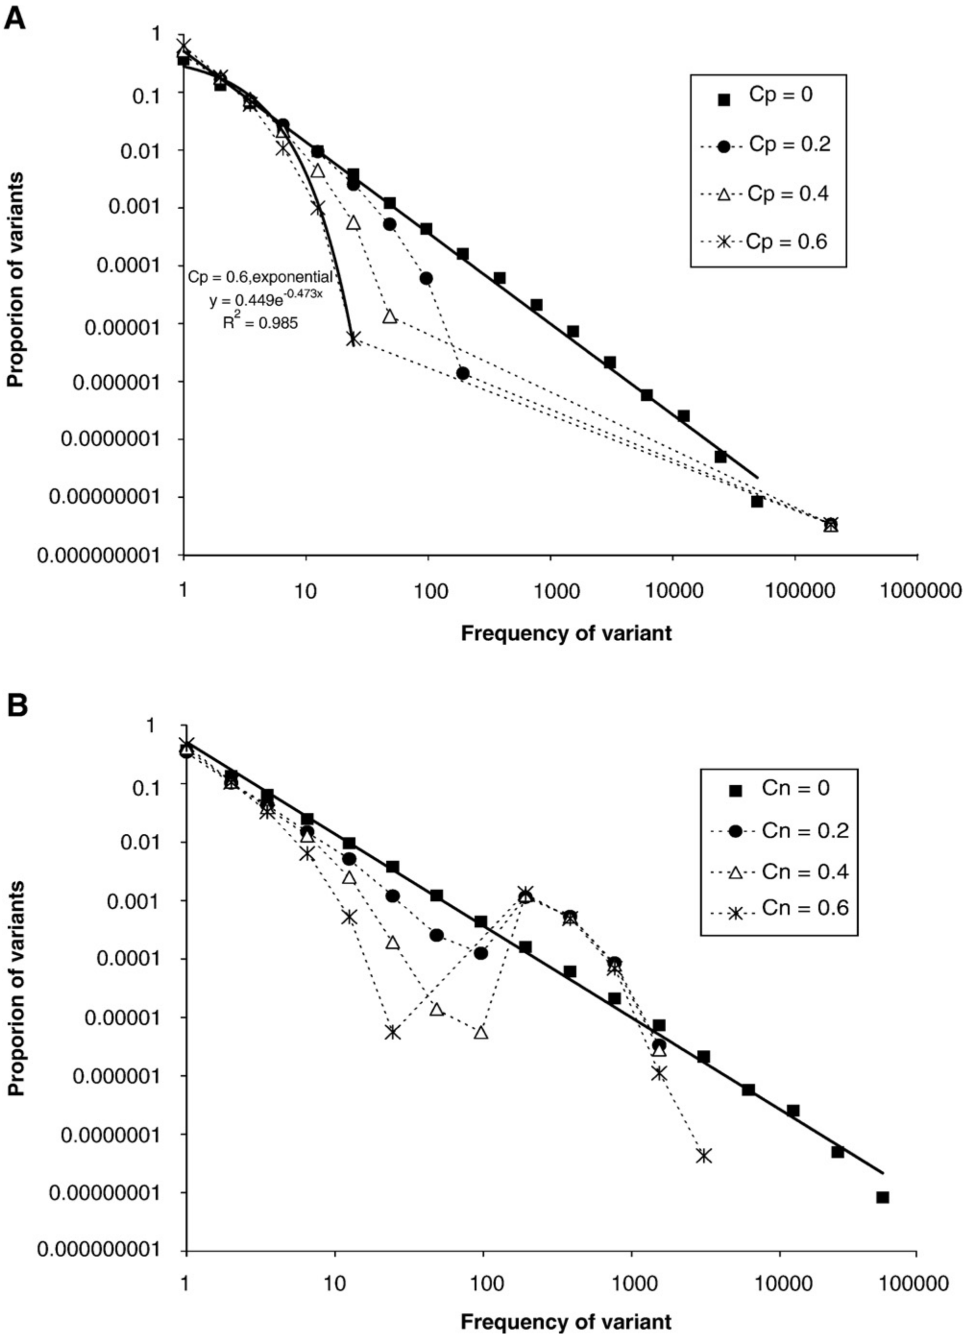

Fig. 2. The effect on the power-law distribution of frequency-dependent copying, showing (A) conformity, where agents select the most common trait with probability c_p, and (B) anti-conformity, where agents select the least common trait with probability c_n. In both cases, the random copying power-law line (c_p=c_n=0) is shown for comparison. All other data points are connected with straight dotted lines to indicate the shape of the distribution. Fig. 2A shows an exponential decay function fitted to all data points of the c_p=0.6 run except the highest frequency data point. All other parameters (N, generations, μ, runs) as for Fig. 1.

is almost identical to random copying (Fig. 1) except that a slightly lower proportion of traits achieve high frequencies. This is because trimming popular traits reduces the maximum frequency that the most common traits can reach, thus reducing the absolute frequency of common traits. Random copying then operates on the other traits to generate a power law.

Rare-trait trimming (B>0; Fig. 3B) also results in a good fit to a power-law distribution (r^2=0.975, α=1.521), with the exception of two data points: there are more low-frequency traits (leftmost data point) and more high-frequency traits (rightmost data point) relative to random copying. The former is caused by mutation introducing novel traits that then go extinct in a single generation (or very few

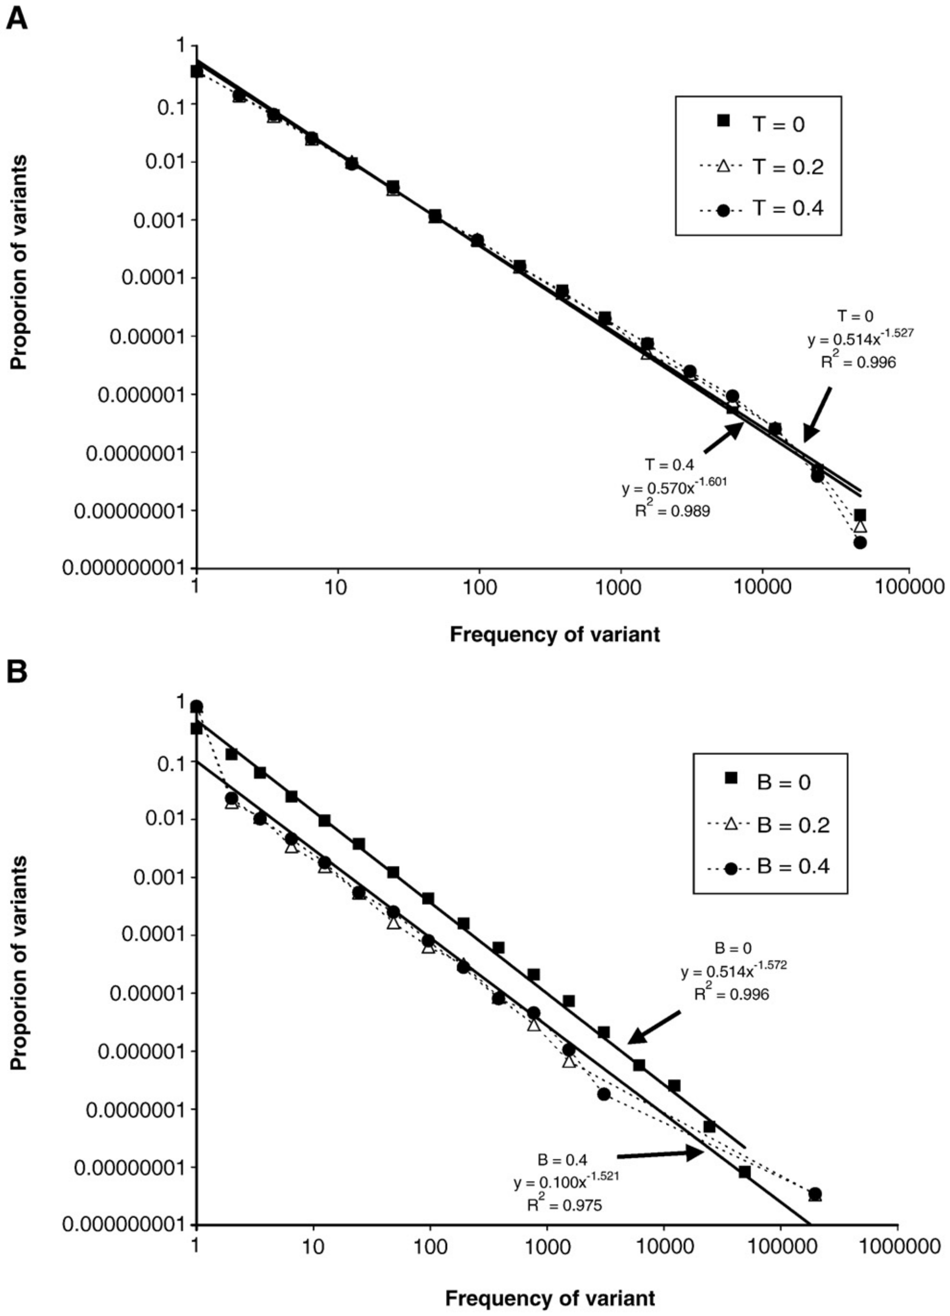

Fig. 3. The effect on the power-law distribution of frequency-dependent trimming, in which agents selectively ignore traits that have frequencies within the (A) top *T* fraction, i.e., common-trait trimming, or (B) bottom *B* fraction, i.e., rare-trait trimming, of the total number of trait copies in the previous generation, copying at random from the remaining traits. In both cases, dotted lines connect data points for $T>0$ and $B>0$. Power functions are fitted to the data points where $T=0.4$ and $B=0.4$, along with the random copying ($T=B=0$) line for comparison. All other parameters as for Fig. 1.

generations) due to rare-trait trimming. The latter is caused by the conformity-like effect of ignoring less-common traits, such that popular traits get even more popular (as in Fig. 2A).

4. Discussion

Previous studies have shown that random copying results in a distinct power-law distribution of cultural traits, and that this distribution characterises several real-life cultural datasets, such as first names, pottery decorations and patent citations (Bentley et al., 2004). Here we used agent-based simulations to explore the effect of frequency-dependent biases on the power-law distribution, given experimental evidence and theoretical models that suggest actual cultural transmission better resembles frequency-dependent biases such as conformity than random copying.

We found that actively selecting traits based on their frequency — either selecting the most popular (conformity) or unpopular (anti-conformity) traits — strongly disrupted the power-law distribution. Conformity resulted in a distribution dominated by a small number of extremely popular traits, resembling a "winner-take-all" distribution, and in which the remaining lower-frequency traits are described by the exponential decay function of independent decisions rather than the power-law rule. Anti-conformity generated an irregular humped distribution in which traits of intermediate frequency are favoured. Generally, the power-law distributions exhibited by the actual cultural datasets analysed to date (e.g., in Bentley et al., 2004) do not appear to be consistent with this form of frequency-dependent copying. However, future studies might reveal these distributions in other cultural datasets, such as the "winner-take-all" cultural markets of the entertainment and sports industries (Frank & Cook, 1995).

In contrast, the more passive frequency-dependent trimming — in which popular or unpopular traits are systematically ignored — generated power-law distributions quite similar to those generated by unbiased random copying, with minor deviations observed at very low and/or very high frequency ranges. Re-examining previously published analyses, we might speculate that female first names (Fig. 4 in Bentley et al., 2004) show evidence of popular trait trimming (our Fig. 3A), given that the most popular female names appear to be less popular than expected under the random copying model. Generally, it may be difficult using real-life datasets to distinguish between purely unbiased random copying and the frequency-dependent trimming that we simulated here due to the potential problem of equifinality: different individual-level assumptions generate identical population-level effects. This problem might be addressed by collecting behavioural data, either from the field or from experiments, that would reveal the cognitive biases that people actually use when deciding which cultural trait to adopt.

The model presented here is an extremely abstract model of cultural change, making such simplifying assumptions as discrete cultural traits, non-overlapping generations, homogenous agent behaviour and a lack of selection. While much real-world cultural change is undoubtedly far more complex, simple models allow us to better understand basic processes of cultural change by separating out the effects of these processes and generating clear predictions that can be tested in quantitative cultural datasets. Indeed, the datasets chosen for previous tests of the random-copying model (e.g., first names, dog breeds, patent citations; Bentley et al., 2004) were chosen precisely because they contained simple, discrete traits, thus allowing the (successful) test of the predictions of the abstract model.

Finally, because of the simplicity and abstractness of our model, it may be fruitful to test its predictions in nonhuman cultural datasets. Many species have been shown to exhibit socially learned patterns of behaviour (i.e., cultural traditions), including species of birds, fish, cetaceans and primates (Whiten et al., 1999; Lycett, Collard & McGrew, 2007; Laland, 2008). Where these socially transmitted behaviours take the form of discrete traits, such as in birdsong (Lachlan & Slater, 2003), similar analyses of frequency distributions to those simulated here may help to reveal the particular learning strategies (Laland, 2004) that are being employed.

Acknowledgments

We thank Alex Bentley, Dan Fessler and two anonymous reviewers for valuable comments on the manuscript.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.evolhumbehav.2008.07.005.

References

Albert, R., & Barabasi, A. -L. (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics*, *74,* 47–97.

Asch, S. E. (1951). Effects of group pressure on the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership and men*. Pittsburgh, PA: Carnegie; p. 177–190.

Baddeley, A. D. (1966). The capacity for generating information by randomization. *Quarterly Journal of Experimental Psychology*, *18,* 119–129.

Baddeley, A. D., Emslie, H., Kolodny, J., & Duncan, J. (1998). Random generation and the executive control of working memory. *Quarterly Journal of Experimental Psychology*, *51A,* 819–852.

Bentley, R. A., Hahn, M. W., & Shennan, S. J. (2004). Random drift and culture change. *Proceedings of the Royal Society B*, *271,* 1443–1450.

Bentley, R. A., & Shennan, S. J. (2003). Cultural evolution and stochastic network growth. *American Antiquity*, *68,* 459–485.

Bond, R. (1995). Group size and conformity. *Group Processes and Intergroup Relations*, *8,* 331–354.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.

Boyd, R., & Richerson, P. J. (2005). *The origin and evolution of cultures*. Oxford: Oxford University Press.

Campbell, D. T. (1965). Variation and selective retention in socio-cultural evolution. In H.R. Barringer, G.I. Blankst en, R.W. Mack (Eds.), *Social change in developing areas*. Cambridge, MA: Schenkman; p. 19–49.

Cavalli-Sforza, L. L., & Feldman, M. W. (1973). Models for cultural inheritance: I. Group mean and within-group variation. *Theoretical Population Biology*, *4,* 42–55.

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution*. Princeton: Princeton University Press.

Crow, J. F., & Kimura, M. (1970). *An introduction to population genetics theory*. New York: Harper & Row.

Durham, W. H. (1992). *Coevolution*. Stanford: Stanford University Press.

Frank, R. H., & Cook, P. J. (1995). *The winner-take-all society*. New York: Free Press.

Hahn, M. W., & Bentley, R. A. (2003). Drift as a mechanism for cultural change. *Proceedings of the Royal Society B*, *270,* S120–S123.

Henrich, J. (2001). Cultural transmission and the diffusion of innovations. *American Anthropologist*, *103,* 992–1013.

Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, *19,* 215–241.

Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, *12,* 123–135.

Herzog, H. A., Bentley, R. A., & Hahn, M. W. (2004). Random drift and large shifts in popularity of dog breeds. *Proceedings of the Royal Society B*, *271,* S353–S356.

Kimura, M. (1983). *The neutral theory of molecular evolution*. Cambridge: Cambridge University Press.

Kreitman, M. (1996). The neutral theory is dead. Long live the neutral theory. *BioEssays*, *18,* 678–683.

Lachlan, R. F., & Slater, P. J. B. (2003). Song learning by chaffinches. *Animal Behaviour*, *65,* 957–969.

Laland, K. N. (2004). Social learning strategies. *Learning and Behavior*, *32,* 4–14.

Laland, K. N. (2008). Animal cultures. *Current Biology*, *18,* R366–R370.

Laland, K. N., Kumm, J., & Feldman, M. W. (1995). Gene-culture coevolutionary theory. *Current Anthropology*, *36,* 131–156.

Lycett, S. J., Collard, M., & McGrew, W. C. (2007). Phylogenetic analyses of behavior support existence of culture among wild chimpanzees. *Proceedings of the National Academy of Sciences U S A*, *104,* 17588–17592.

Mesoudi, A., Whiten, A., & Laland, K. N. (2004). Is human cultural evolution Darwinian? Evidence reviewed from the perspective of *The Origin of Species*. *Evolution*, *58,* 1–11.

Mesoudi, A., Whiten, A., & Laland, K. N. (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, *29,* 329–383.

Neiman, F. D. (1995). Stylistic variation in evolutionary perspective. *American Antiquity*, *60,* 7–36.

Plotkin, H. (1994). *Darwin machines and the nature of knowledge*. London: Penguin.

Shennan, S. J. (2002). *Genes, memes and human history*. London: Thames & Hudson.

Simkin, M. V., & Roychowdhury, V. P. (2003). Read before you cite! *Complex Systems*, *14,* 269–274.

Whiten, A., Goodall, J., McGrew, W. C., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C. E. G., Wrangham, R. W., & Boesch, C. (1999). Cultures in chimpanzees. *Nature*, *399,* 682–685.