Theoretical Population Biology 76 (2009) 59–67

Contents lists available at ScienceDirect

Theoretical Population Biology

journal homepage: www.elsevier.com/locate/tpb

An investigation of the relationship between innovation and cultural diversity

Anne Kandler[a,*], Kevin N. Laland[b]

[a] AHRC CECD, Institute of Archaeology, University College London, 31-34 Gordon Square, London WC1H 0PY, United Kingdom

[b] School of Biology, St. Andrews University, Bute Building, Queens Terrace, St. Andrews, Fife KY16 9TS, United Kingdom

ARTICLE INFO

Article history:
Received 13 November 2008
Available online 22 April 2009

Keywords:
Innovation
Invention
Cultural transmission
Cultural diversity
Conformity
Reaction–diffusion model

ABSTRACT

In this paper we apply reaction–diffusion models to explore the relationship between the rate of behavioural innovation and the level of cultural diversity. We investigate how both independent invention and the modification and refinement of established innovations impact on cultural dynamics and diversity. Further, we analyse these relationships in the presence of biases in cultural learning and find that the introduction of new variants typically increases cultural diversity substantially in the short term, but may decrease long-term diversity. Independent invention generally supports higher levels of cultural diversity than refinement. Repeated patterns of innovation through refinement generate characteristic oscillating trends in diversity, with increasing trends towards greater average diversity observed for medium but not low innovation rates. Conformity weakens the relationship between innovation and diversity. The level of cultural diversity, and pattern of temporal dynamics, potentially provide clues as to the underlying process, which can be used to interpret empirical data.

1. Introduction

Recent decades have witnessed considerable interest in the mathematical modelling of the dynamics of cultural change. A number of researchers have adapted the methods of theoretical population genetics to study the dynamics of cultural change through time, as well as the co-evolution of genes and culture (Boyd and Richerson, 1985; Cavalli-Sforza and Feldman, 1981; Feldman and Cavalli-Sforza, 1976; Feldman and Laland, 1996; Henrich and McElreath, 2003; Richerson and Boyd, 2005; Enquist et al., 2007). Such models include analysis of a variety of other forms of cultural transmission (Boyd and Richerson, 1985; Cavalli-Sforza and Feldman, 1981; Feldman and Cavalli-Sforza, 1976), application of neutral genetic drift models to study the evolution of cultural traits (Bentley et al., 2004), and of phylogenetic methods to reconstruct the history of diverse cultural traits (Gray and Jordan, 2000; Holden, 2002; Holden and Mace, 2003; O'Brien and Lyman, 2003).

While the mathematical modelling of culture is a rapidly developing and highly productive discipline, one topic that has received comparatively little attention is the relationship between innovation and cultural diversity. Nor has this topic been addressed by the diffusion research tradition (Rogers, 2003), which explores how, why, and at what rate new ideas and technology spread through cultures. This is paradoxical for two reasons: first, at least since Boas, widely regarded as the father of American anthropology (e.g. Boas (1911)), understanding why people are different (i.e. the causes of cultural diversity) has been recognized as one of the major objectives of the social sciences. Second, it would be difficult to find another topic in anthropology and archaeology that has played as important a role as innovation in framing arguments about why and how human behaviour changes (O'Brien, in press). Nineteenth century ethnologists, such as Tylor (1871) and Morgan (1877), viewed the production of novelties as the evolutionary driver that propels cultures up a hierarchy of cultural complexity. Innovation was equally important in the work of later cultural evolutionists such as Steward (1955) and White (1959), who regarded it as a key component that a group needs to meet the challenges of its physical and social environment. American culture historians of the twentieth century routinely looked to diffusion and trade as a source of innovations, and hence of cultural change and diversity (e.g. (Ford, 1969)). Even contemporary debates over the legitimacy of cultural phylogenetics (e.g. Borgerhoff Mulder et al. (2006)) hang in part on the extent to which cultural diversity can be explained by divergent tradition, independent invention, or the diffusion of innovation. This issue is also central to attempts to comprehend the processes underlying the human capacity for cumulative culture (Ghirlanda and Enquist, 2007). Accordingly, understanding of the relationship between innovation and diversity would engender widespread interest.

To date, empirical analysis of this question using the archaeological or historical record has been hindered by the difficulty that often only one or two different time periods are considered in a given study, or the fact that the study is focused on the comparison of the temporal evolution of two variants (e.g. the replacement

* Corresponding author.
E-mail address: a.kandler@ucl.ac.uk (A. Kandler).

0040-5809/$ – see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.tpb.2009.04.004

of horses by tractors in agriculture, (Mattingly, 1987)). Here we develop a mathematical analysis of the relationship between innovation and diversity, hoping that this will provide a theoretical framework that will encourage further empirical study. Indeed, our analyses not only shed light on the processes of innovation and cultural change, but also provide methods by which the mechanisms of innovation can be inferred from the degree of diversity.[1]

Our study is designed to address the following issues. First, while it may be intuitive to assume that innovation will generate cultural diversity in the short term, it is far from clear that innovation will lead to enhanced cultural diversity in the longer term, or at equilibrium. In principle, technological innovations could merely displace established technology, without increasing long-term diversity. Accordingly, we ask under what circumstances innovation promotes cultural diversity. Second, we explore whether there is a threshold rate of innovation that is necessary to support the accumulation of cultural diversity. Third, while many innovations can be characterized as refinement or modification of an established variant (Basalla, 1988), other cases might be perceived as entirely independent invention. The latter category include serendipitous or accidental discoveries, for which X-rays, electromagnetism, ozone, photography, dynamite, the gramophone, vaccination, radioactivity, classical conditioning, and penicillin are all examples (Simonton, 1995). We investigate whether these two alternative forms of innovation may impact differentially on cultural diversity. Fourth, in the light of this analysis, we consider whether it might be possible to draw inferences about the innovation mechanism (independent invention versus refinement) on the basis of the observed level of cultural diversity. Fifth, we explore how the relationship between innovation and diversity will be affected by biases in cultural transmission. Some forms of social learning (most obviously, conformist transmission, where individuals preferentially adopt the majority trait) may counteract any positive effect of innovation on cultural diversity.

We apply reaction–diffusion models to investigate how both independent invention and the modification and refinement of established innovations affect cultural diversity, building on earlier studies of language competition (e.g. Kandler and Steele (2008) and Pinasco and Romanelli (2006)), cultural hitchhiking (Ackland et al., 2007), and prestige bias (Ihara, 2008), that exploit similar methods. The models comprise reaction, diffusion and competition components, which collectively are well-suited to capture aspects of the spread of an innovation in a finite population. Using established population ecology metrics, we propose a number of distinct measures of cultural diversity (see Appendix B), illustrating how each is affected by the innovation rate. We begin by assuming that social learning is unbiased (acquired in proportion to the frequency of the variants in the population), but go on to investigate the impact of both frequency-independent and frequency-dependent cultural transmission biases.

2. The model

A detailed mathematical description of the model can be found in Appendix A; here we present an accessible summary. We develop and analyse a model that describes the spread dynamics of n competing (mutually exclusive) variants of a specific cultural trait within a population. Such variants might represent alternative beneficial subsistence techniques, technologies, religious beliefs, or languages. Obviously cumulative increases in "diversity" may occur in cases where innovations do not compete, and our conclusions regarding the relationship between innovation and diversity are restricted to competing cultural variants. Using a continuous differential equation based approach, we determine the temporal and spatial changes in the frequencies of n variants, denoted by $u_1, u_2, \ldots, u_n$. Individuals can adopt only one variant at a time. (We have found that analysis of situations where individuals are allowed to adopt more than one variant at the same time leads to similar results.) We assume a constant and homogeneous environment, which we model through temporally and spatially constant model parameters. Although the model possesses a deterministic nature and therefore ignores drastic and rare events, its results can be regarded as null hypotheses. Failures in the predictions of the model can be understood as evidences of the significant impact of rare changes on the competition dynamic (Ackland et al., 2007).[2]

Variant frequency change is determined by two main components—diffusion and growth. The diffusion component models the spread of an innovation from a specific location in space as a random walk, with density-dependent mixing, equivalent to the random spread of an innovation through direct contact between individuals. The propensities of variants to spread out in space vary and are described by the 'diffusion coefficients' (denoted by d_i). The reaction term describes the increase in frequency of each variant amongst naive individuals (individuals that have not yet adopted the trait), according to a specific 'growth' parameter, where a conventional logistic growth is assumed, as well as representing the effects of competition. Initially we assume that unbiased social learning (or asocial adoption) underpins the adoption process, but later we consider various kinds of cultural transmission bias. Variants differ in their growth propensities, denoted by a_i, depending on the benefits each conveys to its adopters. In addition, growth is influenced by individuals that have already adopted a variant switching to an alternative variant. Variants hinder each other in growth, with the degree of hindrance determined by variant frequencies and 'switch coefficients', denoted c_{ij}, specifying the proportion of adopters of variant i that switch to adopt variant j. The more beneficial a variant, the less likely are its adopters to switch and at the same time the more likely it is the preferred target of switching, which restricts the growth of competing variants. Variants need not be present at the beginning ($t = 0$) of the analysis, and can be invented at later time points t_k. The time between two inventions is modelled by an exponential distributed variable $\tau \sim \exp(\lambda)$, with innovation rate λ.

We consider two classes of innovation, *modification of an existing variant* and *independent invention*, and represent these in the model through choice of appropriate values of the intrinsic growth rate a_i and competition coefficient c_{ij} parameters. While we assume that modification can result in both an improvement and a worsening of the existing variant, here we focus on improvement. Improvement implies that the newly invented variant offers adopters benefits over the original, in which case we assume it will be preferred to the original, and that individuals who have already adopted other variants will be more likely switch to the improved variant than the original. These expectations can be implemented by (1) assuming a higher intrinsic growth coefficient for the improved variant over the original, and (2) imposing the constraints $c_{IO} < c_{OI}$ and $c_{KO} < c_{KI}$, where c_{IO} indicates how much the original variant hinders the improved variant in growth, c_{OI} describes the reverse relation, c_{KI} describes the competition between the improved variant and the other

[1] Currently there is no well established method for measuring cultural trait diversity, nor even a well characterized notion of a cultural trait. It is apparent that the specific definitions and methods deployed may influence the absolute level of diversity observed. We nonetheless assume the existence of a framework regarding how the cultural variants are defined and measured, and do not expect such definitional factors to affect our qualitative conclusions.

[2] Obviously, such failings can also result from inadequate assumptions that fail to capture the underlying basic processes.

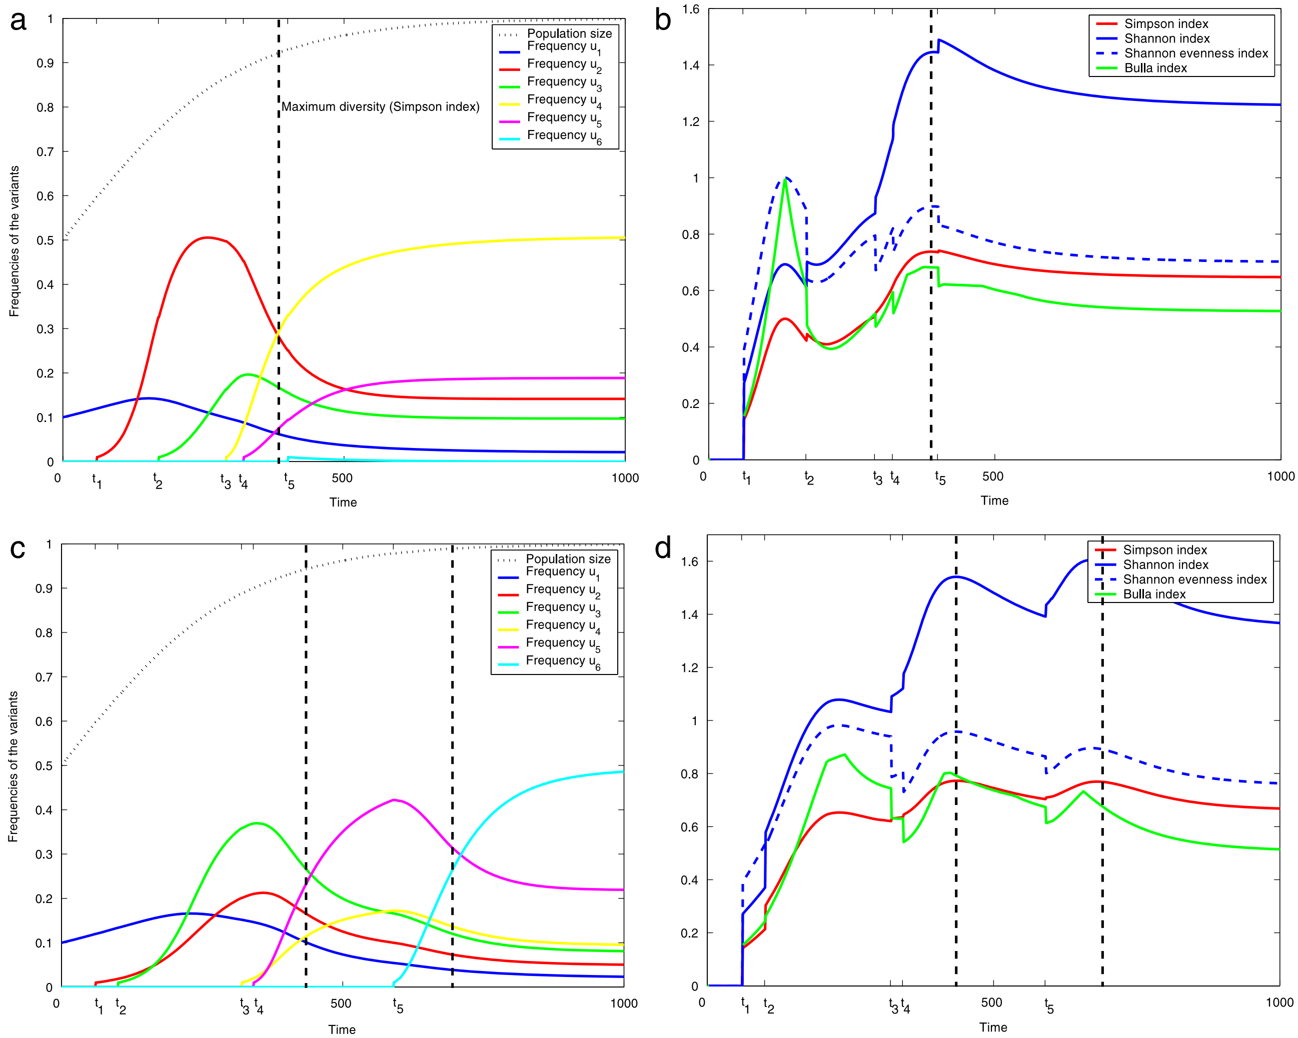

Fig. 1. (a) and (c) Representative frequency time course for six competing variants, with (b) and (d) showing the corresponding diversity measures. (a) and (b) represent modifications that may or may not be improvements (and also to some degree reflects the behaviour of independent inventions), while (c) and (d) represent the case where all newly invented variants are improvements on pre-existing variants. The first variant is invented at $t = 0$ and variant k is invented at random time points $t_k = \sum_{i=1}^{k} \tau_i$ with $\tau_i \sim \exp(0.01)$. The innovation rate $\lambda = 0.01$ leads on average to six innovations during the period [0, 600].

present variants, and c_{KO} the competition between the original variant and the remaining variants.[3] As a result of the described procedure the invented variants can be ordered hierarchically in terms of the benefits they offer to their adopters. In the case of *independent invention*, we assume that a variant is invented without knowledge of established variants, and represent this through allotting parameter values that are randomly selected relative to the variants present in the population.

In our analyses, we consider both the stable equilibria, that is, the long-term ratios of the competing variants, and the temporal dynamics in the approach to equilibria. We show that the stable equilibria are coexistence states, here denoted by $(u_1^*, u_2^*, \ldots, u_n^*)$. Coexistence means that all variants are present in the population, although, depending on the specific circumstances, they may have a negligibly small frequency, and in a stochastic system might plausibly be lost. Interestingly, due to the involved spatial diffusion process the obtained equilibria are spatially uniform,[4] which means the variant's frequency distribution is constant over the entire space. Therefore we focus on the description of the competition dynamic over time, rather than space. However, the use of a spatially explicit model greatly enhances the biological and cultural plausibility of our analysis: independent inventions are not all devised in the same place, but rather are distributed and diffuse through space, while competition between variants depends on local (rather than global) variant frequencies.

Quantitative measures of cultural diversity were developed through borrowing related biodiversity, species richness and evenness measures from biology (see Appendix B). We draw on four established metrics: (i) the Simpson and (ii) Shannon indices, which incorporate variant richness and variant evenness in the evaluation of diversity, (iii) the Shannon evenness index, which considers variant evenness only, and (iv) a diversity measure introduced by Bulla (1994) that is designed to be sensitive to the presence of rare variants.

3 Note that we assume that variants that offer benefits to the user are more likely to be retained than less effective variants. While we acknowledge the possibility that superficially attractive variants could be adopted, but then subsequently be found to be inadequate and rejected, here we prefer to concentrate on the more general case.

4 Spatially dependent equilibria can be obtained by, for instance, assuming a heterogeneous environment where the level of benefit of the variants varies in space or the domain shows a certain non-convex pattern (e.g. a connected island domain).

3. Results

3.1. Innovation and short-term dynamics

Fig. 1 illustrates the change in frequencies over time, and the corresponding impact on diversity, of six competing variants, where variants 2–6 are each a modification of an already existing variant.[5] In order to illustrate the long-term effects of innovation it is convenient to distinguish between two time periods—*an innovation period* $[0, t_{inv}]$ during which innovation occurs and *a stabilization period* $[t_{inv}, t_{stab}]$ where no innovation occurs and the stable long-term frequencies are reached. After stabilization, variant frequencies settle down to the stable coexistence state $(u_1^*, u_2^*, \ldots, u_6^*)$, although it is apparent from Fig. 1 that some variants have a negligible small frequency at equilibrium. In (a) and (b), not all modifications are improvements. Here, variant 4 is the most beneficial variant and establishes itself at the highest frequency at equilibrium. However, the time course of variant 2's frequency illustrates that other variants can increase to equivalent frequencies in the short term. With unbiased social learning the exact times t_k and order of the innovations' inception have no influence on the stable equilibrium, although they do change the course of the competition dynamics.

Fig. 1b shows how the spread of the innovations affects cultural diversity, as measured by our four indices. Reassuringly, all indices produce the same general patterns over time, although their absolute values and micro-dynamics vary considerably. It is immediately apparent that, contrary to the widespread expectation that innovation will inevitably increase diversity, in the short term, cultural diversity does not always increase with innovation, and all diversity measures show marked decreases at some juncture. Indeed, some diversity indices (Shannon evenness index, Bulla index) typically decrease with innovation.

Prior to the invention of the second variant (t_1) all diversity measures are zero, since only one variant is present, and at this juncture an immediate increase in all indices is observable. However, from invention points $t_2, \ldots, t_5$ the indices behave differently. The Shannon evenness index and the Bulla index evaluate only the evenness of the frequency distribution and so typically decrease at every invention point $t_2, \ldots, t_5$ since evenness usually decreases by introducing a new variant at low frequency. In contrast, the other indices show an increase in diversity at invention points, because they incorporate the number of variants, as well as evenness. The indices reach their maxima when variant 2 and variant 4, the most advantageous variants in the sample, have similar frequencies, a pattern that appears general, since this constellation represents the most even state of the frequency distribution. However, from this point all diversity measures decrease, as the dominant variant gradually displaces its competitors. The presence of a dominant variant, by which we mean an innovation that conveys a significantly higher benefit to adopters than other variants, almost always reduces cultural diversity in the longer term.

In the short term, the temporal dynamics for independent invention resembles that for randomly generated modification (Fig. 1a and b). Indeed, since Fig. 1a depicts a situation in which newly invented variants can be both superior and inferior to existing variants, the principal difference between this kind of innovation through modification and independent invention is merely the magnitude of the differences in parameter values between variants. It is only when innovation through modification repeatedly generates improvements that we witness qualitative differences in the patterns of diversity generated by the two types of innovation.

Fig. 1c and d illustrate the situation where all newly invented variants are improvements on pre-existing variants. Here the dynamics exhibit predictable patterns that are highly dependent on the innovation rate. The impact of innovation on long-term diversity is variable across indices, with those measures solely dependent on evenness less affected by new variants in the longer term than other measures. We have established that provided there is sufficient average time between innovations for the most beneficial variants to reach high frequency, a cycling pattern of diversity will emerge. The local maxima of the diversity indices are reached when the two most beneficial variants have similar frequencies (dashed lines).

Fig. 1d illustrates two primary processes that underpin the relationship between innovation and diversity. Discounting evenness measures, the introduction of new variants typically increases diversity substantially in the short term, as the new fit variant increases in frequency; however, once the most beneficial variant becomes dominant, the periods following inventions typically witness reductions in diversity, as the fittest variant out-competes the alternatives. If the time between innovations is comparatively short (a high innovation rate), the dynamics are dominated by the first of these processes, and diversity increases steadily over time. At the other extreme, long periods between innovations (low innovation rate) allow the second process to dominate, and we witness regular cycles in diversity, but with no increase, or only modest increases, in long-term diversity. Between these extremes we witness cycling in diversity associated with a saw-tooth-like increase in diversity with each innovation. We propose that the detection of these characteristic patterns of diversity change in the historical or archaeological record may potentially allow researchers to draw inferences about both the rate of innovation and the extent to which innovations are refinements.

3.2. Innovation and long-term dynamics

We now consider the relationship between innovation rate and cultural diversity at equilibrium. We assumed that the time between two innovations is determined by an exponential distributed variable $\tau \sim \exp(\lambda)$, where the innovation rate λ models how frequent innovations occur. In our model set-up it means that on average t_{inv}/λ variants are invented during the period $[0, t_{inv}]$. Fig. 2a shows the relation between the average level of cultural diversity and innovation rate λ where there is improvement of an existing variant, and Fig. 2b where there is independent invention. The average level is obtained by 10,000 simulations[6] for every value of λ. In both cases we obtain an almost linear relationship: the higher the rate of innovation the higher the cultural diversity. This result confirms the intuitive expectation that an increase in the number of variants is associated with an increase in cultural diversity, although the absolute magnitude of the increase in diversity with innovation rate is frequently

[5] As stated before, our model leads to spatially constant equilibria. Therefore the following figures illustrate the time course of the competition dynamic at a fixed spatial point and for the sake of simplification we call the times t_k invention times of the kth variant whereas in the spatially explicit formulation of the model they are the arrival times of variant k at location x.

[6] The simulations are carried out as follows. In the case of modification, the rank of every newly invented variant in the already existing order of variants is chosen randomly (that means whether the new variant is an improvement or worsening compared to every already existing variant is assigned at random), while the growth and competition coefficients are determined accordingly to the above procedure. In contrast, in the case of independent invention the growth and competition coefficients are modelled by uniform distributed variables. However, importantly, the possible parameter ranges in both cases are the same.

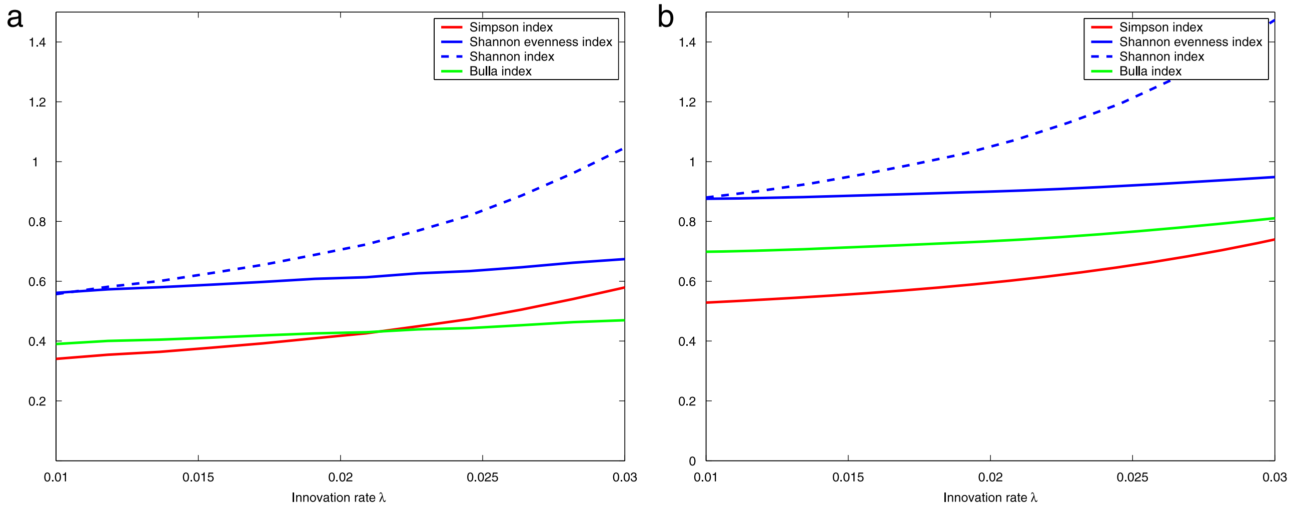

Fig. 2. Relation between the average level of cultural diversity and innovation rate λ for (a) innovations through improvement, and (b) independent invention. For each of 40 values of λ from 0.01 to 0.03, a random number of invented variants were selected, with randomly chosen degrees of advantageousness. Repeating this procedure, and averaging over 10,000 runs, leads to an average level of cultural diversity for the chosen innovation rate. Similar results are obtained when the adoption rates a_i and competition terms c_{ij} are negatively correlated.

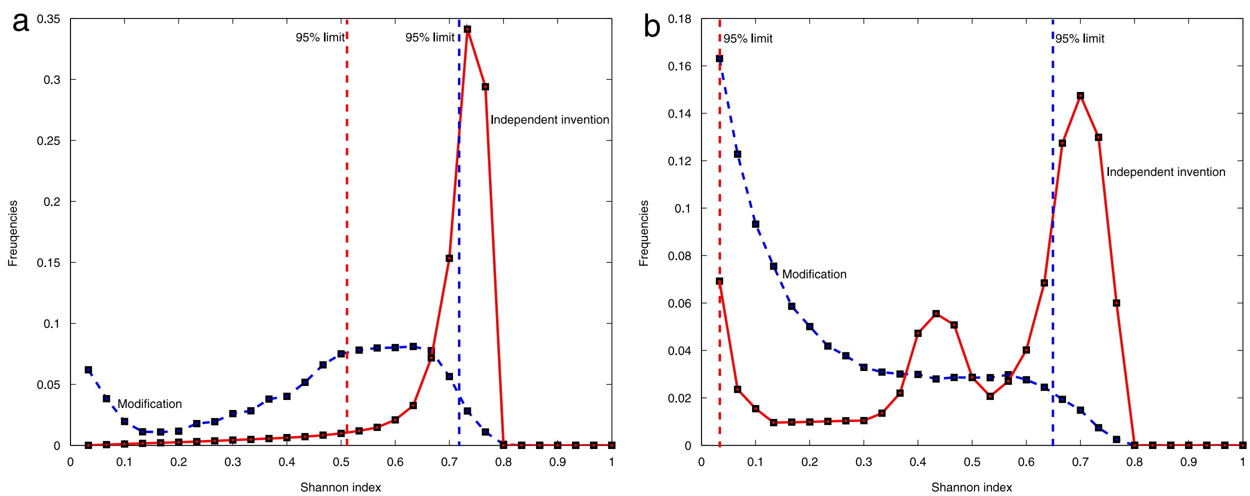

Fig. 3. Frequency distributions of the Shannon index for innovation through modification (dashed lines) and independent invention (solid lines), (a) without conformity, and (b) with strong conformity. We assume a situation where four variants are invented during a given time period. With strong conformity, independent invention favours situations where different numbers of variants have similar frequencies.

surprisingly modest. A doubling of the innovation rate can be associated with an increase in long-term cultural diversity of approximately 6%–32%. Increases on the lower end of the scale are obtained by evenness indices (Shannon evenness index, Bulla index), while increases on the higher end are obtained by the indices that also incorporate the number of variants present.

Independent invention typically generates greater diversity than modification. Over 10,000 runs we observe a maximum error margin of 4% (based on a 95% confidence level) for the four diversity measures, which means that the observed differences between the two types of innovation cannot be attributed to random effects of the Monte Carlo simulation method, and there is a genuine difference in the typical level of diversity associated with the two types of innovation. The higher diversity level associated with independent invention results from the fact that if individuals invent independently the chance of creating a variant that is strongly dominant over all others is lower compared with the situation of improvement, and as a result evenness and cultural diversity increase. Furthermore, we find that the slopes of the diversity measures are frequently steeper for modification than for independent invention. For modified variants we often observe for cases with few variants (indicated by small λ values) that the most advantageous variant has a significant frequency advantage. In contrast, situations with more variants are typically more even because the chance of inventing two comparable beneficial variants is increased. This effect is most apparent in the Shannon evenness index; the more variants are invented the more even becomes the distribution. Conversely, because of the reduced likelihood of a dominant variant in situations of independent invention we do not obtain the same slope.

These findings raise the question of whether it might be possible to distinguish between the two different sources of innovation, modification and independent invention, on the basis of the observed level of cultural diversity. To explore this question we assume a situation where four variants are invented during a given time period. Our aim is to determine the probability of obtaining a particular level of cultural diversity under the assumptions of modification or independent invention.

Fig. 3a shows the frequency distributions of the Shannon index of the equilibria $(u_1^*, u_2^*, u_3^*, u_4^*)$ obtained by 10,000 simulations. We can see that the two sources of innovation result in different characteristic diversity distributions. While independent invention

Fig. 4. Influence of conformity (expressed by the parameter p) on the spread dynamic of variants. (a) A single-variant system. (b) Two variants, where variant 1 is represented by the solid lines and variant 2 by the dashed lines. Here the variants are equally favoured by transmission biases but variant 2 has a competitive advantage to variant 1 ($c_{12} > c_{21}$). (c) Influence of direct and indirect transmission biases (expressed by the parameter b) on the spread dynamic of two variants (again, variant 1 is represented by the solid lines, and variant 2 by the dashed lines). Conformity in the population is at a moderate level but variant 2 has a benefit advantage compared with variant 1 ($c_{12} > c_{21}$). The black lines illustrate the situation where both variants are supported in the same way by direct and indirect biases. Because of its benefit advantage variant 2 achieves a higher frequency than variant 1 at equilibrium. But if variant 1 is favoured by direct and indirect biases (other lines) then it can overcome its benefit disadvantage and establish itself at a higher equilibrium frequency than variant 2. (d) Influence of strong conformity on the spread dynamic of six variants.

favours high-diversity situations where several invented variants are evenly distributed, innovation through improvement leads to a broader distribution of low and intermediate levels of diversity where the most beneficial variant shows a high frequency. The 95% limits of both distributions (where 5% of the probability mass is in the inner tail) indicate that diversity values between 0 and 0.48 can reasonably be attributed to innovation through modification whereas diversity values in the interval [0.72, 1] almost certainly result from independent invention. Thus for most of the range of values of this metric it should potentially be plausible to infer the process underlying innovation with a 5% probability of error. Fig. 3b illustrates the same situation but under the presence of a conformist bias which will be discussed in the next section.

4. Social learning biases

We now model adoption behaviour in more detail by allowing for cultural transmission biases that favour particular variants, exploiting a general formulation devised by Henrich (2001) to capture several transmission biases simultaneously. Here, the adoption rate a is modelled by the function $a = a(u) = (1-p)b + p(u - c_b)$ (Henrich, 2001) and analogously the switch rate c by the function $c = c(u) = [(1-p)d - p(u-c_b)]^+$, where the operator $[\cdot]^+$ stands for the positive part to ensure that the direction of switching is not changed. In both formulations we can clearly distinguish between a constant and a frequency-dependent component. The constant, or frequency-independent, components $(1 - p)b$ and $(1 - p)d$ represent the collective influence of direct bias (the selective copying of pre-existing variants found to be efficacious by individual assessment) and indirect bias (the selective copying of variants from individuals with specific qualities and attributes deemed to make them fit models), with b and d reflecting the population's judgement of the advantageousness of the variant given its intrinsic qualities, and those of its users. The frequency-dependent component $p(u - c_b)$ represents the influence of a conformist bias (individuals preferentially adopt the commonest variant). If the variant's frequency u is above a 'commonness threshold' c_b then the difference $u - c_B$ is positive and conformist bias results in an increase in the adoption rate and a decrease in the switch rate (adopters of a common variant are less likely to switch their variant), while if a variant is rare then conformity will reduce its adoption rate and increase its switch rate. The parameter p is a measure of the relative strength of the frequency-dependent and