Cultural Inheritance in Generalized Darwinism[*]

Christian J. Feldbacher-Escamilla   Karim Baraghith

Spring 2020

Abstract

[237] Generalized Darwinism models cultural development as an evolutionary process, where traits evolve through variation, selection, and inheritance. Inheritance describes either a discrete unit's transmission or a mixing of traits (i.e. *blending inheritance*). In this paper, we compare classical models of cultural evolution (cf. Boyd and Richerson 1988; Mesoudi 2011) and generalized population dynamics (Schurz 2011) with respect to blending inheritance. We identify problems of these models and introduce our model, which combines relevant features of both. Blending is implemented as success-based social learning, which can be shown to be an optimal strategy.

Keywords: cultural evolution, generalized Darwinism, blending inheritance, success, social learning

1   Introduction

This paper deals with a special kind of inheritance in cultural evolution, i.e. within the framework of *generalized Darwinism*. This framework is postulated by scientists and philosophers from different fields of research as a new and interdisciplinary theoretical structure or paradigm (cf. e.g. Richerson and Boyd 2001; Reydon and Scholz 2014). An extensive overview of the generalized-Darwinism approach is, e.g., provided by Schurz (2011). For a strong defense of generalized Darwinism and a carefully worked out core of Darwinian principles see Aldrich et al. (2008). Different approaches and some methodological problems are discussed by Crozier (2008) and Witt (2004). [238] For a critical, but fruitful investigation concerning Darwinian concepts outside of biology, see Reydon and Scholz (2014).

We will focus on cultural inheritance, which differs from biological repro-
duction in relevant aspects. As described perhaps most prominently by Boyd
and Richerson (1988), Mesoudi (2011) or Lewens (2015), cultural inheritance
can be modeled as different forms of social learning. In the following investig-
ation, we will define a specific learning mechanism which consists of a success-
based weighting of variants of cultural behavior which are then *blended* by an
agent who observes them. The underlying assumption is that such variations
typically are not the result of an unweighted averaging of given traits which
then are passed on. Instead, they are guided by principles that make them the
most promising or attractive for social learning.

Our project can be considered a study of philosophy of the special sciences,
providing a rational reconstruction of scientific notions, models, and theor-
ies. Philosophical rational reconstruction as explication in the wide sense con-
sists of two steps (cf. Carnap 1950/1962, §§2–3): identifying the *explicandum* as
clearly as possible, and introducing an *explicatum* to replace the former. In our
case, the concept we are mainly concerned with is *blending inheritance* in cul-
tural evolution. Our reconstruction will be *rational* in the sense that we provide
a justification for considering the model to be adequate. And it is philosoph-
ical, because the reasons we provide are not empirical (e.g. about the empirical
adequacy), but normative ones.

Our investigation is structured as follows: In section 2, we take the first step
of the reconstruction, clarifying the notion of *blending inheritance* and discuss-
ing the main theoretical constraints and arguments for blending put forward in
cultural evolutionary theory: homogeneity (blending decreases the otherwise
too high variation rate due to biases and drift), and fitness enhancement (2.1).
We also describe how the discussion of cultural evolutionary biases is linked
to social learning strategies (2.2). To our knowledge, this relation has not been
noted before and plays an important part in the second step. In section 3, we
finalize the first step by describing two prominent models of cultural inherit-
ance, one model by Boyd and Richerson (1988) and a population dynamical
model described by Schurz (2011). An innovation of the former was the in-
troduction of the distinctive feature of cultural inheritance, namely *blending*
inheritance, *as* social learning (3.1). An advantage of the latter is the simplicity
of implementing a frequency dependency bias, which we will identify as a par-
ticular variant of social learning (3.2). In section 4, we finalize the second step
of the rational reconstruction by combining features of both models. Blending
inheritance is introduced as a form of social learning via success-based weight-
ing. We also provide a normative rationale for the model (4.1). We conclude
the investigation with some provisos and a simulation illustrating the result of
the rational reconstruction (4.2). [239]

2   Blending Inheritance in Cultural Evolution

As indicated in the introduction, we will commence our rational reconstruction
by characterizing different constraints for the notion of *blending inheritance* put

forward in the literature. Afterwards, we will analyze the ingredients of these constraints as a form of social learning, which will be of utmost importance for the second step of our reconstruction.

2.1 Different Explanatory Roles of Blending Inheritance

In this section, we provide a general discussion of the notion of *inheritance* in cultural evolution, which is, due to its specific feature of mixing traits, sometimes also called *blending inheritance*. First, we hint at some historic roots of the concept, then we briefly outline our explication to provide a starting point for the modeling in the subsequent sections.

Mesoudi (2011, p.61) applies the term 'blending inheritance' to a certain microperspective, namely to trait copying individuals *who are exposed to the cultural traits of more than one person, adopting the average of all of those traits*. However, in Darwin's time (as well as contemporarily—cf. Lande 1979), blending inheritance was thought to happen in natural evolution as well, even if Mendelian genetics was accepted (for the following, cf. Mesoudi 2011, pp.41f; and Richerson and Boyd 2005, p.88). According to this hypothesis, offspring constitute an intermediate form of their parents. Darwin himself proposed that inheritance takes an average of the genetic contributions of both parents. However, as Fleeming Jenkin has pointed out in a review of the *Origin* as early as in 1867, such a concept of inheritance would mean that variation would be reduced by half at each new generation, and variety would disappear quickly (if there is no sufficiently high mutation rate). In consequence, Darwin himself abandoned blending as a principle for inheritance in natural evolution and left the problem of inheritance unresolved. Only when the significance of Gregor Mendel's work was properly understood and appreciated, and due to Ronald A. Fisher's population dynamical models, this problem of biological inheritance was tackled. Establishing principles of particulate heredity via the transmission of discrete units (genes) and transmission rules, such as the distinction between dominant and recessive alleles, did the job.

There are theories of cultural evolution which adopt particulate inheritance for the cultural realm, but operationalizing such cultural units, as meme-theorists have tried to do, has been the subject of severe critique (cf. Lewens 2015, pp.26). One of the reasons is that cultural inheritance seems, indeed, to be nonparticulate and blending in many relevant cases.

Not only is it hard to identify and operationalize the notion of *discrete units of cultural inheritance*, but there are also more general arguments in favor of nonparticulate cultural inheritance. The approach of *cultural attractors* of Sperber (1997) and Claidière and Sperber (2007) argues that [240] discrete units of cultural inheritance are misguided and superfluous, since assuming a distribution of psychological dispositions among humans on the natural level suffices to explain similar developments in culture, without there being a need to assume memes as units of reproduction.

Richerson and Boyd (2005, pp.88f) also argue for distinct principles of inheritance in culture as opposed to nature. For natural inheritance, Jenkin argued

that nonparticulate blending leads to too much homogeneity in the sense that variety vanishes, and Fisher's further theoretical framing shows that discrete units of inheritance allow for upholding variety or adequate heterogeneity. On the other hand, in cultural evolution it seems that due to the high variation rate (in a wide sense), assuming discrete units of inheritance such as memes would lead to too much heterogeneity, and that given a high variation rate, nonparticulate blending inheritance allows for adequate homogeneity. So, whereas discrete units in natural inheritance allow for balancing inheritance towards adequate heterogeneity, blending in cultural inheritance allows for balancing inheritance towards adequate homogeneity:

> "We can even imagine that cultural transmission is sufficiently noisy and error prone that blending inheritance would be an advantage in keeping cultural variation from growing disastrously large. In a noisy world, taking the average of many models may be necessary to uncover a reasonable approximation of the true value of a particular trait." (cf. Richerson and Boyd 2005, p.89)

In the same line, Henrich, Boyd and Richerson (2008, misunderstanding 1 and 2) argue against the misunderstanding of cultural evolution that "[r]eplicators are necessary for cumulative, adaptive evolution".

As we will see in the sections on modeling cultural evolution (3 and 4), evolutionary simulations suggest that blending inheritance in fact results in the reduction of cultural variation in the population. Indeed, as discussed by Jenkin for biological inheritance, if blending inheritance were the only process on cultural variation in the population, it would eliminate that variation completely, since intermediate traits of the daughter generations are always coming from more "extreme" ancestral traits. However, it is a high variation or mutation rate (as in natural evolution) and other evolutionary "forces" such as guided variation, content bias, and several kinds of randomization such as drift, which run against a homogenization effect and keep up variation and change over the cultural generations—or, as Mesoudi (2011, p.62) aptly puts it:

> "Obviously, in the real world blending inheritance cannot be the only process operating on cultural evolution, otherwise we would not see the enormous cultural variation [...]: 7.7 million patents, 10,000 religions, 6,800 languages, and so on. There must, then, be other processes at work."

So, in general, *blending inheritance* is characterized as a process in which a single individual adopts the average of two or more demonstrators' continuous traits, whereas in the case of particulate inheritance discrete traits are copied. [241] We will now outline our explication of *blending inheritance* as success-based social learning.

Key to understanding the notion of *blending inheritance* is that it can increase the frequency of the relevant traits, a fact that many authors do not stress enough. This possibility is depicted in figure 1. If we assume two cases of reproduction, one of particulate (left side) and one of blending inheritance (right

side), then the resulting macroevolutionary patterns might exhibit a higher frequency in the case of blending inheritance. We will say a bit more about figure 1 below.

Macrolevel
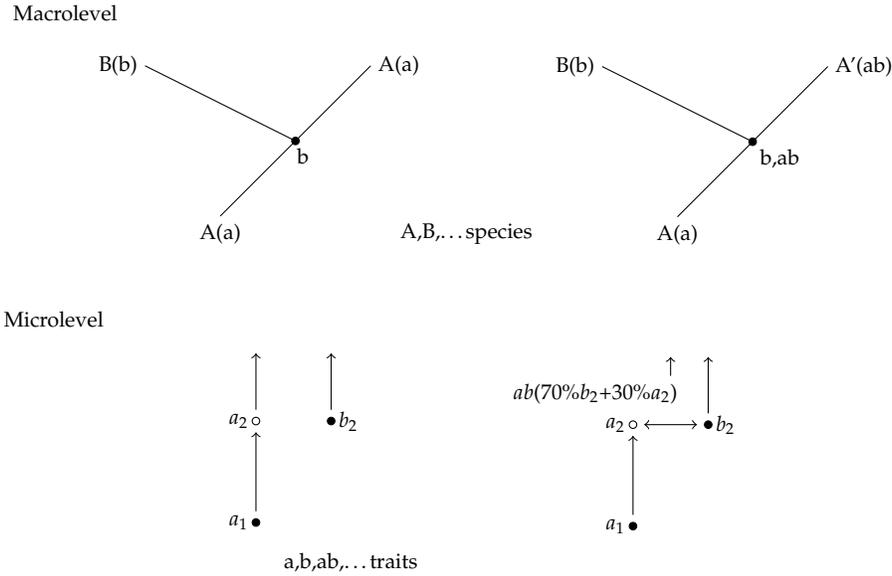
Microlevel

Figure 1: Blending inheritance on a microlevel and the possible results on a macrolevel. If we assume (at the bottom right) that traits a_2 and b_2 blend together under some guided principle of weighting and form a new trait (ab), and that the 70% of b_2 as well as the 30% of a_2 that remain in the new combined trait are not maladaptive, then species A' (top right) should be fitter than in the unblended case on the left. (Indices represent generation numbers and arrows represent transmission.)

In cultural evolutionary modeling, the concept of *traits* describes identifiable units of cultural transmission. The "units" of inheritance and selection are not biological organisms or their genes, but cultural information, skills, or artifacts, which are selected by, vary between, and are inherited through cultural "generations". Accordingly, these generations are not biological lifecycles, but cultural reproductive cycles. Each transmission of socially acquired information from one individual to the next makes up a reproductive cycle, [242] in which a cultural trait is passed on. The units can be represented by discrete entities, as is the case in the models in section 3.2, or by continuous values, as in the models in sections 3.1 and 4. More specifically, in our examples we think of traits as behavioral dispositions. The term *fitness* is well defined in natural evolution, i.e. the number of an organism's reproducing offspring (fertility notion). In cultural evolution, such a straightforward definition is problematic due to the individuation problem of cultural units (cf. Ramsey and De Block 2017). However, as Boyd and Richerson (1988), we will assume that cultural traits are given. We then identify the fitness of such a trait with the coefficient

used to calculate the dynamics of the relative frequencies of that variant across generations. Therefore, it is important to note that with 'fitness' we do not mean biological fitness of the organism bearing that variant. Rather, we mean a coefficient linking the relative frequency of a cultural variant in the set of all relevant alternative cultural variants across generations. Note that we also presuppose a notion of *relevant alternatives* here. Since a cultural variant cannot be blended with every other cultural variant, we are only interested in the relative frequencies of cultural variants which can be blended together. For example, it is possible to mix different styles of piano playing to "create" a new style, but impossible to mix a piano-playing style with a way of cooking (although cooking styles themselves might be mixed with one another). We assume that categories in which cultural variants or traits can be blended together are given by their cultural function in the first place. This cultural evolutionary function plays a significant role in the emergence of higher-order categories or types. A philosophical approach which encapsulates this idea is teleofunctionalism (cf. Millikan 1984, 2005).

An example will illustrate the principle of fitness enhancement via blending along the lines of figure 1: On the lower left side we see trait a, which is passed on from the mother generation (a_1) to the daughter generation (a_2). For illustrative purposes, let us assume that the generations are political election cycles, and the traits are political dispositions, such as being left- or right-wing (cf. Boyd and Richerson 1988, p.70; and also Mesoudi 2011, p.61). More specifically, the traits should be interpreted as manifestations of such dispositions, such as acting x% in accordance with left-wing politics. In this interpretation, a blended trait is a mixed manifestation of two dispositions, e.g., acting 70% according to right-wing and 30% according to left-wing politics. The left side of figure 1 shows a simple case of particulate inheritance with two such variants (a_2, b_2) at generation 2. On the right, blending inheritance is depicted. Here, trait a was not fit for one reason or another, so we do not find it anymore in the daughter generations. However, a new (unblended) variant b arose, as well as a blended variant ab. Let us assume a proportion of 70% (b) to 30% (a) in ab. Let us further assume that the 30% of a were *not* the reason why it died out, meaning that they are not maladaptive given a certain environment. [243] To the contrary: maybe the agent who mixed it into the new combined trait approved it as useful. The resulting macrolevel structure (right side, top) exposes two hypothetical daughter species B and A', where the latter holds of the combined trait ab. Assuming that the 30% of a that remained in ab are useful and even provide some increase in success, the cultural species A' should now be fitter than in the left case. In this way, blending inheritance might prove to be advantageous. It is important to note that this view of blending as a principle of fitness enhancement presupposes that blending is *adaptive*. Trait ab with 30% a and 70% b only increases fitness, if, roughly speaking, ab manifests a in the 30% of cases, in which a is advantageous to b, and ab manifests b in the 70% of cases in which b is advantageous to a. We call this form of blending *adaptive blending*. If ab would manifest a in 30% of the b-advantageous cases, and b in 30% a-advantageous cases as well as 40% b-advantageous cases, blend-

ing would be *maladaptive*. We can also define the notion of *random blending* where *ab* manifests *a* and *b* with random frequency (*a* 30% and *b* 70%). Random blending does not generally enhance fitness. It is important to distinguish these tree forms of blending (adaptive, maladaptive, and random), because our model in section 4 focusses on *adaptive blending* only and shows under which assumptions it is optimal.
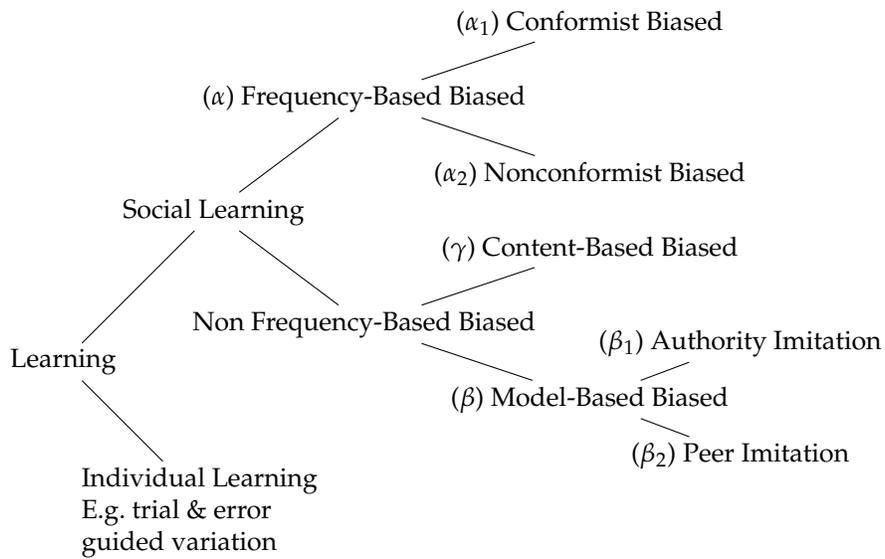
2.2 Identifying Cultural Evolutionary Biases as Social Learning Strategies

According to Boyd and Richerson (1988, p.72), cultural transmission is often employed by individuals who try to estimate which behavior of other individuals in their environment has been favored by selection in previous generations. This form of transmission can be characterized as social learning, as opposed to individual learning, where one achieves information and knowledge via a process of trial and error (e.g. in experiments).
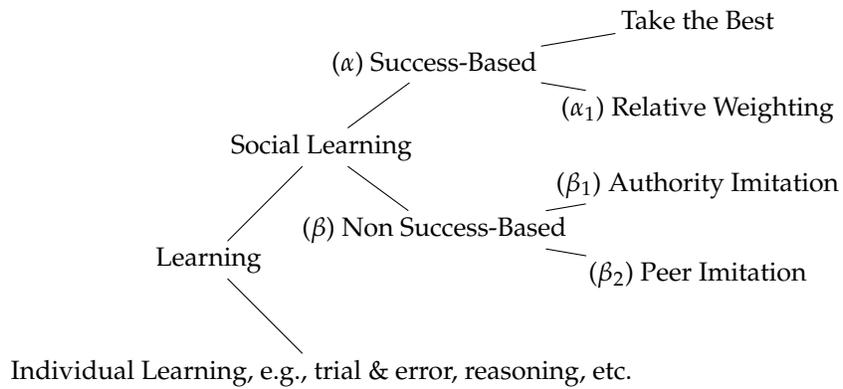
Modeling cultural evolution in the framework of social learning allows us to spell out two basic features of generalized Darwinism, *variation* and *inheritance*. The third feature, *selection*, enriches cultural modeling by introducing "forces of cultural evolution" in the form of *biases*. According to Richerson and Boyd (cf. 2005, p.68), the three most important biases are *content-based biases*, *model-based biases*, and *frequency-based biases*. In Boyd and Richerson (1988, p.135), content-based bias is also referred to as 'direct bias', which means that one cultural variant is more attractive to the learner than others. So, the probability of such a variant being chosen by a social learner is higher than that of its alternatives. Model-based bias is also referred to as 'indirect bias', which means that the choice of a cultural variant by the learner depends on how successful its bearer is in the mother generation. Model-based biases are active, e.g. in authority imitation or peer imitation where, apart from the prestige of the learner's parental models, similarity of the parent to the learner influences success. [244] Finally, frequency-based biases are active in a model of cultural evolution if the probability of choosing a variant among the set of all its alternatives depends on the frequency of the variant in the mother generation. If this probability is positively correlated with frequency, then the bias is a *conformist* one. If they are negatively correlated, then the bias is *nonconformist*. In figure 2, we provide a simplified taxonomy of these biases. They are considered different kinds of social learning, because according to Boyd and Richerson (1988, p.136), "individuals select from among the alternative cultural variants that have been modeled for them rather than choosing among self-generated alternatives".

It is interesting to note that this taxonomy of different forms of biased learning (partially) matches a taxonomy of social learning as used in applied machine learning literature. Figure 3 gives a general overview of these different forms of social learning. [245]

Learning can either be individual or social. Individual learning is trial-

(α₁) Conformist Biased

(α) Frequency-Based Biased

(α₂) Nonconformist Biased

Social Learning

(γ) Content-Based Biased

Non Frequency-Based Biased

(β₁) Authority Imitation

Learning

(β) Model-Based Biased

(β₂) Peer Imitation

Individual Learning
E.g. trial & error
guided variation

Figure 2: An overview of social learning with different biases along the lines of Boyd and Richerson (1988, p.135) and Richerson and Boyd (2005, p.69)

Take the Best

(α) Success-Based

(α₁) Relative Weighting

Social Learning

(β₁) Authority Imitation

(β) Non Success-Based

Learning

(β₂) Peer Imitation

Individual Learning, e.g., trial & error, reasoning, etc.

Figure 3: An overview of social learning strategies as used in machine learning applied in epistemology

and-error learning, reasoning, etc. Social learning strategies are either success-based (α) or not success-based (β). Success-based strategies accept more transmissions from successful parental traits than from non-successful ones. *Take the best*, e.g., simply favors the most successful trait of the mother generation. *Relative weighting*, on the other hand, blends traits according to a weighted average, where the weights are proportional to the traits' past success. In machine learning and in recent approaches of social epistemology, relative weighting is

8

studied under the term *meta-induction*, since success-based weighting can be regarded as induction over success rates (Schurz 2008, 2019).

Comparing the two taxonomies, several points are worth noting: Non-success-based social learning and non-frequency-based biased social learning, do not only cover the mentioned forms of imitation, but also the so-called *fast and frugal heuristics* of Gigerenzer, Todd and The ABC Research Group (1999). In contrast to the biases mentioned above, which are parameters of selection (Boyd and Richerson 1988, p.136), guided variation is a form of individual learning (see figure 2), because it generates alternative cultural traits. It is common in cultural evolution (cf., e.g., Mesoudi 2011, p.63). Sufficiently rational agents will prefer some variants over others. In the context of cumulative evolution, this is not a problem, but an advantage, because fewer dysfunctional mutations are passed on in the evolutionary dynamic. The space of possibilities of variants (many of which are indeed dysfunctional) shrinks, and that leads to an acceleration of the generation of complexity in cultural evolution in comparison to natural evolution.

Furthermore, note that the notion of *success* as it is used in the taxonomy of figure 3 might be different from the notion of *success* in the sense of prestige which is employed by Richerson and Boyd (2005) for describing model-based biased social learning (β in figure 2). In fact, in our model of *blending inheritance* in section 4 we will equate the success of a cultural variant with the frequency of the variant in the mother generation. For this reason, success-based social learning according to the taxonomy of figure 3 is, at a surface level, comparable to frequency-based biased social learning of figure 2 (both αs). However, the important difference between both models of frequency-dependent cultural transmission is that frequency-based biased cultural transmission (figure 2) is *not* blending, since it selects only among the cultural alternatives provided by the mother generation, [246] whereas frequency-based relative weighting (figure 3) *is* blending, since it creates new cultural variants. In this respect, our way of modeling blending is closer to the model of guided variation of Boyd and Richerson (1988, p.136). Therefore, we will concentrate on this model and not on their model on frequency-dependent biased social learning.

In order to embed our model even deeper into recent debates about learning and cultural evolution, let us briefly highlight another model from the literature that resembles our approach in some features, but differs in others. In a series of papers, Griffiths and Kalish (2007) develop an *iterative learning framework with Bayesian Agents*, which combines individual and social learning, occupying a middle ground between the two. Focusing on the evolution of language, the authors present a dynamic analysis of social learning. They understand the process of learning as an agent choosing from a set of hypotheses and data. Applying Bayes' rule (an ability that the agent must be capable of, cf. Griffiths and Kalish 2007, p 472, assumption 3), she can either choose a given hypothesis (which she has acquired via the observation of others) or new data in light of a given hypothesis (cf. Griffiths and Kalish 2007, p.444). Griffiths and Kalish primarily focus on and reflect the strong influence of the (Bayesian) prior of the learner and its effects on the accuracy of transmission

processes. Convergence of the probability that a learner speaks a particular language to the prior probability the learner assigns to that language occurs regardless of the amount of data available to each learner (cf. Griffiths and Kalish 2007, p.444). Their mathematical findings suggest that the influence of inductive biases (individual constraints on learning which influence our conclusions from incomplete knowledge) seem to have a very strong effect on iterated social learning. Furthermore, inductive biases strongly resemble cultural guided variation, since both emphasize the individual aspect of learning (cf. Mesoudi 2011). In contrast to these models, our investigation (cf. section 4) focuses merely on the process of social learning, meaning that our agents lack any form of individual bias. However, as we will see, a purely social learner can still achieve optimal performance.

Having illustrated what blending inheritance is and how it influences the process of cultural evolution, we will now take a closer look at the formal structure of a success-based mechanism of cultural inheritance by studying models of cultural evolution (section 3) and implementing blending in a success-based manner (section 4).

3 Models of Cultural Evolution

Boyd and Richerson (1988) investigate (almost) all forms of biased social learning as described in the taxonomy in figure 2: content-based biased social learning (chpt.5, parameter B, representing an inherent disposition of a cultural trait to be preferred against some other, cf. p.138), model-based biased social learning (chpt.8, via the so-called *runaway processes*), and conformist frequency-based biased social learning (chpt.7, cf., e.g., parameter D, p.209). In the following subsection, we will describe their model of blending inheritance via social learning in detail. [247] Subsequently, we will discuss a population dynamical model which serves as a basis for our model of success-based blending inheritance in section 4.

3.1 Cultural Evolution via Social Learning

Robert Boyd and Peter J. Richerson argue that very often in the cultural realm data, e.g. data about people in a political spectrum, can be represented by a normal distribution (see figure 4). This means that one can take the frequencies, e.g. of the number of left-wing, centered, right-wing people, and summarize them in a probability density of a normal distribution. Thus, one defines the probability Pr of some member of the group being left-wing etc. The normal distribution is characterized by two factors: μ, the mean, which is also the frequencies' median and mode, and σ, the standard deviation (i.e. the average deviation from μ), where σ^2 is considered to be a measure for variance within the sample. The probability that some of the X of generation n, i.e. some of X^n, take value x, e.g. that some member of the group under consideration X^n holds position x in the political spectrum, is:

$$Pr(X^n = x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

$Pr(X^n = x)$

left   right x
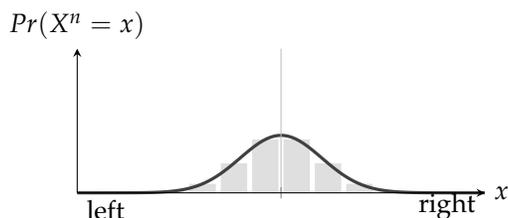
Figure 4: Example of a normal distribution of political attitudes (x) as a cultural trait in generation X^n. Of course, whether political attitudes are normally distributed or not depends on the scale of measurement and the chosen categories (e.g. it depends on whether the frequencies of the categories left, center-left, center, center-right, right—cf. the light-gray bars—or frequencies of other categories of political attitudes result in a Gaussian shape.)

Based on this assumption, Boyd and Richerson propose a model for the transmission of cultural traits by means of social learning (see figure 5). Their idea is that a cultural trait can be transmitted by individual learning (copying that is possibly defective) and social learning about an objective, an ideal state of the system. The model has the following parameters:
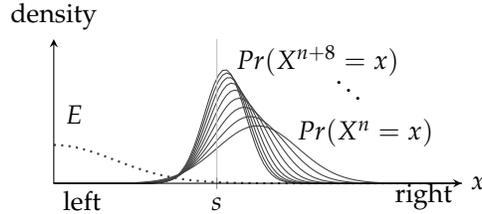
- X^n is the cultural trait to be copied under investigation, expressed as a random variable and defined by $\mu(X^n)$ and $\sigma^2(X^n)$. [248]

- s is the optimal state of the cultural trait; it is supposed to be the objective or goal value, i.e. the value of the distribution that fares best in the system or habitat (cf. $\Gamma(H)$ and H in Boyd and Richerson 1988, p.95). So, in a sense, s encodes selective pressure: the closer an individual is to s, the better it will fare in the system.

- l is a parameter in $[0,1]$ which expresses the individuals' propensity to rely on individual as opposed to social learning (cf. L in Boyd and Richerson 1988, p.95). Individual learning is maximal if $l = 0$ and social learning is maximal, if $l = 1$.

- E is an error distribution which characterizes a deviation during the learning process due to environmental effects, random variation, and estimation errors (cf. Boyd and Richerson 1988, p.71). It is defined by $\mu(E)$ and $\sigma^2(E)$. E might be considered the cultural counterpart of mutation in natural evolution.

- X^{n+1} is the cultural trait of the next generation resulting from individual learning and social learning. It is defined by $\mu(X^{n+1})$ and $\sigma^2(X^{n+1})$.

In short, cultural transmission from X^n to X^{n+1} depends on individual, possibly defective, learning, and cultural learning aiming at s; both are balanced by the parameter l. The influence of individual learning (encoded in X^n) is taken into account via the influence of error, i.e. $\sigma^2(E)$; the influence of social learning (s with error during social learning E) is taken into account via l. Calculating weights by normalizing, we end up with (cf. equ.4.9 in Boyd and Richerson 1988, p.95):

$$X^{n+1} = \frac{\sigma^2(E)}{\sigma^2(E) + l} \cdot X^n + \frac{l}{\sigma^2(E) + l} \cdot (s + E) \tag{1}$$

[249] Given that the learning error is not biased, i.e. $\mu(E) = 0$, Boyd and Richerson show that the resulting cultural trait is a density function with the following properties/values (cf. equ.4.10 in Boyd and Richerson 1988, p.96):

$$\mu(X^{n+1}) = \frac{\sigma^2(E)}{\sigma^2(E) + l} \cdot \mu(X^n) + \frac{l}{\sigma^2(E) + l} \cdot s$$

$$\sigma^2(X^{n+1}) = \left(\frac{\sigma^2(E)}{\sigma^2(E) + l} \right)^2 \cdot \sigma^2(X^n) + \left(\frac{l}{\sigma^2(E) + l} \right)^2 \cdot \sigma^2(E)$$
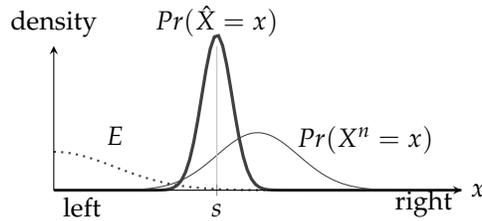

Figure 5: Political attitude with social learning in the model of Boyd and Richerson: Starting from the political attitude of generation n (X^n of figure 4) through social learning with $l > 0$ (here $l = 0.5$), although defective, but unbiased ($\mu(E) = 0$), evolution via guided variation tends towards the best fitted political attitude with value s. (Data generated with Perl.)

Now, given $l > 0$ (and a constant and uniform environment), it follows that iterated individual and social learning will lead to $\mu(X^{n+1}) = s$. Note that the model assumes that individual learning is fully implemented by the first summand in equation 1. Given the assumption that the learning error is not biased (which means that error is not functional), according to this model, s can be only learned if $l > 0$. If error E is functional, then also $l = 0$ allows for learning s. An equilibrium of the system is reached, if $X^{n+1} = X^n$. Since the weights $\frac{\sigma^2(E)}{\sigma^2(E)+l}$ and $\frac{l}{\sigma^2(E)+l}$ add up to 1, the value for the mean in equilibrium state, $\hat{X}$, is calculated as follows:

$$\mu(\hat{X}) = \frac{\sigma^2(E)}{\sigma^2(E) + l} \cdot \mu(\hat{X}) + \frac{l}{\sigma^2(E) + l} \cdot s = s$$

Such an equilibrium state is depicted in figure 6. The equilibrium variance is more complex, but turns out to be:

$$\sigma^2(\hat{X}) = \left(\frac{\sigma^2(E)}{\sigma^2(E)+l}\right)^2 \cdot \sigma^2(\hat{X}) + \left(\frac{l}{\sigma^2(E)+l}\right)^2 \cdot \sigma^2(E) = \frac{\sigma^2(E) \cdot l}{2\sigma^2(E)+l}$$

Figure 6: Political attitude: The equilibrium state $\hat{X}$ centers around the political attitude best fitted with value s. (Data generated with Perl.)

It is important to note that we can map three relevant parameters in this model to the three modules of generalized Darwinism: Selection is taking place via s. Variation comes in two forms: guided variation via l; mutation via E; [250] and the reproduction dynamics is captured via the cultural traits $X^n, X^{n+1}, \ldots$. This model and its interpretation show that a sufficiently high degree of information transmission into the next cultural generation suffices for cumulative cultural evolution. Exact replication of the original is not necessary (cf. Henrich, Boyd and Richerson 2008, misunderstanding 2).

This model was one of the first influential models of cultural evolution. It is relatively simple and allows for an explanation of reproductive success via social learning: Through social learning the mean of a cultural trait ($\mu(\hat{X})$) tends towards the best fitted value (s). However, there are also some restrictions to this model (cf. Schurz 2011, p.220): First, social learning is only successful in the strict sense of $\mu(\hat{X}) = s$ in case that there is no bias in the reproduction error ($\mu(E) = 0$). And second, selective pressure of social learning is held fixed in the model via the cross-generational parameter l. For this reason it is also independent from the reproduction rate, although frequency dependency is an important feature of cultural transmission. Boyd and Richerson (1988) discuss possibilities to relax these constraints. E.g., also in case of a biased reproduction error one can still approach the learning target s, though there is no guarantee to, in fact, learn s. It seems that such an approximation can still be considered as learning success, and hence successful cultural transmission seems to be possible also in case of biased error. Regarding the problem that l is cross-generationally defined, Boyd and Richerson (1988, pp.99f) also discuss models where l changes (due to genetic influence). However, this variation of l is still not dependent on the relative frequency of the cultural traits. Our model

in section 4 aims at adding such frequency dependency to a model of blending. Furthermore, we want to expand Boyd and Richerson's result on cultural transmission: We show that even without an objective learning target s, one can define a frequency-dependent social learning (blending) strategy which is long-run optimal. In the next subsection we will introduce the population dynamical model on which our model of blending inheritance is based.

3.2   Population Dynamics of Cultural Evolution

The first mathematical models of population dynamics trace back to Ronald A. Fisher and John B.S. Haldane (cf. Fisher 1930; and Haldane 1964/2008). A clear and didactically valuable introduction is given by Ridley (2004). Starting with a general scheme, different models for natural and cultural evolution are spelled out and refined. The general scheme has the following ingredients (cf. Schurz 2011, sect.12.3, notation adjusted to the one above):

- $v_1, \ldots, v_k \ldots$ possible variants/values of a system; a set of variants V in a certain environment is called a *population*;

- $Pr(X^n = v_i) \ldots$ relative frequency of X^n taking value v_i

- Generations: $Pr(X^{n+1} = v_i)$: The relative frequency of v_is in the next generation (X^{n+1}).

[251] The simplest models of population dynamics operate only on the variants' frequencies across different generations—these are models of cultural reproduction. By introducing a *selection coefficient s* which constrains the variants' frequencies across different generations, these models are expanded to models of cultural reproduction and selection (Schurz 2011, p.285, p.292). In a further step, one can take variation (in a wide sense) into account via biases in errors. This can be done by introducing a parameter m into the dynamics that represents the mutation rate of a variant. E.g., a mutation of variant v_1 to variant v_2 in 1‰ of the cases could be represented by $m_{v_1 \longrightarrow v_2} = 0.001$; so $m_{v_1 \longrightarrow v_2} \cdot Pr(X^n = v_1)$ is the number of v_1-variants in generation X^n that mutate back to v_2-variants in generation X^{n+1}. Of course, mutation can also take place the other way round, from v_2 to v_1, but we assume that this direction is already subtracted from the higher mutation rate (so, e.g, $m_{v_2 \longrightarrow v_1} = 0$). Furthermore, it is important to note that mutation is understood only as mutation from one variant of a population X to another. The evolution of new variants is not covered by this. This is also the reason why we expand the model in the next section (4) in order to describe such an evolution of new variants, that is, evolution by blending.

In order to implement a mutation mechanism into the dynamics, the number of mutated variants must be subtracted from and added to the frequency of the variant. The resulting formula is as follows:

$$Pr(X^{n+1} = v_i) = \frac{Pr(X^n = v_i) \cdot s_i - \sum\limits_{o \neq i}^{k} Pr(X^n = v_i) \cdot m_{v_i \longrightarrow v_o} + \sum\limits_{o \neq i}^{k} Pr(X^n = v_o) \cdot m_{v_o \longrightarrow v_i}}{\sum\limits_{j=1}^{k} Pr(X^n = v_j) \cdot s_j}$$

The selection coefficient s of this model measures fitness in the sense of selective advantage, and not selective disadvantage (as is the more common use of the term 'selection coefficient'). We choose s as a measure for selective advantage to use terminology coherent to the preceding subsection (where s was interpreted as the relevant parameter for selection).

The question for equilibria for a variant in this model is stated similarly as in the model of Boyd and Richerson by equating $X^{n+1} = X^n = \hat{X}$: Biased error, implemented via a mutation rate m, allows for equilibria where not all the other variants disappear completely. So, in the case of two variants with a positive selection of variant v_1, the frequency of v_1 has to reach only $1 - m_{v_1 \longrightarrow v_2}/(s_1 - s_2)$ in order to be in an equilibrium ($Pr(X^{n+1}) = Pr(X^n)$). If $s_1 = 1$ and $s_2 = 1 - s$, the equilibrium state is at $1 - m_{v_1 \longrightarrow v_2}/s$ (cf. Schurz 2011, p.300). Figure 7 depicts a case of a population dynamics with selection, variation (mutation), and reproduction. The mutation rate of a variant v_i has to be smaller than the selective advantage of that variant. Otherwise it will disappear (*mutate away*) before it can become positively selected. [252]
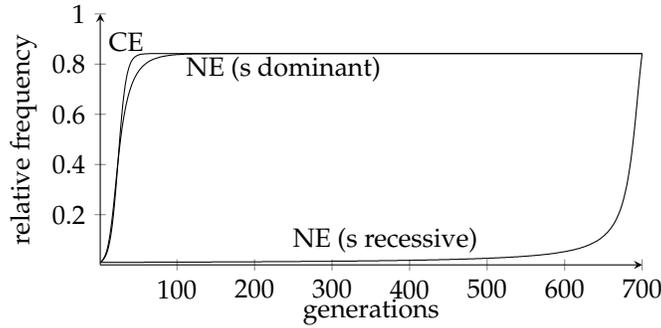
Figure 7: Natural and cultural evolution with reproduction, selection, and mutation in the population dynamical model of Schurz (2011). Parameters: natural evolution, selection of a dominant allele (NE (s dominant)): $s_{dominant} = 1$, $s_{recessive} = 1 - 0.2 = 0.8$, $m_{dominant \longrightarrow recessive} = 5‰$, $Pr(X^0 = v_{dominant}) = 0.01$, $Pr(X^0 = v_{recessive}) = 0.99$; natural evolution, selection of a recessive allele (NE (s recessive)): $s_{dominant} = 1 - 0.2 = 0.8$, $s_{recessive} = 1$, $m_{recessive \longrightarrow dominant} = 1‰$, $Pr(X^0 = v_{dominant}) = 0.99$, $Pr(X^0 = v_{recessive}) = 0.01$; cultural evolution (CE): $s_1 = 1$, $s_2 = 1 - 0.2 = 0.8$, $m_{v_1 \longrightarrow v_2} = 31.6‰$, $Pr(X^0 = v_1) = 0.01$, $Pr(X^0 = v_2) = 0.99$; the equilibrium state is reached at a relative frequency $Pr(\hat{X} = v_1) = 1 - 0.0316/0.2 = 0.842$; only the positively selected variant's progression is depicted; (Data generated with Perl.)

As we have mentioned in the preceding section, in the model of Boyd and Richerson, the parameter l weights the importance of social learning, but is the same for all variants and independent from the variants' reproductive success. In cultural evolution, however, it seems that social learning is correlated with

reproductive success. It is positively correlated when *conformity* (positive frequency dependency) is relevant for fitness; and it is negatively correlated when *originality* (negative frequency dependency) marks a fit variant. Note that frequency dependency does not automatically imply that traits with higher fitness in the sense of higher frequency are also learned more frequently. Rather, frequency dependency means that the fitness coefficients are themselves a function of the frequency of the traits. This is an important difference, because due to such frequency dependency, stable equilibria are no longer guaranteed, but ongoing oscillations are possible (such cases are also covered by our model).

In the above equation such a dependency can be implemented via fine-graining the selection coefficient s_i of a variant v_i by making it dependent on the frequency of the variant ($|v_i|_{X^n}$). By this we finally end up with the following model: [253]

$$Pr(X^{n+1} = v_i) = \frac{Pr(X^n = v_i) \cdot s_i(Pr(X^n = v_i)) \begin{array}{l} - \sum\limits_{i \neq o = 1}^{k} Pr(X^n = v_i) \cdot m_{v_i \longrightarrow v_o} \\ + \sum\limits_{i \neq o = 1}^{k} Pr(X^n = v_o) \cdot m_{v_o \longrightarrow v_i} \end{array}}{\sum\limits_{j=1}^{k} Pr(X^n = v_j) \cdot s_j(Pr(X^n = v_j))} \tag{2}$$

Schurz (2011, chpt.14.1) shows that linear dependence of s_i to $Pr(X^n = v_i)$ does not change the equilibrium state, neither in case of a positive nor in case of a negative dependency—such a case is depicted in figure 8. Things turn out to be different for non-linear dependence: There, negative dependency leads to oscillations, whereas positive dependency might switch the equilibrium state to another extreme.
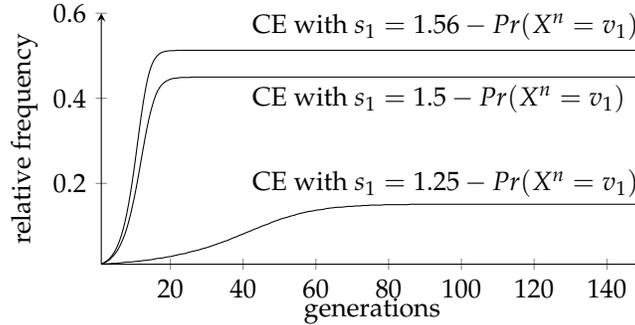
Figure 8: Cultural evolution with reproduction, selection, and mutation in a frequency dependent model. Parameters: s_i is frequency dependent, $Pr(X^0 = v_1) = 0.01$, $Pr(X^0 = v_2) = 0.99$; $m_{v_1 \longrightarrow v_2} = 0.05$; in case of a linear dependence of more than 1.56 times the ancestor frequency v_1 takes over v_2 in the equilibrium state. (Data generated with Perl.)

The models presented so far combine several features: Boyd and Richerson's model allows for *blending inheritance* by the combination of variants

via social learning, but it does not implement frequency dependency. In addition, the population dynamical model presented in equation 2 implements frequency dependency, but does not include blending. In the following part, we will combine both approaches by expanding this model to a model where variants are blended.

4 A Success-Based Model of *Blending Inheritance*

In this section, we will spell out the frequency-dependency of the selection parameter *s* used in the foregoing model in more detail. As we will see soon, this allows for *blending* via social learning. Schurz (2011) mentions this expansion of his model and highlights its significance: [254]

> "Despite the realm of epistemology, meta-induction [understood as success-based blending] is also of utmost importance for cultural evolution. In cultural evolution, innate cognitive modules play the role of non-inductive strategies, individual trial-and-error learning plays the role of object-inductive strategies, and meta-inductive strategies correspond to methods of social learning. The optimality results of meta-induction explain the advantage of creatures who are able to undergo change through cultural evolution." (Schurz 2011, pp.387f, our translation)

The main idea underlying the expansion of the model is to make the selection coefficient for a variant v_i (i.e. s_i) not only dependent on the frequency of the variant ($Pr(X^n = v_i)$), but also dependent on the frequencies of the other variants. Whereas the dependency of s_i on $Pr(X^n = v_i)$ alone is called *reflective frequency dependency*, the latter is called *interactive frequency dependency* (Schurz 2011, p.311). In biology, several models for interactive frequency dependency have been discussed, e.g. the so-called *predator–prey equations* of Alfred J. Lotka and Vito Volterra (note that in predator-prey equation models dying out is possible and thus they allow for changes in the absolute population size, while in the models discussed so far the absolute population size is fixed). Also in the cultural domain, interactive frequency dependent selection has been investigated regarding the development of meaning (cf. Mühlenbernd and Franke 2014, sect.1): In a generalized evolutionary game theory, one might try to spell out evolutionary strategies that interact with one another, resulting in models of interactive frequency dependency. The investigation of Roland Mühlenbernd and Michael Franke already makes use of success-based *imitation or reproduction*. However, a relevant difference to our approach is that we introduce relative-success-based *blending*, which allows for an optimality result for an equilibrium state tending in a similar direction as that of Boyd and Richerson presented above. In the following we will spell out our approach in more detail.

4.1 The Model and Some Analytic Results

Assuming that the fitness or success of a variant v_i consists in its magnitude or relative frequency $Pr(X^n = v_i)$, one can try to *socially learn from* variants by blending the relatively successful ones. To return to the example of Boyd and Richerson, if v_1 and v_2 are the most successful variants centring around the *correct* or *ideal state s*, then blending v_1 and v_2 to $(v_1 + v_2)/2$ might be interpreted as *socially learning* from these variants that also center around the ideal state. Boyd and Richerson interpret *social learning* as erroneous (*E*) taking into account the ideal state (*s*) to some fixed degree (*l*). In our model, we interpret *social learning* as blending those variants (v_i) which were relatively successfully in terms of frequencies ($Pr(X^n = v_i)$). The formal theory of this interpretation stems from mathematical learning theory in general, and from the theory of *online predictions based on expert advice* in particular (cf. Cesa-Bianchi and Lugosi 2006; Schurz 2008, 2019). [255] An important result of this theory is the observation that relative-success-based weighting (cf. α_1 in the taxonomy of figure 3) turns out to be long-run access-optimal in every environment or habitat (even oscillating ones).

The ingredients of our model are the same as that of the population dynamical model of the preceding section. However, we assume that there is a *learning variant v_l* (for some $1 \leq l \leq k$). For this variant, the selection coefficient s_l is made relative-success-dependent on all variant's past frequencies. We will implement this by varying v_l at each generation depending on all variants' past frequencies. The idea is that by such a blending/learning, v_l approximates the most successful variants in the long run. Here are the details of the blending mechanism we propose:

- We define a loss function, which calculates the difference between the relative frequency of a variant from the best fitted variant in a generation n and normalizes it ($\in [0, 1]$):

$$d_n(i) = \frac{|Pr(X^n = v_i) - max(Pr(X^n = v_1), \ldots, Pr(X^n = v_k))|}{max(Pr(X^n = v_1), \ldots, Pr(X^n = v_k))}$$

  (Note that in machine learning literature the difference between a variant's (cumulative) loss and the best variant's (cumulative) loss is also called the "regret" of the variant with respect to the best variant; cf., e.g., Cesa-Bianchi and Lugosi 2006, p.2.)

- Based on this loss function we define a measure for normalized success of a variant up to generation n as the inverse of the average natural loss:

$$as_n(i) = \frac{\sum\limits_{m=1}^{n} 1 - d_n(m)}{i}$$

- Based on the normalized success of a variant up to generation n we define the relative success of a variant with respect to the *social learning* variant

18

v_l up to generation n by cutting off (i.e. setting 0) the normalized success of those variants that fare worse than the *socially learning* variant v_l:

$$rs_n(i) = \begin{cases} as_n(i), \text{ if } i \neq l \text{ and } as_n(i) \geq as_n(l) \\ 0, \text{ otherwise} \end{cases}$$

(What we call "relative success" here is also called "attractivity" in the literature on meta-induction; we opted for the former expression because the latter is also used in Sperber's cultural attraction model and might have caused confusion here; for meta-inductive attractivity cf. Schurz 2008, sect.7.)

- Based on the relative success of a variant up to generation n, we define a weight for the variant for generation $n + 1$ by normalizing the relative success: [256]

$$w_{n+1}(i) = \begin{cases} \dfrac{rs_n(i)}{\sum\limits_{j=1}^{k} rs_n(j)}, \text{ if } \sum\limits_{j=1}^{k} rs_n(j) > 0 \\ \dfrac{1}{k-1}, \text{ otherwise} \end{cases}$$

- And, finally, based on these weights for generation $n + 1$, we define the social learning of variant v_l for generation $n + 1$ as:

$$v_l^{n+1} = \sum_{l \neq j=1}^{k} w_{n+1}(j) \cdot v_j$$

v_l^0 may be defined as blending by unweighted averaging variants:

$$v_l^0 = \frac{\sum\limits_{l \neq j=1}^{k} v_j}{k - 1}$$

All the non-learning variants (v_i with $1 \leq i \neq l \leq k$) remain constant. It is assumed that they can be represented by real numbers as, e.g., political attitudes on a real-valued spectrum. For their dynamics, the equation of the reflective frequency dependent model of Schurz (2011) holds (s_i is only reflective frequency dependent). For the learning variant v_l, the equation must be adapted by considering the interactive frequency-dependency of v_l:

$$Pr(X^{n+1} = v_l^{n+1}) = \frac{Pr(X^n = v_l^n) \cdot s_l(Pr(X^n = v_l^n)) - \sum\limits_{l \neq i=1}^{k} Pr(X^n = v_l^n) \cdot m_{v_l^n \longrightarrow v_i}}{\sum\limits_{i=1}^{k} Pr(X^n = v_i) \cdot s_i(Pr(X^n = v_i))}$$

Since v_l^n and v_l^{n+1} are functions on the frequencies of all variants, the selection coefficient s_l also turns out the be a function of these and by this s_l is interactively frequency dependent.

How do we estimate the equilibrium state for the social learning variant v_l, i.e. what happens if $X^{n+1} = X^n = \hat{X}$? The situation is quite difficult to analyze. However, if, *de facto*, no change takes place any more, i.e., if there really is an equilibrium, then $v_l^{n+1} = v_l^n = \hat{v}_l$. Now, optimality results of social learning via meta-induction state that the normalized success of a learning variant v_l approaches the normalized success of the best variant(s) in the limit if the loss function is convex. Since $d_n(i)$ is a linear combination of convex functions, $d_n(i)$ is also convex, so we can employ the following optimality result (cf. Schurz 2008, p.297; for a simplified proof cf. Feldbacher-Escamilla 2020, appendix):

$$lim_{n \longrightarrow \infty}(as_n(l) - as_n(i)) \geq 0 \quad (1 \leq i \leq k)$$

[257] Since the equilibrium state equals the limiting case, this optimality result holds for the equilibrium state too:

$$Pr(\hat{X} = \hat{v}_l) \geq Pr(\hat{X} = v_i) \quad (1 \leq i \leq k)$$

So, in the equilibrium state the normalized success of the learning variant v_l is at least as high as that of the best fitted non-learning variant. If we identify the best fitted variant in this state with v_b, then we get—in analogy to the result of Boyd and Richerson ($\mu(\hat{X}) = s$)—the limiting result for the learning-variant:

$$\hat{v}_l = v_b$$

However, the model is not restricted to the equilibrium case, with a best fitted variant v_b; v_l proves to be long-run optimal regarding any development of the variants, even if the frequencies of the variants do not converge to a limit, but oscillate (cf. the optimality results in Schurz 2008). Recall the learning target of Boyd and Richerson's model, namely to end up with (ideal) s or approximate it: In contrast to this, our model of blending inheritance (v_l) aims at optimization in terms of approximating the best variants or even outperforming them.

4.2 Some Provisos and a Simulation

We had to make some assumptions in order to transfer the optimality-result for the equilibrium state $\hat{v}_l = v_b$. In the following we will provide a brief discussion of those assumptions.

First, the blended variant is guaranteed to fare well in the long run with respect to all the other variants. Whereas in Boyd and Richerson's model, s can be the *truly maximal/ideal state of the system*, in our model, s matches only the best variants within the system. It is possible that, in reality, another variant would fare much better, but v_l may never approach it, because it only blends accessible variants. So, a wide range of *accessible* variants is presupposed in our model in order for v_l to perform well. Note, however, that the assumption of a constant absolute population size guarantees at least some minimal fitness.

Since in evolutionary theory a comparative stance regarding success is much more common than an absolute stance, this restriction should not be a problem.

Second, blending is very strong regarding the *informational basis*: All past frequencies of accessible variants must be considered in calculating the weights for v_l's blending. Furthermore, since we assume *adaptive* blending, the learning variant also needs to be able to identify relevant features of the environment and partition the frequencies to different types of environments. For real cases, a restriction of such a strong assumption has to be considered.

Third, the optimality of variant v_l with respect to the other variants $v_1, \ldots, v_k$ in terms of relative frequencies is guaranteed only for the long run, i.e. in the limit. In the short run, the performance of v_l depends on the exact evolutionary system under investigation. [258] However, one can derive exact boundaries for the short run performance of v_l (cf. Schurz 2008, p.297): $Pr(X^n = v_{1 \leq i \neq l \leq k}) - Pr(X^n = v_l^n) \leq \sqrt{(k-1)/n}$. Short run performance of relative-success-based blending is guaranteed to be within this bound; however, with the help of simulations one might also study improved performance of such blending. We conclude this section with a simulation to illustrate relative-success-based blending inheritance.

Figure 9 illustrates a simulation of a case of blending via relative-success-based weighting: The learning variant v_l weights the better performing variants according to their past success in terms of relative frequencies. In the first diagram, the development of v_l is illustrated. Regarding relative frequencies, the "cake is sliced", of course, which means that if v_l equals one of the other variants, the success rate in terms of relative frequencies also equals that of the other variant. This fact is depicted in the second diagram. The slice of frequencies is corrected in the third diagram, where the relative frequencies of equal variants (values which, strictly speaking, cannot be distinguished) are added up. There, the learning variant v_l coincides with one of the two variants of the setting (v_1, v_2), except in cases where a shift in the relative-success-rates occurs. In such cases, v_l represents *proper* blending. It is rewarded with a bonus—here modeled with a fixed parameter of $+5\%$ population size compared to the size of the better of the two variants v_1 and v_2. Note that this parameter models the assumption that blending is *adaptive*. The so-defined learning variant v_l can be shown to be *expectation-optimal* (cf. Schurz 2019, sect.6.7.1) also for the case of *random* blending. However, to flesh out such an expansion will be the subject of future investigation.

5   Conclusion

We have examined the phenomenon of blending inheritance within the framework of cultural evolution. Although cultural and natural evolution may share some relevant core properties on a general level of description—the applicability of an *evolutionary algorithm* consisting of variation, selection, and reproduction—they also differ in some crucial and less crucial aspects. One of the differences regards inheritance. Whereas in natural evolution inheritance
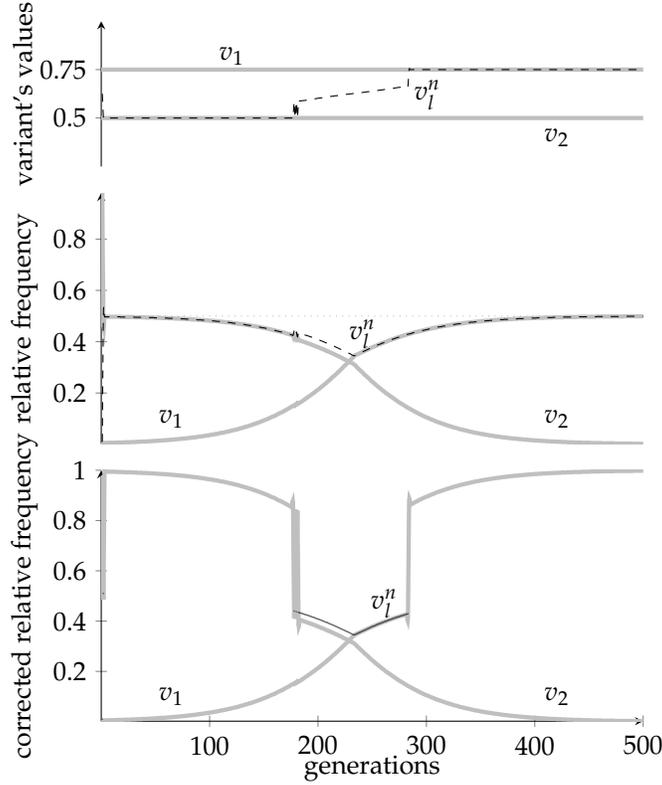
Figure 9: Cultural evolution with reproduction, selection and no mutation with blending inheritance via relative-success-based weighting in an interactive frequency dependent model. Parameters: $s_1 = 1$, $s_2 = 1 - 0.02 = 0.98$, $Pr(X^0 = v_1) = 0.01$, $Pr(X^0 = v_2) = 0.66$; the variants are $v_1 = 0.75$, $v_2 = 0.5$; the mutation rates are 0: $m_{v_1 \longrightarrow v_2} = m_{v_2 \longrightarrow v_1} = 0$; v_l starts with the frequency of v_1, i.e. $Pr(X^0 = v_l^0) = 0.01$; if v_l^n is not blended (unweighted imitation of variant v_1 or v_2), den $Pr(X^n = v_l^n)$ equals the frequency of the imitated variant v_1 or v_2: $Pr(X^n = v_1)$ or $Pr(X^n = v_2)$; if v_l^n is blended, i.e. weighting v_1 and v_2, then v_l^n has a blending advantage/bonus of +5% over the better variant; this implements the assumption that blending is adaptive; the *fallback option* of v_l^n, i.e. its value if none of the variants has an attractivity above 0 is v_l^{n-1}. (Data generated with Perl.)

[259] consists of a transmission of discrete units, in the cultural realm inheritance happens, amongst others, via blending. Following Mesoudi (2011) and the seminal work of Boyd and Richerson (1988), *blending* is a process of mixing ideas and behavior in a non-random manner. Such cultural transmission is often constituted by individuals who try to estimate which behavior of other individuals in their immediate environment may have been favored by selection in previous generations. This fact is captured by several forms of *learning dynamics* in general, and forms of average imitation in particular.

In the second part of the paper, we took a closer look at the formal structure of cultural transmission. To implement a proper formal model which captures

as many facets of cultural transmission as possible and feasible, we proposed the merger of two famous approaches to cultural evolution, the classical statistical model of information transmission by Boyd and Richerson (1988), and population dynamical models as presented e.g. by Schurz (2011). [260]

The first model allows for blending inheritance and the achievement of an ideal state *s* of a system, if such a state exists. The second model allows for so-called *interactive frequency dependency*, a key feature of cultural transmission, but does not implement blending inheritance. A combination of the two models captures as many facets as possible.

Our formal expansion allows not only for *reflective*, but also for *interactive* frequency dependency. Not only is its selection coefficient of any given variant of a cultural trait dependent on the variant's frequency, but its fitness depends also on the frequency of other variants. It is a form of relative-success-based blending, where an agent observes the success rates of other agents, respective to the fitness of their cultural traits, and then combines them into a weighted average. A general result shows that, in the long run, the normalized success of this form of social learning or relative-success-based cultural inheritance of traits is at least as good as the fittest non-learning variant.

To conclude: Blending inheritance allows for an increase in relative frequency of a cultural trait; therefore it is rationally applicable and probably one of the main reasons for the speed of cultural evolution. If we assume that agents blend traits together under a success-guided principle of weighting, this strategy is guaranteed to produce new and—in the long run—more successful traits in cultural evolution.

Acknowledgements

References

Aldrich, Howard E., Hodgson, Geoffrey M., Hull, David L., Knudsen, Thorbjørn, Mokyr, Joel and Vanberg, Viktor J. (2008-10). 'In Defence of Generalized Darwinism'. In: *Journal of Evolutionary Economics* 18.5, pp. 577–596. DOI: 10.1007/s00191-008-0110-z.

Boyd, Robert and Richerson, Peter J. (1988). *Culture and the Evolutionary Process*. Chicago: The University of Chicago Press.

Carnap, Rudolf (1950/1962). *Logical Foundations of Probability*. London: Routledge and Kegan Paul.

Cesa-Bianchi, Nicolo and Lugosi, Gabor (2006). *Prediction, Learning, and Games*. Cambridge: Cambridge University Press. DOI: 10 . 1017 / CBO9780511546921.

Claidière, Nicolas and Sperber, Dan (2007). 'The Role of Attraction in Cultural Evolution'. In: *Journal of Cognition and Culture* 7, pp. 89–111. DOI: 10.1163/156853707X171829.

Crozier, G. K. D. (2008). 'Reconsidering Cultural Selection Theory'. In: *The British Journal for the Philosophy of Science* 59.3, pp. 455–479. DOI: 10.1093/bjps/axn018.

Feldbacher-Escamilla, Christian J. (2020). 'An Optimality-Argument for Equal Weighting'. In: *Synthese* 197.4, pp. 1543–1563. DOI: 10.1007/s11229-018-02028-1.

Fisher, Ronald A. (1930). *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.

Gigerenzer, Gerd, Todd, Peter M. and The ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. Oxford: Oxford University Press.

Griffiths, Thomas L. and Kalish, Michael L. (2007). 'Language Evolution by Iterated Learning With Bayesian Agents'. In: *Cognitive Science* 31, pp. 441–480. DOI: 10.1080/15326900701326576.

Haldane, John B.S. (1964/2008). 'A Defense of Beanbag Genetics'. In: *International Journal of Epidemiology* 37.3, pp. 435–442. DOI: 10.1093/ije/dyn056.

Henrich, Joseph, Boyd, Robert and Richerson, Peter J. (2008-06). 'Five Misunderstandings About Cultural Evolution'. In: *Human Nature* 19.2, pp. 119–137. DOI: 10.1007/s12110-008-9037-1.

Lande, Russell (1979). 'Quantitative Genetic Analysis of Multivariate Evolution, Applied to Brain: Body Size Allometry'. In: *Evolution* 33.1, pp. 402–416. DOI: 10.2307/2407630.

Lewens, Tim (2015). *Cultural Evolution: Conceptual Challenges*. Oxford: Oxford University Press.

Mesoudi, Alex (2011). *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. Chicago: University of Chicago Press.

Millikan, Ruth (1984). *Language, Thought and other Biological Categories*. Cambridge, Massachusetts: The MIT Press.

— (2005). 'On Meaning, Meaning and Meaning'. In: *On Meaning, Meaning and Meaning*. Ed. by Millikan, Ruth. Oxford: Clarendon Press, pp. 53–76.

Mühlenbernd, Roland and Franke, Michael (2014). 'Meaning, Evolution, and the Structure of Society'. In: *Proceedings of the European Conference on Social Intelligence (ECSI-2014)*. Ed. by Herzig, Andreas and Lorini, Emiliano. CEUR Workshop Proceedings 1283, pp. 28–39.

Ramsey, Grant and De Block, Andreas (2017). 'Is Cultural Fitness Hopelessly Confused?' In: *The British Journal for the Philosophy of Science* 68.2, pp. 305–328. DOI: 10.1093/bjps/axv047.

Reydon, Thomas A. C. and Scholz, Markus (2014). 'Searching for Darwinism in Generalized Darwinism'. In: *The British Journal for the Philosophy of Science* 66.3, pp. 561–589. DOI: 10.1093/bjps/axt049.

Richerson, Peter J. and Boyd, Robert (2001). 'Built For Speed, Not For Comfort: Darwinian Theory and Human Culture'. In: *History and Philosophy of the Life Sciences* 23.3/4, pp. 425–465. URL: http://www.jstor.org/stable/23332522.

— (2005). *Not by Genes Alone. How Culture Transformed Human Evolution*. Chicago: The University of Chicago Press.

Ridley, Mark (2004). *Evolution*. Oxford: Blackwell Publishing.

Schurz, Gerhard (2008). 'The Meta-Inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem'. In: *Philosophy of Science* 75.3, pp. 278–305. DOI: 10.1086/592550.

— (2011). *Evolution in Natur und Kultur. Eine Einführung in die verallgemeinerte Evolutionstheorie*. Heidelberg: Spektrum Akademischer Verlag. DOI: 10.1007/978-3-8274-2666-6.

— (2019). *Hume's Problem Solved. The Optimality of Meta-Induction*. Cambridge, Massachusetts: The MIT Press.

Sperber, Dan (1997). 'Selection and Attraction in Cultural Evolution'. In: *Structures and Norms in Science*. Ed. by Dalla Chiara, M.L., Doets, K., Mundici, D. and van Benthem, J. Volume Two of the Tenth International Congress of Logic, Methodology and Philosophy of Science, Florence, August 1995. Dordrecht: Kluwer Academic Publishers, pp. 409–426.

Witt, Ulrich (2004). 'On the Proper Interpretation of 'evolution' in Economics and its Implications for Production Theory'. In: *Journal of Economic Methodology* 11.2, pp. 125–146. DOI: 10.1080/13501780410001694091.