

Title: Combining Conformist and Payoff Bias in Cultural Evolution: An Integrated Model for Human Decision Making

Author: Ze Hong^{ab1}

Author Affiliations:

^a Department of Human Evolutionary Biology, Harvard University, 11 Divinity Avenue, 02138, Cambridge, MA, United States

^b Department of Sociology, Zhejiang University, 866 Yuhangtang Road, 310058, Hangzhou, Zhejiang Province, P.R. China

¹ To whom correspondence should be addressed: ze_hong@g.harvard.edu

Combining Conformist and Payoff Bias in Cultural Evolution: An Integrated Model for Human Decision Making

Abstract

Most existing research on transmission biases in cultural evolution has treated different biases as distinct strategies. Here I present a model that combines both frequency dependent bias (including conformist bias) and payoff bias in a single decision making calculus, and show that such integrated learning strategy may be superior to relying on either bias alone. Natural selection may operate on humans' relative dependence on frequency and payoff information but both are likely to contribute to the spread of variants with high payoffs. Importantly, the magnitude of conformist bias affects the evolutionary dynamics, and I show that an intermediate level of conformity may be most adaptive and may spontaneously evolve as it resists the invasion of low-payoff variants yet enables the fixation of high-payoff variants in the population.

1. Introduction

Unlike most animals, humans obtain a tremendous amount of information from conspecifics which contributes to our ecological dominance (Henrich, 2016; Richerson & Boyd, 2005). The transmission of information in human societies has been extensively studied as an evolutionary process both theoretically (Boyd & Richerson, 1985; Feldman & Cavalli-Sforza, 1976; Kendal et al., 2009) and empirically (Henrich & Henrich, 2010; Mesoudi, 2008). A lot of research has focused on transmission biases, the psychological tendencies of individuals to favor specific cultural variants rather than others (Henrich & McElreath, 2003). These biases can result in evolutionary dynamics that significantly differ from genetic transmission as genetic material can only be passed from parents to offspring whereas cultural information can flow through multiple transmission channels (from non-parents, peers, etc.) (Creanza et al., 2017).

Transmission biases often allow for the adaptive evolution of culture (Kendal et al., 2018) and may themselves be viewed as having a genetic basis and thus be subject to natural selection (Laland, 2004; Mesoudi, 2005). Much effort has been devoted to examining the conditions under which various transmission biases evolve (Kendal et al., 2009; Muthukrishna et al., 2016); among the proposed biases, conformist bias and payoff bias have received particular attention

(Boyd & Richerson, 2009; Denton et al., 2020; McElreath et al., 2008; Whitehead & Richerson, 2009). In the cultural evolution literature, conformist bias refers to a specific kind of frequency-dependent copying strategy where individuals adopt the most common cultural variant with probability that is higher than its actual frequency in the population (Boyd & Richerson, 1985; Henrich & Boyd, 1998). Payoff biased imitation, on the other hand, has been discussed in both economics (Schlag, 1998) and cultural evolution (Boyd & Richerson, 1985, 2009; Mesoudi & O'Brien, 2008), and generally refers to the type of copying strategy where the probability of adopting a cultural variant depends on some observed payoff (the same bias has also been called success bias, see (Baldini, 2012) and "indirect bias" in Boyd & Richerson (1985)'s original formulation). In some recent theoretical models, it has been used to describe the trait-adoption strategy where the probability of adopting a particular cultural variant is positively related (e.g. directly proportional) to its relative payoff (Baldini, 2012; J. Kendal et al., 2009).

Typically, the different transmission biases and the associated learning rules are treated as distinct strategies favored in different environmental contexts. For example, conformist transmission has been shown to be favored (compared to unbiased frequency dependent transmission) when the number of traits involved is large (Nakahashi et al., 2012) or the population size is large (Perreault et al., 2012). Similarly, payoff bias has been suggested to be an adaptive strategy when the high-payoff variant is rare and the payoff information is not very stochastic (Baldini, 2012). Such treatment allows for evolutionary stable strategy (ESS) analysis and has provided much insight for both understanding how our ancestral environmental conditions might have shaped our learning psychology and how such learning psychology might flexibly respond to specific situations that individuals may encounter in their lifetime. In reality, however, humans likely possess a suite of learning strategies and the actual decision making in a given situation may involve more than one strategy. In other words, different types of learning are not psychologically distinct processes (Heyes, 1994; Plotkin, 1988), and instead of employing individual strategies in particular learning instances, humans may combine multiple strategies into a single decision making calculus (Perreault et al., 2012).

However, there are many ways to *combine* or *integrate* different learning strategies into a single strategy (hereafter referred to as "integrated strategy"). Many existing cultural evolutionary models treat these integrated strategies as having a step-like structure; that is,

individuals may utilize strategy 1 by default but will switch to strategy 2 if certain criteria are met. For example, individuals may first attempt payoff biased imitation but will fall back to some frequency dependence if observed payoffs are tied (McElreath et al., 2008). Similarly, Boyd & Richerson (1995) model a situation where individuals first compare the payoffs of two variants at a cost (individual learning) and will imitate (social learning) if the payoff difference is not sufficiently large. In a more general case, Enquist et al. (2007)'s “critical social learning” strategy has the same structure: individuals attempt social learning first and will perform individual learning if socially acquired behavior is deemed unsatisfactory by some standards. Relatedly, in an earlier seminal work on conformist transmission (Henrich & Boyd, 1998), social learning and individual learning are attempted at different stages of individuals' life cycle. In contrast, relatively less attention has been paid to the kind of integrated strategies where information produced by transmission biases is processed simultaneously. Previous work that does take into account multiple transmission biases tends to treat payoff as feedback from the environment and analyzed specific context such as punishment (Henrich & Boyd, 2001) and innovation diffusion (Henrich, 2001).

Existing models have also primarily focused on comparing individual learning and social learning, with the goal of examining the conditions under which either kind of learning is evolutionarily advantageous (Boyd & Richerson, 1996; Nakahashi et al., 2012). It is worth noting that “social learning” is not a single strategy but refers to a multitude of ways in which individuals acquire information from others in the community (Laland, 2004), and much less effort has been devoted to understanding how various social learning biases (such as conformist bias and payoff bias) may interact with one another and their relative importance in influencing trait adoption decisions. Given our species' enormous reliance on social learning, a closer examination of these different types of social learning strategies may be particularly informative.

Empirically, much effort has been devoted to understanding how individuals apply these different learning strategies in experimental setups (McElreath et al., 2005; Morgan et al., 2012; Muthukrishna et al., 2016). In general, payoff bias is observed in learning tasks, including young children and chimpanzees (Vale et al., 2017). On the other hand, human subjects usually respond to frequency information when such information is available, though conformity is not always observed (McElreath et al., 2005). Exactly teasing apart or distinguishing different transmission

biases with experimentation may be difficult due to individual heterogeneity, yet there are good reasons to suspect that individual decision making in real-world situations involve some kind of combination of multiple learning biases, as they often serve the same inferential purpose, especially in the domain of technology.

In this paper, I build upon and complement existing work by proposing a simple model where individuals process frequency and socially acquired payoff information simultaneously in a single decision-making calculus, and examine how the “weights” (assuming they are genetically transmitted) that individuals associate with different information sources evolve. To my knowledge, there has not been any empirical evidence that humans utilize one learning strategy first and then another in a sequential manner when deciding what cultural variant to adopt, and I argue that my model setup is a more realistic description of the actual psychological mechanism of human decision making for two reasons. First, decades of research in cognitive psychology has conclusively demonstrated that humans have rich cognitive structures that process input information in rather sophisticated ways in contrast with blunt stimulus-response behaviorism (Greenwood, 1999; Miller, 2003; Pekala & Pekala, 1991), and that humans are perfectly capable of integrating different kinds of information in a single inferential process to respond flexibly and adaptively to a multi-dimensional environment (Angelaki et al., 2009; Kayser & Shams, 2015). Additionally, I also allow the magnitude of conformity to vary and evolve at the individual level, building upon the classic work of Henrich & Boyd (1998).

Instead of identifying some optimal solution, I aim to illustrate the advantage of utilizing both frequency dependent bias (including conformist bias) and payoff bias and briefly discuss the implications for human social learning. While the ultimate elucidation of the exact mechanisms of human information processing and decision making is likely to require breakthroughs in neurobiology and brain science, theoretical models that take an evolutionary approach can provide potential direction and guidance given that millions of years of evolution presumably has equipped humans with some adaptive design of information acquisition and processing (Richerson, 2019).

Unlike many theoretical models on the evolution of transmission biases, the present model assumes the environment is constant with regard to a given pair of variants in the sense that one variant has higher payoff than the other *on average*, though the actual payoff that each

individual obtains vary probabilistically. Therefore, one cultural variant may be viewed as strictly superior to the other regarding its payoff (see Henrich & Boyd (2001) for some other models with the constant environment assumption). While it is true that the payoff/fitness benefit of many cultural traits depends on environmental states (e.g. variant A confers higher payoff in state 1 but relatively lower payoff in state 2) (Richerson, 2019), our cultural capacity also needs to deal with the myriad of cultural variants whose payoffs do not necessarily depend on environments. In other words, there exist cultural traits where one variant is simply better (has higher payoff) on average than another across different environmental states. This is especially true in the domain of technology; examples include the replacement of stone tools by bronze/iron tools (Edmonds, 2003) and the numerous technological breakthroughs during the industrial revolution (Thackray, 1970). Ultimately, what this assumption says is simply that certain variants confer higher payoff than others over the long run in particular ecological/social environments, although the actual payoff that individuals obtain may vary due to idiosyncratic factors. Indeed, continued increase in adaptiveness¹ would be difficult to achieve if all variants confer exactly the same average payoff. The large repertoire of cultural items in human populations means that naive individuals may often encounter situations where she needs to evaluate alternative cultural variants and decide which one(s) to adopt, and this creates a selective environment in which individuals with a decision-making apparatus that increases their chance of adopting the high-payoff variants would be favored by natural selection.

2. Model and Results

In this stylized model, agents face a decision of adopting one of the two cultural variants (C1 and C2) that have associated payoffs which can be observed with error. I propose a straightforward algebraic way of combining frequency and payoff information into a single probabilistic decision making equation, and examine the evolutionary dynamics of cultural variants under such decision making strategy. I then consider the evolution of the relative importance (referred to as "weights" in the model) that individuals place on observed frequency and payoff under various conditions.

¹ For example, suppose X1/X2, Y1/Y2, and Z1/Z2 denote different components of the same technology; for human populations to progressively achieve higher payoff, evolution needs to figure out the superior variant for each component and "lock onto it". This process will be modelled later with an agent-based simulation approach.

2.1. Baseline Model

For analytic convenience, I take the typical assumptions of asexual reproduction and non-overlapping generations² (Day & Bonduriansky, 2011). Naive agents randomly sample a number of cultural models from the parental generation and make their adoption decision based on both the number of models possessing C1/C2 and the payoff of C1/C2.

Denote the number of C1 and C2 models in the sample n_1 and n_2 , and payoff of C1 and C2 b_1 and b_2 respectively. Assuming the payoff observation error ϵ is normally distributed with mean 0 and variance σ^2 for both variants, define a naive agent's probability of adopting C1 as (note that conformity has not been included yet)

$$Pr(C1 \text{ adopted} | n_1, n_2) = \begin{cases} 1 & \text{if } n_2 = 0 \\ 0 & \text{if } n_1 = 0 \\ \frac{(n_1) \cdot w_n + (b_1 + \epsilon) \cdot w_b}{(n_1 + n_2) \cdot w_n + (b_1 + \epsilon + b_2 + \epsilon) \cdot w_b} & \text{otherwise} \end{cases} \quad (1)$$

Where w_n and w_b represent the weight attached to observed variant frequency and observed payoff. Note that w_n and w_b theoretically can be any real number while in practice their values need to be non-negative to be sensible. Since observed payoff error ϵ has mean 0 it will be omitted in subsequent analytic formulations. This particular way of constructing the probability of adopting C1 ensures that $Pr(C1 \text{ adopted} | n_1, n_2)$ is properly bounded between 0 and 1, and the relative importance of observed frequency and payoff can be flexibly adjusted. Note that in the special case where one of the weights is zero, equation (1) either becomes frequency dependent transmission ($w_b = 0$) or payoff biased transmission ($w_n = 0$).

Note that $\frac{(n_1) \cdot w_n + (b_1 + \epsilon) \cdot w_b}{(n_1 + n_2) \cdot w_n + (b_1 + \epsilon + b_2 + \epsilon) \cdot w_b}$ may be viewed as an individuals' subjective belief of C1's efficacy (i.e. the probability that performing C1 would yield a successful outcome) in a Bayesian framework (anonymized, forthcoming) where individuals update their beliefs regarding the efficacy of various technological variants. The numerator $(n_1) \cdot w_n + (b_1 + \epsilon) \cdot w_b$ would thus denote the amount of evidence that favors C1 and the denominator $(n_1 + n_2) \cdot w_n + (b_1 + \epsilon + b_2 + \epsilon) \cdot w_b$ the total amount of evidence. Interested readers may see Anonymized

²These assumptions will also be relevant in the agent-based simulation later in the paper.

(forthcoming, note the difference in the notation of weights) for an analysis of the population dynamics of technological evolution using the same belief construction method.

In equation (1), the frequency-dependent component of $Pr(C1 \text{ adopted} | n_1, n_2)$ is unbiased (conformist bias will be added later) and the payoff-dependent component follows a version of the "proportional imitation rule" (Schlag, 1998). In this case, naive individuals compare C1 and C2's payoff which proportionally contribute to the overall probability of adopting C1/C2.

First let us examine the change in C1 frequency from one generation to the next under such adoption rule. Let the frequency of C1 at a given time be p . As the individuals choose their models randomly from the parental generation, the number of models with cultural variant C1 should follow a binomial distribution (Boyd & Richerson, 1985). In the next generation, the frequency of C1 p' is therefore

$$p' = p^n + \sum_{n_1=1}^{n-1} \left[\left(\frac{n_1 \cdot w_n + b_1 \cdot w_b}{n \cdot w_n + (b_1 + b_2) \cdot w_b} \right) \cdot \binom{n}{n_1} \cdot p^{n_1} \cdot (1-p)^{n-n_1} \right] \quad (2)$$

where n represents the total number of sampled models ($n = n_1 + n_2$).³ Simplify equation (2), we have

$$p' = \frac{n \cdot p \cdot w_n + (b_1 - (1-p)^n \cdot b_1 + p^n \cdot b_2) \cdot w_b}{n \cdot w_n + (b_1 + b_2) \cdot w_b} \quad (3)$$

In order to identify possible equilibrium, we can simply set $p' = p$. However, analytically solving equation (3) can be unwieldy; according to Abel-Ruffini theorem, there are no solution in radicals for polynomial equations of degree five or higher. Assuming $w_b \neq 0$ (realistically speaking w_b is usually larger than 0) and further simplify and re-arrange equation (3), we get

$$p^* = \frac{b_1 - (1-p^*)^n \cdot b_1 + p^{*n} \cdot b_2}{b_1 + b_2} \quad (4)$$

³ When $w_b = 0$ (i.e. probability of adopting cultural variants only affected by their frequency), equation (2) is a special case of equation (5) in Denten et al. (2020). Put more simply, this equation becomes just $p' = p$.

where p^* denotes the equilibrium C1 frequency, with $p^* = 0$ and $p^* = 1$ as solutions. Note that when the number of sampled models n is relatively large and $0 < p < 1$, we may ignore terms with p^{*n} and $(1 - p^*)^n$ and therefore equation (4) becomes

$$p^* = \frac{b_1}{(b_1 + b_2)} \quad (5)$$

The accuracy of the approximation of equation (5) depends on the magnitude of n , and a thorough exploration of the parameter space to check its validity can be found in Supplemental Material. Equation (5) shows that the frequency of C1 at equilibrium is determined by the relative payoff of the two cultural variants, independent of weights and number of models sampled (given the approximation assumption that n is large). This makes intuitive sense, as unbiased frequency dependent transmission itself does not change the relative frequency of cultural variants. In the following sections, equation (5) will also be used as a baseline condition to contrast with more complex situations that include additional parameters.

Figure 1 provides a graphical illustration of the relationship between current frequency of cultural variant C1 and change in frequency using both analytic computation (equation 2) and agent-based simulation (equation 1). As can be seen, when $n = 10$, equation (5) already provides a pretty good approximation, as the $p' - p$ curve crosses 0 right at $\frac{b_1}{b_1+b_2}$. This stable polymorphic equilibrium ($p' - p$ crossing 0 with negative slope) exists in most cases in addition to the obvious equilibrium states of $p = 0$ and $p = 1$. Note that both $p = 0$ and $p = 1$ are unstable equilibria, meaning they can only be maintained in the absence of innovation/invading variants, and slight deviation would push p away from these equilibrium states. This suggests that in the absence of additional forces, C1 and C2 will co-exist in the population regardless of the initial population composition (so long as the population does not entirely consist of C1 or C2) under unbiased frequency dependence and proportional imitation based on observed payoff⁴. Note the systematic deviation from the predicted equilibrium given by equation (5) when n is small and the payoff difference between the two variants large: most of the $p' - p$ curve in the bottom right plot is below 0, meaning variant C2 is favored. This is because of the particular way

⁴ It is worth pointing out that the evolutionary dynamics as shown in Figure 1 resembles that of anti-conformist transmission, though no conformity related bias has been introduced. This is because the payoff component causes the population to evolve towards $\frac{b_1}{(b_1+b_2)}$, and thus the rare variant can increase in frequency.

adoption probability is constructed as shown in equation (1): the overall adoption probability has both a frequency component and a payoff component, and when n is small, the relative influence of frequency information on the overall probability is smaller and therefore payoff bias will dominate, causing the population to more likely consist of variant with higher payoff (C2)⁵.

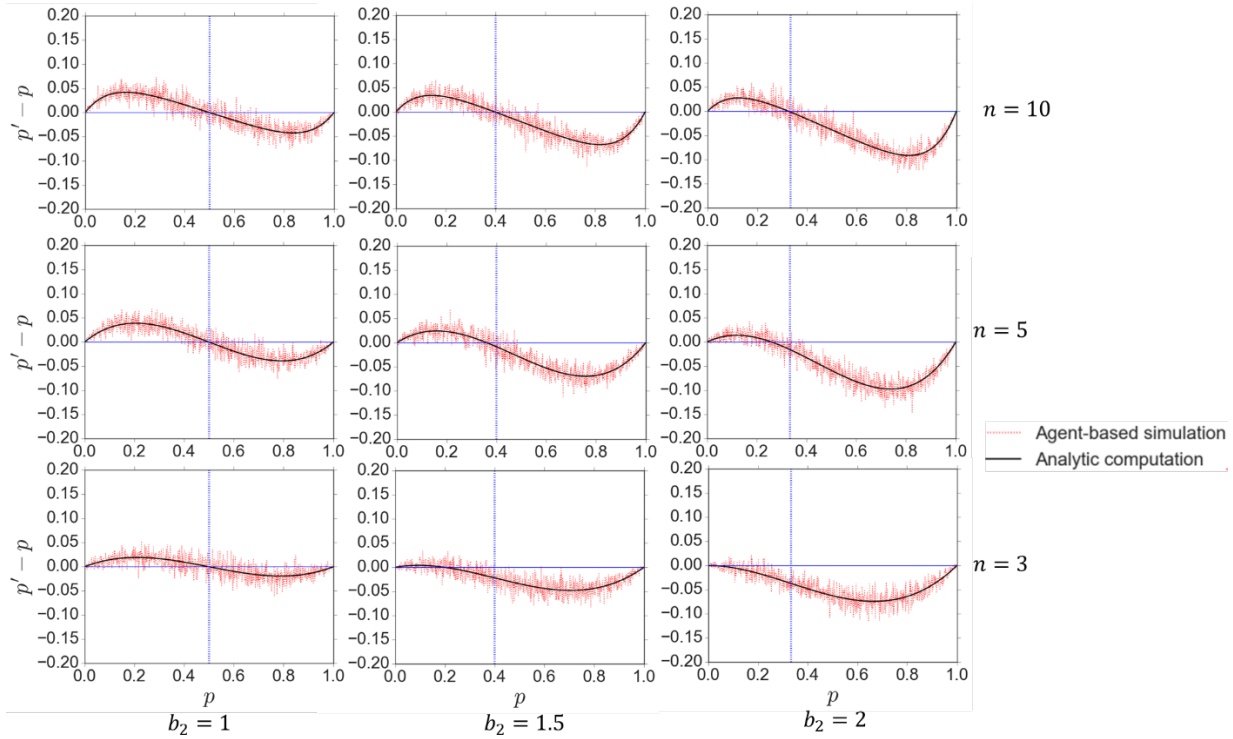


Figure 1: Relationship between current frequency p and change in frequency $p' - p$ (p' denotes the frequency in the next generation) under different parameter combinations. b_1 is fixed at 1, $\sigma^2 = 0.1$ (the variance for the payoff error term in agent-based simulation), and both w_n and w_b are set to be 1. Analytic computation and agent-based simulation are represented by black solid lines and red dotted lines respectively, and the approximated equilibrium values according to equation (5), $(\frac{b_1}{b_1+b_2})$ is marked by dotted blue lines.

2.2 Adding Conformist Bias

We now include a conformist bias parameter. Classically, conformist bias has been modeled as the probability of adopting the most common variant being its actual frequency *plus* a positive

⁵ It is worth re-iterating that this result stems from the construction of equation 1, and may not hold for other formulations.

value (denoted by D)⁶. Following the notation scheme in equation (1), an individual's probability of adopting C1 in the presence of conformist transmission bias thus becomes

$$Pr(C1 \text{ adopted} | n_1, n_2) = \begin{cases} 1 & \text{if } n_2 = 0 \\ 0 & \text{if } n_1 = 0 \\ \frac{(n_1 + D) \cdot w_n + b_1 \cdot w_b}{(n_1 + n_2) \cdot w_n + (b_1 + b_2) \cdot w_b} & \text{if } n_1 \neq 0 \text{ and } n_2 \neq 0 \text{ and } n_1 > n_2 \\ \frac{(n_1 - D) \cdot w_n + b_1 \cdot w_b}{(n_1 + n_2) \cdot w_n + (b_1 + b_2) \cdot w_b} & \text{if } n_1 \neq 0 \text{ and } n_2 \neq 0 \text{ and } n_1 < n_2 \end{cases} \quad (6)$$

Note that in order for $Pr(C1 \text{ adopted} | n_1, n_2)$ to be properly bounded between 0 and 1, D needs to satisfy the condition $-\left(\frac{n}{2} + 1\right) \leq D \leq 1$ when n is even and $-\frac{n+1}{2} \leq D \leq 1$ when n is odd (negative values of D would represent anti-conformist bias), where $n = n_1 + n_2$. If we set $w_b = 0$, i.e. agents make trait adoption decisions only based on frequency information, equation (6) becomes the classic result in cultural evolution literature where Boyd and Richerson (1985) have solved the special case when $n = 3$. Here we are interested in how conformist bias may affect the adoption of cultural variants when both payoff and frequency are taken into the individuals decision making calculus. Again, we first look at the relationship between current frequency and change in frequency under different D conditions. Given the current frequency of C1 being p , the frequency of C1 in the next generation p' can be expressed similarly as in equation (2):

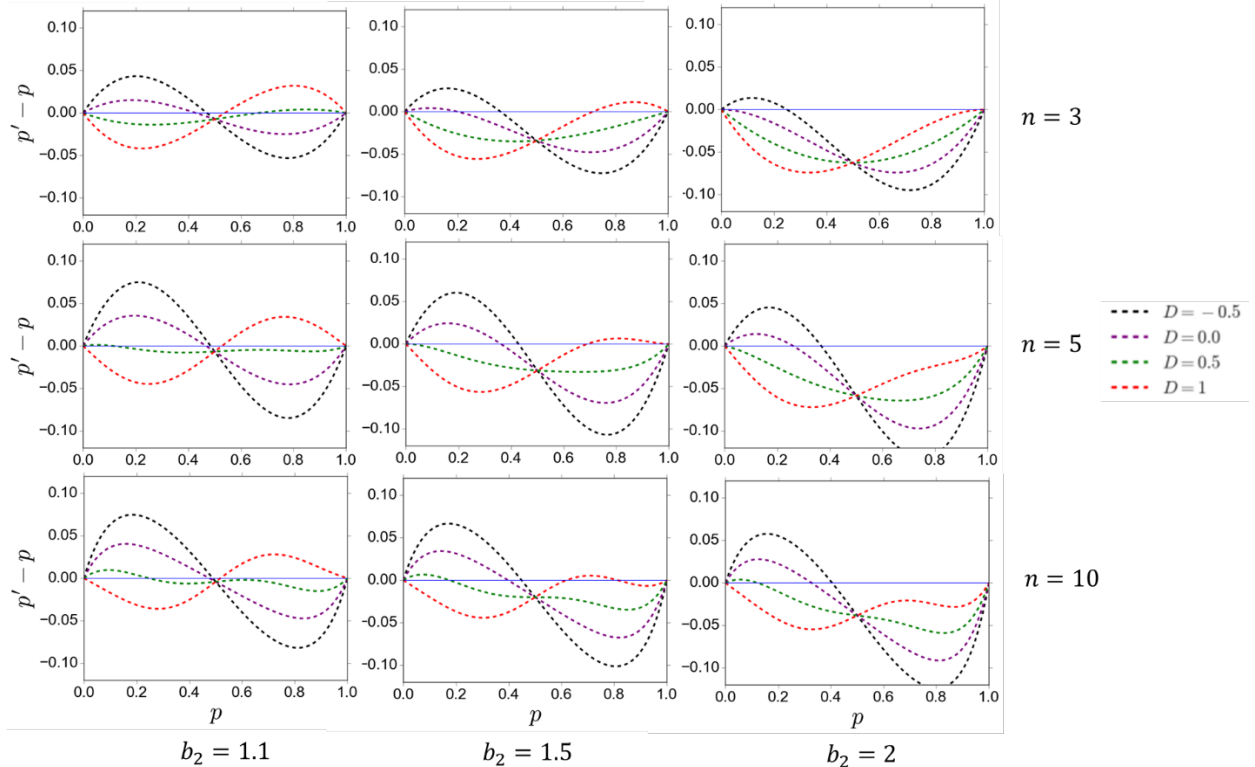
$$p' = \sum_{n_1=0}^n \begin{cases} \frac{n_1}{n} \cdot \binom{n}{n_1} \cdot p^{n_1} \cdot (1-p)^{n-n_1} & \text{if } n_1 = 0 \text{ or } n_1 = n \\ \frac{n_1 \cdot w_n + b_1 \cdot w_b}{n \cdot w_n + (b_1 + b_2) \cdot w_b} \cdot \binom{n}{n_1} \cdot p^{n_1} \cdot (1-p)^{n-n_1} & \text{if } n_1 = n/2 \\ \frac{(n_1 + D) \cdot w_n + b_1 \cdot w_b}{n \cdot w_n + (b_1 + b_2) \cdot w_b} \cdot \binom{n}{n_1} \cdot p^{n_1} \cdot (1-p)^{n-n_1} & \text{if } n_1 > n/2 \\ \frac{(n_1 - D) \cdot w_n + b_1 \cdot w_b}{n \cdot w_n + (b_1 + b_2) \cdot w_b} \cdot \binom{n}{n_1} \cdot p^{n_1} \cdot (1-p)^{n-n_1} & \text{if } n_1 < n/2 \end{cases} \quad (7)$$

where D represents the magnitude of conformist bias. Equation (7) simply says that the frequency of C1 in the next generation may be represented by the sum of probability of adopting C1 under various sample composition conditions. Note that when $n_1 = 0$ or $n_1 = n$, agents

⁶ There has also been some discussion of negative D values (i.e. anticonformity). For example, see Denton et al. (2020).

never get to experience the alternative variant and its payoff, therefore the $\frac{b_1 \cdot w_b}{(b_1 + b_2) \cdot w_b}$ component is absent in the frequency calculation of C1 in the next generation.

As in the case of unbiased transmission, we first examine the change in frequency of C1 $p' - p$ as a function of current frequency of C1, p . Figure 2 shows the expected change in frequency under different conditions, and a few remarkable features should be noted. First, as pointed out previously, there is always a stable polymorphic equilibrium ($p' - p = 0$) when the conformist bias parameter D is zero, which is determined by the relative payoff difference of the two cultural variants as well as the number of models sampled. On the other hand, when the conformist bias parameter D is sufficiently large ($D = 1$)⁷ stable polymorphic equilibrium may not exist and the population will always move towards either $p = 0$ or $p = 1$ depending on the initial frequency (as in the case of $b_2 = 1.1$). The intuition here is that because the most common variant is disproportionately favored under conformist transmission, a strong conformist bias will tend to push the common variant towards fixation.



⁷ Due to the way conformist bias is constructed here (as a single coefficient with maximum $D = 1$), a variant sampled in $\frac{n_1}{n}$ models, where $n_1 > n_2$, can be adopted with a probability of at most $\frac{n_1 + 1}{n}$. Thus, as n becomes large, the effect of conformity becomes weaker

Figure 2: Relationship between current frequency p and change in frequency $p' - p$ in the presence of conformist bias ($D > 0$), anti-conformist bias ($D < 0$), or unbiased frequency-dependent transmission ($D = 0$) under different parameter combinations. b_1 is fixed at 1, and both w_n and w_b are set to be 1. The reference line $y = 0$ is marked by the solid blue line. All values are computed according to equation (7).

What is particularly interesting, however, is that when D is of intermediate magnitude ($D = 0.5$), change in frequency can be entirely negative with $n = 3$ and $n = 5$ in the middle column ($b_2 = 1.5$), meaning that the cultural variant that confers high payoff (C2) will reach fixation regardless of the initial population composition provided that p is not equal to 1, as can be seen in the middle column ($b_2 = 1.5$) and the right column ($b_2 = 2$). This result suggests that moderate conformist transmission bias consistently favors the high payoff variant when individuals take both frequency and payoff into trait adoption decisions. Intermediate conformist bias can both resist the invasion of cultural variant with lower payoff and allow for the spread of cultural variant with higher payoff, because proportional imitation based on relative payoff favors the high-payoff variant but does not push it into fixation as the low-payoff variant still has some probability of being adopted. As the frequency of the high-payoff variant increases, conformist transmission may push it towards fixation once it becomes the more common variant ($p > 0.5$).

2.3 Evolution of Information Weights and Conformist Bias

Since the combination of payoff bias and conformist bias may lead to better population level outcomes (in terms of the adoption of the cultural variant with higher payoff), how would natural selection operate on the relative weight placed on observed frequency and payoff as well as the magnitude of conformist bias? In this section we assume weights (w_n and w_b) and conformist bias (D) are genetically transmitted and allow them to evolve, and track its evolutionary trajectory under various conditions. Notice that cultural variants now affect biological fitness whereas they did not previously. To fully explore the evolutionary dynamics, in particular the co-evolution patterns of different transmission biases, I construct an agent-based simulation that allows both the weights and the conformist bias to evolve by treating them as individual attributes. To increase realism of the setup, instead of solving one problem with a fixed solution (picking the variant with higher payoff from C1/C2), agents need to solve a relatively large number of independent problems; that is, they need pick the superior variant for multiple dichotomous cultural traits whose payoff vary in magnitude probabilistically. These traits could

be different components of the same cultural practice or technology (e.g. different aspects of arrow/bow design) or independent cultural variants that are fitness-relevant. Specifically, in each independent simulation run, each agent possesses a number of traits (20 in the simulation) with the first variant being fixed in the beginning and an “invading” variant being introduced to the population at a small percentage (5% in the simulation) to each trait in a periodic fashion. The payoff of the first variant for each pair (b_1) is set to be 1 for convenience, and the payoff of the invading variant is normally distributed with mean b_2 and variance σ_i^2 . All the following analyses are based on the statistical average of 2000 independent simulation runs to increase the robustness of the conclusions. Note that 1) in any single simulation run the payoff difference for each pair of variants is constant, meaning that for any given trait one variant is either reliably better or reliably worse; however, across different traits the “invading” variant may have payoff higher or lower than the default payoff, and the actual payoff obtained by each individual is also variable, which is modeled as a normal distribution with mean equal to the payoff of cultural variants and variance σ^2 , 2) these different pairs of variants do not appear all at once in the beginning of the simulation but occur periodically over evolutionary time. In the simulation setup below, an “invading” variant is introduced for an existing trait at some fixed interval (10 generations in simulation)⁸, meaning that throughout the simulation (200 generations total) there is a constant supply of new variants and as a result individuals need to solve these new variant adoption problems.

The above setup thus creates a scenario where agents are faced with a number of decision problems that occur at regular intervals in evolutionary time, and each of these problems has multiple potential solutions (in the model, there are two “solutions”, or cultural variants, for each trait adoption decision). The environment is stable/constant in the sense that there is a “correct” solution (the variant with higher average payoff) for all decision problems, yet changing in the sense that new decision problems are periodically introduced rather than all known from the very beginning. Generally speaking, the more ecologically valid conditions are when the invading variants have lower payoff than the existing variants ($b_2 < 0$) as many of the “invading” or new variants are the results of random or unintentional errors, yet for the sake of completeness I also

⁸ Realistically speaking, the time interval at which new variant appears involves some random component, but in the present simulation a fixed interval is used as it does not qualitatively change the nature of the simulation and allows for better visual inspection of evolutionary trends.

explored the opposite conditions when the invading variants on average confer higher payoff ($b_2 > 0$). In the analysis below, I track the temporal changes of variant frequency for C1/C2, the weight on observed frequency (w_n) and the magnitude of conformist bias (D).

Recall that the weights and conformist bias are genetically transmitted under asexual reproduction, and the life cycle of agents are modeled as a simple Wright-Fisher process with selection (Ewens, 2012), where the population is of constant size N and an agent's probability of contributing to the gene pool (reproduction) is proportional to its fitness, which is the sum of the payoffs of the variants that the individual possesses. Each individual's initial information weight on payoff (w_n) is sampled from a uniform distribution (0,2) (their information weight on payoff (w_b) is set to be constant 1), and their magnitude of conformist bias (D) from uniform distribution (-1,1). In all simulation runs the second, invading variant is introduced into the simulation at 5% frequency. Thus, we simulate a situation where individuals start with their default set of variants and new variants “attempt” to invade the population during evolutionary time.

To examine how natural selection affects w_n and D values in such a situation, I ran a large number of simulations to fully explore the parameter space (see Supplemental Material for simulation setup and parameter value details) and Figure 3 shows the evolutionary trajectory of the frequency of C1 (the first variant of the starting trait) as well as the population average w_n and D over time.

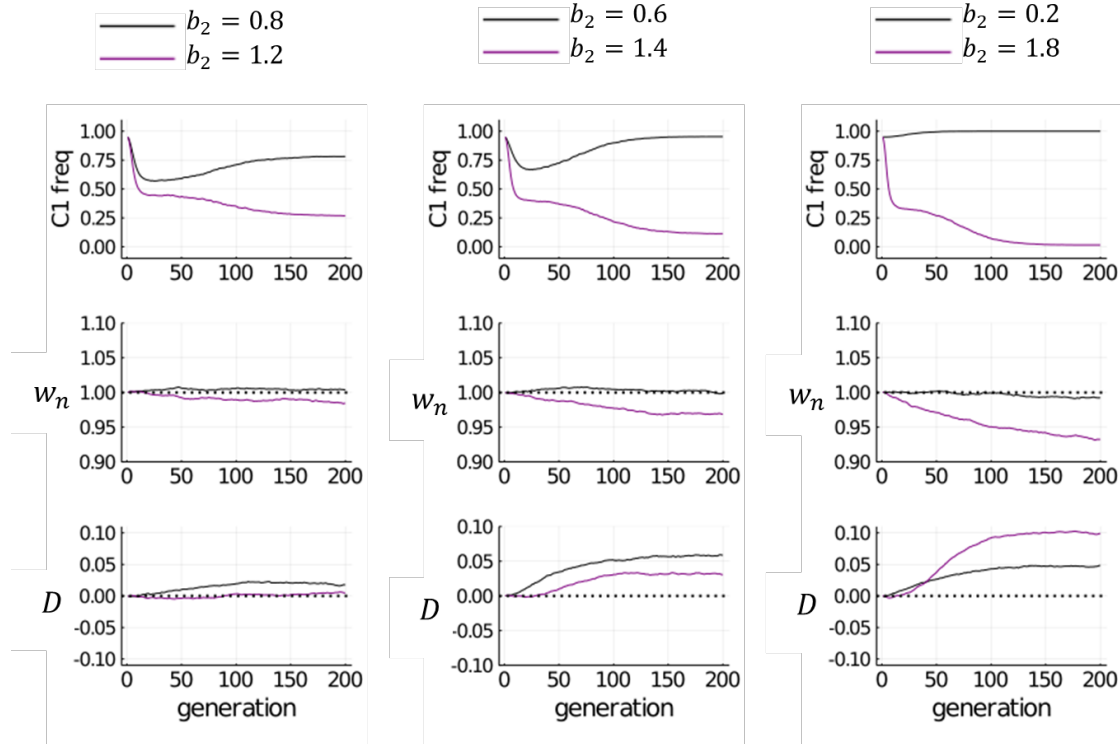


Figure 3: Temporal evolutionary trajectories of C1 frequency as well as weight on frequency information (w_n) and the magnitude of conformist bias (D) under various payoff conditions (average of 2000 independent simulations). $\sigma_i^2 = 0.3$, $\sigma^2 = 0.1$, $N = 1000$, and number of traits each agent possesses = 20 (e.g., each agent needs to make 20 independent variant adoption decisions, with new variants introduced every 10 generation). Dotted lines are references lines ($y = 0$). Other parameter values are specified in Supplemental Material.

As expected, the magnitude of conformist bias parameter D generally increases in most conditions, though a small initial decline may sometimes be observed when the advantageous variant's initial frequency is low. This is because when the advantageous variant (C2) is rare in the beginning of the simulation, individuals with anti-conformist bias (negative D) are more likely to acquire such high-payoff variant, yet once it becomes common in the population a positive conformist bias is favored by natural selection. As can be seen in the case of $b_2 = 1.4$ and $b_2 = 1.8$, the initial decline in D matches the decline in C1 frequency quite well; D starts to increase as soon as the more advantageous variant (C2) crosses over 50% in frequency. Note that the population average D never gets very large and reaches equilibrium at around 100 generations, suggesting that while a positive conformist bias is generally favored, its magnitude is limited in the presence of payoff stochasticity. This is especially true when the payoff difference between the invading variants and the existing variants is not very large, as the more

frequent variant in the population is sometimes not the one with higher payoff, and a very strong conformist bias may mistakenly lead individuals to adopt the low payoff variants. It is important to note, however, that the evolutionary dynamics that we observe is largely due to the stable average payoff difference between the variants for any given trait in a given simulation run (e.g., C1 is either reliably better than C2 or vice versa), and may not generalize to situations where the default variants do not reliably differ from the invading variant with regard to payoff.

We observe that the frequency of C1 (representing the first variant of all traits) all moves in the expected direction over evolutionary time, i.e., when C1 has higher payoff than C2 it remains at relatively high frequency and when C1 has relatively low payoff, its frequency consistently decreases (thus in the long run C2 becomes the dominant variant). Note that large payoff differences ($b_2 = 0.2$ and $b_2 = 1.8$) enables the superior variant (higher payoff on average) to quickly reach fixation. When the average payoff difference between the two variants is small (e.g., $b_2 = 0.8$ and $b_2 = 1.2$), on the other hand, the two variants may stably co-exist for two related reasons: first, the equilibrium that would be reached in the absence of conformist bias may not be sufficiently different from 0.5; secondly, the magnitude of D may never evolved to a sufficiently large degree to be able to push the more frequent variant to fixation.

Regarding the evolution of w_n , there are two points worth noting: first, it changes primarily in the beginning in each condition and then slowly declines over evolutionary time, eventually reaching effective equilibrium (see Supplemental Material for extended simulation run (generation = 400) for condition $b_2 = 0.2$ and $b_2 = 1.8$). This is because the influence of information weights on payoff/fitness is rather limited, especially when the payoff difference between the two alternative variants is small. In our simulation, when the payoff difference between the existing and “invading” variants is larger ($b_2 = 1.4$ and $b_2 = 1.8$), we do observe more/faster change in w_n . Second, the direction of change in w_n depends on the relative frequency of high/low payoff variants in the population. When the invading variants have higher payoff ($b_2 = 1.2, 1.4$ and 1.8), we see a decrease in w_n ; when they have lower payoff ($b_2 = 0.6$ and 0.8), an initial increase in w_n is observed. This is because when high-payoff variants are initially rare, it is better to rely more on payoff information (larger w_b , thus smaller w_n relatively) as the relative payoff advantage is not matched by its frequency in the population. Conversely, when the invading variants are of lower payoffs, stronger reliance on frequency (large w_n) will

be favored. These results are consistent with Baldini (2012)'s conclusion that payoff bias favors rare variants; if the high-payoff variant is rare, then a stronger reliance on payoff (small w_n) would be more adaptive. The magnitude of change in w_n , on the other hand, crucially depends on the population composition with regard to variant frequencies: as can be seen, in the case of $b_2 = 0.2$ there is almost no change in w_n . This is because when the invading variant is vastly inferior regarding its payoff to the existing variant, the population immediately reaches fixation and leaves no time for w_n to evolve. Note that in the most extreme case where all variants are fixed in the population, all individuals would have the exact same fitness regardless of their w_n (or D for that matter). Remarkably, both D and w_n reach effective equilibrium values despite the constant introduction of new traits, suggesting the robustness of these parameter values in continuously identifying adaptive cultural variants over evolutionary time.

It should be noted that the above analyses assumes neither passive mutation nor active innovation of variants (i.e. C1 and C2 never mutate into each other). In the presence of either mutation or innovation that prevents the population from reaching complete fixation, the relative magnitude of w_n will matter insofar as the population is polymorphic. Assuming the population reaches near fixation and the dominant variant has high payoff, a constant supply of low-payoff variant due to mutation/innovation will cause an increase in w_n ; that is to say, individuals that weigh more on frequency information will enjoy a fitness advantage (assuming $D > 0$). On the other hand, if the low-payoff variant reaches near fixation for whatever reason yet cannot drive the high-payoff variant into extinction, then w_n will decrease, meaning those who rely more on payoff would have higher fitness (assuming $D > 0$). The second scenario is less likely because normally we would expect the high-payoff variant to reach near fixation through a number of mechanisms. Therefore, if we consider one pair of dichotomous variants, a stronger reliance on observed frequency (w_n) would be favored on average in settings where the transmission fidelity is not 100% due to transmission errors or certain individuals consciously experimenting different variants. However, we need to keep in mind that 1) as w_n gets larger, the marginal fitness benefit that it confers declines dramatically, and 2) though not as common, there are cases where the low-payoff variant dominates the population and some consideration of the payoff difference would be advantageous.

3 Discussion

3.1 Combining Frequency and Payoff Bias in Social Learning

Although there hasn't been a lack of theorizing of transmission biases in cultural evolution, most published work has treated these biases as distinct strategies and aims to identify evolutionary stable strategies. In this paper I show the evolutionary dynamics of dichotomous cultural traits when individuals combine frequency and payoff information into a single decision-making calculus, and how natural selection may have selected for the weights of the two information input sources as well as the magnitude of conformist bias. In the agent-based simulation, instead of focusing on one pair of variants with fixed payoff differences, it involves a relatively large number of independent traits with stochastic payoffs and individuals need to make trait adoption decisions using their genetically inherited w_n/w_b and D . As such, this simulation offers some insights regarding how different learning biases may co-evolve to robustly solve multiple dynamic problems, and provides mathematical/computational rigor to previous verbal arguments regarding how payoff bias and conformist bias might interact with each other (Sterelny, 2007).

Unbiased frequency-dependent transmission leads the population into a state that resembles Hardy-Weinberg equilibrium in population genetics (Meirmans, 2018) where frequencies of variants remain constant. Conformist bias by definition favors the more common variant and pushes the frequencies of variants towards the absorbing states $p = 0$ and $p = 1$. Intuitively, the caveat of conformist-biased transmission is that a previously adaptive cultural variant may become non-adaptive due to environmental change, and as a result the population may get stuck in a sub-optimal condition. Some theoretical work even suggests that relying on conformist transmission alone may lead to population collapse under certain circumstances (Whitehead & Richerson, 2009).

Payoff bias on the other hand looks like an attractive alternative learning strategy as it directly compares different cultural variants and favors the one with higher payoff. Assuming payoff is statistically associated with fitness⁹ (see Baldini (2012) for instances where payoff may be quite dissociated from fitness), a reliance on payoff information may confer fitness advantage by adopting the high-payoff variant. However, payoff may be noisy and assessing it may involve

⁹ Note that fitness was only incorporated in the agent-based simulation (section 2.3.), and not the first part of the model (section 2.1. and 2.2.).

a cost¹⁰ (Nakahashi et al., 2012). My model shows that a straightforward algebraic combination of these biases may prove a superior strategy: payoff bias tends to increase the frequency of the high-payoff variant, and once it passes 0.5 conformist bias may push it towards fixation.

In the simulation, we observe that an intermediate level of conformist bias robustly evolves as it increases agents' probability of adopting the superior variant that usually gets pushed to frequencies above 50% due to payoff bias. The effect of conformist bias's magnitude on the evolution dynamics is worth reiterating. If it is too small it cannot effectively help the high-payoff variant that is already common in the population to reach fixation, and if it is too large a rare high-payoff variant may never have a chance to pass 0.5 in frequency. In this integrated decision-making calculus, the adaptive synergistic interaction between the two biases requires the magnitude of D to be within a particular range. In fact, D 's magnitude likely matters whenever conformist bias is combined in this way with some additional mechanism that makes the adoption of the variant with higher payoff more likely, as the initial spread of high-payoff variant is always suppressed by strong conformity. We should also be aware of the fact that the conformist bias parameter may be subject to natural selection in contexts other than variants adoption where payoffs are only associated with the variants themselves; for example, a strong conformist bias may be especially advantageous in the domain of normative behaviors, as rare mutants may suffer severe fitness costs in the form of punishment (Nakao & Machery, 2012).

In an earlier seminal paper on conformist transmission, Henrich and Boyd (1998) conclude that maximal conformity is favored under a broad range of conditions. This is not surprising, as the results of simulation studies that explore the evolution of conformity depend sensitively on the precise design of the simulations (Denton et al., 2020). In this case, the different conclusions are due to a crucial difference in the model setup: the payoffs that individuals obtain in Henrich and Boyd (1998) are a result of independent individual learning and as such individuals are always comparing the payoffs of different cultural variants, whereas in my model there is no individual experimentation and payoff information is always obtained culturally, and therefore the payoff of the variant not possessed by one's cultural models is simply not available to the naive individual for payoff evaluation. This means that in the present setup if one variant reaches high frequency most naïve individuals may never experience the

¹⁰ Though such cost is not incorporated in the present model.

payoff of the alternative variant, and in cases where the rare variant happens to have higher payoff, the population may nonetheless push the common variant (with lower payoff) towards fixation in the presence of strong conformist bias.

Therefore, the population level consequences of conformity hinge upon the extent to which accurate payoff information can be obtained either through independent individual learning or some other payoff-revealing mechanisms. While individual trial-and-error learning certainly occurs, I suggest that the payoff information of many cultural items cannot be realistically evaluated and compared by potentially costly individual learning for two reasons. First, naïve individuals may not even be aware of the existence of alternative variants if they do not observe others possessing these variants; second, “trying” both variants and comparing their payoffs may not always be feasible; for example, comparing the efficacy of two illness treatments through individual learning requires one to intentionally get herself ill twice (presumably at different times). In fact, Henrich and Boyd (1998) do show that when individual learning is highly error-prone (small ρ) conformist bias does not evolve to large values. My analysis here thus complements previous work by showing that in cases where payoff information is obtained culturally, a very strong conformist bias may not be optimal at the population level.

In general, integrated strategies, broadly defined, can be more adaptive than individual strategies (Enquist & Ghirlanda, 2007; Laland, 2004). In Enquist et al. (2007)'s approach, for example, “critical social learners” (individuals with the strategy of attempting social learning first and perform individual learning if social learning proves unsatisfactory) are more likely to obtain the more advantageous variant compared to individual learners. However, in the absence of conformist bias the advantageous variant rarely reaches fixation in the population, and a stable polymorphic equilibrium of the two cultural variants exists. It is likely that the critical social learners (as well as individuals with other integrated strategies) may perform even better regarding their probability of obtaining the high payoff variant if they incorporate some form of intermediately strong conformist bias in their trait adoption decision in light of the result of the present model. This is because the conformist component in their decision-making favors the more frequent variant, and large proportion of critical social learners in the population are likely to cause the high payoff variant to be more frequent. It should be kept in mind, however, that

integrated strategies may be more computationally intensive and thus requires costly neural machinery. The present model does not assume any cost in employing integrated strategies, and further theoretical work may take this factor into consideration, ideally in light of neurological mechanisms in human learning.

3.2 How Many Cultural Models to Learn from? The Effect of n

Although most published theoretical models include some mentioning of the effect of n , systematic treatment remains scant (see Denton et al. (2020) and Perreault et al. (2012) for some exceptions). This is probably because the role n plays is different under different sets of assumptions. In my model, n is directly involved in the approximation that leads to equation (5). Here a large n ensures that most naive agents will have both C1 and C2 models in their sample; in other words, a relatively large n ensures that naive agents have a chance to experience the payoff of both C1 and C2 variants. When most naïve learners in the population have a chance to sample both C1 and C2 models, a decreasing n means that the relative influence of the frequency dependent component diminishes as a result of equation (1), except in the most extreme case where $n = 1$ (each agent randomly picks one model from the parental generation) payoff information becomes entirely irrelevant and we end up with a special case of the unbiased frequency-dependent transmission.

Assuming no cost is incurred, sampling more individuals to learn from should always lead to better inferences, particularly from a Bayesian perspective. Perreault et al. (2012) point out that in their model a larger n sometimes leads to worse inference, but suggests that this is an artefact due to agents' relatively simple priors. In reality, people likely do not rigidly sample a fixed number of individuals but rather make trait adoption decisions based on idiosyncratic personal experiences which include both individual learning and social informational inputs. Therefore, n may be determined not by evolved human preferences but by external factors such as population size and interconnectedness, both of which have been shown to be important in human cultural evolution (Henrich, 2004). As such, demographic details of our ancestral population may be needed to better understand the evolutionary dynamics and outcomes of social learning strategies as affected by the number of models picked.

3.3 Evolution of Information Weights and Implications for Modeling Human Decision-making

In this stylized setting, then, how would natural selection operate on the weight of frequency and payoff information respectively when they are allowed to co-evolve with conformist bias? My results show that when one cultural variant on average confers higher payoff than the other variant, the direction and magnitude of change in the relative weights crucially depend on population composition (in particular, whether the high-payoff variant is the common variant or not). Therefore, given that humans likely encounter many such situations where they need to adopt some cultural variant among a number of candidates during their lifetime, it is unlikely that there is a single optimal weight across different learning situations. What is clear, however, is that both w_n and w_b are likely to be positive. In other words, it is better to take both frequency and payoff into consideration, though the degree to which frequency and payoff matters may vary in domain-specific ways.

Intuitively, it seems a bad idea to completely ignore frequency or payoff information when it is available, and humans may adaptively and flexibly evaluate these different types of information depending on their prior beliefs and the specificities of the situation. In a way, what is presented in the paper is a proof-of-concept model showing that *some* way of combining observed frequency and observed payoff in a single decision-making calculus can be more advantageous than strategies discarding either information. My own fieldwork suggests that both types of information often feed into the same inferential process; for example, when evaluating some healing practice, the Wa and Yi people in southwest China frequently use "many people in the community use it" and "it worked on my friends" as reasons for its efficacy (Hong, unpublished). In the domain of technology where means-ends reasoning dominates, transmission biases often manifest themselves as cognitive processes of integrating information from different sources into a single inference, i.e. the cultural variants' efficacy or effectiveness (Hong & Henrich, 2021).

In this model, payoff bias is implemented using the "proportional imitation" rule; Schlag (1998) has shown that this way of incorporating payoff information into decision making outperforms a number of other payoff-based learning strategies. But more generally, why not just compare payoffs when making trait adoption decisions? In addition to the aforementioned noise

and cost, payoff bias may not be the panacea that applies to all learning situations for two other reasons. First, many cultural traits do not have obvious payoffs associated with them. In fact, people often do not understand why particular actions are performed (Henrich, 2016) or the causal mechanisms underlying seemingly purposive actions (Derex et al., 2019). Second, people may simply have the wrong payoff associated with cultural practices. For example, people in small scale societies have tried all kinds of methods to induce rain (Frazer, 1890) and payoff biased imitation would lead one to nowhere as none of the methods had any real influence on weather. Thus, the applicability of payoff bias as a general learning mechanism may be limited and sole reliance on it can be non-adaptive.

Actual human information processing and decision making are complicated and likely affected by a wide range of factors. Future evolutionary theorizing of human social learning may benefit from explicitly considering various types of information feeding into the same computational process and treating human decision making not as rigidly implementing pre-programmed rules but as flexibly influenced by the context. Models constructed with more psychological realism may then be empirically tested and the explanatory as well as predictive power of different learning models could be contrasted and compared to enhance our understanding of the important phenomenon of human social learning.

4 Conclusion

I have presented a model showing that in settings where one cultural variant strictly confers higher payoff than the other variant, combining frequency dependent bias and payoff bias in a single decision-making calculus can be more advantageous than employing either strategy alone, and that intermediate level of conformist bias may be particularly important in balancing resisting the invasion of low-payoff variants and helping the spread of high-payoff variants. The insights here are generally applicable when conformist bias is coupled with some other mechanism that favors the adoption of the high-payoff variant.

Code Availability

The graphical representations of all equations are created using python 3.7. The final agent-based simulation for the evolution of epistemic weights (shown in Figure 3) is created using Julia 1.5.0. All codes are available at https://github.com/kevintoy/epistemic_weight_evo.

Acknowledgement

I thank Dr. Joseph Henrich for his continued encouragement and support for this project, and Peter Park and two anonymous reviewers for their very constructive feedback on earlier drafts of this manuscript.

Conflict of Interest

None.

References

- Angelaki, D. E., Gu, Y., & DeAngelis, G. C. (2009). Multisensory integration: psychophysics, neurophysiology, and computation. In *Current Opinion in Neurobiology*.
<https://doi.org/10.1016/j.conb.2009.06.008>
- Baldini, R. (2012). Success-biased social learning: Cultural and evolutionary dynamics. *Theoretical Population Biology*. <https://doi.org/10.1016/j.tpb.2012.06.005>
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*. [https://doi.org/10.1016/0162-3095\(94\)00073-G](https://doi.org/10.1016/0162-3095(94)00073-G)
- Boyd, R., & Richerson, P. J. (2009). Voting with your feet: Payoff biased migration and the evolution of group beneficial behavior. *Journal of Theoretical Biology*.
<https://doi.org/10.1016/j.jtbi.2008.12.007>
- Boyd, R., & Richerson, P. P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*.
- Creanza, N., Kolodny, O., & Feldman, M. W. (2017). Cultural evolutionary theory: How culture evolves and why it matters. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1620732114>
- Day, T., & Bonduriansky, R. (2011). A Unified Approach to the Evolutionary Consequences of Genetic and Nongenetic Inheritance. *The American Naturalist*.
<https://doi.org/10.1086/660911>

- Denton, K. K., Ram, Y., Liberman, U., & Feldman, M. W. (2020). Cultural evolution of conformity and anticonformity. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.2004102117>
- Dereux, M., Bonnefon, J.-F., Boyd, R., & Mesoudi, A. (2019). Causal understanding is not necessary for the improvement of culturally evolving technology. *Nature Human Behaviour*, 1. <http://www.nature.com/articles/s41562-019-0567-9>
- Edmonds, M. (2003). Stone tools and society: Working stone in neolithic and bronze age Britain. In *Stone Tools and Society: Working Stone in Neolithic and Bronze Age Britain*. <https://doi.org/10.4324/9780203481080>
- Enquist, M., Eriksson, K., & Ghirlanda, S. (2007). Critical Social Learning: A Solution to Rogers's Paradox of Nonadaptive Culture. *American Anthropologist*. <https://doi.org/10.1525/aa.2007.109.4.727>
- Enquist, M., & Ghirlanda, S. (2007). Evolution of social learning does not explain the origin of human cumulative culture. *Journal of Theoretical Biology*, 246(1), 129–135.
- Ewens, W. J. (2012). *Mathematical population genetics 1: theoretical introduction* (Vol. 27). Springer Science & Business Media.
- Feldman, M., & Cavalli-Sforza, L. (1976). Cultural and biological evolutionary processes, selection for a trait under complex transmission. *Theoretical Population Biology*. [https://doi.org/10.1016/0040-5809\(76\)90047-2](https://doi.org/10.1016/0040-5809(76)90047-2)
- Frazer, J. G. (1890). *The Golden Bough: A Study in Comparative Religion, Volume 2* (Vol. 2). Macmillan.
- Greenwood, J. D. (1999). Understanding the “cognitive revolution” in psychology. *Journal of the History of the Behavioral Sciences*. [https://doi.org/10.1002/\(SICI\)1520-6696\(199924\)35:1<1::AID-JHBS1>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1520-6696(199924)35:1<1::AID-JHBS1>3.0.CO;2-4)
- Henrich, J. (2001). Cultural Transmission and the Diffusion of Innovations: Adoption Dynamics Indicate That Biased Cultural Transmission Is the Predominate Force in Behavioral Change. *American Anthropologist*. <https://doi.org/10.1525/aa.2001.103.4.992>

- Henrich, J. (2004). Demography and Cultural Evolution: How Adaptive Cultural Processes Can Produce Maladaptive Losses—The Tasmanian Case. *American Antiquity*.
<https://doi.org/10.2307/4128416>
- Henrich, J. (2016). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
<https://psycnet.apa.org/record/2016-18797-000>
- Henrich, J., & Boyd, R. (1998). The Evolution of Conformist Transmission and the Emergence of Between-Group Differences. *Evolution and Human Behavior*.
[https://doi.org/10.1016/S1090-5138\(98\)00018-X](https://doi.org/10.1016/S1090-5138(98)00018-X)
- Henrich, J., & Boyd, R. (2001). Why people punish defectors: Weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *Journal of Theoretical Biology*. <https://doi.org/10.1006/jtbi.2000.2202>
- Henrich, J., & Henrich, N. (2010). The evolution of cultural adaptations: Fijian food taboos protect against dangerous marine toxins. *Proceedings of the Royal Society B: Biological Sciences*. <https://doi.org/10.1098/rspb.2010.1191>
- Henrich, J., & McElreath, R. (2003). The Evolution of Cultural Evolution. *Evolutionary Anthropology*. <https://doi.org/10.1002/evan.10110>
- Heyes, C. M. (1994). Social learning in animals: Categories and mechanisms. In *Biological Reviews of the Cambridge Philosophical Society*. <https://doi.org/10.1111/j.1469-185X.1994.tb01506.x>
- Hong, Z., & Henrich, J. (2021). The Cultural Evolution of Epistemic Practices. *Human Nature*, 1–30.
- Kayser, C., & Shams, L. (2015). Multisensory Causal Inference in the Brain. In *PLoS Biology*.
<https://doi.org/10.1371/journal.pbio.1002075>
- Kendal, J., Giraldeau, L. A., & Laland, K. (2009). The evolution of social learning rules: Payoff-biased and frequency-dependent biased transmission. *Journal of Theoretical Biology*.
<https://doi.org/10.1016/j.jtbi.2009.05.029>

- Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social Learning Strategies: Bridge-Building between Fields. In *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2018.04.003>
- Laland, K. N. (2004). Social learning strategies. In *Learning and Behavior*. <https://doi.org/10.3758/bf03196002>
- McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences*. <https://doi.org/10.1098/rstb.2008.0131>
- McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten, E., Efferson, C., & Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*. <https://doi.org/10.1016/j.evolhumbehav.2005.04.003>
- Meirmans, P. G. (2018). Hardy-weinberg equilibrium. In *Encyclopedia of Ecology*. <https://doi.org/10.1016/B978-0-12-409548-9.10555-X>
- Mesoudi, A. (2005). The Transmission and Evolution of Human Culture. *Evolution*.
- Mesoudi, A. (2008). An experimental simulation of the “copy-successful-individuals” cultural learning strategy: adaptive landscapes, producer-scrounger dynamics, and informational access costs. *Evolution and Human Behavior*. <https://doi.org/10.1016/j.evolhumbehav.2008.04.005>
- Mesoudi, A., & O’Brien, M. J. (2008). The cultural transmission of great basin projectile-point technology I: An experimental simulation. *American Antiquity*. <https://doi.org/10.1017/S0002731600041263>
- Miller, G. A. (2003). The cognitive revolution: A historical perspective. In *Trends in Cognitive Sciences*. [https://doi.org/10.1016/S1364-6613\(03\)00029-9](https://doi.org/10.1016/S1364-6613(03)00029-9)
- Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., & Laland, K. N. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*. <https://doi.org/10.1098/rspb.2011.1172>

- Muthukrishna, M., Morgan, T. J. H., & Henrich, J. (2016). The when and who of social learning and conformist transmission. *Evolution and Human Behavior*.
<https://doi.org/10.1016/j.evolhumbehav.2015.05.004>
- Nakahashi, W., Wakano, J. Y., & Henrich, J. (2012). Adaptive Social Learning Strategies in Temporally and Spatially Varying Environments: How Temporal vs. Spatial Variation, Number of Cultural Traits, and Costs of Learning Influence the Evolution of Conformist-Biased Transmission, Payoff-Biased Transmissio. *Human Nature*.
<https://doi.org/10.1007/s12110-012-9151-y>
- Nakao, H., & Machery, E. (2012). The evolution of punishment. In *Biology and Philosophy*.
<https://doi.org/10.1007/s10539-012-9341-3>
- Pekala, R. J., & Pekala, R. J. (1991). The Cognitive Revolution in Psychology. In *Quantifying Consciousness*. https://doi.org/10.1007/978-1-4899-0629-8_4
- Perreault, C., Moya, C., & Boyd, R. (2012). A Bayesian approach to the evolution of social learning. *Evolution and Human Behavior*.
<https://doi.org/10.1016/j.evolhumbehav.2011.12.007>
- Plotkin, H. C. (1988). *The role of behavior in evolution*. MIT press.
- Richerson, P. J. (2019). An integrated bayesian theory of phenotypic flexibility. *Behavioural Processes*. <https://doi.org/10.1016/j.beproc.2018.02.002>
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone : how culture transformed human evolution*. University of Chicago Press.
- Schlag, K. H. (1998). Why Imitate, and If So, How? *Journal of Economic Theory*.
<https://doi.org/10.1006/jeth.1997.2347>
- Sterelny, K. (2007). SNAFUS: An Evolutionary Perspective. *Biological Theory*.
<https://doi.org/10.1162/biot.2007.2.3.317>
- Thackray, A. (1970). Science and Technology in the Industrial Revolution. *History of Science*.
<https://doi.org/10.1177/007327537000900104>
- Vale, G. L., Flynn, E. G., Kendal, J., Rawlings, B., Hopper, L. M., Schapiro, S. J., Lambeth, S.

P., & Kendal, R. L. (2017). Testing differential use of payoff-biased social learning strategies in children and chimpanzees. *Proceedings of the Royal Society B: Biological Sciences*. <https://doi.org/10.1098/rspb.2017.1751>

Whitehead, H., & Richerson, P. J. (2009). The evolution of conformist social learning can cause population collapse in realistically variable environments. *Evolution and Human Behavior*. <https://doi.org/10.1016/j.evolhumbehav.2009.02.003>