PNAS RESEARCH ARTICLE | EVOLUTION ANTHROPOLOGY

Weak individual preferences stabilize culture

Alberto Acerbi[a,1] and Benoît de Courson[b]

Edited by Marcus Feldman, Stanford University, Stanford, CA; received June 20, 2024; accepted January 9, 2025

Outcomes in the cultural arena are due to many factors but are there general rules that can suggest what makes some cultural traits successful and others not? Research in cultural evolution theory distinguishes factors related to social influence (such as copying from the majority, or from certain individuals) from factors related to individual, nonsocially influenced, propensities such as evolved cognitive predispositions, or physical, biological, and environmental constraints. Here, we show, using analytical and individual-based models, that individual preferences, even when weak, determine the equilibrium point of cultural dynamics when acting together with nondirectional social influence in three out of four cases we study. The results have implications regarding the importance of keeping into account individual-level, nonsocial, factors, when studying cultural evolution, as well as regarding the interpretation of cross-cultural regularities, that must be expected, but can be product of weak directional forces, intensified by social influence.

cultural evolution | social influence | cultural transmission | modelling

Significance

We readily attribute cultural phenomena to social influence forces, like conformity or copying prestigious people. However, when social influence forces act together with stable nonsocial forces that influence preferences for certain cultural traits, our models show that the latter often determine which cultural trait will be successful. Stable nonsocial preferences, differently from social forces, act consistently in the same direction, so that, even when weak, they sway social influence forces in the long run. This simple effect has overarching consequences that have been so far overlooked. Stable cognitive, biological, and environmental constraints are often driving forces of cultural evolution; cross-cultural regularities do not necessitate to be supported by strong constraints.

On the first Thursday of March, the United Kingdom and Ireland celebrate World Book Day, and, in many primary schools, children dress up as their favorite book characters. When children turn seven or eight, parents start to recognize in the courtyard the familiar figures of the Hogwarts School of Witchcraft and Wizardry. Children in the United Kingdom and Ireland go through generational waves, where each cohort "rediscovers" Harry Potter. Why are some things successful in spreading widely and stably, such as the fictional world of Harry Potter in the last twenty years, while others are not?

An intuitive distinction concerns the effect of social influence versus stable, not-socially influenced, preferences that make some features of the content of the traits more appealing. The "rediscovery" of Harry Potter is due to parents, elder siblings, and early-adopter peers, from which children learn, as a minimum, about its existence. At the same time, the content of Harry Potter's stories should be attractive enough to reinforce social influence, in order to be stable through years, and in many different countries. Stable, nonsocial preferences and social influence are likely to act together, to a different degree. For some traits, however, individual preferences for the content seem more important: Western children, on average, prefer pizza to boiled spinach; the great majority of cultures use, on some occasions, masks, or make-up for faces (1). For others, it may be the opposite: Hugging or kissing can be used as greetings in some societies, but considered inappropriate in others where handshakes, or bowing, are used; beanie hats and skinny jeans come and go.

This intuitive distinction reflects important practical differences. Social factors can be leveraged to promote behavioral change (whether for the bad or for the good), while stable preferences for content features tend to be more difficult to overcome, as parents or educators trying to have children eating spinach instead of pizza know well. Social factors intuitively should result in more cross-cultural diversity, where relatively unconnected subpopulations converge on different cultural configurations (see, e.g., the experimental work in ref. 2) while more stable nonsocial preferences should attract the same subpopulations toward similar outcomes.

The same distinction is used in evolutionary approaches to the study of culture. Epistemic vigilance distinguishes, for example, between the evaluation of the "source" and of the "content" of communicated information (3). In the cultural evolution framework, different mechanisms have been proposed as reflecting social influence, usually under the general label of indirect-biased transmission (4) or context-biased social learning strategies (5): These mechanisms act by selecting among different cultural traits not drawing upon their content, but on features of the individuals holding those traits. On the other side, direct-biased transmission, or content-biased social learning, refers to the selection of traits based on their intrinsic features. Preferences for content also act outside cultural selection: We can adopt traits via individual learning, sometimes labeled as guided variation in cultural evolution framework (4). More generally, various individual processes can make

Author affiliations: [a]Department of Sociology and Social Research, University of Trento, 38122 Trento, Italy; and [b]Max Planck Institu te for the Study of Crime, Security and Law, 79100 Freiburg, Germany

Author contributions: A.A. designed research; A.A. and B.d.C. performed research; A.A. and B.d.C. analyzed data; and A.A. and B.d.C. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: alberto.acerbi@unitn.it.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2412380122/-/DCSupplemental.

Published February 21, 2025.

us converge with higher probability to some traits, or particular configurations of traits: In cultural evolution terminology, they are referred to as convergent transformation (6, 7), or content-based attraction (8).

In what follows, we present analytical models of cultural evolution that consider both stable nonsocial preferences and social influence, both with binary and continuous traits. While in the cultural evolution framework content biases include also culturally variable, possibly short-lived, preferences (e.g., a preference for a tool can change when a more effective one is found), here we have in mind a restricted version, that refers to stable and possibly universal preferences. These may be preferences for features that appeal to evolved cognitive predispositions, but also broader biological preferences or physical and environmental constraints. In detail, we are using the label "individual preferences" for mechanisms that are *not socially influenced*, that is, they do not depend directly on features of the larger population; *directional*, that is, point toward a particular trait, or traits configuration; and, finally, *stable*: They do not change, at least at the time scale of the models.

It is important to point out that, because of the above definition, we explicitly model our individual preferences not as a selection process (a process that is determined also by aspects of the structure of the population, i.e., the frequency of the trait), but as a purely individual process. The individual preferences modeled here are equivalent, in the binary case, to biased mutation as in ref. 9; in the continuous case, to guided variation (4), biased or convergent transformation (6, 7), or content-based attraction (8).

In opposition, in our models, "social influence" identifies mechanisms that select traits to copy based on features of the population/source, that are not associated with traits, and thus change accordingly to changes in the population. The models consider two of the most studied mechanisms of social influence: conformity (or frequency-based indirect bias) and demonstrator-based indirect bias. Conformity is defined as a disproportionate tendency to copy from the majority, and it is implemented, in the binary case, following ref. 9 for the individual-based version, and ref. 4 for the analytical model. Conformity with continuous traits is rarely modeled: (but see refs. 10 and 11). Demonstrator-based indirect biases instead do not depend on the frequency of traits, but on features of the demonstrator. A classic example is prestige bias, or a tendency to copy preferentially from individuals that are considered to be "high-status" (12), but any tendency that makes individuals preferentially choose some demonstrators because of features independent from the copied traits would fit the description, such as copying preferentially younger (or older) demonstrators. The individual-based models implementing demonstrator-based indirect biases are inspired, with modification, by refs. 9 and 13, while for the analytical treatment we follow refs. 14 and 15.

Surprisingly, there is not a univocal understanding of what are exactly the consequences of dynamics driven by stable, nonsocial, preferences for content or by social influence. Many studies in cultural evolution focus on social influence, possibly because heuristics like conformity or prestige bias can produce population-level adaptive behaviors that go beyond individual cognition, a process that is considered central in cultural evolution (16). However, this focus has sometimes overshadowed the role of stable nonsocial forces, and other researchers have instead highlighted the importance of weak but stable preferences for content as the main way to support cultural transmission and hence stabilize traditions (17). Some other models have considered the joint effect of content biases, intended as a selection processes, and social forces, focusing for example on the evolution of altruism (18). Others examined the relationship between individual learning processes and social forces, showing for example that the combination of both results in a better adaptive outcome than using separately either of them (19) or that dynamics of diffusion of innovation suggests social forces have a predominant role (20).

Models inspired by cultural attraction theory, focusing on the effect of stable, nonsocial, preferences, are closer to our intent. Ref. 21 showed that when both social influence and preferences for content act, social influence fully determines the outcome. However, their model assumes two preferences for content, and social influence acting stably toward one of the two. In response to this work, ref. 22 presented a model where social influence and the target for content preferences are separated, and they show how the final equilibrium point depends on the relative strength of social influence and content preferences (see also ref. 8).

In the models above, the target of social influence is however fixed and linked to a particular trait, or trait configuration. In our model, instead, we consider "pure" social influence, as determined *only* by the context, be it the frequency of traits in the population, or some demonstrators' features, independent from the copied trait. This makes clear the specific role of social influence. In reality, we expect that social influence would preferentially target certain traits and not others. As we examine in the *Discussion*; however, we believe that this reinforces our point, underscoring the importance of trait features for cultural dynamics.

Models and Results

Frequency-Based Social Influence. The analytical model assumes a large population of individuals. In the binary trait model, individuals possess a cultural trait A or B, and $-p$ denotes the frequency of A. For consistency across our different models, we use the variable p to quantify the trait. In the binary trait, trait A is equivalent to $p = 1$, and trait B to $p = 0$. Traits A and B can be thought as two mutually exclusive traits that are already present in the population and can be chosen by the individuals, like in the pizza and boiled spinach example, but also traits that represent choices that can easily be taken by single individuals without the necessity of inventing them anew. Examples include choosing to paint faces with direct versus averted eye gaze, as in ref. 23, with a cognitive individual preference pointing to direct eye gaze; using tonal or atonal structure in music, as in ref. 24, with a cognitive individual preference for tonal music; or transmitting stories' details with negative versus positive information, as in ref. 25, with a cognitive individual preference for negative information. When examining frequency-based social influence (e.g., conformity), following ref. 4, at each time step, three demonstrators are randomly chosen for each individual. If all have the same trait (three As or three Bs), the individual copies it automatically. In the other cases, the majority trait (i.e., the one possessed by two demonstrators) is adopted with a probability equal to *2/3 + D/3*. The parameter D goes between 0 and 1, regulating the strength of conformity. With $D = 0$ (no conformity), the probability of copying the majority trait is *2/3*, equivalent to unbiased copying, and with $D = 1$ (maximum conformity) individuals always copy the majority trait. At each time step, a fraction $D\overline{p}(1-\overline{p})(2\overline{p}-1)$ of the population, where $\overline{p}$ is the frequency of trait A, switches to trait A (4). In other words, the effect of conformity is proportional to the strength of conformity D, to the variance of trait in the population $\overline{p}(1-\overline{p})$, and to the margin of majority of the dominant trait $(2\overline{p}-1)$. In the models, D is fixed to 1. In addition, individuals have a nonsocial preference for a trait: At each time step, with a probability a, individuals switch to trait A.

In our analytical model, we assume that the population is infinite, so that the dynamics are deterministic. For simplicity, we assume that the two effects act simultaneously and independently. The results align with our individual-based simulations where populations are finite and where social influence and individual preferences act sequentially, which shows that the model is robust to relaxing our simplifying assumptions (*SI Appendix*). The dynamics of the system can be summarized by the same equation as the one governing conformist transmission in ref. 4, adding a term representing the individual preference (to maintain realism, we impose that $0 \leq \overline{p}_t \leq 1$, which has no consequence at equilibrium):

$$\overline{p}_{t+1} = \overline{p}_t + D\overline{p}_t\left(1-\overline{p}_t\right)\left(2\overline{p}_t - 1\right) + \alpha\left(1-\overline{p}_t\right),$$

When only conformity acts ($\alpha = 0$), the trait that is initially in majority will fixate. Fig. 1 shows the dynamics for all the range of $\overline{p}$ and for low values of α. When $\alpha > 1/8$, A always fixates. When $\alpha < 1/8$, we have two stable equilibria: Either the fixation of A or a strong majority of B (>¾). The one reached depends on the initial value of $\overline{p}$, above or below the "unstable equilibria" line in Fig. 1, which separates the two basins of attraction. Like in ref. 4, individuals only observe three demonstrators. With more demonstrators, equations have degree higher than two and the model thus becomes much more complex, or even intractable, but the value of α needed to guarantee that A fixates would be higher, keeping constant the strength of conformity. (A python notebook to reproduce Figs. 1 and 2 below is available at https://github.com/albertoacerbi/attraction_social.)

In the continuous trait models, individuals have a continuous trait p, and they are subject to three forces. First, social influence: Each individual samples three individuals and updates its cultural trait in the direction of the mean trait of the three demonstrators, by keeping a fraction (1 − β) of its current trait value and adding a fraction β of the mean trait value of three randomly sampled demonstrators, whose traits are denoted p_1, p_2, and p_3. Second, individuals have a nonsocial preference for the trait value 1: At each time step, they reduce their distance to trait value 1 by a fraction α. Third, a random value ϵ is added, drawn from a normal distribution with mean 0 and variance σ_e^2, representing copying error or stochasticity.

In summary, we have, for each individual, at each time step:

$$p_{t+1} = (1-\beta)p_t + \beta\frac{p_1 + p_2 + p_3}{3} + \left(1-p_t\right)\alpha + \epsilon,$$

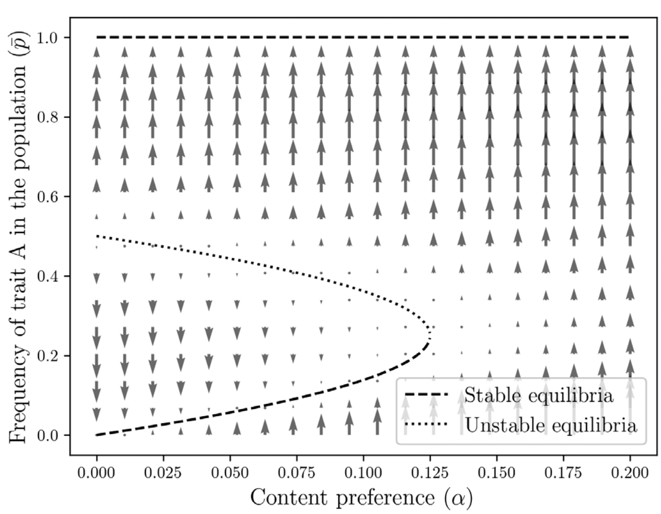

Fig. 1. Vector field for the frequency-based social influence model with discrete traits. The trait A always fixates, unless $\overline{p}$ is initially below the unstable equilibrium line and α < 1/8 (0.125). For $\alpha = 1/8$, we observe a saddle-node bifurcation: The stable and unstable equilibria collide and disappear. Here, as elsewhere, we use $D = 1$ (maximum conformity strength).

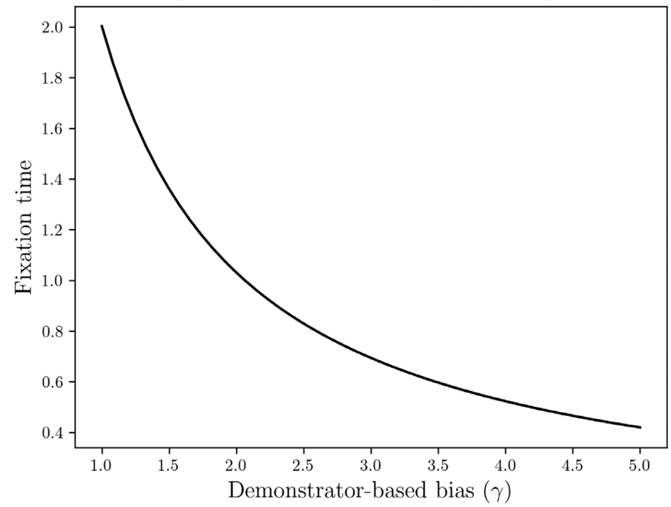

Fig. 2. Time to fixation of the favored trait as a function of the strength of demonstrator-based social influence. The stronger the demonstrator-based bias (γ), the faster the fixation of the trait favored by the individual preference (here α = 0.1 and $\overline{p}$ is initially 0.5).

Again, we assume an infinite population for simplicity, and our results align with our finite population simulations. In this case, the mean $\overline{p}$ of the trait in the population is in expectation unaffected by the copying process, so that, with $\alpha > 0$, the only equilibrium is $\overline{p} = 1$, i.e., given some individual preference for the trait, the population converges on it. At equilibrium, the population is normally distributed around the trait value 1, with a variance equal to $\frac{\sigma_e^2}{2\alpha + 2\beta}$ (*SI Appendix*).

Demonstrator-Based Social Influence. As above, in the binary trait model, we consider a population where individuals possess a cultural trait A or B, and $\overline{p}$ denotes the frequency of A. Now, individuals choose only one demonstrator when they update their trait. We assume that a fraction C_s of the population (e.g., prestigious individuals) has a probability C_{copy} times higher than the rest of the population to be selected as demonstrators. As before, agents have a preference for trait A: With probability α, they switch to the trait A.

This setting is more subject to randomness: with $\alpha = 0$, the dynamics resemble a random walk. For this reason, we do not assume an infinite population here, and build instead a stochastic model. When the population is large, the system can be approximated by a Wright–Fisher diffusion (15, 25). There, both the time t and the proportion $\overline{p}$ of individuals possessing the trait A become continuous variables, which allows us to make use of differential equations. Then, we can study the long-term behavior of the system, and the time to reach equilibrium (*SI Appendix*).

The system is governed by the stochastic differential equation:

$$d\overline{p}_t = \alpha\left(1-\overline{p}_t\right)dt + \gamma\sqrt{\overline{p}_t\left(1-\overline{p}_t\right)}\,dB_t,$$

where B_t is the standard Brownian motion and γ measures the strength of the demonstrator-based bias, encompassing both C_s and C_{copy} (*SI Appendix*). In other words, the system is subject to two forces, the first one being directional, proportional to the individual preference α, and the other one being nondirectional, analogous to genetic drift, proportional to the demonstrator-based bias strength γ.

As shown in ref. 26, in the long run, the trait A always fixates: $\overline{p} = 1$ is the only absorbing state of the equation. A possibly more surprising result is that the stronger the demonstrator-based bias (either because individuals that are preferentially copied are rarer, or because they are more influential), the quicker the fixation (Fig. 2). In our model, individuals that are preferentially copied do not have, in expectation, a different trait than the rest of the population. As visible in the above equation, the demonstrator-based bias has therefore no directional effect, but only increases the volatility by shrinking the pool of demonstrators, which makes it easier for the favored trait to fixate.

If the trait is instead continuous, we can apply the same reasoning as in the frequency-based social influence case with $\beta = 0$ (*SI Appendix*). As long as $\alpha > 0$, individuals cluster around the trait 1, with a variance proportional to the magnitude of copying error and inversely proportional to α.

Discussion

The results of the models show that, in most scenarios we studied, stable, individual preferences stabilize culture at the point where they are directed to. Besides the case of conformity acting on binary choice (see more below), weak individual preferences (any α higher than 0) suffice. Social influence mechanisms—both based on frequency (conformity) and on demonstrators' features (e.g., prestige bias)—are, in our models, independent of traits' features and nondirectional. When they act together with individual preferences, the latter are therefore the only directional forces, and their existence is sufficient to determine the cultural evolutionary outcome of the system.

This message suffers one exception. When most of the population holds the nonpreferred trait in a binary choice, and conformity acts, the preference for the (minority) traits needs to be sufficiently strong to overcome the majority. In fact, conformity is nondirectional with respect to the traits, but reinforces existing majorities. In cultural evolution, it has long been recognized that a conformist bias can make cultural traits persistent and maintain between-group cultural variation (refs. 27 and 28, but see ref. 29). If individual preferences for trait A build gradually, our model suggests that cultural change could happen suddenly: Trait A would stay rare for a while, then suddenly spread as α crosses 1/8 (Fig. 1). This value of 1/8 applies in the setting we study, with three demonstrators and maximum conformity; it would be higher with more demonstrators, and smaller with weaker conformity. When conformity and individual preferences act together, a subtle change in preferences can be enough to trigger a sudden cultural shift. Our model thus provides a possible parsimonious explanation for "tipping points" in cultural evolution (30, 31).

In the case of demonstrator-based bias (e.g., prestige), interestingly, increasing social influence not only does not hamper the favored trait fixation but even accelerates the process (Fig. 2). This counterintuitive has also recently been found independently (32) and interpreted there as a "rich-get-richer" process. We prove mathematically (see *SI Appendix*, p. 8) that the mechanism at play is that a stronger prestige bias accelerates the fixation simply by reducing the effective population size and thereby adding more stochasticity.

As we mentioned above, conformity with continuous traits has been rarely modeled. In our implementation conformity with continuous traits produces large differences in results with respect to conformity with binary traits. The latter, usually considered in cultural evolution theory, is a process, driving population toward the majority trait, while the former only reduces variation in the population. Future works may explore different implementations of conformity with continuous traits or with multiple ordinal discrete traits (see, e.g., ref. 33).

It is important to notice that our results depend on the fact that the modeled social influence mechanisms are fully detached from the content of the traits. In reality, we expect that, for example, prestigious individuals would possess, on average, more adaptive traits than individuals chosen at random, or that the majority targeted by conformist copying would effectively pool information from individual learning (4). However, we believe that this reinforces our point, as it underscores again how social influence needs to be guided by individual decisions or preferences for features of traits to be effective.

Also, in the models presented here, the individual preference is uniform in the population, i.e., only one preference was considered, but the same logic can be applied to more realistic situations with various individual preferences acting in different directions, where we would expect them to homogenize population toward the stronger nonsocial preferences.

As mentioned in the Introduction, few works have explicitly addressed these questions, but the results presented here are consistent with suggestions coming from other cultural evolution models and give them a more general background. Acerbi et al. (6), for example, found that convergent transformation drives cultural dynamics when acting together with unbiased copying (and, similarly to here, the more faithful the copying is, the stronger the effect of convergent transformation). Morgan et al. (10) found that even weak priors render conformity unable to stabilize traditions and determine the outcome, in most conditions for binary choices, and always for continuous choices.

The immediate take-home message of these results is that, if our question is why some things are culturally successful and others are not, stable nonsocial forces, even when weak, need to be taken into account. This holds also in case of social norms, where there are benefits to coordinating with other group members, or sanctions in deviating from other group members: Models including these dynamics have recently shown that when stable, directional, factors are present, different populations converge on the same norm, at least when norms are continuous (34), echoing our results in the frequency-based social influence model.

A less obvious take-home message concerns the interpretation of cross-cultural regularities. The existence of human universals (35) is sometimes interpreted as supporting the existence of strong cognitive evolved dispositions, or strong ecological constraints, and indicating a somehow limited role of culture. On the other end of the spectrum, sociocultural anthropologists have tended to diminish the importance of cross-cultural regularities to stress the importance of culture. While everyone would agree this is a false dichotomy (see, e.g., ref. 36), these results suggest a way to understand why it is so: Weak directional, nonsocial forces, as long as they are stable enough, can produce strong regularities. These can be (possibly weak) cognitive priors, physical affordances, relatively stable ecological conditions (such as the availability of certain materials), and so on. Conceptually, it is important to think of social influence and individual preferences not as opposing forces. Nondirectional social influence provides strength to the weak but directional preferences for certain contents. In other words, culture magnifies individual-level tendencies, allowing them to become stable at population level. This can be clearly seen, in our model, in the case of the demonstrator-based bias, where the stronger is the social influence, the faster is the convergence toward the equilibrium to which the individual preferences point.

Data, Materials, and Software Availability. R and python code data have been deposited in attraction_social (GitHub) (https://github.com/albertoacerbi/attraction_social) (37).

ACKNOWLEDGMENTS. We would like to thank Thomas Morgan, Karthik Panchanathan, and the anonymous reviewers for feedback on an earlier version of the draft. James Winters suggested implementing copying probabilities in the demonstrator-biased simulation with a softmax function. Marius Mercier helped spot typos in the preprint. José Moran and Emilien Jemelen helped with the analytical model.

1. D. Sperber, L. A. Hirschfeld, The cognitive foundations of cultural stability and diversity. *Trends Cogn. Sci.* **8**, 40–46 (2004).
2. M. J. Salganik, P. S. Dodds, D. J. Watts, Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**, 854–856 (2006).
3. D. Sperber *et al.*, Epistemic vigilance. *Mind Lang.* **25**, 359–393 (2010).
4. R. Boyd, P. J. Richerson, *Culture and the Evolutionary Process* (University of Chicago Press, 1985).
5. R. L. Kendal *et al.*, Social learning strategies: Bridge-building between fields. *Trends Cogn. Sci.* **22**, 651–665 (2018).
6. A. Acerbi, M. Charbonneau, H. Miton, T. Scott-Phillips, Culture without copying or selection. *Evol. Hum. Sci.* **3**, e50 (2021).
7. A. Mesoudi, Cultural selection and biased transformation: Two dynamics of cultural evolution. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **376**, 20200053 (2021).
8. N. Claidière *et al.*, Convergent transformation and selection in cultural evolution. *Evol. Hum. Behav.* **39**, 191–202 (2018).
9. A. Acerbi, A. Mesoudi, M. Smolla, *Individual-Based Models of Cultural Evolution: A Step-by-Step Guide Using R* (Routledge, 2022).
10. T. J. H. Morgan, B. Thompson, Biased transformation erases traditions sustained by conformist transmission. *Biol. Lett.* **16**, 20200660 (2020).
11. P. E. Smaldino, J. M. Epstein, Social conformity despite individual preferences for distinctiveness. *R. Soc. Open Sci.* **2**, 140437 (2015).
12. Á. V. Jiménez, A. Mesoudi, Prestige-biased social learning: Current evidence and outstanding questions. *Palgrave Commun.* **5**, 1–12 (2019).
13. A. Mesoudi, The cultural dynamics of copycat suicide. *PLoS One* **4**, e7252 (2009).
14. J. Moran, A. Fosset, M. Benzaquen, J.-P. Bouchaud, Schrödinger's ants: A continuous description of Kirman's recruitment model. *J. Phys. Complexity* **1**, 035002 (2020).
15. A. Etheridge, *Some Mathematical Models from Population Genetics: École d'Été de Probabilités de Saint-Flour XXXIX-2009* (Springer, 2011).
16. J. Henrich, *The Secret of Our Success* (Princeton University Press, 2016).
17. O. Morin, *How Traditions Live and Die* (Oxford University Press, 2015).
18. K. K. Denton, Y. Ram, M. W. Feldman, Conformity and content-biased cultural transmission in the evolution of altruism. *Theor. Popul. Biol.* **143**, 52–61 (2022).
19. Z. Hong, Combining conformist and payoff bias in cultural evolution: An integrated model for human decision-making. *Hum. Nat.* **33**, 463–484 (2022).
20. J. Henrich, Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *Am. Anthropol.* **103**, 992–1013 (2001).
21. J. Henrich, R. Boyd, On modeling cognition and culture: Why cultural evolution does not require replication of representations. *J. Cogn. Cult.* **2**, 87–112 (2002).
22. N. Claidière, D. Sperber, The role of attraction in cultural evolution. *J. Cognit. Cult.* **7**, 89–111 (2007).
23. O. Morin, How portraits turned their eyes upon us: Visual preferences and demographic change in cultural evolution. *Evol. Hum. Behav.* **34**, 222–229 (2013).
24. S. A. Mehr *et al.*, Universality and diversity in human song. *Science* **366**, eaax0868 (2019).
25. K. Bebbington, C. MacLeod, T. M. Ellison, N. Fay, The sky is falling: Evidence of a negativity bias in the social transmission of information. *Evolut. Hum. Behav.* **38**, 92–101 (2017).
26. W. Feller, "Diffusion processes in genetics" in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (University of California Press, 1951), pp. 227–247.
27. A. Mesoudi, Migration, acculturation, and the maintenance of between-group cultural variation. *PLoS One* **13**, e0205573 (2018).
28. J. Henrich, R. Boyd, The evolution of conformist transmission and the emergence of between-group differences. *Evol. Hum. Behav.* **19**, 215–241 (1998).
29. K. K. Denton, Y. Ram, U. Liberman, M. W. Feldman, Cultural evolution of conformity and anticonformity. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 13603–13614 (2020).
30. D. Centola, J. Becker, D. Brackbill, A. Baronchelli, Experimental evidence for tipping points in social convention. *Science* **360**, 1116–1119 (2018).
31. O. Kolodny, N. Creanza, M. W. Feldman, Evolution in leaps: The punctuated accumulation and loss of cultural innovations. *Proc. Natl. Acad. Sci. U.S.A.* **112**, E6762–E6769 (2015).
32. S. Egozi, Y. Ram, Prestige bias in cultural evolutionary dynamics. *R. Soc. Open Sci.* **11**, 230650 (2024).
33. E. H. Mora, K. K. Denton, M. E. Palmer, M. W. Feldman, Conformity to continuous and discrete ordinal traits. *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2417078122 (2025).
34. M. Yan, S. Mathew, R. Boyd, "Doing what others do" does not stabilize continuous norms. *PNAS Nexus* **2**, pgad054 (2023).
35. D. E. Brown, *Human Universals* (McGraw Hill, 1991).
36. H. C. Barrett, *The Shape of Thought* (Oxford University Press, 2015).
37. A. Acerbi, B. de Courson, attraction_social. Github. https://github.com/albertoacerbi/attraction_social. Deposited 24 January 2025.