The logic of fashion cycles

Alberto Acerbi[1,*], Stefano Ghirlanda[1,2], Magnus Enquist[1,3]
1 Centre for the Study of Cultural Evolution, University of Stockholm, Sweden
2 Honors Academy and Department of Psychology, Brooklyn College, New York, US
3 Department of Zoology, University of Stockholm, Sweden
∗ E-mail: alberto.acerbi@gmail.com

Abstract

Many cultural traits exhibit volatile dynamics, commonly dubbed fashions or fads. Here we show that realistic fashion-like dynamics emerge spontaneously if individuals can copy others' preferences for cultural traits as well as traits themselves. We demonstrate this dynamics in simple mathematical models of the diffusion, and subsequent abandonment, of a single cultural trait which individuals may or may not prefer. We then simulate the coevolution between many cultural traits and the associated preferences, reproducing power-law frequency distributions of cultural traits (most traits are adopted by few individuals for a short time, and very few by many for a long time), as well as correlations between the rate of increase and the rate of decrease of traits (traits that increase rapidly in popularity are also abandoned quickly and vice versa). We also establish that alternative theories, that fashions result from individuals signaling their social status, or from individuals randomly copying each other, do not satisfactorily reproduce these empirical observations.

Introduction

While some cultural traits, once introduced in a population, tend to become to varying degrees stable part of the cultural repertoire of that population, others exhibit peculiar volatile dynamics, commonly dubbed fads or fashions. Well documented examples are as diverse as skirt lengths [1], pop songs [2], first names [3–5], dog breeds [6], pottery decorations in the archaeological record [7], and keywords in academics vocabulary [8]. Such fluctuations in popularity are not mainly due to intrinsic characteristics of the traits—there seems to be nothing intrinsically advantageous about, say, wearing purple one year but not the next—but they are likely to reflect forces that are internal to cultural dynamics.

Since at least the 18^{th} century, fashions have been considered a product of social stratification [9–13]. According to such "status" models, a fashion arises because individuals of low social status copy those of perceived high status. When a trait becomes popular, however, high-status individuals quickly abandon the trait to differentiate themselves from low-status individuals. As a consequence, low-status individuals abandon the traits too, bringing the fashion cycle to an end (for recent computational models see [14–16]).

Recently, however, a "neutral" model of cultural change [4] has been proposed as a more parsimonious explanation of fashions. In such a model fashions arise as by-products of individuals copying each other randomly (akin to changes in allele frequency in neutral genetic evolution [17]). The main appeal of this model is its simplicity, and the fact that it reproduces realistic turnover rates as well as empirical frequency distributions of cultural traits in several domains. These distributions closely approximate power-law or log-normal curves [2, 4, 8], meaning that only very few cultural traits become very common, while the vast majority remains rare (Figure 1A).

While the neutral model provides a powerful starting point to study cultural change [4, 20], there are many reasons to believe that fashions are not driven by random choices. Individuals express strong positive and negative preferences for cultural traits, and prefer to copy some models rather than others [21–23]. Status models recognize these factors, at the cost of postulating that high-status individuals are anti-conformist while everyone else is conformist. Status models also postulate a given subdivision between high- and low-status individuals. A model that could explain why some individuals have higher status would yield greater insight into fashion dynamics.

We present here an alternative model of fashion and fads that builds on our previous work on the effects of repeated cultural transmission on the frequency of cultural traits [24–26]. We study the possibility that fashions and fads arise because individuals can copy each other's preferences for cultural traits, in addition to the traits themselves. We call this the "preference model" of fashion. Because preferences determine which traits appear most attractive to individuals, they act as "regulatory" traits in the sense that are both socially learned and influence the outcome of social learning. This highlights an important difference between cultural and genetic evolution: in the former the rules of transmission (e.g., whether to copy or not) may be modified by the cultural process itself [24–27].

We first consider the coevolution of one cultural trait and the preference for such trait, showing that the the repeated cultural transmission of trait and preference is enough to generate a fashion cycle in which the trait first

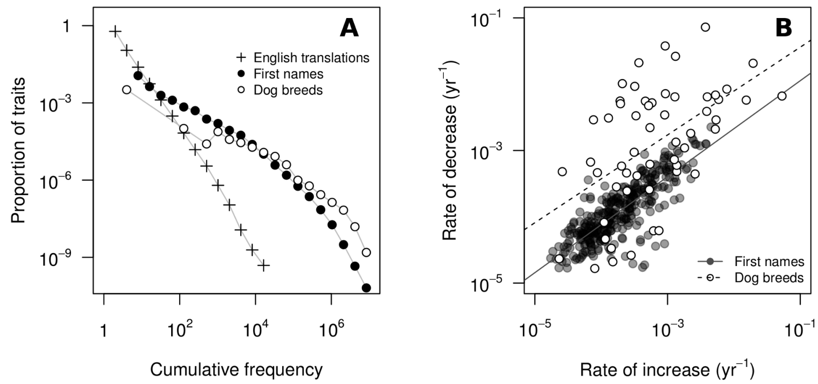

Figure 1. Empirical characteristics of fashion cycles. (A) Frequency distributions of cultural traits often follow power-law or log-normal distributions, i.e., the vast majority of traits remains very rare, while a small minority become very popular. Crosses: Number of times a foreign author has been translated into English (with permission from UNESCO's *Index Translationum*) [18]. Closed circles: number of times a first name has been given to a newborn in the U.S.A., 1880–2006 [19]. Open circles: number of dogs of 154 breeds registered with the American Kennel Club, 1926–2005 (courtesy of H. Herzog) [6]. (B) Correlation between the rates of increase and decrease in the popularity of U.S.A. first names (Pearson's $r = 0.82$, $p < 10^{-10}$, $N = 339$, only names reaching a frequency of at least 0.1% are included, see [5, 19]) and dog breeds (Pearson's $r = 0.53$, $p < 10^{-4}$, $N = 55$).

becomes popular and then disappears from the population. We then consider the simultaneous coevolution of many trait-preference pairs, showing that cultural transmission alone produces fashion dynamics exhibiting two key properties of actual fashion cycles: the power-law distribution of frequency of cultural variants mentioned above, and the finding that cultural traits that increase rapidly in popularity are also abandoned quickly, while slow increase in popularity correlates with slow decrease. This has been shown for first names in the U.S. and France [5], and here we report that it also holds for the popularity of dog breeds, based on an analysis of data on dogs of 154 breeds registered with the American Kennel Club, 1926–2005 (Figure 1B). Lastly, we show that the neither the neutral model of cultural change, nor models based on status can account for the same breadth of empirical data.

Methods

Model 1: One cultural trait

Our core idea is that preferences for cultural traits, besides influencing the adoption and abandonment of traits, are themselves cultural traits that can be adopted or abandoned through social learning. We start exploring this idea modeling the coevolution of one cultural trait and the preference for such trait. Individuals can be of one of four types: 0, lacking both trait and preference; T, possessing the trait only; P, possessing the preference only; PT, possessing both. Individuals meet randomly in pairwise social interactions in which one individual (the *observer*) may copy another (the *model*). The trait and the preference may be copied independently of each other. The probability that copying occurs between any two cultural types is u, apart in the following cases:

- When the model has the trait, an observer with the preference is more likely to copy the model than an observer without the preference. We thus assume that copying between P observers and T or PT models, and between PT observers and T models, occurs with increased probability $w > u$.
- Copying between T observers and 0 models occurs with increased probability $v > u$, meaning that individuals with the trait, but without the preference, have a higher probability of abandoning the trait upon meeting individuals with neither trait nor preference.

Note that copying may result in the observer abandoning the trait or preference, if these are absent in the model.

Our aim is to track the frequency of the four cultural types 0, P, T, and PT over time to understand the cultural dynamics generated by our assumptions. Taking into account all possible transitions between types (e.g., a 0 observer who meets a T model becomes T with probability u), we arrive at the following equations (see Supporting Information for details):

$$\dot{x}_0 = (v-u)x_0x_T - 2u(1-u)x_0x_{PT} + [w(1-w) + u(1-u)]x_Px_T \tag{1}$$

$$\dot{x}_P = -(w-u)x_Px_{PT} - \left[1-u^2-(1-w)^2\right]x_Tx_P + 2u(1-u)x_0x_{PT} \tag{2}$$

$$\dot{x}_T = (w-u)x_Tx_{PT} - (v-u)x_Tx_0 + 2u(1-u)x_0x_{PT} - \left[1-w^2-(1-u)^2\right]x_Tx_P \tag{3}$$

$$\dot{x}_{PT} = (w-u)x_{PT}(x_P - x_T) - 2u(1-u)x_0x_{PT} + [w(1-w) + u(1-u)]x_Px_T \tag{4}$$

where x_i is the frequency of type i, and $\dot{x}_i$ its rate of change (derivative). We now study the cultural dynamics emerging from these equations and relate it to the empirical findings considered above. We consider a generic initial condition in which all types have nonzero frequency. We note from the outset that the frequency of the preference can never increase. Indeed, writing such frequency as $g = x_P + x_{PT}$ and summing equations (2) and (4), we obtain

$$\dot{g} = -(w-u)gx_T \tag{5}$$

which is always negative so long as $x_T > 0$ and $g > 0$. Thus this model cannot explain how a population comes to prefer a given cultural trait. We address how this can happen later in the paper. Here we study cultural dynamics assuming that the preference has reached a frequency $g(0) > 0$ (see also Supporting Information).

Equations (1–4) imply that the trait eventually disappears, while the preference may persist at a very low frequency (see Supporting Information). Fashion cycles are possible, however, because trait frequency can increase for some time before starting to drop. Let $f = x_T + x_{PT}$ be the frequency of the trait in the population. Its dynamics, obtained summing equations (3) and (4), is

$$\dot{f} = (w-u)x_P f - (v-u)x_0 x_T \tag{6}$$

Hence trait frequency increases as long as

$$(w-u)\, x_P f > (v-u)\, x_0 x_T \tag{7}$$

This condition is easily understood noting that, on the l.h.s., $x_P f$ is the rate at which individuals with the preference, and lacking the trait, meet individuals with the trait, while, on the r.h.s., $x_T x_0$ is the rate at which individuals with the trait, and lacking the preference, meet individuals without trait nor preference. In the first kind of encounters trait frequency increases with probability $w-u$, while in the second kind of encounters trait frequency decreases with probability $v-u$. We show in Supporting Information that a fashion cycle occurs provided the initial frequency $g(0)$ is higher than a threshold value (graphed in Figure 2A), which is a function of the combination of system parameters given by

$$\phi = \frac{w-u}{v-u} \tag{8}$$

If the initial preference is lower than the threshold, trait frequency steadily decreases without showing the rise-and-fall pattern characteristic of fashion

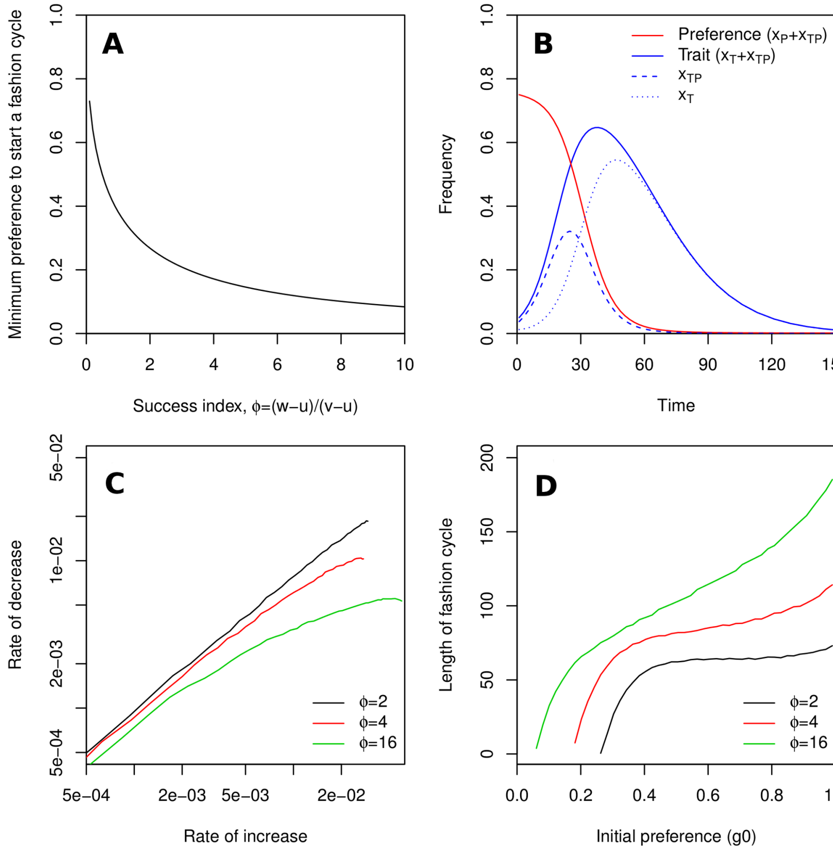

Figure 2. Characteristics of fashion cycles generated by Model 1 (equations 1-4). (A) Minimum initial frequency, $g(0)$ necessary to start a fashion cycle for system parameters u, v, w such that $\frac{w-u}{v-u} = \phi$. The curve has equation $g_{\min}(0) = 1 + \frac{1}{2}\phi(1 - \sqrt{1 - 4/\phi})$ (Supporting Information). (B) Example of fashion cycle starting from initial preference frequency $g(0) = 0.75$ and initial trait frequency $f(0) = 0.05$, with parameters $u = 0.1$, $v = 0.15$, $w = 0.3$. (C) Correlation between rate of increase and decrease of traits. (D) Duration of fashion cycles. The initial frequency of the preference interacts non-linearly with system parameters: only when ϕ is high very long cycles can occur, given a high initial preference. Maximum trait frequency is approximately proportional to the initial preference, and does not depend strongly on ϕ (Figure S3).

cycles. The initial frequency of the trait plays no role in determining the success of the trait: even a trait that is introduced at a high frequency (e.g., through promotional sales or other marketing strategies) will disappear quickly unless it is preferred by sufficiently many individuals.

Figure 2B shows an example of the dynamics generated by this model, assuming that the frequency of the trait is initially low, $f(0) = 0.05$, and that the preference is initially high, $g(0) = 0.75$. As anticipated, the preference steadily decreases, while the trait exhibits a cycle of initial diffusion and eventual abandonment. The cycle results from the superposition of two cycles, first involving the growth and decline of cultural type x_{PT} and then of type x_T. When a cycle occurs, we observe a strong correlation between the rate of increase in trait frequency and the rate of subsequent decrease (Figure 2C), mirroring empirical observations (Figure 1B).

The dynamics of Model 1 (Figure 2B) suggests a theory of the fashion cycle as composed of four phases (Figure 3):

1. The preference for a trait arises in the population (Figure 3A). As already mentioned, this phase is not captured by the model above, and is studied later in the paper.

2. If the preference becomes sufficiently common, the trait itself starts to spread and PT individuals, possessing both the preference and the trait, increase in frequency as the many P individuals meet the few T individuals (Figure 3B, and dashed line in Figure 2B). The reason is that, when the preference is common, T individuals are more effective cultural models than P individuals, i.e., encounters between these two types result in a net flow of P into PT. The magnitude of such flow is regulated by the difference $w - u$ in our model.

3. When PT individuals become common, the situation changes as T individuals gain an advantage in social interactions. Interactions between T and PT, in fact, favors PT individuals becoming T rather than vice-versa, leading to the eventual disappearance of the preference from the population (Figure 3C, and dotted line in Figure 2B).

4. Once the preference has disappeared, the trait tends to disappear, too, because individuals with the trait (both T and PT) are no longer copied, and because T individuals eventually relinquish the trait upon meeting 0 individuals (Figure 3D).

The somewhat surprising shift in the dynamics in phase 3, that explains how a fashion begins to fade, has been studied in previous work exploring

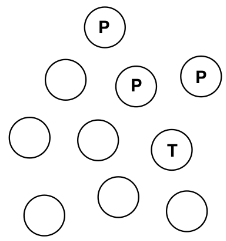

Phase 1. The preference (P) for a trait (T) arises in some individuals. If these are effective cultural models, the preference can spread.

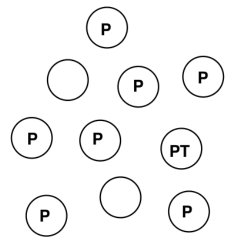

Phase 2. When the preference becomes common, the trait can spread, because PT (or T) individuals are now more effective cultural models than P individuals.

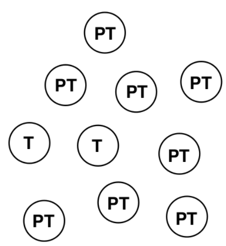

Phase 3. When PT individuals are common, T individuals are more effective cultural models.This favors the disappearance of the preference.

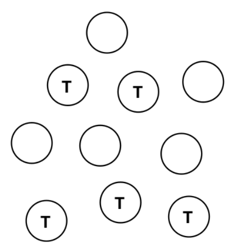

Phase 4. Once the preference has disappeared, the trait tends to disappear, too, because individuals with the trait are no longer effective cultural models.

Figure 3. Schematic illustration of the phases of the fashion cycle of a cultural trait.

what makes an individual an effective cultural model, i.e., someone who significantly influences others [24–26]. The basic idea is that, when a cultural trait is common, individuals with a low preference for the trait are favored as cultural models because they are less likely to copy others, and thus they display to observers a stable set of traits that can be transmitted repeatedly (see also [27, 28]). In line with this logic, T individuals increase in frequency at the expense of PT individuals, when then the latter are common, because the PT type is more likely to copy the T type than vice-versa.

The correlation between rates of increase and decrease of trait frequency arises because the speed at which the successive phases of the fashion cycle unfold is largely determined by the initial frequency of the preference. The

key observation is that transformations between cultural types occur at rates proportional to the abundance of the cultural types themselves. Thus a large initial preference produces quickly relatively many PT individuals, as there are many P observer who can adopt the trait from T and PT models (transition from phase 2 to phase 3 in Figure 3, or dashed blue line in Figure 2B). As P individuals turn quickly into PT, transitions from PT to T increase in frequency and quickly convert PT into T individuals (transition from phase 3 to phase 4 in Figure 3, or dotted blue line in Figure 2B), which eventually leads to the abandonment of the trait. When the initial preference is lower, all these processes occur more slowly because the level reached by the variables x_P, x_{PT}, and x_T is lower.

In conclusion, this simple model shows that the cultural transmission of traits and preferences is sufficient to generate fashion cycles with intuitively appealing features, including the observed correlation between rates of increase and decrease in trait popularity. The model cannot, however, explain how preferences for traits can reach high levels, and cannot explain the observed frequency distribution of cultural traits. To address these issues, we generalize the model to include many cultural traits and preferences.

Model 2: Many cultural traits

To describe individuals who can bear many cultural traits and preferences, we introduce the variables q_i ($i = 1, \dots, n$) that encode whether the individual lacks ($q_i = 0$) or possesses ($q_i = 1$) any of n traits, and by variables p_i that encode the individual's preferences for the traits. Preferences range continuously from -1 (strong dislike) to 1 (strong liking). Our core assumption is that an observer's probability to copy a cultural model is an increasing function of how much the observer prefers the model's traits. We compute the overall preference of observer o for model m as

$$P_{om} = \sum_{i=1}^{n} p_{oi} q_{mi} \tag{9}$$

and we define the probability that cultural transmission occurs from the model to the observer as

$$c_{om} = \frac{1}{1 + e^{-P_{om}}} \tag{10}$$

This expression ranges from near 0, when the model has many traits that the observer dislikes ($P_{om} \ll 0$), to near 1, when the model has many preferred traits ($P_{om} \gg 0$). When the observer is indifferent to the model's traits

($P_{om} = 0$), the probability of cultural transmission is $c_{om} = 0.5$. When an observer copies a model, she chooses a trait-preference randomly, then copies each member of the pair independently of the other with probability c_{om}. Equations (9) and (10) are similar to successful models of human decision making, such as discrete choice models in econometry [29] and neural network models of behavior [30], in which decisions arise from individuals attributing different weights (preferences) to pieces of information from their environment.

Model 2 reduces to Model 1 when there is only one cultural trait and preferences only have two values. For example, choosing -1 and 1 as possible preferences results in $u = 1/(1 + e^1) \simeq 0.27$ for the probability that an observer without the preference copies the model, and $w = 1/(1 + e^{-1}) \simeq 0.73$ for the probability that an observer with the preference copies a model with the trait. In Model 2 we do not need to assume a special rule for interactions between observers with the trait, but lacking the preference, and models with neither trait nor preference (parameter v in Model 1). Rather, when many traits are present, those for which preferences are low are abandoned spontaneously as a result of competition with other traits for the opportunity of being copied.

Neutral model and status model

We simulate a neutral model of cultural evolution that mirrors closely existing ones [4], but allows individuals to interact many times during their life [27] and to accumulate cultural traits [31, 32]. In this model there are no preferences for cultural traits, hence we fix the probability of cultural transmission at $c_{om} = 0.5$. When an observer copies a model, she copies a randomly chosen trait (the observer can also abandon a trait, if the trait randomly chosen is possessed by the observer but not by the model).

We also simulate a model of cultural transmission based on status. Individual are endowed with an additional variable that describes their social status (S). Cultural transmission occurs with a constant probability $c_{om} = 0.5$ as in the neutral model, but the outcome depends on the relative status of observer and model. If $S_m > S_o$, the observer copies from the model a randomly chosen trait. If $S_m \leq S_o$, the observer abandons a randomly chosen trait possessed by both observer and model. The status of newborn individuals is assigned randomly. We also simulated alternative status models, either with binary status (95% of low-status individuals and 5% "high-status") or varying the threshold for abandoning or copying traits, with qualitatively similar results not reported here.

Results

The preference, neutral and status models all exhibit dynamics in which trait frequencies change with time, and even traits that reach high frequencies eventually disappear (Figure 4). The preference and neutral models generate realistic frequency distributions, while the status model appears incapable of generating traits that persist at high frequency for a long time (Figure 1A and Figure 5A,C,E, and Figure S4 for trait lifespan distributions). The reason is that as soon as a trait becomes common, high-status individuals abandon it, which in turns triggers abandonment from low-status individuals. Thus the status model may account for brief fads, but not for "classic" styles such as the four-in-hand tie knot popular since the early 20$^{\mathrm{th}}$ century [33], or English names such as Mary and John [19].

The preference and status models exhibit correlations between the rate of increase and decrease of traits similar to those observed in empirical data (Figure 1B and Figure 5B,D,F). Such correlations do not exist in the neutral model, because the frequency of a trait at time $t+1$ depends solely on the frequency at time t and not on the history of the trait. In the preference model, on the other hand, the frequency of a trait at time $t+1$ depends on the frequency of the trait at time t *and* on the preference value at time t (cf. Figure 2B).

The many-trait preference model (Model 2) addresses the main shortcoming of the model with one trait (Model 1), showing that the preference for a trait can spread if it happens to be associated with effective cultural models (Figure 6A). As these individuals are copied more often than others, the preference spreads in the population. After a preference has been established, fashion cycles in Model 2 follow the same logic as in Model 1. Figure 6B shows that a trait increases in frequency when the associated preference has increased, and that the preference starts falling as the trait becomes common (cf. Figure 2B; see Figure S5 for examples of the preference-frequency dynamics of individual traits). Cross-correlation analysis of trait and preference dynamics shows that preferences anticipate frequencies by, on average, 20% of a trait's lifetime, with an average cross-correlation of 0.60 ($p < 10^{-10}$, $t = 51.51$, $N = 178$, two-tailed one-sample t test).

Discussion

We have shown that the social transmission of preferences for cultural traits is sufficient to generate realistic fashion cycles, resulting in a better fit with

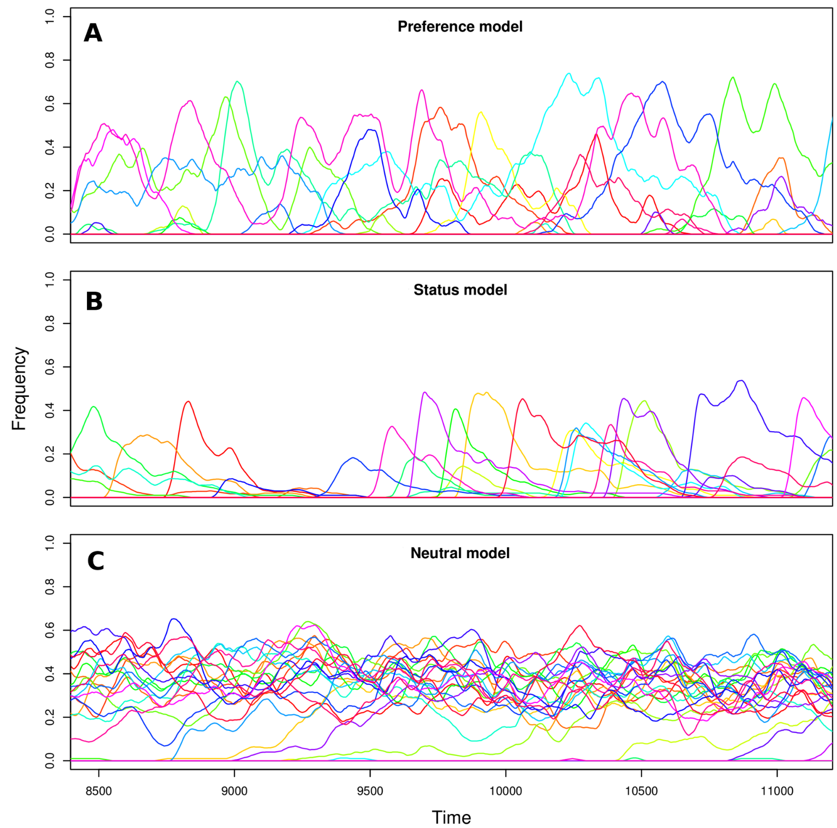

Figure 4. Changes of trait frequency over time in the multi-trait models of fashion. We simulate $N = 100$ individuals interacting randomly in discrete time steps. Cultural traits are continuously introduced with a probability $\mu_q = 0.001$ per individual per time step (a new trait is introduced, on average, every 10 time steps). Individuals may spontaneously change their preferences for existing traits, resetting them to random values between -1 and 1 at a rate of $\mu_p = 0.001$ per individual per time step. Each individual has a probability of dying of $\tau = 0.01$ per time step (average lifetime is 100 time steps). A dying individual is replaced by a naive individual who possesses no cultural traits ($q_i = 0$, $i = 1, \ldots, n$) and is maximally open to learning for others ($p_i = 1$, $i = 1, \ldots, n$). Simulations last 20,000 time steps (only 3,000 are shown in the figure, and a maximum of 40 traits is shown to maintain legibility).

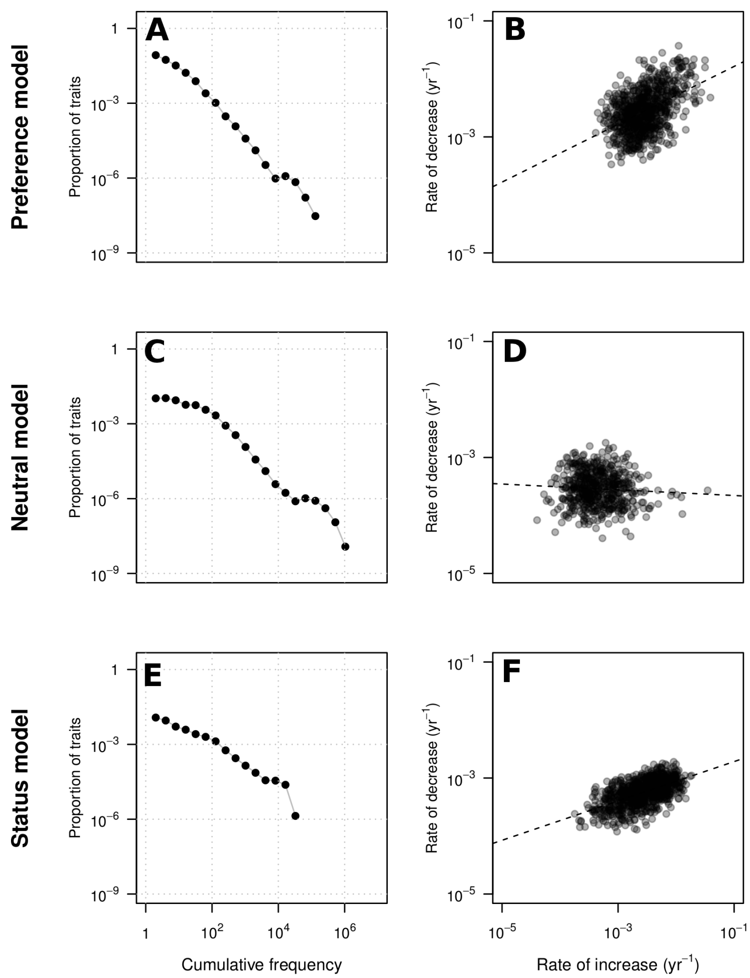

Figure 5. Characteristics of cultural dynamics in the multi-trait models of fashion. (A,C,E) Cumulative frequencies of traits (cf. Figure 1A for empirical data). (B,D,F) Correlations between rates of increase and decrease of traits (cf. Figure 1B for empirical data): preference model (Pearson's $r = 0.50$, $p < 10^{-10}$, $N = 789$) neutral model (Pearson's $r = -0.07$, $p = 0.12$, $N = 558$), and status model (Pearson's $r = 0.58$, $p < 10^{-10}$, $N = 871$). Data from 5 simulations with the same parameters as in Figure 4. Simulated time steps have been converted to years assuming an average lifetime of 70 years.

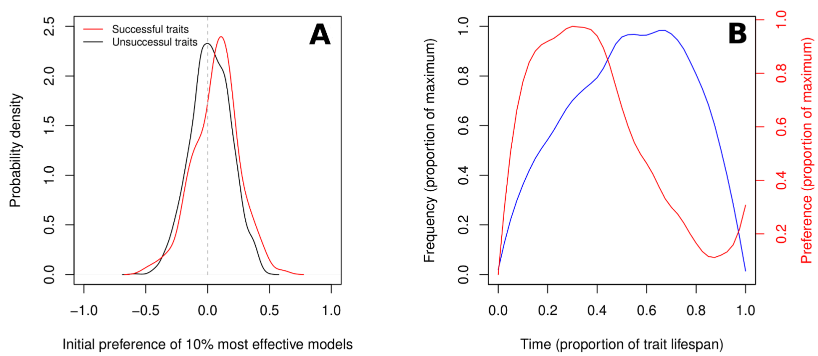

Figure 6. Dynamics of preference-trait coevolution in Model 2. (A) Initial preferences of the 10% most effective cultural models at the start of a trait's lifetime, for traits that will be successful (red, mean=0.08) or unsuccessful (black, mean=0.02, $p < 0.001$, two-tailed Wilcoxon rank-sum test, the graph shows estimated probability densities). Traits are considered successful if they reach a frequency of at least 0.1. For successful traits, initial preference is calculated as the average preference from the time a trait is introduced until the trait reaches 0.1 frequency. For unsuccessful traits, initial preference is calculated as the average over the mean time required for successful traits to reach 0.1 frequency. (B) Black line: Trait frequency as fraction of the maximum frequency. Red line: Average preference as fraction of the maximum preference. The average is computed over all traits reaching at least a frequency of 0.1, and lasting in the population for at least 50 time steps. Time is measured as proportion of the trait's lifetime. Simulation parameters as in Figure 4.

empirical data than the neutral and status models of cultural change (cf. Figures 1 and 5).

Our model also overcomes theoretical shortcomings of previous models. In respect to the neutral model, it recognizes that people have preferences for cultural traits and individuals and that those preferences influence the copying process [21–23]. In respect to the status model, it suggests how social status can emerge from cultural evolution itself. In our model, influential individuals are those who possess many traits that others prefer and, at the same time, have low preferences for widespread traits. This echoes the concept of "high-status" in status models (high-status individuals are copied by others, but do not themselves copy others), without assuming either a pre-existent social stratification or anti-conformist behavior in high-status individuals. In the preference model, moreover, an individual's status can change if traits or preferences in the population change.

Alternative, more elaborated, versions of the status model could produce results different from the ones shown here, and possibly a better fit of the empirical data. For example, it has been suggested that individuals adopt cultural traits to signal their identity and differentiate themselves from individuals of other social groups, without necessarily implying a "vertical" (high/low) status hierarchy [34, 35]. Our goal here is certainly not to deny the importance of status or identity signaling, but to show, as explained above, that at least certain features of status-driven cultural dynamics can be reproduced within a more parsimonious set-up.

Finally, it is worth to note that the quantity ϕ in equation (8) can be considered a "success index" akin to genetic fitness, in that it allows to predict whether a trait will spread. While it is impossible to define a generally valid success index for cultural evolution [27], ratios similar to ϕ may predict the success of cultural traits when individuals have many opportunities for social learning [24, 25, 27]. Numerators in such ratios measure the ease with which a trait is transmitted. Here, the difference $w - u$ measures the advantage, as cultural models, of individuals with the trait. More broadly, transmissibility may relate to specific characteristics of traits that make them easy to learn or otherwise acquire (e.g., affordable or widely marketed clothing items), or to their psychological appeal (often referred in cultural evolution literature as content-bias, see [36]), and so on. For example, it has been shown that traits that elicit emotional reactions in general [37], or particular emotion like disgust [38] tend to be transmitted more than traits that do not.

On the other hand, since individuals have many occasions to learn new traits and replace the existing ones, denominators in indices of cultural success measure the resistance of traits in being relinquished. Here, the differ-

ence $v - u$ measures how easily individuals lacking the preference abandon a trait. This suggests that traits that are particularly memorable, useful, or otherwise durable (e.g., tattoos) should enjoy an advantage in cultural evolution. Here we have shown that such ratios are not only theoretically important, but yield insight into actual cultural processes.

Acknowledgments

Work supported by the Uniquely Human project funded by the Swedish Research Council. SG is on leave of absence from the University of Bologna.

References

1. Belleau BD (1987) Cyclical fashion movement: Women's day dresses: 1860 - 1980. Clothing and Textiles Research Journal 5: 15–20.

2. Bentley RA, Lipo CP, Herzog HA, Hahn M (2007) Regular rates of popular culture change reflect random copying. Evolution and Human Behavior 28: 151–158.

3. Shepard D (1972) Fashion in christian names. In: Willis G, Midgley D, editors, Fashion marketing: An anthology of viewpoints and perspectives, London: George Allen & Unwin.

4. Bentley RA, Hahn M, Shennan SJ (2004) Random drift and culture change. Proceedings Royal Society London B 271: 1443–1450.

5. Berger J, Le Mens G (2009) How adoption spread affects the abandonment of cultural tastes. PNAS 106: 8146–8150.

6. Herzog HA, Bentley RA, Hahn MW (2004) Random drift and large shifts in popularity of dog breeds. Proceedings Royal Society London B 271: S353–S356.

7. Neiman FD (1995) Stylistic variation in evolutionary perspective: Inferences from decorative diversity and interassemblage distance in illinois woodland ceramic assemblages. American Antiquity 60: 7–36.

8. Bentley RA (2008) Random drift versus selection in academic vocabulary: An evolutionary analysis of published keywords. PLOS One 3: e3057.

9. Smith A (1759/2000) The theory of moral sentiments. Amherst, NY: Prometeus Books.

10. Kant I (1798/2006) Anthropology from a pragmatic point of view. Cambridge, UK: Cambridge University Press.

11. Simmel G (1904) Fashion. International Quarterly 10: 130-155.

12. Sapir E (1931) Fashion. In: Enclyclopedia of the Social Sciences, New York: Macmillan, volume VI. pp. 139-144.
13. Svendsen LFH (2006) Fashion: A philosophy. London: Reaktion Books.
14. Pesendorfer W (1995) Design innovation and fashion cycles. The American Economic Review 85: 771-792.
15. Bergman M (2003) When a fad ends: An agent-based model of imitative behavior. In: Computing in Economics and Finance.
16. Tassier T (2004) A model of fads, fashion, and group formation. Complexity 9: 51–61.
17. Crow JF, Kimura M (1970) An Introduction to Population Genetics Theory. New York: Harper and Row.
18. UNESCO (2004). Index translationum. http://databases.unesco.org/xtrans/stat.
19. Social Security Administration of the USA (2009). Popular baby names. http://www.ssa.gov/OACT/babynames.
20. Mesoudi A, Lycett SJ (2009) Random copying, frequency-dependent copying and culture change. Evolution and Human Behavior 30: 41 – 48.
21. Chaiken S, Wood W, Eagly AH (1996) Principles of persuasion. In: Higgins ET, Kruglanski A, editors, Social psychology: Handbook of basic mechanisms and processes, New York: Guilford Press.
22. Haugtvedt CP, Herr P, Kardes F, editors (2008) The Handbook of Consumer Psychology. New York and London: Lawrence Erlbaum Associates.
23. Rendell L, Fogarty L, Hoppitt WJE, Morgan TJH, Webster MM, et al. (2011) Cognitive culture: theoretical and empirical insights into social learning strategies. Trends in Cognitive Sciences 15: 68–76.
24. Ghirlanda S, Enquist M, Nakamaru M (2006) Cultural evolution develops its own rules. the rise of conservatism and persuasion. Current Anthropology 47: 1027–1034.
25. Acerbi A, Enquist M, Ghirlanda S (2009) Cultural evolution and individual development of openness and conservatism. PNAS 106: 18931–18935.
26. Ghirlanda S, Acerbi A, Enquist M, Nakamaru M (2010) The sometimes evitable route to conservatism and persuasiveness. Current Anthropology 51: 271-272.
27. Strimling P, Enquist M, Eriksson K (2009) Repeated learning makes cultural evolution unique. PNAS 106: 13870–13874.
28. Tanaka MM, Kendal JL, Laland KN (2009) From traditional medicine to witchcraft: Why medical treatments are not always efficacious. PLOS One 4: e5192.

29. Train K (2003) Discrete choice models with simulation. Cambridge, UK: Cambridge University Press.

30. Enquist M, Ghirlanda S (2005) Neural Networks and Animal Behavior. Princeton University Press.

31. Strimling P, Sjostrand J, Enquist M, Eriksson K (2009) Accumulation of independent cultural traits. Theoretical Population Biology 76: 77–83.

32. Lehmann L, Aoki K, Feldman MW (2011) On the number of independent cultural traits carried by individuals and populations. Phil Trans R Soc 366: 424–435.

33. Fink T, Mao Y (1999) The 85 ways to tie a tie. London: Fourth Estate.

34. Davis F (1992) Fashion, Culture and identity. Chicago: Chicago University Press.

35. Berger J, Heat C (2008) Who drives divergence? identity-signaling, outgroup dissimilarity, and the abandonment of cultural tastes. Journal of Personality and Social Psychology 95: 593–607.

36. Boyd R, Richerson PJ (1985) Culture and the Evolutionary Process. Chicago: Chicago University Press.

37. Berger J, Milkman K (2012) Social transmission, emotion, and the virality of online content. Journal of Marketing Research forthcoming.

38. Heat C, Bell C, Sternberg E (2001) Emotional selection in memes: The case of urban legends. Journal of Personality and Social Psychology 81: 1028–1041.

Supporting information for "The logic of fashion cycles," A. Acerbi, S. Ghirlanda, and M. Enquist.

Model 1

Derivation

Equations (1–4) in the main text are derived considering the possible outcomes of interactions of observers of type i with models of type j, with $i, j \in \{0, P, T, PT\}$. According to the rules given in the main text, the probability that an observer of type i copies a model of type j is given by the entries in the following array:

	Model:			
Observer:	0	P	T	PT
0		u	u	u
P	u		$\boldsymbol{w}$	$\boldsymbol{w}$
T	$\boldsymbol{v}$	u		u
PT	u	u	$\boldsymbol{w}$	

where the entries in boldface reflect our assumptions on the effect of preference and lack of preference on cultural transmission. Interactions of models and observers of the same type are not considered as they do change the frequency of types.

Transition rates between types are constructed as follows. Suppose, for instance, that an observer of type PT meets a model of type T. Trait and preference are copied independently with probability w. Thus with probability $w(1-w)$ the observer copies the model's preference value but not its trait value, resulting in the observer changing from PT to T. With probability $w(1-w)$ the observer copies the model's trait value but not its preference value, which results in no change in the observer's type. With probability w^2 the observer copies both the trait and preference values, resulting in a change from type PT to T. Hence in interactions between PT observers and T models there is overall a probability $w(1-w) + w^2 = w$ that the observer changes from PT to T, and a probability $1-w$ that the observer does not change. Since encounters between PT and T occur at a rate of $x_{PT}x_T$, the overall rate at which such transitions occur is $w\, x_{PT}x_T$. Table S1 shows the rates of all possible transitions, calculated in this same way (the one just calculated is entry 20). Equations (1–4) in the main text

	Observer	Model	Outcome	Rate
1.	0	P	P	$u\, x_0 x_P$
2.	0	T	T	$u\, x_0 x_T$
3.	0	PT	P	$u(1-u)\, x_0 x_{PT}$
4.	0	PT	T	$u(1-u)\, x_0 x_{PT}$
5.	0	PT	PT	$u^2\, x_0 x_{PT}$
6.	P	0	0	$u\, x_P x_0$
7.	P	T	0	$w(1-w)\, x_P x_T$
8.	P	T	T	$w^2\, x_P x_T$
9.	P	T	PT	$w(1-w)\, x_P x_T$
10.	P	PT	PT	$w\, x_P x_{PT}$
11.	T	0	0	$v\, x_0 x_T$
12.	T	P	0	$u(1-u)\, x_P x_T$
13.	T	P	PT	$u(1-u)\, xP x_T$
14.	T	P	P	$u^2\, x_T x_P$
15.	T	PT	PT	$u\, x_T x_{PT}$
16.	PT	0	0	$u^2\, x_0 x_{PT}$
17.	PT	0	T	$u(1-u)\, x_0 x_{PT}$
18.	PT	0	P	$u(1-u)\, x_0 x_{PT}$
19.	PT	P	P	$u\, x_P x_{PT}$
20.	PT	T	T	$w\, x_T x_{PT}$

Table S1. Model 1 transitions.

follow from calculating the net effect of these transitions as follows:

$$\dot{x}_0 = -(1:5) + 6 + 7 + 11 + 12 + 16 \tag{11}$$
$$\dot{x}_P = -(6:10) + 1 + 3 + 14 + 18 + 19 \tag{12}$$
$$\dot{x}_T = -(11:15) + 2 + 4 + 8 + 17 + 20 \tag{13}$$
$$\dot{x}_{PT} = -(16:20) + 5 + 9 + 10 + 13 + 15 \tag{14}$$

where numbers refer to lines in Table S1 and $a : b$ indicates the range from a to b inclusive.

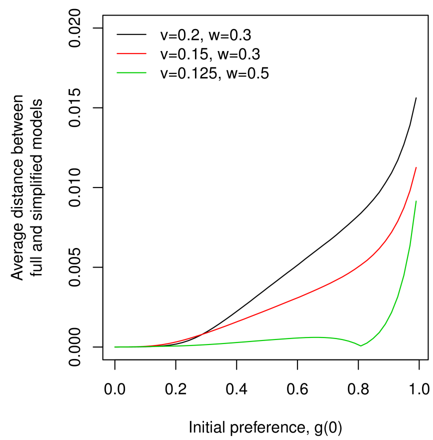

Figure S1. Absolute difference between trait frequency ($f = x_T + x_{PT}$) according to the full Model 1 (equations 1–4 in the main text) and the simplified model in equations (6,7,9) as a function of initial preference, $g(0)$, and for different combinations of v and w parameters ($u = 1$). Trait frequency in the simplified model lies within 2% of the frequency given by the full model.

Analysis

We study the model using a mix of analytical and numerical methods as follows. We noticed numerically that the equation for $\dot{x_0}$ (equation 1 in the main text) can be simplified as

$$\dot{x}_0 = (v - u)\, x_0 x_T \tag{15}$$

as the sum of the two other terms is between 10 and 100 times smaller than $(v - u)\, x_0 x_T$ over a range of initial conditions and parameter values (Figure S1). This simplification allows us to write a closed system for x_0 and $g = x_P + x_{PT}$:

$$\dot{x}_0 = (v - u)\, x_0\, (1 - g - x_0) \tag{16}$$

$$\dot{g} = -(w - u)\, g\, (1 - g - x_0) \tag{17}$$

where we have eliminated x_T through the identity

$$x_T + x_0 + g = 1 \tag{18}$$

The description of the system is completed by the equation for the trait frequency $f = x_T + x_{PT}$ (equation 6 in the main text), which in terms of the variables x_0, g, and f is rewritten as:

$$\dot{f} = -(v - u)\, x_0(1 - g - x_0) - (w - u)\, f(1 - f - x_0) \tag{19}$$

Equilibria

We denote equilibrium values by a superscript *. Equilibria of equations (16–17) are of the form

$$x_0^\star + g^\star = x_0^\star + x_P^\star + x_{PT}^\star = 1 \tag{20}$$

implying

$$x_T^\star = 0 \tag{21}$$

Given equation (20) and the equilibrium condition $\dot{f} = 0$, equation (19 implies

$$f^\star(1 - f^\star - x_0^\star) = 0 \tag{22}$$

which, using equation (21) and expanding $f^\star = x_T^\star + x_{PT}^\star$ yields

$$x_{PT}^\star x_P^\star = 0 \tag{23}$$

Implying that either $x_P^\star = 0$ or $x_{PT}^\star = 0$, or both. Equation (4) in the main text, however, implies $x_{PT}^\star = 0$ for all $x_0^\star > 0$. In conclusion, equilibria of the model are of the form:

$$x_0^\star + x_P^\star = 1 \qquad x_{PT}^\star = x_T^\star = 0 \tag{24}$$

hence the trait cannot persist in the population, but the preference can. The following analysis characterizes system trajectories and shows that $x_P^\star$ is generally small.

System trajectories

Because x_0 and g are frequencies, they must be ≥ 0 and must sum to a number less than 1. Thus equations (16) and (17) hold in the triangle defined by $x_0 \geq 0$, $g \geq 0$, $1 - x_0 - g \geq 0$. In its interior, $\dot{g} < 0$ and $\dot{x}_0 > 0$ always. Taking the ratio of equations 16 and 17 we obtain a differential equation for the shape of the system trajectory as a curve $g = g(x_0)$:

$$\frac{dg(x_0)}{dx_0} = \frac{dg}{dt}\frac{dt}{dx_0} = -\frac{w-u}{v-u}\frac{g}{x_0} \tag{25}$$

The combination of parameters $(w-u)/(v-u)$ recurs often, hence we define

$$\phi = \frac{w-u}{v-u} \tag{26}$$

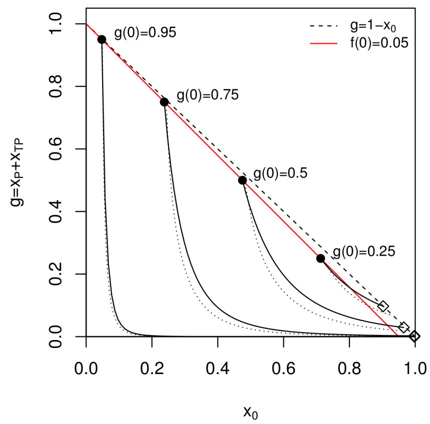

Figure S2. Model 1 trajectories. Sample model trajectories in the (x_0, g) plane of the simplified system in equations 6–7 (Supporting Information), for different initial frequencies of the preference, $g(0)$, and initial frequency of the trait $f(0) = 0.05$. Trajectories start at the closed circle and end at the open diamond. The dashed line is the line $g = 1 - x_0$, which delimits the state space together with the lines $x_0 = 0$ and $g = 0$. The red line is the locus of all starting conditions with $f(0) = 0.05$ (assuming the trait and the preference are initially distributed independently). The dotted lines are trajectories of the full system, equations 1–4 in the main text, showing the quality of our approximation.

The solution of (25) is thus

$$g(x_0) = C\, x_0^{-\phi} \tag{27}$$

where the constant C is set from initial conditions as

$$C = g(0)\,(x_0(0))^{\phi} \tag{28}$$

Examples of these trajectories are in Figure S2 given the initial condition $f(0) = 0.05$ and varying $g(0)$ as indicated. All trajectories end on the line $g = 1 - x_0$, meaning that $x_T^\star = 0$; as showed in the previous section, $x_{PT}^\star = 0$ as well. The value of $x_0^\star$, and thus of $x_P^\star = 1 - x_0^\star$ is the solution of

$$C(x_0^\star)^{-\phi} + x_0^\star - 1 = 0 \tag{29}$$

with the highest value (there are two solutions, as it is apparent from Figure S2). This equation usually yields very low values of x_P (Figure S3).

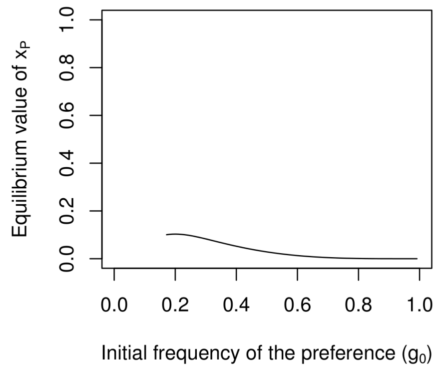

Figure S3. Frequency of the preference for a trait at the end of a fashion cycle ($x^\star_P = 1 - x^\star_0$ in equation 15 (Supporting Information) as a function of initial frequency of the preference, $g(0)$. All parameters as in Figure S2.

Existence of fashion cycles

We have established in the previous sections that all fashions eventually die in this model. A fashion cycle occurs when, before disappearing, a trait initially increases in frequency. We determine here, for given parameters u, v, and w the minimum initial frequency of the preference, $g_{\min}(0)$, for a cycle to occur.

Given initial conditions $f(0)$ and $g(0)$ the condition for f to initially increase is, from equation (19):

$$-(v-u)\,x_0(0)(1-g(0)-x_0(0)) + (w-u)\,f(0)(1-f(0)-x_0(0)) > 0 \quad (30)$$

Assuming that trait and preference are initially distributed independently, implying the initial conditions $x_{PT}(0) = f(0)g(0)$, $x_T(0) = f(0)(1-g(0))$, and $x_P(0) = g(0)(1-f(0))$, we have

$$x_0(0) = 1 - f(0) - g(0) + f(0)g(0) \quad (31)$$

Substituting this expression in equation (30) we get the condition

$$g^2_{\min}(0) - (2+\phi)g_{\min}(0) + 1 < 0 \quad (32)$$

where $\phi = \frac{w-u}{v-u}$ as defined in equation (26). The solution is

$$g_{\min}(0) > 1 + \frac{\phi}{2}\left(1 - \sqrt{1 + \frac{4}{\phi}}\right) \quad (33)$$

(The second solution of the quadratic equation (32) is of no interest as it is always > 1). Note that this condition does not rely on the simplified dynamics of x_0 in (16) and, moreover, it is independent of $f(0)$. Thus, for a given value of ϕ, the initial preference determines whether a fashion cycle occurs irrespective of the initial frequency of the trait.

Maximum frequency

A differential equation for trait frequency as a function of x_0 is obtained proceeding as in equation (25):

$$\frac{df(x_0)}{dx_0} = -1 + \phi \frac{f(1-f-x_0)}{x_0(1-g(x_0)-x_0)} \tag{34}$$

This equation, however, does not appear to have a closed form solution as equation (25), hence it is not possible to determine analytically the maximum frequency attained. It is possible to obtain lower and upper bounds through the identity $x_T \leq f \leq x_T + g$, accurate to about 10% for small $g(0)$ to near-perfect for large $g(0)$, but the formulae are not telling and we prefer to present numerical results. Figure S5 shows that the maximum frequency grows approximately linearly with $g(0)$, and is not greatly influenced by ϕ.

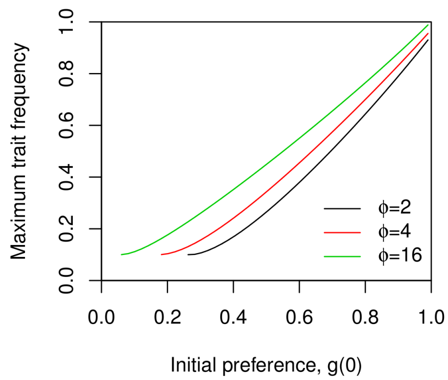

Figure S4. Maximum frequency attained during a trait's fashion cycle, as a function of initial preference, $g(0)$, and system parameters, ϕ. Initial frequency is $f(0) = .0.05$.

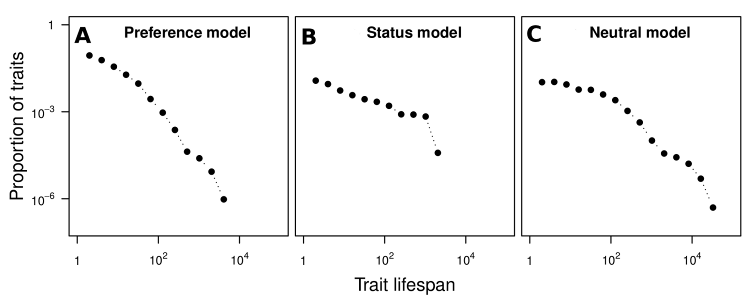

Figure S5. Distribution of trait lifespans for the preference, status and neutral models for the simulations described in the main text.

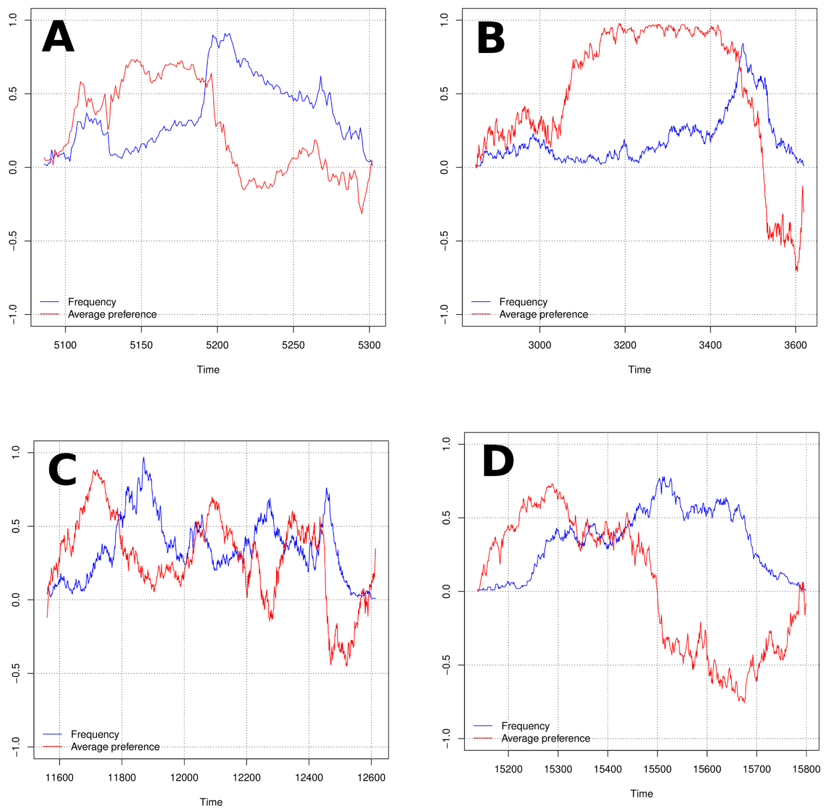

Figure S6. Examples of frequency-preference dynamics in simulations of the multi-trait preference model. (A) The preference (red) for a trait rises in the population, which causes a rise in frequency (blue). As the trait becomes common, the preference falls and eventually the trait declines in frequency. See main text for discussion. (B) Another example of the same dynamics, showing that the latency between rise in preference and rise in frequency may be long. (C) A trait undergoing multiple fashion cycles. Trait revival is possible by either chance fluctuation or because effective models adopt the trait again. (D) A trait that remains popular for some time, despite not being preferred. This may happen because a common trait is likely to be possessed by successful cultural models, hence it can be copied even if it does not contribute to the model's success. A real-life example may be common names such as George, who may not be perceived as particularly catchy but is nevertheless associated with successful individuals such as George Washington, George Harrison, or George Clooney.