

Evolution of music by public choice

Robert M. MacCallum^{a,1}, Matthias Mauch^{b,2}, Austin Burt^a, and Armand M. Leroi^a

^aDepartment of Life Sciences, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom; and ^bMedia Interaction Group, National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba, Ibaraki 305-8568, Japan

Edited* by Richard E. Lenski, Michigan State University, East Lansing, MI, and approved May 10, 2012 (received for review February 27, 2012)

Music evolves as composers, performers, and consumers favor some musical variants over others. To investigate the role of consumer selection, we constructed a Darwinian music engine consisting of a population of short audio loops that sexually reproduce and mutate. This population evolved for 2,513 generations under the selective influence of 6,931 consumers who rated the loops' aesthetic qualities. We found that the loops quickly evolved into music attributable, in part, to the evolution of aesthetically pleasing chords and rhythms. Later, however, evolution slowed. Applying the Price equation, a general description of evolutionary processes, we found that this stasis was mostly attributable to a decrease in the fidelity of transmission. Our experiment shows how cultural dynamics can be explained in terms of competing evolutionary forces.

culture | algorithm

The music made by the world's cultures is immensely diverse (1, 2). Because music is transmitted from one musician to another, and frequently modified in transmission, this diversity must arise from descent by modification rather like the diversity of living things, languages, and other cultural artifacts (3). What drives this process? It is often supposed that the music we listen to is primarily the product of aesthetic decisions made by "producers" (i.e., composers, performers) (4). Early Greek texts speak of specialist composers/performers, and the rudiments of formal musical theory, at least 2,500 y ago (5), and specialist composer/performers are found in many other societies as well (6). However, the reproduction, spread, and persistence of particular songs must also depend on the preferences of "consumers" (i.e., the people who listen to them) (7). These preferences are also clearly a selective process and, like any selective process, can have a creative role (8). Disentangling the roles of composers, producers, and consumers in shaping musical diversity is difficult in existing musical cultures.

For this reason, inspired by studies of experimental evolution in microbes (9, 10), digital organisms (11–16), and previous work on evolutionary music and art (17–19), we developed an artificial system for studying musical evolution called "DarwinTunes." Evolutionary music studies, to date, have either attempted to automate the selection process (20) or have focused on the development of single-user composition (21) and performance (22) aids. In DarwinTunes, however, fitness is defined by the aesthetic tastes of the public at large or some more selected body of musical consumers. As such, it implements aesthetic evolution in a complex natural environment. Using DarwinTunes then, we ask: Is it possible to make music without a composer? If so, what kind of music is made? What limits the evolution of music?

DarwinTunes: A Musical Variation-Selection Engine

In brief, DarwinTunes works as follows: An algorithm maintains a population of tree-like digital genomes, each of which encodes a computer program. When a program is executed, a short, seamlessly looping polyphonic sound sequence, a loop, is produced deterministically. Each genome/program specifies note placement, instrumentation, and performance parameters; however, tempo, meter, and tuning system are fixed for all loops. No human-derived sounds, rhythms, or melodies are provided as input to the algorithm. During the experiments, loops periodically

replicate to produce new loops. The daughter loops are not, however, identical to their parents for two reasons. First, in a process analogous to recombination, the genome of each daughter loop is formed from the random combination of its two parents' genomes. Second, in a process analogous to mutation, each daughter also contains new, random genetic material. These two processes mimic the fusion of existing, and invention of novel musical motifs, rhythms, and harmonies that can be heard in musical evolution (6). The only selective pressure in DarwinTunes comes from a population of consumers who listen to samples of the loops via a Web interface and rate them for their appeal. These ratings are then the basis of a fitness function that determines which loops in a given generation will be allowed to mate and reproduce. We therefore expect that the frequency of musical traits will evolve under the influence of this selective process rather as trait frequencies in organisms do under the influence of natural selection.

The processes underlying a single DarwinTunes population are shown in Fig. 1A. At any given time, a DarwinTunes population has 100 loops, each of which is 8 s long. Consumers rate them on a five-point scale ("I can't stand it" to "I love it") as they are streamed in random order. When 20 loops have been rated, truncation selection is applied whereby the best 10 loops are paired, recombine, and have two daughters each. These daughters replace their parents that die. In our first experiment, designated experimental population 1 (EP1), we began with 100 loops that had been generated from two random founders and that were then allowed to evolve for 100 generations without selection to maximize the stock of standing variation in the starting population. Public consumers were then recruited to rate the loops. Because consumers did not know each other's ratings, there is no possibility of social influence on rating (7). In all, 6,931 consumers made 85,533 ratings over the course of 2,513 generations of evolution during which 50,480 loops were born. We recorded the ratings, number and identity of offspring, and genome of each loop. Thus, we can describe the evolutionary dynamics of these populations in detail.

Results

Selection Rapidly Evolves Music from Noise but Then Stops. As EP1 evolved, it seemed to us that the loops were becoming more pleasing to listen to, and that we were, in fact, evolving music from noise (audio examples have been archived at doi:10.5061/dryad.h0228 and can be heard at <http://soundcloud.com/uncoolbob/sets/darwintunes/>). To test this objectively, we carried

Author contributions: R.M.M. and A.M.L. designed research; R.M.M. and M.M. performed research; R.M.M. and M.M. contributed new reagents/analytic tools; R.M.M., M.M., A.B., and A.M.L. analyzed data; and A.M.L. wrote the paper.

The authors declare no conflict of interest.

*This Direct Submission article had a prearranged editor.

Freely available online through the PNAS open access option.

Data deposition: The data reported in this paper have been deposited in the Dryad Repository, <http://datadryad.org> (DOI no. 10.5061/dryad.h0228).

¹To whom correspondence should be addressed. E-mail: uncoolbob@gmail.com.

²Present address: Centre for Digital Music, School of Electronic Engineering and Computer Science, Queen Mary, University of London, London E1 4NS, United Kingdom.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1203182109/-DCSupplemental.


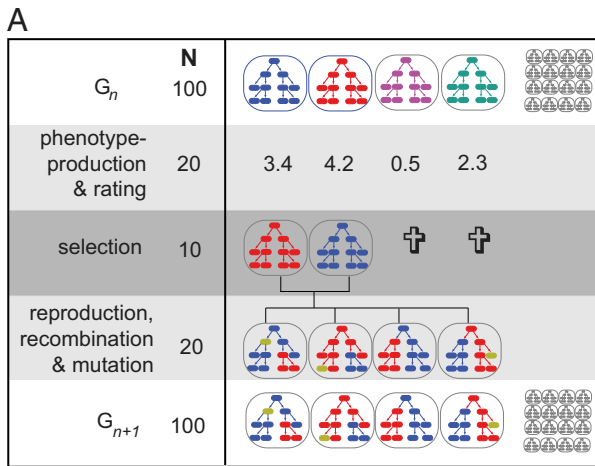


Fig. 1. (A) Evolutionary processes in DarwinTunes. Songs are represented as tree-like structures of code. Each generation starts with 100 songs; however, for clarity, we only follow one-fifth of them. Twenty songs are randomly presented to listeners for rating, and the remaining 80 survive until the next generation; thus, at any time, the population contains songs of varying age. Of the 20 rated songs, the 10 best reproduce and the 10 worst die. Reproductives are paired and produce four progeny to replace themselves and the dead in the next generation. The daughters' genomes are formed from their parents' genomes, subject to recombination and mutation. (B) Evolution of musical appeal. During the evolution of our populations, listeners could only listen to, and rate, songs that belonged to one or, at best, consecutive generations. Here, they were asked to listen to, and rate, a random sample of all the songs that had previously evolved in the public population, EP1. Thus, these ratings can be used to estimate the mean absolute musical appeal, M , of the population at any time. To describe the evolution of M , we fitted an exponential function. Because the parameter that describes the rate of increase of M is significantly greater than zero, M increases over the course of the experiment (*SI Appendix, B.1*).

out a new experiment. We randomly sampled 2,000 of the 50,480 loops produced at any time during EP1's evolution and, via a Web interface, asked public consumers to rate them as before. Because consumers heard and rated loops sampled from the entire evolutionary trajectory in this experiment, their ratings can be used to estimate the mean absolute musical appeal, M , of the population at any time. This is analogous to bacterial experiments in which the fitness of an evolved strain is compared directly with that of its ancestor (10). Fig. 1B shows that M increased rapidly for the first 500–600 generations but then came to equilibrium. Thus, in our system, musical quality evolves, but it seems that it does not do so indefinitely.

What makes the loops of later generations so much more pleasing? The aesthetic value of a given piece of music depends on many different features, such as consonance, rhythm, and melody (23). In recent years, music information retrieval (MIR) technology has permitted the automatic detection of some of these features (24–26); reasoning that our raters listen to, and like, Western popular music, we measured the phenotypes of our loops using two MIR algorithms designed to detect features in this music. The first, Chordino, detects the presence of chords commonly used in the Western repertoire (27). The fit of a loop to Chordino's canonical chord models is given by a log-likelihood value C_L and is an estimate of the clarity of the chordal structure. The second, Rhythm Patterns (28), extracts a rhythmic signature, from which we derive a complexity measure, R . To validate these algorithms, we tested them on a standardized test set of specifically generated loops (*SI Appendix, A.3*).

To examine the evolution of musical qualities in EP1, we measured C_L and R for every loop. We found that, like musical appeal, these traits increased rapidly over the first 500–600 generations but then appear to fluctuate around a long-term mean (Fig. 2A and B). Given these dynamics, and because C_L and R are measured without error, we are able to model their evolution using a discrete version of the Ornstein–Uhlenbeck (O-U) process, according to which the change in the mean of a character from one generation to the next is anticorrelated to how far it is from a long-term mean:

$$\Delta \bar{z} = a(u - \bar{z}) + \epsilon,$$

where $\Delta \bar{z}$ is the difference between the means of each offspring and parental generation, $\bar{z}_o - \bar{z}_p$; a is a constant such that $a > 0$; u is the long-term mean; and ϵ is a normally distributed random variable with a mean = 0. For both C_L and R , the confidence limits on the long-term mean do not include the initial values ($P = 1.0 \times 10^{-6}$ and $P = 2.0 \times 10^{-7}$, respectively), confirming the visual impression that C_L and R increased significantly over the course of the experiment (Fig. 2A and B and *SI Appendix, Table S6*).

Because musical appeal and its components all increase, they are probably being selected. However, the trajectory of a DarwinTunes population, like that of any evolving population, depends not only on selection but on stochastic sampling, the analog of genetic drift. In experimental evolution, replicable responses are a signature of selection (10, 29). Hence, to determine whether the increases in chordality and rhythmicity are idiosyncratic to the preferences of the particular set of consumers contributing to EP1, or perhaps are attributable to chance correlations between these characters and other characters that were the true targets of selection, we repeated the experiment in a more controlled setting. To do this, we cloned additional populations from the same base population that EP1 started with and asked undergraduates to rate them. These populations, designated EP2 and EP3, were allowed to evolve independently for about 400 generations and received an average of 10,683 ratings. We found that C_L and R also increased rapidly in these populations, again to a plateau (*SI Appendix, B.1* and *SI Appendix, Fig. S5*). As controls, we generated 1,000 additional populations with the same origin as the experimental populations and subject to the same variational processes and demography for 400 generations but differing from them in that ratings were assigned randomly rather than by consumers. We found that mean C_L and R of the selected populations were significantly higher than those of the unselected control populations by generation 100 (Fig. 2C and D and *SI Appendix, B.2*). We also used the control populations to examine whether C_L and R are intrinsically related to each other and found that they are weakly correlated, $r = 0.26(\pm 0.016)$ [mean ($\pm 95\%$ confidence interval)] (*SI Appendix, B.3*). Thus, although selection on

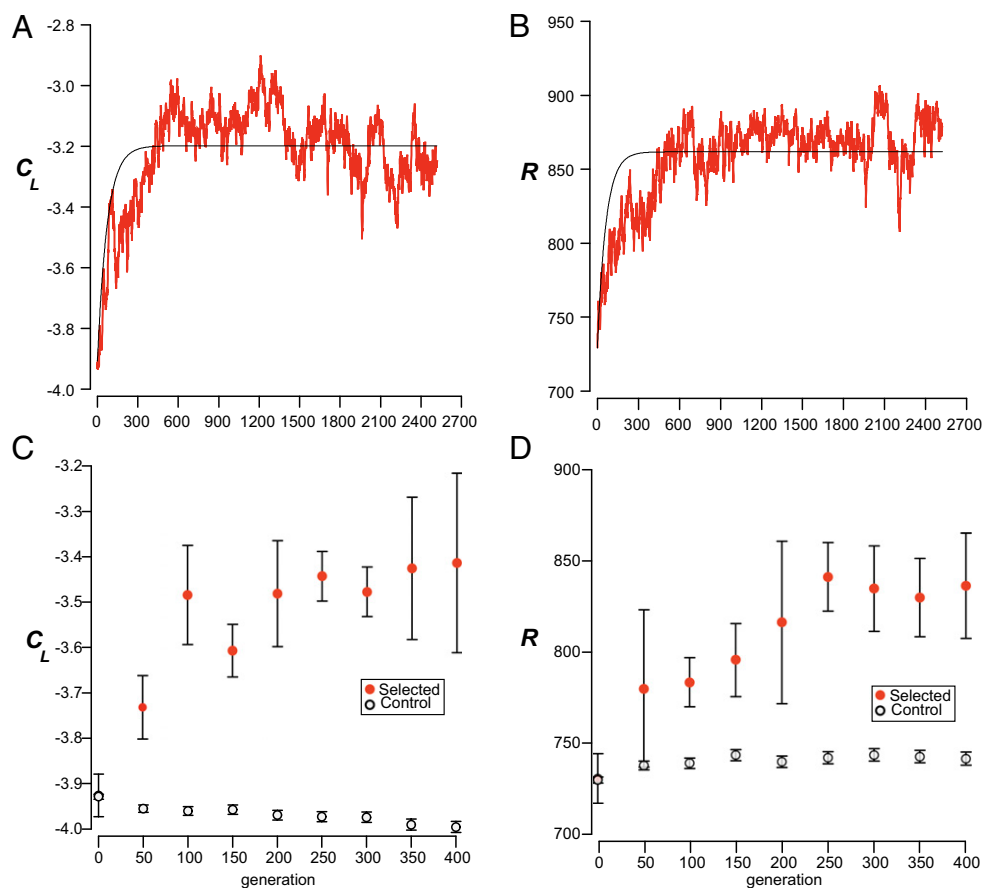


Fig. 2. Evolution of musical attributes. (A) Evolution of chordal clarity, C_L , in EP1. (B) Evolution of rhythm, R , in the public population, EP1. Both features were fitted with an O-U model that includes a stochastic component. In the fits shown, the stochastic parameter, s , was set to zero for the sake of clarity; however, during model fitting, s was included as a freely varying parameter. *SI Appendix, Fig. S5* shows equivalent plots for the replicate populations. (C) Evolution of C_L in three selected populations (EP1–EP3) and 1,000 unselected control populations over 400 generations. (D) Evolution of R in three selected populations (EP1–EP3) and 1,000 unselected control populations over 400 generations. Error bars represent 95% confidence intervals estimated by a linear mixed model. By generation 100, both C_L and R are significantly elevated in the selected populations, compared with the control unselected populations.

one of these features may influence the evolution of the other, they are largely independent. We cannot, however, preclude the possibility that either feature is highly correlated with unmeasured traits that are more direct targets of selection.

Variation and Adaptation in DarwinTunes Populations. The increase in C_L and R implies that selection is directional. Thus, why do our populations stop evolving? Remarkably, it is not merely that these traits cease to evolve: Musical appeal itself does also. This pattern of fast-slow evolution or even stasis is often seen in biological populations, whether in the laboratory, wild, or fossil record. Stasis can result from several different population genetic forces; however, it has often been difficult to distinguish among them (10, 30–32). Because we know the complete histories of the DarwinTunes populations, we can study the forces driving their evolutionary dynamics in detail. We first considered the possibility that DarwinTunes populations have arrived at an adaptive peak, such that selection, which was previously directional, now stabilizes the population means. To investigate this, we estimated selective surfaces using multivariate cubic-spline regression (33) and plotted adaptive walks on them. Fig. 3A shows that EP1 has a single adaptive peak near high R and C_L and that although it walks erratically up the slope toward the peak, it does not reach it. Very rhythmic loops (very high R) may be less fit than slightly less rhythmic ones; even so, it is clear that EP1 has stopped evolving at least 1 SD in each dimension away

from its adaptive peak; thus, stasis is not attributable to an absence of selection. Interestingly, the topology of the EP1 adaptive landscape suggests that R and C_L have a synergistic effect on fitness: high C_L loops are especially fit when they have a high R as well; a model with $C_L \times R$ interaction explains significantly more of the variation than one without it. A similar interaction is found in EP2 but not EP3 (*SI Appendix, B.4*).

We next considered the possibility that the populations have simply run out of genetic variation and that they have become fixed for all beneficial variants. Fig. 3B and C show the frequency distributions of C_L and R over the evolution of EP1. The rapid progress of the population before generation 1,000 is associated with a decrease in frequency of loops with the lowest chordal clarity and rhythmic complexity, likely attributable to selection. However, as the population continues to evolve, new low C_L and R loops are reintroduced by mutation or recombination, and throughout the evolution of the populations, many loops have higher C_L and R values than the long-term O-U mean. Thus, the lack of progressive evolution after about generation 500 is not attributable to fixation of high C_L and R variants and complete exhaustion of genetic variation. This is also true for EP2 and EP3 (*SI Appendix, Fig. S8*).

Applying the Price Equation. To probe the forces acting on these populations further, we made use of the Price equation (34–37). The Price equation, a general description of all evolutionary

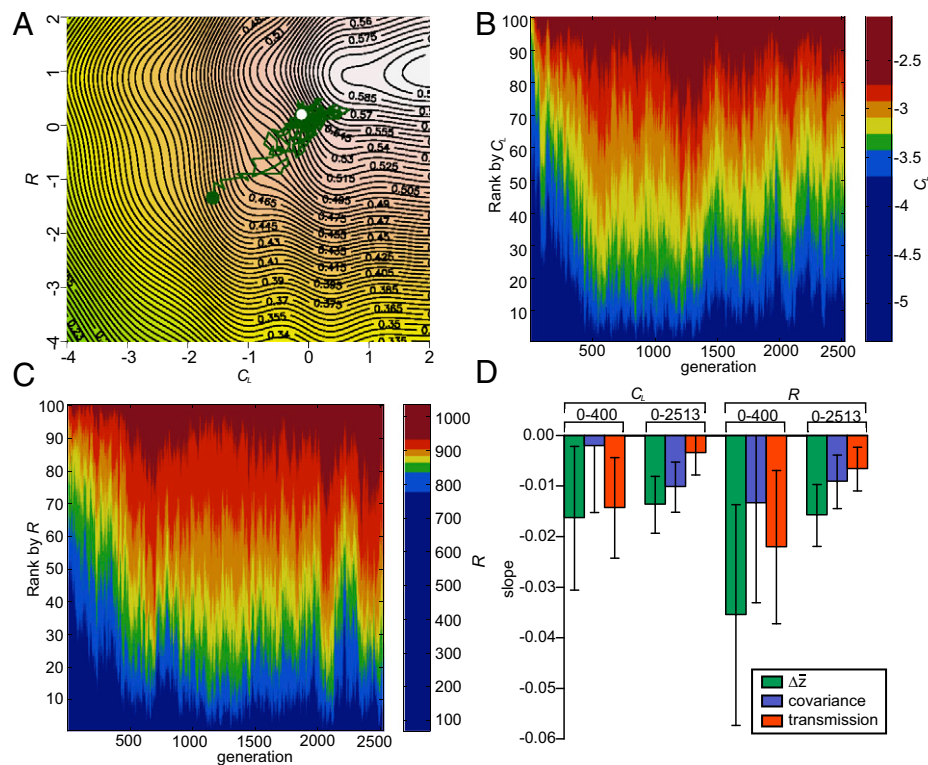


Fig. 3. Explaining the evolutionary dynamics of the public population, EP1. (A) Adaptive surfaces and evolutionary trajectory. Data have been normalized to mean = 0 and SD = 1. The contours show fitness (probability of reproduction) and are based on cubic-spline regressions using data from all generations. The dark green lines show the progress of the bivariate means of the populations binned over 10 generation intervals for clarity, and the green and white circles are the start and end generations, respectively. This shows that the last generation and many previous generations do not approach the adaptive peak; thus, the failure of the population to progress cannot be attributable to stabilizing selection. EP1 has a single adaptive peak; a model with $C_L \times R$ interaction explains significantly more of the variation than one without ($P = 1.8 \times 10^{-5}$, log-likelihood test). Similar analyses of EP2 and EP3 can be found in *SI Appendix, Fig. S7*. (B and C) Frequency distributions of C_L and R in the public population, EP1, over generations, unnormalized data. These show that at no point does the population become fixed for high C_L or R variants; thus, the failure of the population to progress cannot be attributable to complete exhaustion of variation in these traits. Similar analyses of EP2 and EP3 can be found in *SI Appendix, Fig. S8*. (D) Change in Price parameters as a function of the current value, z , in EP1. As evolution proceeds, the change in mean from one generation to the next, $\Delta \bar{z}$, declines. This can be attributable to a decline in the covariance term, the transmission term, or both. In the first 400 generations, for both C_L and R , only the transmission term shows a significant decline, suggesting that the initial decline in the rate of evolution is attributable to an increased mutational or recombinational load. This is comparable to what is seen in the replicate populations, EP2 and EP3, over the same time period. Considering all 2,513 generations, however, the covariance term also declines, suggesting that either the intensity of selection or variability also contributes to population stasis in the long term. The latter proves to be the case (*SI Appendix, B.6*). Error bars are twice the SE of the estimate.

processes, decomposes the mean response to selection in a given generation, $\Delta \bar{z}$, into a covariance term that describes the effect of selection and a transmission term that describes the effect of inheritance:

$$\begin{aligned} \Delta \bar{z} &= \text{covariance term} + \text{transmission term} \\ &= \text{cov}(w/\bar{w}, z) + E((w/\bar{w})\Delta z) \end{aligned}$$

where z is the phenotype of an individual, \bar{z} is the mean phenotype of the population; w is the fitness of an individual (how many offspring it has); \bar{w} is the mean fitness of the population; $\Delta z = z_o - z_p$, where z_p is the phenotype of an individual and z_o is the mean phenotype of its offspring; and $\Delta \bar{z}$ has been defined previously. The covariance term in any generation is the product of the population variance, $\text{var}(z)$, and the strength of directional selection, which, in turn, can be estimated as the slope of a linear regression of the fitness of parents on a phenotype, $\beta(w, z)$. The transmission term is based on the phenotypic similarity of parents to their offspring, and thus estimates the fidelity of transmission: When it is zero, inheritance is perfect; when it is negative, offspring have a lower phenotype than their parents; and when it is positive, offspring have a greater phenotype than their parents.

At evolutionary equilibrium, $\Delta \bar{z} = 0$, the covariance and transmission terms must be equal in magnitude but opposite in sign. Given that our populations appear to be at equilibrium, one or both of these terms must have changed during their evolution. However, which term changed? As noted above, in an O-U process, the expected change from one generation to the next is a linear function of the current value with a negative slope (i.e., changes are expected to be positive when the current value is below the long-term mean and negative when it is above the long-term mean). We now decompose $\Delta \bar{z}$ into covariance and transmission terms, and we test whether either changes as a function of the mean, \bar{z} . Considering only the first 400 generations, before C_L and R approach equilibrium, the change in trait value attributable to selection (i.e., the covariance term) is independent of the current value in all cases but the amount by which offspring differ from their parents (i.e., the transmission term) becomes increasingly negative as the current value increases (Fig. 3D and *SI Appendix, B.6*). This indicates that the fidelity of transmission becomes an increasing impediment to progress as adaptation proceeds. It is this factor that causes evolution to slow down as C_L and R increase over the first 400 generations.

Transmission Fidelity and the Limits of Directional Evolution. In organisms, a decrease in the fidelity of transmission could be attributable to an increase in environmental variance, recombination pressure, or mutation pressure (38). Because the genome for any DarwinTunes loop produces an identical sound file on all computers, there is no environmental variance; a decrease in the fidelity of transmission must therefore be attributable to an increase in recombination, mutation pressure, or both. Because the genomic rates of recombination and mutation were constant throughout the experiment, this increase cannot be attributable to an increase in the frequency of recombination or mutation but must be attributable to increasingly deleterious phenotypic effects. Recombination could have increasingly deleterious effects if, as the population evolves, high fitness comes to depend on particular genomic configurations that can be broken up by sex; in other words, fitness epistasis increases (39). If so, this explanation would be analogous to the epistatic effects that Dobzhansky and Muller thought were responsible for hybrid sterility and lethality (40). As noted above, there is some evidence for synergistic fitness epistasis between R and C_L . Perhaps loops with pleasing combinations of R and C_L are selected but then quickly broken up by recombination. If so, this would imply that these traits are controlled by different regions of the loops' genomes, but we do not know this, and rhythm and chordal clarity may themselves be influenced by multiple interacting loci. Alternatively, mutations may become increasingly deleterious as the populations become more adapted for the same reasons that R. A. Fisher inferred they do in organisms: the increasing vulnerability of complex, fine-tuned structures to change (41, 42). We cannot distinguish between these explanations for decrease in transmission fidelity in our populations, but further experiments may do so.

Curiously, if we consider all 2,513 generations of EP1, we get a different picture in which the transmission term is no longer significant for C_L and the covariance terms for both C_L and R show a significant decline (Fig. 3D). To investigate this further, we decomposed the covariance term into the strength of selection, $\beta(w, z)$, and the variance of the trait, $\text{var}(z)$. The slopes $\beta(w, C_L)$ and $\beta(w, R)$ are significantly positive in all cases, showing directly that both C_L and R were under directional selection (SI Appendix, B.6). As the population mean increases, $\beta(w, C_L)$ remains constant, whereas $\beta(w, R)$ increases significantly; thus, consistent with our impression from the adaptive landscapes, the long-term stasis of neither trait is attributable to a decline in the strength of directional selection. By contrast, both $\text{var}(C_L)$ and $\text{var}(R)$ decline as the population mean increases, implying that the long-term stasis of this population is at least partly attributable to a decrease in the amount of phenotypic variance present (SI Appendix, B.6). Thus, although recombination or mutation pressure limits adaptive evolution in the short term, in the longer term, even a subtle decline in the amount of genetic variation can do so as well.

Because C_L and R have increased as a result of selection, they must be contributing to the overall increase in musical appeal (M) (Fig. 1). However, music has many dimensions, and we only measured two. We used single and multiple linear regression analysis to estimate how much of the overall increase in M is attributable to the features we measured. We find that C_L alone is responsible for 3.0% of the increase in M and R alone is responsible for 2.8%, whereas, together, they account for 4.2%, leaving 95.8% unexplained (SI Appendix, B.7); thus, other features must also contribute to the evolution of appealing music in these populations. In the future, we will be able to examine these with an expanded MIR toolkit.

Discussion and Conclusions

We have shown that a simple Darwinian process can produce music. In recent years, inspired by cultural transmission theory

(43, 44), the evolution of material artifacts, such as stone tools, has been analyzed in terms of selection-variation processes (45–47). Our results suggest that the evolution of music can be viewed in the same way. Although our system is an artificial one, it may shed light on the evolution of real musical cultures. In our system, musical appeal increased rapidly but then stopped, and this stasis is mostly attributable to a decrease in the fidelity of transmission. Analogously, musical styles in premodern societies appear to be very conservative; some may be thousands of years old (3, 6, 48). Given that many such societies lacked symbolic or mechanical means of transmitting their music, we speculate that the cause of stasis is the same: low transmission fidelity that, in such musical cultures, arises from the errors introduced as musicians teach and learn complex musical themes. Our results may also explain the dynamics of many adaptive systems. Populations of genetic algorithms, digital organisms, and real organisms often converge to an evolutionary equilibrium (10, 11, 30–32). Such equilibria are typically interpreted as a sign that the population has arrived at an adaptive optimum or else has exhausted its selectable variation. A decrease in transmission fidelity, of the sort seen here, is another explanation for evolutionary stasis, one that may be particularly important as more complex, or at least more finely adapted, structures evolve.

Our experiment demonstrates the creative role of consumer selection in shaping the music we listen to. However, the evolution of music in human societies is certainly shaped by other forces as well. Humans do compose music before releasing it for public consumption (6), and consumers do not choose the music they like entirely on the basis on aesthetic quality but are also influenced by the preferences of others (7). Thus, musical evolution is the result of selection at multiple levels: within individuals (producers), among individuals (consumers), and among groups (social networks of consumers). The theoretical framework we have used here can be extended to accommodate these forces because the Price equation is particularly suited to partitioning the effects of selection at different levels (35, 37, 49). The DarwinTunes system can, similarly, be extended to accommodate these additional selective forces by allowing individual consumers to select among variants (i.e., compose) before releasing them into the population or by allowing consumers to see each other's preferences. The relative importance of selection at these different levels—producer, consumer, and consumer-group—in shaping the evolution of the world's music is unknown and may vary among societies. Western societies have long had specialist guilds of composers and performers; however, in other cultures, participation is more widespread [e.g., early 20th century Andaman Islanders (50)]. The ability to download, manipulate, and distribute music via social-networking sites has democratized the production of music and may change the balance of these forces again. In partitioning these selective forces, our analysis points the way to the future evolutionary dynamics of digital culture (51).

Materials and Methods

The architecture of the DarwinTunes evolutionary music engine, the structure of the selection experiments, the experimental populations, and the control populations are found in SI Appendix, A. These also contain the details of the rerating experiment, the audio feature extraction algorithms, and their validation. Our methods of fitting functions to evolutionary trajectories are described in SI Appendix, B.

ACKNOWLEDGMENTS. We thank the participants in this study, particularly the undergraduates at Imperial College; B. North for statistical advice; Ricardo Azevedo, Steve Frank, William Rice, Jennifer Jacquet, Katy Nolan, Dan Stowel, Frans Wiering, and Geraint Wiggins for comments on the manuscript; the Imperial College London open access publication fund; and Imperial College London and the School of Computer Science and Electronic Engineering, Queen Mary University of London, for the use of high-performance computing facilities.

1. Nettl B (2005) *The Study of Ethnomusicology: Thirty-One Issues and Concepts* (Univ of Illinois Press, Urbana, IL).
2. Malm WP (1977) *Music Cultures of the Pacific, the Near East, and Asia. Prentice Hall History of Music Series* (Prentice Hall, Englewood Cliffs, NJ).
3. Leroi AM, Swire J (2006) The recovery of the past. *World of Music* 48(3):43–54.
4. Scruton R (1997) *The Aesthetics of Music* (Clarendon, Oxford).
5. Mathiesen TJ (1999) *Apollo's Lyre: Greek Music and Music Theory in Antiquity and the Middle Ages. Publications of the Center for the History of Music Theory and Literature* (Univ of Nebraska Press, Lincoln, NE).
6. Merriam AP (1964) *The Anthropology of Music* (Northwestern Univ Press, Evanston, IL).
7. Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311:854–856.
8. Campbell DT (1974) Evolutionary epistemology. *The Philosophy of Karl Popper* (Open Court Press, LaSalle, IL), Vol 1.
9. Lenski RE, Travisano M (1994) Dynamics of adaptation and diversification: A 10,000-generation experiment with bacterial populations. *Proc Natl Acad Sci USA* 91:6808–6814.
10. Elena SF, Lenski RE (2003) Evolution experiments with microorganisms: The dynamics and genetic bases of adaptation. *Nat Rev Genet* 4:457–469.
11. Holland J (1992) Genetic algorithms. *Sci Am*, 267, pp 66–72.
12. Wilke CO, Wang JL, Ofria C, Lenski RE, Adami C (2001) Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* 412:331–333.
13. Lenski RE, Ofria C, Pennock RT, Adami C (2003) The evolutionary origin of complex features. *Nature* 423:139–144.
14. Ofria C, Wilke CO (2004) Avida: A software platform for research in computational evolutionary biology. *Artif Life* 10:191–229.
15. Adami C (2006) Digital genetics: Unravelling the genetic basis of evolution. *Nat Rev Genet* 7:109–118.
16. Wischmann S, Floreano D, Keller L (2012) Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots. *Proc Natl Acad Sci USA* 109:864–868.
17. Sims K (1991) Artificial evolution for computer graphics. *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 1991)*, ed Thomas JJ (Association for Computing Machinery, New York), pp 319–328. Available at <http://dl.acm.org/citation.cfm?doi=122718.122752>.
18. Miranda ER, Biles JA, eds (2007) *Evolutionary Computer Music* (Springer, London).
19. Romero J, Machado P, eds (2007) *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music. Natural Computing Series* (Springer, Heidelberg).
20. Wolkowicz J, Heywood M, Keselj V (2009) Evolving indirectly represented melodies with corpus-based fitness evaluation. *Proceedings of the EvoWorkshops 2009 on Applications of Evolutionary Computing: EvoCOMNET, EvoENVIRONMENT, EvoFIN, EvoGAMES, EvoHOT, EvoIASP, EvoINTERACTION, EvoMUSART, EvoNUM, EvoSTOC, EvoTRANSLLOG (EvoWorkshops 2009)*, eds Giacobini M, et al. (Springer, Heidelberg), pp 603–608. Available at <http://dl.acm.org/citation.cfm?id=1533647>.
21. Marques VM, Reis C, Machado JA (2010) Interactive evolutionary computation in music. *2010 IEEE International Conference on Systems Man and Cybernetics (SMC) (IEEE)*, pp 3501–3507. Available at http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5642417&tag=1.
22. Biles J (1994) Genjam: A genetic algorithm for generating jazz solos. *Proceedings of the 1994 International Computer Music Conference*, pp 131–137. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.55.6146>.
23. Patel AD (2008) *Music, Language, and the Brain* (Oxford Univ Press, New York).
24. Casey M, et al. (2008) Content-based music information retrieval: Current directions and future challenges. *Proc IEEE* 96:668–696.
25. Dixon S, Mauch M, Anglade A (2011) Probabilistic and logic-based modelling of harmony. *CMRR'10: Proceedings of the Seventh international Conference on Exploring Music Contents*, eds Ystad S, Aramaki M, Kronland-Martinet R, Jensen K (Springer, Berlin, Heidelberg), pp 1–19. Available at <http://dl.acm.org/citation.cfm?id=2040791>.
26. Davies MEP, Plumbley MD, Eck D (2009) Towards a musical beat emphasis function. *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA 2009)*. Available at <http://dx.doi.org/10.1109/ASPAA.2009.5346462>.
27. Mauch M, Dixon S (2010) Simultaneous estimation of chords and musical context from audio. *IEEE Transactions on Audio Speech and Language Processing* 18:1280–1289.
28. Lidy T, Rauber A (2005) Evaluation of feature extractors and psycho-acoustic transformations for music genre classification. *Proceedings of the Sixth International Conference on Music Information Retrieval (ISMIR 2005)*, pp 34–41. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.59.9623>.
29. Rose MR, et al. (2005) The effects of evolution are local: Evidence from experimental evolution in drosophila. *Integr Comp Biol* 45:486–491.
30. Eldredge N, et al. (2005) The dynamics of evolutionary stasis. *Paleobiology* 31:133–145.
31. Blows MW, Hoffmann AA (2005) A reassessment of genetic limits to evolutionary change. *Ecology* 86:1371–1384.
32. Ostrowski EA, Ofria C, Lenski RE (2007) Ecological specialization and adaptive decay in digital organisms. *Am Nat* 169:1–20.
33. Schluter D, Nychka D (1994) Exploring fitness surfaces. *Am Nat* 143:597–616.
34. Price GR (1970) Selection and covariance. *Nature* 227:520–521.
35. Price GR (1972) Extension of covariance selection mathematics. *Ann Hum Genet* 35:485–490.
36. Frank SA (1995) George Price's contributions to evolutionary genetics. *J Theor Biol* 175:373–388.
37. Frank SA (1997) The Price equation, Fisher's fundamental theorem, kin selection, and causal analysis. *Evolution* 51:1712–1729.
38. Burt A (1995) The evolution of fitness. *Evolution* 49:1–8.
39. Barton N, Partridge L (2000) Limits to natural selection. *Bioessays* 22:1075–1084.
40. Coyne J, Orr H (2004) *Speciation* (Sinauer, Sunderland, MA).
41. Fisher RA (1930) *The Genetical Theory of Natural Selection* (Clarendon, Oxford).
42. Lenski RE, Ofria C, Collier TC, Adami C (1999) Genome complexity, robustness and genetic interactions in digital organisms. *Nature* 400:661–664.
43. Cavalli-Sforza LL, Feldman MW (1981) *Cultural Transmission and Evolution: A Quantitative Approach. Monographs in Population Biology* (Princeton Univ Press, Princeton).
44. Boyd R, Richerson PJ (1985) *Culture and the Evolutionary Process* (Univ of Chicago Press, Chicago).
45. O'Brien MJ, Lyman RL (2000) *Applying Evolutionary Archaeology: A Systematic Approach* (Kluwer Academic, Plenum, New York).
46. Shennan S (2002) *Genes, Memes, and Human History: Darwinian Archaeology and Cultural Evolution* (Thames and Hudson, London).
47. Mesoudi A, Whiten A, Laland KN (2004) Perspective: Is human cultural evolution Darwinian? Evidence reviewed from the perspective of the Origin of Species. *Evolution* 58:1–11.
48. Lomax A, Berkowitz N (1972) The evolutionary taxonomy of culture. *Science* 177:228–239.
49. Frank SA (1998) *Foundations of Social Evolution. Monographs in Behavior and Ecology* (Princeton Univ Press, Princeton).
50. Raddiffe-Brown AR (1933) *The Andaman Islanders* (Cambridge Univ Press, Cambridge, UK).
51. Michel JB, et al.; Google Books Team (2011) Quantitative analysis of culture using millions of digitized books. *Science* 331:176–182.