

Too many love songs: Sexual selection and the evolution of communication

Gregory M. Werner and Peter M. Todd

Center for Adaptive Behavior and Cognition
Max Planck Institute for Psychological Research
Leopoldstrasse 24, 80802 Munich Germany
gwerner@, ptodd@mpipf-muenchen.mpg.de

Abstract

Communication signals in many animal species (including humans) show a surprising amount of variety both across time and at any one instant in a population. Traditional accounts and simulation models of the evolution of communication offer little explanation of this diversity. Sexual selection of signals used to attract mates, and the coevolving preferences used to judge those signals, can instead provide a convincing mechanism. Here we demonstrate that a wide variety of “songs” can evolve when male organisms sing their songs to females who judge each male’s output and decide whether or not to mate with him based on their own coevolved aesthetics. Evolved variety and rate of innovation are greatest when females combine inherited song preferences with a desire to be surprised. If females choose mates from a small pool of candidates, diversity and rate of change are also increased. Such diversity of communication signals may have implications for the evolution of brains as well.

1 Introduction

Why are there so many love songs on the radio? To the sometimes slight extent that they do, why do these songs change from year to year? Why do birds bother to sing so many notes? What would be the disadvantage of just a single long, loud blast? In short, why is there so much diversity in communication signals, both within and between generations?

In species with highly evolved, elaborate communication systems, there is often a great diversity of signals used within a given population, and between populations (including successive generations and recently-diverged species) over time. Humans of course have an unmatched capacity to generate novel signals (Pinker, 1994). Many songbirds have repertoires of dozens of distinct song types, a few species can sing hundreds of different songs, and the brown thrasher checks in with a remarkable

repertoire size of over 2000 (Catchpole & Slater, 1995). Moreover, any one male of a given songbird species will typically sing a different repertoire from other conspecific males. Moving from air to ocean, cephalopods (particularly cuttlefishes, octopuses, and squids) also use a surprising variety of signal types, with some species using as many as 35 different displays in a wide range of combinations and sequences (Hanlon & Messenger, 1996). In all of these cases, the reason for extensive signal diversity remains a mystery: Catchpole and Slater (1995, p. 187) say that “At first sight, the diversity of modes of singing amongst birds is so great that it defies explanation,” while Hanlon and Messenger (1996, p. 131), feeling that birds are easier to understand than their favorite animals, wonder, “does the great variety of signals [in cephalopods] serve as a measure of the signaller’s fitness, as in bird song?”

Traditional reasons given for the evolution of communication cannot provide the whole answer to the questions surrounding signal diversity. If communication is viewed as a means of transferring veridical information from one organism to another (see Hauser, 1996), we would expect repeated communications of the same information (by one individual or within a population) to be performed in a similar manner to avoid misinterpretation by the receiver. In the particular case of accurate species identification for mating purposes, there should also be little variation between signals of conspecifics. If communication is seen instead as a way to manipulate the behavior of another organism (which can include non-veridical deceit—see Dawkins & Krebs, 1978), the signal used in any particular case should be the single one found to be most effective. And if communication is considered a means of altruistically benefiting one’s genetic relatives (Ackley & Littman, 1994), we would expect convergence onto stable (but possibly family-specific) ways to help one another.

What then can drive the evolution of a large variety of elaborate communication signals? In this paper, we explore a particularly powerful force that can engender such diversity: sexual selection acting via coevolving

mate preferences and traits. Specifically, we develop a simulation model that demonstrates how, when communicative signals are used by males to attract females as mates, sexual selection can drive the evolution of a variety of male songs and female song preferences. However, this evolution is likely to stagnate unless the females choose songs based not just on their evolved preferences, but also on a desire to be surprised by what they hear. Loosely speaking, when females can be bored by the same old song, males must strive to provide the females with something new in order to assure their own mating success. As a consequence, a variety of male songs evolves, both within a single generation, and across successive generations over time. (To explain the rapid cultural cycles of human love songs, we must resort to learning—see section 4—but here we show that evolution alone can generate these other major sources of signal diversity.)

In the next section, we consider past approaches to modeling the evolution of communication, and show how our current perspective can more readily explain the appearance and maintenance of signal diversity. Our simulation method here is an extension of our previous work on sexual selection, from a simple two-dimensional phenotype space to a multidimensional behavioral trait (a signal). In section 3, we describe how this method of modeling sexual selection is applied to simple songs and preferences, and show the results in terms of evolved song variation. We conclude with a consideration of the implications of sexual selection for the evolution of communication, and indicate the further directions in which this research itself can evolve.

2 Past approaches to modeling the evolution of communication

2.1 *Communication for mutual benefit or manipulation*

As indicated in the previous section, many functions have been proposed for evolved communication systems, and the most prevalent of these have been modeled recently using the techniques of individual-based evolutionary simulations. While none of the earlier models aimed specifically at exploring the mystery of signal diversity, their results do provide hints towards the explanation that we develop here, as we will show in this section. (Non-individual-based models, such as theoretical population genetics models, do not give us much insight into the evolution of diversity, because of the strong mathematically-required restrictions they place on the possible evolved signals.)

In most earlier simulations of the evolution of communication, the function of communicating has been taken to be the dissemination of information that will benefit the survival prospects of either the sender or the sender's

relatives. One of the first such studies was MacLennan's (1990, 1992) series of synthetic ethology experiments, which investigated the ways in which meaningful signals could arise in a breeding population of "simorgs." Meaningful signals in this case are those that tell an organism how to behave in response to an unknown aspect of the environment, so as to increase its chances of survival.

To allow such informative signals to evolve, MacLennan divides up his artificial world into local environments and restricts direct knowledge of each local environment to a single simorg, thereby "permit[ting] some simorgs to 'see' things that others cannot; otherwise there would be no advantage in communicating" (MacLennan, 1992, p. 639). Both signaller and receiver benefit in terms of a fitness gain if they successfully transmit information about a local environment. This is in marked contrast to the use of signals as a way of attracting other organisms (as potential mates), rather than informing them—in our simulation, there is nothing to "see" but the signal itself.

MacLennan is interested in the grounding of meaningful symbols through their attachment to states of the world. But it is this very process of grounding that ties the signals down and prevents them from evolving into more various and elaborate forms over time. By freeing signals from any concrete reference, other than the fact that a particular organism is able to produce that signal, sexual selection allows them to change continuously and fairly rapidly over time. (MacLennan's early simulations used only 8 world states and correspondingly only 8 possible signals, hardly allowing the evolution of much variety; but the number of signals used would always be tied down by the number of world states in his environment. When learning was added, individual signals often took on more than one meaning, indicating greater communication variety within the population, but even this diversity was static over time.)

Ackley and Littman (1994) explored altruism as a possible function of communication. In their simulation, local populations of organisms could evolve to signal one another about the common features of their local environment, again under the assumption that not every organism could see every feature. In contrast to MacLennan's world, the signallers in Ackley and Littman's model did not receive any fitness benefit from their selfless act of shouting—only the receivers would benefit. But since the receivers, all living nearby, were also all likely to be closely genetically related to the signaller, and since local populations competed with their neighboring populations through the occasional exchange of offspring, shouting out the right signals could benefit the signaller's genetic representation through kin selection.

Again, the pressure in this world to be informative all but eliminated any opportunity to be interesting or novel—the evolved signals were short and simple, and largely homogeneous within any given local environment.

They varied to a certain degree between localities, of course, at least early in any simulation run; but as more and more successful signalling strategies evolved at the local level, they could spread to take over the entire global world as well through the low-level migration process. Thus early diversity would evolve to stable homogeneity in most cases. Some variety could still emerge over time, as parasites periodically evolved to take advantage of current communication patterns and quickly spread across the world, only to be replaced at a later date by a new strain of altruistic communicators resistant to that breed of parasitism.

This kind of coevolutionary change between parasites and hosts resembles that which occurs in our sexual selection simulations between males and females, but at a much slower time scale, and without engendering the same kind of within-generation signal diversity. Note that in Nature, parasite-host coevolution may occur faster than male-female coevolution because of the rapid generation turnover in parasites, and indeed parasite-host coevolution may be the reason for the phenotypic variety created by sexual recombination. While this form of coevolution can possibly affect the evolution of signals that indicate parasite load (Hamilton & Zuk, 1982), it seems less likely that parasite selection pressure could directly foster the signal diversity within populations or across generations that is our main interest here.

Much closer to our current concerns is (perhaps not surprisingly) an earlier simulation created by the first author: Werner and Dyer’s (1992) evolution of signals generated by females to guide blind males towards them for mating. The communication in this case functioned to allow reproduction between the roving males and the stationary, signalling females they succeeded in finding; thus we can say that this communication evolved via sexual selection rather than natural selection (as in MacLennan’s case, or kin selection in Ackley & Littman’s system). We could even say that the females’ signals are evolving to be attractive to the males, and that the males are deciding which females are singing the most attractive songs (by “following” the songs they hear and mating with whichever one is most effective in literally attracting them).

As a consequence of this form of sexual selection, we can see the beginnings of signal diversity in this simulation: signal “dialects” appeared in some runs, leading Werner and Dyer to speculate that these “communication protocols could provide a natural way of establishing genetic barriers that spontaneously emerge” (Werner & Dyer, 1992, p. 685) and lead to distinct coexisting species. But these spontaneous breeding barriers emerge only slowly in the rather diffuse sexual selection operating in this model; in the simulations we report here, we greatly increase the power of sexual selection by allowing individuals to sample several potential mates, rather

than the few that a male might stumble across in Werner and Dyer’s original setup. And this early simulation, like MacLennan’s, does not generate diversity across time: once the population settles on a particular signal protocol, it wavers little from that solution.

2.2 *Communication for mate attraction*

To generate signal diversity both across time and at any given instant—that is, both diachronically and synchronically—we must somehow combine the power of Werner and Dyer’s simulation to generate a variety of signals within one generation (albeit in a limited fashion) with the ability of Ackley and Littman’s model to engender signal change from one generation to the next (albeit at a slow pace). Sexual selection through mate choice allows the former, leading a population to adopt a variety of sub-species signalling protocols (see Todd & Miller, 1991, for a simulation model of this effect for simple phenotypic traits). We need some force to push a population out of its attained stable pattern of speciation, though, the role that parasites played in Ackley and Littman’s world. In sexual selection, this can be achieved through directional mate preferences (Kirkpatrick, 1987; Miller & Todd, 1993, 1995), which for example cause females always to look for brighter, or more colorful, or more ornamented males, and thereby push a population to continue evolving. For the evolution of communication, as we will see, this constant striving force can be effected through neophilia: females always looking for signals that are novel and unexpected.

Sexual selection has been implicated in the evolution of communication signals, particularly birdsong, ever since Darwin’s (1871) introduction of the concept and his original proposal for the role of female choice in the evolution of elaborate male song (see Catchpole & Slater, 1995, chapter 7). In ethology, much research has been devoted in particular to the diversity-related issue of male song repertoire size in songbirds, seeking to identify the function of song variety at the individual level (as opposed to the population or multi-generation level that we address here; see Catchpole & Slater, 1995, chapter 8). The idea that female mate choice based on preferences for novelty in male song could lead to larger male song repertoires can be traced back to Hartshorn (1973) (following Darwin’s lead a century earlier). It has more recently been argued by Searcy (1992; see also Hauser, 1996) based on the notion of dishabituation—that is, females exposed to the same song repeatedly will habituate to it (become bored) and respond less, but if a male can sing her different songs, this will cause dishabituation and increase her likelihood to respond to his overtures. This learning-based hypothesis remains contentious (Catchpole & Slater, 1995, pp. 179-182); here we avoid the question of the effect of learning on individual male repertoires, and instead focus on the corre-

sponding mystery of the evolution of differences *between* individual males' song output.

This "good taste" hypothesis, that males who can sing many songs are simply more attractive to the tastes of females, stands in contrast to the "good genes" models of the evolution of large song repertoires. These latter theories propose instead that the number of songs a male sings is an indication of some underlying aspect of his genetic quality, for instance his strength, or longevity, or parasite resistance (Hamilton & Zuk, 1982), or food-finding ability (because singing takes a lot of energy; see Hauser, 1996, and Andersson, 1994, for discussion of these and other possible signs of quality). But the overall impression from a number of studies is that support for these "good genes" models for the evolution of song variety is equivocal at best (Hauser, 1996; Andersson, 1994; Catchpole & Slater, 1995).

Our simulations reported here support the "good taste" explanation for the diversity of song, extending the unit of analysis from the single singing male to a whole population evolving over time. (While we restrict each male to a single song at present, we could modify the simulation to see how the variety in each individual song repertoire can evolve as well.) The males in our model have nothing to offer (and nothing to be judged upon) other than the pretty song they sing, and females choose them solely because of that song and the hope that their own (male) offspring will sing as well and attract more mates as a consequence (the "sexy son" effect, which can lead to runaway sexual selection for particular traits and the preferences for those traits—see Weatherhead & Robertson, 1979; Andersson, 1994).

3 Coevolving hopeful singers and music critics

In our first attempt to evolve communication signal diversity, we coevolved artificial neural network "males" who produced rhythmic "songs" along with picky neural network "females" who judged those songs and used them to decide whom to mate with. In females, the neural network mapped inputs from an "ear" to output units that indicated her decision to mate or abstain. In males, the circuitry produced a sequence of sounds in response to the presence of a female. When this model was run, the male neural networks produced complex output patterns, typically containing several concurrent, not-quite-repeating patterns. These songs changed dramatically over evolutionary time, driven by the preferences of the female networks—but the patterns proved to be very difficult to analyze for complexity, diversity, or change. It was clear that the songs were evolving, but not clear how. It was time to call in more rudimentary musicians.

Each of the "dumbed-down" males we next turned to has genes that directly encode the notes of his song (rather than a song-generating network). Each male song

(and hence genotype) consists of 32 notes, each of which can be a single pitch selected from a two-octave (24 pitch) range. Females' genes now encode a transition matrix which is used to rate transitions from one note to another in male songs. This matrix is an N -by- N table, where N is the number of possible pitches the males can produce (24 in these experiments). Each entry in this table represents the female's expectation of the probability of one pitch following another in a song. For instance, entry {4, 11} (or C-G in our two-octave case) in a particular female's table captures how often she thinks pitch 11 will follow pitch 4, on average, in male songs. Given these expectations, females can decide how well they like a particular song in different ways, as we will see in the next subsection. Whatever method she uses, as she listens to a male, the female considers the transition from the previous note's pitch to the current note's pitch for each note in a song, gives each transition a score based on her transition table, and sums those scores to come up with her final evaluation of the male and his serenade.

Each female listens to the songs of a certain number of males who are randomly selected to be in her "courting choir." All females hear the same number of males, and the size of the courting choir—that is, a female's sample size—is specified for each evolutionary run. After listening to all the males in her potential-mate choir, the female selects the one that she most preferred (i.e. the one with the highest score) as her mate. This female choice process ensures that all females will have exactly one mate, but males can have a range of mates from 0 (if his song is unpopular with everyone) to something close to the courting choir size (if he has a platinum hit that is selected by all the females who listen to him). Each female has one child per generation created via crossover and mutation with her chosen male mate. This temporarily puts the population at about 50% above a specified "carrying capacity" (target population size). We then kill off approximately a third of the individuals, bringing the population back to a predetermined carrying capacity. This whole process is repeated for some desired number of generations.

3.1 Different ways females can rate their mates

We employed three different methods for scoring the male songs using these tables. In the first method, the female simply scores each transition as it occurs in the song by immediately looking up how much she expected that particular transition and adding it to the running total score for the song. Thus, those songs that contain more of the individual transitions that the female expects (for example, songs with many C-G transitions, if she expects C's to be followed by G's very often) will be scored higher by her, and she will prefer to mate with the males who sing these songs. We call this the *local transition preference* scoring method.

In the second method, the female listens to a whole song first, counting the number of each type of transition that occurs in the song (for example, she might tally up G's following C's four times in the song, and other notes following C's two times). Then from these counts she constructs a transition matrix for that particular individual song (for example, with an entry of .66 for the C-G transition, because that is what occurred two-thirds of the time after a C in this song). Finally, she compares that song's transition table with her expected (preferred) transition table, and the closer the two tables match (on an entry-by-entry basis), the higher score and preference she gives to that song.

Thus this method means that a female will prefer songs that match the overall statistical pattern of transitions in her transition table. We call this the *global transition preference* scoring method. Continuing with our example, if the female has a value of .75 stored in her own transition table for the C-G transition, she will like songs most that have a C-G transition exactly three-fourths of the time (along with other C-x transitions, where x is any note other than G, for the other quarter of the time that C appears). In contrast, with local transition scoring, she would prefer C-G transitions after every C, because they give a higher local score than any other transition from C.

The third scoring method produced females that enjoy being surprised. The female listens to each transition in the song individually as in the first method, looks up how much she expected that transition, and subtracts this probability value from the probability she attached to the transition she most expected to hear. Consider our female from the previous paragraph again. Whenever she hears a C in a male's song, she most expects a G to follow it (75% of the time). Imagine she instead hears a C-E transition in a song. This transition is a surprise to her, because it violates the C-G transition expectation—and so she likes this song more as a consequence.

But how much of a surprise was this note, and how much does it increase her preference for this song? To find out, the female critic first looks up the C-E transition in her table, and finds she expected that transition 15% of the time. Thus, this C-E transition was not a complete surprise, since she had some previous expectation for it, but it was a reasonably large one. We quantify the surprise level with a score of $.75 - .15 = .6$ for that transition (that is, $\text{prob}(C-G) - \text{prob}(C-E)$). This expected-minus-actual-transition-probability score is summed up for all the transitions in the current song, and the final sum registers how much surprise the female experienced, and therefore how much she preferred that song. Not surprisingly, we call this the *surprise preference* scoring method. Note that it will not result in the males singing random songs—in order to get a high surprise score, a song must first build up expectations, by making transi-

tions to notes that have highly expected notes following them, and then violate those expectations, by not using the highly expected note. Thus there is a constant tug-of-war between doing what is expected and what is unexpected in each song.

The first two preference scoring methods can be considered forms of non-directional mate preferences: evolved male songs that match evolved female expectations most precisely (either locally or globally) will receive the most mating interest. The third surprise preference scoring method, however, is a type of directional mate preference. Rather than rewarding male songs that match female expectations, surprising songs that are some ways off from the evolved female transition tables in song space will be sought after. Thus we expected to see less movement through song space for the local and global transition preferences (though possibly more speciation), and more constant change when surprise preferences were used.

We also expected that surprise scoring would create greater diversity within any given generation than would preferences based on matching local or global expectations, because there are more ways to violate expectations (causing surprise) than to meet them. Note that this is different from the kinds of directional preferences we have previously considered (Miller & Todd, 1993), where only a single preferred direction was indicated (e.g. a greener vs. a bluer patch of plumage). In those cases, the population could evolve to all head in one direction in phenotype space; here, the population will be more likely to scatter in many directions in phenotype space. (This is similar to apostatic selection for multiple phenotypes caused by pressure to evade easy detection predators—see Driver & Humphries, 1988.) But this is not an unreasonable effect to expect, because it can be a direct outcome of low-level dishabituation or novelty-seeking processes.

We varied two additional parameters beyond the female preference scoring method, to test their effects on the evolved diversity of songs. First, we controlled the number of males a female listens to before selecting a mate—that is, the size of her courting choir (2 or 20). This parameter is essentially a “volume knob” on the overall impact of sexual selection in the simulation—if females can only sample one male, then there is no sexual selection taking place, while the greater number of males she can listen to before choosing a mate, the stronger will be the selective force of her preferences. We expected that smaller sample sizes would lead to greater diversity than larger sample sizes, but that larger sample sizes might support a number of distinct “species” of songs. Smaller samples should, on average, give males a better chance of reproducing even if their song is not close to what is desired by females, because each male in a female's small sample set faces less competition than if she sampled a large number of males. On the other hand,

we believed that large sample sizes would quickly draw males close to what was desired by females, but that the preferences could aggregate in distinct clusters. So while smaller samples could lead to a larger spread around a central average preference, larger samples could lead to more than one tight cluster (i.e., distinct species). Finally, the greater selection pressure on males caused by larger sample sizes should mean that there will be more rapid change in male songs across generations as well.

The last parameter we considered was whether female expectation transition tables were fixed across time (that is, female offspring contain exact copies of their mother’s transition table) or allowed to coevolve with the male songs. We expected that coevolving preferences would allow more change (or diversity) in songs over time because the targets for the males would themselves be moving. In a system without coevolution, male songs will tend to converge on the female preferences and stay there, providing little evolutionary movement. Sexual selection via coevolving traits and preferences is the whole premise of our argument in this paper, so this comparison is a crucial one.

3.2 Resulting song change over time

We ran populations of 1000 individuals for 1000 generations in 12 different conditions: all combinations of 3 preference scoring methods, 2 sample sizes, and fixed vs. coevolving preferences. In each case, we initiated the males (i.e., their songs) randomly, and the first generation of female transition tables were set with probabilities calculated from a collection of simple folk-tune melodies. This way we could ensure that female preferences in our simulations at least started out with some resemblance to human melodic preferences; however, once evolution started moving the preferences and songs around, any hope of the population’s aesthetics matching human aesthetics would quickly be lost. Thus, we could not listen to the system and readily judge its progress; we had to resort to more objective measures (further motivating the simplified form of song and preference representation described in the previous section).

To measure evolving song change over time—diachronic diversity—we use a “progress chart” technique modified from Cliff and Miller’s (1995) work on measuring coevolutionary progress in pursuit-evasion games. This “ancestral distance map” (Cliff & Miller use Hamming distance) plots generations G in time from left to right (from generation $G=0$ to $G=1000$), and generations G' backwards in time (relative to each generation G) from top to bottom (from generation $G'=G-1$ to generation $G'=G-999$). At each point (G, G') in the triangular region so formed, we plot the difference between the modal male song (i.e. the most common note at each of the 32 positions) at generation G and that at generation G' , with difference measured as the number of positions where the




Figure 1: Change in modal song from current generation G (left to right) to all previous generations G' (from $G-1$ at top to $G-999$ at bottom). Here a coevolving surprise-preference sample-size-2 population shows continuous rapid change over time.

two songs differ. This difference score, from 0 to 32, is indicated by the darkness of the plotted point, with greater differences mapping onto lighter points. (See Figure 1.)

The top row of points in this type of plot shows the difference between the modal song at any particular generation and the modal song of the previous generation, while the rightmost column of points shows the difference between the modal song at generation $G=1000$ and the modal songs of all generations before that (from $G'=999$ at the top, to $G'=1$ at the bottom). The faster the population evolves and the modal song changes over time, the more of this plot will be filled with points registering the maximum difference of 32—that is, the bigger the light-colored regions will be.

Using this technique, we compared the rate of change of population modal songs over time for our 13 different conditions. Our results mostly matched our expectations, but there was a surprise: listening to only 2 males yielded much faster evolutionary change than choosing from 20 males. This was the largest effect on rate of change, and goes against the selection pressure argument we put forth in the previous section. Instead, this effect could occur because with bigger sample sizes, traits could match preferences much more closely, and so little movement of either would be necessitated over time; in addition, because both parents are more closely matched with a bigger sample size, their offspring will also resem-




Figure 2: Change in modal song for a non-coevolving global-preference sample-size-20 population, showing little change over time.

ble them more closely, further slowing down change from one generation to the next.

Surprise scoring yielded greater change than either global or local transition scoring. Local scoring, in fact, made the population converge rather rapidly to the locally-preferred song transitions, so that male songs often degenerated to repetition of a single note or alternation between two notes. (This also gave these runs very low within-generation synchronic diversity scores, so we did not analyze this type of preference further.) Finally, coevolution led to faster change than fixed female preferences, at least when surprise scoring was used—the situation is less clear with global transition scoring, which we are investigating further.

We can easily visualize the difference between rate of change in the fastest case and its parametric “opposite” (i.e. changing all the parameters), which is one of the slowest cases, by plotting their progress charts. In Figure 1, we show modal song change for a coevolving surprise-scoring small-sample (i.e. sample-size 2) population. The relatively small region of dark points, indicating small changes between past and present generations, is dominated by a large light region, indicating large changes over time. (Remember that the maximum possible difference between any two modal melodies is 32, so the nearly-white region indicates that the something close to the maximum number of changed positions has been reached—more change is still occurring in that region, but our distance measure has hit ceiling and does

not reflect the further movement.) In Figure 2, we plot the chart for a fixed-preference global-transition-scoring large-sample (20) population. Here the differences between present and past modal songs are mostly small (dark points), meaning that little change has occurred over time. The light band along the diagonal indicates that there was a lot of change in the first few generations, as the initially random male songs were most strongly winnowed down, but after that little more transpired.

3.3 Resulting song diversity within populations

To measure the synchronic diversity of songs within a population at any particular generation, we computed the set of differences (again 0-32) between every pair of males’ songs in the population. This set of differences could be plotted as a histogram for any given generation, with highly converged, low-diversity populations having histograms skewed toward low values, and unconverged, high-diversity populations having histograms skewed towards high values. Furthermore, populations with two or more distinct “species” of songs will show up as multiple peaks in the histogram (representing the distributions of between-species and within-species distances). To explore how this within-generation diversity changes across generations, we change each histogram into a one-dimensional density plot by essentially viewing it “from overhead” and representing high regions in the histogram with dark points in the plot. Then we can line up these one-dimensional density plots next to each other, generation by generation, to make a two-dimensional plot of the changing synchronic diversity in the population over time (see Figure 3). Now we have a plot with generation G along one dimension, and distance between each male song and the modal song along the other dimension, with the darkness of each point indicating the number of males who are that different from the population’s current modal song.

We used this visualization method to compare the evolving synchronic diversity of songs in populations in 8 conditions (leaving out the degenerate hyper-converged local transition score populations). Again our expectations were mostly met, and again the largest effect came from the size of the female’s sample set. Sampling 2 males preserved diversity in the population to a much greater degree than sampling 20 males; in the former case, most males retained 10-20 different notes from the modal song after 1000 generations, while in the latter, most males had only one or two notes different. Coevolution yielded greater synchronic diversity than fixed female preferences, but to a lesser degree than sample size (e.g. about 18 notes different from the modal song for the coevolving surprise sample-2 population versus about 11 notes different for the fixed surprise sample-2 population after 1000 generations). The preference scoring method (surprise versus global transition scoring) showed little




Figure 3: Diversity of songs in each generation G , from $G=1$ at top to $G=1000$ at bottom. Each point shows the number of pairs of songs that have a certain number of notes different between them. Here, diversity is preserved in a coevolving surprise-preference sample-size-2 population.

consistent effect on within-generation diversity, however.

We show the difference between the case with the greatest synchronic diversity and its parametric opposite with one of the lowest diversities in Figures 3 and 4 respectively. Figure 3 displays the song diversity in a coevolving surprise-scoring sample-2 population over time, starting at generation 0 at the top of the graph and proceeding to generation 1000 at the bottom. Diversity starts out maximal in the early generations (when the random initial male songs were all very far from the modal song), and declines somewhat over time. But even after 1000 generations, most male songs have about 20 notes out of 32 that are different from their population’s modal song. In contrast, the fixed global-transition-scoring sample-20 population in Figure 4 converges from its initial diversity to population-wide homogeneity very rapidly. Within 150 generations, most males sing songs that are only slight (3-position) variations on the population modal song, and this clustering even gets slightly tighter over time. But this tight clustering from the large sample-size, when combined with the directional-selection effects of surprise preferences, can lead new song “species” to emerge and differentiate from each other over time. Figure 5 shows this effect, indicating that diversity across the whole population can be replaced by diversity between subpopulations.




Figure 4: Diversity of songs in a non-coevolving global-preference sample-size-20 population, showing loss of diversity over time.

4 Implications and further work

What does all this mean? Without sexual selection, as we saw in section 2, simulation models have evolved little diversity in communication signals. When instead we replace natural selection with sexual selection, signal diversity within and across generations blossoms. Our simulations here lend strong support for the role of coevolving songs and directional (surprise-based) preferences in maintaining diversity over time (Figure 3), and in continuously altering that diversity as time goes by (Figure 1). With non-coevolving, non-directional preferences, progress is slower (Figure 2) and diversity collapses (Figure 4). The number of mates sampled is the selective-force amplifier for these effects: small sample sizes promote diversity and change, while large sample sizes encourage fulfilled desires and, counterintuitively, population conformity.

It is interesting to compare this sample size effect to that seen with tournament selection in standard genetic algorithms (as analyzed in Goldberg & Deb, 1991): there, the larger the set of competing individuals, the greater is the selection pressure and speed of evolution. The difference is that in tournament selection, individuals are being selected by a fixed global fitness function applying across all tournaments, while in our female choice situation, each competition is decided by a different female and her individual preferences. Thus holding many large tournaments will find the best individuals in the whole population according to some single criterion, and




Figure 5: Diversity of songs in a coevolving surprise-preference sample-size-20 population, showing evolution of two tightly-clustered song “species” between generations 600 and 800.

the population as a whole can move quickly to the optimum. But correspondingly large mate-sample sets in the form of sexual selection incorporated here find the best individuals for each idiosyncratic female preference, leading to little cohesive population movement.

Overall, then, we have shown that sexual selection via coevolving male-generated mate-attracting signals and female-operated mate-assessing directional preferences can lead to the maintenance and continual turnover of signal diversity over time. But we have not yet satisfactorily answered the question of where such signal diversity comes from in the first place—why don’t we just have *one* love song at a time, changing from year to year? In other words, how can we ever progress from one song to many? Our simulation does not yet address this question: we have started with an initial population of many different male songs, and seen how that diversity changes over time. What we need to do next is to start with a population of males who all sing the same song initially, and see how the different female choice configurations we have described here alter *that* population. We suspect that, once again, coevolving surprise-based preferences with small sample sizes will first diversify the population, and then continue to alter that diversity across successive generations. (We can also start out with a converged female population all sharing the same preferences for song traits, here in terms of note transitions, which are *not* present in the initial male population, and

see how male song evolution adapts to those preferences over time. This situation matches the idea of pre-existing sensory biases on the part of females, which can drive the evolution of new male traits, rather than the traits and preferences both emerging over time together.)

Another issue we are beginning to explore is how to create the female expectations in the first place: where should their transition tables come from? In our current system, females inherit their transition tables from their mother and father (after the females in the initial generation were loaded with transition expectations computed from real song examples, as mentioned in section 4.2). Because of this, “surprising” note transitions can only be surprising relative to a particular female’s inherited expectations. But certainly for humans, and for other animals as well, expectations are built up through experience and learning within one’s lifetime (see Bharucha & Todd, 1989). So instead we can let a female learn expectations about note transitions based on a set of songs from her current generation, or from the previous generation, as if she has heard those songs and picked up knowledge of her “culture” from them. Then she will be surprised when she hears something new that toys with these learned expectations, building them up and then violating them. We expect that using learning to create the note transition expectations, rather than evolving them, will allow the population to “change its tune” even more rapidly than the cases we have described in this paper, because the expectations will be able to shift just as rapidly as the songs themselves—learning operates faster than selection. Thus, love songs this year may be different from the love songs of yesteryear because, in part, everyone listening to them gets bored, and selects instead for novelty and surprise.

Furthermore, we could allow learning in the females to occur at an even faster time-scale, so that instead of habituating to songs heard too many times last week, each female could habituate to notes and phrases heard too many times within the current male’s song. In this case, females would seek novelty and expectation-violation *within* each song they hear. To sing preferred songs, males will have to balance the amount of repetition and newness in one song just right. We expect that this will lead to increased complexity of the internal structure of the songs themselves (not just of the population of songs), allowing us to explore the other great mystery about elaborate communication signals.

One final question must be considered: So what? Why is it important to understand the sources of diversity in communication signals? Isn’t that just the way the cultural world is, full of all sorts of inexplicable variety? The answer is no, that cannot be just the way the world is. The diversity of signals does not come for free: to be able to generate and evaluate this range of possibilities, there must be correspondingly extensive behav-

ioral mechanisms as well. So the evolution of communication signal diversity must be linked to the evolution of brains and behavior more generally. Indeed, Miller (1993) has proposed that the very size and structure of the hyper-encephalized human brain has been most strongly (and most recently) shaped by a process of runaway sexual selection, new neural circuits being added by the pressure to produce and evaluate ever more elaborate mate-attracting cultural displays, including music and language. If this is true, our vast cultural diversity reflects the very essence of our evolved, sexually-selected human nature, and we should expect the never-ending flow of ever-changing love songs to be with us for a long time to come.

5 References

- Ackley, D.H., and Littman, M.L. (1994). Altruism in the evolution of communication. In R.A. Brooks and P. Maes (Eds.), *Artificial Life V* (pp. 40-48). Cambridge, MA: MIT Press/Bradford Books.
- Andersson, M. (1994). *Sexual selection*. Princeton, NJ: Princeton University Press.
- Bharucha, J.J., and Todd, P.M. (1989). Modeling the perception of tonal structure with neural nets. *Computer Music Journal*, 13(4), 44-53.
- Catchpole, C.K., and Slater, P.J.B. (1995). *Bird song: Biological themes and variations*. Cambridge, UK: Cambridge University Press.
- Cliff, D., and Miller, G.F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Moran, A. Moreno, J.J. Merelo and P. Cachon (Eds.), *Advances in artificial life: Proceedings of the Third European Conference on Artificial Life* (pp. 200-218). Lecture Notes in Artificial Intelligence 929. Berlin: Springer-Verlag.
- Darwin, C. (1871). *The descent of man and selection in relation to sex*. London: John Murray.
- Driver, P.M., and Humphries, D.A. (1988). *Protean behavior: The biology of unpredictability*. Oxford, UK: Oxford University Press.
- Goldberg, D.E., and Deb, K. (1991). A comparative analysis of selection schemes used in genetic algorithms. In G. Rawlins (Ed.), *Foundations of genetic algorithms*. San Mateo, CA: Morgan Kaufman.
- Hamilton, W.D., and Zuk, M. (1982). Heritable true fitness and bright birds: A role for parasites? *Science*, 218, 384-387.
- Hanlon, R.T., and Messenger, J.B. (1996). *Cephalopod behavior*. Cambridge, UK: Cambridge University Press.
- Hartshorn, C. (1973). *Born to sing*. New York: Harper & Row.
- Hauser, M. (1996). *The evolution of communication*. Cambridge, MA: MIT Press/Bradford Books.
- Kirkpatrick, M. (1987). The evolutionary forces acting on female preferences in polygynous animals. In J.W. Bradbury and M.B. Andersson (Eds.), *Sexual selection: Testing the alternatives* (pp. 67-82). New York: Wiley.
- Krebs, J.R., and Dawkins, R. (1984). Animal signals: Mind-reading and manipulation. In J.R. Krebs and N.B. Davies (Eds.), *Behavioral ecology* (pp. 380-402). Sunderland, MA: Sinauer Associates.
- MacLennan, B.J. (1990). *Evolution of communication in a population of simple machines*. Technical report CS-90-99, Computer Science Department, University of Tennessee, Knoxville TN.
- MacLennan, B.J. (1992). Synthetic Ethology: An Approach to the Study of Communication. In C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (Eds.), *Artificial Life II* (pp. 631-658). Reading, MA: Addison-Wesley.
- Miller, G.F. (1993). *Evolution of the human brain through runaway sexual selection: The mind as a protean courtship device*. Unpublished PhD thesis, Stanford University Department of Psychology. (Available through UMI Microfilms.)
- Miller, G.F., and Todd, P.M. (1993). Evolutionary wanderlust: Sexual selection with directional mate preferences. In J. A. Meyer, H. L. Roitblat, and S. W. Wilson (Eds.), *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 21-30) Cambridge, MA: MIT Press/Bradford Books.
- Miller, G.F., and Todd, P.M. (1995). The role of mate choice in biocomputation: Sexual selection as a process of search, optimization, and diversification. In W. Banzhaf and F.H. Eeckman (Eds.), *Evolution and biocomputation: Computational models of evolution. Lecture notes in computer science 899* (pp. 169-204). Berlin: Springer-Verlag.
- Pinker, S. (1994). *The language instinct*. New York: Morrow.
- Searcy, W.A. (1992). Song repertoire and mate choice in birds. *American Zoologist*, 32, 71-80.
- Todd, P.M., and Miller, G.F. (1991). On the sympatric origin of species: Mercurial mating in the Quicksilver model. In R.K. Belew and L.B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms* (pp. 547-554). San Mateo, CA: Morgan Kaufman.
- Weatherhead, P.J., and Roberston, R.J. (1979). Offspring quality and the polygyny threshold: The "sexy son" hypothesis. *American Naturalist*, 113, 201-208.
- Werner, G.M., and Dyer, M.G. (1992). Evolution of Communication in Artificial Organisms. In C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (Eds.), *Artificial Life II* (pp. 659-687). Reading, MA: Addison-Wesley.