Relaxed selection and the role of epigenesis in the evolution of language

Terrence W. Deacon University of California, Berkeley E-mail: deacon@berkeley.edu

ABSTRACT It is generally assumed that there is a positive correlation between the complexity of functional adaptive systems and the intensity and duration of natural selection driving their evolution. Although the role of selection is unquestionable, especially with respect to functional correspondences between organism and environment, the correlation with functional complexity is not so clear. The recent resurgence of interest in the contribution of epigenetic processes to the course of evolutionary change, popularly known as evodevo, has begun to focus attention on other potential sources of complex integration: i.e. the self-organizing and intra-selection processes that are recruited by evolution to serve epigenetic functions. I review evidence for a general evolutionary logic that I call the "Lazy Gene" effect, which suggests that genes will tend to offload control of morphogenetic processes to epigenetic mechanisms in evolution whenever reliable extragenomic constraints or influences can induce the same effect. This is because such extragenomic factors will produce relaxed selection with respect to these traits and their differentiation. But this reduction in the genetic constraints increases the probability that dynamically generated combinatorial effects of epigenetic interactions will emerge, increasing the probability of generating synergistic interactions among these epigenetic processes, and—if this proves to be adaptive—a more complex phenotype. I explore an example of this effect that is particularly relevant to language evolution: a case of birdsong change under the influence of artificial breeding for an unrelated trait. The song of the domesticated Bengalese Finch shows an increased complexity and variability of structure, increased complexity of its neural production, and an increased role for social learning, compared to its wild conspecific, the White-Rump Munia. This has evolved in the absence of explicit selection on singing. These effects are consistent with the predictions of the lazy gene hypothesis. Comparing this to behavioral and neurological features associated with language demonstrates many parallels, suggesting that relaxed selection likely contributed to the human capacity for speech and language. By implication, the extensive offloading of language maintenance onto social transmission processes is evidence that the human language capacity is emergent from epigenetic interactions with this extrinsic inheritance, and not a product of genetically encoded language-specific information.

The evodevo perspective Historically, theories attempting to explain the evolution of brain organization and neural complexity have not ascribed a significant role to developmental processes, other than carrying out the implementation of genetic instructions for neural differentiation and adapting neural circuits to receptor and effector systems of the body. This is understandable from the point of view of what might be called classical neoDarwinism. To the extent that the mechanism of natural selection is classically framed in post hoc terms and new variations of structure and function are understood as arising irrespective of any eventual functional consequence, the specific mechanisms producing these variations tend to be considered irrelevant with respect to their adaptive consequences. This classical perspective has, however, come to be seen as a special case of a more general paradigm, popularly known as "evodevo." In this augmentation of neoDarwinian theory, constraints on the development and differentiation of organism features are understood to introduce an independent source of structure into the

evolutionary process. This approach does not assume any violation of the neoDarwinian doctrine that the ultimate sources of the variations presented to natural selection only derive from chance genetic mutations and sexual recombination. But, it additionally recognizes that those variations, which ultimately get expressed at the whole organism level, are products of developmental mechanisms that are themselves highly structured, indirect, and systemically interdependent. In this respect, past adaptations, and specifically the mechanisms underlying their developmental expression, serve as the context in which unprecedented genetic variations get expressed. Thus, otherwise randomly generated genetic variations do not necessarily have uncorrelated developmental consequences. To invert a famous aphorism about computing: genetic garbage in, does not necessarily produce phenotypic garbage out. More importantly, these developmental constraints and biases contribute an organizing influence on the variant traits presented to natural selection; an influence that is independent of genetic variation, and which indirectly reflects constraints of past adaptations. One of the more counterintuitive consequences of this epigenetic contribution to evolutionary design is that its effects can become more marked under conditions of relaxed selection,1 precisely because the systematic biases these mechanisms contribute are not constrained by the demands of adaptive fit. This introduces a non-Darwinian mechanism, which acts as a complement to natural selection in the generation of functional synergies, because its form-generating properties derive from the self-organizing tendencies of molecular and cellular interactions rather than from relationships to environmental conditions. Paradoxically, this suggests that selection may actually hinder the evolutionary "exploration" of alternative functional synergies, and that the relaxation of selection may play an important role in the evolution of increased functional complexity. This possibility will be explored below. In as much as the architecture of the central nervous system is generated via the operation of many more levels of developmental interactions than most other tissues—including contributions from such uniquely neuronal mechanisms as activity-based modification of connectivity and axonal extension linking distal cell populations—there is an enormous opportunity for developmental mechanisms to contribute to brain evolution. This greater degree of developmental and functional interdependency should make brain evolution particularly susceptible to epigenetic influences. This susceptibility is particularly relevant for understanding some of the more enigmatic and unprecedented events in brain evolution, such as the evolution of language. But to understand how epigenetic processes can contribute to patterns of evolution, even in the absence of the effects of natural selection (or, as we will see below, precisely because of its absence), we first need to step back from the neuroanatomical details to reconsider more generally how such entanglements of evolution with epigenesis and behavior can occur and how this may alter the simple predictions of neoDarwinian theory. In this chapter, then, we begin with a brief overview of the contemporary view of the role of development in brain evolution and outline the historical development of ideas about this entanglement of levels of formative processes in biology. Next we consider some counterintuitive implications of the role of developmental processes in evolution, as contrasted to natural selection processes, and offer a few paradigm examples of these general principles at work at multiple levels of organic structure and function. In the final sections of the chapter we will turn to a particularly counter-intuitive example of increased neural and functional complexity emerging in the absence of natural selection, and finally show how many features of human cognition and language function also exhibit analogous features. There is little doubt that many features of human brain evolution relevant to the language capacity have evolved under the

influence of natural selection, and identifying these effects and their causes has been the focus of nearly all theories of human brain evolution (Deacon 1997). But here we focus on the possibility that certain, perhaps crucial, contributors to the human difference may have arisen as a consequence of epigenetic constraints and biases rather than selection favoring adaptive improvement. Indeed, an adequate account of language origins depends on it. By refocusing attention on this otherwise ignored component of the evolutionary process, I do not intend to suggest that natural selection mechanisms are incorrect or subordinate to these epigenetic contributions. Instead, I believe that epigenetic constraints and biases contribute critical formative principles that must be considered an integral component of a complete natural selection theory. Developmental influences affecting mammalian brain evolution The starting point for any exploration of brain evolution is the recognition that brain organization in all bilateral animals shares deep commonalities, and that these are largely the result of highly conserved developmental mechanisms. The segmental organization of the brain is a consequence of the neuromeric partitioning of the embryonic neural tube by rostro-caudally distributed bands of gene expression that determine regional cell fates. At this level of differentiation one can observe a remarkably conserved plan of organization characteristic of many animal phyla. Thus, the homeotic gene expression in developing flies and mice exhibit a generic segmental organization that likely antedates their common ancestor. Vertebrate brain organization is a variation on this general theme, with more subdomains and dorso-ventral differentiation contributed by paralogues of these homeotic genes, creating mosaic patterns of concentric and intersecting expression domains (Rubenstein et al. 1994; Lumsden and Krumlauf 1996). Within vertebrates (though excluding the most primitive transitional forms) this patterning is ubiquitous, resulting in a shared brain plan in which there are local regional structural variations, but effectively no significant variation in the topology of major brain regions, such as the classic major segmental divisions, such as telencephalon, diencephalon, mesencephalon, etc.. The variations that are most common involve the relative sizes and degree of differentiation of the further subdivisions of these major structures (for a general overview see Butler and Hodos, 1996; Striedter, 2005). The major clades of the vertebrate radiation are distinguished by sharing distinctive patterns of subdivision organization within this global segmental architecture. The degree of structural variation of brains within any given clade of species is also roughly correlated with its phyletic age (Striedter 1997). Consequently, there is considerable diversity of regional brain structure size and specialization in different fish species due to the cumulative effect of long separate evolutionary histories in distinct ecological domains demanding highly specialized sensory-motor adaptations (figure 1). This size correlation of brain structure with functional importance or with peripheral organ development has led to a general expectation that relative deviation of subdivision size and differentiation, with respect to the average for the clade, reflects a computational rule of brain evolution, in which differential elaboration of a given structure with respect to others roughly reflects relative computational power dedicated to that function.

So, while there can be orders of magnitude difference in the sizes of mammal brains, their regional organization generally shares highly invariant patterns of relative growth. Among mammalian species, for instance, there are remarkable quantitative correlations linking the sizes of different major brain segments. Allometric studies of primate and insectivore brains, for example, have shown that the relative sizes of most of the major divisions of the brain are predictable from the single variable of brain size (see for example an early study

by Sacher 1970; and a recent approach by Finlay and Darlington 1995). There appear to be subtle clade-specific trends (e.g., between insectivores and primates, and between prosimians and anthropoid primates) but remarkably high correlations within mammalian orders despite large body and brain size ranges. These striking quantitative regularities almost certainly reflect highly conserved mechanisms for the control of early stem cell multiplication and neuronal differentiation throughout the brain.

But this comparative uniformity of regional growth patterns in mammal brains has not limited the capacity for mammals to evolve highly distinctive specializations of sensory and motor adaptation, which are reflected both in structural size effects and in relative differentiation of the relevant brain structures. Though much of this specialization can be attributed to divergent organization of peripheral structures, such as the eyes, the olfactory system, or musculature, there is also considerable brain reorganization possible at the subdivision level. The most obvious expression of this is in the distribution of functionally distinct subdivisions of the cerebral cortex (see a recent review by Krubitzer and Kaas 2005; see figure 3). Cerebral cortex is, however, peculiar to mammals and is one of the more anomalous variants of vertebrate telencephalic organization (Karten 1969). Although the vertebrate telencephalon is the outgrowth of a surface sheet of cells, in most vertebrates the multiplication of telencephalic stem cells lining the ventricular wall and their eventual production of neurons within this embryonic forebrain structure tends to enlarge the ventral and lateral walls of each telencephalic "bubble" so that they expand into the ventricular space to form nucleated structures. Only in mammals does the dorsal half expand in a highly laminar fashion, due to the mostly radial migration of newly produced neuroblasts through layers of cells produced in previous waves of neurogenesis. So unlike the telencephalic nuclei of other vertebrate brains, the resulting laminar and columnar organization of mammalian cerebral cortex maintains a uniform sheet topology that is well suited to maintain discretely separated map-like representations of its input-output connections with other structures; particularly with sensory receptor systems like the retina, the skin surface, and the organ of corti within the cochlea.

Whereas mammal brains lack the large-scale quantitative segmental adaptations characteristic of much older vertebrate clades, they do exhibit considerable quantitative variation of cortical subregions. In this way, despite the relatively limited phylogenetic depth of the mammalian radiation, mammal brains have evolved highly distinctive and divergent adaptive organizations. This evolvability is aided by two additional features characteristic of mammalian brains. The first of these is the comparatively larger size of many mammalian brains, both compared to other species' averages of gross brain sizes and with respect to average brain/body proportions; or encephalization. Since significant differences in vertebrate brain sizes are not correlated with a corresponding enlargement of neuronal soma sizes, larger brains have more opportunity for regional differentiation (Killackey 1990; Deacon 1990). This general rule is reflected across brain structures in all vertebrate species, but is particularly evident when comparing numbers of functionally distinguishable cortical regions in mammalian brains of different sizes. Moreover, the segmental scaling trend for mammalian brains is not isometric with respect to gross brain size. The telencephalon scales with positive allometry with respect to most other brain structures (e.g., diencephalon and brain stem) and so with size increase the relative proportion of the brain comprising cerebral cortex increases disproportionately. Thus the "opportunity" for regional differentiation of cerebral cortex significantly increases with size

as well. The second of these features of mammalian brains that contributes to their heightened evolvability, despite their conservative global organization, is the potential for developmental processes—such as cell-cell interaction and functional activity—to play a role in subregional organization of connections and neuronal differentiation. Afferentefferent determined modifiability of connectivity pattern and even of neuronal survival is a generic feature of the developing vertebrate CNS (see Holliday and Hamburger 1976; Purves and Lichtman 1980; Oppenheim 1981; Cowan et al. 1984; Wilczynski 1984; Finlay et al. 1987; Deacon 1990). It is, for example, responsible for the matching of central versus peripheral cell populations and establishing topographic mapping between them, in the case of spinal neurons and peripheral target cells, such as muscle cells (Purves 1988). But mammalian cerebral cortex offers a highly responsive forebrain arena for this mechanism to influence global brain function as well. Thus both the relative sizes and the organization of cortical regions associated with distinct peripheral specializations and analytic processes can be significantly influenced by the numbers of afferent axons projecting to them and the level and patterning of signals conveyed by these projections (Deacon 1990). This is demonstrated by the responsiveness of cortical patterns to manipulations of projections and afferent signals during brain development. Over the past three decades developmental manipulations have demonstrated afferent-efferent dependent differentiation of cortical functional differentiation in many cortical domains (e.g. Frost 1981; Stanfield and O'Leary 1985; Sur et al. 1988; and recent studies reviewed in Kubitzer and Kaas 2005) The interaction between extrinsic activity-dependent influences and intrinsic genetic mechanisms that establish cell fate and guide axon guidance growth have demonstrated that cortical regionalization is a multilevel process. A once lively debate between researchers who argued that cortical area distinctions were expressions of a locally controlled protomap (e.g. Rakic 1988) versus those who argued that they were the result of competition between peripheral influences conveyed by afferent axons (e.g. O'Leary 1989), has mostly been resolved by recognizing their necessary interdependence (Mallamaci and Stoykova 2006; O'Leary and Nakagawa 2002). Relatively conserved "generic" cell differentiation and axon guidance mechanisms are generated by the overlap of gene expression and growth factor gradients across cortex. These constrain and bias the subsequent topographic biases of axonal invasion, which in turn bias activity-dependent competitive interactions between afferent axons to determine the final fine scale topography and the relative sizes of cortical subdivisions. Evidence contributing to this general conclusion comes from many different research paradigms. Cross-species fetal neuronal transplantation studies have shown that initial axon guidance mechanisms are largely shared across species (e.g. Balaban 1997; Isacson and Deacon 1996). Manipulations of gene expression and regional growth factor production in the developing cortex has also shown that although there may not be protomap determination of functional area positioning in cerebral cortex, the relative concentration gradients of these molecular cues distributed in different and overlapping patterns set up a continuous "grid" defined by the different relative concentrations of these morphogens (Mallamaci and Stoykova 2006; O'Leary and Nakagawa 2002). As a result, experimental modification of these gradients by increasing or blocking expression of a particular molecular cue can distort the eventual functional topography, and ectopic introduction of one of these morphogens can fractionate this topology, producing mirrorimage duplication of distinctive cortical areas. Finally, very early loss of specific sensory afferents or peripheral degeneration of specific receptor systems can also cause significant changes in cortical area differentiation, including significant size effects and the failure to form distinctive boundaries between

cortical areas (e.g. Krubitzer and Kass 2005). So we can roughly summarize the process of cortical regional differentiation in terms of three general levels of developmental processes that ensue largely independent of cortical stem cell production and neurogenesis (which do not predetermine cortical sub-regionalization). The first phase is due to locally generated molecular gradients and their overlapping patterns of distribution, that function analogous to a coordinate grid. The second phase is characterized by afferent invasion from thalamic neurons influenced by afferent population differences. The third phase is characterized by synaptic competition between afferent axons, which is strongly influenced by activity-dependent signal-correlation mechanisms. By this means the sizes, organization, and functioning of distal brain and peripheral structures contributes to central organization. Additionally, the particularities of peripheral receptor organization and the invariants of stimulation imposed by extrinsic sources (such as the systematic visual map differences of the two retinas) also influences the organization of cortical function and differentiation. To summarize: a significant fraction of the structural information embodied in mammalian brains emerges from self-organizing and Darwinian epigenetic processes, and is not "coded" in the genes. The role of genetic information might thus be described as guaranteeing the maintenance of the boundary conditions necessary to support the reliable emergence of these spontaneous organizing processes. This mediating role of Darwinian-like epigenetic processes is depicted in Figure 4.

Can behavior influence brain evolution? Epigenetic influences on normal brain development, and especially peripheral stimulation and environmental invariants, open the possibility that behavior itself might directly influence the course of brain evolution. This possibility has been considered of central significance by many researchers interested in human brain evolution over the last century (including the present author, whose 1997 book The Symbolic Species explored this possibility). The classic attempt to explain how behavioral interaction with the environment could shape the course of evolution was the use-inheritance theory of Jean Baptiste de Lamarck (1809). He argued that animals acquired adaptations in response to environmental conditions and that these adaptations were inherited by offspring. He thus reasoned that acquired traits could be combined, adjusted, and honed by feedback trial and error and in this way convey considerable adaptive flexibility and exploration of alternatives, potentially aiding both optimality and integration of interdependent systems. Lamarck's ideas were later championed by Herbert Spencer during Darwin's time, but evidence against this mechanism of inheritance was presented by August Weismann and others toward the end of the 19th century, and additional contrary evidence for the transmission of acquired characters has accumulated since (though limited inheritance of other acquired traits has also been demonstrated). In the late 19th century, a non-Lamarckian mechanism was described that appeared able to produce an analogous evolutionary effect. Almost simultaneously Baldwin (1896) and two other researchers, Conwy Lloyd Morgan (1896) and Henry Osborne (1896), independently described a theoretical mechanism whereby acquired traits might influence natural selection in such a way as to increase the probability that parallel inherited traits might emerge spontaneously. Each of these theorists independently published papers that proposed such a mechanism. They argued that although Lamarckian inheritance is unlikely, Darwinian mechanisms might be able to produce analogous consequences by virtue of the fact that learning and plasticity can change the effects of selection imposed by the environment. Though this theoretical mechanism has come to be called the Baldwin effect, Baldwin himself called the effect "organic selection." He argued that the Lamarckian-like

consequence could be produced if an acquired trait, such as a learned habit, could block the effects of natural selection. The ability to use the flexibility of developmental and learning processes to persist and reproduce in conditions that are not optimal preserves a lineage. This would protect the organism's lineage from being eliminated in the reproductive lottery. This "umbrella effect" would also thereby increase the probability that some members of the lineage would additionally acquire a fortuitous variation that would more efficiently accomplish that which behavioral flexibility previously accomplished. Over time this would aid the accumulation of inherited variants that would lighten the load imposed by trial and error learning and enable fortuitous mutations to arise and replace the original acquired capacity. Baldwin called this process organic selection because it occurs at the organism level as a result of what an individual organism does during its lifetime. It is not something that is initially transmitted genetically and yet it can indirectly bias the direction of evolution, according to Baldwin, in such a direction as to lead it to become more genetically inherited. About 50 years later a developmental biologist, Conrad Waddington (1942), proposed a related process that he called "genetic assimilation." This process also appeared to produce an apparent path from environmentally acquired traits to congenitally generated traits. But in this case it wasn't just a theoretical possibility. Waddington (1953) demonstrated by experiment that stress-induced trait-expression could eventually come to breed true. He showed that a trait that in one generation was only expressed under special atypical environmental conditions could become an ineluctably expressed trait independent of the environment, merely by selectively breeding for that environmentally induced trait. Waddington's experiments with fruit flies raised in atypical environments (e.g., unusually high temperatures) showed that traits initially expressed by some individuals only when in these special environments could nevertheless be selectively bred to become expressed under normal conditions. He argued that this involved a kind of co-assortment effect in which multiple genetic contributors are collected together over generations by selectively breeding individuals with these environmental sensitivities, producing a synergistic epigenetic effect that independently parallels the environmental contribution. Recent, genetic analysis of this effect largely supports this combinatorial genetic-epigenetic explanation (e.g., Rutherford and Lindquist 1998). It is a common error to see these two theoretical mechanisms—organic selection and genetic assimilation—as variants of the same process (Deacon 2003). But aside from the fact that only Waddington's claims had direct empirical support, organic selection and genetic assimilation only share a superficial similarity. And neither actually provides even a pseudo-Lamarckian mechanism.

Crucial to both mechanisms is the fact that many expressed physical and behavioral traits may have no reproductive consequence for generations. For example, nose shape, chin shape, earlobe shape, and many other human physiognomic traits may be below the threshold of having a reproductive consequences and so are effectively neutral with respect to natural selection. Neutral traits can almost freely vary in form over the course of evolution until for some reason they rise above this threshold. Until this happens, however, their probability of expression also drifts randomly from individual to individual and generation to generation. With respect to the differential susceptibility of traits to selection, Waddington's and Baldwin's mechanisms are in fact inverses of each other. Waddington's mechanism assumes that certain previously neutral traits become newly subject to natural selection due to changing environmental conditions. In contrast, Baldwin's mechanism assumes that reducing selection due to behavioral or physiological flexibility can create

"space" for better variants to emerge. This difference has significant implications for the two mechanisms and what they are likely to produce over generations. The basis of Waddington's effect is an exposure (or "unmasking") of features that were previously hidden from selection, and he showed that this produces an organizing effect influencing who mates with whom by virtue of these unmasked traits. The basis of Baldwin's claimed effect is that learning and other types of environmentally responsive plasticity shield the organism from selection. This hides (or masks) many variations that were previously exposed to selection, with the result that less and less is selectively stabilized, more drift is possible, and thus more variation eventually shows up. This difference matters. The first step in gaining perspective on these mechanisms is to abandon the Lamarckian framing of the issue. Rather than attempting to frame a Darwinian means to achieve a Lamarckian end, we can simply ask what Baldwin's mechanism would likely produce. Many problems with the Baldwin Effect have been identified over the years but are often not generally acknowledged. For example, the phenotypic plasticity that it depends upon will increase variability but will not necessarily favor selection for new adaptations to replace it. The very plasticity that favors an increase in variation will also mask (i.e. partially shield the organism from) the very forces of selection that would be necessary to shape up or stabilize any of the congenitally produced surrogates that by chance could have supplanted this acquired adaptation. This shielding would not only inhibit their evolution, it would likely also degrade any existing congenital analogues to the acquired adaptation due to a failure to eliminate mutations affecting it. So reduced selection is a double-edged sword. In 1953 the evolutionary biologist George Gaylord Simpson wrote a critique of the Baldwin Effect (in which the phrase "Baldwin Effect" was coined) that pursued some of the implications of this problem. Simpson argued that indeed it did not violate the strictures of strict neoDarwinian logic, but that it may be only trivially important in evolution because it can only occur in very special conditions. These constraints have been more rigorously explored by computer simulation studies (e.g., Yamauchi 2004). First, it requires that there is a high cost-benefit ratio between the acquired and highly canalized alternatives of a trait. For example, if trial and error learning is associated with high error costs (e.g., likely getting eaten by predators) then an inherited stereotypic behavior that is easy to produce (e.g., a tendency to automatically hide when startled) and not particularly costly (e.g., doesn't impair exploration for food or mates) will tend to be favored. This can be viewed as merely two competing adaptations, though it puts considerable weight on the probability of such a coincidental mutation occurring. But even if such a coincidental variant were to arise, the cost-differential between the acquired and inherited trait would be self-limiting for selection. Because of the flexibility of an acquired trait, if a highly heritable variant came to co-exist alongside it, the effect will be to reduce the costs of its acquisition, and so will reduce selection favoring it. The result is that evolution will stall somewhere in between. A second condition that limits the Baldwin effect is more troublesome. For the replacement effect to work there also has to be a high correlation between the acquired trait (also sometimes called a phenocopy, following Waddington) and the ineluctable highly canalized version. This is most likely when alleles of the same genes are involved in producing both the acquired and congenitally generated adaptations (a condition that is probably highly unlikely for complex traits). As a result, it is most likely in cases of simple genotype-phenotype correspondences and less likely where multigenic and epistatic factors are involved. A third restrictive condition, related to the first, is that the presence of an acquired adaptation sufficient to preserve the linage that expresses it, creates conditions that limit forces of selection that would be necessary to select among initially crude

versions of an alternative congenitally produced adaptation to make it progressively more suited to its context. This is especially significant if the adaptation in question must be integrated with other functions or responsive to variable environmental conditions. Thus models of the Baldwin effect that postulate "hopeful monster" mutations (i.e. which produce a fully functional "innate" alternative phenotype) have shown the most promise, but make unrealistic epigenetic assumptions. Despite these constraints on a Lamarckian outcome, however, traits becoming shielded from selection by virtue of acquired adaptations should be quite common. The Baldwinian mechanism is therefore important even if its predicted effects might not be what Baldwin envisioned (Deacon 2003). Are some of these other consequences of the Baldwinian masking effect relevant to evolution? The answer is yes, and as I will argue below, they may be far more significant than Baldwin could have imagined. They just turn out to be, paradoxically, the reverse of what Baldwin claimed, and yet perhaps more important in evolution than Baldwin's predicted outcome. The evolutionary cycle of duplication-masking-degeneration-complementation There is a parallel to the Baldwinian masking of the effects of selection in a wellstudied genomic effect: gene duplication. Gene duplication is a common occurrence in the evolution of genomes (Ohno 1970; Ohta 1994; Van de Peer, Y. et al. 2001). It is probably the major source of new genes in the course of evolution. It is also a major means by which cooperative protein complexes arise in evolution (Orgel 1977; Zhang 2003). Thus, multiple occurrences of gene duplication over the course of evolution have produced large "families" of structurally and functionally related genes. Indeed, most genes can be recognized as members of larger families of genes sharing a common ancestral gene (Walsh 1995; Zhang 2003). During gene duplication, a length of DNA is literally copied and spliced into the chromosome nearby, possibly as a result of uneven cross-over events during meiotic replication, viral gene insertion and excision, or some other intrinsic or extrinsic mechanism that modifies gene replication. The result of such events is that a nucleotide sequence may be duplicated that contains intact regulatory and coding segments for production of a functional protein. The functional consequence is that there is now two ways of producing the same phenotypic effect. This redundancy can provide masking of selection on the duplicate gene's function, much as learning can mask selection on phenotypic adaptations in the case of the Baldwin effect. Thus if one of two duplicated genes acquires a mutation that alters its protein product in a way that modifies or degrades its function, this mutation won't necessarily impact the reproduction of the organism so long as the other copy remains intact. Moreover, the now mutated gene can continue to acquire mutational changes without negatively impacting organisms that inherit it, so long as this slightly modified phenotypic contribution is not somehow deleterious. Such mutations will thus be effectively or nearly neutral. The typical consequence of this sort of neutrality with respect to selection can be described as a "random walk" away from the original function. The result is the accumulation of arbitrary sequence changes at the genetic level and a progressively degraded or de-differentiated contribution to the phenotype. Presumably, persistent shielding from any selection effects will eventually lead to complete loss of function, as in pseudogenes that no longer produce any RNA transcripts. Accumulation of a very large number of mutations, or of mutations that stop the translation of its sequence information, can in this way ultimately produce complete loss of function. This has been the fate of a very large number of genes in the human genome which were once associated with a more acute olfactory sense (e.g., Rouquier 1998). It is estimated that a typical mammal has on the order of 1300 genes encoding distinct olfactory receptor molecules. This number was radically reduced both in primate and in human evolution.

The human species has had roughly 60% of these genes degrade to become pseudogenes (Gilad et al. 2003). This clearly reflects very weak, to nearly nonexistent, selection to maintain the numbers and diversity of these receptor molecules. But degradation to pseudogene status is not inevitable for gene duplicates. Precisely because gene duplication can involve an already functional segment of DNA, slight degradations of its sequence will only incrementally alter the structure of its coded protein. So long as the changes do not involve an essential binding site or some other critical structure, its functional links to other molecular components are also likely to degrade in non-catastrophic ways. This often means a progressive loss of specificity, with some functional associations being lost but other related interactions becoming possible that were previously prevented by the stabilizing selection that sustained the original function. In this sense, the progressive drift in genetic sequence is at the same time an exploration of the "phase space" of interactions "near" an original function. And whereas a single protein may require structural compromises to accommodate its multiple associations to others, multiple variant forms may provide a "have your cake and eat it too" option with each variant able to evolve greater specificity for one or another of these capacities. In other words, the duplication, masking, and random walk can provide a kind of exploration of the space of possible synergistic relationships that lie, in effect, in the "function space" just adjacent to an existing function. This is a recipe for increasing functional complexity (Lynch and Conery 2003). If the prevalence of gene duplication in animal and plant genomes is any indication, thhe probability that a given duplication will achieve functional integration is far from zero. Gene families, consisting of large numbers of paralogous genes (e.g. derived from a common ancestral gene, are widespread in complex organisms, and are often responsible for similar or even synergistic phenotypic functions. One incidental advantage for genomic research has been that identification of a functional correlate of one genetic sequence often provides a probe sequence that can be used for searching out other members of its family that have related functions. To illustrate this, let me describe two well-known examples. The first is the globin gene family, and specifically the hemoglobins. The hemoglobin protein complex contained in red blood cells comprises two varieties of the hemoglobin protein—alpha and beta hemoglobin—each coded by a distinct gene. The structure of the protein makes it possible to bind a special molecular formation within which an iron atom is suspended. It is this iron atom that provides the oxygen binding capacity. Two alpha and two beta hemoglobin proteins fit together to form a tetrahedral complex made possible due to the complementary shapes of the molecular surfaces forming the interior of the tetramer. The two forms of hemoglobin arose from a gene duplication event, and the ancestral hemoglobin itself arose as one of two duplicates from the common ancestral gene for both hemoglobin and myoglobin. The alpha and beta hemoglobin duplicates each acquired independent changes in shape but only minimal changes in oxygen binding capacity in their separate divergent "degradation." Changes that increased the stability of tetrameric binding appear to have been favored by natural selection with respect to one another, probably because of the superior oxygen transport capacity of the tetrahedral form. In other words, in their random walks through different 3-dimensional configurations, the duplicates retained their oxygen binding function while effectively "sampling" functional consequences of this secondary feature of molecular shape. This particular combination of alpha and beta hemoglobins is not, however, present at all stages of the mammalian life-cycle. In the fetus of a placental mammal, additional variant beta-hemoglobin forms are expressed, three of which are termed gamma, delta, and epsilon hemoglobin. These variants are expressed at different stages of gestation, and are each

coded by a different variant duplicate of the beta form of the gene, with the entire family present in a continuous segment of the chromosome. These beta hemoglobin duplication events, which occurred during the course of placental mammal evolution, have also given rise to two pseudo beta hemoglobin genes, which no longer produce a corresponding protein. In effect, these variants acquired mutations that inactivated gene translation in their random walk away from the original sequence. The remaining four beta-hemoglobin genes are expressed at slightly different times during development in the order epsilon-gamma-delta-beta. The functional value of this is related to the fetus's need to acquire oxygen from mother's hemoglobin and yet still transfer it from blood to somatic cells. So in order to be able to "steal" oxygen from maternal hemoglobin, fetal hemoglobin requires a slightly higher oxygen binding affinity than mother's hemoglobin. It then must diffuse oxygen out of its own hemoglobin into its tissues, which ideally requires yet a higher oxygen binding affinity than its own hemoglobin. Between these values there is an optimal balance, but this will change as the fetus grows larger in size and its oxygen needs change. The result is that the different beta hemoglobin variants expressed during different phases of gestation each have aslightly different affinity for oxygen that allows the fetus to progressively adapt to this challenge, until at birth the beta hemoglobin becomes the predominant form produced. So in this case, analogous to the shape complementarities "discovered" consequent to alpha/beta duplication, these parallel random walks of beta-hemoglobin gene duplicates led to synergies of timing and molecular affinities, as certain variant mutations which modified the different redundant genes' oxygen-binding properties became subject to selection with respect to each other in the context of internal gestation. Perhaps the most dramatic example of the duplication, masking, random walk, and functional complementation effect is demonstrated by duplication of genes that code for proteins that bind to DNA and regulate the expression of yet other genes. One consequence of this hierarchic recursive genetic relationship is that changes in one gene can influence a large number of other genes in concert. So the functional divergence and interaction effects that result from duplication of such regulatory genes can be global and systemic. The classic example of regulatory gene duplication effects involves a family of genes containing a motif called the homeobox. In the fruit fly they are called HOM genes, and their homologues in mammals are called Hox genes (though there are a very large number of more distantly related regulatory genes as well). These underwent a number of duplications in the common ancestry leading up to the separation of the arthropod and vertebrate lineages. Because these genes affect coordinated expression of large suites of other genes (many of which also have further regulatory functions) they play a role in producing slightly variant forms of whole body structures in these animals. This was first demonstrated by recognizing that mutations of these genes produce systematic variations of body segments in flies causing out-of-place expression of structures that normally are segment-specific, such as legs expressed where antennae are normally produced. The discovery that the theme-and-variation logic of the different insect body segments was correlated with the expression of different HOM gene duplicates in that segment, has revolutionized the study of development, and served as the keystone insight solidifying the value of the evodevo paradigm. But the fact that homeotic gene duplication expresses itself as organ duplication demonstrates that the logic of duplication, masking, divergence, and complementation is general. In arthropods such as centipedes the corresponding organs (e.g., legs) of adjacent segments are highly similar, but since adjacent legs serve almost identical functions they can also partially mask selection on the functional specificity of one another. This reduction of the effects of stabilizing selection can lead to drift of features on

one segment away from those on another. The structural-functional redundancy provided by adjacent segments minimizes the probability of catastrophic loss of function, and also increases the likelihood that complementary functions might develop on other segments. In various arthropods, such as grasshoppers, spiders, lobsters, flies, and so forth, the different appendages with leg-like form have evolved into specialized antennae, spinnerets, claws, and many other structures sharing the same jointed architecture, but modified to serve quite distinct functions.

In each of these cases, and despite their different levels of function, the redundancy of function that results from duplications significantly reduces the improbability of evolving synergistic functional linkages. Because they share a common ancestral function, randomly variant duplicated features of an organism potentially "explore" the diverse dimensions of the original function. Their underlying commonalities also increase the probability that variant duplicates will fractionate the original function, each assuming greater roles with on but not another aspect of the original thereby increasing synergistic organization. By fractionating an originally unitary and poorly optimized adaptation, variant duplicates can also be independently expressed in response to critical variables, and thus provide greater flexibility of function with respect to uncertain environments. This interplay between duplicated genetic epigenetic factors borrows features from both Baldwinian and Waddingtonian mechanisms, and yet it does not involve, even superficially, a Lamarckian logic. The way that duplication reduces the constraints of natural selection on a particular structure or function is analogous to the way acquired adaptations produce what Baldwin thought of as protection from selective elimination. But unlike Baldwin's hypothesized effect, this more often contributes to degradation of epigenetic constraint than the reverse. The way that the resulting functional interactions exploit combinatorial relationships that were previously hidden (or inaccessible because intense selection prevented variation) is analogous to Waddington's logic of canalization, in which epigenetic interdependencies can emerge to become selected in their own right. Like Waddington's notion of a "phenocopy," a novel functional capacity that emerges from complementary combinatorial relationships can become selectively favored for the synergy that results. Together these effects not only "explore" adjacent functional possibilities and "capture" novel higher-order synergistic relationships, but they provide an evolutionary cycle that can generate progressively more complex forms of adaptation, as each stabilized synergistic relationship can supply the substrates for new duplication effects. Extrinsic factors and the lazy gene hypothesis But this interplay between Baldwinian and Waddingtonian mechanisms suggest an even more general application of this principle, which is particularly relevant to brain evolution. This is the possibility that redundancy and masking effects can be generated extrinsically, and maintained irrespective of specific genetic inheritance, by virtue of redistributing selection fractionally onto highly diverse, and previously independent, genetic loci and extra-genomic mechanisms. A classic example that bridges between genetic and environmental duplication-masking effects is the evolution of ascorbic acid (vitamin C) dependency in anthropoid primates. Monkeys and apes, including humans, are among some of the very few mammals that must obtain ascorbic acid from dietary sources (Chatterjee 1973). Most mammals modification with a number of whole gene deletion and duplication effects distinguishing each cluster. This has produced a more cryptic segmental organization of the body plan in vertebrates in contrast to arthropods. But still the theme-and-variation logic is visible in parallels between limbs and their digits, vertebrae, ribs, teeth, and many other segmental features.

synthesize their own ascorbic acid. This is the case for rats. In 1994 a group of Japanese researchers (Nishikimi et al. 1994) sequenced the gene on chromosome 8 of the rat that codes for the final catalyst in the metabolic pathway that endogenously produces ascorbic acid (called L-gulano-lactone oxidase, abbreviated GULO). They then used the sequence from this gene to probe the genomes of other species. One of the first species they probed was Homo sapiens. What they found was surprising. Although humans are unable to synthesize their own ascorbic acid, the human genome includes a pseudogene that is homologous to the rat GULO. The GULO pseudogene in humans has accumulated considerable mutational damage, including the deletion of large coding regions (exons) and the random insertion of "stop" codons (see Figure 5). This is evidence that it has long been freed from the stabilizing influence of natural selection, and that the sequence has effectively taken a random walk, resulting in complete loss of function. So what masked its functionality, and allowed it to degrade to this extent?

Figure 5. Comparison of the GULO gene in rat and human. The rat gene produces a protein that catalyzes the final step in ascorbic acid (vitamin-C) synthesis. The human pseudogene for GULO has incorporated many deleterious mutations, including deletions of entire exons and the insertion of stop codons.

Although phylogenetic analysis of the variants of the GULO pseudogene in other anthropoid primates is still incomplete, it is likely that all share a GULO pseudogene with divergent mutations. A reasonable estimate of the date in the evolution of primates when

this gene began to accumulate damaging mutations is suggested by the comparative fossil evidence. Changes in eyes and teeth of fossil primates suggest that a shift to diurnal foraging and a shift from insectivory to frugivory took place roughly 35 million years ago in the lineage leading to anthropoids. The evolutionary implication is that at this point regular foraging on fruit introduced a semi-reliable extrinsic source of ascorbic acid into the diet. Under these conditions there would be no selective disadvantage of inheriting or transmitting a non-functional variant of the GULO gene. Selection would be masked by an acquired behavioral adaptation and by the ascorbic acid rich niche that was thereby created. The eventual complete loss of function of the GULO gene would lead to the equivalent of an evolutionary addiction to foods providing ascorbic acid (depicted in Figure 6).

Figure 6. Schematic depiction of the stages in the evolution of primate dependence on environmental sources of ascorbic acid (e.g. via frugivory). In the third stage (far right) the GULO gene has become entirely non-functional and selection has shifted to any loci in the genome that produce phenotypic effects that help maintain reliable access to extrinsic ascorbic acid (e.g. digestive adaptations and 3-color vision; depicted in green and blue).

This example of extrinsic ascorbic acid dependency provides us with an opportunity to look at some of the secondary consequences of this degradation. The masking of selection maintaining this enzyme in turn would have unmasked selection on a variety of other traits that help guarantee the availability of this now essential nutrient. Quite possibly, the evolution of 3-color vision and various tooth and digestive adaptations were also related to this evolutionary addiction, as the need for extrinsic ascorbic acid would become a new selection pressure. In other words, the behavioral flexibility that initially allowed primates to regularly forage on fruit and led to the spontaneous degradation of

the ascorbic acid synthesis pathway, eventually "addicted" anthropoid primates to a behavioral niche in which fruit acquisition and digestion were critical. This addiction would have unmasked selection on many diverse traits that coincidentally supported this "addiction," Including, for example, the capacity to judge the ripeness of fruit, forage on the outer limbs of trees, find the sugar-rich and slightly acidic content of fruit attractive, and metabolize the sugars and tolerate the ethanol that ripe fruits contain. All of these could be considered part of an adaptive suite for guaranteeing the presence of ascorbic acid. This process shares a number of features in common with the evolutionary process now generally referred to as niche construction (Odling-Smee et al. 2003; and see also West et al. 1988). One of the classic examples of niche construction is the production of beaver dams. Beaver dams produce an artificial aquatic environment. Beaver bodies currently exhibit adaptations to this aquatic niche as do a variety of beaver behavioral tendencies. So the artificial niche created by beaver behaviors has played a role in the evolution of many other beaver adaptations, in the same way that independent environmental features play a role in selectively favoring the preservation of other traits. This self-made environment introduces a sort of short circuit in the cause-effect cycle of natural selection. Beavers today are essentially unable to successfully reproduce unless they live in such a beaver-made niche. The development of stone tool technology in our own evolutionary prehistory was a niche construction event with similarly ramifying consequences. For example, the use of stone tools in hominid ancestry appears to have mitigated the intense masticatory processing of vegetable foods that characterized the australopithecines. Within a half million years of the appearance of the first stone tools, there was a radical reduction in the large flat molars, thick enamel, robust face and jaw structure, and powerful jaw muscles of the australopithecines. As I will argue below, this initial step of niche construction was only the first of many niche construction innovations that radically altered the forces of natural selection affecting our ancestors and their cognitive traits. This artificial niche did not merely create new selection pressures, it also eliminated many others. As I will argue below, both effects are relevant to explaining the anomalous trajectory of human evolution. To provide a mnemonic aphorism that characterizes the special evodevo logic that s implicit in all these examples, I call it the Lazy Gene hypothesis. The Lazy Gene hypothesis suggests that we should not assume genetic micromanagement of epigenetic processes, but rather only genetic regulation of the boundary conditions affecting processes that have the potential of arising by self-organization, self-assembly, or other extra-genomic dynamical processes. This view shares some features in common with the perspective called Developmental Systems Theory (e.g. see Oyama et al 2001; and also Johnston, this volume), which argues that the inheritance of extra-genomic factors is as important as genetic inheritance. One difference, however, is that the lazy gene approach assumes that genetic information plays a critical role maintaining the conditions for inheritance, without which most of these various extra-genomic influences would not be as reliably available. In other words, there is an important sense in which genetic information is primarily responsible for inheritance even if it is not the sole or even the major source of the morphogenetic processes contributing to epigenesis. More importantly, the lazy gene

approach attempts to explain how such extra-genomic factors might evolve to become reliable epigenetic mechanisms in the first place, even though their ultimate origin might be extrinsic to the organism. The most relevant (and also counter-intuitive) extra-genomic influences are those that arise from systemic effects, often involving what are often described as selforganizing processes, since these do not arise from specific structural or molecular effects, but rather from interaction dynamics. A classic example of a self-organized developmental process is the mathematically regular Fibonacci spirals observed in pinecones, sunflowers, celery stalks, and a myriad of other plants. This complex and highly regular pattern contributes to optimal spacing around a stalk and to scalable selfsimilar growth patterns over many orders of magnitude. However, its characteristic 3-5, 5-8, 8-13, 13-21, 21-33 … pattern of interlocking spirals (see Figure 7) is not explicitly coded in the genome. It is a self-organizing effect that emerges from cell-cell interactions in response to differential sensitivity to growth hormone expression. Only this hormonal response is passed from generation to generation by genetic inheritance. Indeed, many of the form-generating processes of embryogenesis involve the self-organizing consequences of recurrent cell-cell interactions and the spontaneous geometry of molecular diffusion processes. Like the duplication-masking effects at the genomic, phenotypic, and behavioral levels, such emergent organizing constraints will likewise promote degradation of redundant genomic influences and increasing dependency on and synergistic adaptation to these extrinsic factors.

Figure 7. Examples of spiral phylotaxis exhibiting highly precise Fibonacci relationships. Green arrows in the magnified image of a growing apical meristem demonstrate the sequence of eruption of new growth sites, determined by concentration of diffusible growth factors.

The lazy gene hypothesis is an epigenetic parsimony principle. Epigenetic biases and constraints that are reliably present or that reliably emerge during development, due to adaptive flexibility or extra-genomic influences, will tend to mask selection maintaining corresponding genetically inherited information. Thus the genome will tend to offload morphogenetic control, in the course of evolution, in a way that takes advantage of the emergent regularities that characterize many epigenetic processes. This does not mean, however, that we should treat genetic and epigenetic inheritance as simply parallel 21

inheritance systems that are interchangeable, as is sometimes argued by proponents of Developmental Systems Theory. Although gene expression depends upon epigenetic processes and epigenetic processes depend on conditions produced by genes, the genetic information is embodied in a structural artifact whereas epigenetic information is dynamical in origin and must emerge anew in each new developing organism. In this sense, the lazy gene argument recognizes that the ultimate source of genetic information is vested in self-organizing epigenetic interactions, but it also recognizes that because these effects are highly context sensitive, only an independent non-labile non-dynamical means for "remembering" these conditions from generation to generation (e.g. genetic information) is able to guarantee their reliable emergence in development. In this way we avoid the greedy reductionism of the new synthesis gene-centered view but do not confuse the complementary roles of genetic and epigenetic inheritance. On the one hand, this tendency to spontaneously offload morphogenetic control in evolution obviates the possibility that such emergent effects will ever be fully replaced by genetic prespecification, as Baldwin believed. On the other hand, it suggests a powerful non-Darwinain mechanism that may contribute to the evolution of complex functional synergies and to the emergence of highly distributed control of development. Because spontaneous self-organizing biases arise irrespective of natural selection, the two processes can complement one another. Natural selection captures and stabilizes the conditions for useful self-organizing processes, if they emerge, but stabilizing selection will tend to inhibit the expression of variant or alternative self-organizational tendencies. Without the selective elimination of most variants of the information embodied in the physical structure of genes, the fragile conditions conducive to self-organizing epigenetic effects would quickly be lost. But novel self-organizing effects can only emerge if these strictures periodically loosen. Relaxed selection and brain complexity: a birdsong example With respect to brain evolution and development, a spontaneous tendency for genomes to offload epigenetic control to any and all reliably present organizing influences outside the genome has interesting implications. The brain, more than any other organ, is likely to be strongly coupled with extrinsic regularities in the world. Also, because of the brain's highly reentrant network architecture and activity-dependent modifiability, nervous system functions are likely to generate complex emergent higherorder regularities. Brain evolution should therefore be highly susceptible to lazy gene effects. The effect of relaxed selection on brain organization and control of complex behavior is dramatically demonstrated by a recent series of studies comparing a wild finch, the White-Rump Munia, to a long domesticated breed of the same species, the Bengalese Finch (Okanoya 2004). The Bengalese Finch has been bred for coloration in Japan for about 250 years, but not for singing. And yet, surprisingly, it appears that in the absence of breeding for singing, being shielded from natural and sexual selection has produced increased song complexity, greater involvement of social learning in song development, and more diverse neural control of singing behaviors, as compared to its wild cousin. Since it is generally believed that complexity of singing behavior in songbirds is a consequence of sexual selection for male display (Darwin 1871; Catchpole & Slater 1996; Zahavi & Zahavi 1997; Okanoya 2002), this result appears paradoxical.

The song of the Bengalese Finch is complex and variable, and exhibits what amounts to a form of song "syntax" (here understood only to refer to combinatorial re-ordering possibilities). The specific details of song structure, including specific song elements, and variants of song syntax are acquired by early social learning from adult singers. It learns its song shortly after hatching, as is typical of many songbirds. And like other songbird species that acquire their songs via social learning, the Bengalese Finch makes use of multiple brain areas for song acquisition and production. In contrast, its wild cousin has a simpler and much more stereotypic song, for which learning plays a relatively lesser role. This raises a troubling question: How could an increase in the complexity of both song structure and of neural systems for producing song have evolved in the absence of overt selection acting on these traits? More generally, it raises questions concerning the evolution of the wide range and phylogenetic distribution of species differences in the flexibility, the extent of social learning, and the recruitment of brain structures in singing behavior. It also may shed light on some surprising seemingly incidental effects of domestication (e.g. see Belyaev 1979; Trut 1999). Whereas such an effect makes little sense for classical neoDarwinian accounts, the lazy gene perspective developed here offers a plausible mechanism for the shift away from congenitally prespecified stereotypic song to socially acquired flexible singing behavior, in the absence of selection. I have argued that prolonged domestication masked selection that in the wild had maintained highly canalized control over song structure and production (Deacon 2006). Relaxing these environmental pressures and undermining sexual selection effects led to spontaneous degradation of previously strong genetic constraints and opened the door to increased epigenetic variability and conditionality. This diminution of bottom-up constraints allowed a wider range of neural substrates and sensory experiences to influence singing behavior. But more importantly, I will argue that this relaxation-degradation effect and the upward shift of morphogenetic control that it produces, can under certain circumstances also lead to an increase in complexity of both brain and behavior. The Bengalese Finch example provides a model of how this process might begin. Homologous nuclei and connections are present in both social learners of song and stereotypic singers, as well as in both males and females (though typically only males sing). The differences mostly involving the relative size of telencephalic nuclei, the level of their interconnectivity, and the differential expression of molecular features (MacDougall-Shackleton & Ball 1999; Tobari et al. 2006). So the potential substrates for complex learned song are present generally, but they serve other functions in birds that do not sing or do so stereotypically, and probably did not play a role in complex vocal behavior in the common ancestor to songbirds, and their other avian cousins (Jarvis 2004). As a consequence of having their mate choice determined by breeders on the basis of coloration patterns, and irrespective of singing, selection in this species was relaxed on all aspects of song production. In the absence of either sexual selection for the display characteristic of song, or selection with respect to species identification or predator avoidance, we can assume that genetic influences constraining song production became subject to random drift. Under these circumstances any mutations affecting genes that contribute to the control of song structure, without imposing serious costs on other essential systems, would not be eliminated from the genome. The result, over time, would

be a kind of focal degeneration of song-constraining genetic influences in the domesticated lineage. As with other cases of genetic drift, the expected consequence would be an increase in phenotypic variability (i.e. song variability) and a decrease in its autonomy from other related influences. Analyzing the ethogram of the two breeds, Okanoya and colleagues observed evidence of increased song variability in the domesticated finch compared to the wild finch, but not merely an increase in variability. The songs can be analyzed into relatively discrete song elements, sometimes referred to as "notes." Transitions between characteristic notes have different probabilities characteristic of a particular song type or singer. The White-Rump Munia song is characterized by a limited and relatively invariant set of notes, as well as a very distinct distribution of transition probabilities between the different notes. In other words, most transitions are of either very high probability or very low probability (including many of near zero probability), resulting in highly predictable sequences of notes. Moreover, there are few if any loop-backs in this ethogram, which results in a very linear progression of notes from start to finish of a singing bout, and seldom any recurring note cycles. Different individuals within the species sing very similar versions of the same song. In contrast, the song ethogram of the Bengalese Finch includes both a larger repertoire of notes, many more non-zero transition probabilities between notes, and fewer very high probability transitions. This results in a far more complex ethogram, with many more transitions possible, including many possible loop-backs. Individuals thus sing a more varied song and the individual differences are also considerably greater. The Bengalese Finch song variability is not, however, just randomly distributed. It is strongly affected by experience listening to model singers during development. Interestingly, there is also evidence of individuals who learn from two different adult singers who produce a hybrid song as an adult, borrowing from each of its models. And Bengalese Finches can also learn to sing the White-Rump Munia song, if raised hearing only that song. The reverse in not true. White-Rump Munia chicks do not appear to be influenced by what they hear, or whether they hear another Munia singing at all. Thus if cross-fostered with Bengalese Finch parents they nevertheless mature to sing the typical Munia song. The effect of genetic drift on ethogram structure can be expected to decrease stereotypy in two respects. First, note structure should simplify and note specificity should become more variable both within and across individuals. Second, the specificity of the transition probabilities between notes should also be relaxed, causing probability differences to regress towards the mean, eventually making all transitions more equiprobable. Both effects are characteristic of the difference between Munia and Bengalese Finch song. A recent computer simulation test of this hypothesis (Ritchie & Kirby 2007) demonstrated that analogous relaxation of selection on "song" production of a population of algorithmic "agents," resulted in an analogous song complexification effect, as well as a serendipitous coupling between song structure and learning mechanism. These agent were initially selected so that "reproduction" of new agents from the prior generation was initially predicated on intense selection for song matching, and thus produced highly invariant songs. Increased complexity following relaxation of this restriction could be understood as merely increased variability, but the increased role of learning complicates

this interpretation, as does the fact that in living birds a vastly complicated genetic, epigenetic, and neurological system. In a songbird species that does not learn its regional variant of the species' song, there are relatively few forebrain structures playing a critical role in song production. A motor nucleus located poteriorly in the forebrain called the robust nucleus of the archopallium (RA) plays the major role in production of song in songbirds (e.g., Nottebohm et al. 1976). Damage to it can severely simplify or even eliminate singing ability in most bird species. In species with highly stereotypic songs, however, damage to other forebrain nuclei, generally has minimal detrimental effects on singing so long as it spares RA (except where it produces deafness prior to song development near maturity). This is not the case with species that acquire their regional song variant via social learning, and which have more variation in song style across individuals. In such species, numerous forebrain nuclei and interconnections also play supportive roles both in learning and in the production of songs. Included among these structures are auditory areas critical for "recording" the songs sung by nearby adult singers and for remembering the perceptual models against which to compare one's own singing. In addition, there are a handful of more anteriorly located forebrain neostriatal nuclei (e.g., area X, MAN, Nif) that play critical roles in learning to replicate the acquired song. Finally, a premotor structure, the caudal region of the hyperstriatum vertrale (Hvc) plays a crucial role in both integrating these influences and contributing higher-order regulatory control over RA and thus critically important in song complexity and flexibility (see figure 8).

Figure 8. Comparison of the general organization of the songbird brain singing circuits in species/breeds that do not learn their songs (left) and that acquire songs through social learning (right). Red arrows indicate efferent circuits Birds that have strict innate song templates (like the Munia) have minimal forebrain involvement in song control, but birds that acquire their song socially (like the Bengalese Finch) utilize numerous forebrain nuclei and large numbers of fiber tracts. Note especially the relatively greater involvement of association connections (green lines).

The Bengalese Finch and White-Rump Munia fit on either end of this spectrum. Whereas Munia song is only disrupted by damage to the RA nucleus in the forebrain, Bengalese Finch song is drastically simplified by damage to Hvc, resulting in a highly 25

stereotypic song. And song learning is significantly impaired by damage to auditory or rostral striatal regions of the forebrain. While it is not difficult to imagine how domestication might produce increased behavioral variability, it is less obvious how the effects of drift could explain this increase in the complexity of neural control. Perhaps one reason that this seems counterintuitive is that we tend to think of increasing complexity as the addition of something. Thus the above difference could be described as adding each of these new structures to the song control system of the bird brain. However, this is misleading in two senses. First, these additional structures function in multiple contexts in every songbird brain, irrespective of the kind of song, and connections between them, some of which include RA, exist even in species where they play no role in singing. They are not specifically limited to "song circuits," but are to some extent specialized to each provide a distinct computational contribution to the learning, planning, motor coordination, perceptual memory, etc., involved in many different behaviors. Second (assuming that bird and mammal brains involve roughly analogous developmental processes) in the immature brain there is likely to be less exclusive specificity of connectivity between these and other forebrain nuclei, which is only later reduced and/or limited in its influence. So we can also describe this effect of domestication as disinhibiting the influence of these connections, and/or less effectively and less selectively culling the diversity of immature connections. Additionally, we might expect that the circuitry within RA of the domesticated finch is also less differentiated and thus less constraining of motor output. In other words, drift should allow a slight degeneration of the genetic specificity of certain features of brain development, and the result would not so much appear as loss of function, but rather degeneracy leading to reduction of constraint on factors controlling production. In summary, reduced selection can be expected to produce progressive despecialization of the circuits that contribute tight constraints on motor patterning that specify song structure. With this relaxation of innate biases, other previously inhibited or ineffectual influences, contributed by other brain systems, such as the trace of early auditory experience, could begin to play a larger role in biasing song formation. Both the weakening of constraints intrinsic to RA circuits and the possible mature persistence of significant immature projection patterns from other forebrain nuclei could thus contribute to a more distributed and thus multi-faceted control of singing and song structure. This pattern of redistributed control of function among brain systems is loosely analogous to the redistribution of the control over ascorbic maintenance that resulted from degradation of the GULO gene in primate evolution (described above), in the sense that degradation of this highly constrained mechanism opened the door to many other contributions from diverse and previously irrelevant mechanisms, such as specialized digestive and visual adaptations. However, in the case of the Bengalese Finch, there is no selection to shape up these more indirect influences. Their emergent roles in song production are merely incidental to their other functional and developmental features. Nevertheless, despite the non-functional and purely serendipitous changes in song control, and the fact that they are due entirely to degenerate genetic influences and the associated de-specialization of certain epigenetic processes, the result is appropriately described as an increase in complexity for the following reasons.

1. There is a considerable increase in the number of song variables able to be modified and an expansion of the range of extrinsic factors with respect to which these song features can be adjusted. 2. The resulting behavioral function is not merely more variable, it is also more conditionally responsive to contextual influences, including social influences. 3. The prior adaptive functions of the many neural systems that ultimately came to play a role in singing provided a large and previously untapped domain of functional overlap where new interrelationships contribute considerably increased structural complexity. Of course, specificity of these prior functions and the constrained variation of song in the wild were adaptive. So the domesticated species' less precise and more flexible behaviors would likely not serve it well if released again into the same ancestral environment. In this respect, increased complexity is not necessarily an adaptive advantage, though increased flexibility might make Bengalese Finches better able to adapt to non-native environments. Of course, singing behavior is only one of many systems that is likely to be dedifferentiated by domestication. Since domestication globally reduces selection on a large number of behavioral systems and physical capacities it should be expected to lead to widespread patterns of de-differentiation involving many non-essential systems, not just the genes and corresponding brain structures involved in song control. In the Bengalese Finch these appear to extend to mate choice, nesting behavior, offspring care, and toleration of environmental disturbance, among other features (Okanoya, personal communication). Considering that the widespread genetic drift associated with long-term domestication is likely to be expressed developmentally as less differentiated adult features, it also should not surprise us that domestication tends to produce animals that appear superficially altricial in both behavioral and physiognomic characteristics (e.g. see Beylaev 1979; Trut 1999). Neural developmental degeneracy and the evolution of language The process of relaxing selection on neuro-behavoral traits in domesticated birds, described above, provides a suggestive analogy for rethinking the evolution of a similarly complexified neuro-behavioral trait in humans: language. It is almost universally assumed that the human language capacity evolved under the influence of intense selection favoring the evolution of supportive neural mechanisms. There are, however, serious difficulties in using the simple neoDarwinian paradigm to explain some of the more enigmatic features of this uniquely human adaptation. First there are the many functional discontinuities with respect to other species' vocal and behavioral communication systems. Second, there is the unprecedented complexity of language and the apparent miracle of rapid language acquisition in infancy. And third, there is the extensive synergy of diverse brain systems involved in the comprehension and production of language. These curious attributes have frustrated attempts to move beyond highly generic claims about selection for "increased neural complexity" or appeals to "big bang" lucky mutations that simply posit the introduction of these unprecedented faculties. But reflecting on these challenges in the context of the "lazy gene" hypothesis provides many striking parallels. Although it seems undeniable that many aspects of the language adaptation have been honed by natural selection, there are many other aspects

involving individual and cultural variability, flexibility of form and expression, complexity of behavioral output, dependence on social learning, and highly distributed synergistic neurological organization, that are reminiscent of the relaxed selection effects described above for Bengalese Finch singing. The most obvious place to start to explore these parallels is to ask whether there are aspects of the human language adaptation that are similar to the reorganization of song control in the Bengalese Finch. Potential examples might include the following features of language and its correlates: First, there appears to have been significant loss, simplification, and co-option of many species typical stereotypic calls, which are characteristic of the communication of our close relatives the chimpanzees. With the exception of laughter, sobbing, and shrieks of fright, we inherit a very limited repertoire of this sort of prespecified vocalizations specialized for distinct social messages or objects in our environment. Since stereotypic calls are neither produced by prelinguistic infants nor in cases of global aphasia, it is unlikely that they are merely suppressed or superseded by language use. More likely their neural substrates have been subject to degeneration, perhaps analogous to the degeneration of song stereotypy in the Bengalese Finch. There has also been an analogous degeneration of the vocalization transition biases that are typical of calls (e.g., as is still present in the stereotypic 'ha-ha-ha" of laughter), allowing spoken language to take advantage of the ability to combine almost any oralvocal articulations that are mechanically possible for the production of words and sentences. There has been a corresponding decoupling of the correlations between particular vocalizations and arousal states that are also be characteristic of stereotypic calls. So except for the few preserved human calls, there is effectively a completely arbitrary link between any specific vocalization and emotional state. What might be termed "evolutionary disinhibition" is dramatically exemplified by a related human-unique behavior: infant babbling (See Goldstein and Schwade, this volume). In the context of other species' vocal predispositions, infant babbling can be compared to a neurological disinhibition effect. A young child begins to babble in low arousal contexts, when stereotypic calls like crying are unlikely, and prior to its use for communication. The organization of babbled "syllables" is largely unconstrained except perhaps by mechanical limitations, though later it tends to show mimicry of the tonal characteristic of the surrounding speakers' language. Finally, some features characteristic of other species' call mechanisms seem to be incorporated into speech, but in quite a different way. Speech prosody seems to co-opt many of these call arousal correlations, subordinated to the articulation of speech phonemes. The prosody of speech appears to have borrowed many of the tonal and rhythmic features that are associated with specific emotional states in ways analogous to these features in the calls of other species. So in human speech we communicate our emotional tone, our interest, and our attention to things we consider important using generic vocal changes that are more universal than any sound-meaning couplings in language. So not only has there been partial degeneracy of this call system but, again analogous to the new conditionality of vocalization in the Bengalese Finch example, this has enabled two independent mechanisms to work in concert.

Not only does language require vocal flexibility and decoupling of specific vocalizations from specific cognitive-emotional states, but it also requires coupling with auditory analysis and memory and with motor skill learning systems. The functional linkage of these otherwise largely autonomous neural subsystems in the evolution of the human brain is also paralleled by the effects of relaxed selection on the brain systems controlling song in the Bengalese Finch. The neuroanatomical contrast between songlearners' brains and stereotypic singers among songbirds is, in many respects, strikingly analogous (perhaps even in the evolutionary sense) to the neuroanatomical differences in the substrates for stereotypic calls and language (see figure 9).

Figure 9. Comparison of the general organization of neural circuits controlling stereotypic calls (e.g. laughing, sobbing) and language in humans. Analogous to the Munia-Bengalese finch comparison, the difference shows minimal forebrain involvement (e.g. few cortical areas are involved in innate call production), but highly distributed forebrain (especially cerebral cortex) involvement in language processes. The depiction of language circuits is extremely simplified and only show major cortico-cortical pathways, though cortico-subcortical connections are also extensive. Color coding as in Figure 8.

A generation of using lesion and stimulation techniques to study call production systems in primate brains has mapped out an elaborate limbic-midbrain system, in which distinct basal and midline forebrain structures and pathways associated with specific drives and arousal states are also centers from which electrical stimulation can induce the production of stereotyipic calls typically associated with those states. Cerebral cortical sites associated with call production are limited almost exclusively to the anterior midline, in peri-limbic regions such as the anterior cingulate cortex (Jürgens 1979). Motor and premotor cortex appears uninvolved (Jürgens et al. 1982). In general, these forebrain sites are few, overlapping, and highly interconnected. In contrast, language is significantly dependent on cerebral cortical systems that are widely separated between frontal motor and premotor, temporal auditory, parietal multi-sensory, and prefrontal executive systems, as well as with striatal and thalamic structures to which each of these cortical regions are linked (Deacon, 1997). Even the cerebellum appears to play a role in cognitive aspects of language processing. Thus language is characterized by a largely independent and quite widely distributed network of forebrain systems, which in other

species do not share this kind of close synergistic contribution to a single cognitivebehavioral function. Of course one of the most significant features of language, also paralleled by the Bengalese Finch example, is its critical dependence on social transmission. The level of this social influence of the local speech environment on human neurological function is unparalleled in any other species. Although we take this for granted, it is hard to overemphasize the extent to which the epigenetic honing of this function has been "offloaded" to this extrinsic source of structure. Like the addiction to ascorbic acid, this has almost certainly had a secondary reorganizing effect on many other systems, which now serve to help stabilize and maintain this extrinsic epigenetic information source. It is perhaps the most extreme development of the trend in large brained species to rely on extrinsically driven activity-dependent shaping of neural circuits to complete neural circuit specification. Unlike the Bengalese Finch example, however, the story of human brain and language evolution is not merely one of relaxed selection. Relaxed selection may have played an important role in allowing complex functional synergies to emerge, both among brain systems now contributing to language and between these brain systems and socially transmitted information. But once these capacities became available and the reproductive advantages of symbolically transmitted knowledge and social organization became reliable there would certainly have been selection favoring not only the stabilization of this unprecedented neural synergy but for optimizing its functionality (Deacon 1997). So while we may be able to attribute many features of this unprecedented capacity to relaxation of selection and the complexity afforded by degeneracy of previously more autonomously specialized systems, we need to also consider how this may have unmasked selection on other neurological features.
Conclusions: This analysis has suggested that certain critical human language adaptations may be consequences of the evolutionary degradation of genetic and neurological substrates due to the removal of selective pressures. This appears to be the opposite of the view of most students of human mental evolution who hold that our special propensity for language can be understood as an 'adaptation,' honed by natural selection. What I have loosely called the lazy gene effect is merely an amplification of the evodevo paradigm for integrating the formative influences contributed by epigenetic mechanisms with those that can be ascribed to natural selection. It can be seen as the exception that proves the rule that natural selection and epigenetic processes both contribute to the form that adaptations can take. Specifically, the reduction of natural selection is shown to release what might be described as "latent" potentials implicit in the epigenetic mechanisms that translate genetic information into phenotypic consequences. These mechanisms should be quite general, applying to many cases of the evolution of complex functional synergies at all levels of scale. The interplay between the genetically degenerate effects of relaxed selection, which can expose novel epigenetic synergies, and the sculpting effects of selection, which can preserve useful emergent consequences, can thus be recognized as a major engine driving the evolution of functional complexity. This is particularly relevant to the evolution of nervous system complexity because of the potential for brain

development to entrain many levels of morphological biases, including those imposed from environmental and social sources. The neurological specializations for language that aid in its transmission, acquisition, and use exhibit many features that may be better understood in these non-Darwinian terms. The advantage of considering this possibility is that it refocuses attention on the multitude of processes involved in neural development that can potentially provide sources of information contributing to language organization, and the particular processes by which language competence is acquired. From this vantage point, many converging factors appear likely to have played a formative role in the emergence of this unique social-cognitive-biological phenomenon. By turning our attention away from the search for hypothetical selection pressures favoring the evolution of language, to the exploration of epigenetic mechanisms that were modified and coöpted for this novel function we can avoid many of the pitfalls of inventing evolutionary scenarios to explain the many complex idiosyncrasies of language. More importantly it opens up a whole new domain of experimental inquiry into what was once considered terra incognita. For linguists the lazy gene approach to language adaptation suggests some quite counterintuitive hypotheses. Languages are vast systems of socially transmitted and maintained behavioral algorithms. In this respect, there is an enormous capacity for masking effects. To whatever extent that intrinsic semiotic constraints, neural processing limitations, and social transmission effects might contribute emergent self-organizing and selection effects influencing language structure, we should expect them to impede the evolution of corresponding innate language universals. This will be difficult news for theorists who have argued that innate grammatical biases could have arisen due to the Baldwin Effect (e.g. Pinker 1994), because as we've seen the reverse effect is more likely. However, as highly canalized adaptations for social communication become increasingly degraded in the context of socially transmitted language, the increasing dependence on social transmission will "unmask" selection on any mental capacities that enhance social transmission. This will have the effect of redistributing selection to diverse genetic loci and extrinsic factors, in much the same way as happened with the loss of ascorbic acid synthesis capacity. The inevitable diversity, indirectness, and combinatorial synergies of such acquisition and transmission biases will make this a far more biologically messy explanation of language universals than postulating an innate language faculty, but such highly distributed control of language genesis may be the explanation for its remarkable robustness despite its unprecedented complexity and flexibility. Finally, for paleontologists and anthropologists, this analysis brings our attention to one final question about our genome and our brains in general. Are we genetically augmented apes with numerous genetic and neural improvements that make us better than our cousins? Or is it more parsimonious to recognize that many of these presumed "adaptations" are better understood as consequences of degenerative processes due to a reduction of selection pressures? Almost certainly, the unique combination of traits that support our language capacities and our social dependence are a consequence of both kinds of influences. We are in many regards a self-domesticated species. Would it be too humbling to see ourselves as a somewhat genetically degenerate, neurologically de-differentiated ape? Reframing humanness in biologically degenerate terms is not, as we have seen, to deny

that we are in many respects more complex, both neurologically and behaviorally than other ape species. Moreover, the dedifferentiating effects of domestication may explain certain other enigmatic features of human nature, such as our cultural variability and even our fascination with art and music. Perhaps our great leap forward required our taking first a few steps back.