

# 1

## The evolution, development, and modification of behavior

Organisms are machines designed by their evolution to play a certain role. The role, and the stage – the environment – where it is played, is called the organism’s *niche*. For example, most cats – tigers, leopards, mountain lions – play the role of solitary hunters. Wolves and wild dogs are social hunters; antelope are social grazers; and so on. The basis for the modern idea of niche is Charles Darwin’s discussion of an organism’s “place in the economy of nature.”<sup>1</sup>

A niche defines the pattern of behavior – the *adaptive repertoire* – compatible with an organism’s survival and reproduction. A niche doesn’t tell an organism how to behave. It just punishes it – by death or reproductive failure – for doing the wrong thing. A niche is a filter not a creator. Niches are best defined by example. It is pretty obvious that the talents required of a good leopard are quite different from those needed by an effective antelope. For leopards, powerful means of attack, a digestive system attuned to meat, and a visual system adapted to attend to one thing at a time work well. But a prey animal like an antelope needs a good way to evade attack, a lengthy gut able to cope with poor herbivore diet, and a visual system able to detect threat from any quarter. Hence, the claws and teeth of the leopard, its forward-facing eyes and short digestive tract, as well as the rapid running and maneuvering of the antelope, its lengthy gut, and sideways-facing eyes – all have an obvious *functional* explanation.

The behavioral adaptations required by different niches are usually less obvious than morphological (form) differences, especially if they involve the ways that past experience affects present potential, that is, differences in *learning*. The match between adaptation and niche is no less close because it is hard to see, however.

For simple niches, such as those filled by most nonsocial invertebrates, a set of built-in responses to commonly encountered environments suffices to get the organism to its evolutionary goal, which, for bugs as much as billionaires, is survival and reproduction (Darwinian *fitness*). The animal need only avoid bad things and approach good ones, all signified by signals innately coded. Stimulus–response mechanisms, plus some sensitivity

<sup>1</sup> Darwin, Charles. *On the origin of species* (1859 and five subsequent editions). London: John Murray, 1859; many internet versions are available. See also Wallace, Alfred Russel. *Contributions to the theory of natural selection* (Google Books) (2nd edn.). Macmillan and Company, 1870.

to rates of change, are sufficient for a wide range of surprisingly intelligent behavior. I discuss adaptive mechanisms that require little or no dependence on history in Chapters 2, 3, and 4.

As the niche grows more complex, adaptive behavior depends more and more on the animal's past. Simple mechanisms are still needed, no matter how complex the niche – even human beings need reflexes, for example. But, in addition, more complex, history-dependent processes are required. Greater flexibility carries with it two kinds of cost: First, the animal must *have* a past if its behavior is to be guided by it. This implies a lengthening of infancy and adolescence – delayed reproductive maturity. Waiting to breed puts the individual at a reproductive-fitness disadvantage compared to others quicker on the draw. It is sometimes better to be quick and dumb rather than intelligent and slow.

Second, there is a growing bookkeeping cost. The behaviors acquired through past experience, and some “representation” of the environments in which they are appropriate, must be “stored” in such a way that the animal has ready access to the most appropriate action. The scare quotes are a warning. An organism's history may affect future behavior via processes that little resemble human storage systems, be they books, photos, files, rolodexes, or random-access memories.

Representing data in the most flexible and economical way is a problem that also confronts human data systems. Much work in computer science is concerned with “database management,” as this is termed. Still-evolving search engines like Google and Bing daily demonstrate the enormous power of efficient data management. Early learning theories underestimated the information-processing task implied by the behavior even of plants and insects, much less mammals and birds.

Situations rarely recur in precisely the same form, but only some of the differences are important for action. Hence, the animal's representation of past environments must also allow it to behave appropriately in environments similar to those it has already encountered. Just what *similar* means, and how it is determined both by the animal's evolutionary history and its own experience, is one of the most intriguing questions in animal behavior. These issues are taken up in Chapter 11.

An animal's past experience can affect its future in a variety of ways. The simplest way to make sense of these is the conventional dichotomy between *learned* and *innate* behavior. Innate behavior is completely independent of experience, and learned behavior is, well, learned. Of course, nothing is truly innate, in the sense of being independent of any experience, but many things are almost independent of any particular kind of experience. For example, many small invertebrates avoid light; they need no special training, no nasty shock in a lighted place, to show this pattern. Most mammalian reflexes are like this: As soon as an infant can move at all, it will automatically withdraw its hand from a pinprick. The knee jerk to a tap, pupillary contraction to a bright light, and many other reflexes all develop in a variety of environments, common to all normal members of the human species. I discuss reflexes in Chapter 3.

But there are many effects of experience that do not fit the innate–learned dichotomy. For example, age slows responses and stiffens joints, fatigue weakens muscles, hunger

(food deprivation) and thirst change preferences in systematic, reversible ways, and so on; a number of other, developmental effects will be discussed shortly. None of these corresponds to the usual meaning of the term learning, which refers to a more specific and only partly reversible change, most easily seen in relation to a positive or negative outcome: The animal learns where food is to be found or to avoid the predator. But learning also occurs even when no obvious reinforcement is involved, although it is harder to detect than when it is tied to something that can be presented or removed.

This book is primarily concerned with learning in this broad sense, but the category is not exact – simply because we do not really know what learning is. *Learning is not a technical term.* The more subtle and complex the task, the more likely we are to call it “learning” rather than “reflex,” “habituation,” or “instinct.” There is probably no single process that underlies learning in this sense. Experience can change behavior in many ways that manifestly do not involve learning, as well as in ways where we are not sure. In other words, there is no hard-and-fast line separating learning from other kinds of behavioral change. There is no neat dichotomy between “learned versus innate” behavior; rather, there is a spectrum of ways in which past experience affects future behavior.

### Canalization

The innate–learned dichotomy nevertheless refers to a useful distinction better expressed as *canalization*.<sup>2</sup> A structure or behavior is said to be canalized if its development is almost independent of a particular experience or environment. Features like the four-chambered heart of mammals, or bilateral symmetry – or most reflexes – are strongly canalized, in the sense that just about any environment that allows the organism to survive at all will allow them to develop. A trait such as competence in the English language, or the ability to do algebra, is not canalized at all, because it is critically dependent on a particular history. Competence in *some* language is an intermediate case: Evidently just about any linguistic environment is sufficient to produce language learning in a normal human infant, even without explicit instruction. In a similar way, male chaffinches and white-crowned sparrows will develop some adult song if they can listen to a model at the critical time in their first year of life. But the kind of song they develop depends on the model, as well as the species. Language and song development are canalized, but the particular song or language to be learned is not. A language may even develop on its own. Striking proof of this appeared a few years ago in a school for the deaf in Nicaragua.<sup>3</sup> Taught in school only to lip read, the pupils nevertheless developed entirely on their own a highly structured sign language. The icons increasingly found in digital devices to signal various “apps” and services may perhaps evolve into a kind of *lingua digitalia* common to all users.

<sup>2</sup> The term canalization was coined by British biologist Conrad Waddington (1905–75) in 1942. A review of work up to 2005 is Flatt, T. The evolutionary genetics of canalization. *The Quarterly Review of Biology*, September 2005, 80, 3.

<sup>3</sup> A linguistic Big Bang, by Lawrence Osborne. *New York Times*, October 24, 1999.

What an animal learns, and the way that it learns it, is much affected by its niche. Because niches differ in many respects, so do learning mechanisms. Since niches do not differ in every respect, there are also similarities among learning mechanisms.

Space and time are common to all niches. In consequence, a wide range of animal species adapt to the temporal and spatial properties of the environment in similar ways. There are also some general rules that apply across niches: Old information is generally less useful than new information; consequently animals forget, and they forget less about things they have learned recently. Conversely, the environment of an animal around the time of birth usually has a special significance, and things learned at that time may be especially resistant to change. Experience on first exposure to a new environment is also likely to be especially well remembered. Food, water, sex, and habitat are vitally important to all species. Hence these things are better remembered than neutral events and have special properties as guides of behavior.

This book is mainly concerned with the way that animals adapt to these things that are common to all niches. The major emphasis is on adaptation to rewards and punishments.

Students of animal behavior – *ethologists* and comparative psychologists – tend to be interested in features of adaptive behavior that differ among niches. Learning psychologists are more interested in behavioral mechanisms that are common to all niches. In years past psychologists rallied around the search for “general laws of learning.”<sup>4</sup> The discovery of types of learning specific to particular situations or species gradually made this position untenable. There *are* general laws, but they seem to reflect commonalities among niches or general features of all information-processing systems, rather than a common plan of construction – as the earlier view implied. In biological terms, the resemblances are a mixture of *convergence* and *homology* rather than the pure homology implied by a general law.

### Explanation

Animals and people seem to have purposes, beliefs, attitudes, and desires; they seem to know some things and not others, to want some things and reject others, and so on. They are what philosophers call *intentional systems*.<sup>5</sup> Intentionality may seem to set psychology apart from the physical and biological sciences. After all, the chemist does not worry about the beliefs of his compounds nor is the physicist concerned about the quirks of quarks. Does this mean that psychology is not scientific? Does it mean that it is different in kind from the physical sciences? Not at all; the difference is in the richness of behavior of the things studied, their sensitivity to their environment, and the dependence of present behavior on past experience. The language of intentionality is simply the everyday way that we deal

<sup>4</sup> See, for example, Bower, G. H., & Hilgard, E. R. *Theories of learning* (5th edn.). Englewood Cliffs, NJ: Prentice-Hall, 1981; Dickinson, A. *Contemporary animal learning theory*. Cambridge, UK/New York: Cambridge University Press, 1980. A classic account is Estes, W. K., Koch, S., MacCorquodale, K., Meehl, P. E., Mueller, C. G., Schoenfeld, W. N., & Verplanck, W. S. *Modern learning theory*. New York: Appleton-Century-Crofts, 1954.

<sup>5</sup> Dennett, D. Intentional systems. *Journal of Philosophy*, February 25, 1971, LXVIII.

with complex historical systems. I typed the first edition of this book with the aid of a microcomputer (remember those?) that had a primitive operating system called “CP/M.” Look how the instruction manual (remember *those!*) refers to CP/M and its associated programs: “CP/M could not find a disk . . .” “PIP assumes that a . . . character” “CP/M does not know that . . .” “Seven commands are recognized by CP/M.” No one assumes that there is a little man or woman, complete with “real” knowledge, beliefs, desires, and understanding, inhabiting the microchips. Anything that responds to varied stimuli in varied ways, especially if its behavior depends upon past history and appears goal-oriented, is understood at a commonsense level in intentional terms.

Chess-playing programs offer examples. A good one elicits precisely the same kinds of comment we would use for a human player: “It is attacking the queen,” “It’s trying to get control of the center of the board,” and so on. Yet no one doubts that the underlying program provides a perfectly precise and mechanical account of the machine’s behavior.

So are there really such “things” as beliefs, desires, and so forth? If so, do machines possess them? There are two schools of thought on these questions: The first answers “yes” to the first question – “beliefs,” etc. are real; and “no,” or at least “probably not” – machines don’t have them. But there is a more scientifically useful view, which is that these questions are unhelpful, not to say irrelevant.

Why? Attitudes and beliefs are often thought to cause actions: People do what they do *because* they believe what they believe. This approach is inapplicable to the behavior of animals, of course (does the spider curl up at your touch because he’s afraid?). But it also has both experimental and theoretical limitations. The experimental problem derives from the difficulty of separating correlation from causation. This is an old question:<sup>6</sup> Do we run because we are afraid, or are we afraid because we run? “Fear” is a property of the subject’s internal state. It is inferred from his<sup>7</sup> behavior, but it is not something external that the experimenter can manipulate directly. Consequently, one can never be certain that the running and the fear are not both simultaneously caused by something else (like some idiot waving a gun). The problem is not insuperable. There are ways that intentional terms like “fear,” “hope,” and so on can be made methodologically respectable and tied to observables.

But the scientific/theoretical question is whether the labor involved is worth it. The whole enterprise rests on the presupposition that familiar intentional terms such as “fear,” “belief,” “attitude,” and the like form the very best basis for theoretical psychology. Yet clever computer programs show that such terms represent a primitive kind of explanation at best. They enable someone ignorant of the details of the program to make some sort of sense of what the machine is doing. But full understanding rarely reveals anything in the code that corresponds directly to intentional terms, useful though they may be in the absence of anything better. It is rash, therefore, to base a program of psychological research on the assumption that intentional terms represent fundamental causal processes.

<sup>6</sup> See the James–Lange theory: [http://en.wikipedia.org/wiki/James%E2%80%93Lange\\_theory](http://en.wikipedia.org/wiki/James%E2%80%93Lange_theory).

<sup>7</sup> I always use “he” and “his” in a generic sense, preferring this usage to the cumbersome “he/she” or the politically inspired “she,” for reasons of euphony and historical continuity.

Thus, my answer to the question “Are there really such things as beliefs, desires, and so on?” is “Maybe . . . but are they scientifically useful?” Obviously these terms are helpful ways of coping with some complex systems. But the “really” question is metaphysical. There is no reason to suppose that these terms will prove especially useful in unraveling the causes – mechanisms – of behavior, which are what we are really interested in. Hence, I will be concerned with the behaviors of people and animals, measured pretty much in physical terms – that is, with a minimum of interpretation. This is *methodological behaviorism*, which is usually contrasted with *radical behaviorism*, the position advocated most forcefully by the hugely influential B. F. Skinner (1904–90).<sup>8</sup> Radical behaviorism asserts that it is unnecessary to go significantly beyond the level of behavioral description to account for all behavior. The position made some sense in reaction against a dominant mentalism, but makes none now. To pursue the computer analogy, it is like asserting that the chess-playing program can be explained entirely in terms of its inputs and outputs and direct (stimulus–response) links between them with no contribution from the lines of code that reside in its memory and link input and output.

The approach that underlies this book is called *theoretical behaviorism* because theories – mechanisms, models – are generally what we will be looking for. It is not actively hostile to intentional terms; it’s simply agnostic. It does not presume that psychological theory should be based on intentional language.

Of course, human psychology must eventually come up with an explanation for why intentional terms are so useful and ubiquitous as makeshift explanations. Perhaps the answer is that intentional accounts are just fuzzy functional explanations, and thus the best that one can do without detailed knowledge of behavioral mechanisms. More on function next.

### ***Function, causation, and teleonomy***

The ultimate explanation of the chess-playing program is of course the program itself: the individual instructions that determine each move as a function of prior moves by both players. But it is usually convenient, when designing such a program as well as when trying to understand it, to divide it into two parts: a part that generates potential moves, and a part that evaluates each move in terms of a set of criteria. The dichotomy between *variation* and *selection* was proposed by Darwin and Wallace as part of their theory of evolution by natural selection, but the distinction is more general: All adaptive, purposive behavior can be analyzed in this way. The dichotomy leads to two kinds of explanation for adaptive behavior. *Causal* or *mechanistic* explanations define both the rules by which behaviors are generated (rules of variation) and the rules by which adaptive variants are selected (selection rules). *Functional* explanations just specify the inferred *goal* of the behavior. Mechanistic

<sup>8</sup> See, for example, Skinner, B. F. The generic nature of the concepts of stimulus and response. *Journal of General Psychology*, 1935, 12, 40–65; Skinner, B. F. *The behavior of organisms*. New York: Appleton-Century-Crofts, 1938; Skinner, B. F. Are theories of learning necessary? *Psychological Review*, 1950, 57, 193–216; Skinner, B. F. *Science and human behavior*. New York: Macmillan, 1953. See history and review in my *The New Behaviorism* (2nd edn.). New York: Psychology Press, 2014.

accounts deal only in antecedent causes; functional accounts in terms of final outcomes. Thus the form of the shark is explained functionally by its hydrodynamic efficiency, the taking by a chess program of its opponent's queen in terms of the final checkmate.

As we will see, the selection rules for learning cannot be stated as explicitly as the rule of natural selection. Indeed, even that rule is in some ways less clear than it was in the days before we were aware of the problem of the unit of selection: Individual organisms – phenotypes – succeed or fail to reproduce, but it is individual genes that are passed on. What, then, is selected? Consequently, functional explanations for adaptive behavior are often stated in terms of goals, purposes, or reinforcers (rewards and punishments), which act as guides of behavior. These notions can be formalized in terms of some kind of *optimality theory* that makes goals explicit and shows how conflicting goals are to be reconciled. The general idea is that animals act so as to maximize something that relates to inclusive fitness, such as net rate of food acquisition, number of offspring, or territory size. I return to the relation between optimality accounts and selection rules in a moment.

Functional explanations can, in principle, be reduced to mechanistic ones: Given perfect understanding of the principles of genetics and development, and complete information about evolutionary history, we can, in principle, reconstruct the process by which the shark achieved its efficient form. The biologist Colin Pittendrigh<sup>9</sup> suggested the label *teleonomic* (as opposed to *teleological*) for such accounts. Teleological explanations are problematic because they imply final causation – the shark's streamlining is teleologically explained by Mother Nature's hydrodynamic foresight. Teleonomic accounts relate form and hydrodynamics through the mechanisms of variation and natural selection. Teleonomic functional accounts are philosophically respectable; teleological ones are not. In practice, of course, the necessary detailed information about mechanisms is often lacking so that we must settle for functional accounts and hope that they are teleonomic, that a causal account will eventually appear.

Functional explanations have sometimes been criticized as being “Just-so” stories, because they are so flexible – adaptive significance, or an unsuspected reward or punishment, can be conjured up to explain almost anything. There are two answers to this criticism: Functional explanations often lead to mechanistic explanations; and functional explanations can sometimes make sense of relationships that cannot be explained in any other way.

Functional accounts are often way stations to mechanistic explanations. In studies of learning they help identify important variables and draw attention to the constraints that limit animals' ability to attain functional goals. These constraints, in turn, provide clues to underlying mechanisms. For example, mammals and birds can easily learn to use stimuli as guides to the availability of food; a hungry pigeon has no difficulty learning that a peck on a red disk yields food whereas a peck on a blue disk does not. But they are much less capable of using past stimuli as guides. In the delayed-match-to-sample task, one of two

<sup>9</sup> Pittendrigh, C. S. Adaptation, natural selection, and behavior. In A. Roe & G. G. Simpson (Eds.), *Behavior and evolution*. New Haven, CT: Yale University Press, 1958.



stimuli is briefly presented, then after some delay both are presented, and a response to the one that matches the first is rewarded. Delays of more than a few seconds between sample and choice presentations gravely impair most animals' ability to choose correctly. This is a memory constraint. Other psychological constraints have to do with animals' ability to process information, and with their perceptual and motor abilities. Identification of limitations of this sort is the first step toward understanding behavioral mechanisms.

In addition to internal (psychological) constraints, there are also constraints imposed by the environment. For example, the animal cannot do more than one thing at a time, so that total amount of activity is limited; spatial arrangements limit the order in which food sites can be visited and the time between visits. Reinforcement schedules, either natural (as in picking up grain, one peck per grain, or in natural replenishment processes) or artificial (ratio and interval schedules, for example), further constrain the distribution of activities. Functional explanations, precisely expressed in the form of optimality theory, allow, indeed force, one to take account of these external constraints.

Functional explanations do one thing that no mechanistic explanation can: They can explain similar outcomes produced by different means. For example, the eyes of vertebrates and cephalopods are very similar in many ways: Both have lenses, a retina, and some means of limiting the amount of light that can enter. This convergence cannot be explained by a common ancestry or any similarity of developmental mechanisms. The only explanation we can offer for this astonishing similarity is the common function of these organs as optical image-formers. Because convergence is such a common phenomenon in evolutionary biology, it is no wonder that functional explanations are so common and so powerful there.

Experimental psychologists have usually disdained functional accounts, although they often come in by the back door, in the form of vaguely expressed reinforcement theories. Indeed, one of our most influential figures boasts in his memoirs that in planning his major work he deliberately avoided any discussion of adaptiveness. Too bad! Looking at behavior in terms both of its adaptive (evolutionary) function and in relation to current goals (reinforcers) is useful in identifying important variables and in distinguishing environmental from psychological constraints.

The idea that organisms attain goals, either through natural selection for the best form of wing or individual reinforcement of the most effective foraging strategy, derives naturally from the selection/variation idea: A wide range of variants occurs, the best (in terms of flight efficiency or eating frequency) are preferentially selected, the next round of variants contains a few that do even better, and so on. This process will, indeed, lead to better adaptation only if three things are true: We have the selection rule right – that better fliers really have more offspring. The right variants occur. And a selected behavior must also persist. Some kind of memory is required. Memory is the ontogenetic equivalent of heredity. Heritability of fitness differences is essential to the effectiveness of natural selection. Some kind of memory for behavioral variants selected by environmental reinforcement is essential to learning. In other words, an animal may fail to behave in what seems to us the optimal fashion either if we have misread the selection rule, because it never generates the necessary



behavioral variant or because memory is lacking. The most efficient foraging strategy cannot be selected (reinforced) if it never occurs or if it occurs but is not remembered.

Memory also constrains behavior in ways that prevent animals from developing certain kinds of foraging patterns – patterns that require memorization of complicated sequences, for example.<sup>10</sup> These patterns will not be repeated, even in situations where they would be optimal. Thus, failures to optimize are, if anything, even more informative than successes, because they offer clues to the underlying behavioral mechanisms. Optimality theories to explain how animals adapt to reward and punishment are discussed in Chapter 8 et seq.

### Evolution and development

Development, from embryo to adult, of the individual organism – its *ontogeny* – is the product of past evolution and also limits future evolutionary possibilities. The earlier a feature appears during development, for example, the less likely its modification by future evolution. In recent decades, a whole new field concerned with evolution and development – *evo-devo* – has grown up to address this issue. Unlike human machines, natural machines – animals and plants – manufacture themselves.<sup>11</sup> The process of development limits their potential and often incorporates the effects of experience in ways that contrast with, and thus help define, learning.

Organisms change throughout their lifetimes, and the processes by which they change are the outcome of past evolution. As Darwin pointed out, organisms bear their evolutionary history both in their structure and in the manner of its development. Rudimentary organs provide some striking examples. The human vermiform appendix, the rudimentary breasts of male mammals, the vestigial second lung of snakes, the teeth of fetal whales that vanish in the adult, the uncut teeth of unborn calves – none has any function in the adult, yet they remain: “They may be compared with the letters in a word, still retained in the spelling, but become useless in the pronunciation, but which serve as a clue for its derivation.”<sup>12</sup> There are behavioral parallels in the useless “grass-flattening” of domestic dogs on the living-room carpet, and exaggerated fears (of the dark, or of strangers, for example) in human children. In many cases these vestigial behaviors disappear with age, as in some of Darwin’s examples.

Vestigial and transient morphological features illustrate the half-truth that ontogeny recapitulates phylogeny, that is, the idea that the stages through which an organism passes, from embryo to blastocyst to adult, represent a history of the race in abbreviated form. Gill slits in the human fetus were once taken to mean that the fetus at that stage resembles the ancient fish from which mammals are descended. The actual relations between ontogeny

<sup>10</sup> Machado, A. Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior*, 1992, 58, 241–63, showed that pigeons readily learn simple but not complex sequences.

<sup>11</sup> Many people, beginning, perhaps with Samuel Butler in an article “Darwin among the Machines,” published anonymously in *The Press* newspaper in 1863 in New Zealand and later incorporated into his prophetic book *Erewhon* (1872), have claimed to see quite exact parallels between Darwinian and machine evolution, however. See, also Staddon, J. E. R. *Adaptive dynamics: the theoretical analysis of behavior*. Cambridge, MA: Massachusetts Institute of Technology Press/Bradford, 2001, Chapter 1.

<sup>12</sup> Darwin, C. *The origin of species*. Oxford: Oxford University Press, 1951 (reprinted from the 6th edn., 1872), p. 525.

and phylogeny are more complicated of course. They derive from the fact that evolution acts via the mechanisms of development.

The German biologist Ernst Haeckel (1834–1919) was one of the first (Darwin preceded him) to point out the relation between development and evolution, although his views are in many ways too simple and lend themselves to the “ontogeny recapitulates phylogeny” parody. S. J. Gould summarized the modern view as follows: “Evolution occurs when ontogeny is altered in one of two ways: when new characters are introduced at any stage of development with varying effects upon subsequent stages, or when characters already present undergo changes in developmental timing.”<sup>13</sup>

Development can be compared to a railroad switchyard in which incoming cars on a single track are sorted by a branching arrangement of switchoffs so that each car arrives at a different destination. Soon after conception the organism is essentially undifferentiated and “pluripotent,” that is, many things are possible (these are the “stem cells” that have played a role in recent political controversy about research ethics). With progressive cell divisions, there is increasing differentiation and the options for further development are reduced: The railroad car has passed through several switchoffs and is closer to its final destination. This process of progressively finer differentiation, and the concomitant reduction in future options, takes place throughout life. Eventually the car enters the final stretch of track that terminates in death – which is not so much a wearing out as the largely predetermined end of a course charted by prior evolution. Typical life-span, like other characteristics, is determined by its costs and benefits, weighed in the delicate balance of natural selection.

Genes determine the direction of the successive switches that occur throughout ontogeny. We don’t yet know exactly how this works. Nevertheless, one thing is clear: The genetic changes that provide the raw material for evolution act not directly on morphology or behavior, but on the course of development – a stage may be added or missed entirely, stages may be accelerated or retarded.

These changes in the path of development are the raw material for the formation of new species. For example, if the reproductive system matures relatively faster than the rest of the body, the result may be a sexually mature “larval” animal, as in the case of the Mexican axolotl (*Ambystoma tigrinum*), a salamander that can become sexually mature while still a tadpole. Continued selection might well fix a change of this sort, so that the terrestrial stage is completely abolished and a new species of entirely aquatic amphibian is the result.

It is easy to see that this process will leave traces of a species’ past evolutionary history along the path of development of an individual organism. For example, a fish such as the angelfish (*Pterophyllum scalare*), is strongly laterally compressed in the adult, while a flatfish such as the flounder (*Bothus lunatus*), is vertically compressed and has lost bilateral symmetry by having both eyes on the same side of the head. Nevertheless, the immature forms of both species are quite normal looking, with the “typical” elongated, bilaterally symmetrical fish shape. Presumably the abnormal body form arose via genetic changes that acted to modify growth gradients in the ancestral species at a relatively late stage of

<sup>13</sup> Gould, S. J. *Ontogeny and phylogeny*. Cambridge, MA: Harvard University Press, 1977, p. 4.