
Understanding Adaptive Behavior and Thought as Purposeful Evolution: Combining Bernard and Darwin

It is wonderful what the principle of selection by man, that is the picking out of individuals with any desired quality, and breeding from them, and again picking out, can do. Even breeders have been astounded at their own results. . . . Man, by his power of accumulating variations, adapts living beings to his wants—may be said to make the wool of one sheep good for carpets, of another for cloth, &c.

—Charles Darwin (from an 1858 letter to American biologist Asa Gray; reprinted in Bajema 1983, pp. 191–192)

Three Lessons of Biology for Behavioral Science

A major theme of this book as elaborated in the preceding chapters can be summarized by extracting a few important lessons about what biology has taught us about the what, why, and how of animate behavior.

The lesson inspired by Claude Bernard and introduced in chapter 4 is that the functioning of physiological systems can be understood as the means by which an organism controls its internal environment. But since physiological control is achieved by internal processes normally hidden from view, this lesson is more relevant to physiology and medicine than it is to behavioral science.

Instead, it is the *extended* Bernardian lesson that makes sense of observable behavior, grounded on Bernard's basic insight, further developed in the mid-twentieth century by cyberneticians, and systematized into a unified working theory of animate behavior by William T. Powers and his associates. Presented in chapters 5 and 6, the extended Bernardian lesson informs us that a living organism acts to control aspects of its external environment. And since an organism can know its environment only through its perceptual systems (including vision, hearing, touch, and other

sensory modalities), animate behavior can be understood as the *control of perception*. The extended Bernardian lesson is concerned with the proximate (here and now) causes of behavior, and when augmented by perceptual control theory it establishes working models of behavior that are both physical and purposeful.

But controlling a perception requires the existence of an *intended* perception, that is, a goal, standard, or reference level with which to compare perception. This is where the basic Darwinian lesson becomes relevant. Chapters 7 and 8 informed us that the goals an organism pursues are not chosen at random. Neither are they determined in any direct, one-way causal manner by the organism's current environment. Instead, an organism's basic goals were selected during its evolutionary past for the effects they had on survival and reproductive success. It is not just a lucky coincidence that a male robin does all it can to maximize its distance from hawks while minimizing its distance from earthworms and female robins, since previous robins that didn't do likewise left few if any descendants. Human behavior is much more complex than that of other animals. But there are nonetheless good Darwinian reasons why men are much more inclined than women toward casual sex with a variety of partners, and why fast-food restaurants are able to attract millions of paying customers with their offerings of quick and conveniently packaged sugar, fat, and salt. Where the extended Bernardian lesson is concerned with the proximate causes of behavior, the basic Darwinian lesson has to do with the ultimate, evolutionary causes of behavior.

The third lesson of biology for behavioral science is the extended Darwinian lesson, presented in chapter 9. The basic Darwinian lesson draws its explanatory power from the cumulative variation and selection of organisms over long periods of phylogenetic time, resulting in the evolution of adaptive structures and behaviors. In contrast, the extended Darwinian lesson points out processes of cumulative variation and selection occurring *within* organisms over the much shorter span of their lives. Although the best currently understood example of cumulative within-organism variation and selection is the functioning of the mammalian immune system, growing evidence suggests that the brain also employs cumulative variation and selection to arrive at creative thoughts, innovative behaviors, and problem solutions.

How Evolution Can Be Purposeful

All three lessons (extended Bernardian, basic Darwinian, and extended Darwinian) are essential for making sense of animate behavior. But by combining the extended Bernardian and extended Darwinian lessons we gain special insight into why and how the goals of an organism change during its lifetime and how an organism is able to adapt its perceptual-behavioral systems to achieve these new goals.

Last month a teenage boy was spending several hours a week in the gym trying to get in shape for the upcoming football season; today he no longer pumps iron but spends hours with his guitar. Last year a middle-aged housewife was content to remain at home performing domestic tasks, but now works long days selling real estate and has developed impressive computing, financial, and interpersonal skills that she did not have before. These changes in goals and abilities, resulting from the process we referred to as reorganization, require a directed, purposeful Darwinian process involving the cumulative variation and selection of lower-level goals to achieve higher-level ones. This combined lesson includes aspects of Bernard's and Darwin's insights and involves proximate and ultimate causes of behavior.

But to refer to a Darwinian process as "purposeful" or "directed" might seem to indicate a basic misunderstanding of the process itself. After all, Darwin proposed his theory of evolution by natural selection to explain how species could change over time and new ones appear without the involvement of a supernatural designer or preordained cosmic plan. This is why Richard Dawkins (1986, p. 5) described biological evolution as a blind watchmaker:

All appearances to the contrary, the only watchmaker in nature is the blind forces of physics, albeit deployed in a very special way. A true watchmaker has foresight: he designs his cogs and springs, and plans their interconnections, with a future purpose in his mind's eye. Natural selection, the blind, unconscious, automatic process which Darwin discovered, and which we now know is the explanation for the existence and apparently purposeful form of all life, has no purpose in mind. It has no mind and no mind's eye. It does not plan for the future. It has no vision, no foresight, no sight at all. If it can be said to play the role of watchmaker in nature, it is the *blind* watchmaker.

Biological evolution may have no goal or ultimate purpose and in this sense it *is* blind. But this does not mean that cumulative variation and

selection cannot be used by organisms in purposeful ways. It could even be argued (as some do) that natural selection was (and still is) God's way of creating and modifying life on our planet. A truly omniscient God would be able to foresee the organisms that would evolve from such a process even if we (and Richard Dawkins) cannot, making the emergence of our species part of the God's overall plan. Still, the great strength of Darwin's theory (and what makes it a scientific theory) is that it provides an explanation for life in all its diverse forms without requiring the involvement of any such supernatural designer or the occurrence of miracles.

But it doesn't require a god to use the Darwinian process in a purposeful way. In fact, even one of the simplest forms of life is able to do so. *Escherichia coli* is a bacterium that lives in a liquid environment (such as the contents of your stomach) and can either swim in a more or less straight line or tumble randomly in one spot. If it senses that it is getting closer to food it will continue on its straight course. But if it finds that it is not getting closer to food, it will stop, tumble a while, and head off in a new, randomly generated direction. If the new heading brings the bacterium closer to food it will continue on this course; but it will stop and tumble again if the direction turns out to be no better than the previous one. Although this method of locomotion may initially appear quite crude, it turns out to be a remarkably adept and virtually foolproof way for the bacterium to get where it needs to go (see Koshland 1980, pp. 14–15). The reader can see just how effective it can be by trying out the *E. coli* program for either IBM-compatible or Macintosh personal computers available on the Web at www.uiuc.edu/ph/www/g-czikotwd/.

E. coli's method of locomotion is of particular interest as an example of one form of purposeful evolution. What evolves in this sense is not a new organism but rather a sequence of swimming directions that is effective in leading the bacterium to food. When the heading is not taking it closer to a food source, it has no clue which way to turn since it has no sense of vision or other means of determining the location of food at a distance. So it simply varies its orientation randomly and tries a new heading. Although it has no guarantee that the new direction will be any better than the previous one, if it isn't better the bacterium can try yet another one and another, until eventually it is able to home in on a meal. So by randomly varying its direction, quickly eliminating those that do not

take it closer to food and selecting those that do, *E. coli* is able to use a simple yet effective process to accomplish its goal. This is a form of purposeful behavior that combines Bernardian (control) and Darwinian (cumulative variation and selection) processes.

Not to be outdone by the lowly bacteria tumbling in our tummies, humans have also made use of various forms of purposeful evolution. One of the first was breeding plants and animals. Ever since the development of agriculture, humans have been selecting plants and animals with desirable characteristics for propagating more plants and animals. Since breeders usually have no idea what genes are responsible for the characteristics they desire in crops and livestock, all they can do is select and breed those plants or animals that are in some way better than others. Natural selection may have no purpose, but artificial selection of plants and animals involves a purposeful selector.

The last decade of the twentieth century has seen the development of some very promising high-tech forms of purposeful evolution. Computer scientists have developed a technique called *genetic programming* in which pairs of randomly generated computer programs “mate” with each other, and their resulting “offspring” (programs that resemble but are not identical to their parents) are either selected for another round of mating or eliminated according to how close they come to fulfilling the criteria of the human programmer (see Koza 1992, 1994).

In chemistry, techniques referred to as *directed molecular evolution* (Joyce 1992) and *combinatorial chemistry* (Hall 1997; Plunkett & Ellman 1997) have been developed in which a multitude of different molecules are generated and screened for desired properties, such as their ability to bind to other molecules or be biologically active in medicinally useful ways. Thus new drugs can be created using a form of purposeful variation and selection without having to know the structure of the compound or why it behaves the way it does (see Cziko 1995, chapters 13 & 14, for additional information on these and other forms of purposeful evolution).

In each of these cases, a type of directed or purposeful evolution is used to achieve a goal that cannot be achieved with already acquired knowledge. This requires a search using blind variation and selection, as you would have to do if you wanted to open a lock and possessed a large set of keys but didn't know which one fit the lock. Opening the lock is your

goal, but since you don't know which key will work, you have no choice but to proceed by the trial-and-error-elimination method of the evolutionary process. Even if you are able to eliminate certain keys that are obviously too large, too small, or of the wrong shape, you will still have to employ blind variation and selection among the remaining keys. Although I refer to this process as a form of guided or purposeful evolution, it is important to recognize that the variations (trials) generated are not guided (although they may be usefully constrained). Rather, the process is purposeful insofar as a reference level serves as a selection criterion for which certain variations (trials) are retained and others are eliminated.

As we understand the normal process of biological evolution, there is no reference level, no selection by a purposeful agent. Rather, organisms that are more successful in surviving and reproducing come to dominate their populations while those that are less successful are eventually eliminated. Darwin referred to this as *natural* selection to contrast it with the *artificial* selection made by agriculturists in selecting plants and animals for breeding. But whereas artificial selection is purposeful (believe it or not, someone really did want to produce those grotesque goldfish you can see at any pet shop, with the swollen bodies and puffy sacks for eyes, and went through a lot of trouble to do so), natural selection is not, although it resulted in the evolution of purposeful behavior, such as artificial selection performed by humans. Still, artificial (purposeful) and natural (non-purposeful) selection are similar in that the same processes of cumulative blind variation and selection (either by a purposeful agent or by inanimate physical processes) combine to generate entities that are adapted to some selection criteria (faster wild antelopes as lions and other predators eliminate slower ones; more productive domestic dairy cows as farmers purposefully breed animals that produce the most milk).

Problems of Learning

This concept of purposeful evolution based on the combination of Bernardian and Darwinian processes provides a key for understanding how it is that organisms change their behavior over time in adaptive ways, what is usually referred to as learning.

What does a theory of adaptive behavioral change have to explain? First, it must explain how an organism can come to perform an adaptive behavior that it could not do previously. This could be as simple as a rat learning to push a lever to obtain food in a Skinner box, or as complicated as a college student learning to solve differential equations. It may involve long hours of practice and gradually improving performance, such as learning to play a musical instrument or speak a foreign language. Or it may appear quite suddenly with no previous observable behavior or practice, as when someone suddenly comes up with a new idea for an invention. We will refer to this as the new knowledge problem.

Second, we must account for how it is that new behaviors can remain adaptive under changing environmental conditions. We saw in chapter 6 that these changing conditions and the new disturbances they impose mean that learning cannot be the acquisition of invariant motor responses to stimuli. Instead, an organism's actions must continually vary to bring about desired results. No matter how many times you may have driven your car from home to your place of work, you cannot make the trip using the same pattern of arm and leg movements that you used on any previous trip. Continually changing traffic, weather, and road conditions would make any such fixed pattern of actions ineffective in getting to work (not to mention dangerous if not fatal). This behavioral flexibility in the achievement of goals is not limited to humans but is characteristic of all animate behavior (recall from chapter 7 the varied behaviors undertaken by the burying beetle to bury small animal corpses on which to lay its eggs). We will refer to this as the behavioral flexibility problem.

There are two general approaches to dealing with the new knowledge problem as it relates to learning. The first is to appeal to innate knowledge as the source of what appears to be new knowledge. For example, during the first four years of life a human child makes amazing progress in acquiring the language of its caretakers. This involves learning the sounds of the language (phonology), its grammatical structure (syntax), and the meanings of words and phrases (semantics). The most widely accepted account of this remarkable feat (although one that is contested by many, including yours truly) is that this knowledge is essentially innate, or "hard-wired" into the child's brain. This innatist approach to the problem as it

applies to humans essentially denies that new knowledge is actually acquired or created by an individual, so what looks like the acquisition of new knowledge is actually the growth or maturation of old knowledge. This is essentially the position taken by the influential linguist Noam Chomsky, to whom we return in chapter 11, where we examine more closely his decidedly un-Darwinian view of innate knowledge.

The other approach is to recognize that the acquisition of genuinely new knowledge is possible and that real learning does take place. Some attempts to solve the new knowledge problem, as noted in chapter 9, made explicit use of variation and selection, but to date they appear at best incomplete and at worst misguided and misleading. Among them is Skinner's theory of operant conditioning.

As discussed in chapters 3 and 7, that theory had at least three major flaws. First, it considered animate behavior as caused by the environment instead of the means by which aspects of the environment are controlled by the behaving organism. Second, all learning involved overt responses that were varied and then selected (or not) depending on whether or not the responses were followed by a reinforcing event (such as the presentation of food). The theory thus had no room for learning based on mental or internal processes that were not accompanied by overt behavior. Third, Skinner denied that internal purposes had a real role in behavior. He therefore could not account for how organisms were able to vary their behavior to achieve repeatable effects on their environment. He did recognize the power of variation and selection to account for new knowledge as reflected in new adaptive forms of behavior. But his inability to see behavior as purposeful, and his obsession with overt behavior to the exclusion of cognitive processes, made his attempt to incorporate within-organism Darwinism into a theory of learning a rather resounding (if nonetheless quite influential) failure.

Among incomplete Darwinian approaches to learning are theories of cognitive and neural variation and selection we have seen in the previous chapter. These attempts to apply Darwinian thinking to the new knowledge problem recognize that truly new, adaptive forms of knowledge (cognitive, perceptual, and behavioral) must rely on some process of cumulative variation and selection. Cognitive theories describe the variation and selection of ideas or thought patterns, and neural theories attempt

to account for adaptive changes in the structure of the brain that are believed to underlie all forms of learning. But these theories are incomplete not only because of current limitations to our knowledge of the structure and functioning of the brain, but also because they fail to account for the purposeful nature of these Darwinian-based changes. Since all such adaptive changes allow the organism to control some aspect of its environment that it could not control before (or at least not as efficiently or precisely), such changes have to be understood as purposeful rather than as effects directly caused by environmental factors. In other words, these theories respect the extended Darwinian lesson, but they do not take into account the extended Bernardian lesson.

Moving on, it turns out that none of the major learning theories or their variations successfully deals with the behavioral flexibility problem. This is because they all embrace simple one-way causality from stimulus to response or (as more fashionable these days) from stimulus to cognitive computation to response. But any theory that posits behavior as an end product (output or response) that is elicited by an input (stimulus or perception) with or without intervening cognitive processes is inherently incapable of accounting for the continuous variations in behavior that we observe in the service of achieving goals in the face of continually changing disturbances. Thus a theory that attempts to explain learning as acquisition of a repertoire of responses must fail.

But by combining the insights of Bernard and Darwin we can arrive at an account of learning that solves both the behavioral flexibility and new knowledge problems.

With respect to the behavioral flexibility problem, chapter 6 showed how perceptual control theory gives us a working model for how organisms are able constantly to vary their behavior to achieve goals despite disturbances. An experienced driver can keep his car on the road and in the proper lane while maintaining a relatively constant speed, in spite of varying wind, road cambers, curves, and hills. He must constantly vary his behavior with respect to the steering wheel and accelerator pedal (and perhaps brake pedal) to achieve these effects, and he is able to do this because he has developed, through experience, the necessary perceptual control systems.

Now imagine that a teenager, having learned to drive during the summer, encounters winter driving conditions for the first time. These conditions present new disturbances for which his skills are inadequate. The first time on snow, he will likely accelerate, brake, and make turns too abruptly, resulting in skidding and (one hopes temporary) loss of control of the car. To maintain control of his vehicle in these new conditions, he must adapt his existing control systems, that is, reorganize them in ways so that he will be able to drive safely on winter roads.

Our driver can reorganize his currently existing network of driving-related control systems in a number of ways. The first involves resetting one or more reference levels. For example, under dry conditions negotiating a street corner at 15 mph may be quite safe. But this speed could be dangerous or even impossible to maintain safely while turning on a snow- or ice-covered street. Assuming that the driver's higher-level goal is to negotiate the turn successfully, he will have to reset his lower-level reference level to a lower speed. But since he doesn't actually know what speed is possible to maintain while turning on snow, this resetting will necessarily involve some degree of trial and error (variation and selection).

Another way that control systems can be modified is by reorganizing perceptual functions. This can occur in one of at least two ways. Trying out various combinations of lower-order perceptions can create new higher-level ones. For example, our driver will have to learn to recognize conditions that require reduced speed while turning. During warm weather he may have paid attention only to whether large objects (such as a person or another vehicle) lay before him in the road; now he must become perceptive to indications of the presence of snow or ice on the road surface. The second way to change one's perceptual function is to make it more or less sensitive to certain aspects of the environment. This is technically known as the *gain* of the perceptual function. The driver may have to learn to develop greater sensitivity to the beginning of the car's skid to take prompt corrective actions.

The third major way in which the reorganization of control systems can take place involves modification of their output functions. First it must be recalled that outputs of an internal control system are not motor commands resulting in a specific action or muscle twitch. Instead, they

serve as reference levels for lower-order control systems, and the particular action that results will depend on both this specified reference level and current environmental conditions. Like perceptual input functions, output functions vary with respect to their gain, so that a given error signal (that is, the discrepancy between a reference level and perceptual signal; see chapter 6) may result in output signals of different strengths. In our driving example, we might expect that certain output gains would have to be reduced to avoid too-quick steering, accelerating, or braking behaviors that could cause the car to skid on snow or ice.

Output functions may also change with respect to the particular lower-level reference signals they influence. Consider someone who has always driven a car with an automatic transmission but who now wants to drive one with a manual stick shift. Previously, accelerating from standing to highway cruising speed simply required depressing the accelerator with the right foot and waiting for the desired speed to be attained. Now it requires accomplishing additional lower-level goals involving the left foot and right hand as they operate the clutch and change gears until cruising speed is reached.

This account of learning as the reorganization of perceptual control systems leads to an interesting concept of learning. All traditional learning theories see learning as a modification of one-way cause-effect (stimulus-response or stimulus-computation-response) associations. Recall from chapter 3 that Pavlov understood learning as the association of new stimuli with old responses, as when his dog learned to salivate to the sound of a bell after the bell had preceded several times the introduction of food into the animal's mouth. Skinner (and Thorndike before him) was interested in how new responses to old stimuli were acquired, as when a rat learns to press a lever to obtain food. In marked contrast to both Pavlov and Skinner's stimulus-response theories of learning (and contrasting as well to stimulus-computation-response learning theories of current cognitive science), perceptual control theory sees learning as involving modification of *perceptual* associations, not stimulus-response associations.

To explore this idea, consider a chef who wants to develop a new shrimp entrée to serve at his restaurant. He has a definite goal in mind of what he is trying to achieve in terms of taste, appearance, and consistency (these are the higher-level perceptual goals), but he doesn't yet know what

combination of lower-level goals (that is, perceptions) will lead to his higher-level goals. How many shallots should he mince (should the amount fill four or six tablespoons?)? How long to sauté (how much time should be seen to elapse on the timer?) and at what temperature (how high should the flame be under the pan?)? How well should the shrimp be cooked before adding the wine (offering little resistance to a probing fork or a bit more?)? His cooking experience may well suggest answers to many of these questions. But if he is developing a new dish, the chef is going to have to spend some time experimenting to find the right combination of lower-level perceptions that leads to the desired higher-level perception of a new culinary masterpiece. From this perspective, learning involves discovering new relationships among perceptions (and, of course, being able control them against disturbances), *not* the association of new stimuli with old responses (as in Pavlovian classical or respondent conditioning) or the association of new responses to old stimuli (as in Skinner's operant conditioning).

This culinary experimentation is, of course, an instance of Darwinian variation and selection. Since the chef does not yet know what combination of lower-level perceptions will lead to his desired dish, he will have to use some cumulative trial and error elimination to find out. Four tablespoons of shallots made the dish too bland when first tried, and six made it too spicy on the second attempt. So try five tablespoons and see what happens. Or perhaps stay with six and add a bit more wine. This process of within-organism variation and selection provides an answer to how new knowledge is possible. It does not involve variation and selection of specific overt responses as Skinner believed, but rather the variation and selection of controlled lower-level perceptions, eliminating those that do not lead to the desired higher-level goal and retaining those that do.

In further contrast to Skinner's theory, variation and selection of lower-level perceptions in the service of higher-level ones need not involve overt behavior, at least not with humans and some other primates. Instead, we can use our mental models of how our physical and social worlds work to try out combinations of lower-level reference perceptions and imagine their effects on higher-level ones. So if I am having a dinner party and inviting ten guests—some of whom get along well together while some others don't—I can imagine different seating plans (variations) before the

guests arrive and eliminate potentially troublesome arrangements (those that put suspected antagonists within striking distance of each other) until I come up with a plan that seems best (selection). This process is an instance of Donald Campbell's vicarious variation and selection described in chapter 9.

A good deal of what we call thinking—at least thinking that involves problem solving, invention, and creativity—may actually be vicarious variation and selection of perceptual control systems. Such a concept lets us understand how cognitive processes involved in thinking can be purposeful even when they are not accompanied by concurrent purposeful behavior. And if thinking alone cannot generate solutions, we can assist it with other forms of substitute variation and selection: writing down our ideas on paper, using computers to run simulations of candidate solutions, or discussing the problem (proposing solutions, eliminating bad ones, and keeping the best) with other individuals.

A New Conception of Learning

By combining the extended Bernardian lesson (that organisms vary their behavior to control their perceptions) and the extended Darwinian lesson (that organisms make use of variation and selection to gain control of aspects of their environment) we arrive at a new conception of learning. Learning is no longer the association of new stimuli to old responses, or acquisition of new responses to old stimuli, but rather acquisition of new means of perceptual control by reorganizing existing perceptual control systems by within-organism variation and selection. In much the same way that *E. coli* randomly changes its direction when it senses that it is not moving closer to food, all learning requires an organism to make some change to its current organization of perceptual control systems when there is some chronic error between perception and reference level. And whereas previous learning experiences may usefully constrain the variations that are tried (an automobile mechanic is not likely to change the air pressure in a car's tires to see if doing so will make it start), acquisition of new knowledge requires at least some blind variation to explore and discover new useful relationships between combinations of perceptual variables.

The knowledge gained by such a process is discovering what combinations of lower-level perceptions are successful in bringing about desired higher-level perceptions controlling these lower-level perceptions against disturbances. A chef does not measure a cup of water by holding a container under an open faucet for a fixed amount of time, as this would lead to varying amounts of water due to the fluctuating pressure of the water supply line. Instead, he keeps the container under the faucet until the water level reaches the one-cup mark, no matter how long it may take. By successfully controlling this and other lower-level perceptual variables, he is able to prepare the entrée he has in mind, that is, match his higher-level reference perception. By extending this form of purposeful evolution to the mental realm when no overt behavior is involved, we obtain a new framework for understanding cognitive processes. Cognition is no longer seen as planning responses to certain stimuli, but rather as Darwinian reorganization of Bernardian perceptual control systems to control new aspects of the environment.

This view can be used to develop a general framework of knowledge and its acquisition. Within such a framework are three principal types of knowledge. First is the biologically based knowledge that we and all other organisms are essentially born with. This may be all the knowledge that a single-cell organism will ever have and it is reflected both in its structure and instinctive behavior. The way we see colors (or how we see at all) is a form of this knowledge and it cannot be changed, although certain experiences are necessary for it to develop, such as growing up and interacting in a world with visible light. This knowledge is derived from the cumulative among-organism selection of the fittest, as originally proposed by Darwin.

Second is the knowledge that some organisms acquire during their lifetimes. It results from the interaction of one's biological endowment with one's particular experiences, and it is limited in important ways by one's biologically provided knowledge. Humans can learn only certain types of languages. We cannot learn to make visual distinctions between two ultraviolet patterns the way bees can. Rats fail to learn certain tasks requiring visual discrimination, but can learn similar tasks involving their keen sense of smell. Such knowledge is similar to biologically based knowledge in that it also depends on variation and selection. But it is different in two

key respects: it involves the within-organism variation and selection of modifications to perceptual control systems, and it is driven by the organism's internal goals. So unlike biological evolution, the knowledge that an organism acquires during its lifetime results from a purposeful form of continuing variation and selection.

Finally, some organisms, especially, humans, seem to acquire knowledge from others. But this is actually a special case of the second form since it is acquired as the result of one's individual experiences interacting with biological knowledge. We may be able to make use of the trial-and-error experience of others by observing their (successful and unsuccessful) behavior, or talking with them or reading their books. But we cannot simply absorb this knowledge in the way that a blank computer diskette can receive the information stored on another. Instead, it could be reasonably argued that the knowledge we derive from others' experiences also requires some degree of within-organism variation and selection (see Cziko 1995, chapter 10). I may observe how an expert skier moves his skis and holds his body as he descends the slope. He may even give me instructions and sell me his book and video on skiing. But although this information may facilitate my development as a skier, it cannot replace the need for me to reorganize my perceptual control systems, eliminate those modifications that leave me sitting in the snow, and retain those that keep my posterior above my skis.

The view of learning provided by combining the extended Bernardian and extended Darwinian lessons has important implications for all forms of education and training, a topic to which we will return in the final chapter.