

*THE UNIT OF SELECTION:
WHAT DO REINFORCERS REINFORCE?*

JOHN W. DONAHOE, DAVID C. PALMER, AND JOSÉ E. BURGOS

UNIVERSITY OF MASSACHUSETTS AT AMHERST, SMITH COLLEGE, AND
UNIVERSIDAD CENTRAL DE VENEZUELA AND UNIVERSIDAD CATOLICA DE VENEZUELA

We begin by stipulating three central points upon which we and the broad consensus of commentators agree. First, the effects of reinforcers on behavior can be readily demonstrated under conditions in which the antecedents of behavior are not identified. That is, orderly functional relations emerge between operants and reinforcers when the experimental analysis of the effects of antecedents is impossible or impractical. Second, response-contingent reinforcers most commonly alter the control of responses by their antecedents. That is, discriminative control of responding is “practically inevitable” (Skinner, 1937, p. 273). Third, the manipulated antecedents of behavior are typically environmental events in cases that are amenable to experimental analysis, but may include covert or private (intraorganismic) events as well. Covert events (characterized behaviorally or neurally) invariably accompany environment–behavior relations and indispensably contribute to scientific interpretation.

Where we differ from some commentators (but fewer than we might have supposed) is with respect to whether control by antecedents is inescapable, including on those occasions in which the antecedents cannot be manipulated (i.e., those in which functional relations between antecedents and behavior

cannot be experimentally evaluated). Stated most simply, our view is this: Because response–reinforcer contiguities have been shown to alter the control of behavior by its antecedents whenever circumstances permit the antecedents to be manipulated, such control may be presumed to exist under conditions that do not permit the experimental analysis of the effects of antecedents. In our view, to hold otherwise is tantamount to claiming that events do not exist when we are not sensing them, and Skinner long ago rejected this sterile form of operationalism (Skinner, 1945). Parenthetically, when many, including ourselves, find that Skinner’s writings facilitate theoretical developments whereas others find that they inhibit such efforts, these differing effects can hardly be attributed to the words alone; the differing histories of the readers must bear some of the responsibility.

Space does not permit us to respond to all of the worthwhile points raised in the commentaries. In that regard, it is well to remember that the original review in this sequence (Shull, 1995) was prompted by a book-length manuscript, *Learning and Complex Behavior* (Donahoe & Palmer, 1994; *LCB*), and that many of the questions posed in the commentaries are explicitly addressed there. What we have offered in *LCB* is a *beginning* of an integration of behavior analysis and neuroscience. Nevertheless, even at this early juncture, we do not believe that any of the issues raised in the commentaries fundamentally challenge the central themes of our biobehavioral approach or necessarily exceed its grasp. This is not to say that all issues related to some of the specific points have been satisfactorily or completely addressed, only that their answers do not seem, in principle, to be beyond the ken of the approach outlined in *LCB*. In what follows, we consider three general issues that are associated with many of the specific points made in the commentar-

The preparation of this reply was supported in part by a faculty research grant from the Graduate School of the University of Massachusetts at Amherst. The authors thank Vivian P. Dorsel for comments on an earlier version of the manuscript and J. Russel Keath for discussions concerning the cellular mechanisms of plasticity.

Correspondence and requests for reprints may be addressed to John W. Donahoe, Department of Psychology, Program in Neuroscience and Behavior, University of Massachusetts, Amherst, Massachusetts 01003 (E-mail: jdonahoe@psych.umass.edu); David C. Palmer, Department of Psychology, Clark Science Center, Smith College, Northampton, Massachusetts 01063 (E-mail: dcpalmer@science.smith.edu); or José E. Burgos, Consejo de Estudios de Post-grado, Facultad de Humanidades y Educación, Universidad Central de Venezuela, Caracas, Venezuela (E-mail: jburgos@zeus.ucab.edu.ve).

ies: (a) appropriate units of selection and levels of analysis for behavioral research, (b) the effects of selection by reinforcement, and (c) the relation of neural networks to neuroscience. The commentaries most clearly relevant to the various points are indicated within parentheses throughout the reply. (Where the views of the commentators are generally consistent with ours, their names are simply listed. Where they have identified an issue without necessarily advocating a particular position, their names are preceded by *viz.*, *i.e.*, to wit. Where they have taken a position that appears to differ substantially from ours, their names are preceded by *cf.*, *i.e.*, compare. In many cases, the remarks of the commentators on a topic are multifaceted and resist simple categorization. Accordingly, readers are encouraged to evaluate each of the cited commentaries for themselves.)

UNITS OF SELECTION AND LEVELS OF ANALYSIS

As noted previously, our position is that antecedent–behavior relations are appropriate units of selection for behavior analysis, where the antecedent takes the form of an environmental stimulus in most experimental analyses (*viz.*, Field). In short, reinforcers alter the strength of antecedent–behavior relations, not behavior alone. This view, if sustained, opens the way for a parsimonious account of a wide variety of phenomena that have been found with both discriminative and nondiscriminative (free-operant) procedures and operant and respondent (Pavlovian) procedures. (The competence of the neural mechanisms proposed to mediate these phenomena and of the neural networks intended to simulate these mechanisms are separate matters that are treated in the second and third major sections of our reply.) Antecedent–behavior relations are the focus of selection by reinforcement just as phenotypic characteristics are the focus of natural selection. At the neural level, synaptic efficacies along the pathways mediating antecedent–behavior relations are the physical substrate of selection by reinforcing dopaminergic systems just as, at the genetic level, genes are the substrate of natural selection by the ancestral environment (*cf.* Hutchison).

The primacy of antecedent–behavior relations is implicit in the work of both the experimental analyst and the applied behavior analyst. The experimentalist studying the effects of response–reinforcer contingencies holds constant the environment so that the effects of the manipulated contingency might be observed unaffected by variations in the environment. The applied analyst seeking to reduce phobic behavior extinguishes or counterconditions fearful responses in the contexts in which those responses have been conditioned. “Fear” is not reduced in general, but with respect to certain stimuli. However, in neither case are we arguing that knowledge of the environment is sufficient for understanding the behavior of interest. The organism, through the effects of its selection history on its nervous system, makes a necessary contribution to the environment–behavior relation (Galbicka; McIlvane & Dube; Staddon), and commentators agree that our simulations possess this characteristic (Kemp; Williams).

In order for the effects of selection by reinforcement to be understood, the “natural lines of fracture” must be honored at all levels of organization (Churchland & Sejnowski, 1992)—behavioral, neural, and cellular (Field; Galbicka; Hutchison; see Palmer, Donahoe, & Crowley, 1985). As an example, what is considered a unitary response at the behavioral level (e.g., lever pressing) is an expression of the concerted firing of a large population of cortical and motor neurons at the neural level (e.g., Georgopoulos, 1990, *in press*). It is not that one level of analysis is “right” and the other “wrong,” but that one or the other level is more or less appropriate for the phenomenon under study. At the behavioral level, the neural activity underlying the response is invisible to the reinforcing environment. No matter what population of neural activity produces the behavioral response, the reinforcer strengthens the environmental control of that response and, in so doing, changes the synaptic efficacies that mediate it. As Skinner noted, neuroscience fills in the “gaps” (Skinner, 1974, p. 237) between the environment and behavior (McIlvane & Dube).

Our position regarding the relation between different levels of analysis is as follows:

The observability of a response is not determined by its intensity or magnitude, but by the characteristics or tools of the observer. Many venerable dependent variables . . . , such as changes in skin conductivity, blood pressure, heart rate, etc., are measurable only through instrumental amplification. . . . We must avoid the temptation to think of covert behavior as a *kind* of behavior, with properties essentially different from overt behavior. Rather, all behavior lies on a continuum of observability. . . . Our subject matter—behavior—is not defined by its magnitude or by the ability of observers to agree on its occurrence. Rather it is any activity of the organism that can enter into orderly relationships with environmental events. . . . In a real organism the eyeblink is an overt behavior and, as such, can be measured, recorded, and agreed upon by disinterested observers. However, the activity of [motor neurons] is observable, if at all, only by the most intrusive of neurophysiological techniques. Since the activity of [motor neurons] determines the activity of the eyelid there is no reason to distinguish the two events: They are two dimensions of a single response. However, the two events play different roles. The environment only sees the overt response, and so can arrange contingencies of selection only for the eyeblink. However, the enduring changes brought about by the contingency surely happen at [motor neurons], and elsewhere in the network. The probability of overt responses is altered by contingencies of reinforcement only because the nervous system is. The overt response is no more the “real” response than its neural precursors are. . . . We believe that we are justified in considering covert events . . . in our interpretation of complex behavior provided that we do not introduce *ad hoc* principles that are not founded in the experimental analysis of overt, measurable, quantifiable behavior. . . . Inferences about covert events should *follow from* behavioral laws, not serve to mask their inadequacy. (*LCB*, pp. 275–277)

Four reasons are given in *LCB* for integrating the experimental analysis of behavior and neuroscience (pp. 5–8): (a) As just noted, distinctions between different levels of analysis—or, better perhaps, different scales of measurement (Philip Heline, personal communication)—are epistemological, not ontological. That is, the distinction is between different ways of studying the same nature, not different natures. (b) The precision of functional relations uncovered at one behavioral level can sometimes be improved

when they are supplemented (not replaced) by observations at finer levels of analysis, whether microbehavioral or neural. (c) Some antecedent–behavior relations exemplify nonlinearly separable functions—such as the XOR problem in artificial intelligence (*viz.*, Hutchison)—that require the introduction of mediating events if ad hoc formulations are to be avoided. (d) And, finally, we were motivated by “a simple curiosity about the processes underlying the functional relations between the environment and behavior” (*LCB*, p. 8). As noted by some commentators (McIlvane & Dube) and by ourselves in other venues (e.g., Donahoe & Palmer, 1989), Skinner encouraged the ultimate integration of behavior and neuroscience. To wit,

The physiologist of the future will tell us all that can be known about what is happening inside the behaving organism. His account will be an important advance over a behavioral analysis, because the latter is necessarily “historical”—that is to say, it is confined to functional relations showing temporal gaps. . . . What he discovers cannot invalidate the laws of a science of behavior, but it will make the picture of human action more nearly complete. (Skinner, 1974, pp. 236–237)

And, much earlier (Skinner, 1938),

I am not overlooking the advance that is made in the unification of knowledge when terms at one level of analysis are defined (“explained”) at a lower level. (p. 428)

I agree with Carmichael [1936] that “those concepts which do not make physiological formulation impossible and which are amenable to growing physiological knowledge are preferable, other things being equal, to those that are not so amenable.” (p. 440)

Moreover, the interchange between behavior analysis and neuroscience should not be a one-way street. Again, as Skinner (1938) also noted, neuroscience benefits from a science of behavior at least as much as a science of behavior benefits from neuroscience. No one who has read the musings of some eminent neuroscientists on the subject matter of behavior analysis can quarrel with that claim (e.g., Eccles, 1994; Griffin, 1992; Sperry, 1990).

As but one example of the potential benefits of the integration of behavioral and neural observations, consider the following ex-

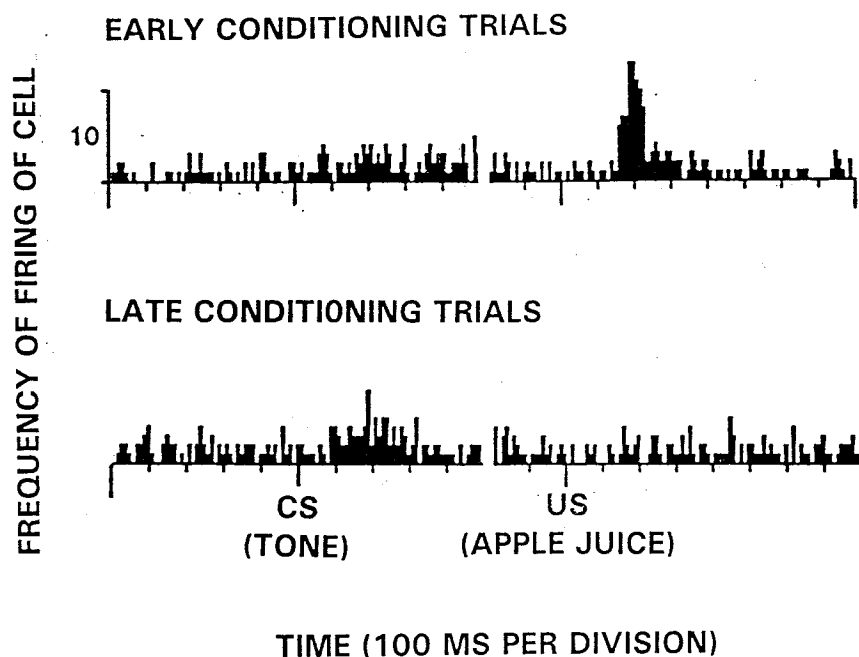


Fig. 1. Neural basis of the discrepancy requirement for conditioning. Shown is the frequency of firing of a dopaminergic neuron in the midbrain (ventral tegmental area, or VTA) of a monkey during early trials (top) and late trials (bottom) of a Pavlovian, or respondent, conditioning procedure. Note that the neuron increased in its frequency of firing after the onset of the unconditioned stimulus (US) during early trials but after the onset of the conditioned stimulus (CS) during later trials. Adapted from Schultz (in press).

periment that came to our attention after the publication of *LCB*. Prior behavioral work had indicated that, in addition to temporal contiguity, putative reinforcing stimuli were effective only if, speaking nontechnically, the reinforcer was “surprising” (Kamin, 1968; Rescorla & Wagner, 1972). *Surprise* may be given a technical definition at the behavioral level as a discrepancy between the response evoked by the reinforcing stimulus (the unconditioned response, or UR) and the level of that *same* response in the environment in which the operant and reinforcer occur (the conditioned response, or CR) (cf. Vaughan). In these terms, a putative reinforcer strengthens environment–behavior relations when there is a contiguous CR–UR discrepancy (Donahoe, Crowley, Millard, & Stickney, 1982; cf. Rescorla, 1968). Although some data directly substantiate this behavioral formulation of discrepancy (e.g., Donegan & Wagner, 1987; Terry & Wagner, 1975), many preparations for the study of respondent conditioning (e.g., conditioned suppression; Estes & Skinner, 1941) and almost all for the study of

operant conditioning do not permit direct measurement of the CR and UR (Galbicka; but see Shapiro & Miller, 1965). In addition, the behavioral measures are, at best, *correlated* with the underlying neural activity that mediates conditioning. That is, no one seriously contends that behaviorally defined CR–UR discrepancies *cause* the changes in synaptic efficacies that mediate observed environment–behavior relations.

Direct neural measures of the discrepancy are now provided by studies conducted in Wolfram Schultz’s laboratory (Schultz, in press; Schultz et al., 1995). The activity of dopaminergic cells in the ventral tegmental area (VTA) was monitored when a reinforcing stimulus of apple juice was presented to a restrained monkey. (The VTA was proposed in *LCB* as the source of a nonspecifically projecting dopaminergic system that modulated changes in synaptic efficacies.) As shown in Figure 1, during the first trial in which apple juice was given preceded by a brief stimulus (a tone), VTA neurons fired rapidly upon the onset of the apple juice. However, after a

number of trials in which the apple juice was preceded by the tone, these same neurons fired to the onset of the tone but *not* to the apple juice. Thus, the basis of the reinforcing effect of the apple juice was eliminated in an environment in which the tone was present. At the neural level, a discrepancy no longer existed between the conditioned activity of VTA neurons to tone and the initial unconditioned activity of VTA neurons to apple juice.

On the general topic of levels of analysis (scales of measurement), we have nothing to add to Skinner's position as supplemented by normative philosophy of science: (a) General principles arise as inductions from the experimental analysis of particular public observations. (b) Such principles are validated by independent observations at their own level of analysis, and, when validated, are said to explain the observations. (c) The principles themselves are explained by observations at levels lower than those at which the principle was formulated. To hold otherwise invites the circular reasoning that bedevils normative psychology (*LCB*, pp. 9–10). The unending nature of reductive scientific explanation is not a source of embarrassment; it is simply the way science proceeds (cf. Field).

In behavior analysis, the levels-of-analysis issue most often surfaces in the context of the molar–molecular debate. We have only two comments beyond those contained in the target article: (a) It is undeniably correct that orderly relations exist between variables defined over appreciable intervals of time (Field) (e.g., as, in polydipsia, between the amount of water consumed and the rate of food presentation over a session). However, it is also true that orderly relations exist between individual events within the session (e.g., between the onset of a bout of drinking and the delivery of a piece of food). Order exists at both levels simultaneously; which orderly relation is most useful depends on the question at hand. There are no inherently molar or molecular levels of analysis; these are relative terms, either of which can be applied to the very same observation depending on the frame of reference at that moment. (b) It is quite true that molecular explanations of molar regularities in behavior are incomplete (e.g., the regularities that are observed when behavior is exposed to a

schedule of reinforcement) (viz., Williams). However, the distinction is not between an incomplete molecular account and a complete molar account. Both accounts are incomplete (e.g., B. Williams & Bell, 1996). Whatever the eventual resolution of an integrated molar–molecular account, it would be well to avoid locutions that risk confusing the effects of three-term contingencies with the explanations of those effects. For example, consider the statement: “A critical issue is the nature of the choice rule that *governs* response selection” (Williams, p. 257, emphasis added). Even if an adequate molar principle of choice were formulated, it would, at most, “govern” the behavior of the scientist, not the subject. A molar choice rule may be a valid induction from observations of behavior, but the moment-to-moment contiguities of environment, behavior, and reinforcer are the events that make contact with the organism (Galbicka; Hutchison; Marr; Vaughan).

THE EFFECTS OF SELECTION BY REINFORCEMENT

The approach to conditioning developed in *LCB* and elsewhere (e.g., Donahoe, Burgos, & Palmer, 1993; Donahoe et al., 1982) proposes a single principle of selection by reinforcement that is competent to yield a variety of differing effects as its cumulative product. Which effects are produced depend on the specific environmental and behavioral events that are reliably contiguous with the reinforcer (i.e., with the stimulus evoking a behavioral discrepancy). It is in this sense that we have proposed a *unified* reinforcement principle. However, we were remiss in not explicitly stating that the proposed neural mechanisms implementing that functional principle apply to only those antecedent–behavior relations that are mediated by neural systems whose synaptic efficacies are directly modifiable by the nonspecific dopaminergic system (viz., Dworkin & Branch). Because the nonspecific dopaminergic system does not directly innervate the autonomic nervous system, our position is consistent with early speculations by Skinner (1937) regarding the inability to condition autonomic responses using operant reinforcement contingencies (see Donahoe, in press-b).

Two caveats are in order: First, the foregoing does not preclude the possibility of indirectly affecting autonomically mediated responses by instituting operant contingencies for response systems that are affected by the nonspecific dopaminergic system (cf. Smith, 1954). For example, chewing and tongue movements in dogs could be operantly conditioned with the collateral effect of evoking salivation through the oral stimulation occasioned by such responses. Second, although a dopaminergic neuromodulator may not affect autonomically mediated antecedent-behavior relations, other neuromodulators (perhaps present tonically rather than phasically) may be implicated (cf. Bear & Singer, 1986). The inability of autonomically mediated responses to be modified directly by operant contingencies illustrates a case in which the interpretation of anomalous findings at the behavioral level (some responses are affected by their consequences, whereas others are not) may be resolvable by information from another level of analysis.

The fact that autonomically mediated responses may not be acquired through the same neural systems as responses mediated by the dopaminergically modulated system has no necessary effect on the province of a unified reinforcement principle: This functional principle continues to apply to all responses that are acquired through operant and respondent contingencies. Autonomically mediated responses are simply excluded as candidates for direct selection by operant contingencies, but they are nevertheless acquired when operant contingencies are implemented. (A possible underlying reason for their exclusion is presented later.) The contention that both operants and respondents are selected during operant contingencies is noncontroversial, and is explicitly endorsed by a number of the commentators (Galbicka; Dworkin & Branch; see *LCB*, pp. 44–45). Indeed, one of the most important implications of the unified reinforcement principle is that the outcome of conditioning “depends on interactions, if any, among stimuli and responses that are candidates for inclusion in the selected environment-behavior relation” (*LCB*, p. 52, see especially Figure 2.11). As an illustration, pecking by pigeons is readily conditioned to a spatially localized stimulus using food as a reinforcer because the operant and

the reinforcer-elicited response facilitate one another: The responses that are reliably present in the vicinity of the discrepancy are pecking due to the operant peck-food contingency and pecking elicited by the presentation of food (see Jenkins & Moore, 1973). In addition, any autonomically mediated responses that are evoked by food (e.g., changes in heart rate) are also conditioned at the same time. Contrariwise, withholding pecking a stimulus that is followed by food should be problematic, as it is in the negative automaintenance procedure (D. Williams & Williams, 1969), because the response elicited by food is now incompatible with the response that preceded food (*LCB*, p. 52; Donahoe & Wessells, 1980, p. 161). Interactions between the various responses selected by the reinforcer provide the basis for interpreting a number of phenomena about which questions were raised in the commentaries.

We turn now to some of these phenomena and how they may be interpreted within the framework of a unified reinforcement principle. These accounts are, in many cases, verbal interpretations and, as such, require additional experimental work and formal interpretation (e.g., via neural network simulations). Our goal here is the modest one of demonstrating that these phenomena are neither inconsistent with nor necessarily beyond the reach of the approach taken in *LCB*. The place in *LCB* where the phenomena are discussed is indicated as each point is considered.

Motivation and Internal Stimuli

Our approach to conditioning assigns an important role to internal (private) stimuli (viz., Field; McIlvane & Dube), including those produced by motivational operations (viz., Michael). Clearly, such stimuli are essential to any comprehensive account of behavior. The emphasis upon the control of behavior by environmental stimuli, both in *LCB* and in the target article, occurred because such stimuli are more directly manipulable than internal stimuli and, hence, are more readily subjected to experimental analysis. And, more fundamentally, private events are ultimately traceable to the action of events that originate in the environment. It is primarily with regard to *interpretation* rather than experimental analysis that internal stimuli

make an appearance in *LCB*. There, internal stimuli are produced by feedback from neural and motor activity and are realized as recurrent connections within neural networks (e.g., *LCB*, p. 276).

In our formulation, deprivation may affect behavior in several ways (*LCB*, pp. 35–36). For one, by depriving an organism of contact with a stimulus, that stimulus typically becomes a more vigorous elicitor of behavior. As such, the stimulus is able to function as a more effective reinforcer because its presentation evokes a larger behavioral discrepancy. Further, discriminative and occasion-setting functions of deprivation, and motivating operations in general, may be readily implemented in neural networks. A discriminative function is enabled to the extent that the motivating operation differentially activates some units within a network; a motivating function is enabled to the extent that the motivating operation nondifferentially activates a range of units within the network. In either case, the activation levels of units within the network are changed, thereby changing which connections are eligible for modification by the reinforcer (neuromodulator). (Recall that the dopaminergic reinforcer increases the connection weights between *all* coactive units.) Through both means, the motivating operation may have a pervasive effect on behavior and on neural networks intended to simulate behavior. This is not the place to describe potential neural mechanisms that might implement such motivational effects, but several candidates are available among the other nonspecific projection systems whose nuclei are located in the midbrain and brainstem.

Reinforcing Operations

Stimuli are proposed to function as reinforcers to the extent that, on the behavioral level, they provoke a behavioral discrepancy and, on the neural level, they cause dopamine to be liberated in synaptic clefts between coactive pre- and postsynaptic neurons. The three ways in which stimuli can function as reinforcers (viz., Dworkin & Branch) can all be interpreted in terms of this account. First, stimuli, such as cocaine, that function unconditionally as reinforcers cause the release of dopamine (and other neuromodulators) and block its reuptake. This increases

the concentration of dopamine in the synaptic cleft and, thereby, promotes reinforcement at the cellular level. Second, deprivation enhances the reinforcing function of stimuli through increasing the ability of stimuli to cause the release of dopamine by fibers originating from the VTA. Among others, this effect has been demonstrated using *in vivo* microdialysis to detect the effects of deprivation on the release of dopamine in the nucleus accumbens and frontal cortex (e.g., Hernandez & Hoebel, 1990; Hoebel, 1988). Third, pairing a stimulus with a known reinforcer can establish that stimulus as a conditioned reinforcer through strengthening feedback pathways from the frontal cortex and neostriatum to VTA (*LCB*, pp. 96–101). The proposed neural mechanisms of conditioned reinforcement merit a somewhat more detailed description because they are pertinent to replies to several queries.

Conditioned reinforcement. Unconditioned reinforcers are proposed to have two effects. First, and most obviously, they strengthen the responses that precede them. This is the defining feature of reinforcement on the behavioral level and is accomplished, on the neural level, through the release of dopamine by VTA neurons, with the consequent strengthening of synaptic efficacies between neurons arising from sensory areas and projecting to motor areas (see Figure 2). Second, at the same time that environment–behavior relations are strengthened, synaptic efficacies are being increased along other pathways leading from frontal cortex (and neostriatum) back to VTA (*LCB*, pp. 96–101; Donahoe, *in press-b*). As these feedback pathways are strengthened, the stimulus paired with the unconditioned reinforcer becomes able, *by itself*, to activate the VTA. In short, conditioned reinforcers exploit feedback pathways to activate the VTA. Once these feedback pathways become functional, both unconditioned and conditioned reinforcers cause dopamine to be liberated and synaptic efficacies to be changed via a common cellular mechanism. Computer simulations of the effects of the feedback pathways that mediate conditioned reinforcement indicate that conditioning in multilayered networks, in which unavoidable and appreciable delays occur between the coactivity of units along pathways that mediate environment–behavior relations




Fig. 2. Schematic diagram of the neural circuits for conditioned reinforcement. Unconditioned reinforcers, such as food, stimulate afferent pathways that activate the ventral tegmental area (VTA) directly, thereby strengthening synaptic efficacies via the diffusely projecting, dopaminergic system. Synaptic efficacies are strengthened along two sets of pathways: (a) those mediating reinforced behavior and (b) those leading back to the VTA (curved line with two right arrows). During the course of conditioning, these feedback pathways become capable of implementing conditioned reinforcement. After Donahoe and Palmer (1994).

and the production of the unconditioned reinforcer by the operant, cannot occur without such feedback circuits (*LCB*, see Figure 4.8, p. 101). These feedback pathways are said to implement *internal reinforcement*. Because feedback pathways from the cortex and neostriatum to the autonomic system appear to be absent, the inability to engage the internal-reinforcement mechanism may account for the failure of autonomic responses to be conditioned with operant contingencies: The delay between the occurrence of the autonomic response and the occurrence of the external (unconditioned) reinforcer may exceed the temporal requirements of the cellular mechanisms of reinforcement.

Devaluation

Much experimental work within the associationist tradition has been directed at parsing the “associative structure” into its constituent parts (Colwill & Rescorla, 1986; Mackintosh, 1983; Rescorla, 1991). Thus, Pavlovian conditioning is said to involve the acquisition of stimulus–reinforcer associations,

whereas operant conditioning is said to involve, in addition, response–reinforcer associations. This enterprise cannot be comprehensively evaluated here. In our view, it has been extremely productive experimentally, but is conceptually flawed (e.g., Donahoe, 1984). Devaluation is among the phenomena that have been taken as evidence for the formation of response–reinforcer associations, or response–outcome associations as they are often called in this literature. In the typical devaluation experiment, a response is acquired using one reinforcer and then some manipulation is applied to the reinforcer that is designed to alter—to devalue—its reinforcing effect (*viz.*, Williams). Later, when returned to the environment in which the response was originally acquired, responding is found to be weakened relative to appropriate control levels. How is this phenomenon to be understood in terms of the account provided in *LCB*?

Our interpretation exploits the fact that the CR is acquired before the operant in our neural network simulations under conditions

that are behaviorally and neurally plausible and very general (*LCB*, pp. 58–60). If the neural processes leading to CRs are acquired before operants and if CR processes are discriminable, then the internal stimuli produced by CR processes are in a favorable temporal relation to the operant to acquire discriminative control over the operant. (The relation of this proposal to Hull's r_g-s_g mechanism has not escaped us; Hull, 1934; *LCB*, p. 122.) In this way, the operant comes under the control of *both* environmental stimuli and the internal CR-produced stimuli that precede the operant. (It is not clear how response–reinforcer associations can control the emission of an operant, because such associations, of necessity, *follow* the operant. One would be forced to argue that response–reinforcer associations are evoked by stimuli that precede the operant, in which case we are back to the three-term contingency of Skinner.) Any subsequent manipulation that alters the discriminative control exerted by these CR(UR)-produced stimuli could also weaken the operant to the extent that CR-produced stimuli evoke neural and behavioral responses that interfere with the operant (*LCB*, pp. 108–109). This interpretation of devaluation is related to two-factor theory (cf. Williams), but does not ascribe a motivational function to CR-produced stimuli and does not consider such stimuli to be necessary for the occurrence of the operant (cf. Trapold & Overmier, 1972). Although efforts have been made to eliminate interpretations in which CR-produced stimuli have discriminative control over operants (e.g., Rescorla & Colwill, 1989), the present proposal remains viable (*LCB*, pp. 122–123).

Sensory Processes and Representations

At the behavioral level, stimuli are defined functionally (i.e., they are events that control behavior as the result of natural selection or selection by reinforcement). At the neural level, they may also be defined functionally (i.e., they are presynaptic events, including environmental events, that control the firing of postsynaptic neurons). Stimuli are also defined functionally as interpreted by selection networks (cf. Dworkin & Branch). Environmental events are simulated by the activation of one or more input units with the activity of different sets of units corresponding to dif-

ferent events such as lights or tones. No effort has been made to capture the richness of receptor structure and transduction process in our simulations, although potentially both could be required to simulate the particulars of some phenomenon. The conjunction of environmental events is simulated in selection networks by simultaneously activating more than one set of input units and, under certain conditions such as configural conditioning, by increasing connection weights from multiple sets of input units to some interior units in the stimulus association portion of the network (see Figure 2 of the target article; *LCB*, pp. 196–198). In the latter case, the reinforcing stimulus, through its effects on hippocampal output (see below), causes a subpopulation of units to form in the sensory association subnetwork that can be said to “represent” (viz., Williams) the conjunction of stimuli (e.g., a light–tone compound stimulus). Note that such functionally complex units are acquired on-line as needed; they are not the product of a priori “handcrafting” the network for that particular condition (cf. Gluck & Bower, 1988).

Although subpopulations of units within the sensory association subnetwork may be said to form a “representation” of complex stimulus configurations, we explicitly disavowed the term in *LCB*: We described this practice as “misleading” because the pattern of neural activity initiated in the nervous system by a stimulus did not have the dimensions of the stimulus, did not include the motor components of the neural activity initiated by the stimulus, and was too static a conception. For us, a stimulus initiated a variable but distinctive pattern of neural activity—“a probabilistic ensemble of unit activity and not a ‘thing’” (p. 264). Note also that these ensembles of sensory association units are directly facilitated by the action of the reinforcer. As shown in Figure 2 of the target article, dopaminergic fibers from VTA innervate the hippocampus, more particularly pathways originating in the hippocampus and projecting nonspecifically to sensory association cortex. Research from Stein’s laboratory (e.g., Stein & Belluzzi, 1988, 1989; Stein, Xue, & Belluzzi, 1993, 1994) has shown that dopamine modulates hippocampal output and, in that way, may affect perceptual processes. Thus, dopamine coordinates the selection of

environment–behavior relations *and* the complex concatenations of stimuli that often guide complex behavior. In short, dopamine was proposed to have both a reinforcing function and a salience-enhancing function in the nervous system and in our simulations (*viz.*, Dworkin & Branch).

Aversive Stimuli

A number of commentators raised questions about the treatment of aversive stimuli within the context of a unified reinforcement principle (*viz.*, Dworkin & Branch; Field; Vaughan). Only one issue is addressed here, but others have been treated elsewhere (e.g., Donahoe & Wessells, 1980, pp. 207–210; Weisman & Litner, 1969). How can the same principle be consistent with both response strengthening (reinforcement) and response weakening (punishment)?

Aversive stimuli are stimuli that, by definition, evoke escape or withdrawal responses. Recall that reinforcers have two effects with operant contingencies in our formulation: (a) Reinforcers lead to the acquisition of both the operant (R) and the reinforcer-elicited response (UR). (b) The conditioned response (CR) is acquired *before* the R. In the case of an aversive eliciting stimulus, the UR is withdrawal and successfully competes with the operant, thereby preventing the operant from being strengthened by the aversive elicitor, which would otherwise be the case (*LCB*, pp. 114–115). According to the proposal, punishment is produced by the more rapid acquisition of conditioned withdrawal responses than operants. This account parallels a core feature of Skinner's (1953) treatment of punishment, in which aversive stimuli were held to weaken only indirectly the responses that produced them.

Note also that the interpretation of punishment provided by the unified reinforcement principle predicts that aversive stimuli should be able to function as reinforcers under some circumstances: If the operant that precedes the aversive stimulus is topographically similar to the responses elicited by the aversive stimulus, then the aversive stimulus should strengthen the operant. In fact, something closely resembling this phenomenon has been found. Especially with monkeys, electric shock functions as a reinforcer for some shock-elicited responses such as pulling on a

leash, biting, or hand movements directed toward objects in the environment (e.g., Branch & Dworkin, 1981; Morse & Kelleher, 1977). Although related effects have been reported with rats (*LCB*, p. 115), the effect is strikingly more pronounced and stable in monkeys. It is of interest that dopaminergic projections (Berger, Gaspar, & Verney, 1991) and projections from prefrontal cortex to VTA (Goldman-Rakic, Chafee, & Friedman, 1993) are much more extensive in primates than in rodents. The availability of more extensive circuitry for mediating internal reinforcement would reduce the difference between the rate of acquisition of CRs and Rs and, in so doing, enable aversive stimuli to more readily function as reinforcers for some operants. The greater demands placed on the internal reinforcement mechanism by the need to modify the connectivity of the larger, more deeply layered brains of primates (including ourselves) may have exacted a price—vulnerability to the temptations of aversive stimuli. (The foregoing is not intended as a comprehensive account of the behavioral and neural processes engaged by aversive stimuli. For example, other neural systems, most notably those involving the amygdala and its interactions with the VTA, undoubtedly play a major role; Davis, Rainie, & Cassell, 1994; LeDoux, 1995.)

NEURAL NETWORKS AND NEUROSCIENCE

The final section of our reply is concerned with the relation between behavior analysis and neuroscience, particularly with respect to the interpretation of behavior by means of neural networks. This is a vast topic about which the commentators have made many important points. Here, we briefly consider only two issues: How, and to what extent, should neural network research be informed by neuroscience? How can the structure of neural networks (i.e., their architecture or “neuroanatomy”) be addressed in a principled manner?

Before considering these questions, we wish to identify ourselves with the proposition that behavior analysis is a science whose independence is unthreatened by its relation to neuroscience (Hutchison; Marr; McIlvane & Dube; Moore). This is also the view main-

tained in *LCB*. Indeed, "behavior analysis will be just as true and valid without a specification of the underlying neural mechanisms" (Moore, p. 244). Further, simulation research informed exclusively by behavioral studies is a valuable interpretive tool whether in the form of neural networks (viz., Hutchison) or otherwise (e.g., Epstein, 1984; Field, Tonneau, Ahearn, & Himeline, 1996; Hinson & Staddon, 1983a, 1983b). Nevertheless, for reasons outlined in *LCB* (see also Donahoe & Palmer, 1989), we believe that much is gained when simulations are also informed by neuroscience. The integration of behavior analysis and neuroscience is not only "in the long run essential" (Hutchison, p. 224), but is also encouraged in the nearer term by recent progress in neuroscience and network research (e.g., Barto & Sutton, in press). At the birth of behavior analysis, Skinner was wise to postpone building bridges to neuroscience; neither discipline was on a sufficiently firm footing to sustain the effort. Now, however, both sciences have matured to the point that such an effort benefits the progress of each.

*Neuroscience Constrains and Informs
Neural Network Research*

Neural networks are generally regarded as computationally universal. That is, they can compute any input–output function given unrestricted freedom in the choice of network architecture and connection weights (viz., Staddon). If the parameters that determine architectures and connection weights may vary indefinitely, then neural networks are little more than nonlinear multiple-regression techniques (Werbos, 1974), and the fact that they may simulate some environment–behavior relation is not a cause for celebration. However, if the characteristics of neural networks are constrained by independent experimental analyses of neuroscience as well as behavior, then the set of possible networks is reduced and we may regard a successful simulation as something more than an exercise in curve fitting. It is a scientific interpretation.

Genuine neuroscience-based constraints on neural networks are fast emerging, but few current simulations, including our own, are as tightly constrained as one might wish. For example, selection networks are constrained by neuroanatomical information about non-

specific dopaminergic projections to frontal lobes and hippocampus, and by cellular information about the effects of dopamine on synaptic efficacies in motor association and sensory association cortex. Even though these constraints are quite "gentle," critical pieces of the relevant neuroanatomical information are less than 10 years old. In addition, more detailed information about the connectivity of the nervous system (e.g., reciprocal connections between various cortical regions) has not yet been implemented, although the learning algorithm is quite capable of modifying such connections. The learning algorithm employed in our simulations is constrained by information about the cellular processes that are involved in long-term potentiation (LTP). Critical aspects of this information are less than 3 years old. And, our simulations capture only the net effects of these processes, not the intracellular events that produce them. Although the fruits of neuroscientific research must more fully inform neural network research, no simulation will ever implement all of the *potentially* relevant neuroscience, nor should it. Skinner's admonition to seek that level of analysis which most parsimoniously reveals orderly functional relations applies with the same force to neuroscience as to behavior analysis (Skinner, 1950).

The basis for our learning algorithm, derived from the literature on LTP, merits a few additional comments. We appear to have miscommunicated our views to some commentators, and others may have been similarly misled. A partial description of the proposed process follows: Initially, a presynaptic neuron releases glutamate, an excitatory neurotransmitter, that causes a burst of firing in the postsynaptic neuron through its effect on non-NMDA glutamate receptors. Bursting depolarizes the membrane of the postsynaptic cell long enough to allow magnesium to migrate from the NMDA glutamate channel, thereby permitting calcium to enter the cell. This sequence of events is, by itself, insufficient to produce long-lasting LTP, although it may produce shorter term effects. Only when the initial calcium influx is accompanied by intracellular events initiated by the neuromodulator dopamine are long-lasting changes induced in the postsynaptic non-NMDA glutamate receptors that

support LTP. These intracellular events critically include the catalysis of cAMP through stimulation of adenylyl cyclase via G proteins, although this process was not described in the target article (but see *LCB*, pp. 66–67). If, following the induction of long-lasting LTP, glutamate alone is presented, then the synaptic efficacy decreases, as reported by Stein and as required by our learning algorithm (see Donahoe *et al.*, 1993, Equation 5, p. 40).

Our learning algorithm is constrained by information from neuroscience and is not an inference from behavior (*viz.*, Kemp; Stein), although the effects of the algorithm must be consistent with observations of behavior. The learning algorithm is also not related to the neurally implausible back-propagation algorithm (*cf.* Dworkin & Branch) that is used in much connectionist research in cognitive psychology (see Donahoe & Palmer, 1989). We employ an algorithm that is the product of converging evidence at the cellular level, and a number of others have independently arrived at a similar conclusion (*e.g.*, Chioda & Berger, 1986; Frey, *in press*; Frey, Huang, & Kandel, 1993; Groves *et al.*, 1995; Wickens & Kötter, 1995). In agreement with Stein, whose experimental findings (Stein & Belluzzi, 1989) were greeted by us with joy because they opened the way for the coordinated modification by reinforcers of synapses in sensory association and motor association cortex, we believe that theorizing about synaptic mechanisms must be guided by neuroscience. We are avid consumers of that research, and we welcome efforts to determine how *in vitro* reinforcement (IVR) may be brought under stimulus control (Kemp; Stein).

Network Architecture and Neuroscience

Neural network research has been concerned primarily with the learning algorithm. The architecture of the neural network has usually been handcrafted by the researcher to carry out the task at hand when acted upon by the learning algorithm. However, because architecture determines the environment–behavior relations that networks can mediate (*e.g.*, Barnes & Hampson, *in press*), the choice of architecture must also be addressed in a principled manner. (Network architecture is an instance of the more general prob-

lem of structure with which behavior analysis has long grappled, and whose functional approach to structure separates our discipline from most of psychology and linguistics.) An overview of our program of simulation research on structure and function is shown in Figure 3 (Donahoe, *in press-a*). Three central areas are identified: behavioral research on the effects of the relations between environmental events and individual behavior (the environmental algorithm; see Palmer, *in press*), biobehavioral research on the selection of environment–behavior relations by the individual environment (the learning algorithm), and, most relevant for present purposes, biological research on selection by the ancestral environment of simulated genes that direct the structure of the network (the genetic algorithm). The genetic algorithm is not further characterized here except to indicate that it is informed by research on genetics and neural development.

One of us (Burgos, 1996, *in press*) has focused his recent simulation research on the genetic algorithm; primarily through that work, issues may be addressed about how our approach scales upward to larger networks and fares with respect to speed of learning (*viz.*, Dworkin & Branch; Hutchison; Marr). Findings with our hybrid genetic-learning algorithm indicate that connection weights of relatively large networks (over 100 units) may be appropriately modified by the learning algorithm to acquire the selected environment–behavior relations within a relatively small number (5 to 25) of reinforcers. Environment–behavior relations simulated by the environmental algorithm include conditioning with various CS–US intervals, multiple CSs, and various types of discriminations. Although the networks that have been evolved remain much less complex than the neuroanatomy of even a portion of the real nervous system (*viz.*, Marr), these preliminary results suggest that the approach is limited more by the demands for computing power than by an inability to accommodate complex environment–behavior relations.

We conclude by expressing our deep appreciation to the commentators. At first blush, we might have wished that all had immediately come to the conclusion expressed in Dinsmoor's title, "I agree." However, we would have been deprived of the opportunity

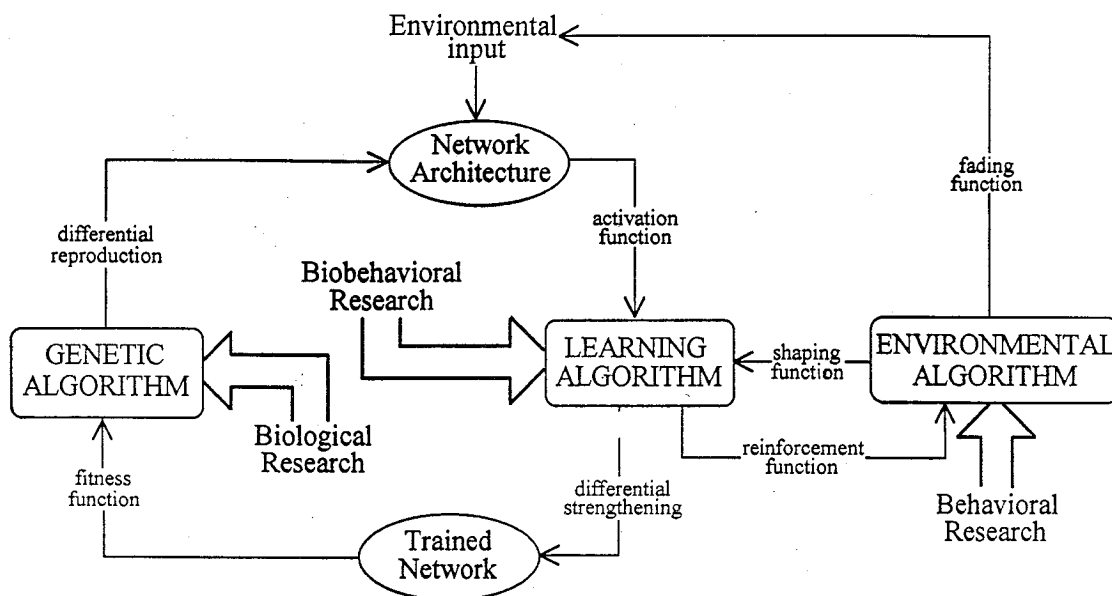


Fig. 3. Overview of a biobehavioral approach to experimental and simulation research. An organism, whose nervous system is simulated by a network architecture, is acted upon by its environment. The relation between environmental stimuli and the organism's behavior is the province of behavior analysis and is simulated by the environmental algorithm. The contingencies of reinforcement select environment-behavior relations and are the joint provinces of behavior analysis and neuroscience. These interactions and their effects are simulated by the learning algorithm. This process produces a population of organisms whose members may differ in their performance (i.e., their fitness). In accordance with the contingencies of survival, the genes that guide the architectures of the neural networks are differentially reproduced in proportion to the relative fitness of the organism. This process is the province of biological research, especially genetics and neurodevelopment, and is simulated by the genetic algorithm. Over repeated cycles of selection by the individual environment (simulated by the environmental and learning algorithms) and by the ancestral environment (simulated by the genetic algorithm), a population of adapted artificial neural networks can be produced. From Donahoe (in press-a).

to struggle with many good questions and to clarify and expand our views. It is our belief that most agree with at least the *promise* of an integration of behavior analysis and neuroscience. The task of providing an integrated account of environment-behavior relations is at a very early stage of development, and requires the concerted efforts of many if its promise is to be fulfilled. Finally, the foregoing technical discussion of selection by reinforcement should not overshadow the fact that less than half of *LCB* is devoted to these matters; the great majority of the book explores the implications of the approach for complex human behavior—equivalence classes, attending, perceiving, remembering, and verbal behavior.

REFERENCES

- Barnes, D., & Hampson, P. J. (in press). Connectionist models of arbitrarily applicable relational responding: A possible role for the hippocampal system. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations*. Amsterdam: Elsevier.
- Barto, A. G., & Sutton, R. S. (in press). Reinforcement learning in artificial intelligence. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations*. Amsterdam: Elsevier.
- Bear, M. F., & Singer, W. (1986). Modulation of visual cortical plasticity by acetylcholine and noradrenaline. *Nature*, *320*, 172-176.
- Berger, A. J., Gaspar, P. L., & Verney, R. N. (1991). Dopaminergic innervation of the cerebral cortex: Unexpected differences between rodents and primates. *Trends in Neurosciences*, *14*, 21-27.
- Branch, M. N., & Dworkin, S. I. (1981). Effects of ratio contingencies on responding maintained by schedules of electric shock presentation (response-produced shock). *Journal of the Experimental Analysis of Behavior*, *36*, 191-206.
- Burgos, J. E. (1996). *Computational explorations of the evolution of artificial neural networks in Pavlovian environments*. Unpublished doctoral dissertation, University of Massachusetts, Amherst.
- Burgos, J. E. (in press). Evolution of Pavlovian conditioning. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neu-*

- ral-network models of cognition: Biobehavioral foundations.* Amsterdam: Elsevier.
- Chioda, L. A., & Berger, T. W. (1986). Interactions between dopamine and amino-acid induced excitation and inhibition in the striatum. *Brain Research*, *375*, 198–203.
- Churchland, P. J., & Sejnowski, T. J. (1992). *The computational brain.* Cambridge, MA: MIT Press.
- Colwill, R. M., & Rescorla, R. A. (1986). Associative structures in instrumental learning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 20, pp. 55–104). New York: Academic Press.
- Davis, M., Rainie, D., & Cassell, M. (1994). Neurotransmission in the rat amygdala related to fear and anxiety. *Trends in Neuroscience*, *17*, 208–214.
- Donahoe, J. W. (1984). Commentary: Skinner—The Darwin of ontogeny. *The Behavioral and Brain Sciences*, *7*, 287–288.
- Donahoe, J. W. (in press-a). The necessity of neural networks. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations.* Amsterdam: Elsevier.
- Donahoe, J. W. (in press-b). Selection networks: Simulation of plasticity through reinforcement learning. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations.* Amsterdam: Elsevier.
- Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). Selectionist approach to reinforcement. *Journal of the Experimental Analysis of Behavior*, *60*, 17–40.
- Donahoe, J. W., Crowley, M. A., Millard, W. J., & Stickney, K. A. (1982). A unified principle of reinforcement. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior* (Vol. 2, pp. 493–521). Cambridge, MA: Ballinger.
- Donahoe, J. W., & Palmer, D. C. (1989). The interpretation of complex human behavior: Some reactions to *Parallel Distributed Processing*. *Journal of the Experimental Analysis of Behavior*, *51*, 399–416.
- Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior.* Boston: Allyn & Bacon.
- Donahoe, J. W., & Wessells, M. G. (1980). *Learning, language, and memory.* New York: Harper & Row.
- Donegan, N. H., & Wagner, A. R. (1987). Conditioned diminution and facilitation of the UR: A sometimes opponent-process interpretation. In I. Gormezano, W. F. Prokasy, & R. F. Thompson (Eds.), *Classical conditioning: III. Behavioral, neurophysiological, and neurochemical studies in the rabbit* (pp. 339–369). Hillsdale, NJ: Erlbaum.
- Eccles, J. C. (1994). *How the self controls the brain.* New York: Springer-Verlag.
- Epstein, R. (1984). Simulation research in the analysis of behavior. *Behaviorism*, *12*, 41–59.
- Estes, W. K., & Skinner, B. F. (1941). Some quantitative properties of anxiety. *Journal of Experimental Psychology*, *29*, 390–400.
- Field, D. P., Tonneau, F., Ahearn, W., & Hinline, P. N. (1996). Preference between variable-ratio and fixed-ratio schedules: Local and extended relations. *Journal of the Experimental Analysis of Behavior*, *66*, 283–295.
- Frey, U. (in press). Cellular mechanisms of long-term potentiation: Late maintenance. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations.* Amsterdam: Elsevier.
- Frey, U., Huang, Y.-Y., & Kandel, E. R. (1993). Effects of cAMP simulate a late stage of LTP in hippocampus CA1 neurons. *Science*, *260*, 1661–1664.
- Georgopoulos, A. P. (1990). Neurophysiology of reaching. In M. Jeannerod (Ed.), *Attention and performance XIII* (pp. 227–263). Hillsdale, NJ: Erlbaum.
- Georgopoulos, A. P. (in press). Motor cortex: Neural and computational studies. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations.* Amsterdam: Elsevier.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227–247.
- Goldman-Rakic, P. E., Chafee, M., & Friedman, H. (1993). Allocation of function in distributed circuits. In T. Ono, L. R. Squire, M. E. Raichle, D. I. Perrett, & M. Fukada (Eds.), *Brain mechanisms of perception and memory: From neuron to behavior* (pp. 445–456). New York: Oxford University Press.
- Griffin, D. R. (1992). *Animal minds.* Chicago: University of Chicago Press.
- Groves, P. M., Garcia-Munoz, M., Linder, J. C., Manley, M. S., Martone, M. E., & Young, S. J. (1995). Elements of the intrinsic organization and information processing in the neostriatum. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 49–96). Cambridge, MA: MIT Press.
- Hernandez, L., & Hoebel, B. G. (1990). Feeding can enhance dopamine turnover in the prefrontal cortex. *Brain Research Bulletin*, *25*, 975–979.
- Hinson, J. M., & Staddon, J. E. R. (1983a). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, *39*, 25–47.
- Hinson, J. M., & Staddon, J. E. R. (1983b). Matching, maximizing, and hillclimbing. *Journal of the Experimental Analysis of Behavior*, *40*, 321–331.
- Hoebel, B. G. (1988). Neuroscience and motivation: Pathways and peptides that define motivational systems. In R. A. Atkinson (Ed.), *Stevens' handbook of experimental psychology* (pp. 547–625). New York: Wiley.
- Hull, C. L. (1934). The concept of the habit-family hierarchy and maze learning. *Psychological Review*, *41*, 33–54.
- Jenkins, H. M., & Moore, B. R. (1973). The form of the auto-shaped response with food and water reinforcers. *Journal of the Experimental Analysis of Behavior*, *20*, 163–181.
- Kamin, L. J. (1968). Attention-like processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the prediction of behavior* (pp. 9–31). Miami, FL: University of Miami Press.
- LeDoux, J. E. (1995). Emotions: Clues from the brain. *Annual Review of Psychology*, *46*, 209–235.
- Mackintosh, N. J. (1983). *Conditioning and associative learning.* New York: Oxford University Press.
- Morse, W. H., & Kelleher, R. T. (1977). Determinants of reinforcement and punishment. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 174–200). Englewood Cliffs, NJ: Prentice Hall.
- Palmer, D. C. (in press). Selectionist constraints on neural networks. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations.* Amsterdam: Elsevier.
- Palmer, D. C., Donahoe, J. W., & Crowley, M. A. (1985).

- Discriminated interresponse times. *Journal of the Experimental Analysis of Behavior*, 44, 301–313.
- Rescorla, R. A. (1968). Conditioned inhibition of fear. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 65–89). Halifax, Nova Scotia: Dalhousie University Press.
- Rescorla, R. A. (1991). Associative relations in instrumental learning: The eighteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, 43B, 1–23.
- Rescorla, R. A., & Colwill, R. M. (1989). Associations with anticipated and obtained outcomes in instrumental learning. *Animal Learning & Behavior*, 17, 291–303.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning* (Vol. 2, pp. 64–99). New York: Appleton-Century-Crofts.
- Schultz, W. (in press). Adaptive dopaminergic neurons report the appetitive value of environmental stimuli. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations*. Amsterdam: Elsevier.
- Schultz, W., Romo, R., Ljungberg, T., Mirenovicz, J., Hollerman, J. R., & Dickinson, A. (1995). Reward-related signals carried by dopamine neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233–248). Cambridge, MA: MIT Press.
- Shapiro, M. M., & Miller, T. M. (1965). On the relationship between conditioned and discriminative stimuli and between instrumental and consummatory responses. In W. F. Prokasy (Ed.), *Classical conditioning* (pp. 269–301). New York: Appleton-Century-Crofts.
- Shull, R. L. (1995). Interpreting cognitive phenomena: Review of Donahoe and Palmer's *Learning and Complex Behavior*. *Journal of the Experimental Analysis of Behavior*, 63, 347–358.
- Skinner, B. F. (1937). Two types of conditioned reflex: A reply to Konorski and Miller. *Journal of General Psychology*, 16, 272–279.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1945). The operational analysis of psychological terms. *Psychological Review*, 52, 270–277.
- Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, 57, 193–216.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. (1974). *About behaviorism*. New York: Random House.
- Smith, K. (1954). Conditioning as an artifact. *Psychological Review*, 61, 217–225.
- Sperry, R. W. (1990). *Brain circuits and the functions of the mind*. New York: Cambridge University Press.
- Stein, L., & Belluzzi, J. D. (1988). Operant conditioning of individual neurons. In M. L. Commons, R. M. Church, J. R. Stellar, & A. R. Wagner (Eds.), *Quantitative analyses of behavior* (Vol. 7, pp. 249–264). Hillsdale, NJ: Erlbaum.
- Stein, L., & Belluzzi, J. D. (1989). Cellular investigations of behavioral reinforcement. *Neuroscience and Biobehavioral Reviews*, 13, 69–80.
- Stein, L., Xue, B. G., & Belluzzi, J. D. (1993). A cellular analogue of operant conditioning. *Journal of the Experimental Analysis of Behavior*, 60, 41–53.
- Stein, L., Xue, B. G., & Belluzzi, J. D. (1994). In vitro reinforcement of hippocampal bursting: A search for Skinner's atoms of behavior. *Journal of the Experimental Analysis of Behavior*, 61, 155–168.
- Terry, W. S., & Wagner, A. R. (1975). Short-term memory for "surprising" versus "expected" unconditioned stimuli in Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 104, 122–133.
- Trapold, M. A., & Overmier, J. B. (1972). The second learning process in instrumental learning. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory* (pp. 427–452). Englewood Cliffs, NJ: Prentice Hall.
- Weisman, R. G., & Litner, J. S. (1969). Positive conditioned reinforcement of Sidman avoidance in rats. *Journal of Comparative and Physiological Psychology*, 68, 597–603.
- Werbos, P. J. (1974). *Beyond regression: New tools for prediction and analysis in the behavioral sciences*. Unpublished doctoral dissertation, Harvard University, Cambridge, MA.
- Wickens, J., & Kötter, R. (1995). Cellular models of reinforcement. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 187–214). Cambridge, MA: MIT Press.
- Williams, B. A., & Bell, M. C. (1996). Changeover behavior and preference in concurrent schedules. *Journal of the Experimental Analysis of Behavior*, 65, 513–526.
- Williams, D. R., & Williams, H. (1969). Auto-maintenance in the pigeon: Sustained pecking despite contingent nonreinforcement. *Journal of the Experimental Analysis of Behavior*, 12, 511–520.