

*SOME THOUGHTS ON THE S-R ISSUE AND  
THE RELATION BETWEEN BEHAVIOR ANALYSIS AND  
BEHAVIORAL NEUROSCIENCE*

JAY MOORE

UNIVERSITY OF WISCONSIN-MILWAUKEE

In a recent review of Donahoe and Palmer's *Learning and Complex Behavior* (1994), Shull (1995, pp. 353–354) questioned whether the emphasis of adaptive neural networks in their biobehavioral approach constitutes a return to an S-R psychology that is inconsistent with the behavior-analytic conception of operant behavior. Donahoe, Palmer, and Burgos believe that it does not constitute such a return, and seek to allay any concerns by clarifying certain aspects of the biobehavioral approach. I would like to examine two of the many highly interesting issues raised in their discussion. The first is fairly general: the relation between behavior analysis and behavioral neuroscience. The second follows from the first, but is more specific: the relations among the interpretation of behavior, the experimental analysis of behavior, and levels of analysis. I am not sure I can resolve any unresolved matters, but perhaps I can contribute to their discussion from a perspective that might enable someone else to do so.

*The Relation Between Behavior Analysis and Behavioral Neuroscience*

It seems to me that the relation between behavior analysis and behavioral neuroscience is often misunderstood; I suggest that they are complementary. To be sure, Skinner's *Behavior of Organisms* (1938) was a "declaration of independence from physiology" (Skinner, 1995, p. 157) in that it called for behavior to be regarded as a subject matter in its own right. Nevertheless, a knowledge of physiology is certainly not irrelevant to the study of behavior, and a thoroughgoing behavior analysis does not imply such irrelevance. Skinner's argument, in 1938 and subsequently, was that behavior is not to be

treated merely as evidence for inferences about entities from neural, mental, or conceptual dimensions, and that such entities are not to be given special causal status. The contribution of physiology will be to show how exposure to such factors as contingencies of reinforcement changes an organism, and how it is that the changed organism then behaves differently at a later date. The information from physiology will come from methods appropriate to it as an independent discipline, and not as inferences from the very same behavior that it is supposed to explain (e.g., Skinner, 1974, pp. 218–223).

Thus, it seems to me that a comprehensive science of behavior is appropriately concerned with two issues. The first is: How is an organism's behavior functionally related to its environment? The second is: How do the neurophysiological systems of the organism mediate those functional relations? The first issue is the province of behavior analysis. It contributes to a science of behavior by analyzing the control exerted by contingencies operating at the phylogenetic, ontogenic, and cultural levels. The second issue is the province of behavioral neuroscience. As Skinner noted,

A behavioral analysis has two necessary but unfortunate gaps—the spatial gap between behavior and the variables of which it is a function and the temporal gap between the actions performed upon an organism and the often deferred changes in its behavior. These gaps can be filled only by neuroscience, and the sooner they are filled, the better. (Catania & Harnad, 1988, p. 470)

Given that behavior analysis and behavioral neuroscience are viewed as complementary, the basis for the complementarity still needs to be clarified. I suggest that it is pragmatic, rather than logical or reductive. That is, once we know how physiological inner states are functionally related to behavior, then predic-

Address correspondence and reprint requests to Jay Moore, Department of Psychology, University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201.

tions about an organism's behavior (or even interventions aimed at control) may be based on information about the current status of those inner states, rather than on a possibly inadequate specification of the history responsible for the states (e.g., Skinner, 1953, p. 34; 1969, p. 283; 1974, p. 221). If we raise the question of whether psychology needs to consider underlying neural mechanisms to become adequate, we have departed from our pragmatic concerns and have made an excursion into an epistemological or reductionistic dimension, where one sort of knowledge is valid if and only if it can be grounded on another.

In any event, if we continue on this theme of the nature of the relation between behavior analysis and traditional biological science, we can also see that a parallel exists between the history of evolutionary theory and that of behavior analysis (e.g., Catania, 1987, p. 255; Donahoe, Burgos, & Palmer, 1993, pp. 18–19; Skinner in Catania & Harnad, 1988, p. 111). The history of evolutionary theory suggests that Darwin's original ideas about evolution and Mendel's statistical laws became fully influential only some 60 years after they were first presented, when biochemists and others provided a persuasive account of how genetic mechanisms and DNA mediated evolution through natural selection. These events are often referred to as the "modern synthesis." The parallel is that behavior analysis now awaits its own modern synthesis, in which a specification of the neural mechanisms underlying the processes by which reinforcement selects behavior will presumably provide the answers for behavior analysis that are analogous to the answers that the gene and DNA provided for Darwin's and Mendel's theories.

I raise the point about the parallel histories because, given Donahoe and Palmer's (1994) biobehavioral approach to complex behavior, it is important to ask whether adaptive neural networks are the best candidate for the neural mechanisms and for filling the "two gaps" inherent in a behavior-analytic account. Donahoe et al. argue that adaptive neural networks are precisely such a candidate. This stance raises the issue to which we may now turn.

*The Relations Among the Interpretation of Behavior, the Experimental Analysis of Behavior, and Levels of Analysis*

Donahoe (1993, p. 453) defines *interpretation* as the use of principles derived from experimental analyses and constrained by formal (i.e., logical or mathematical) considerations to provide an account of events that occur under conditions that preclude controlled experimental analysis. Donahoe further suggests that the greater part of the scientific enterprise is interpretation, and that the greater part of Skinner's writings are interpretations rather than experimental analyses (e.g., Skinner, 1957). In this regard, Donahoe et al. clearly view adaptive neural networks as appropriate and meaningful *interpretive* devices for filling the two gaps in a behavior-analytic account:

Understanding is achieved through scientific interpretations that are constrained by experimental analyses of behavior and neuroscience. The most compelling interpretations promise to be those that trace the cumulative effects of reinforcement through formal techniques, such as adaptive neural networks, as a supplement to purely verbal accounts. (p. 193)

I see two critical matters in this approach. The first concerns levels of analysis: Are appeals to adaptive neural networks interpretive in the traditional behavior-analytic sense? Behavior-analytic *interpretations* traditionally account for behavioral events in terms such as discriminative stimuli, responses, and reinforcers, but without identifying those elements through controlled, formal experimental analysis. Behavior-analytic *explanations* traditionally remain at the level of behavior rather than, say, at the level of physiology. However, if interpretations are taken as explanatory activities, should they also remain at the level of behavior rather than at the level of neurophysiology? Is anything important lost when interpretations are at the level of underlying neural mechanisms, even if they are formally constrained by reputable principles?

The second critical matter concerns confirmation. Interpretations ordinarily apply principles that have been confirmed elsewhere, in independent analyses. However, many of the principles of the adaptive neural networks are

not yet confirmed through experimental analysis of the relevant physiology. For example, is there existing, independent physiological evidence that accounts for variability in the topography of an operant, as well as moment-to-moment variability in its rate or in its emission in the presence or absence of a discriminative stimulus? To be sure, Donahoe et al. readily acknowledge that further information confirming any contribution of adaptive neural networks will ultimately have to come from neuroscience itself. Nevertheless, I confess that I do not know the answers to the thorny questions noted above, and I anxiously await their resolution.

#### *Summary and Conclusions*

In conclusion, despite behavior analysts' convictions about the value of behavior analysis, the rest of the scientific community appears to be unconvinced. One need only compare the number of behavior analysts that are currently on the faculty of the most highly ranked, prestigious universities, to, say, the number of cognitive psychologists. What can be done to increase the acceptance of behavior analysis?

One possibility is to develop better techniques of prediction and control. However, in so doing we may fall victim to the fallacy of the "better mousetrap." The Dvorak keyboard has a demonstrably better arrangement of keys for a typewriter or computer terminal than does the traditional QWERTY keyboard, but it has not yet become dominant. Behavior analysis offers a range of effective techniques for prediction and control that have been available for many years, but it has not become dominant either. Thus, the answer seems to lie beyond the merit of techniques that are relevant to prediction and control. If you build a better mousetrap, the world does not necessarily beat a path to your door.

Another possibility, and the one that is relevant to the present discussion, would be to specify the underlying neural mechanisms that mediate the selection of behavior through reinforcement. Perhaps we should regard the question of the acceptance of behavior analysis as a behavioral question rather than as a question about logical validity. Behavior analysis will be just as true and valid without a specification of the underlying neu-

ral mechanisms, just as the work of Mendel is just as true without a specification of genetic structure and DNA. If the question is a behavioral one, then perhaps we need to view the problem as a shaping problem, and start at the level of the subjects whose behavior we want to shape. If people are more willing to accept behavior analysis when the underlying neural mechanisms that fill the gaps are specified, then perhaps we should consider providing the specification. To be sure, knowledge of the underlying mechanisms would also yield pragmatic benefits. This knowledge would ultimately open new avenues for the control of behavior, as in therapeutic interventions, although we are admittedly still a long way from interventions at this level.

In any case, I know of no better candidate at present for the underlying neural mechanisms than adaptive neural networks. Thus, they seem to be worth pursuing. It seems to me that we can do so with the recognition that models are a concrete context within which to examine a given phenomenon and assess the grain of truth of some statements that relate to the phenomenon in question. We need not embrace instrumentalist concerns that neural models, any more than any other kind of theory, are propaedeutic to scientific knowledge.

#### REFERENCES

- Catania, A. C. (1987). Some Darwinian lessons for behavior analysis: A review of Bowler's *The Eclipse of Darwinism*. *Journal of the Experimental Analysis of Behavior*, 47, 249-257.
- Catania, A. C., & Harnad, S. R. (Eds.). (1988). *The selection of behavior: The operant behaviorism of B. F. Skinner—Comments and controversies*. Cambridge: Cambridge University Press.
- Donahoe, J. W. (1993). The unconventional wisdom of B. F. Skinner: The analysis-interpretation distinction. *Journal of the Experimental Analysis of Behavior*, 60, 453-456.
- Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). A selectionist approach to reinforcement. *Journal of the Experimental Analysis of Behavior*, 60, 17-40.
- Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior*. Boston: Allyn and Bacon.
- 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. (1938). *The behavior of organisms: An experimental analysis*. New York: Appleton-Century-Crofts.

- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1969). *Contingencies of reinforcement*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1974). *About behaviorism*. New York: Knopf.
- Skinner, B. F. (1995). *The Behavior of Organisms* at fifty. In J. T. Todd & E. K. Morris (Eds.), *Modern perspectives on B. F. Skinner and contemporary behaviorism* (pp. 149–161). Westport, CT: Greenwood.

---

*THEORY AND BEHAVIOR ANALYSIS:  
COMMENTARY ON DONAHOE, PALMER, AND BURGOS*

J. E. R. STADDON

DUKE UNIVERSITY

---

The target article raises a number of interesting issues and comes to several conclusions with which most can readily agree. Operant and Pavlovian conditioning are measured with different procedures but are not completely different processes; Skinner's goal of explanation at the level of moment-by-moment behavior is a desirable one; and neurophysiology does not invalidate behavioral laws. I can add only a couple of comments.

First, although Skinner often urged moment-by-moment analysis ("Farewell my lovely!" and so forth), his consistent antagonism to real theory inhibited theories at that level. Because only "laws" (like Weber's law) seemed to be acceptable in behavior analysis, theory has for years been stuck at the level of molar laws. This development was not, as Skinner complained, a reaction against his ideas, but was in fact the only path he left open. After all, if all theory that "appeals to events taking place somewhere else, at some other level of observation, described in different terms, and measured, if at all, in different dimensions" (Skinner, 1950, p. 193) is prohibited, but we want to explain things anyway, then molar laws are all that is left. Skin-

ner was not worried by the fact that his prescription would have ruled out most of the great theoretical developments in physics and biology, from the atomic theory and the theory of the circulation of the blood through genetics and the wave theory of light. Almost every important theoretical advance in science has postulated "events taking place somewhere else [or] at some other level of observation." Donahoe et al. are quite right to insist on the necessity for real-time theory, but are wrong to credit Skinner with sympathetic anticipation of their proposal. Far from promoting the solution, Skinner's stance on this issue was part of the problem.

My second comment concerns the main point of the target article: whether reinforcement acts to strengthen responding or stimulus-response connections. This seems to be a straightforward empirical issue: Is operant learning context dependent or not? In other words, after training does responding decrease when the context is changed, or not? Is there a generalization gradient? The answer obviously is, "Almost always." With very few exceptions, operant learning in mammals and birds is subject to stimulus generalization decrement. Therefore, reinforcement must act not just on the response but also on its connection with context. Nevertheless, not all organisms show context dependence. The

---

Correspondence concerning this article should be addressed to J.E.R. Staddon, Department of Psychology: Experimental, Duke University, Durham, North Carolina 27708 (E-mail: staddon@psych.duke.edu).

operant-like behavior of orienting microorganisms, for example, doesn't seem to be under stimulus control (Staddon, 1983), and it is perfectly possible to design an operant mechanism that is context independent (Staddon & Zhang, 1991). So the question is certainly worth asking.

Finally a comment on the authors' question "Are neural networks capable of simulating the effects of nondifferential as well as differential operant contingencies?" (p. 202). As McCulloch and Pitts (1943/1965) showed many years ago, even very simple neural networks are general computing devices of the same order as the Turing machine, and hence are capable of simulating any well-defined process. The scientific issue, therefore, is not whether a given data set can be simulated by a neural network (it can), but whether a given simulation is the simplest and best—truest—model for that data set. What is

"true"? Francis Bacon quotes Jestling Pilate asking that question in another context, but Pilate "stayed not for an answer," perhaps because it is not a question that has (as the mathematicians say) a "closed-form solution."

#### REFERENCES

- McCulloch, W. S., & Pitts, W. H. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133. (Reprinted in W. S. McCulloch, *Embodiments of mind*, 1965, Cambridge, MA: MIT Press)
- Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, 57, 193–216.
- Staddon, J. E. R. (1983). *Adaptive behavior and learning*. New York: Cambridge University Press.
- Staddon, J. E. R., & Zhang, Y. (1991). On the assignment-of-credit problem in operant learning. In M. L. Commons, S. Grossberg, & J. E. R. Staddon (Eds.), *Neural networks of conditioning and action, the XIIIth Harvard Symposium* (pp. 279–293). Hillsdale, NJ: Erlbaum.

---

### BIOLOGICAL SUBSTRATES OF OPERANT CONDITIONING AND THE OPERANT-RESPONDENT DISTINCTION

LARRY STEIN

UNIVERSITY OF CALIFORNIA AT IRVINE

---

At the outset I should identify myself as a fellow advocate of the views of Donahoe and his colleagues—as someone who shares their selectionistic approach to behavior, admires their work, and embraces their positions on

---

This research was supported by grants from the National Institute on Drug Abuse (DA-05107 and DA-07747) and the Air Force Office of Scientific Research (89-0213). I thank my colleague James D. Belluzzi for invaluable contributions on a daily basis to all aspects of the experimental and theoretical work, and A. Harry Klopf and David W. Self for stimulating discussions over the years. Bao G. Xue and Mark Estacion provided neurophysiological advice and skillful technical assistance.

Address correspondence to Larry Stein, Department of Pharmacology, University of California, Irvine, California 92697.

many specific issues. In particular, I enthusiastically endorse the main organizing idea: that complex behavior "is best understood as the cumulative product of the action over time of relatively simple biobehavioral processes, *especially selection by reinforcement*" (p. 193, emphasis mine). And with regard to the important issue of the nature and complexity of the reinforced response, Donahoe et al. and I hold the same minority position. Together we reject the common supposition that the "whole" response (or its complex neural substrate) can be identified as the functional unit of reinforcement. Rather, we assume that the unit of reinforcement is some sort of infinitesimal response element or be-

havioral atom (Skinner, 1953), whose biological substrate is presumably cellular or sub-cellular in scale.

Nevertheless, our positions differ substantially with respect to the set of questions collectively termed “the S-R issue” in this essay. The central hypothesis of the authors—as aptly restated by Shull (1995)—“is that the fundamental effect of reinforcement is to select an environment–behavior relationship rather than to increase the emission rate of the reinforced response” (p. 353). Accordingly, the behavioral atoms of Donahoe and Palmer (1994) are conceived as elementary stimulus–response units, whose connection weights can be increased by the release of the reinforcement transmitter dopamine into the synapses of *coactive* pre- and postsynaptic elements. Furthermore, the unified reinforcement principle in their formulation applies equally to operant and Pavlovian conditioning. For these reasons, the Donahoe–Palmer hypothesis was characterized by Shull (1995) as “unconventional in some interesting respects” (p. 353). After all, the operant–respondent distinction is commonly thought to be at the heart of Skinner’s thinking on these matters, and Skinner (1953) has indicated explicitly that “Operant behavior, in short, is *emitted*, rather than *elicited*” (p. 107). Indeed, in later writings, Skinner (1981) even proposed that the evolution of operant conditioning must itself have required the parallel evolution and availability of “a supply of behavior . . . which has little or no relation to [eliciting or releasing] stimuli” (p. 501).

In response, Donahoe et al. acknowledge that, indeed—after Skinner and others had firmly established control by consequences as an empirical fact—the defining postulate of classical S-R psychology (that eliciting antecedents control all behavior) became untenable. But how can this conclusion be squared with the hypothesis that “what is selected is always an environment–behavior relation, never a response alone” (p. 196)? “The apparent incongruity,” the authors argue, “arises from a confusion of levels of analysis. . . . control by consequences (as opposed to antecedents) stands as a behavioral law, but we propose (at another level of analysis) that the effects of those consequences are implemented by changes in synaptic efficacies” (p. 196).

Although Donahoe et al. repeatedly assert “that it is a mistake to categorize accounts at the behavioral level by one’s view of the underlying biology” (p. 197), they nevertheless seek to reinterpret apparently contradictory neurophysiological findings from my own laboratory (e.g., Stein, 1994; Stein & Belluzzi, 1988, 1989; Stein, Xue, & Belluzzi, 1993, 1994). My colleagues and I found that the spontaneous bursting rate of individual cells in hippocampal slices was progressively increased in a dose-related manner by locally applied, burst-contingent microinjections of dopamine or other reinforcing transmitters or drugs. General pharmacological stimulation of bursting could be ruled out because the same injections given independently of bursting were ineffective, and because contingent (or noncontingent) injections of glutamate—an excitatory transmitter not associated with behavioral reinforcement—failed to increase and even suppressed hippocampal bursting. At the same time, the glutamate injections sharply increased the frequency of solitary spikes. We have interpreted these findings to mean that dopamine and other chemicals can increase the spontaneous bursting activity of neurons by a novel cellular mechanism (in vitro reinforcement or IVR) that is analogous to the strengthening of emitted behavior by reinforcing consequences. In short, IVR was conceptualized as a cellular analogue of operant conditioning (Stein et al., 1993).

Donahoe et al. offer “an alternative interpretation of these same facts . . . that is consistent with our view that reinforcers affect input–output relations and not output alone” (pp. 196–197). Their reinterpretation is based on the premise that IVR should not be viewed as a novel mechanism, but rather that it arises as a variation or manifestation of long-term potentiation (LTP, a well-established model of synaptic plasticity; see review of Bliss & Collingridge, 1993). According to this alternative explanation, the increased bursting induced by chemical reinforcement in the brain-slice experiments “may reflect a heightened sensitivity of the postsynaptic neuron to the release of the neurotransmitter glutamate by presynaptic neurons” (p. 197). Unfortunately, this ingenious and admirably detailed idea is probably incorrect; as already noted, applications of glutamate over a wide

dose range do not increase hippocampal bursting rates (as the alternative interpretation implies they should), but instead strongly suppress them.

In my opinion, IVR and LTP are not mere variations of a common mechanism of synaptic plasticity—they are separate processes. If so, it is possible that their dissimilar neurophysiological properties, when expressed in behavior, may underlie in part the operant-responder distinction. Although calcium-dependent signaling mechanisms are involved in both processes, IVR and LTP appear to depend on different types of  $\text{Ca}^{2+}$  channels. The relevant  $\text{Ca}^{2+}$  channel (NMDA channel) for LTP in the CA1 area is well established (Bliss & Collingridge, 1993); this receptor-operated  $\text{Ca}^{2+}$  channel is activated by glutamate in conjunction with membrane depolarization. On the other hand, the L-type  $\text{Ca}^{2+}$  channel seems to be critical for IVR (Stein et al., 1994); this voltage-gated  $\text{Ca}^{2+}$  channel is activated by membrane depolarization, but only if the channel protein is phosphorylated.

Evidence that implicates the L-type channel in IVR includes the following: (a) L-type channels control the generation of calcium spikes in hippocampal CA1 and CA3 neurons (Kostyuk, 1989), (b) L-type channels are located in hippocampal CA1 cell bodies and cluster in high density at the base of major dendrites (Westenbroek, Ahljianian, & Catterall, 1990), (c) influx of  $\text{Ca}^{2+}$  through hippocampal L-type channels regulates gene transcription through a distinct signaling pathway (Bading, Guity, & Greenberg, 1993), thus providing a possible mechanism for long-term reinforcement effects, and (d) L-type channels must be phosphorylated in order to open when the cell membrane is depolarized (Armstrong, 1989); this property could reasonably provide the hippocampal pyramidal cell with a mechanism for modulating calcium fluxes in response to external (reinforcing) signals. The tentative identification of the L-type channel as a significant protein target of the cellular reinforcement process has suggested a plausible and testable molecular hypothesis of IVR and (by extrapolation) operant conditioning (Stein, 1994).

My molecular model is based on the premise that the phosphorylation of L-type  $\text{Ca}^{2+}$

channels is the ultimate step in the reinforcement of hippocampal bursting responses.<sup>1</sup> More precisely, I propose that the normally rapid dephosphorylation of L-type channels after nonreinforced bursts is prevented by burst-contingent applications of dopamine or other reinforcing agents. Hippocampal bursts are made up of a few initial sodium spikes followed by a succession of calcium spikes (Jensen, Azouz, & Yaari, 1994; Schwartzkroin & Slawsky, 1977; Wong & Prince, 1978). The latter are mediated by voltage-gated L-type channels, which open in response to membrane depolarization (especially that produced by the  $\text{Na}^+$  spikes) if the  $\text{Ca}^{2+}$  channel protein is phosphorylated. The enzymatic addition or removal of phosphate esters alters the confirmation and thus changes the activity state of many nerve cell proteins (Nestler & Greengard, 1984).

#### *Molecular Hypothesis of In Vitro Reinforcement*

*Nonreinforced bursting.* Following a burst of nonreinforced calcium spikes, a “protective” intracellular cascade is activated to reduce the probability of further bursting (Figure 1) (Armstrong, 1989). This arrangement is thought to be self-protective, because each burst of calcium spikes introduces  $\text{Ca}^{2+}$  into the cell, and because high levels of intracellular  $\text{Ca}^{2+}$  are toxic. The burst-induced rise in intracellular  $\text{Ca}^{2+}$  activates the calcium-dependent enzyme calcineurin, which rapidly dephosphorylates the recently active  $\text{Ca}^{2+}$  channels that participated in the burst. Calcineurin can itself dephosphorylate the channel protein, but it acts mainly indirectly via inactivation of a key inhibitor (DARPP-32) of the principal dephosphorylating enzyme (phosphatase-1) of L-type channels. Gluta-

<sup>1</sup> Hippocampal pyramidal cells are not thought to be the sole targets of the in vitro reinforcement process. Following Skinner (1953), we have proposed that operant behavior arises from the collective action of a population of infinitesimal response elements or behavioral atoms that serve as the functional units of reinforcement (Stein & Belluzzi, 1982, 1988; Stein et al., 1994). At the biological level, atoms of behavior are assumed to arise from the bursting activity of specialized “reinforceable” neurons (e.g., hippocampal CA1 cells) that are localized in the target fields of dopamine, opioid peptide, or cannabinoid reinforcement systems. Reinforceable neurons thus have a wide distribution in the brain.

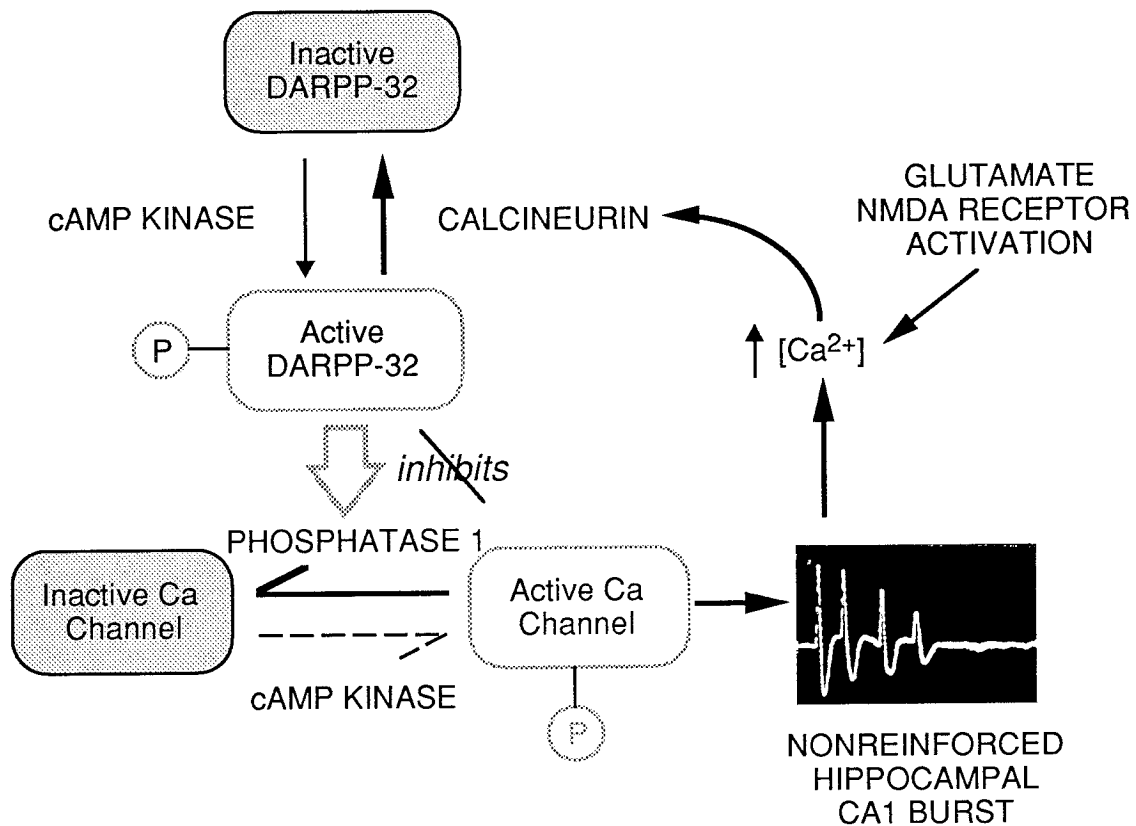


Fig. 1. Inactivation (dephosphorylation) of calcium channels by nonreinforced bursting or glutamate. Elevation of intracellular  $\text{Ca}^{2+}$  by a nonreinforced burst of calcium spikes or application of glutamate activates the calcium-dependent phosphatase, calcineurin, which in turn dephosphorylates and thus inactivates DARPP-32 (dopamine- and cAMP-regulated phosphoprotein). Such inactivation of DARPP-32 indirectly dephosphorylates the L-type calcium channels by releasing the principal dephosphorylating enzyme (phosphatase-1) of these channels from inhibition. Phosphorylation is denoted by the letter P in a circle.

mate also elevates intracellular  $\text{Ca}^{2+}$  by stimulation of NMDA receptors. Effective doses of this transmitter would therefore activate the calcineurin pathway, and the resulting dephosphorylation of  $\text{Ca}^{2+}$  channels should reduce hippocampal bursting rates, as we in fact have found (Stein et al., 1993). The glutamate-induced increase in the frequency of single spikes (i.e.,  $\text{Na}^+$  spikes), which we also observed, may be explained by the simultaneous stimulation of non-NMDA (AMPA) receptors that control sodium channels.

*Reinforced bursting.* According to my hypothesis, if bursting responses are closely followed by stimulation of dopamine D1/D5 receptors, the protective calcineurin pathway will be overridden and dephosphorylation of the recently activated  $\text{Ca}^{2+}$  channels is thereby

prevented (Figure 2).<sup>2</sup> The central feature of any useful hypothesis of IVR must be a plausible molecular explanation of the burst-dependent nature of dopamine's reinforcing action. As in the case of behavioral reinforcers, response-independent applications of dopamine are not reinforcing; hence, in our account, they should not prevent the dephosphorylation of  $\text{Ca}^{2+}$  channels. What is needed most critically to complete the hypothesis is a molecular coincidence detector, able to respond selectively to the conjunction of burst-

<sup>2</sup> Activation of dopamine D2, D3, and possibly D4 receptors is also reinforcing (Stein et al., 1994), presumably via an alternative pathway for the phosphorylation of calcium channels involving the arachidonic acid cascade (Piomelli et al., 1991), but this aspect of the hypothesis is beyond the scope of the present paper.

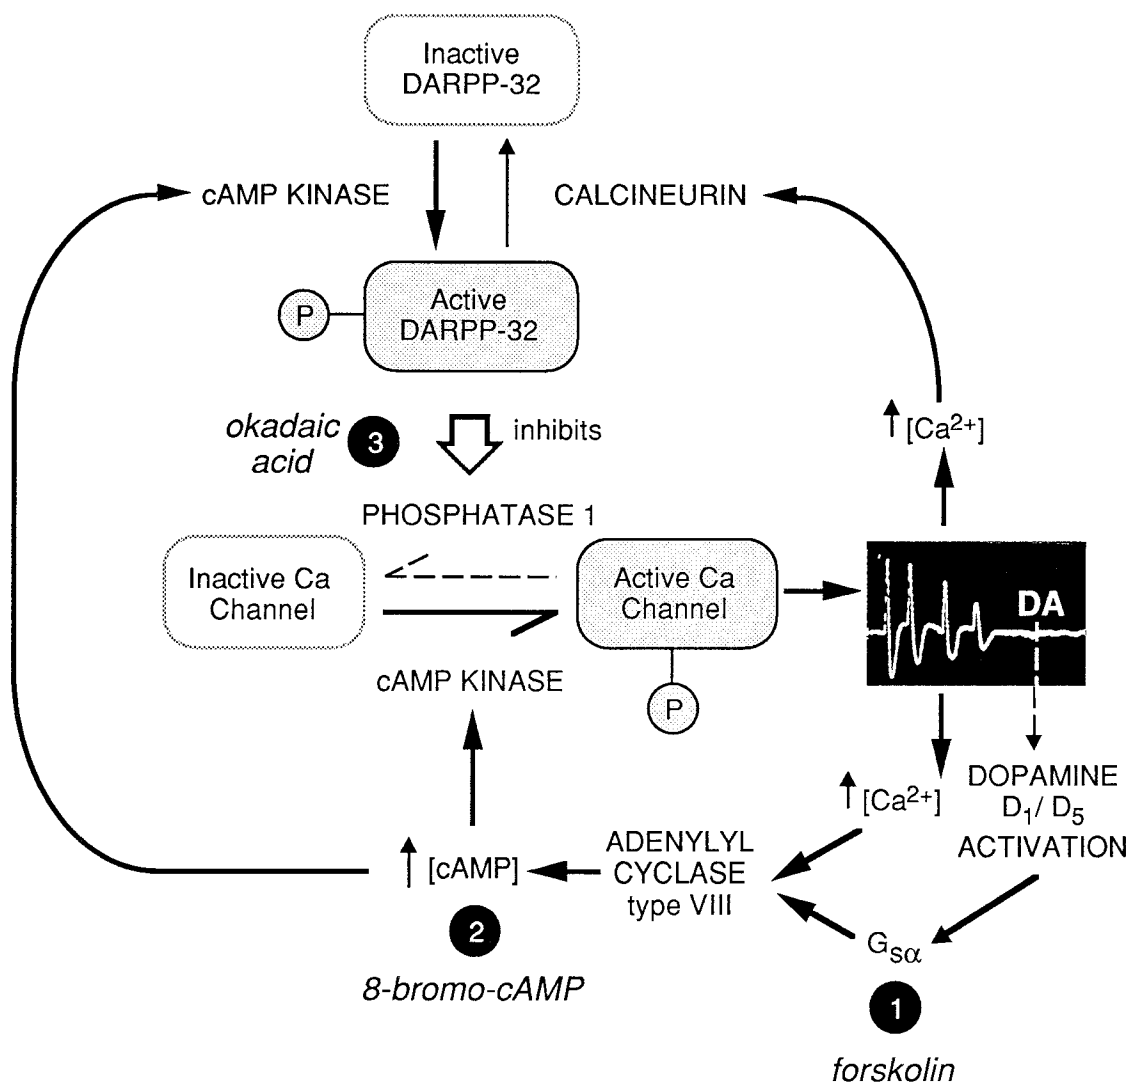


Fig. 2. Molecular hypothesis of in vitro reinforcement: Burst-contingent application of dopamine (DA) overrides the calcineurin cascade shown in Figure 1 and thus prevents the dephosphorylation of recently active L-type calcium channels. The conjunction of burst-induced  $\text{Ca}^{2+}$  elevation and dopamine D<sub>1</sub>/D<sub>5</sub> receptor stimulation (via the stimulatory G protein subunit,  $G_{s\alpha}$ ) synergistically activates type VIII adenylyl cyclase, causing a sharp rise in cAMP. The resulting activation of cAMP kinase overrides  $\text{Ca}^{2+}$ -channel dephosphorylation both directly, and also indirectly, by activation of the potent phosphatase-1 inhibitor, DARPP-32. For further explanation, see text.

ing activity and dopamine receptor activation. I propose that the enzyme type VIII adenylyl cyclase performs this function.

Adenylyl cyclase, the enzyme that synthesizes cyclic AMP (cAMP), is the prototypical second messenger generator; indeed, the concept of signaling by second messengers originated with the discovery of the role of cAMP (for reviews, see Nestler & Greengard, 1984, and Cooper, Mons, & Karpen, 1995).

The ubiquitous cAMP-dependent protein kinase (cAMP kinase) pathway accounts for the hormonal control of many cellular events; included among these is the phosphorylation of L-type  $\text{Ca}^{2+}$  channels. Eight types of adenylyl cyclases have been cloned to date, and each has been shown to be regulated by a variety of influences. Surprisingly, one of these adenylyl cyclases—type VIII—exhibits the precise biochemical properties and re-




Fig. 3. Synergistic action of Ca<sup>2+</sup>-calmodulin and the activated  $\alpha$ -subunit of stimulatory G protein (G<sub>s $\alpha$</sub> ) on type VIII adenylyl cyclase. Bars show mean cAMP accumulation from embryonal cells expressing rat type VIII adenylyl cyclase after no treatment (basal), addition of each regulatory molecule alone, or combined treatment (both). After Cali et al. (1994).

gional and subcellular localization required for our molecular coincidence detector. The highest concentrations of type VIII immunoreactivity are found postsynaptically in hippocampal CA1 dendritic spines "in intimate association with sites of calcium ion entry into the cell" (Cooper et al., 1995, p. 421); furthermore, type VIII is the only member of the adenylyl cyclase family that responds synergistically to Ca<sup>2+</sup> and dopamine (via the stimulatory G protein, G<sub>s</sub>) (Cali, Zwaagstra, Mons, Cooper, & Krupinski, 1994) (Figure 3). Thus, it can be anticipated that the conjunction of burst-induced Ca<sup>2+</sup> elevation and dopamine D1/D5 receptor stimulation would readily and selectively activate type VIII adenylyl cyclase; the cAMP generated would need to diffuse only a short distance before activating its enzymatic target (cAMP kinase), known to be anchored in high concentrations in the same dendritic spines. Such activation would override the calcineurin cascade and prevent the dephosphorylation of the Ca<sup>2+</sup> channels in those spines (Figure 2). Finally, the important negative action of burst-independent dopamine is nicely explained by the fact that, in the absence of elevated Ca<sup>2+</sup>, the response of type VIII adenylyl cyclase to D1/D5 activation will be inadequate.

The model is being tested at various steps of the proposed cascade. Three such IVR

tests (each denoted by a black circled number in Figure 2) are in progress. In the first test, microinjections of forskolin (a G<sub>s</sub>-mimicking activator of adenylyl cyclase) are substituted for dopamine as reinforcement for CA1 bursting. It is anticipated that forskolin will function as a typical *in vitro* reinforcer and will produce burst-contingent, but not burst-independent, increases in hippocampal bursting. Similar IVR experiments also are being performed with a membrane-permeant cAMP analogue (8-Br-cAMP, Test 2) and a phosphatase-1 inhibitor (okadaic acid, Test 3). Both 8-Br-cAMP and okadaic acid exert their effects at late stages of the phosphorylation cascade and thus bypass the coincidence-detecting action of type VIII adenylyl cyclase, which occurs at an earlier step. Hence, according to the model, these agents should not require a contemporaneous Ca<sup>2+</sup> signal to be effective, and both should facilitate bursting whether administered on a burst-contingent or burst-independent schedule. So far, all three tests have yielded promising results.

What are the main implications of this body of work for the S-R issue and the reinforcement hypothesis of Donahoe et al.? If eventually validated and extended to other target regions of the brain's reinforcement systems, the molecular model would provide a detailed explanation of dopamine's reinforcing action at the cellular and subcellular levels. The key feature is the precise specification of a coincidence-detecting molecule (e.g., type VIII adenylyl cyclase) that reacts selectively and exclusively "to the contiguity between the bursting of the postsynaptic neuron and the introduction of the neuromodulator" (p. 197). This two-term molecular interpretation of reinforcement thus seems quite consistent with Skinner's views on control by consequences and the operant-response (emission-elicitation) distinction, and it may indeed "more transparently parallel behavioral laws than the accounts we have offered" (p. 196).

Having said this, I should like to close with an expression of support for the emphasis Donahoe et al. correctly place on the difficult problem of operant discrimination and the role of the discriminative stimulus. Skinner (1953) himself has pointed out that operant behavior ubiquitously and "almost necessari-

ly” comes under the control of discriminative stimuli, according to a “three-term” contingency (p. 108). Nevertheless, he continues:

The relation between the discriminative operant and its controlling stimulus is very different from elicitation. Stimulus and response occur in the same order as in the reflex, but this does not warrant the inclusion of both types in a single “stimulus-response” formula. The discriminative stimulus does not elicit a response, it simply alters a probability of occurrence. (p. 110)

I am intrigued by the thought that my molecular hypothesis might be extended along the lines followed by Donahoe et al. to reinterpret our IVR experiments, but then used instead to suggest a biological explanation of operant discrimination. The broadened hypothesis focuses on the previously noted facts that (a) the initial spikes in a burst are non-reinforceable Na<sup>+</sup> spikes, and (b) these spikes can be generated by glutaminergic non-NMDA synaptic activation. The initial sodium spikes, of course, are not by themselves able to trigger the calcium spikes that complete the burst; in addition to depolarization of the postsynaptic membrane, a sufficient number of phosphorylated Ca<sup>2+</sup> channels must be available in the target zone. I propose that the initial sodium spikes in the burst may serve as the biological mediators of discriminative stimuli. Discriminative stimuli gradually assert their powerful control over behavior after many instances of differential reinforcement. The paired presentation of discriminative and reinforcing stimuli leads inevitably to their strong association through Pavlovian conditioning. Accordingly, the extended model must ensure that the sodium spikes are intimately associated with dopamine reinforcement, but the connection should be mediated by a Pavlovian (LTP-like) rather than an operant (IVR-like) mechanism. Donahoe et al. may have identified just such a process based on the LTP research of Frey, Huang, and Kandel (1993). Incorporation of their idea into the model provides it with a plausible Pavlovian-like mechanism by which dopamine D1/D5 activation could act in conjunction with NMDA stimulation to enhance the sodium spiking of the postsynaptic cell to coactive presynaptic (non-NMDA) glutaminergic input. If so, it perhaps would be interesting to employ these mechanisms in a

computer simulation of discriminated behavior.

#### REFERENCES

- Armstrong, D. L. (1989). Calcium channel regulation by calcineurin, a Ca<sup>2+</sup>-activated phosphatase in mammalian brain. *Trends in Neuroscience*, *12*, 117–122.
- Bading, H., Guity, D. D., & Greenberg, M. E. (1993). Regulation of gene expression in hippocampal neurons by distinct calcium signalling pathways. *Science*, *260*, 181–186.
- Bliss, T. V. P., & Collingridge, G. L. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, *361*, 31–39.
- Cali, J. J., Zwaagstra, J. C., Mons, N., Cooper, D. M. F., & Krupinski, J. (1994). Type VIII adenylyl cyclase. *Journal of Biological Chemistry*, *269*, 12190–12195.
- Cooper, D. M. F., Mons, N., & Karpen, J. W. (1995). Adenylyl cyclases and the interaction between calcium and cAMP signalling. *Nature*, *374*, 421–424.
- Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior*. Boston: Allyn & Bacon.
- Frey, U., Huang, Y.-Y., & Kandel, E. R. (1993). Effects of cAMP simulate a late stage of LTP in hippocampal CA1 neurons. *Science*, *260*, 1661–1664.
- Jensen, M. S., Azouz, R., & Yaari, Y. (1994). Variant firing patterns in rat hippocampal pyramidal cells modulated by extracellular potassium. *Journal of Neurophysiology*, *71*, 831–839.
- Kostyuk, P. G. (1989). Diversity of calcium ion channels in cellular membranes. *Neuroscience*, *28*, 253–261.
- Nestler, E. J., & Greengard, P. (1984). *Protein phosphorylation in the nervous system*. New York: Wiley.
- Piomelli, D., Pilon, C., Giros, B., Sokoloff, P., Martres, M.-P., & Schwartz, J.-C. (1991). Dopamine activation of the arachidonic acid cascade as a basis for D1/D2 receptor synergism. *Nature*, *353*, 164–167.
- Schwartzkroin, P. A., & Slawsky, M. (1977). Probable calcium spikes in hippocampal neurons. *Brain Research*, *135*, 157–161.
- 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. (1953). *Science and human behavior*. New York: Free Press.
- Skinner, B. F. (1981). Selection by consequences. *Science*, *213*, 501–504.
- Stein, L. (1994). In vitro reinforcement of hippocampal bursting—possible cellular and molecular mechanism of drug reward. *Regulatory Peptides*, *54*, 285–286.
- Stein, L., & Belluzzi, J. D. (1982). Beyond the reflex arc: A neuronal model of operant conditioning. In A. Morrison & P. L. Strick (Eds.), *Changing concepts of the nervous system* (pp. 651–655). New York: Academic 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 & 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.
- Westenbroek, R. E., Ahljianian, M. K., & Catterall, W. A. (1990). Clustering of L-type  $\text{Ca}^{2+}$  channels at the base of major dendrites in hippocampal pyramidal neurons. *Nature*, *347*, 281–284.
- Wong, R. K. S., & Prince, D. A. (1978). Participation of calcium spikes during intrinsic burst firing in hippocampal neurons. *Brain Research*, *159*, 385–390.

---

### MELIORATION AND CONTIGUITY

WILLIAM VAUGHAN, JR.

CHEBEAGUE ISLAND, MAINE

---

Donahoe, Palmer, and Burgos make a number of arguments: Molar relations should be understood as the outcome of local processes; reinforcement is not simply the strengthening of responses but also involves the stimuli present at that time; operant and classical conditioning are not distinct, but are separated only on the basis of what kinds of events are reliably present when reinforcement is presented; modeling (in this case by means of a neural network) can be productive in terms of integrating a number of behavioral phenomena.

A number of these issues tie in with an account of melioration (Herrnstein & Vaughan, 1980) in terms of strengthening by contiguity. Consistent with Donahoe et al., I believe that it is possible to derive melioration from the more basic processes advocated by Skinner. In an experiment using concurrently available alternatives, an alternative can gain value by pairing with reinforcement, whether the reinforcement is response produced or not (e.g., using concurrent variable-time [VT] VT schedules and only requiring a changeover response); time spent without reinforcement in the presence of that alternative drives its value toward zero. From these assumptions,

one can deduce that the value of an alternative is a strictly monotonically increasing function of rate of reinforcement in its presence. Given two or more such alternatives, changeover responses can then be viewed as increasing or decreasing in strength, depending on whether they make a transition from a lower to a higher, or from a higher to a lower, situation. This strengthening model (presented in Vaughan, 1982), in a nutshell, allows one to deduce the process of melioration, and in turn account for behavior on concurrent variable-interval (VI) VI, concurrent VI variable-ratio (VR), and concurrent VR VR schedules. The fact that a changeover delay is often required to prevent rapid alternation, with a duration similar to the duration of unsignaled delays that will reduce responding to a low level (Williams, 1976), suggests that the strength of changeovers is being maintained by the transitions from one conditioned reinforcer to another, rather than by food presentations on the alternative being changed to.

On the other hand, Donahoe et al.'s argument that operant and classical conditioning are the same processes, distinguished only by what event is reliably contiguous with reinforcement, may require some modification. For example, consider the Rescorla-Wagner model (Rescorla & Wagner, 1972),

---

Correspondence concerning this article should be addressed to William Vaughan, Jr., South Road, Chebeague Island, Maine 04017 (E-mail: wvaughan@nllis.net).

which states that the presentation of a reinforcer changes the value of stimuli present:

$$\Delta V_A = \alpha_A \beta_1 (\lambda_1 - V_{AX}) \quad (1)$$

Here,  $\Delta V_A$  is the change in value of Stimulus A,  $\alpha_A$  is a measure of the salience of A,  $\beta_1$  is a rate parameter related to this reinforcer,  $\lambda_1$  is the asymptotic level that the value of A can reach, and  $V_{AX}$  is the current value of A along with background Stimulus X. Let me first suggest a slightly simpler form of Equation 1:

$$\Delta V_A = f(V_B - V_A). \quad (2)$$

Here,  $V_A$  is the value of Stimulus A, and  $V_B$  is the value of reinforcement (the stimulus paired with A). In the case of the Rescorla-Wagner model, this is the value of all stimuli present when reinforcement occurs. That property (along with  $\alpha$  and  $\beta$ ) could be incorporated into Equation 2 but would serve no useful purpose here. The function  $f$  is simply assumed to be strictly monotonically increasing, with  $f(0) = 0$ .

If a pigeon pecks a key and there are no transitions, extinction takes place. For simplicity, one may assume one form of equation for both classical and operant extinction, which operates in the absence of transitions, but only while Stimulus (or Response) C is present:

$$dV_C/dt = f(V'_C - V_C). \quad (3)$$

Here,  $V'_C$  is what may be called the unconditioned value of Stimulus or Response C. In the case of a keylight or pecking a lighted key,  $V'_C$  is assumed to equal approximately zero, whereas in the case of food or a response whose operant level is above zero it would be positive; for shock it would be negative. If a pigeon is responding on two keys alternately, only the key currently being pecked is assumed to be governed by Equations 2 and 3; the other is not considered to be present. In this case there are also two other classes of responses, changeovers from Side A to Side B and changeovers in the opposite direction.

Suppose one wants to generalize Equation 2 to account for strength of responding. One might hypothesize an equation analogous to Equation 2, letting  $V_{RA}$  represent the value of a response during Stimulus A:

$$\Delta V_{RA} = f(V_B - V_{RA}). \quad (4)$$

Assume that a response with positive value

(or strength) occurs with some frequency that is a monotonic function of that strength, and that responses with zero or less than zero strength do not occur. It turns out that there are data that Equation 4 cannot handle. Dinsmoor (1962), for example, found that rats responded to escape a stimulus correlated with shocks and produced one not correlated with shocks. If  $V_{RA}$  is the value of responding and  $V_B$  is the value of the stimulus without shocks, such an equation would not predict any increase in the value of responding above zero, because the value of the stimulus produced by a response is zero. Consider now:

$$\Delta V_{RA} = f[(V_B - V_A) - V_{RA}]. \quad (5)$$

According to Equation 5, if Response RA is contiguous with a transition from Stimulus A to Stimulus B and there is a change in value upon making that transition, then the value of Response RA will approach the magnitude of that *change* in value,  $V_B - V_A$ , rather than the magnitude of the value of the stimulus being changed to. In Dinsmoor's (1962) experiment, for example, Stimulus A would have negative value, due to its pairings with shock (as a result of Equation 2). A response would produce a transition from A to B, a stimulus without shocks, and the value of that transition would be positive. Response RA, then, would gain positive value due to being contiguous with that positive transition.

By this account, then, classical and operant conditioning are closely related but distinct processes. It is logically possible that Equations 2 and 3 account for the value of a stimulus paired with reinforcement, and Equation 5 (in conjunction with Equation 3) accounts for the value of a response paired with a transition from one situation to another.

This conception is similar to that of Baum (1973), who treated reinforcement and punishment as cases of "situation transition." According to Baum (p. 151), reinforcement consists in the transition from a lower valued situation to a higher valued one, and punishment is a transition in the opposite direction. This analysis differs from his in that situations are assumed to gain and lose value by a process of classical conditioning, such as Equations 2 and 3 specify; Baum assumed that situations had no hedonic value but only discriminative values.

This local analysis of melioration is largely consistent with the position of Donahoe et al. Like theirs, this approach treats strength of responding as part of the three-term contingency. On the other hand, this analysis assumes two equations for learning, one for classical and one for operant behavior. Perhaps Equation 5 could be generalized to cover classical conditioning, in which case we would be back to a single process, but the equation would represent a break from the Rescorla–Wagner model.

## REFERENCES

- Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 20, 137–153.
- Dinsmoor, J. A. (1962). Variable-interval escape from stimuli accompanied by shocks. *Journal of the Experimental Analysis of Behavior*, 5, 41–47.
- Herrnstein, R. J., & Vaughan, W., Jr. (1980). Melioration and behavioral allocation. In J. E. R. Staddon (Ed.), *Limits to action* (pp. 143–176). New York: Academic Press.
- 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: II. Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Vaughan, W., Jr. (1982). Choice and the Rescorla-Wagner model. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 263–279). Cambridge, MA: Ballinger.
- Williams, B. A. (1976). The effects of unsignalled delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 26, 441–449.

---

WHAT IS LEARNED?  
REVISITING AN OLD ISSUE

BEN A. WILLIAMS

UNIVERSITY OF CALIFORNIA, SAN DIEGO

---

The authors of this provocative article argue that an S-R approach to behavior is an implicit assumption of connectionist network models of behavior. More exactly, their S-R model is better depicted as an S-O-R model, because a large part of variability in behavior comes not from variation in the stimulus input but from differences in activity, including spontaneous activity, between the nodes of the intervening network. As the authors make clear, these assumptions make it difficult to distinguish their S-R account from the traditional operant analysis in terms of the three-term contingency.

Although the authors make clear that their

type of S-R analysis does not necessarily require a causal explanation of behavior in terms of the particular stimulus that elicits each response, their approach does share with other S-R approaches (e.g., Thorndike, Hull) the assumption that the function of the reinforcer is to provide catalysis of S-R associations but not to enter into the associative relation itself. Donahoe et al.'s model utilizes the release of dopamine as the agent that increases the connection weights between different elements of the network. This so-called reinforcer, like the reinforcer in traditional S-R theory, remains outside of the associative network.

A critical issue posed for Donahoe et al. is whether a satisfactory model of conditioning can be constructed that omits any role for response–reinforcer associative relations. Al-

---

Correspondence concerning this article should be addressed to Ben A. Williams, Psychology Department, University of California San Diego, La Jolla, California 92093-0109 (E-mail: bawilliams@ucsd.edu).

though Skinner himself eschewed any such analysis, other prominent behavior analysts have questioned whether the role of the reinforcer can be understood simply in terms of catalysis of the S-R connections. For example, Catania (1984, p. 77) writes, "The consequences of responding are critical to our understanding of learning not because learning follows from them but because they are what is learned." Subsequent empirical developments in associative learning theory have strongly supported Catania's perspective that the response-reinforcer relation is an essential part of what is learned in operant behavior. Much of this work comes from studies of reinforcer devaluation. For example, Colwill and Rescorla (1986) trained separate responses using different reinforcers for each response, and then independently (with the responses no longer available) devalued one of the reinforcers by pairing it with poison. The subject was then returned to the training situation with both responses freely available but during extinction. The results were that only the response that led to the devalued reinforcer in the past was decreased in strength. The specificity of the effect seems most easily interpretable in terms of specific response-reinforcer associations. Donahoe et al. are clearly aware of the importance of these devaluation studies, in that Donahoe and Palmer (1994, pp. 108-109) provide an explanation of the basic effect. However, their treatment of the issue commits them to an interpretation in terms of two-factor theory, in which the animal's conditioned anticipation of the reward produces stimuli that assume discriminative control. Although such an interpretation has not been totally ruled out, variations in the procedure (see Colwill & Rescorla, 1986) that cause the same external stimulus to be present for both responses (e.g., with a bidirectional lever, with different reinforcers contingent on the different directions of the response) do seem to cause the two-factor perspective to be strained in its application.

Perhaps the most direct evidence that the discriminative stimulus serves to cue the response-reinforcer relation in effect, and thus cannot be reduced to being a direct elicitor of responding, comes from studies of stimulus control that use blocking as a tool for analysis (Colwill & Rescorla, 1990, Experi-

ment 3; Rescorla, 1990). In the pretraining phase of these studies, each of two responses led to one of two different reinforcers in the presence of a discriminative stimulus (A), and neither reinforcer was available during the absence of the stimulus. In the compound conditioning phase, two new stimuli were added to form two separate stimulus compounds, BA and CA. In the presence of BA, the response-reinforcer relations were the same as those used during pretraining, whereas in the presence of CA, the opposite response-reinforcer relations were in effect. Then Elements B and C were tested alone to determine if they had acquired stimulus control over responding. In both experiments, Element B failed to acquire stimulus control, indicating blocking, but Element C did acquire control. The most plausible interpretation of this pattern of results is that stimulus control is determined by whether the stimulus provides new information about the response-reinforcer relation. The challenge I would pose for Donahoe et al. is to demonstrate that their connectionist model can account for this array of results. Given that the reinforcer itself is not represented in the network, it is not obvious how such an account would proceed.

The issue of what terms enter into associative relations is also important to the authors' distinction between Pavlovian and instrumental conditioning. They are certainly correct that responses necessarily are occurring in the presence of stimuli at the time of reinforcement and that the major difference between the two kinds of contingencies is the degree of constraint on the co-occurrence of the different events. But a large amount of evidence argues that the associative relation in Pavlovian conditioning cannot be reduced to S-R connections. For example, if a tone is paired with food in an autoshaping procedure, it elicits orientation but no pecking. If the tone is then made contingent on the prior occurrence of a keylight, pecking to the keylight will occur. What appears to have occurred is that the tone has assumed the status of a substitute for food, and the behavior it elicits is some combination of the unconditioned effects of the reinforcer and of the particular stimulus that serves as the CS. A major assumption of S-R theory is that the behavior to a stimulus is what is learned, so

that the interpretation of Pavlovian conditioning in terms of sheer stimulus substitution omits a critical ingredient of the S-R associative relation. The evidence for a stimulus substitution interpretation, instead of an S-R association, is, to this reader, compelling (see Mackintosh, 1983, for a review and analysis).

Quite apart from the issue of what is learned, there are several other issues that prompt Donahoe et al. to endorse a theoretical position that is open to challenge. Whether these positions are intrinsic to their connectionist model is unclear. The most general is their commitment to a moment-by-moment analysis of phenomena such as the matching law in terms of changing stimulus states, even when the available evidence seems to support the molar interpretation. Because the molar-molecular issue has been debated extensively in the literature, it may be useful to consider some of its implications for their connectionist model.

A critical issue is the nature of the choice rule that governs response selection. S-R analyses implicitly assume that the rule is maximizing, or winner-take-all, in that whichever response has the greatest associative connection, given the stimulus set, will be evoked. In contrast is the matching rule, which assumes that the different responses are emitted proportionally to their strength. The issue is not yet resolved (see Williams, 1994, for a discussion), but Donahoe et al. argue that the evidence favors the momentary maximizing approach when the two different perspectives are pitted against each other. I disagree with this interpretation of the evidence. Although it is true that momentary contingencies, defined by fluctuating stimulus states, can control behavior when they are sufficiently strong, it does not appear to be the case that the molar phenomena that have dominated operant research over the past 30 years are easily reduced to the molecular effects. The results of Williams (1991) provide an example. In that study rats were trained on a discrete-trial probability learning task in which the local contingencies always favored win-shift and lose-shift response patterns (i.e., alternation), but were asymmetric for the two responses, causing the molar reinforcement rates for the two responses to differ. The results were that some modicum of control by the local contingencies did occur with short

intertrial intervals, but these contingencies were discriminated very poorly with longer intertrial intervals. Most important, approximation to the matching law was much closer with the longer ITIs, when the local contingencies had little control over behavior. The issue posed for Donahoe et al. is how to provide a moment-by-moment analysis of behavior, in this case choice behavior, when the contingencies associated with the local stimulus environment apparently cannot account for the obtained pattern of behavior.

An especially challenging example of the inadequacy of molecular models is provided by Neuringer (1992). Rats' responding was reinforced on separate VI schedules when that behavior met a criterion for repetition or for randomness. The result was approximate matching of the two types of behavior to their relative reinforcement frequencies. Given that one of the response alternatives could not be predicted on a moment-by-moment basis, and that prior work (Page & Neuringer, 1985) had made a strong case that learning to emit random behavior could not be the result of a memory-based strategy, the functional equivalence of such behavior with normal behavior with known controlling stimuli poses a major conceptual challenge for any type of S-R analysis.

A final issue that deserves consideration is the status of associationism in general. Gallistel (1990) has provided a strong cogent argument that learning does not consist of the formation of associative links, but instead consists of representations of the structure of the environment. In other words, the animal acquires veridical knowledge of the spatial and temporal properties of important events, which is mediated only indirectly by how those events make contact with behavior via a reinforcement contingency. For example, in his view, the matching law is the result of the animal having knowledge of the different distributions of interreinforcement intervals contingent on each response alternative, and is not the result of the number of pairings of the response and reinforcer, or of the probability that a response will be followed by a reinforcer. Substantial evidence (Gibbon, Church, Fairhurst, & Kacelnik, 1988; Mark & Gallistel, 1994; Williams, 1993) now exists in favor of such an analysis. Obviously, it is premature to pass judgment on such an impor-

tant overarching issue. Whether associative models can deal with the evidence compiled by Gallistel (1990) remains an open question, and only direct simulations of those findings will provide an answer. It is my hope that Donahoe et al. will take up this challenge and the others previously described. Connectionist networks will rise or fall as psychological explanations to the extent that they can simulate the properties of real behavioral events.

#### REFERENCES

- Catania, A. C. (1984). *Learning* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall.
- 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.
- Colwill, R. M., & Rescorla, R. A. (1990). Evidence for the hierarchical structure of instrumental learning. *Animal Learning & Behavior*, 18, 71–82.
- Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior*. Boston: Allyn and Bacon.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gibbon, J., Church, R. M., Fairhurst, S., & Kacelnik, A. (1988). Scalar expectancy theory and choice between delayed rewards. *Psychological Review*, 95, 102–114.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford: Oxford University Press.
- Mark, T. A., & Gallistel, C. R. (1994). Kinetics of matching. *Journal of Experimental Psychology: Animal Behavior Processes*, 20, 79–95.
- Neuringer, A. (1992). Choosing to vary and repeat. *Psychological Science*, 3, 246–250.
- Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, 11, 429–452.
- Rescorla, R. A. (1990). The role of information about the response-outcome relation in instrumental discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 16, 262–270.
- Williams, B. A. (1991). Choice as a function of local versus molar reinforcement contingencies. *Journal of the Experimental Analysis of Behavior*, 56, 455–473.
- Williams, B. A. (1993). Molar versus local reinforcement probability as determinants of stimulus value. *Journal of the Experimental Analysis of Behavior*, 59, 163–172.
- Williams, B. A. (1994). Reinforcement and choice. In N. J. Mackintosh (Ed.), *Animal learning and cognition* (pp. 81–108). San Diego: Academic Press.