TIMING WITHOUT A TIMER

John W. Donahoe and José E. Burgos

UNIVERSITY OF MASSACHUSETTS AT AMHERST AND
UNIVERSIDAD CATÓLICA ANDRÉS BELLO AND
UNIVERSIDAD CENTRAL DE VENEZUELA

Inferred-process theories propose various internal timing mechanisms to provide discriminative stimuli for temporal control of behavior. Biobehaviorally informed computer simulations indicate that timing may emerge from the action of reinforcement on neural circuits without the need to postulate timers of any sort.

Key words: timing, neural networks, neuroscience, fixed-interval schedule

Three separable aspects of the Staddon-Higa article may be distinguished: (a) a general critique of accounts of temporal regularities in behavior that appeal to inferred entities such as ''internal clocks'' or ''pacemakers,'' (b) a proposal for an alternative approach to temporal control—the multiple-time-scale (MTS) model (Staddon & Higa, 1996), and (c) a comparison of MTS with a particular inferred-process theory—scalar expectancy theory (SET) (Gibbon, 1977; Gibbon & Church, 1990). The first and second aspects of the article are addressed in this commentary. (As a preliminary matter, we use the phrase *temporal control* as shorthand for an orderly relation between behavior and time measured from some event. Clearly, time itself is not a stimulus. There is no environmental or neural event that uniquely and uniformly covaries with time from an *arbitrary* stimulus and, hence, no possibility that natural selection could have produced a neural system for its detection.)

Inferred-Process Theories

Staddon and Higa reiterate the behavior-analytic critique of attempts to explain regularities in behavior by invoking structures and processes that are based solely on inferences from behavior. In the case of SET, functional relations between behavioral measures (rate and choice) and time from the onsets of stimuli serve as the basis for inferring that behavior is controlled by the output of an internal pacemaker. The recourse to inferred processes is a hallmark of cognitive theorizing generally (e.g., Atkinson & Shiffrin, 1968), and SET fits squarely within that tradition. In contrast, behavior analysis regards theories employing inferred processes as ''having all the virtues of theft over honest toil,'' as Bertrand Russell remarked in another context. Staddon and Higa note that theories of this sort postulate that which they seek to explain. Such theories endow the organism with the very capabilities that are needed to encompass the relation in question. In so doing, inferred-process theories embrace an essentialist view of science (Palmer & Donahoe, 1992) that impedes genuine explanation and that inadvertently encourages circular reasoning and the nominal fallacy (Donahoe & Palmer, 1994, pp. 9, 152; Skinner, 1938).

Temporal control presents special temptations to inferred processes: Behavior changes in an orderly fashion in the absence of a correlated change in the external environment. Consider responding during a fixed-interval schedule. Following the reinforcing stimulus, it is reasonable to regard responding as under the control of the discriminative effects of the reinforcer because food never occurs during this period. The sequence of events instantiates a three-term contingency that is known to establish a stimulus as discriminative for infrequent responding. This account

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 01002 (E-mail: jdonahoe@psych.umass.edu) or to José E. Burgos, Apartado 48276, Los Chaquaramos, Caracas, 10041, Venezuela (E-mail: jburgos@ucab.edu.ve). A compiled version of the learning algorithm and copies of the files used to simulate fixed-interval behavior are on deposit at the Cambridge Center for Behavioral Studies, 1770 Massachusetts Avenue, 123, Cambridge, Massachusetts 02140.

of the postreinforcement pause in fixed-interval schedules meets the criteria for scientific interpretation: *Independent* evidence demonstrates that responding declines in frequency under conditions that are satisfied by the interpreted situation (see Donahoe & Palmer, 1994, pp. 126–129). Moreover, the interpretation can be experimentally evaluated because the event of concern—the reinforcing stimulus—can be manipulated. For example, the reinforcer can be introduced as a probe stimulus into the fixed interval at a time when responding would otherwise be substantial. Under these circumstances, the reinforcing stimulus brings responding to a halt, thereby documenting its discriminative effects. (Reinforcing stimuli have other than discriminative effects, but this complication does not detract from the point being made here; cf. Papini & Hollingsworth, 1998.)

Responses that occur somewhat later in the fixed interval present quite a different challenge. Independent experimental analysis does not indicate, in general, that putative discriminative stimuli can acquire control over responses when the two events are separated by such long intervals. Moreover, the time between these later responses and the subsequent reinforcer is too long for the behavior to be maintained by that reinforcer. Again, independent experimental analysis does not confirm, in general, that behavior can be strengthened by putative reinforcers that occur after such intervals (but see Lattal & Gleeson, 1990). The interpretive challenge is complicated further by the fact that environmental stimuli within the test chamber remain constant throughout the fixed interval. These invariant stimuli cannot support changes in responding in any obvious fashion (cf. Donahoe, Palmer, & Burgos, 1997; Shull, 1995). The question remains, how can we interpret behavior that covaries with time in a constant environment?

Two complementary strategies for addressing temporal control are consistent with behavior analysis: (a) The behavior that occurs in the situation of interest may be further analyzed in that situation. (b) A subset of the events that occur in the original situation may be analyzed in other situations. In either case, the additional information is then used to interpret behavior in the original situation (see Donahoe, 1993). As an illustration of the first strategy, other characteristics of the operant can be measured, such as interresponse times (e.g., Anger, 1956; Williams, 1968). These measures reveal the discriminative effects of stimuli that are occasioned by immediately preceding responses (Morse, 1966). As another example, responses in addition to the operant may be concurrently measured during the interval, responses such as the adjunctive behavior of drinking (Roper, 1978) or general activity (Killeen, Hanson, & Osborne, 1978; see also Domjan, 1994; Silva, Timberlake, & Cevik, 1998). Temporally constrained adjunctive behavior is a potential source of immediately preceding stimuli that can exert discriminative control over the operant. Indeed, such stimuli have been exploited in the interpretation of temporal control by the behavioral theory of timing (BeT) (Killeen & Fetterman, 1988) and related formulations (Machado, 1997). Staddon and Higa ironically remark that interpretations based on the discriminative effects of adjunctive behavior address the temporal problem "in a less ad hoc (and more easily disprovable) way" (p. 215) than models such as SET. We prefer to characterize such behavioral models as incomplete rather than incorrect. Scientific interpretations can never be wholly wrong if they are based on valid experimental analysis. Indeed, this is the great advantage of scientific interpretations over inferred-process theories such as SET and, we believe, MTS as well.

The second strategy for understanding temporal control is to identify possible variables that may have acted in the original situation and then to subject them to experimental analysis in a different situation. The new observations may be made at any level—behavioral, neural, cellular, biochemical—without violating the basic strategy. At the behavioral level, it is often assumed that some stimulus, such as an "aftereffect" of the reinforcer, changes progressively during the fixed interval and that the operant is controlled by this aftereffect in proportion to its magnitude (cf. Capaldi, 1966). To determine whether events correlated with time since a reinforcer can serve a discriminative function, Ferster and Skinner (1957) introduced into a fixed-interval schedule an environmental stimulus whose characteristics varied with time since the preceding reinforcer. To wit,

the response key grew brighter as time passed. When the values of this stimulus (called an added clock) were later manipulated, responding changed in a manner that was consistent with the presumed role of the unobserved stimuli within a standard fixed-interval schedule. (This finding is a necessary but not a sufficient condition for the proposition that otherwise comparable, but unobserved, stimuli are discriminated in fixed-interval schedules without added clocks.) As a second behavioral example, a stimulus of either of two durations may be presented followed by a two-choice concurrent schedule in which the reinforced component is a function of the duration of the first stimulus. To wit, after a short stimulus duration on a middle key, pecking the left concurrent key is reinforced, whereas after a long duration, pecking the right key is reinforced. If the duration of the first stimulus is correlated with distinctive stimuli that change with time, then those stimuli should acquire control over concurrent responding, as, indeed, they do (e.g., Rilling, 1967). The formation of a concurrent temporal discrimination under these circumstances is consistent with the interpretation of fixed-interval performance as a product of control by unobserved stimuli that vary with the passage of time.

Neural processes are particularly likely to be invoked in discussions of temporal control. Because the external environment is constant, internal time-dependent stimuli are proposed to serve as the controlling stimuli. SET appeals to internal clocks, and even BeT relies on a pacemaker to coordinate transitions between behavioral states. ''For our characterization of the pacemaker we are immediately driven inside the organism. . . . The pacemaker itself is almost certainly a biological oscillator'' (Killeen & Fetterman, 1988, p. 289). The tactic of moving to a finer level of analysis when the explanatory variance at the behavioral level has been exhausted is consistent with the behavior-analytic approach:

> 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 [the physiologist] 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, commenting on the interval ''between the stimulating action of the environment and the response of the organism, . . . only brain science can fill these gaps. In doing so, it completes the account; it does not give a different account of the same thing'' (Skinner, 1953, p. 18; see also Michael, 1998, p. 160).

Staddon and Higa criticize the notions of timers and pacemakers (rightly so in our view) but then introduce in their place a series of hypothetical units with decreasing rates of habituation. According to MTS, temporal discriminations are dependent on comparisons between a ''long-term memory'' of the output of this cascade—stored at times when the operant was reinforced in the past—and the present value of the output (''short-term memory''). If the difference between the two memories is less than some threshold value (i.e., if the present time is similar to past times when the response was reinforced), then the operant is emitted. In evaluating MTS and other proposals, recall that an account meets the demands of scientific interpretation only if it draws upon events and processes that have been subjected to independent experimental analysis. What experimental evidence at the neural level substantiates the existence of clocks, pacemakers, and habituation units (or systems of such units) that operate in the behaviorally relevant range of seconds and minutes? To our knowledge, none exists. And others concur: ''There are no physiological data that support the existence of pacemakers and counters processing temporal information on the order of [even] hundreds of milliseconds'' (Buonomano & Merzenich, 1997, p. 139). The habituation units of MTS are no more the product of independent experimental analysis than are the pacemakers of SET. Habituation units are pacemakers under another name.[1] Thus, it is possible to

[1] The reference cited in support of biological pacemakers (Spitzer & Sejnowski, 1997) does not appear to be relevant to timing mechanisms of the type proposed in SET and MTS. The finding that units ''deeper'' in the brain are less likely to be activated as the time between stimulus presentations decreases (Williamson, described

agree with Staddon and Higa (p. 245) that "the jury is still out on whether [pacemaker] theories or the MTS theory have the stronger claim to biological plausibility" without endorsing either theory. Of course, circuits have been identified that respond differentially to very small differences in the time of arrival of biologically important stimuli within the same modality, such as those that mediate echolocation in bats (see Carr, 1993). Also, neural circuits have been identified that mediate circadian rhythms, although these are strongly entrained by environmental stimuli (see Rusak, Abe, Mason, Piggons, & Yang, 1993). Finally, circuits have been identified that temporally coordinate repetitive motor patterns, such as the tongue movements that accompany drinking (see Marder & Calabrese, 1996). Note, however, that all of these circuits mediate temporal relations between events that are biologically important over the evolutionary history of the species. None of these circuits mediates temporal control by *arbitrary* stimuli over the time ranges that are encountered in behavioral experiments. The remainder of this commentary describes an approach to temporal control that is more nearly consistent with the demands of experimental analysis.

Temporal Control as an Emergent Effect of Reinforcement

Biobehaviorally informed computer simulations indicate that behavior that covaries with time in a constant environment can emerge from the action of reinforcement on neural networks without invoking timers, pacemakers, or differentially habituating units (Burgos, 1995, 1997; Burgos & Donahoe, in press; see also Buonomano & Merzenich, 1995). The neural network consisted of input units (activated by simulated environmental events), sensory-association units, motor-association units, and output units (simulating the behavior of the network). Connection weights (simulating synaptic efficacies) to motor and motor-association units were modified by a nonspecifically projecting system (simulating the dopaminergic neuromodulatory system subserving reinforcement). Connection weights to sensory-association units were modified by a second nonspecifically projecting system (simulating projections from the hippocampus to sensory-association cortex) that, in turn, was also modulated by the reinforcing system. When input units were stimulated by constant environmental stimuli, activity propagated through the network and potentially activated the output unit simulating the operant. Connection weights were modified (updated) in successive time steps that simulated the continuous passage of real time. The weights increased whenever pre- and postsynaptic units were coactivated and a reinforcing signal was present. The weights decreased whenever the pre- and postsynaptic units were coactive but the reinforcing signal was absent (see Donahoe, Burgos, & Palmer, 1993, and Donahoe, 1997, for additional information about the learning algorithm).

Figure 1 shows a cumulative record of activations of the operant unit over a series of 10 time-step fixed intervals (Burgos & Donahoe, in press). In the training regimen, a reinforcing stimulus was presented when the operant unit had an activation greater than zero on the 10th time step. After the reinforcer, all activations were returned to their spontaneous levels and the next 10 time-step fixed interval began. With training, the operant unit was activated at very low levels immediately after the reinforcer, but was increasingly activated as the time to the next reinforcer approached. Thus, the simulated behavior bore an orderly relation to time within the 10 time-step fixed interval, but there was no timing mechanism. How did this

in Glanz, 1998) is consistent with MTS, but it can be accommodated without recourse to progressively more slowly habituating units—an assumption for which there is no direct experimental evidence. There is direct evidence, however, that all neurons have a refractory period after firing, and this fact is sufficient to account for the findings as presented. Let p be the probability that a neuron is in a state in which it can be activated by its inputs (i.e., not in a refractory period) and let there be n neurons in a series. In order for the nth neuron in the series to be activated by stimulating the first neuron, all neurons in the chain must become activated. Thus, the probability that the nth neuron is not activated (p_{na}) is $1 - p^n$ (ignoring complications such as spontaneous firing and parallel processing, neither of which affect the basic point). The probability p_{na} decreases as a function of n and as a function of p, which in turn decreases as the rate of stimulation increases. (Given any fixed time interval, the proportion of that interval in which an element is in a refractory period increases as the rate of stimulation increases.) In short, the finding is equally consistent with a series of units all of which have the *same* habituation rate.

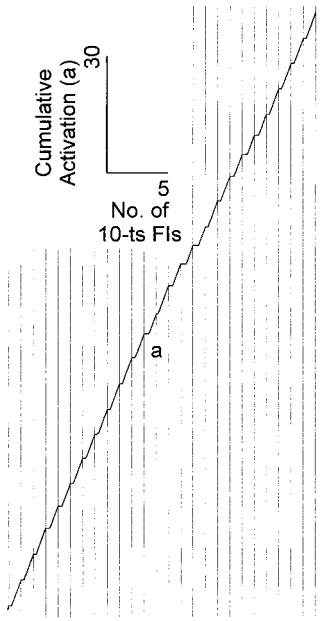

Fig. 1. Cumulative record of the activations ($0 < a < 1$) of an operant unit trained with a 10 time-step (*ts*) fixed-interval (FI) schedule of reinforcement. Vertical lines indicate boundaries between successive fixed intervals. Performance is shown for FI 276 through 300 of training. (Adapted from Burgos & Donahoe, in press.)

come about? The process whereby temporal control developed may be sketched as follows: Early in training, the reinforcer strengthened connection weights that caused the operant to be activated on the next time step in which the stimulus occurred (i.e., the onset of the next fixed interval). However, early activations of the operant unit were not reinforced and, therefore, the set of connection weights that mediated early responses weakened. With continued training, the reinforcing signal came to select a set of connection weights that permitted the constant environment to activate the operant unit only later in the fixed interval. An examination of the cumulative record of activations in Figure 1 provides some sense of this process. At point *a,* the postreinforcement pause was unusually brief. These nonreinforced responses caused the set of connection weights that supported early responding to weaken. As a result, postreinforcement pauses were unusually long during the immediately following fixed intervals. The temporal pattern of responding reveals a dynamic process of acquisition and extinction across successive fixed intervals.

We wish to make it clear that we do not regard the foregoing as a full simulation of behavior reinforced on a fixed-interval schedule. Many other variables would have to be taken into account in a complete simulation. However, the simulation does demonstrate that temporal control can develop in a constant environment as an emergent effect of reinforcement without recourse to pacemakers and the like. Nothing in the network resembles a timer, and the parameter values for all the units were identical to those used in previous simulations (e.g., Donahoe et al., 1993). Instead of pacemakers or differentially habituating units, reinforcement simply selected a set of connection weights such that only prolonged exposure to the constant environment was sufficient to strongly activate the operant. After exposure to the contingencies of reinforcement, the network acted *as if* there were a timer, but in fact none existed. Moreover, the characteristics of the network were constrained by information obtained from independent experimental analyses carried out at the neural level. These included the general architecture of the network, the effects of the hippocampal and the dopaminergic systems on synaptic efficacies in sensory association and motor cortex, the effect of dopamine on long-lasting long-term potentiation, temporal summation of presynaptic inputs, and competition between presynaptic inputs for control of postsynaptic receptors among others (see Donahoe, 1997). A more adequate simulation would be more fully informed by such information. We are only at the beginning of biobehaviorally constrained interpretations of complex contingencies of reinforcement.

The fruitfulness of a biobehavioral approach is further illustrated by other results that demonstrate a relation between the time of maximal responsiveness and the variance about that time. Simulations were conducted in which networks were trained in Pavlovian procedures to respond to interstimulus intervals (ISIs) of different durations: four, eight, or 16 time steps (Burgos, 1995, 1997). These simulations used a hybrid genetic-developmental-learning algorithm. The genetic algorithm simulated evolution through natural selection to determine the network architec-

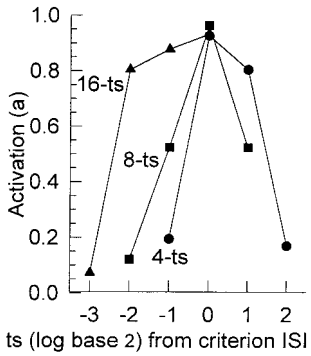

Fig. 2. Mean activations of the output units to probe stimuli of various durations. Different networks were given 200 trials of Pavlovian training with the criterion interstimulus intervals (ISIs) of either four, eight, or 16 time steps (*ts*). The ISI of the probe stimulus that was the same as the training ISI is designated as zero on the *x* axis. The differences between the ISIs of the probe stimuli and the criterion ISI are expressed in logarithmic units (base 2). (Adapted from Burgos, 1995, 1997.)

ture. The developmental algorithm simulated the epigenetic processes that affect the initial connectivity of the networks produced by the genetic algorithm. The learning algorithm modified the connection weights of the networks produced by the combined actions of the genetic and developmental algorithms. The final networks were the result of as many as 100 generations, with each generation consisting of successive implementations of the three algorithms.

Figure 2 shows the activation levels of conditioned response units during probe stimuli presented at the conclusion of the evolutionary process. The intervals for which selection occurred during evolution (i.e., the criterion ISIs) were four, eight, or 16 time steps. The value of zero on the *x* axis indicates that the ISI of the probe stimulus was the same as the value of the criterion ISI during evolution. The durations of the other probe stimuli were either shorter (negative values) or longer (positive values) than the criterion ISI. The values shown are the activation levels at the time steps immediately prior to the end of the probe stimulus. Two aspects of the results are pertinent here. First, regardless of the criterion ISI, architectures and connection weights were selected that permitted the networks to respond maximally to the selected ISI. Thus, the contingencies of selection were sufficient without recourse to either pacemakers or habituation units. Second, the variance in responding increased as the criterion ISI increased. Increased variance in temporal control occurred automatically as an emergent product of selection. (Interpretations of responding to ISI values greater than those encountered during the selection process are problematic; cf. Staddon & Higa, p. 226.) In short, Weber-law-like temporal control appeared without the need to appeal to pacemakers, hypothetical units with differential habituation rates, or transformations applied to linear processes.

Conclusion

The core problem of temporal control is to explain how behavior can vary with time in a constant environment. The problem arises because experimental analysis has shown that operants are almost invariably discriminated (i.e., under the control of some antecedent) and that the controlling stimulus must occur almost immediately before the response in question. What antecedent stimulus is responsible for temporal control? SET hypothesizes a pacemaker and associated cognitive machinery to supply the missing events. MTS hypothesizes a cascade of differentially habituating units together with various memory and comparison processes. BeT looks toward stimuli produced by temporally constrained adjunctive behavior, with a pacemaker to coordinate transitions between behaviorally defined states. What we suggest is that timing can occur without a timer of any sort. Reinforcers simply select a set of synaptic efficacies such that, when the neural system is acted upon by a constant environment, the operant is emitted with changing probability. Before the action of reinforcement, no timer (of whatever form) existed. After reinforcement, no timer exists. The system simply functions *as if* there were a timer. Natural selection has given organisms a network of richly interconnected neurons with the capability for their synapses to be modified by reinforcement, and this appears to be sufficient for the acquisition of temporal control.

REFERENCES

Anger, D. (1956). The dependence of interresponse times upon the relative reinforcement of different in-

terresponse times. *Journal of Experimental Psychology, 52,* 145–161.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 89–195). San Francisco: Freeman.

Buonomano, D. V., & Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science, 267,* 1026–1028.

Buonomano, D. V., & Merzenich, M. M. (1997). Temporal information processing: A computational role for paired-pulse facilitation and slow inhibition. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations* (pp. 129–139). Amsterdam: Elsevier Science Press.

Burgos, J. E. (1995). *Evolving artificial neural networks in Pavlovian environments.* Unpublished doctoral dissertation, University of Massachusetts at Amherst.

Burgos, J. E. (1997). Evolving artifical neural networks in Pavlovian environments. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations* (pp. 58–79). Amsterdam: Elsevier Science Press.

Burgos, J. E., & Donahoe, J. W. (in press). A formal interpretation of the structure-function distinction. In D. Blackman & J. C. Leslie (Eds.), *Current issues in experimental and applied behavior analysis.* Reno, NV: Context Press.

Capaldi, E. J. (1966). Partial reinforcement: A hypothesis of sequential effects. *Psychological Review, 73,* 459–477.

Carr, L. E. (1993). Processing of temporal information in the brain. *Annual Review of Neuroscience, 16,* 223–243.

Domjan, M. (1994). Formulation of a behavior system for sexual conditioning. *Psychonomic Bulletin & Review, 1,* 421–428.

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. (1997). Selection networks: Simulation of plasticity through reinforcement learning. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations* (pp. 336–357). Amsterdam: Elsevier Science Press.

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 & Bacon.

Donahoe, J. W., Palmer, D. C., & Burgos, J. E. (1997). The S-R issue: Its status in behavior analysis and in Donahoe and Palmer's *Learning and Complex Behavior. Journal of the Experimental Analysis of Behavior, 67,* 193–211.

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement.* Englewood Cliffs, NJ: Prentice Hall.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279–325.

Gibbon, J., & Church, R. M. (1990). Representation of time. *Cognition, 37,* 23–54.

Glanz, J. (1998). Magnetic brain imaging traces a stairway to memory. *Science, 280,* 37.

Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review, 95,* 274–293.

Killeen, P. R., Hanson, S. J., & Osborne, S. R. (1978). Arousal: Its genesis and manifestation as response rate. *Psychological Review, 85,* 571–581.

Lattal, K. A., & Gleeson, S. (1990). Response acquisition with delayed reinforcement. *Journal of the Experimental Analysis of Behavior, 16,* 27–39.

Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review, 104,* 241–265.

Marder, E., & Calabrese, R. L. (1996). Principles of rhythmic motor pattern generation. *Physiological Reviews, 76,* 687–717.

Michael, J. (1998). The current status and future directions of the analysis of verbal behavior: Comments on the comments. *The Analysis of Verbal Behavior, 15,* 157–161.

Morse, W. H. (1966). Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52–108). Englewood Cliffs, NJ: Prentice Hall.

Palmer, D. C., & Donahoe, J. W. (1992). Selectionism and essentialism in behaviorism and cognitive science. *American Psychologist, 47,* 1344–1358.

Papini, M. R., & Hollingsworth, P. R. (1998). Role of nonreinforcement in the fixed-interval performance of pigeons. *Psychonomic Bulletin & Review, 5,* 84–90.

Rilling, M. E. (1967). Number of responses as a stimulus in fixed-interval and fixed-ratio schedules. *Journal of Comparative and Physiological Psychology, 63,* 60–65.

Roper, T. J. (1978). Diversity and substitutability of adjunctive activities under fixed-interval schedules of food reinforcement. *Journal of the Experimental Analysis of Behavior, 30,* 83–96.

Rusak, B., Abe, H., Mason, R., Piggons, H. D., & Yang, S-W. (1993). Neurophysiological analysis of circadian rhythm entrainment. *Journal of Biological Rhythms, 8,* 39–45.

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.

Silva, F. J., Timberlake, W., & Cevik, M. O. (1998). A behavior systems approach to the expression of backward associations. *Learning and Motivation, 29,* 1–22.

Skinner, B. F. (1938). *The behavior of organisms.* New York: Appleton-Century-Crofts.

Skinner, B. F. (1953). *Science and human behavior.* New York: Macmillan.

Skinner, B. F. (1974). *About behaviorism.* New York: Random House.

Spitzer, N. C., & Sejnowski, T. J. (1997). Biological information processing: Bits of progress. *Science, 277,* 1060–1061.

Staddon, J. E. R., & Higa, J. J. (1996). Multiple time scales in simple habituation. *Psychological Review, 103,* 720–733.

Williams, D. R. (1968). The structure of response rate. *Journal of the Experimental Analysis of Behavior, 11,* 251–258.