

OPERANT CONDITIONING

J. E. R. Staddon and D. T. Cerutti

*Department of Psychological and Brain Sciences, Duke University,
Durham, North Carolina 27708-0086; e-mail: staddon@psych.duke.edu,
cerutti@psych.duke.edu*

Key Words interval timing, choice, concurrent schedules, matching law, self-control

■ **Abstract** Operant behavior is behavior “controlled” by its consequences. In practice, operant conditioning is the study of reversible behavior maintained by reinforcement schedules. We review empirical studies and theoretical approaches to two large classes of operant behavior: interval timing and choice. We discuss cognitive versus behavioral approaches to timing, the “gap” experiment and its implications, proportional timing and Weber’s law, temporal dynamics and linear waiting, and the problem of simple chain-interval schedules. We review the long history of research on operant choice: the matching law, its extensions and problems, concurrent chain schedules, and self-control. We point out how linear waiting may be involved in timing, choice, and reinforcement schedules generally. There are prospects for a unified approach to all these areas.

CONTENTS

INTRODUCTION	116
INTERVAL TIMING	118
WEBER’S LAW, PROPORTIONAL TIMING AND TIMESCALE INVARIANCE	119
Cognitive and Behavioral Approaches to Timing	121
The Gap Experiment	123
Temporal Dynamics: Linear Waiting	125
Chain Schedules	126
CHOICE	133
Concurrent Schedules	134
Concurrent-Chain Schedules	136
Self-Control	137
CONCLUSION	139

INTRODUCTION

The term *operant conditioning*¹ was coined by B. F. Skinner in 1937 in the context of reflex physiology, to differentiate what he was interested in—behavior that affects the environment—from the reflex-related subject matter of the Pavlovians. The term was novel, but its referent was not entirely new. *Operant behavior*, though defined by Skinner as behavior “controlled by its consequences” is in practice little different from what had previously been termed “instrumental learning” and what most people would call habit. Any well-trained “operant” is in effect a habit. What was truly new was Skinner’s method of automated training with intermittent reinforcement and the subject matter of *reinforcement schedules* to which it led. Skinner and his colleagues and students discovered in the ensuing decades a completely unsuspected range of powerful and orderly schedule effects that provided new tools for understanding learning processes and new phenomena to challenge theory.

A reinforcement schedule is any procedure that delivers a *reinforcer* to an organism according to some well-defined rule. The usual reinforcer is food for a hungry rat or pigeon; the usual schedule is one that delivers the reinforcer for a switch closure caused by a peck or lever press. Reinforcement schedules have also been used with human subjects, and the results are broadly similar to the results with animals. However, for ethical and practical reasons, relatively weak reinforcers must be used—and the range of behavioral strategies people can adopt is of course greater than in the case of animals. This review is restricted to work with animals.

Two types of reinforcement schedule have excited the most interest. Most popular are time-based schedules such as fixed and variable interval, in which the reinforcer is delivered after a fixed or variable time period after a *time marker* (usually the preceding reinforcer). *Ratio* schedules require a fixed or variable number of responses before a reinforcer is delivered.

Trial-by-trial versions of all these *free-operant* procedures exist. For example, a version of the fixed-interval schedule specifically adapted to the study of interval timing is the peak-interval procedure, which adds to the fixed interval an intertrial interval (ITI) preceding each trial and a percentage of extra-long “empty” trials in which no food is given.

For theoretical reasons, Skinner believed that operant behavior ought to involve a response that can easily be repeated, such as pressing a lever, for rats, or pecking

¹The first and only previous *Annual Review* contribution on this topic was as part of a 1965 article, “Learning, Operant Conditioning and Verbal Learning” by Blough & Millward. Since then there have been (by our estimate) seven articles on learning or learning theory in animals, six on the neurobiology of learning, and three on human learning and memory, but this is the first full *Annual Review* article on operant conditioning. We therefore include rather more old citations than is customary (for more on the history and philosophy of Skinnerian behaviorism, both pro and con, see Baum 1994, Rachlin 1991, Sidman 1960, Staddon 2001b, and Zuriff 1985).

an illuminated disk (key) for pigeons. The rate of such behavior was thought to be important as a measure of response strength (Skinner 1938, 1966, 1986; Killeen & Hall 2001). The current status of this assumption is one of the topics of this review. True or not, the emphasis on response rate has resulted in a dearth of experimental work by operant conditioners on nonrecurrent behavior such as movement in space.

Operant conditioning differs from other kinds of learning research in one important respect. The focus has been almost exclusively on what is called *reversible* behavior, that is, behavior in which the steady-state pattern under a given schedule is stable, meaning that in a sequence of conditions, XAXBXC. . . , where each condition is maintained for enough days that the pattern of behavior is locally stable, behavior under schedule X shows a pattern after one or two repetitions of X that is always the same. For example, the first time an animal is exposed to a fixed-interval schedule, after several daily sessions most animals show a “scalloped” pattern of responding (call it pattern A): a pause after each food delivery—also called *wait time* or *latency*—followed by responding at an accelerated rate until the next food delivery. However, some animals show negligible wait time and a steady rate (pattern B). If all are now trained on some other procedure—a variable-interval schedule, for example—and then after several sessions are returned to the fixed-interval schedule, almost all the animals will revert to pattern A. Thus, pattern A is the stable pattern. Pattern B, which may persist under unchanging conditions but does not recur after one or more intervening conditions, is sometimes termed *metastable* (Staddon 1965). The vast majority of published studies in operant conditioning are on behavior that is stable in this sense.

Although the theoretical issue is not a difficult one, there has been some confusion about what the idea of stability (reversibility) in behavior means. It should be obvious that the animal that shows pattern A after the second exposure to procedure X is not the same animal as when it showed pattern A on the first exposure. Its experimental history is different after the second exposure than after the first. If the animal has any kind of memory, therefore, its internal state² following the second exposure is likely to be different than after the first exposure, even though the observed behavior is the same. The behavior is reversible; the organism’s internal state in general is not. The problems involved in studying nonreversible phenomena in individual organisms have been spelled out elsewhere (e.g., Staddon 2001a, Ch. 1); this review is mainly concerned with the reversible aspects of behavior.

Once the microscope was invented, microorganisms became a new field of investigation. Once automated operant conditioning was invented, reinforcement schedules became an independent subject of inquiry. In addition to being of great interest in their own right, schedules have also been used to study topics defined in more abstract ways such as timing and choice. These two areas constitute the majority of experimental papers in operant conditioning with animal subjects during

²By “internal” we mean not “physiological” but “hidden.” The idea is simply that the organism’s future behavior depends on variables not all of which are revealed in its current behavior (cf. Staddon 2001b, Ch. 7).

the past two decades. Great progress has been made in understanding free-operant choice behavior and interval timing. Yet several theories of choice still compete for consensus, and much the same is true of interval timing. In this review we attempt to summarize the current state of knowledge in these two areas, to suggest how common principles may apply in both, and to show how these principles may also apply to reinforcement schedule behavior considered as a topic in its own right.

INTERVAL TIMING

Interval timing is defined in several ways. The simplest is to define it as covariation between a dependent measure such as wait time and an independent measure such as interreinforcement interval (on fixed interval) or trial time-to-reinforcement (on the peak procedure). When interreinforcement interval is doubled, then after a learning period wait time also approximately doubles (*proportional timing*). This is an example of what is sometimes called a *time production* procedure: The organism produces an approximation to the to-be-timed interval. There are also explicit *time discrimination* procedures in which on each trial the subject is exposed to a stimulus and is then required to respond differentially depending on its absolute (Church & Deluty 1977, Stubbs 1968) or even relative (Fetterman et al. 1989) duration. For example, in *temporal bisection*, the subject (e.g., a rat) experiences either a 10-s or a 2-s stimulus, *L* or *S*. After the stimulus goes off, the subject is confronted with two choices. If the stimulus was *L*, a press on the left lever yields food; if *S*, a right press gives food; errors produce a brief time-out. Once the animal has learned, stimuli of intermediate duration are presented in lieu of *S* and *L* on test trials. The question is, how will the subject distribute its responses? In particular, at what intermediate duration will it be indifferent between the two choices? [Answer: typically in the vicinity of the geometric mean, i.e., $\sqrt{(L.S)}$ — 4.47 for 2 and 10.]

Wait time is a latency; hence (it might be objected) it may vary on time-production procedures like fixed interval because of factors other than timing—such as degree of hunger (food deprivation). Using a time-discrimination procedure avoids this problem. It can also be mitigated by using the peak procedure and looking at performance during “empty” trials. “Filled” trials terminate with food reinforcement after (say) *T* s. “Empty” trials, typically $3T$ s long, contain no food and end with the onset of the ITI. During empty trials the animal therefore learns to wait, then respond, then stop (more or less) until the end of the trial (Catania 1970). The mean of the distribution of response rates averaged over empty trials (*peak time*) is then perhaps a better measure of timing than wait time because motivational variables are assumed to affect only the height and spread of the response-rate distribution, not its mean. This assumption is only partially true (Grace & Nevin 2000, MacEwen & Killeen 1991, Plowright et al. 2000).

There is still some debate about the actual pattern of behavior on the peak procedure in each individual trial. Is it just wait, respond at a constant rate, then wait again? Or is there some residual responding after the “stop” [yes, usually (e.g.,

Church et al. 1991)]? Is the response rate between start and stop really constant or are there two or more identifiable rates (Cheng & Westwood 1993, Meck et al. 1984)? Nevertheless, the method is still widely used, particularly by researchers in the cognitive/psychophysical tradition. The idea behind this approach is that interval timing is akin to sensory processes such as the perception of sound intensity (loudness) or luminance (brightness). As there is an ear for hearing and an eye for seeing, so (it is assumed) there must be a (real, physiological) clock for timing. Treisman (1963) proposed the idea of an internal pacemaker-driven clock in the context of human psychophysics. Gibbon (1977) further developed the approach and applied it to animal interval-timing experiments.

WEBER'S LAW, PROPORTIONAL TIMING AND TIMESCALE INVARIANCE

The major similarity between acknowledged sensory processes, such as brightness perception, and interval timing is *Weber's law*. Peak time on the peak procedure is not only proportional to time-to-food (T), its coefficient of variation (standard deviation divided by mean) is approximately constant, a result similar to Weber's law obeyed by most sensory dimensions. This property has been called *scalar timing* (Gibbon 1977). Most recently, Gallistel & Gibbon (2000) have proposed a grand principle of *timescale invariance*, the idea that the frequency distribution of any given temporal measure (the idea is assumed to apply generally, though in fact most experimental tests have used peak time) scales with the to-be-timed-interval. Thus, given the normalized peak-time distribution for $T = 60$ s, say; if the x -axis is divided by 2, it will match the distribution for $T = 30$ s. In other words, the frequency distribution for the temporal dependent variable, normalized on both axes, is asserted to be invariant.

Timescale invariance is in effect a combination of Weber's law and proportional timing. Like those principles, it is only approximately true. There are three kinds of evidence that limit its generality. The simplest is the steady-state pattern of responding (key-pecking or lever-pressing) observed on fixed-interval reinforcement schedules. This pattern should be the same at all fixed-interval values, but it is not. Gallistel & Gibbon wrote, "When responding on such a schedule, animals pause after each reinforcement and then resume responding after some interval has elapsed. It was generally supposed that the animals' rate of responding accelerated throughout the remainder of the interval leading up to reinforcement. In fact, however, conditioned responding in this paradigm . . . is a two-state variable (slow, sporadic pecking vs. rapid, steady pecking), with one transition per interreinforcement interval (Schneider 1969)" (p. 293).

This conclusion over-generalizes Schneider's result. Reacting to reports of "break-and-run" fixed-interval performance under some conditions, Schneider sought to characterize this feature more objectively than the simple inspection of cumulative records. He found a way to identify the point of maximum acceleration in the fixed-interval "scallop" by using an iterative technique analogous to

attaching an elastic band to the beginning of an interval and the end point of the cumulative record, then pushing a pin, representing the break point, against the middle of the band until the two resulting straight-line segments best fit the cumulative record (there are other ways to achieve the same result that do not fix the end points of the two line-segments). The postreinforcement time (x -coordinate) of the pin then gives the break point for that interval. Schneider showed that the break point is an orderly dependent measure: Break point is roughly 0.67 of interval duration, with standard deviation proportional to the mean (the Weber-law or scalar property).

This finding is by no means the same as the idea that the fixed-interval scallop is “a two-state variable” (Hanson & Killeen 1981). Schneider showed that a two-state model is an adequate approximation; he did not show that it is the best or truest approximation. A three- or four-line approximation (i.e., two or more pins) might well have fit significantly better than the two-line version. To show that the process is two-state, Schneider would have had to show that adding additional segments produced negligibly better fit to the data.

The frequent assertion that the fixed-interval scallop is always an artifact of averaging flies in the face of raw cumulative-record data—the many nonaveraged individual fixed-interval cumulative records in Ferster & Skinner (1957, e.g., pp. 159, 160, 162), which show clear curvature, particularly at longer fixed-interval values ($> \sim 2$ min). The issue for timescale invariance, therefore, is whether the shape, or relative frequency of different-shaped records, is the same at different absolute intervals.

The evidence is that there is more, and more frequent, curvature at longer intervals. Schneider’s data show this effect. In Schneider’s Figure 3, for example, the time to shift from low to high rate is clearly longer at longer intervals than shorter ones. On fixed-interval schedules, apparently, absolute duration does affect the pattern of responding. (A possible reason for this dependence of the scallop on fixed-interval value is described in Staddon 2001a, p. 317. The basic idea is that greater curvature at longer fixed-interval values follows from two things: a linear increase in response probability across the interval, combined with a nonlinear, negatively accelerated, relation between overall response rate and reinforcement rate.) If there is a reliable difference in the shape, or distribution of shapes, of cumulative records at long and short fixed-interval values, the timescale-invariance principle is violated.

A second dataset that does not agree with timescale invariance is an extensive set of studies on the peak procedure by Zeiler & Powell (1994; see also Hanson & Killeen 1981), who looked explicitly at the effect of interval duration on various measures of interval timing. They conclude, “Quantitative properties of temporal control depended on whether the aspect of behavior considered was initial pause duration, the point of maximum acceleration in responding [break point], the point of maximum deceleration, the point at which responding stopped, or several different statistical derivations of a point of maximum responding Existing theory does not explain why Weber’s law [the scalar property] so rarely fit the results”

(p. 1; see also Lowe et al. 1979, Wearden 1985 for other exceptions to proportionality between temporal measures of behavior and interval duration). Like Schneider (1969) and Hanson & Killeen (1981), Zeiler & Powell found that the breakpoint measure was proportional to interval duration, with scalar variance (constant coefficient of variation), and thus consistent with timescale invariance, but no other measure fit the rule.

Moreover, the fit of the breakpoint measure is problematic because it is not a direct measure of behavior but is itself the result of a statistical fitting procedure. It is possible, therefore, that the fit of breakpoint to timescale invariance owes as much to the statistical method used to arrive at it as to the intrinsic properties of temporal control. Even if this caveat turns out to be false, the fact that every other measure studied by Zeiler & Powell failed to conform to timescale invariance surely rules it out as a general principle of interval timing.

The third and most direct test of the timescale invariance idea is an extensive series of time-discrimination experiments carried out by Dreyfus et al. (1988) and Stubbs et al. (1994). The usual procedure in these experiments was for pigeons to peck a center response key to produce a red light of one duration that is followed immediately by a green light of another duration. When the green center-key light goes off, two yellow side-keys light up. The animals are reinforced with food for pecking the left side-key if the red light was longer, the right side-key if the green light was longer.

The experimental question is, how does discrimination accuracy depend on relative and absolute duration of the two stimuli? Timescale invariance predicts that accuracy depends only on the ratio of red and green durations: For example, accuracy should be the same following the sequence red:10, green:20 as the sequence red:30, green:60, but it is not. Pigeons are better able to discriminate between the two short durations than the two long ones, even though their ratio is the same. Dreyfus et al. and Stubbs et al. present a plethora of quantitative data of the same sort, all showing that time discrimination depends on absolute as well as relative duration.

Timescale invariance is empirically indistinguishable from Weber's law as it applies to time, combined with the idea of proportional timing: The mean of a temporal dependent variable is proportional to the temporal independent variable. But Weber's law and proportional timing are dissociable—it is possible to have proportional timing without conforming to Weber's law and vice versa (cf. Hanson & Killeen 1981, Zeiler & Powell 1994), and in any case both are only approximately true. Timescale invariance therefore does not qualify as a principle in its own right.

Cognitive and Behavioral Approaches to Timing

The cognitive approach to timing dates from the late 1970s. It emphasizes the psychophysical properties of the timing process and the use of temporal dependent variables as measures of (for example) drug effects and the effects of physiological interventions. It de-emphasizes proximal environmental causes. Yet when

timing (then called temporal control; see Zeiler 1977 for an early review) was first discovered by operant conditioners (Pavlov had studied essentially the same phenomenon—*delay conditioning*—many years earlier), the focus was on the *time marker*, the stimulus that triggered the temporally correlated behavior. (That is one virtue of the term *control*: It emphasizes the fact that interval timing behavior is usually not free-running. It must be cued by some aspect of the environment.) On so-called spaced-responding schedules, for example, the response is the time marker: The subject must learn to space its responses more than T s apart to get food. On fixed-interval schedules the time marker is reinforcer delivery; on the peak procedure it is the stimulus events associated with trial onset. This dependence on a time marker is especially obvious on time-production procedures, but on time-discrimination procedures the subject's choice behavior must also be under the control of stimuli associated with the onset and offset of the sample duration.

Not all stimuli are equally effective as time markers. For example, an early study by Staddon & Innis (1966a; see also 1969) showed that if, on alternate fixed intervals, 50% of reinforcers (F) are omitted and replaced by a neutral stimulus (N) of the same duration, wait time following N is much shorter than after F (the *reinforcement-omission effect*). Moreover, this difference persists indefinitely. Despite the fact that F and N have the same temporal relationship to the reinforcer, F is much more effective as a time marker than N. No exactly comparable experiment has been done using the peak procedure, partly because the time marker there involves ITI offset/trial onset rather than the reinforcer delivery, so that there is no simple manipulation equivalent to reinforcement omission.

These effects do not depend on the type of behavior controlled by the time marker. On fixed-interval schedules the time marker is in effect inhibitory: Responding is suppressed during the wait time and then occurs at an accelerating rate. Other experiments (Staddon 1970, 1972), however, showed that given the appropriate schedule, the time marker can control a burst of responding (rather than a wait) of a duration proportional to the schedule parameters (*temporal go-no-go* schedules) and later experiments have shown that the place of responding can be controlled by time since trial onset in the so-called tri-peak procedure (Matell & Meck 1999).

A theoretical review (Staddon 1974) concluded, "Temporal control by a given time marker depends on the properties of recall and attention, that is, on the same variables that affect attention to compound stimuli and recall in memory experiments such as delayed matching-to-sample." By far the most important variable seems to be "the *value* of the time-marker stimulus—Stimuli of high value . . . are more salient . . ." (p. 389), although the full range of properties that determine time-marker effectiveness is yet to be explored.

Reinforcement omission experiments are *transfer tests*, that is, tests to identify the effective stimulus. They pinpoint the stimulus property controlling interval timing—the effective time marker—by selectively eliminating candidate properties. For example, in a definitive experiment, Kello (1972) showed that on fixed interval the wait time is longest following standard reinforcer delivery (food hopper

activated with food, hopper light on, house light off, etc.). Omission of any of those elements caused the wait time to decrease, a result consistent with the hypothesis that reinforcer delivery acquires inhibitory temporal control over the wait time. The only thing that makes this situation different from the usual generalization experiment is that the effects of reinforcement omission are relatively permanent. In the usual generalization experiment, delivery of the reinforcer according to the same schedule in the presence of both the training stimulus and the test stimuli would soon lead all to be responded to in the same way. Not so with temporal control: As we just saw, even though N and F events have the same temporal relationship to the next food delivery, animals never learn to respond similarly after both. The only exception is when the fixed-interval is relatively short, on the order of 20 s or less (Starr & Staddon 1974). Under these conditions pigeons are able to use a brief neutral stimulus as a time marker on fixed interval.

The Gap Experiment

The closest equivalent to fixed-interval reinforcement–omission using the peak procedure is the so-called gap experiment (Roberts 1981). In the standard gap paradigm the sequence of stimuli in a training trial (no gap stimulus) consists of three successive stimuli: the intertrial interval stimulus (ITI), the fixed-duration trial stimulus (S), and food reinforcement (F), which ends each training trial. The sequence is thus ITI, S, F, ITI. Training trials are typically interspersed with empty probe trials that last longer than reinforced trials but end with an ITI only and no reinforcement. The stimulus sequence on such trials is ITI, S, ITI, but the S is two or three times longer than on training trials. After performance has stabilized, gap trials are introduced into some or all of the probe trials. On gap trials the ITI stimulus reappears for a while in the middle of the trial stimulus. The sequence on gap trials is therefore ITI, S, ITI, S, ITI. Gap trials do not end in reinforcement.

What is the effective time marker (i.e., the stimulus that exerts temporal control) in such an experiment? ITI offset/trial onset is the best temporal predictor of reinforcement: Its time to food is shorter and less variable than any other experimental event. Most but not all ITIs follow reinforcement, and the ITI itself is often variable in duration and relatively long. So reinforcer delivery is a poor temporal predictor. The time marker therefore has something to do with the transition between ITI and trial onset, between ITI and S. Gap trials also involve presentation of the ITI stimulus, albeit with a different duration and within-trial location than the usual ITI, but the similarities to a regular trial are obvious. The gap experiment is therefore a sort of generalization (of temporal control) experiment. Buhusi & Meck (2000) presented gap stimuli more or less similar to the ITI stimulus during probe trials and found results resembling generalization decrement, in agreement with this analysis.

However, the gap procedure was not originally thought of as a generalization test, nor is it particularly well designed for that purpose. The gap procedure arose

directly from the cognitive idea that interval timing behavior is driven by an internal clock (Church 1978). From this point of view it is perfectly natural to inquire about the conditions under which the clock can be started or stopped. If the to-be-timed interval is interrupted—a gap—will the clock restart when the trial stimulus returns (reset)? Will it continue running during the gap and afterwards? Or will it stop and then restart (stop)?

“Reset” corresponds to the maximum rightward shift (from trial onset) of the response-rate peak from its usual position t s after trial onset to $t + G_E$, where G_E is the offset time (end) of the gap stimulus. Conversely, no effect (clock keeps running) leaves the peak unchanged at t , and “stop and restart” is an intermediate result, a peak shift to $G_E - G_B + t$, where G_B is the time of onset (beginning) of the gap stimulus.

Both gap duration and placement within a trial have been varied. The results that have been obtained so far are rather complex (cf. Buhusi & Meck 2000, Cabeza de Vaca et al. 1994, Matell & Meck 1999). In general, the longer the gap and the later it appears in the trial, the greater the rightward peak shift. All these effects can be interpreted in clock terms, but the clock view provides no real explanation for them, because it does not specify which one will occur under a given set of conditions. The results of gap experiments can be understood in a qualitative way in terms of the similarity of the gap presentation to events associated with trial onset; the more similar, the closer the effect will be to reset, i.e., the onset of a new trial. Another resemblance between gap results and the results of reinforcement-omission experiments is that the effects of the gap are also permanent: Behavior on later trials usually does not differ from behavior on the first few (Roberts 1981). These effects have been successfully simulated quantitatively by a neural network timing model (Hopson 1999, 2002) that includes the assumption that the effects of time-marker presentation decay with time (Cabeza de Vaca et al. 1994).

The original temporal control studies were strictly empirical but tacitly accepted something like the psychophysical view of timing. Time was assumed to be a sensory modality like any other, so the experimental task was simply to explore the different kinds of effect, excitatory, inhibitory, discriminatory, that could come under temporal control. The psychophysical view was formalized by Gibbon (1977) in the context of animal studies, and this led to a static information-processing model, *scalar expectancy theory* (SET: Gibbon & Church 1984, Meck 1983, Roberts 1983), which comprised a pacemaker-driven clock, working and reference memories, a comparator, and various thresholds. A later dynamic version added memory for individual trials (see Gallistel 1990 for a review). This approach led to a long series of experimental studies exploring the clocklike properties of interval timing (see Gallistel & Gibbon 2000, Staddon & Higa 1999 for reviews), but none of these studies attempted to test the assumptions of the SET approach in a direct way.

SET was for many years the dominant theoretical approach to interval timing. In recent years, however, its limitations, of parsimony and predictive range, have

become apparent and there are now a number of competitors such as the behavioral theory of timing (Killeen & Fetterman 1988, MacEwen & Killeen 1991, Machado 1997), spectral timing theory (Grossberg & Schmajuk 1989), neural network models (Church & Broadbent 1990, Hopson 1999, Dragoi et al. 2002), and the habituation-based multiple time scale theory (MTS: Staddon & Higa 1999, Staddon et al. 2002). There is as yet no consensus on the best theory.

Temporal Dynamics: Linear Waiting

A separate series of experiments in the temporal-control tradition, beginning in the late 1980s, studied the real-time dynamics of interval timing (e.g., Higa et al. 1991, Lejeune et al. 1997, Wynne & Staddon 1988; see Staddon 2001a for a review). These experiments have led to a simple empirical principle that may have wide application. Most of these experiments used the simplest possible timing schedule, a response-initiated delay (RID) schedule³. In this schedule the animal (e.g., a pigeon) can respond at any time, t , after food. The response changes the key color and food is delivered after a further T s. Time t is under the control of the animal; time T is determined by the experimenter. These experiments have shown that wait time on these and similar schedules (such as fixed interval) is strongly determined by the duration of the previous interfood interval (IFI). For example, wait time will track a cyclic sequence of IFIs, intercalated at a random point in a sequence of fixed ($t + T = \text{constant}$) intervals, with a lag of one interval; a single short IFI is followed by a short wait time in the next interval (the effect of a single long interval is smaller), and so on (see Staddon et al. 2002 for a review and other examples of temporal tracking). To a first approximation, these results are consistent with a linear relation between wait time in IFI $N + 1$ and the duration of IFI N :

$$t(N + 1) = a[T(N) + t(N)] + b = aI(N) + b, \quad (1)$$

where I is the IFI, a is a constant less than one, and b is usually negligible. This relation has been termed *linear waiting* (Wynne & Staddon 1988). The principle is an approximation: an expanded model, incorporating the multiple time scale theory, allows the principle to account for the slower effects of increases as opposed to decreases in IFI (see Staddon et al. 2002).

Most importantly for this discussion, the linear waiting principle appears to be obligatory. That is, organisms seem to follow the linear waiting rule even if they delay or even prevent reinforcer delivery by doing so. The simplest example is the RID schedule itself. Wynne & Staddon (1988) showed that it makes no difference whether the experimenter holds delay time T constant or the sum of $t + T$ constant ($t + T = K$): Equation 1 holds in both cases, even though the optimal (reinforcement-rate-maximizing) strategy in the first case is for the animal to set

³When there is no response-produced stimulus change, this procedure is also called a conjunctive fixed-ratio fixed-time schedule (Shull 1970).

t equal to zero, whereas in the second case reinforcement rate is maximized so long as $t < K$. Using a version of RID in which T in interval $N + 1$ depended on the value of t in the preceding interval, Wynne & Staddon also demonstrated two kinds of instability predicted by linear waiting.

The fact that linear waiting is obligatory allows us to look for its effects on schedules other than the simple RID schedule. The most obvious application is to ratio schedules. The time to emit a fixed number of responses is approximately constant; hence the delay to food after the first response in each interval is also approximately constant on fixed ratio (FR), as on fixed- T RID (Powell 1968). Thus, the optimal strategy on FR, as on fixed- T RID, is to respond immediately after food. However, in both cases animals wait before responding and, as one might expect based on the assumption of a roughly constant interresponse time on all ratio schedules, the duration of the wait on FR is proportional to the ratio requirement (Powell 1968), although longer than on a comparable chain-type schedule with the same interreinforcement time (Crossman et al. 1974). The phenomenon of *ratio strain*—the appearance of long pauses and even extinction on high ratio schedules (Ferster & Skinner 1957)—may also have something to do with obligatory linear waiting.

Chain Schedules

A *chain schedule* is one in which a stimulus change, rather than primary reinforcement, is scheduled. Thus, a chain fixed-interval–fixed-interval schedule is one in which, for example, food reinforcement is followed by the onset of a red key light in the presence of which, after a fixed interval, a response produces a change to green. In the presence of green, food delivery is scheduled according to another fixed interval. RID schedules resemble two-link chain schedules. The first link is time t , before the animal responds; the second link is time T , after a response. We may expect, therefore, that waiting time in the first link of a two-link schedule will depend on the duration of the second link. We describe two results consistent with this conjecture and then discuss some exceptions.

Davison (1974) studied a two-link chain fixed-interval–fixed-interval schedule. Each cycle of the schedule began with a red key. Responding was reinforced, on fixed-interval I_1 s, by a change in key color from red to white. In the presence of white, food reinforcement was delivered according to fixed-interval I_2 s, followed by reappearance of the red key. Davison varied I_1 and I_2 and collected steady-state rate, pause, and link-duration data. He reported that when programmed second-link duration was long in relation to the first-link duration, pause in the first link sometimes exceeded the programmed link duration. The linear waiting predictions for this procedure can therefore be most easily derived for those conditions where the second link is held constant and the first link duration is varied (because under these conditions, the first-link pause was always less than the programmed first-link duration). The prediction for the terminal link is

$$t_2 = aI_2, \quad (2)$$

where a is the proportionality constant, I_2 is the duration of the terminal-link fixed-interval, and t_2 is the pause in the terminal link. Because I_2 is constant in this phase, t_2 is also constant. The pause in the initial link is given by

$$t_1 = a(I_1 + I_2) = aI_1 + aI_2, \quad (3)$$

where I_1 is the duration of the first link. Because I_2 is constant, Equation 3 is a straight line with slope a and positive y -intercept aI_2 .

Linear waiting theory can be tested with Davison's data by plotting, for every condition, t_1 and t_2 versus time-to-reinforcement (TTR); that is, plot pause in each link against TTR for that link in every condition. Linear waiting makes a straightforward prediction: All the data points for both links should lie on the same straight line through the origin (assuming that $b \rightarrow 0$). We show this plot in Figure 1. There is some variability, because the data points are individual subjects, not averages, but points from first and second links fit the same line, and the deviations do not seem to be systematic.

A study by Innis et al. (1993) provides a dynamic test of the linear waiting hypothesis as applied to chain schedules. Innis et al. studied two-link chain schedules with one link of fixed duration and the other varying from reinforcer to reinforcer




Figure 1 Steady-state pause duration plotted against actual time to reinforcement in the first and second links of a two-link chain schedule. Each data point is from a single pigeon in one experimental condition (three data points from an incomplete condition are omitted). (From Davison 1974, Table 1)

according to a triangular cycle. The dependent measure was pause in each link. Their Figure 3, for example, shows the programmed and actual values of the second link of the *constant-cycle* procedure (i.e., the first link was a constant 20 s; the second link varied from 5 to 35 s according to the triangular cycle) as well as the average pause, which clearly tracks the change in second-link duration with a lag of one interval. They found similar results for the reverse procedure, *cycle-constant*, in which the first link varied cyclically and the second link was constant. The tracking was a little better in the first procedure than in the second, but in both cases first-link pause was determined primarily by TTR.

There are some data suggesting that linear waiting is not the only factor that determines responding on simple chain schedules. In the four conditions of Davison's experiment in which the programmed durations of the first and second links added to a constant (120 s)—which implies a constant first-link pause according to linear waiting—pause in the first link covaried with first-link duration, although the data are noisy.

The alternative to the linear waiting account of responding on chain schedules is an account in terms of *conditioned reinforcement* (also called secondary reinforcement)—the idea that a stimulus paired with a primary reinforcer acquires some independent reinforcing power. This idea is also the organizing principle behind most theories of free-operant choice. There are some data that seem to imply a response-strengthening effect quite apart from the linear waiting effect, but they do not always hold up under closer inspection. Catania et al. (1980) reported that “higher rates of pecking were maintained by pigeons in the middle component of three-component chained fixed-interval schedules than in that component of the corresponding multiple schedule (two extinction components followed by a fixed-interval component)” (p. 213), but the effect was surprisingly small, given that no responding at all was required in the first two components. Moreover, results of a more critical control condition, chain versus tandem (rather than multiple) schedule, were the opposite: Rate was generally higher in the middle tandem component than in the second link of the chain. (A tandem schedule is one with the same response contingencies as a chain but with the same stimulus present throughout.)

Royalty et al. (1987) introduced a delay into the peck-stimulus-change contingency of a three-link variable-interval chain schedule and found large decreases in response rate [wait time (WT) was not reported] in both first and second links. They concluded that “because the effect of delaying stimulus change was comparable to the effect of delaying primary reinforcement in a simple variable-interval schedule . . . the results provide strong evidence for the concept of conditioned reinforcement” (p. 41). The implications of the Royalty et al. data for linear waiting are unclear, however, (*a*) because the linear waiting hypothesis does not deal with the assignment-of-credit problem, that is, the selection of the appropriate response by the schedule. Linear waiting makes predictions about response timing—when the operant response occurs—but not about which response will occur. Response-reinforcer contiguity may be essential for the selection of the operant response

in each chain link (as it clearly is during “shaping”), and diminishing contiguity may reduce response rate, but contiguity may play little or no role in the timing of the response. The idea of conditioned reinforcement may well apply to the first function but not to the second. (b) Moreover, Royalty et al. did not report *obtained time-to-reinforcement* data; the effect of the imposed delay may therefore have been via an increase in component duration rather than directly on response rate.

Williams & Royalty (1990) explicitly compared conditioned reinforcement and time to reinforcement as explanations for chain schedule performance in three-link chains and concluded “that time to reinforcement itself accounts for little if any variance in initial-link responding” (p. 381) but not timing, which was not measured. However, these data are from chain schedules with both variable-interval and fixed-interval links, rather than fixed-interval only, and with respect to response rate rather than pause measures. In a later paper Williams qualified this claim: “The effects of stimuli in a chain schedule are due partly to the time to food correlated with the stimuli and partly to the time to the next conditioned reinforcer in the sequence” (1997, p. 145).

The conclusion seems to be that linear waiting plays a relatively major, and conditioned reinforcement (however defined) a relatively minor, role in the determination of response timing on chain fixed-interval schedules. Linear waiting also provides the best available account of a striking, unsolved problem with chain schedules: the fact that in chains with several links, pigeon subjects may respond at a low level or even quit completely in early links (Catania 1979, Gollub 1977). On fixed-interval chain schedules with five or more links, responding in the early links begins to extinguish and the overall reinforcement rate falls well below the maximum possible—even if the programmed interreinforcement interval is relatively short (e.g., $6 \times 15 = 90$ s). If the same stimulus is present in all links (tandem schedule), or if the six different stimuli are presented in random order (scrambled-stimuli chains), performance is maintained in all links and the overall reinforcement rate is close to the maximum possible ($6I$, where I is the interval length). Other studies have reported very weak responding in early components of a simple chain fixed-interval schedule (e.g., Catania et al. 1980, Davison 1974, Williams 1994; review in Kelleher & Gollub 1962). These studies found that chains with as few as three fixed-interval 60-s links (Kelleher & Fry 1962) occasionally produce extreme pausing in the first link. No formal theory of the kind that has proliferated to explain behavior on concurrent chain schedules (discussed below) has been offered to account for these strange results, even though they have been well known for many years.

The informal suggestion is that the low or zero response rates maintained by early components of a multi-link chain are a consequence of the same discrimination process that leads to extinction in the absence of primary reinforcement. Conversely, the stimulus at the end of the chain that is actually paired with primary reinforcement is assumed to be a conditioned reinforcer; stimuli in the middle sustain responding because they lead to production of a conditioned reinforcer

(Catania et al. 1980, Kelleher & Gollub 1962). Pairing also explains why behavior is maintained on tandem and scrambled-stimuli chains (Kelleher & Fry 1962). In both cases the stimuli early in the chain are either invariably (tandem) or occasionally (scrambled-stimulus) paired with primary reinforcement.

There are problems with the conditioned-reinforcement approach, however. It can explain responding in link two of a three-link chain but not in link one, which should be an extinction stimulus. The explanatory problem gets worse when more links are added. There is no well-defined principle to tell us when a stimulus changes from being a conditioned reinforcer, to a stimulus in whose presence responding is maintained by a conditioned reinforcer, to an extinction stimulus. What determines the stimulus property? Is it stimulus number, stimulus duration or the durations of stimuli later in the chain? Perhaps there is some balance between contrast/extinction, which depresses responding in early links, and conditioned reinforcement, which is supposed to (but sometimes does not) elevate responding in later links? No well-defined compound theory has been offered, even though there are several quantitative theories for multiple-schedule contrast (e.g., Herrnstein 1970, Nevin 1974, Staddon 1982; see review in Williams 1988). There are also data that cast doubt even on the idea that late-link stimuli have a rate-enhancing effect. In the Catania et al. (1980) study, for example, four of five pigeons responded faster in the middle link of a three-link tandem schedule than the comparable chain.

The lack of formal theories for performance on simple chains is matched by a dearth of data. Some pause data are presented in the study by Davison (1974) on pigeons in a two-link fixed-interval chain. The paper attempted to fit Herrnstein's (1970) matching law between response rates and link duration. The match was poor: The pigeon's rates fell more than predicted when the terminal links (contiguous with primary reinforcement) of the chain were long, but Davison did find that "the terminal link schedule clearly changes the pause in the initial link, longer terminal-link intervals giving longer initial-link pauses" (1974, p. 326). Davison's abstract concludes, "Data on pauses during the interval schedules showed that, in most conditions, the pause duration was a linear function of the interval length, and greater in the initial link than in the terminal link" (p. 323). In short, the pause (time-to-first-response) data were more lawful than response-rate data.

Linear waiting provides a simple explanation for excessive pausing on multi-link chain fixed-interval schedules. Suppose the chief function of the link stimuli on chain schedules is simply to signal changing times to primary reinforcement⁴.

⁴This idea surfaced very early in the history of research on equal-link chain fixed-interval schedules, but because of the presumed importance of conditioned reinforcement, it was the time to reinforcement from link stimulus offset, rather than onset that was thought to be important. Thus, Gollub (1977), echoing his 1958 Ph.D. dissertation in the subsequent Kelleher & Gollub (1962) review, wrote, "In chained schedules with more than two components . . . the extent to which responding is sustained in the initial components . . . depends on the time that elapses from the end of the components to food reinforcement" (p. 291).

Thus, in a three-link fixed-interval chain, with link duration I , the TTR signaled by the end of reinforcement (or by the onset of the first link) is $3I$. The onset of the next link signals a TTR of $2I$ and the terminal, third, link signals a TTR of I . The assumptions of linear waiting as applied to this situation are that pausing (time to first response) in each link is determined entirely by TTR and that the wait time in interval $N+1$ is a linear function of the TTR in the preceding interval.

To see the implications of this process, consider again a three-link chain schedule with $I=1$ (arbitrary time units). The performance to be expected depends entirely on the value of the proportionality constant, a , that sets the fraction of time-to-primary-reinforcement that the animal waits (for simplicity we can neglect b ; the logic of the argument is unaffected). All is well so long as a is less than one-third. If a is exactly 0.333, then for unit link duration the pause in the third link is 0.33, in the second link 0.67, and in the first link 1.0. However, if a is larger, for instance 0.5, the three pauses become 0.5, 1.0, and 1.5; that is, the pause in the first link is now longer than the programmed interval, which means the TTR in the first link will be longer than 3 the next time around, so the pause will increase further, and so on until the process stabilizes (which it always does: First-link pause never goes to ∞).

The steady-state wait times in each link predicted for a five-link chain, with unit-duration links, for two values of a are shown in Figure 2. In both cases wait times in the early links are very much longer than the programmed link duration. Clearly, this process has the potential to produce very large pauses in the early links of multilink-chain fixed-interval schedules and so may account for the data Catania (1979) and others have reported.

Gollub in his dissertation research (1958) noticed the additivity of this sequential pausing. Kelleher & Gollub (1962) in their subsequent review wrote, "No two pauses in [simple fixed interval] can both postpone food-delivery; however, pauses in different components of [a] five-component chain will postpone food-delivery additively" (p. 566). However, this additivity was only one of a number of processes suggested to account for the long pauses in early chain fixed-interval links, and its quantitative implications were never explored.

Note that the linear waiting hypothesis also accounts for the relative stability of tandem schedules and chain schedules with scrambled components. In the tandem schedule, reinforcement constitutes the only available time marker. Given that responding after the pause continues at a relatively high rate until the next time marker, Equation 1 (with b assumed negligible) and a little algebra shows that the steady-state postreinforcement pause for a tandem schedule with unit links will be

$$t = \frac{a(N-1)}{1-a}, \quad \text{if } t \geq 1, \quad (4)$$

where N is the number of links and a is the pause fraction. In the absence of any time markers, pauses in links after the first are necessarily short, so the experienced link duration equals the programmed duration. Thus, the total interfood-reinforcement




Figure 2 Wait time (pause, time to first response) in each equal-duration link of a five-link chain schedule (as a multiple of the programmed link duration) as predicted by the linear-waiting hypothesis. The two curves are for two values of parameter a in Equation 1 ($b = 0$). Note the very long pauses predicted in early links—almost two orders of magnitude greater than the programmed interval in the first link for $a = 0.67$. (From Mazur 2001)

interval will be $t + N - 1$ ($t \geq 1$): the pause in the first link (which will be longer than the programmed link duration for $N > 1/a$) plus the programmed durations of the succeeding links. For the case of $a = 0.67$ and unit link duration, which yielded a steady-state interfood interval (IFI) of 84 for the five-link chain schedule, the tandem yields 12. For $a = 0.5$, the two values are approximately 16 and 8.

The long waits in early links shown in Figure 2 depend critically on the value of a . If, as experience suggests (there has been no formal study), a tends to increase slowly with training, we might expect the long pausing in initial links to take some time to develop, which apparently it does (Gollub 1958).

On the scrambled-stimuli chain each stimulus occasionally ends in reinforcement, so each signals a time-to-reinforcement (TTR)⁵ of I , and pause in each link should be less than the link duration—yielding a total IFI of approximately N ,

⁵Interpreted as time to the first reinforcement opportunity.

i.e., 5 for the example in the figure. These predictions yield the order IFI in the chain > tandem > scrambled, but parametric data are not available for precise comparison. We do not know whether an N -link scrambled schedule typically stabilizes at a shorter IFI than the comparable tandem schedule, for example. Nor do we know whether steady-state pause in successive links of a multilink chain falls off in the exponential fashion shown in Figure 2.

In the final section we explore the implications of linear waiting for studies of free-operant choice behavior.

CHOICE

Although we can devote only limited space to it, choice is one of the major research topics in operant conditioning (see Mazur 2001, p. 96 for recent statistics). Choice is not something that can be directly observed. The subject does this or that and, in consequence, is said to choose. The term has unfortunate overtones of conscious deliberation and weighing of alternatives for which the behavior itself—response A or response B—provides no direct evidence. One result has been the assumption that the proper framework for all so-called choice studies is in terms of response strength and the value of the choice alternatives. Another is the assumption that procedures that are very different are nevertheless studying the same thing.

For example, in a classic series of experiments, Kahneman & Tversky (e.g., 1979) asked a number of human subjects to make a single choice of the following sort: between \$400 for sure and a 50% chance of \$1000. Most went for the sure thing, even though the expected value of the gamble is higher. This is termed *risk aversion*, and the same term has been applied to free-operant “choice” experiments. In one such experiment an animal subject must choose repeatedly between a response leading to a fixed amount of food and one leading equiprobably to either a large or a small amount with the same average value. Here the animals tend to be either indifferent or risk averse, preferring the fixed alternative (Staddon & Innis 1966b, Bateson & Kacelnik 1995, Kacelnik & Bateson 1996).

In a second example pigeons responded repeatedly to two keys associated with equal variable-interval schedules. A successful response on the left key, for example, is reinforced by a change in the color of the pecked key (the other key light goes off). In the presence of this second stimulus, food is delivered according to a fixed-interval schedule (fixed-interval X). The first stimulus, which is usually the same on both keys, is termed the *initial link*; the second stimulus is the *terminal link*. Pecks on the right key lead in the same way to food reinforcement on variable-interval X . (This is termed a *concurrent-chain* schedule.) In this case subjects overwhelmingly prefer the initial-link choice leading to the variable-interval terminal link; that is, they are apparently risk seeking rather than risk averse (Killeen 1968).

The fact that these three experiments (Kahneman & Tversky and the two free-operant studies) all produce different results is sometimes thought to pose a serious research problem, but, we contend, the problem is only in the use of the term *choice*

for all three. The procedures (not to mention the subjects) are in fact very different, and in operant conditioning the devil is very much in the details. Apparently trivial procedural differences can sometimes lead to wildly different behavioral outcomes. Use of the term *choice* as if it denoted a unitary subject matter is therefore highly misleading. We also question the idea that the results of choice experiments are always best explained in terms of response strength and stimulus value.

Concurrent Schedules

Bearing these caveats in mind, let's look briefly at the extensive history of free-operant choice research. In Herrnstein's seminal experiment (1961; see Davison & McCarthy 1988, Williams 1988 for reviews; for collected papers see Rachlin & Laibson 1997) hungry pigeons pecked at two side-by-side response keys, one associated with variable-interval v_1 s and the other with variable-interval v_2 s (*concurrent variable-interval–variable-interval* schedule). After several experimental sessions and a range of v_1 and v_2 values chosen so that the overall programmed reinforcement rate was constant ($1/v_1 + 1/v_2 = \text{constant}$), the result was matching between steady-state relative response rates and relative *obtained* reinforcement rates:

$$\frac{x}{y} = \frac{R(x)}{R(y)}, \quad (5)$$

where x and y are the response rates on the two alternatives and $R(x)$ and $R(y)$ are the rates of obtained reinforcement for them. This relation has become known as Herrnstein's matching law. Although the obtained reinforcement rates are dependent on the response rates that produce them, the matching relation is not forced, because x and y can vary over quite a wide range without much effect on $R(x)$ and $R(y)$.

Because of the negative feedback relation intrinsic to variable-interval schedules (the less you respond, the higher the probability of payoff), the matching law on concurrent variable-interval–variable-interval is consistent with reinforcement maximization (Staddon & Motheral 1978), although the maximum of the function relating overall payoff, $R(x) + R(y)$, to relative responding, $x/(x + y)$, is pretty flat. However, little else on these schedules fits the maximization idea. As noted above, even responding on simple fixed- T response-initiated delay (RID) schedules violates maximization. Matching is also highly overdetermined, in the sense that almost any learning rule consistent with the law of effect—an increase in reinforcement probability causes an increase in response probability—will yield either simple matching (Equation 5) or its power-law generalization (Baum 1974, Hinson & Staddon 1983, Lander & Irwin 1968, Staddon 1968). Matching by itself therefore reveals relatively little about the dynamic processes operating in the responding subject (but see Davison & Baum 2000). Despite this limitation, the strikingly regular functional relations characteristic of free-operant choice studies have attracted a great deal of experimental and theoretical attention.

Herrnstein (1970) proposed that Equation 5 can be derived from the function relating steady-state response rate, x , and reinforcement rate, $R(x)$, to each response key considered separately. This function is negatively accelerated and well approximated by a hyperbola:

$$x = \frac{kR(x)}{R(x) + R_0}, \quad (6)$$

where k is a constant and R_0 represents the effects of all other reinforcers in the situation. The denominator and parameter k cancel in the ratio x/y , yielding Equation 5 for the choice situation.

There are numerous empirical details that are not accounted for by this formulation: systematic deviations from matching [undermatching and overmatching (Baum 1974)] as a function of different types of variable-interval schedules, dependence of simple matching on use of a *changeover delay*, extensions to concurrent-chain schedules, and so on. For example, if animals are pretrained with two alternatives presented separately, so that they do not learn to switch between them, when given the opportunity to respond to both, they fixate on the richer one rather than matching [extreme overmatching (Donahoe & Palmer 1994, pp. 112–113; Gallistel & Gibbon 2000, pp. 321–322)]. (Fixation—extreme overmatching—is, trivially, matching, of course but if only fixation were observed, the idea of matching would never have arisen. Matching implies partial, not exclusive, preference.) Conversely, in the absence of a changeover delay, pigeons will often just alternate between two unequal variable-interval choices [extreme undermatching (Shull & Pliskoff 1967)]. In short, matching requires exactly the right amount of switching. Nevertheless, Herrnstein's idea of deriving behavior in choice experiments from the laws that govern responding to the choice alternatives in isolation is clearly worth pursuing.

In any event, Herrnstein's approach—molar data, predominantly variable-interval schedules, rate measures—set the basic pattern for subsequent operant choice research. It fits the basic presuppositions of the field: that choice is about *response strength*, that response strength is equivalent to response probability, and that response rate is a valid proxy for probability (e.g., Skinner 1938, 1966, 1986; Killeen & Hall 2001). (For typical studies in this tradition see, e.g., Fantino 1981; Grace 1994; Herrnstein 1961, 1964, 1970; Rachlin et al. 1976; see also Shimp 1969, 2001.)

We can also look at concurrent schedules in terms of linear waiting. Although published evidence is skimpy, recent unpublished data (Cerutti & Staddon 2002) show that even on variable-interval schedules (which necessarily always contain a few very short interfood intervals), postfood wait time and changeover time covary with mean interfood time. It has also long been known that Equation 6 can be derived from two time-based assumptions: that the number of responses emitted is proportional to the number of reinforcers received multiplied by the available time and that available time is limited by the time taken up by each response (Staddon 1977, Equations 23–25). Moreover, if we define mean interresponse time as the

reciprocal of mean response rate,⁶ x , and mean interfood interval is the reciprocal of obtained reinforcement rate, $R(x)$, then linear waiting yields

$$1/x = a/R(x) + b,$$

where a and b are linear waiting constants. Rearranging yields

$$x = \frac{\frac{1}{b}R(x)}{\frac{a}{b} + R(x)}, \quad (7)$$

where $1/b = k$ and $a/b = R_0$ in Equation 6. Both these derivations of the hyperbola in Equation 6 from a linear relation in the time domain imply a correlation between parameters k and R_0 in Equation 6 under parametric experimental variation of parameter b by (for example) varying response effort or, possibly, hunger motivation. Such covariation has been occasionally but not universally reported (Dallery et al. 2000, Heyman & Monaghan 1987, McDowell & Dallery 1999).

Concurrent-Chain Schedules

Organisms can be trained to choose between sources of primary reinforcement (concurrent schedules) or between stimuli that signal the occurrence of primary reinforcement (*conditioned reinforcement*: concurrent chain schedules). Many experimental and theoretical papers on conditioned reinforcement in pigeons and rats have been published since the early 1960s using some version of the concurrent chains procedure of Autor (1960, 1969). These studies have demonstrated a number of functional relations between rate measures and have led to several closely related theoretical proposals such as a version of the matching law, incentive theory, delay-reduction theory, and hyperbolic value-addition (e.g., Fantino 1969a,b; Grace 1994; Herrnstein 1964; Killeen 1982; Killeen & Fantino 1990; Mazur 1997, 2001; Williams 1988, 1994, 1997). Nevertheless, there is as yet no theoretical consensus on how best to describe choice between sources of conditioned reinforcement, and no one has proposed an integrated theoretical account of simple chain and concurrent chain schedules.

Molar response rate does not capture the essential feature of behavior on fixed-interval schedules: the systematic pattern of rate-change in each interfood interval, the “scallop.” Hence, the emphasis on molar response rate as a dependent variable has meant that work on concurrent schedules has emphasized variable or random intervals over fixed intervals. We lack any theoretical account of concurrent fixed-interval–fixed-interval and fixed-interval–variable-interval schedules. However, a recent study by Shull et al. (2001; see also Shull 1979) suggests that response rate may not capture what is going on even on simple variable-interval schedules, where the time to initiate bouts of relatively fixed-rate responding seems to be a

⁶It is not of course: The reciprocal of the mean IRT is the harmonic mean rate. In practice, “mean response rate” usually means arithmetic mean, but note that harmonic mean rate usually works better for choice data than the arithmetic mean (cf. Killeen 1968).

more sensitive dependent measure than overall response rate. More attention to the role of temporal variables in choice is called for.

We conclude with a brief account of how linear waiting may be involved in several well-established phenomena of concurrent-chain schedules: preference for variable-interval versus fixed-interval terminal links, effect of initial-link duration, and finally, so-called self-control experiments.

PREFERENCE FOR VARIABLE-INTERVAL VERSUS FIXED-INTERVAL TERMINAL LINKS On concurrent-chain schedules with equal variable-interval initial links, animals show a strong preference for the initial link leading to a variable-interval terminal link over the terminal-link alternative with an equal arithmetic-mean fixed interval. This result is usually interpreted as a manifestation of nonarithmetic (e.g., harmonic) reinforcement-rate averaging (Killeen 1968), but it can also be interpreted as linear waiting. Minimum TTR is necessarily much less on the variable-interval than on the fixed-interval side, because some variable intervals are short. If wait time is determined by minimum TTR—hence shorter wait times on the variable-interval side—and ratios of wait times and overall response rates are (inversely) correlated (Cerutti & Staddon 2002), the result will be an apparent bias in favor of the variable-interval choice.

EFFECT OF INITIAL-LINK DURATION Preference for a given pair of terminal-link schedules depends on initial link duration. For example, pigeons may approximately match initial-link relative response rates to terminal-link relative reinforcement rates when the initial links are 60 s and the terminal links range from 15 to 45 s (Herrnstein 1964), but they will undermatch when the initial-link schedule is increased to, for example, 180 s. This effect is what led to Fantino's delay-reduction modification of Herrnstein's matching law (see Fantino et al. 1993 for a review). However, the same qualitative prediction follows from linear waiting: Increasing initial-link duration reduces the proportional TTR difference between the two choices. Hence the ratio of WTs or of initial-link response rates for the two choices should also approach unity, which is undermatching. Several other well-studied theories of concurrent choice, such as delay reduction and hyperbolic value addition, also explain these results.

Self-Control

The prototypical self-control experiment has a subject choosing between two outcomes: not-so-good cookie now or a good cookie after some delay (Rachlin & Green 1972; see Logue 1988 for a review; Mischel et al. 1989 reviewed human studies). Typically, the subject chooses the immediate, small reward, but if both delays are increased by the same amount, D , he will learn to choose the larger reward, providing D is long enough. Why? The standard answer is derived from Herrnstein's matching analysis (Herrnstein 1981) and is called *hyperbolic discounting* (see Mazur 2001 for a review and Ainslie 1992 and Rachlin 2000 for

longer accounts). The idea is that the expected value of each reward is inversely related to the time at which it is expected according to a hyperbolic function:

$$V_i = \frac{A_i}{1 + kD_i}, \quad (8)$$

where A_i is the *undiscounted value* of the reward, D_i is the delay until reward is received, i denotes the large or small reward, and k is a fitted constant.

Now suppose we set D_L and D_S to values such that the animal shows a preference for the shorter, sooner reward. This would be the case ($k = 1$) if $A_L = 6$, $A_S = 2$, $D_L = 6$ s, and $D_S = 1$ s: $V_L = 0.86$ and $V_S = 1$ —preference for the small, less-delayed reward. If 10 s is added to both delays, so that $D_L = 16$ s and $D_S = 11$ s, the values are $V_L = 0.35$ and $V_S = 0.17$ —preference for the larger reward. Thus, Equation 8 predicts that added delay—sometimes awkwardly termed *pre-commitment*—should enhance self-control, which it does.

The most dramatic prediction from this analysis was made and confirmed by Mazur (1987, 2001) in an experiment that used an *adjusting-delay* procedure (also termed *titration*). “A response on the center key started each trial, and then a pigeon chose either a standard alternative (by pecking the red key) or an adjusting alternative (by pecking the green key) . . . the standard alternative delivered 2 s of access to grain after a 10-s delay, and the adjusting alternative delivered 6 s of access to grain after an adjusting delay” (2001, p. 97). The adjusting delay increased (on the next trial) when it was chosen and decreased when the standard alternative was chosen. (See Mazur 2001 for other procedural details.) The relevant independent variable is TTR. The discounted value of each choice is given by Equation 8. When the subject is indifferent does not discriminate between the two choices, $V_L = V_S$. Equating Equation 8 for the large and small choices yields

$$D_L = \frac{A_L}{A_S} \cdot D_S + \frac{A_L - A_S}{kA_S}; \quad (9)$$

that is, an *indifference curve* that is a linear function relating D_L and D_S , with slope $A_L/A_S > 1$ and a positive intercept. The data (Mazur 1987; 2001, Figure 2) are consistent with this prediction, but the intercept is small.

It is also possible to look at this situation in terms of linear waiting. One assumption is necessary: that the waiting fraction, a , in Equation 1 is smaller when the upcoming reinforcer is large than when it is small (Powell 1969 and Perone & Courtney 1992 showed this for fixed-ratio schedules; Howerton & Meltzer 1983, for fixed-interval). Given this assumption, the linear waiting analysis is even simpler than hyperbolic discounting. The idea is that the subject will appear to be indifferent when the wait times to the two alternatives are equal. According to linear waiting, the wait time for the small alternative is given by

$$t_S = a_S D_S + b_S, \quad (10)$$

where b_S is a small positive intercept and $a_S > a_L$. Equating the wait times for small

and large alternatives yields

$$D_L = \frac{a_S}{a_L} \cdot D_S + \frac{b_S - b_L}{a_L}, \quad (11)$$

which is also a linear function with slope > 1 and a small positive intercept.

Equations 9 and 11 are identical in form. Thus, the linear waiting and hyperbolic discounting models are almost indistinguishable in terms of these data. However, the linear waiting approach has three potential advantages: Parameters a and b can be independently measured by making appropriate measurements in a control study that retains the reinforcement-delay properties of the self-control experiments without the choice contingency; the linear waiting approach lacks the fitted parameter k in Equation 9; and linear waiting also applies to a wide range of time-production experiments not covered by the hyperbolic discounting approach.

CONCLUSION

Temporal control may be involved in unsuspected ways in a wide variety of operant conditioning procedures. A renewed emphasis on the causal factors operating in reinforcement schedules may help to unify research that has hitherto been defined in terms of more abstract topics like timing and choice.

ACKNOWLEDGMENTS

We thank Catalin Buhusi and Jim Mazur for comments on an earlier version and the NIMH for research support over many years.

The *Annual Review of Psychology* is online at <http://psych.annualreviews.org>

LITERATURE CITED

- Ainslie G. 1992. *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*. Cambridge, MA: Harvard Univ. Press
- Autor SM. 1960. *The strength of conditioned reinforcers as a function of frequency and probability of reinforcement*. PhD thesis. Harvard Univ., Cambridge, MA
- Autor SM. 1969. The strength of conditioned reinforcers and a function of frequency and probability of reinforcement. In *Conditioned Reinforcement*, ed. DP Hendry, pp. 127–62. Homewood, IL: Dorsey
- Bateson M, Kacelnik A. 1995. Preferences for fixed and variable food sources: variability in amount and delay. *J. Exp. Anal. Behav.* 63:313–29
- Baum WM. 1974. On two types of deviation from the matching law: bias and undermatching. *J. Exp. Anal. Behav.* 22:231–42
- Baum WM. 1994. *Understanding Behaviorism: Science, Behavior and Culture*. New York: HarperCollins
- Blough DS, Millward RB. 1965. Learning: operant conditioning and verbal learning. *Annu. Rev. Psychol.* 17:63–94
- Buhusi CV, Meck WH. 2000. Timing for the absence of the stimulus: the gap paradigm

- reversed. *J. Exp. Psychol.: Anim. Behav. Process.* 26:305–22
- Cabeza de Vaca S, Brown BL, Hemmes NS. 1994. Internal clock and memory processes in animal timing. *J. Exp. Psychol.: Anim. Behav. Process.* 20:184–98
- Catania AC. 1970. Reinforcement schedules and psychophysical judgments: a study of some temporal properties of behavior. In *The Theory of Reinforcement Schedules*, ed. WN Schoenfeld, pp. 1–42. New York: Appleton-Century-Crofts
- Catania AC. 1979. *Learning*. Englewood Cliffs, NJ: Prentice-Hall
- Catania AC, Yohalem R, Silverman PJ. 1980. Contingency and stimulus change in chained schedules of reinforcement. *J. Exp. Anal. Behav.* 5:167–73
- Cerutti DT, Staddon JER. 2002. The temporal dynamics of choice: concurrent and concurrent-chain interval schedules. Submitted
- Cheng K, Westwood R. 1993. Analysis of single trials in pigeons' timing performance. *J. Exp. Psychol.: Anim. Behav. Process.* 19:56–67
- Church RM. 1978. The internal clock. In *Cognitive Processes in Animal Behavior*, ed. SH Hulse, H Fowler, WK Honig, pp. 277–310. Hillsdale, NJ: Erlbaum
- Church RM, Broadbent HA. 1990. Alternative representations of time, number and rate. *Cognition* 37:55–81
- Church RM, Deluty MZ. 1977. Bisection of temporal intervals. *J. Exp. Psychol.: Anim. Behav. Process.* 3:216–28
- Church RM, Miller KD, Meck WH. 1991. Symmetrical and asymmetrical sources of variance in temporal generalization. *Anim. Learn. Behav.* 19:135–55
- Crossman EK, Heaps RS, Nunes DL, Alferink LA. 1974. The effects of number of responses on pause length with temporal variables controlled. *J. Exp. Anal. Behav.* 22:115–20
- Dallery J, McDowell JJ, Lancaster JS. 2000. Falsification of matching theory's account of single-alternative responding: Herrnstein's K varies with sucrose concentration. *J. Exp. Anal. Behav.* 73:23–43
- Davison M. 1974. A functional analysis of chained fixed-interval schedule performance. *J. Exp. Anal. Behav.* 21:323–30
- Davison M, Baum W. 2000. Choice in a variable environment: Every reinforcer counts. *J. Exp. Anal. Behav.* 74:1–24
- Davison M, McCarthy D. 1988. *The Matching Law: A Research Review*. Hillsdale, NJ: Erlbaum
- Donahoe JW, Palmer DC. 1994. *Learning and Complex Behavior*. Boston: Allyn & Bacon
- Dragoi V, Staddon JER, Palmer RG, Buhusi VC. 2002. Interval timing as an emergent learning property. *Psychol. Rev.* In press
- Dreyfus LR, Fetterman JG, Smith LD, Stubbs DA. 1988. Discrimination of temporal relations by pigeons. *J. Exp. Psychol.: Anim. Behav. Process.* 14:349–67
- Fantino E. 1969a. Choice and rate of reinforcement. *J. Exp. Anal. Behav.* 12:723–30
- Fantino E. 1969b. Conditioned reinforcement, choice, and the psychological distance to reward. In *Conditioned Reinforcement*, ed. DP Hendry, pp. 163–91. Homewood, IL: Dorsey
- Fantino E. 1981. Contiguity, response strength, and the delay-reduction hypothesis. In *Advances in Analysis of Behavior: Predictability, Correlation, and Contiguity*, ed. P Harzem, M Zeiler, 2:169–201. Chichester, UK: Wiley
- Fantino E, Preston RA, Dunn R. 1993. Delay reduction: current status. *J. Exp. Anal. Behav.* 60:159–69
- Ferster CB, Skinner BF. 1957. *Schedules of Reinforcement*. New York: Appleton-Century-Crofts
- Fetterman JG, Dreyfus LR, Stubbs DA. 1989. Discrimination of duration ratios. *J. Exp. Psychol.: Anim. Behav. Process.* 15:253–63
- Gallistel CR. 1990. *The Organization of Learning*. Cambridge, MA: MIT/Bradford
- Gallistel CR, Gibbon J. 2000. Time, rate, and conditioning. *Psychol. Rev.* 107:289–344
- Gibbon J. 1977. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* 84:279–325
- Gibbon J, Church RM. 1984. Sources of variance in an information processing theory of

- timing. In *Animal Cognition*, ed. HL Roitblat, TG Bever, HS Terrace. Hillsdale, NJ: Erlbaum.
- Gollub LR. 1958. *The chaining of fixed-interval schedules*. Unpublished doctoral dissertation, Harvard Univ.
- Gollub L. 1977. Conditioned reinforcement: schedule effects. See Honig & Staddon 1977, pp. 288–312
- Grace RC. 1994. A contextual choice model of concurrent-chains choice. *J. Exp. Anal. Behav.* 61:113–29
- Grace RC, Nevin JA. 2000. Response strength and temporal control in fixed-interval schedules. *Anim. Learn. Behav.* 28:313–31
- Grossberg S, Schmajuk NA. 1989. Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Netw.* 2:79–102
- Hanson SJ, Killeen PR. 1981. Measurement and modeling of behavior under fixed-interval schedules of reinforcement. *J. Exp. Psychol.: Anim. Behav. Process.* 7:129–39
- Herrnstein RJ. 1961. Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* 4:267–72
- Herrnstein RJ. 1964. Secondary reinforcement and rate of primary reinforcement. *J. Exp. Anal. Behav.* 7:27–36
- Herrnstein RJ. 1970. On the law of effect. *J. Exp. Anal. Behav.* 13:243–66
- Herrnstein RJ. 1981. Self control as response strength. In *Recent Developments in the Quantification of Steady-State Operant Behavior*, ed. CM Bradshaw, CP Lowe, F Szabadi, pp. 3–20. Amsterdam: Elsevier/North-Holland
- Heyman GM, Monaghan MM. 1987. Effects of changes in response requirements and deprivation on the parameters of the matching law equation: new data and review. *J. Exp. Psychol.: Anim. Behav. Process.* 13:384–94
- Higa JJ, Wynne CDL, Staddon JER. 1991. Dynamics of time discrimination. *J. Exp. Psychol.: Anim. Behav. Process.* 17:281–91
- Hinson JM, Staddon JER. 1983. Matching, maximizing and hill climbing. *J. Exp. Anal. Behav.* 40:321–31
- Honig WK, Staddon JER, eds. 1977. *Handbook of Operant Behavior*. Englewood Cliffs, NJ: Prentice-Hall
- Hopson JW. 1999. Gap timing and the spectral timing model. *Behav. Process.* 45:23–31
- Hopson JW. 2002. *Timing without a clock: learning models as interval timing models*. PhD thesis. Duke Univ., Durham, NC
- Howerton L, Meltzer D. 1983. Pigeons' FI behavior following signaled reinforcement duration. *Bull. Psychon. Soc.* 21:161–63
- Innis NK, Mitchell S, Staddon JER. 1993. Temporal control on interval schedules: What determines the postreinforcement pause? *J. Exp. Anal. Behav.* 60:293–311
- Kacelnik A, Bateson M. 1996. Risky theories—the effects of variance on foraging decisions. *Am. Zool.* 36:402–34
- Kahneman D, Tversky A. 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47:263–91
- Kelleher RT, Fry WT. 1962. Stimulus functions in chained and fixed-interval schedules. *J. Exp. Anal. Behav.* 5:167–73
- Kelleher RT, Gollub LR. 1962. A review of positive conditioned reinforcement. *J. Exp. Anal. Behav.* 5:541–97
- Kello JE. 1972. The reinforcement-omission effect on fixed-interval schedules: frustration or inhibition? *Learn. Motiv.* 3:138–47
- Killeen PR. 1968. On the measurement of reinforcement frequency in the study of preference. *J. Exp. Anal. Behav.* 11:263–69
- Killeen PR. 1982. Incentive theory: II. Models for choice. *J. Exp. Anal. Behav.* 38:217–32
- Killeen PR, Fantino E. 1990. Unification of models for choice between delayed reinforcers. *J. Exp. Anal. Behav.* 53:189–200
- Killeen PR, Fetterman JG. 1988. A behavioral theory of timing. *Psychol. Rev.* 95:274–95
- Killeen PR, Hall SS. 2001. The principal components of response strength. *J. Exp. Anal. Behav.* 75:111–34

- Lander DG, Irwin RJ. 1968. Multiple schedules: effects of the distribution of reinforcements between components on the distribution of responses between components. *J. Exp. Anal. Behav.* 11:517–24
- Lejeune H, Ferrara A, Simons F, Wearden JH. 1997. Adjusting to changes in the time of reinforcement: peak-interval transitions in rats. *J. Exp. Psychol.: Anim. Behav. Process.* 23:211–321
- Logue AW. 1988. Research on self-control: an integrating framework. *Behav. Brain Sci.* 11:665–709
- Lowe CF, Harzem P, Spencer PT. 1979. Temporal control of behavior and the power law. *J. Exp. Anal. Behav.* 31:333–43
- MacEwen D, Killeen P. 1991. The effects of rate and amount on the speed of the pacemaker in pigeons' timing behavior. *Anim. Learn. Behav.* 19:164–70
- Machado A. 1997. Learning the temporal dynamics of behavior. *Psychol. Rev.* 104:241–65
- Matell MS, Meck WH. 1999. Reinforcement-induced within-trial resetting of an internal clock. *Behav. Process.* 45:159–71
- Mazur JE. 1987. An adjusting procedure for studying delayed reinforcement. In *Quantitative Analyses of Behavior*, Vol. 5. *The Effects of Delay and Intervening Events on Reinforcement Value*, ed. ML Commons, JE Mazur, JA Nevin, H Rachlin, pp. 55–73. Mahwah, NJ: Erlbaum
- Mazur JE. 1997. Choice, delay, probability, and conditioned reinforcement. *Anim. Learn. Behav.* 25:131–47
- Mazur JE. 2001. Hyperbolic value addition and general models of animal choice. *Psychol. Rev.* 108:96–112
- McDowell JJ, Dallery J. 1999. Falsification of matching theory: changes in the asymptote of Herrnstein's hyperbola as a function of water deprivation. *J. Exp. Anal. Behav.* 72:251–68
- Meck WH. 1983. Selective adjustment of the speed of an internal clock and memory processes. *J. Exp. Psychol.: Anim. Behav. Process.* 9:171–201
- Meck WH, Komeily-Zadeh FN, Church RM. 1984. Two-step acquisition: modification of an internal clock's criterion. *J. Exp. Psychol.: Anim. Behav. Process.* 10:297–306
- Mischel W, Shoda Y, Rodriguez M. 1989. Delay of gratification for children. *Science* 244:933–38
- Nevin JA. 1974. Response strength in multiple schedules. *J. Exp. Anal. Behav.* 21:389–408
- Perone M, Courtney K. 1992. Fixed-ratio pausing: joint effects of past reinforcer magnitude and stimuli correlated with upcoming magnitude. *J. Exp. Anal. Behav.* 57:33–46
- Plowright CMS, Church D, Behnke P, Silverman A. 2000. Time estimation by pigeons on a fixed interval: the effect of pre-feeding. *Behav. Process.* 52:43–48
- Powell RW. 1968. The effect of small sequential changes in fixed-ratio size upon the post-reinforcement pause. *J. Exp. Anal. Behav.* 11:589–93
- Powell RW. 1969. The effect of reinforcement magnitude upon responding under fixed-ratio schedules. *J. Exp. Anal. Behav.* 12:605–8
- Rachlin H. 1991. *Introduction to Modern Behaviorism*. New York: Freeman
- Rachlin H. 2000. *The Science of Self-Control*. Cambridge, MA: Harvard Univ. Press
- Rachlin H, Green L. 1972. Commitment, choice and self-control. *J. Exp. Anal. Behav.* 17:15–22
- Rachlin H, Green L, Kagel JH, Battalio RC. 1976. Economic demand theory and psychological studies of choice. In *The Psychology of Learning and Motivation*, ed. GH Bower, 10:129–54. New York: Academic
- Rachlin H, Laibson DI, eds. 1997. *The Matching Law: Papers in Psychology and Economics*. Cambridge, MA: Harvard Univ. Press
- Roberts S. 1981. Isolation of an internal clock. *J. Exp. Psychol.: Anim. Behav. Process.* 7:242–68
- Roberts S. 1983. Properties and function of an internal clock. In *Animal Cognition and Behavior*, ed. R Melgren, pp. 345–97. Amsterdam: North-Holland

- Royalty P, Williams B, Fantino E. 1987. Effects of delayed reinforcement in chain schedules. *J. Exp. Anal. Behav.* 47:41–56
- Schneider BA. 1969. A two-state analysis of fixed-interval responding in pigeons. *J. Exp. Anal. Behav.* 12:677–87
- Shimp CP. 1969. The concurrent reinforcement of two interresponse times: the relative frequency of an interresponse time equals its relative harmonic length. *J. Exp. Anal. Behav.* 1:403–11
- Shimp CP. 2001. Behavior as a social construction. *Behav. Process.* 54:11–32
- Shull RL. 1970. The response-reinforcement dependency in fixed-interval schedules of reinforcement. *J. Exp. Anal. Behav.* 14:55–60
- Shull RL. 1979. The postreinforcement pause: some implications for the correlational law of effect. In *Reinforcement and the Organization of Behavior*, ed. MD Zeiler, P Harzem, pp. 193–221. New York: Academic
- Shull RL, Gaynor ST, Grimes JA. 2001. Response rate viewed as engagement bouts: effects of relative reinforcement and schedule type. *J. Exp. Anal. Behav.* 75:247–74
- Shull RL, Pliskoff SS. 1967. Changeover delay and concurrent schedules: some effects on relative performance measures. *J. Exp. Anal. Behav.* 10:517–27
- Sidman M. 1960. *Tactics of Scientific Research: Evaluating Experimental Data in Psychology*. New York: Basic Books
- Skinner BF. 1937. Two types of conditioned reflex: a reply to Konorski and Miller. *J. Gen. Psychol.* 16:272–79
- Skinner BF. 1938. *The Behavior of Organisms*. New York: Appleton-Century
- Skinner BF. 1966. Operant behavior. In *Operant Behavior: Areas of Research and Application*, ed. WK Honig, pp. 12–32. New York: Appleton-Century-Crofts
- Skinner BF. 1986. Some thoughts about the future. *J. Exp. Anal. Behav.* 45:229–35
- Staddon JER. 1965. Some properties of spaced responding in pigeons. *J. Exp. Anal. Behav.* 8:19–27
- Staddon JER. 1968. Spaced responding and choice: a preliminary analysis. *J. Exp. Anal. Behav.* 11:669–82
- Staddon JER. 1970. Temporal effects of reinforcement: a negative “frustration” effect. *Learn. Motiv.* 1:227–47
- Staddon JER. 1972. Reinforcement omission on temporal go–no-go schedules. *J. Exp. Anal. Behav.* 18:223–29
- Staddon JER. 1974. Temporal control, attention and memory. *Psychol. Rev.* 81:375–91
- Staddon JER. 1977. On Herrnstein’s equation and related forms. *J. Exp. Anal. Behav.* 28:163–70
- Staddon JER. 1982. Behavioral competition, contrast, and matching. In *Quantitative Analyses of Behavior*, Vol. 2. *Quantitative Analyses of Operant Behavior: Matching and Maximizing Accounts*, ed. ML Commons, RJ Herrnstein, H Rachlin, pp. 243–61. Cambridge, MA: Ballinger. 5 Vols.
- Staddon JER. 2001a. *Adaptive Dynamics: The Theoretical Analysis of Behavior*. Cambridge, MA: MIT/Bradford. 423 pp.
- Staddon JER. 2001b. *The New Behaviorism: Mind, Mechanism and Society*. Philadelphia: Psychol. Press. 211 pp.
- Staddon JER, Chelaru IM, Higa JJ. 2002. A tuned-trace theory of interval-timing dynamics. *J. Exp. Anal. Behav.* 77:105–24
- Staddon JER, Higa JJ. 1999. Time and memory: towards a pacemaker-free theory of interval timing. *J. Exp. Anal. Behav.* 71:215–51
- Staddon JER, Innis NK. 1966a. An effect analogous to “frustration” on interval reinforcement schedules. *Psychon. Sci.* 4:287–88
- Staddon JER, Innis NK. 1966b. Preference for fixed vs. variable amounts of reward. *Psychon. Sci.* 4:193–94
- Staddon JER, Innis NK. 1969. Reinforcement omission on fixed-interval schedules. *J. Exp. Anal. Behav.* 12:689–700
- Staddon JER, Motheral S. 1978. On matching and maximizing in operant choice experiments. *Psychol. Rev.* 85:436–44
- Starr B, Staddon JER. 1974. Temporal control on fixed-interval schedules: signal properties

- of reinforcement and blackout. *J. Exp. Anal. Behav.* 22:535–45
- Stubbs A. 1968. The discrimination of stimulus duration by pigeons. *J. Exp. Anal. Behav.* 11:223–38
- Stubbs DA, Dreyfus LR, Fetterman JG, Boynton DM, Locklin N, Smith LD. 1994. Duration comparison: relative stimulus differences, stimulus age and stimulus predictiveness. *J. Exp. Anal. Behav.* 62:15–32
- Treisman M. 1963. Temporal discrimination and the indifference interval: implications for a model of the “internal clock.” *Psychol. Monogr.* 77(756): entire issue
- Wearden JH. 1985. The power law and Weber’s law in fixed-interval post-reinforcement pausing. *Q. J. Exp. Psychol. B* 37:191–211
- Williams BA. 1988. Reinforcement, choice, and response strength. In *Stevens’ Handbook of Experimental Psychology*, ed. RC Atkinson, RJ Herrnstein, G Lindzey, RD Luce, pp. 167–244. New York: Wiley. 2nd ed.
- Williams BA. 1994. Conditioned reinforcement: neglected or outmoded explanatory construct? *Psychon. Bull. Rev.* 1:457–75
- Williams BA. 1997. Conditioned reinforcement dynamics in three-link chained schedules. *J. Exp. Anal. Behav.* 67:145–59
- Williams BA, Royalty P. 1990. Conditioned reinforcement versus time to primary reinforcement in chain schedules. *J. Exp. Anal. Behav.* 53:381–93
- Wynne CDL, Staddon JER. 1988. Typical delay determines waiting time on periodic-food schedules: static and dynamic tests. *J. Exp. Anal. Behav.* 50:197–210
- Zeiler MD. 1977. Schedules of reinforcement: the controlling variables. See Honig & Staddon 1977, pp. 201–32
- Zeiler MD, Powell DG. 1994. Temporal control in fixed-interval schedules. *J. Exp. Anal. Behav.* 61:1–9
- Zuriff G. 1985. *Behaviorism: A Conceptual Reconstruction*. New York: Columbia Univ. Press