The mechanics of reinforcement

PETER R. KILLEEN and LEWIS A. BIZO
Arizona State University, Tempe, Arizona

Mathematical principles of reinforcement were developed in order to (1) account for the interaction of target responding and other behavior; (2) provide a simple graphical representation; (3) deal with measurement artifacts; and (4) permit a coherent transition from a statics to a dynamics of behavior. Rats and pigeons were trained to make a target response while general activity was measured with a stabilimeter. The course of behavioral change was represented as a trajectory through a two-dimensional behavior space. The trajectories rotated toward or away from the target dimension as the coupling between the target response and the incentive was varied. Higher rates of reinforcement expanded the trajectories; satiation and extinction contracted them. Concavity in some trajectories provided data for a dynamic generalization of the model.

Motivation, association, and response constraints are central phenomena in learning and performance. A recent theory of reinforcement—a *mechanics of reinforcement* (Killeen, 1992)—offers a formal description of these processes in terms of three principles. This paper summarizes the principles, elaborates them so as to generate clear predictions, and reports new data that bear on their evaluation. In particular, the action of the principles are represented in a *behavior space*, in which changes in the relation between target or operant behaviors and other activities appear as behavioral trajectories.

The Principles

Activation. The role of activation or arousal in conditioning is rarely disputed, even more rarely engaged (but see Bouton, 1993; Gibbon, 1995; Hogan, 1997; Silva & Pear, 1995; White & Milner, 1992). Skinner (1948) avoided it and attributed the hyperactivity of pigeons that were fed independently of their responding to the adventitious conditioning of responses that occurred just before reinforcement. But misattribution is an unlikely cause of the activity, as pigeons can readily report whether or not their behavior causes reinforcement (Killeen, 1978). Staddon and Simmelhag (1971) observed the putatively *superstitious* responding and found that it often occurs after, not before, reinforcement—just the wrong place for conditioning by contiguity. They extended Darwin's distinction between the evolutionary agencies of variation and selection to learned behaviors and proposed *principles of variation* that include all factors that originate behavior and *principles of selection* that include all factors that select among those responses made available by the former. *Activation*, or arousal, is our principle of variation, and *cou-*

pling is our principle of selection. This paper will analyze the effects on behavior when these two factors are independently manipulated.

The *first principle* states that *the delivery of incentives increases the activity of an organism*. Killeen (1975) reported levels of general activity in pigeons under a wide range of reinforcement rates. When corrected for ceilings on response rates, the activity levels were proportional to the rate of reinforcement (see Figure 1). Killeen, Hanson, and Osborne (1978) derived a model of incentive motivation that predicted the change in arousal levels as a function of changes in the rate of reinforcement. They fed pigeons once every day and measured the resulting activity, which averaged 360 responses/reinforcer. They showed that the activation cumulates according to an exponentially weighted moving average, whose output is the *arousal level*, A :

$$A = aR \quad (1)$$

where R is the rate of reinforcement. Equation 1 predicted Killeen's (1975) data (Figure 1), with $a = 360$.

Killeen (1998) extended the notion of arousal and its accumulation to other contexts, including classical and avoidance conditioning, in which the phenomena of pseudoconditioning and warm-up are its manifestations. The next step is to develop the theory, so as to deal generally with temporal constraints on responding.

Constraints on responding. Constraints are limitations—things organisms can't do no matter how powerful the motivation or how effective the conditioning. They are the complements of predispositions—things organisms do with seemingly little motivation or conditioning. Skinner (1938) represented this difference in sensitivities to reinforcement in his *extinction ratio*. Most generally, Seligman (1970) placed responses on a continuum of preparedness, ranging from contra-prepared through neutral to prepared. Here the focus is on temporal constraints—the increasing difficulty of making a response as a function of the ongoing rate of responding. The *second principle* can be succinctly stated: *Responses compete for expression*.

This research was supported by NSF Grants IBN-9408022 and BNS 9021562 and by NIMH Award MH01293 K05. We thank Geof White, John Wixted, and an anonymous reviewer for their generous help in configuring the manuscript. Address correspondence to P. R. Killeen, Department of Psychology, Arizona State University, Tempe, AZ 85287-1104 (e-mail: killeen@asu.edu).

Figure 1. Changes in the arousal level of pigeons, as inferred from asymptotes of response rates, when the pigeons are fed on periodic schedules at different rates. The behavior measured is general activity, and the arousal level is inferred from the scale factor of the general gamma distribution fit to the data. The data are from Killeen (1975).

If response competition is not taken into account, Equation 1 will overestimate the rate of responding: Responses, including responses of the same type, impede the emission of another response. If delta (δ) seconds are required to make a response, the rate of that response can obviously be no greater than $1/\delta$. It is less obvious how rates change as they approach that maximum. The model outlined here arrives at the same solution as the one presented by Killeen (1994a, 1994b) and by Staddon (1977). Let b equal the proportion of time occupied by responding; the proportion of time left available for additional responding is $1 - b$. The reduced opportunity to emit an additional response at higher rates of responding attenuates the force of motivation (A) by this factor:

$$b = A(1 - b). \quad (2)$$

Equation 2 states that the ability of response rates to change decreases proportionately as rates approach their ceiling (here, that is $b = 1.0$). It is as though we were trying to compress a gas; the closer b gets to its ceiling, the more force is necessary to hold it at the level b . Equation 2 thus corrects the time base for the time occupied by responding; it may be written as

$$b = \frac{A}{1 + A}, \quad (3)$$

which shows that response strength is a hyperbolic function of arousal level. Whereas Equation 1 gives the amount of behavior that is *evoked* by an incentive, Equation 3 gives the amount of behavior that is able to be *emitted*. In Appendix C, this equation is derived as the equilibrium solution to the equation of motion for behavior.

Coupling. The third principle is *coupling*, our principle of selection. Coupling occurs when an incentive occupies the same memory window as a response and is roughly

synonymous with the strengthening of an association. Association has always been the agent of choice for bringing about learning (Wasserman & Miller, 1997). It has two avatars: In classical, or Pavlovian, conditioning, the pairing of an arbitrary stimulus (conditioned stimulus, CS) with a biologically potent stimulus (unconditioned stimulus, US) changes the subject's response to the CS. In instrumental conditioning, the pairing of an arbitrary stimulus (S^D) and response (R) with a biologically potent stimulus (S^R) changes the subject's response to the S^D , in particular changing the frequency of R . Instrumental conditioning is a kind of compound conditioning, in which the subject must supply one of the elements.

Equations 1–3 represent undirected force—incentive motivation. The force is directed by the association of the incentive with particular stimuli and responses. It is the combination of excitation and association that constitutes *reinforcement*. According to Killeen's (1994a) principles, coupling is tightest (and, thus, association greatest) when an incentive occupies the same memory window as a response. Incentives that are not coupled to a particular stimulus or response arouse an animal but are unlikely to reinforce an instrumental response of interest to the experimenter (the *target response*). Instead, substantial adjunctive, superstitious, or frustrative behaviors may occur, which are often interpreted as hallmarks of arousal. But aroused organisms do not emit such responses in situations in which the contingencies of reinforcement focus the force of the incentive on the target response.

How can the nature and contents of the subject's memory be determined, in order to most effectively pair reinforcement with target responses? Operations such as priming (see, e.g., Brodbeck, 1997), elicitation (Locurto, Terrace, & Gibbon, 1981), "putting-through [the paces]," and shaping get responses—or their approximations—into memory. Their traces will decay with time or as new items (stimuli or responses) are added to memory. To determine the decay function, Killeen (1994a) reinforced pigeons' interresponse times (IRT's) according to rules that stipulated memory windows (or, equivalently, memory discount rates) of various sizes. The top panel of Figure 2 shows a sequence of IRTs, represented as columns, and a weighting function with a decay rate of 0.25 per item, representing the animal's window on the past. By this account, the animal's memory for its recent temporal pattern of responding is

$$\overline{\text{IRT}} = \sum_{n=1}^{\infty} w_n \text{IRT}_n,$$

with the weights w_n decaying exponentially into the past. It is clear that, if the experimenter discounts the past more heavily (e.g., by a reinforcement criterion that attends to only the most recent IRT), data that are salient to the organism will be left out. Conversely, if the experimenter includes too much in the window (e.g., by weighting the most recent 20 IRTs equally), insufficient weight will be given to the most recent responses. As one might expect,

Figure 2. Top panel: The columns depict a random sequence of interresponse times (IRTs), and the curve depicts an exponentially decaying discount function with a rate constant of $1/4$, the area under which is 1.0. The animal's characterization of its response rate is given by the sum of the products of the weighting function and the value of the IRT. Bottom panel: The slopes of the learning curves for 4 pigeons, plotted as a function of the experimenter's rate of discounting the recent history of responding (alpha). The tuning curve through the data is given by the theory; its maximum is at the imputed value of the subject's memory discount rate (λ). Recovered values range from 0.23 to 0.37 for individual subjects; the curve drawn through the average data has a peak at $\alpha = 0.275$. The data are from Killeen (1994a).

learning is fastest when the experimenter's criterion discounts the animal's past behavior at the same rate as does the animal. The experimental discount rates are represented by the values of the abscissae in the bottom panel of Figure 2. In some conditions, Killeen probabilistically presented reinforcement when the above weighted sum was in the top 20% of the animal's repertoire. In other cases, he presented reinforcement when the weighted sum was in the bottom 20% of the animal's repertoire. The changes in learning speed were measured as the slope of learning curves, both ascending (driving response rates faster) and

descending (driving response rates slower). These averaged slopes are displayed in the bottom panel of Figure 2 as a function of the experimenter's discount rate (α).

The correlation between these two memory windows—the experimenter's, characterized by α , and the subject's, characterized by λ —is called the *coupling coefficient*, ζ (zeta). As the coupling coefficient approaches 1.0, the learning curves approach their maximum. The coupling coefficient, and thus the rates of learning, are lower when the experimenter either underestimates the animal's discount rate ($\alpha < \lambda$) or overestimates it ($\alpha > \lambda$). This is shown in the bottom panel of Figure 2, where the curves give the predicted value of ζ when the animal's discount function is assumed to be exponential with $\lambda = 0.275$.

The theoretical coupling coefficient ζ tells us the proportion of memory that is typically filled by target responses at the moment of reinforcement. Because this class is strengthened as a function of its representation in memory, a resulting positive feedback loop drives behavior toward equilibrium. Different reinforcement schedules are characterized by different values of ζ . Knowing the subject's memory decay rate (λ) it is possible to calculate ζ for various experimental arrangements and schedules of reinforcement. For instance, under variable ratio schedules (i.e., where each response is reinforced with probability p),

$$\zeta = \lambda / (\lambda + p). \quad (4)$$

Under interval schedules, memory is less likely to be filled by target responses at the time of reinforcement, because pausing before the final response is reinforced to the same extent as is responding before the final response. The desultory character of the penultimate responses under interval schedules drives behavior toward an equilibrium rate of responding that is lower than that for ratio schedules.

Zeta is theoretically derived, and, although its value can be arbitrarily changed by the experimenter, it takes a while for behavior to follow suit. A term is needed that will refer to the proportion of target responses in an animal's repertoire at any point in time; we represent that with the letter C . When behavior has come to equilibrium, we predict that $C = \zeta$. We call C the *empirical coupling coefficient*. This distinction is not very important for this paper but will permit us to develop a dynamic model that predicts just how C will track arbitrary changes in ζ .

To predict the strength of a target response, b_T , multiply Equation 3 by C :

$$b_T = \frac{CA}{1 + A}. \quad (5)$$

To convert Equation 5 into a response *rate*, designated by a capital B_T , divide both sides by the time required for a response (δ_T). (If an organism spends half its time responding [$b_T = 0.5$], and if each response requires a quarter of a second [$\delta_T = 0.25$ sec], the animal is emitting

Figure 3. Average response rates of pigeons on a sequence of variable ratio (VR) schedules. The figure is reprinted from Bizo and Killeen (1997) with the permission of the American Psychological Association. The curve is drawn by Equation 6, using the coupling coefficient for VR schedules given by Equation 4 and the schedule feedback function for ratio schedules. The parameters are 0.36 sec for δ , 191 sec/reinf for a , and 0.9 for the memory decay rate λ .

$B_T = b_T/\delta_T = 0.5/0.25 = 2$ responses per second.) This carries us to the fundamental model of Killeen's (1994a, 1995) behavioral mechanics:

$$B_T = \frac{CA}{\delta_T(1+A)}, \quad \delta_T > 0, \quad (6)$$

where B_T is the target response rate. At high levels of arousal, response rates approach their ceiling (C/δ_T); as arousal or coupling falls to zero, response rates follow suit. Equation 6 incorporates motivation ($A = aR$), association (C), and constraints on responding (δ) in order to predict rates in a variety of situations. For example, inserting Equation 1 yields Herrnstein's (1979) hyperbola, which has provided a very robust account of behavior under schedules in which the rate of reinforcement is controlled. The parameters are interpreted differently, however, with a equalling the reciprocal of Herrnstein's R_O —the hypothesized reinforcement for other behaviors—and C/δ equalling his k —the total amount of behavior in the context. Under schedules in which the rate of reinforcement varies with the rate of responding, the schedule feedback function must be inserted in Equation 1. For ratio schedules, it is $R = B/N$, where N is the number of responses required for reinforcement. Then, Equation 6, along with ζ for ratio schedules (Equation 4), will draw the curve shown in Figure 3.

Changes in Arousal

The principles of reinforcement can be elaborated so as to account for changes in arousal within and between sessions. Such changes can be generated by a host of factors relating to the incentive, such as its type and size, and to the organism, such as its satiation and prior experience in the experimental context.

Warm-up. Many experiments show an increase in subjects' response rates through the early part of a session, even after many sessions of conditioning. Killeen (1998) analyzed this warm-up process and showed that an early model of it (Killeen et al., 1978) continues to provide a reasonable description. The model assumes that some fraction of the arousal that occurs during the session is conditioned to the experimental context and that the time course of this conditioned arousal follows that of initial acquisition, growing to asymptote during the first few minutes of the session. The mathematical representation of that process is a factor of A in the elaborated model.

Satiation. When an organism is deprived of food, the exigency of hunger grows slowly at first and more vigorously as deprivation continues. Conversely, arousal level is decreased by the satiation of the organism within experimental sessions. The expansion of a to include satiation is detailed in Killeen (1995). Changes in arousal as animals satiate can be predicted as a function of time in session, by including as parameters the effects of the amount or quality of an incentive (see, e.g., Weingarten, Duong, & Elston, 1996), the hunger drive, the initial deprivation level, and the threshold level of motivation that is required for responding to be initiated.

Figure 4 illustrates the manner in which the response rate changes within sessions as a function of different reinforcer types and amounts. The fitted functions are given by the generalization of Equation 6 described by Killeen (1995) and require a priori specification of the values of several factors, such as the amount of food consumed per reinforcer, the crop or stomach size of the animals, and several additional free parameters. Such a requirement for detail compromises simplicity.

The fresh approach offered in this paper finds a way around such complexification. Our main purpose is to show that changes in behavior that result from changes

Figure 4. Within-session changes in responding on a VI 60-sec schedule for 1 pigeon, given different amounts of food as a reinforcer. The data are from Bizo, Bogdanov, and Killeen (in press), and the curves are from an instantiation of the detailed model (Killeen, 1995).

Figure 5. An illustration of behavior represented in state space. The ordinates are the rates of emission of target responses, such as keypecks. The abscissae are the rates of emissions of all other responses. The three position vectors represent three different behavior states. This proportion of target responses is given by the coupling coefficient C , and the slope of the vectors by the ratio $C/(1 - C)$. The slope of vector 2 is greater than that of vector 1, indicating a larger value for C in that state, whereas the total amount of behavior (the length of the vector) is approximately the same. For vector 3, C has remained approximately constant, whereas the total amount of behavior has decreased, indicating decreased motivation.

in arousal or coupling can be simply understood in terms of movement in a behavior space. This relation is demonstrated by the data described in Experiments 1A and 1B, in which both operant responding and general activity were measured during acquisition. To motivate those simple experiments, we first describe the framework in which they will be placed.

Behavioral State Space

Figure 5 shows three locations in a *behavior space*, within which data represent *states* of behavior—rates of emitting the responses associated with each of the axes. Equation 5 may be written as a vector: $\mathbf{b} = \mathbf{CA}/(1 + A)$. This is a position vector. When we refer to *trajectories* in behavior space, we are describing the motion of the tip of this vector. Manipulations of arousal expand and contract the vectors; manipulations of coupling rotate the vectors. In Figure 5, State 2 has rotated counterclockwise from State 1, indicating that coupling to the target response (C) has increased, while arousal level has remained constant. This may be due to learning, to a shift to reinforcement schedules that have characteristically greater coupling (e.g., to ratio schedules, as opposed to interval schedules), or to changes in the probability of reinforcement with time or stimulus change. State 3 shows a proportional decrease in rates, which suggests a decrease in arousal, with little change in coupling. The next section provides the mathematical substrate for this representation.

Other investigators have used the concept *state* as an important theoretical variable, but they usually identify it with a broad class of responses. For instance, Anderson and Shettleworth (1977) noted that that description “is particularly appropriate in the present case because groups of activities rather than single activities are involved” (p. 47). In like manner, Timberlake (1993, 1994) identified a hierarchy of states, from the most general (systems and subsystems), through more specific (modes and modules), down to the elemental action patterns. Timberlake’s behavior system theory may provide a general framework for our more particular analysis—the semantics for our syntactics. Properties that are assumed by our principle may hold only within (or, in other cases, only across) various levels of his hierarchy.

Equations of motion. Equation 6 dealt with nontarget responses implicitly, assuming that their occurrence neither fostered nor interfered with the target response. In this section, the nontarget responses are made explicit. They are called *other* responses and indicated by the subscript O . We can generalize Equation 2 by expanding b as the target (T) plus other responses, ($b_T + b_O$):

$$b_T + b_O = A(1 - b_T - b_O). \quad (7)$$

Solving for target response strength yields

$$b_T = \frac{A}{1 + A} - b_O,$$

and it is then a short step to predict response rates:

$$B_T = \frac{A}{\delta_T(1 + A)} - \frac{\delta_O}{\delta_T} B_O; \quad \delta_T, \delta_O > 0. \quad (8)$$

When everything on the right-hand side is constant except other behavior (B_O), Equation 8 describes a straight line with a negative slope. It shows that, with C implicit and thus free to vary, target response rates will be complementary to other response rates. When such negative covariation is due to limits on the time available to emit responses, it is called a *restriction effect* (Allison, 1981, 1983, 1993; Staddon 1979, 1988).

Equation 8 represents both restriction and motivational effects; when subjects are relatively unmotivated (when A is small), they will respond well below their ceiling rates, as the intercepts of this negative diagonal are hyperbolic functions of arousal level. Rearranging Equation 8 shows that, when response rates are scaled by their durations (δ_i), the magnitude of the vector is a hyperbolic function of arousal level:

$$\delta_T B_T + \delta_O B_O = b_T + b_O = \frac{A}{1 + A}. \quad (9)$$

Not only do responses compete with one another, they may give rise to false alarms. For instance, the rat may stand on the lever in order to sniff the houselight, in which case the exploratory behavior is also recorded as a target response; conversely, a target response might, through its force, activate a stabilimeter, or it might give rise to a

second target response (because of *key-bounce*). These cases are analyzed in Appendix A, where it is shown that such artifacts will affect the values of the parameters in these equations but will not change their form.

The parameter C is not present in Equations 8 or 9, because its role is to account for the proportion of behavior that is focused on the operandum, but now that is accounted for explicitly by the introduction of B_O . Note that setting the rate of other behaviors equal to zero in Equation 8 does not return us to Equation 6. Setting the rate of other behaviors to zero is a definite assertion about the lack of competition: The proportion of target behavior (C) is $b_T/(b_O + b_T) = 1 - [(1 - A)/A]b_O$. Setting $b_O = 0$, thus, forces $C = 1$ —that is, perfect coupling of the incentive with the target response—and leaves us with a higher predicted target rate than does Equation 6 (which assumes that the fraction $(1 - C)$ of the force of reinforcement is spent on other behavior). If $b_O = 0$, the coupling coefficient C must be 1, and the two equations are consistent. Behaviorally, this corresponds to restricting the alternative behaviors, in which case there would be an increase in target responses. Equation 6 is a general statement, whereas Equation 8 adds the particulars of our knowledge of B_O .

Just as target and other responses are complementary and exhaustive, coupling of incentives to them is also complementary. Equation 6 can be written for B_O by substituting $1 - C$ for C . Then, divide one by the other to eliminate A and to obtain

$$B_T = \left(\frac{\delta_O}{\delta_T} \right) \left(\frac{C}{1 - C} \right) B_O; \quad 0 < C < 1, \delta_T, \delta_O > 0. \quad (10)$$

Equation 10 shows that, with A implicit and thus free to vary, response rates should fall on a straight line, with a slope proportional to $C/(1 - C)$. This is the ratio of target to nontarget responses typically in the memory window at the moment of reinforcement. Equation 10 is useful for the analysis of motivational effects, such as satiation, in which arousal is varying continuously through the session. When only coupling is varied, as in initial conditioning, or when there are systematic changes in the location of reinforcement over time, the slope of the vector should vary with it, and the resulting locus of target rates should fall along the line given by Equation 8 (see, e.g., Figure 5, States 1 and 2). When both vary, the behavior follows more complex trajectories.

In a study of the interaction of adjunctive behaviors, Reid and Dale (1985) found an increase in schedule-induced drinking by rats when the amount of food was increased between sessions, but within sessions they found the linear relation between drinking time and *head in feeder* time that is predicted by Equation 8. They concluded that "(1) Food presentation facilitates food-related behavior through elicitation and anticipation; and (2) food related behaviors are reciprocally, linearly related" (p. 147).

These are what we have called activation and constraint effects, respectively.

The coupling coefficient plays a dual role. When derived from the properties of reinforcement schedules, ζ (zeta) may be used to predict response rates at equilibrium (i.e., when the proportion of target behavior in the repertoire has stabilized). A ray with the slope $\zeta/(1 - \zeta)$ is the attractor of the behavioral trajectories—the ray along which behavior will settle asymptotically (Killeen, 1994a, Appendix D). But, in transition, that proportion will be changing, and C must be inferred from the actual locus of the trajectories in their state space. For ratio schedules, ζ can be specified a priori, but for interval schedules its exact value depends on the probability that the response occurring just before the reinforced response is also a target response, and that will evolve as learning progresses. This positive feedback loop is what causes learning curves to accelerate, but it is also responsible for amplifying small instabilities into unstable asymptotic performances.

If coupling is perfect (but see, e.g., Davison & Jenkins, 1985, and Appendix A), doubling the rate of reinforcement for a response, R_i , will approximately double the coupling that that response receives, relative to other responses. If C is written as R_T and $1 - C$ as R_O , Equation 10 is consistent with the matching relation. If coupling is less than perfect, some of those reinforcers will be misattributed, and subjects will undermatch.

Equation 9 presumes that all relevant behaviors are measured and that A captures the salient motivations. If there are systematic changes in the coupling of responses to other reinforcers throughout the session, either the data must be analyzed in a three-dimensional chart (where this theory, like Staddon's (1979), predicts that the locus of points will lie on a plane) or the trajectory will rotate (as happens in Figure 8). Behavior spaces will be used to represent the results of the following experiments, which are designed to test these models. It is predicted that arousal manipulations will primarily affect the magnitude of the vectors, whereas contingency manipulations will primarily affect their angle. The metric of these spaces is discussed in greater detail in Appendix B.

Motion Through Behavior Space

In Experiment 1A (for rats) and Experiment 1B (for pigeons), acquisition of a target response (leverpressing or keypecking) was recorded during the early stages of conditioning, under conditions in which habituation, satiation, and changes in conditioned arousal and coupling were expected to influence performance during the course of the experimental sessions. General activity was also recorded with a stabilimeter throughout the session. By plotting the rate of target behavior against the rate of activity, it was possible to derive a behavior space in which different behavior states reflect changes in arousal and coupling (as is shown in Figure 5).

Figure 6. Leverpresses and general activity as a function of trials. Response totals for each trial were averaged across rats; across blocks of 5 trials for the conditions habituation (diamonds, bottom panel), fixed time 30 sec (circles), and continuous reinforcement (triangles); and across blocks of 10 trials for the fixed interval 30-sec condition (squares). The top panel shows the data for the rats that initially experienced the delivery of a pellet every 30 sec, and the bottom panel shows the data for the rats that received 20 min pre-exposure to the experimental chamber prior to the first pellet delivery.

EXPERIMENT 1A Acquisition in Rats

In Experiment 1A, one group of rats was given free food in one long experimental session, followed by periods of continuous reinforcement and periodic reinforcement of leverpresses. A second group of rats experienced the same procedure, but after 20 min of prior exposure to the experimental chamber.

Method

Subjects. Eight experimentally naive female hooded rats (*Rattus norvegicus*, Long-Evans strain) were housed in groups of 4 with a reversed 12:12-h light:dark cycle, with dark beginning at 6 a.m. The rats were deprived to approximately 80% of their ad-lib weight by providing 6–12 g of Teklad rodent diet after all rats had completed the day's experimental session. The rats had free access to water.

Group NH comprised Rats 13, 14, 15, and 16, and Group NH comprised Rats 17, 18, 19, and 20.

Apparatus. The experimental chamber, measuring 27 cm high $\times$ 30 cm wide $\times$ 25 cm front, was lodged inside a Lehigh Valley sound-attenuating box. The chamber contained a 5-cm wide response lever, centered 4 cm from the side wall of the chamber and 5 cm from the chamber floor, which required 0.4 N force to activate a micro-switch. A centrally located pellet dispenser delivered 45 mg Noyes rat pellets. A house light was illuminated throughout the experimental session. The floor of the experimental chamber was connected to a Lafayette stabilimeter pickup (Model 86010, Lafayette Instruments; gain set to 6, activity set to rapid). Activity events recorded by the stabilimeter are called *movements*, although it is not assumed that they represent a modal action pattern (see Appendix A). A ventilation fan mounted in the side wall of the experimental chamber provided air and masking noise.

Procedure. Rats from Group NH were placed in the experimental chamber for one long session. On being placed in the chamber, they were given 25 pellets every 30 sec, independently of their behavior.

Figure 7. Leverpresses as a function of activity. The data, from Experiment 1 (see Figure 6), are plotted as an implicit function of time. The top panel shows the responses per trial for the rats that initially experienced the delivery of a pellet every 30 sec and the bottom panel shows the data from the rats that received 20 min preexposure to the experimental chamber prior to the first pellet delivery. The filled symbols indicate the origin of the trajectories.

This is called a fixed time 30-sec (FT 30) schedule. Next, a single leverpress was required for each pellet (continuous reinforcement, CRF). After 10 pellets, a period of 30 sec had to elapse since the previous reinforcer before a leverpress would be reinforced (a fixed interval 30-sec schedule, FI 30). This lasted until 205 pellets had been delivered. Rats from Group H were run on the same schedule, except that they were given 20 min preexposure to the experimental chamber before initiating the FT 30 schedule. Both groups were then given three additional sessions of FI 30. Each interval is called a *trial*; there were 100 trials per session in these last three sessions, each trial separated by a 1-sec intertrial interval (ITI), signaled by a darkening of the house light.

Results

The data in the top panel of Figure 6 are the number of leverpresses and movements during each trial for Group NH, averaged across subjects and blocks of intervals. The rats' leverpressing (filled symbols) increased with CRF and increased further under FI 30 until the rats had received about 120 pellets, whereafter it began to decrease. General activity (open symbols) started at a high level and fell quickly during the first 25 trials and more slowly thereafter.

The data in the bottom panel of Figure 6 are the number of leverpresses (filled symbols) and movements on each trial for Group H, averaged across subjects and blocks of trials. Notice that the initial decrease in movements is the same in this group, despite absence of food, and in Group NH. For both groups, leverpressing increased across the first 120 trials and decreased thereafter. General activity started high, decreased through the period of habituation, and stabilized at an asymptote of about 15 responses per trial.

Discussion

Initial exposure to the experimental chamber evokes a substantial amount of activity that decreases over the first 15 min in the chamber. The rats were observed to circle the chamber, sniff at the bottom of the walls, and rear in the corners during this time. We believe that the decrease in this activity is habituation of exploration (Forster, 1995). The decrease in this initial activity followed the same time course for both groups, although it occurred at a higher level for rats that were being reinforced at that time.

A new perspective is provided by the state space analysis. Figure 7 shows the leverpress rates from Figure 6, plotted as a function of the activity rates on the 1st day. Under FT 30—noncontingent delivery of reinforcement every 30 sec—the coupling coefficient is near zero, and the data lie along a ray from the origin that falls very close to the x-axis. This is consistent with Equation 10. Immediately on initiation of the FI contingency, the data rotate up to an intermediate position in the state space (from the last open circle to the filled square), as expected. The slope of the average vector through the FI 30 data is steeper for the animals that had been first habituated to the chamber (Group H, bottom panel) than for those who had not (Group NH, top panel): Group H's rates of leverpressing were slightly higher, and their rate of general activity significantly lower, than those for Group NH. The early habituation thus focused more of the rats' behavior on the lever. This may have happened because there was less general activity in the habituation group that was available to be adventitiously captured by reinforcement at the beginning of the conditioning phase, as is shown by the locus of the circles in the top and bottom panels.

Figure 8 presents the data from the 2nd and 4th days of conditioning, along with an ellipse representing the FI data from the first session (from Figure 7). Individual variability is portrayed by the axes of the ellipse, which equal the standard deviations of the rates. The data in Figure 8 are averaged over both groups. Notice how the trajectory of the data from the 4th day of FI 30 falls above and runs parallel to that from the 2nd. Those from the 3rd day (not shown) fall between these two data sets. This movement away from the origin indicates that the total amount of behavior is increasing, whereas the movement from right to left within sessions indicates an increase in the proportion of behavior that is dedicated to leverpressing. The increase in total behavior is interpreted as the

Figure 8. The top panel shows rats' leverpresses expressed as a function of activity on subsequent days. Response totals across a trial were averaged across all 8 rats and across blocks of 10 trials for the 2nd and 4th days' exposure to fixed interval (FI) 30. The filled symbols signify the start of the trajectory. The ellipse shows the locus of the states on the 1st day's exposure to FI 30, with the minor and major axes representing the standard deviations of rates on that day. The bottom panel shows a more traditional plot of leverpress and general activity totals as a function of trials.

conditioning of arousal to the context (Killeen, 1998), a process that requires multiple sessions—just as the extinction of that conditioned arousal requires multiple sessions, as is shown by the persistence of spontaneous recovery (Bouton, 1994; Mazur, 1996).

At the start of a session, rats vigorously explore the chamber, and this is reflected by the start-up transient (the leader entering from the right) in both Figures 7 and 8. Observations of the rats indicate that the initial transient was due to exploration and its habituation. A similar warm-up that detracted from avoidance responding during the first 15 min of a session was noted by Hineline (1978) and by others.

The time course of the effects of conditioning are clearer in the traditional graph and must be inferred from the distance between data points in the state space. Conversely, the covariation of leverpressing and other responses are manifest in the state space and must be inferred from the traditional graph.

In Experiment 1A, arousal level and coupling were not systematically manipulated, in order to monitor the changes in the factors that naturally accompany the early stages of acquisition. Coupling is varied more substantially in Experiment 2 and motivational level in Experiment 3. But first, the initial conditioning of pigeons' keypecking is analyzed for similar patterns.

EXPERIMENT 1B Acquisition in Pigeons

Method

Subjects. Eight experimentally naive common pigeons, *Columba livia*, were food deprived to $80\% \pm 10$ g of their ad-lib weights. The birds were housed in a room with a 12:12-h light:dark cycle of illumination, with dawn at 7 a.m. Supplementary mixed grain was provided at the end of each day, in order to maintain the bird's weights.

Apparatus. The Lehigh Valley experimental chamber was 29 cm high, 31 cm wide, and 35 cm front. The floor rested on springs and was connected to a Lafayette stabilimeter pickup (Model 86010, gain set to 4.5, activity set to slow). A response key requiring 0.22 N force for activation was centrally mounted on the interface panel. A central house light could illuminate the chamber. A magazine aperture provided a 3-sec access to milo grain, the reinforcer. A photocell mounted in the bottom of the magazine aperture could record when the pigeon placed its head into the magazine opening. White noise was provided by a speaker located behind the interface.

Procedure. The pigeons were hopper trained and then autoshaped to respond to a white key. Autoshaping consisted of response-independent presentation of a 15-sec white key light, followed by 4-sec of timed access to food. These trials were separated by 90-sec ITIs. Training was terminated after six keypecks, which always occurred within two 60-min sessions. The pigeons were then given a three-session exposure to an FI 30 schedule. Trials were separated

Figure 9. Behavior space for the data from Experiments 1A and 1B, averaged over sessions and subjects. The pigeon data are shown by the disks, with pecking rates increasing linearly from the first to the third session, and other activity showing a complementary decrease. The vectors for the rats rates rotate with the increase in coupling caused by the schedule change from fixed time to fixed interval (FI) and then expand with additional training on the FI schedule.

by a 20-sec ITI in blackout, with 60 trials to a session. The number of keypecks and stabilimeter activations were averaged across subjects and blocks of 6 trials.

Results and Discussion

There was a slight counterclockwise rotation of response rates within each of the three sessions and a more evident rotation from one session to the next: As the target response rate increased, the rate of other responses decreased proportionately (disks, Figure 9). These results are contrasted with those from Experiment 1A in Figure 9, which gives the session averages over subjects from all conditions, excluding the first 10 trials of each session for rats (squares).

The pigeon data are consistent with our expectation that conditioning will increase the value of the coupling coefficient, as they travel up the negative constraint line from the first through the third sessions. There is no expansion away from the origin during these conditions. The pigeons had all experienced hopper training and several sessions of autoshaping before these data were collected, which probably brought the excitatory conditioning of the experimental context to asymptote.

These data may be contrasted with those provided by the rats, which show a clear rotation of the vectors between FT and FI conditions but no change in coupling after the 1st day of conditioning. The 2nd and 3rd days' data show an expansion of the vector, with no further rotation, which we interpret as reflecting the cumulative conditioning of arousal to the experimental context.

EXPERIMENT 2

Testing Equation 8 by Varying Coupling

The purpose of this experiment is to demonstrate that changing the contingencies that define the target response will affect the coupling coefficient and, thereby, rotate the locus of the states, in accordance with Equation 8. The previous experiment studied concurrent target responses and other responses. It is more typical in the literature for concurrent responses to be two target responses of similar topography occurring on separate operanda, for which there is no crosstalk (Appendix A)—unless the animal can manage to reach both at the same time. Equation 8 predicts a simple linear relationship between responses, and this should hold, independent of their homogeneity over time. This section analyzes the adequacy of that linear relation in a context in which the coupling varies as a regular function of time throughout the session.

Method

Subjects. The subjects were 5 experimentally naive common pigeons, *Columba livia*, food deprived to 80% $\pm$ 10 g of their ad-lib weights.

Apparatus. The experimental chamber was a 31 cm high, 30 cm wide, and 35 cm front compartment made by Lehigh Valley. Four response keys 2.5 cm in diameter were mounted on the interface panel in the shape of a diamond: the left (red) key was 18 cm above the floor and 10 cm from the left wall; the right (green) key was

18 cm above the floor and 26 cm from the left wall; the top (red) key was 22 cm from the left wall and 22 cm above the floor; and the bottom (green) key was 22 cm from the left wall and 14 cm above the floor. A force of 0.27 N was required to activate the keys. A magazine aperture provided 2.2-sec access to milo grain. An infrared activity monitor (Coulbourn, Model E24-61) was mounted on the ceiling of the experimental chamber with its sensor 16 cm from the interface panel. White noise was provided by a speaker located behind the interface.

Procedure. The pigeons had experienced approximately 25 sessions of training on a procedure that began with the illumination of the left and right keys (Condition LR). During the first 25 sec of a 50-sec trial, responding to the left key was reinforced according to a variable interval (VI) 45-sec schedule and responding to the right key was in extinction; during the second half of the trial, responding to the right key was reinforced according to a VI 45-sec schedule and responding to the left key was in extinction. The Catania and Reynolds (1968) VI schedules in the two halves of the trial were independent. Each peck caused a 50-msec blink of the key that was pecked. Reinforcers scheduled for delivery in a component but not delivered were held over until the next trial. Each pigeon experienced 75 trials per session, each separated by a 10-sec ITI in blackout. Pigeons were given 60 sessions of this condition, but no data are reported from it here, as the activity monitor had not yet been connected.

In the next phase, the top and the bottom keys were used, rather than the left and right keys (Condition TB), with the pigeons given 26 sessions of retraining, with the final 5 of those sessions providing the data shown in the left column of Figure 10. Responding was recorded only during trials in which the subject did not receive a reinforcer. The pigeons were then returned to condition LR for 15 sessions, with the final two of those sessions providing the data shown in the right column of Figure 10.

Results and Discussion

The top panels of Figure 10 show the probability of responding on the *late* key as a function of time through the trial, with probability calculated as the relative number of responses on one key divided by the total number of responses. The middle panels show the probability of responding on the *early* key, plotted as a function of the probability of responding on the *late* key. For the left column, the *early* key was the top-center one and the *late* key was the bottom-center one, whereas, for the right panel, the *early* key was on the left and the *late* key on the right. The top panels show that the motivational force is coupled exclusively to the *early* key responses at the beginning of the trial and predominantly to the *late* key responses by the end of the trial. Operationally the coupling varies as a step-function of time halfway between those endpoints (25 sec); however, the temporal location of this point is uncertain for the animals, and its variability from one trial to the next gives rise to ogival psychometric functions, seen in the top panel (see, e.g., Bizo & White, 1994; Killeen, Fetterman, & Bizo, 1997).

Equation 8 tells us that the locus of the data in the middle panels should be a straight line decreasing from left to right. The regression lines are consistent with this prediction; in both cases, their slopes are about $-\frac{3}{4}$, showing a longer response time on the top and left keys than on the bottom and right keys. Under well-controlled conditions, the prediction of Equation 8 is sustained. There

Figure 10. Top panels: Probability of responding on the *late* key during each second of a 50-sec trial, in an experiment in which pecking was reinforced on the *early* key according to a variable interval 45-sec schedule during the first half of the trial and on the *late* key according to the same schedule during the second half. Left column of panels: *Early* and *late* keys are mounted vertically. Right column of panels: *Early* and *late* keys are mounted horizontally. Middle panels: Response rate on the *early* key, as a function of the response rate on the *late* key. Bottom panels: Activity as a function of responding on *early* and *late* keys. All of the straight lines are regressions consistent with Equation 8.

is some concavity in the right panel, possibly because of the prevalence of switching back and forth between keys during the middle of the trials; this topography has a longer response time, and thus, fewer change-over responses can be made than single-key responses in the same time.

Responding on the top key moved the pigeons forward out of the most sensitive area of the infrared pick-up, so that movement, as measured by this device, covaried

with the target responses, with top-key responses interfering with movement monitoring and bottom-key responses contributing to them. This is visible in the bottom left panel, which shows that, as rate on the top key (plotted on the x-axis) increased, activity decreased, whereas the reverse was the case for the right key. The proximity of the data to their regressions validates the extension of the model for crosstalk in the transducers (Appendix A).

Figure 11. Satiation trajectory for 6 pigeons, recorded in 10-min bins, showing the changes in pecking rate as a function of changes in activity rate. The filled triangle denotes the start of the session.

As a check on this analysis, the keys were moved back to a side-by-side configuration, and, with additional training, we obtained the data displayed in the right column of Figure 10. The bottom panel shows that this change reduced the extent to which one response was recorded as the other and returned the regressions to the horizontal, as expected in this more symmetric situation.

EXPERIMENT 3

Testing Equation 10 by Varying Motivation

The following experiment examined the trajectories resulting from a progressive satiation of pigeons on aperiodic schedules. As long as coupling is kept constant, Equation 10 predicts that the data will fall along a ray from the origin. Tandem ratio schedules were employed in order to increase the consistency of the coupling between incentive and behavior, and these were suffixed to variable interval schedules, which ensured relative constancy in the delivery of reinforcements over time.

Method

Subjects and Apparatus. Six new pigeons, with experience responding on a VI 60-sec schedule, were maintained at 80% of their ad-lib weight. The same apparatus, with the stabilimeter gain set to 4.0, was employed.

Procedure. The pigeons were trained on a Tandem (VI 4 min, FR 4) schedule for four sessions, during which reinforcement was 2-sec timed access to milo: In this schedule, after reinforcement was set up by the VI component, the fourth response would collect it. On the 5th day, the hopper duration was increased to 10 sec, the VI was reduced to 90 sec, and the session was extended to 220 min. The VI schedules were 20 interval Catania and Reynolds (1968) distributions.

Results and Discussion

Figure 11 shows the response rate in 10-min bins throughout the course of the single day of satiation. Under these contingencies the trajectory of keypecking as a

function of general activity follows Equation 10, which predicts a straight-line decrease to the origin.

These results—a symmetric decrease in the strength of two operants as motivation is decreased—replicates the proportional decrease in concurrent time allocation under extinction conditions, found by Buckner, Green, and Myerson (1993), and the proportional decrease in keypecking under extinction from ratio schedules, found by Myerson and Hale (1988). Proportional declines were predicted for extinction by the kinetic model (Myerson & Miezin, 1980), because, in that model, changes in preference were proportional to rates of reinforcement (see Appendix C); when those go to zero in extinction, there can be no further change in preference.

These predictions require that the motivation for other responses is the same as that for the target response—in this case, arousal because of the delivery of food to a hungry organism. To the extent that there are other reasons for moving around the chamber, satiation and extinction trajectories will intercept the x-axis to the right of the origin.

EXPERIMENT 4

Testing Equation 10 by Varying Rate of Reinforcement

The arousal of the organism is a product of the motivational level (a) and the rate of reinforcement (R ; cf. Equation 1). In Experiment 4, R was manipulated to assay the effect on response trajectories. As in Experiment 3, tandem VI-FR schedules were employed: After the VI timed out, an additional four responses were required for reinforcement. Very long VI values were employed, in order to cause response rates to vary over a substantial range (see Catania & Reynolds, 1968).

Method

Subjects and Apparatus. Six pigeons with extensive and varied experimental histories were used. The apparatus was the same as that used in Experiment 2. The stabilimeter gain was set to 1.5.

Figure 12. Transitions to the variable interval (VI) 4-min schedules (up-triangles) and the VI 64-min schedules (down-triangles). The large triangles denote the terminus of each trajectory.

Procedure. The pigeons were trained on a 15 interval constant-probability VI 1-min schedule for 20 sessions. They were then exposed to a Tandem (VI 4 min, FR 4) schedule for 6 sessions, during which reinforcement was a 2-sec timed access to milo. After reinforcement was set up by the VI scheduler, the fourth response would collect it. The VI was then decreased to a VI 64 min (3 sessions) and then set back to a VI 4 min (3 sessions).

Results

Figure 12 shows the rates of keypecking and movement, averaged over subjects. The data for the first (transitional) session in a new condition are reported in 10-min bins (small triangles), whereas the rates for the last session at each condition are averaged over the session (large triangles). The down-triangles show the slowing of rate on first exposure to the lean VI schedule, ending at the large triangle at coordinates (13,36), which gives the last day's performance on the VI 64 min. The up-triangles show the speeding of rate upon return to the rich VI 4 min schedule, ending with the large triangle at (49,76), the last day at the VI 4 min.

Discussion

The trajectories of responding as the rate of reinforcement is varied over a 16-fold range are generally consistent with our prediction of a proportional relationship. This indicates that, as expected, changes in rates of reinforcement primarily affect motivation. There is a slight concavity—a hysteresis effect—visible in the trajectories, however. This could be due to a differential susceptibility to conditioning and to extinction of different responses. Seligman (1970) called such a difference one of *preparedness*. Figure 12 shows that keypecking extinguished more slowly and recovered more quickly than other behaviors—it was a more *prepared* response. This hysteresis, although subtle, provides the text for a more complete dynamics of action, described in Appendix C.

GENERAL DISCUSSION

The experiments reported here demonstrate a new method of analyzing behavior that (1) simplifies a detailed model of equilibrium performance and (2) paves the way for a dynamic analysis. Whether other behaviors are measured with symmetric operanda (two keys) or with asymmetric ones (stabilimeters and levers), the basic predictions retain their functional forms: motivational operations expand or contract the behavioral vectors (Experiments 3 and 4), whereas changes in association (Experiment 2) rotate them. Redundancy and crosstalk in the measurement of these responses do not change the predictions. Extraneous sources of reinforcement, such as exploration of the chamber, are associated with characteristic trajectories.

Niels Bohr noted the complementary relation between simplicity and specificity (French & Kennedy, 1985), a trade-off that confronts any theoretician. The models that describe changes in behavior as a function of changes in motivation are complicated (Killeen, 1995, 1998). We sought a simpler representation, both in the service of

communication and in the interests of establishing a context in which a mechanics might be more readily developed. By plotting one behavior against others parametrically, it was possible to avoid the specification of numerous constants and parameters. The realization of that simplicity was contingent on the demonstration, made in Appendix A, that redundant and misattributed responses would not change the character of the representation. It is also contingent on the metric of behavior space. Responses generally come in units, called action pat dix B). If this were not the case, the lumping of other responses on a single axis would not be possible, because, as the composition of the other responses changed, their extension on that axis would not be conserved (they would have to be combined with the Pythagorean rule).

The behavior space representation provides a convenient diagnostic of the nature of changes in behavior: Figure 9 showed that, for rats, changes in coupling/association were complete within one session, whereas arousal levels continued to increase over several sessions. Conversely, arousal levels were invariant for pigeons, whereas the proportion of their behavior that was emitted as the target response continued to grow over sessions. These differences were probably due to the training regimens rather than to species differences, but that remains an empirical question.

In most experiments, both coupling and arousal will change simultaneously, causing behavioral trajectories to describe arbitrary figures through their space. But where two of the three factors—arousal, constraint, and coupling—can be held constant, then Equations 8 and 10 predict orderly changes in behavior, such as those seen in Figures 9, 10, and 11. Figure 12 revealed small but systematic deviations from the predictions. Although these could be due to concurrent changes in coupling and arousal, they may also be due to different momenta for the target and other responses. These alternate explanations may be tested, as is shown in Appendix C.

Tracking changes in behavior over minutes rather than over days limits the data base for each point in those graphs. Sampling error was decreased by increasing the number of subjects whose data contributed to each point. Similar behavior spaces for individual subjects show significantly greater variability. The complementary relation between simplicity and specificity holds for data, as well as for models.

A formal system such as this bears comparison with Hull's (1943; Wearden, 1994) and, in some perspectives, shares guilt by association. Hull believed his attempts at quantification would be the most enduring aspect of his work (Amsel & Rashotte, 1984), but, alas, interest focused on qualitative tests, and theory development became broad rather than deep. Behavioral mechanics has available three crucial resources that Hull lacked—an extensive empirical literature that shapes the assumptions, theoretical developments that guide the models, and efficient computers for testing and rejecting alternate models. Our future endeavor is to replicate the hysteresis suggested

in Figure 12 (and other data sets) and to describe that process with models that complement those of Myerson and Miezen (1980) and of Nevin (1992). This entails a dynamical representation (outlined in Appendix C), of which Equation 5 is the equilibrium solution.

Much is left unexamined with this type of approach, as is noted by Timberlake and Silva (1994) and by Zeiler (1996), among others. Nonetheless, Hull's hypothetico-deductive approach, developed patiently with renewed attention to quantification, has much yet to tell us about behavior and its modification.

REFERENCES

- ALLISON, J. (1981). Economics and operant conditioning. In P. Harzem & M. D. Zeiler (Eds.), *Advances in the analysis of behavior: Vol. 2, Predictability, correlation, and contiguity* (pp. 321-353). Chichester, U.K.: Wiley.
- ALLISON, J. (1983). *Behavioral economics*. New York: Praeger.
- ALLISON, J. (1993). Response deprivation, reinforcement, and economics. *Journal of the Experimental Analysis of Behavior*, **60**, 129-140.
- AMSEL, A., & RASHOTTE, M. E. (1984). *Mechanisms of adaptive behavior: Clark L. Hull's theoretical papers, with commentary*. New York: Columbia University Press.
- ANDERSON, M. C., & SHETTLEWORTH, S. J. (1977). Behavioral adaptation to fixed-interval and fixed-time food delivery in golden hamsters. *Journal of the Experimental Analysis of Behavior*, **25**, 33-49.
- ARABIE, P. (1991). Was Euclid an unnecessarily sophisticated psychologist? *Psychometrika*, **56**, 567-587.
- BAUM, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, **22**, 231-242.
- BIZO, L. A., BOGDANOV, S. V., & KILLEEN, P. R. (in press). Satiation causes within-session decreases in instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes*.
- BIZO, L. A., & WHITE, K. G. (1994). Pacemaker rate in the behavioral theory of timing. *Journal of Experimental Psychology: Animal Behavior Processes*, **20**, 308-321.
- BOUTON, M. E. (1993). Context, time, and memory retrieval in extinction and the interference paradigms of Pavlovian learning. *Psychological Bulletin*, **114**, 80-99.
- BOUTON, M. E. (1994). Conditioning, remembering, forgetting. *Journal of Experimental Psychology: Animal Behavior Processes*, **20**, 219-231.
- BRODBECK, D. R. (1997). Picture fragment completion: Priming in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes*, **23**, 461-468.
- BUCKNER, R. L., GREEN, L., & MYERSON, J. (1993). Short-term and long-term effects of reinforcers on choice. *Journal of the Experimental Analysis of Behavior*, **59**, 293-307.
- CATANIA, A. C., & REYNOLDS, G. S. (1968). A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, **11**, 327-383.
- DAVISON, M., & JENKINS, P. E. (1985). Stimulus discriminability, contingency discriminability, and schedule performance. *Animal Learning & Behavior*, **13**, 77-84.
- DAVISON, M., & JONES, B. M. (1995). A quantitative analysis of extreme choice. *Journal of the Experimental Analysis of Behavior*, **64**, 147-162.
- FORSTER, F. C. (1995). Prior and concurrent learning increases object exploration in the rat (*Rattus-norvegicus*). *Australian Journal of Psychology*, **47**, 147-151.
- FRENCH, A. P., & KENNEDY, P. J. (Eds.) (1985). *Niels Bohr: A centenary volume*. Cambridge, MA: Harvard University Press.
- GALLO, A., DUCHATELLE, E., ELKHESSAIMI, A., LEPAPE, G., & DESPORTES, J. P. (1995). Topographic analysis of the rat's bar behaviour in the Skinner box. *Behavioural Processes*, **33**, 319-327.
- GIBBON, J. (1995). Dynamics of time matching: Arousal makes better seem worse. *Psychonomic Bulletin & Review*, **2**, 208-215.
- HARPER, D. N., & MCLEAN, A. P. (1992). Resistance to change and the law of effect. *Journal of the Experimental Analysis of Behavior*, **57**, 317-337.
- HERRNSTEIN, R. J. (1979). Derivatives of matching. *Psychological Review*, **86**, 486-495.
- HINELINE, P. N. (1978). Warmup in avoidance as a function of time since prior training. *Journal of the Experimental Analysis of Behavior*, **29**, 87-103.
- HOGAN, J. A. (1997). Energy models of motivation: A reconsideration. *Applied Animal Behavior Science*, **53**, 89-105.
- HULL, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.
- KILLEEN, P. R. (1975). On the temporal control of behavior. *Psychological Review*, **82**, 89-115.
- KILLEEN, P. R. (1978). Superstition: A matter of bias, not detectability. *Science*, **199**, 88-90.
- KILLEEN, P. R. (1992). Mechanics of the animate. *Journal of the Experimental Analysis of Behavior*, **57**, 429-463.
- KILLEEN, P. R. (1994a). Mathematical principles of reinforcement. *Behavioral & Brain Sciences*, **17**, 105-135.
- KILLEEN, P. R. (1994b). Rats, responses and reinforcers: Using a little psychology on our subjects. *Behavioral & Brain Sciences*, **17**, 157-172.
- KILLEEN, P. R. (1995). Economics, ecologies, and mechanics: The dynamics of responding under conditions of varying motivation. *Journal of the Experimental Analysis of Behavior*, **64**, 405-431.
- KILLEEN, P. R. (1998). The first principle of reinforcement. In C. Wynne & J. E. R. Staddon (Eds.), *Models of action* (pp. 127-156). Mahwah, NJ: Erlbaum.
- KILLEEN, P. R., FETTERMAN, J. G., & BIZO, L. A. (1997). Time's causes. In C. M. Bradshaw & E. Szabadi (Eds.), *Time and behaviour: Psychological and neurobiological analyses* (pp. 79-131). Amsterdam: Elsevier.
- KILLEEN, P. R., HANSON, S. J., & OSBORNE, S. R. (1978). Arousal: Its genesis and manifestation as response rate. *Psychological Review*, **85**, 571-581.
- LOCURTO, C. M., TERRACE, H. S., & GIBBON, J. (Eds.) (1981). *Auto-shaping and conditioning theory*. New York: Academic Press.
- MAZUR, J. E. (1996). Past experience, recency, and spontaneous recovery in choice behavior. *Animal Learning & Behavior*, **24**, 1-10.
- MYERSON, J., & HALE, S. (1988). Choice in transition: A comparison of melioration and the kinetic model. *Journal of the Experimental Analysis of Behavior*, **49**, 291-302.
- MYERSON, J., & MIEZIN, F. M. (1980). The kinetics of choice: An operant systems analysis. *Psychological Review*, **87**, 160-174.
- NEVIN, J. A. (1988). Behavioral momentum and the partial reinforcement effect. *Psychological Bulletin*, **103**, 44-56.
- NEVIN, J. A. (1992). An integrative model for the study of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, **57**, 301-316.
- NEVIN, J. A., TOTA, M., TORQUATO, R. D., & SHULL, R. L. (1990). Alternative reinforcement increases resistance to change: Pavlovian or operant contingencies? *Journal of the Experimental Analysis of Behavior*, **53**, 359-379.
- REID, A. R., & DALE, R. H. I. (1985). Dynamic effects of food magnitude on interim-terminal interaction. *Journal of the Experimental Analysis of Behavior*, **39**, 135-148.
- SCHWENDIMAN, J. W. (1997, May). *Extreme choice: Contingency-discriminability or generalized matching?* Poster presented at the Society for Quantitative Analysis of Behavior, Chicago.
- SELIGMAN, M. E. P. (1970). On the generality of the laws of learning. *Psychological Review*, **77**, 406-418.
- SILVA, F. J., & PEAR, J. J. (1995). Stereotypy of spatial movements during noncontingent and contingent reinforcement. *Animal Learning & Behavior*, **23**, 245-255.
- SKINNER, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- SKINNER, B. F. (1948). Superstition in the pigeon. *Journal of Experimental Psychology*, **38**, 168-172.
- STADDON, J. E. R. (1977). On Herrnstein's equation and related forms. *Journal of the Experimental Analysis of Behavior*, **28**, 163-170.

- STADDON, J. E. R. (1979). Operant behavior as adaptation to constraint. *Journal of Experimental Psychology: General*, **108**, 48-67.
- STADDON, J. E. R. (1988). Quasi-dynamic choice models: Melioration and ratio invariance. *Journal of the Experimental Analysis of Behavior*, **49**, 383-405.
- STADDON, J. E. R., & SIMMELHAG, V. (1971). The "superstition" experiment: A re-examination of its implications for the principles of adaptive behavior. *Psychological Review*, **78**, 3-43.
- TEN CATE, C., & BALLINTJN, M. (1996). Dove coos and flashed lights: Interruptibility of "song" in a nonsongbird. *Journal of Comparative Psychology*, **110**, 267-275.
- TIMBERLAKE, W. (1993). Behavior systems and reinforcement: An integrative approach. *Journal of the Experimental Analysis of Behavior*, **60**, 105-128.
- TIMBERLAKE, W. (1994). Behavior systems, associationism, and Pavlovian conditioning. *Psychonomic Bulletin & Review*, **1**, 405-420.
- TIMBERLAKE, W., & SILVA, F. J. (1994). Observation of behavior, inference of function, and the study of learning. *Psychonomic Bulletin & Review*, **1**, 73-88.
- WASSERMAN, E. A., & MILLER, R. R. (1997). *What's elementary about associative learning?* *Annual Review of Psychology* (Vol. 48, pp. 573-607). Palo Alto, CA: Annual Reviews.
- WEARDEN, J. H. (1994). Fifty years on: The new "principles of behavior"? *Behavioral & Brain Sciences*, **17**, 155.
- WEINGARTEN, H. P., DUONG, A., & ELSTON, D. (1996). Interpretation of sham feeding data: Curve-shift studies. *American Journal of Physiology: Regulatory Integrative & Comparative Physiology*, **40**, R1009-R1016.
- WHITE, N. F., & MILNER, P. M. (1992). The psychobiology of reinforcers. *Annual Review of Psychology*, **43**, 443-471.
- ZEILER, M. D. (1996). What behaviors do. *Behavioral & Brain Sciences*, **19**, 549-550.

APPENDIX A Detecting Behavior

Response transducers act as filters, differentially emphasizing actions that are components of different responses. The settings on stabilimeters are arbitrary: Each report from them does not indicate that a single behaviorally meaningful act—an action pattern—has occurred. The same is true for leverpressing (Gallo, Duchatelle, Elkhessaimi, Lepape, & Desportes, 1995). At high sensitivity settings of the transducers, each action pattern will give rise to multiple (*redundant*) signals, whereas, at low settings, some responses will be missed. In addition, the stabilimeter may be activated by the target response, and the operandum may pick up other responses (*crosstalk*). What do these problems of sensitivity and selectivity in our transducers do to the simple relationships predicted by Equations 8 and 10? To answer this question, a continuous version of signal detection theory (SDT) may be used to analyze their performance. The rate at which our instruments report target behavior is

$$\hat{B}_T = \gamma_{\cdot T \cdot | T} B_T + \gamma_{\cdot T \cdot | O} B_O, \quad (A1)$$

where $\hat{B}_T$ is the measured rate, $\gamma_{\cdot T \cdot | T}$ is the average number of target responses reported for each that occurred (the *gain* for target responses), and $\gamma_{\cdot T \cdot | O}$ is the average number of target responses reported for each *other* response that occurred (the *gain* for crosstalk to the target channel). Similarly,

$$\hat{B}_O = \gamma_{\cdot O \cdot | O} B_O + \gamma_{\cdot O \cdot | T} B_T. \quad (A2)$$

It is obvious that, in an ideal transducer system, the gain for veridical reporting will be 1.0— $\gamma_{\cdot T \cdot | T} = \gamma_{\cdot O \cdot | O} = 1$ —whereas the gain for crosstalk will be 0— $\gamma_{\cdot T \cdot | O} = \gamma_{\cdot O \cdot | T} = 0$. In this ideal case, the measured responses are in perfect correspondence with the emitted responses, and $\hat{B}_i = B_i$.

For brief epochs τ in which only one response is likely, this treatment relates to the classic matrix of SDT, as shown in Table A1, where it is assumed that the number of counts is a random variable $G_{\cdot i \cdot | j}$, with mean given by the gain parameters $\gamma_{\cdot i \cdot | j}$, and $F(\gamma)$ is a distribution, such as the Poisson.

Equations A1 and A2 may be solved for the true response rate, given the reported rates and the gain parameters:

$$B_T = \left(\hat{B}_T - \frac{\gamma_{\cdot O \cdot | T} \hat{B}_O}{\gamma_{\cdot O \cdot | O}} \right) \left(\gamma_{\cdot T \cdot | T} - \frac{\gamma_{\cdot O \cdot | T} \gamma_{\cdot T \cdot | O}}{\gamma_{\cdot O \cdot | O}} \right)^{-1},$$

$$0 \leq \gamma_{\cdot i \cdot | j}, 0 < \gamma_{\cdot i \cdot | i}. \quad (A3)$$

A symmetric equation, A4, may be written for B_O . Note that, when the crosstalk is zero, $B_i = \hat{B}_i / \gamma_{\cdot i \cdot | i}$: the true rate of target (and other) responses is the reported rate divided by the average number of reports that each response occasions.

Substitute the Equations A3 and A4 into Equation 8 in the main text and, after some algebra, find that:

$$\hat{B}_T = k \left[\frac{k'A}{1+A} - k''\hat{B}_O \right], \quad (A5)$$

with

$$k = [\delta_T \gamma_{\cdot O \cdot | O} - \delta_O \gamma_{\cdot T \cdot | O}]^{-1},$$

$$k' = \gamma_{\cdot T \cdot | T} \gamma_{\cdot O \cdot | O} - \gamma_{\cdot O \cdot | T} \gamma_{\cdot T \cdot | O},$$

$$k'' = \delta_O \gamma_{\cdot T \cdot | T} - \delta_T \gamma_{\cdot O \cdot | T},$$

and where

$$0 \leq \gamma_{\cdot i \cdot | j}, 0 < \gamma_{\cdot i \cdot | i}.$$

The important point about Equation A5 is that it is essentially the same as Equation 5: the only difference is the presence of combinations of constants multiplying the terms. All of the implications drawn from Equation 5 hold when there is noise and redundancy in our measurement of behavior.

Most experimenters will not have independent measures of the gain parameters. In this case, the recovered parameters (C , A , and δ_i) will not precisely reflect their actual values. As long as those parameters are constant, however, the rate of target responding will be proportional to the rate of other responses, and the plot of one against the other will describe a ray emanating from the origin. If sources of reinforcement for target or other responses change during the course of a session, the ray should rotate as the effective value of C changes. If alternate sources of reinforcement exist, such as general exploration, the x -intercept of the trajectory will be to the right of the origin.

This development assumes that the measurement of behavioral allocation C is debased by the instrumentation error represented by the gain parameters. But there is also slippage between the experimenter's allocation of reinforcers and the animal's perception of them. For instance, the multiple pecks at 3–4 Hz that characterize pigeon's keypecking may indicate that some of the recorded responses are a part of a reflexive tattoo, rather than

Table A1
The Types and Probabilities of Reports
Conditioned on Their Origin

Report	Origin	
	Target Response	Other Response
Target	True Positive: $F(\gamma_{\cdot T \cdot T})$	False Positive: $F(\gamma_{\cdot T \cdot O})$
Other	False Negative: $F(\gamma_{\cdot O \cdot T})$	True Negative: $F(\gamma_{\cdot O \cdot O})$

Figure A1. A close look at a behavioral trajectory. Because it is difficult or impossible for animals to emit fractions of an action pattern, behavior space is quantal in nature.

being unique action patterns. In that case, $\gamma_{T,T} > 1$. For another instance, an experimenter may schedule 75% of the available reinforcers for responses on Key 1, but, without a *change-over delay*, the animals may switch rapidly back and forth, so that some of the reinforcers scheduled for Key 1 also strengthen responding on Key 2. A precise model of coupling will permit us to predict this, by assigning appropriate theoretical coupling coefficients (represented as ζ) to the various sources of incentives. But a simpler approach notes that the subject's *misallocation* of behavior is no different in its effects than the instruments': the *response* that is reinforced is recorded by one operandum but may actually involve behavior on both operanda (e.g., the contingencies may reinforce switching). In the simple case in which there is no *key bounce* ($\gamma_{T,T} = 1$) and the response times for the two keys are equal, Equation 10 may be written as

$$\hat{B}_1 = \hat{B}_2 \left[\frac{\zeta + (1-\zeta)\gamma_{1/2}}{(1-\zeta) + \zeta\gamma_{2/1}} \right], \quad (A6)$$

where ζ is the proportion of reinforcers scheduled for Key 1. Davison and Jones (1995; Baum, 1974) also assume that some proportion of the reinforcers delivered for one response can strengthen other responses and, from that assumption, develop a model of concurrent choice that is superior to the generalized matching account (cf. Schwendiman, 1997). Their model is equivalent to Equation A6. If the crosstalk parameters are zero, it reduces to simple matching.

It is routine to assume that each peck of a key or press of the lever constitutes a unit of behavior, but this may not be the case. The utility of the present analysis is that it permits us to analyze data in a principled fashion, even when the operant response may be more or less or other than that which activates the transducer.

Summary

If the crosstalk gains are greater than zero, or the *sametalk* gains are other than 1.0, the slopes of Equations 8 and 10 will be different from the ratio of response durations (Equation 8) or from that ratio times the ratio of the coupling coefficients (Equation 10), but the basic linearity will be retained. Although the model is developed here for instrument/interface errors, it can serve as a model for subject/experimenter classification errors, in which case it reduces to a standard model of choice (Equation A6).

APPENDIX B The Metric of Behavior-Space

Because measurable responses come in minimal quantal sizes (see, e.g., Ten Cate & Ballintijn, 1996), the metric of behavioral state space is *city block*, not Euclidean: Organisms cannot get from point 1 to point 2 by alternating fractions of a target response with fractions of other behaviors. In this kind of metric space, the hypotenuse is always a staircase, as is shown in Figure A1. No matter how small the risers and treads, the distance along this path will always equal the sum of the coordinates. Note that the number of segments (including the first tread) from the origin to $B_T = 4$ is 8. It is this inability to fly as the crow that gives this " l_1 " metric the informal name *city block*.

This quantal nature is not changed if responses are divided by time to graph a rate, nor if they are multiplied by the duration of the responses to normalize the graph. The coordinates may then appear as fractions, but the inevitable discontinuity in the trajectories will remain. When data are aggregated over the course of a session, the wrinkles may be smaller than the resolution of the printer that draws the trajectories, but macroscopic implications remain: The iso-arousal trajectory followed when coupling is varied is a straight line (Equation 8 and Figures 5 and 10), not a segment of an ellipse (or of a circle, in normalized coordinates).

Given the importance of metrics other than the Euclidean in the analysis of perceptual similarity spaces, Arabie (1991, p. 581) observed that "it is too late to fault Euclid, but one can question the sophistication of behavioral scientists who have been too willing to assume the ubiquity of Euclidean geometry."

APPENDIX C The Dynamics of Behavior

The dynamic models in this paper start with models of behavior after it has come to equilibrium (statics) and attempt to preserve them as equilibrium solutions of dynamic models that are developed for behavior in transition. The models in the body of this paper are a special, simple case of dynamical systems in which changes in responses during transitions are perfectly correlated. But, if some responses are more sensitive to variations in motivation or contingency, a more general treatment is required. That is outlined here. In the process, parameters are introduced (e.g., mass and friction), whose psychological correlates will be discussed later. Despite the seeming complexity of the development, it eventuates that the trajectories in state space are governed by a single (composite) parameter—the relative masses of the responses—and reduces to the treatment given in the body of the text when those masses are equal.

Consider first a simple system in which a force A works on a body of mass m , whose velocity is b . The force A may be opposed by a frictional force $-\mu b$ (μ is the *dynamic coefficient of friction*, here denoting the resistance of a response rate to increase). Then, with the net force $F = A - \mu b$ and from $F = m(db/dt)$, it follows that

$$A - \mu b = m \frac{db}{dt}. \quad (A7)$$

At equilibrium, the differential on the right-hand side (rhs) is zero, because response rates have stopped changing, so that $b = A/\mu$: Response rate is proportional to arousal level (which gives back Equation 2, with $\mu = 1/(1 - b)$; If the force is changed to a new value A' , integration of A7 shows that response rates will follow a smooth approach to their new equilibrium:

Figure A2. The data of Figure 12, reanalyzed under the assumption that keypecking has half the behavioral mass of typical other behaviors. This ratio gives the exponent of the curves; for equal masses, the curves would converge into a straight line.

$$b' = \frac{1}{\mu} \left[A' + (A - A')e^{-\mu t/m} \right]. \quad (\text{A8})$$

For example, in extinction, $A' = 0$, and Equation A8 is an exponential decay function.

Relation to Nevin's Theory of Behavioral Momentum

The rhs of A7 is what both Nevin (e.g., 1988, 1992; Nevin, Tota, Torquato, & Shull, 1990) and we call *behavioral momentum*. The present model of it agrees with Nevin's theory and data in many ways and diverges in some other ways. Nevin's modulus operandi has been to characterize resistance to change by comparing the proportional change in two behaviors as A changes because of satiation, extinction, or other causes; this tactic is the same as ours. There are two sources of resistance to change in Equation 7—friction, represented by μ , and behavioral inertia, represented by m . Nevin has not introduced a frictional term in his model; he interprets differential resistance as being due solely to different behavioral masses. But it is important to have such a term; without it, if A goes to zero, there can be no further change in behavior. Nonzero values for μ permit behavior to undergo extinction: when A goes to zero in Equation A7, the rate of responding must decrease (the only force on the left-hand side (lhs) is negative). Finally, the frictional force always drives rates slower, whereas the behavioral mass opposes any change in rate, increasing or decreasing. A more detailed comparison of these models is beyond this paper; it is hoped that the introduction of the frictional term will generalize Nevin's model sufficiently to account for a few anomalous findings (see, e.g., Harper & McLean, 1992), as Nevin's approach constitutes the most important dynamic theory in the field, in terms of continued empirical tests and development and in terms of applications. (See *Journal of Applied Behavior Analysis*, 29 [1996].)

Ceilings on Responding

Equation A7 ignores the ceiling on responding. This constraint becomes insuperable as response rates approach their ceiling (here, 1.0), so that Equation A7 becomes

$$A - \frac{\mu b}{1 - b} = \frac{m db}{dt}. \quad (\text{A9})$$

As response rates approach their ceiling (here, 1.0), the force necessary to overcome them (A) increases without bound. Now the force opposing A involves both the frictional resistance μb and the competition for expression, $1/(1 - b)$. At low response rates, the latter is approximately 1, returning us to A7. At equilibrium, when rates are no longer changing ($db/dt = 0$), it is $A = \mu b/(1 - b)$, whose solution is

$$b = \frac{A}{\mu + A},$$

which returns Equation 3, with $\mu = 1$. Thus, these expanded models, which include mass (m), differential resistance to change (μ), and constraints on responding, resolve to the simpler treatment, given in the body of the paper, at equilibrium.

If the force is changed to a new value A' , response rates will change as the integral of Equation A9. The dependent variable cannot be isolated, but numerical simulations show a smooth approach to the new equilibrium.

Behavioral mass, present in Equations A7–A9, does not appear in the equilibrium solutions; it modulates the speed of approach to those equilibria but not their locations. Slow changes in the force of arousal should give near-equilibrium changes in behavior that conform to Equation 10 (see, e.g., Figure 11). Faster changes in arousal level will give concave or convex convergence on that line, depending on the relative masses of the behaviors and the direction of change in arousal. There is a hint of such curvilinearity in Figure 12.

Behavioral Allometry

Transition analysis can be simplified by several gambits: collecting data where they are not too close to their ceiling, which permits use of Equation A8; graphing one behavior in terms of the other (rather than in terms of time); and rescaling the axes. First, it is shown that the trajectories converge on the equilibrium solutions given in the text. Rewrite Equation A8:

$$y = \frac{C}{\mu} A' (1 - e^{-\mu t/m_T}) + \frac{C}{\mu} A e^{-\mu t/m_T}.$$

Note that, when $t = 0$, $y = CA/\mu = y_0$. Then for target responses,

$$y = \frac{C}{\mu} A' (1 - e^{-\mu t/m_T}) + y_0 e^{-\mu t/m_T}, \quad (\text{A10a})$$

and for other responses,

$$x = \frac{1-C}{\mu} A' (1 - e^{-\mu t/m_0}) + X_0 e^{-\mu t/m_0}. \quad (\text{A10b})$$

Eliminate the unknown asymptotic level of arousal A' to obtain

$$y = \left(\frac{C}{1-C} \right) \left(\frac{1 - e^{-\mu t/m_T}}{1 - e^{-\mu t/m_0}} \right) (x - X_0 e^{-\mu t/m_0}) + y_0 e^{-\mu t/m_T}, \quad t > 0. \quad (\text{A11})$$

As $t \rightarrow \infty$, y approaches proportionality $[C/(1 - C)]$ with other behavior, thus satisfying Equation 10. If the behavioral masses are equal ($m_T = m_0$), that equation is also satisfied at all time values. But if the masses differ, then Equation A11 describes a curvilinear approach from one equilibrium to the next. Depending on the size of μ/m_i and the speed and extent of the change in motivation, the trajectories may hug the equilibrium line so closely that deviations are not noticeable or may bulge out from that line, so that Equation 10 appears violated.

A simple rendition of this approach to equilibrium is obtained by rescaling the axes. Note that as $t \rightarrow \infty$, $y \rightarrow CA'/\mu = y_\infty$,

and $x \rightarrow (1 - C)A'/\mu = x_\infty$. Substitute into Equation A10 and solve to eliminate μt :

$$\frac{y - y_\infty}{y_0 - y_\infty} = \left[\frac{x - x_\infty}{x_0 - x_\infty} \right]^{m_O/m_T}. \quad (\text{A12})$$

This shows that the proportional distance of target responses from their asymptote is a power function of the proportional distance of *other* behaviors from their asymptote. The power is m_O/m_T , the ratio of the masses of the responses. If both responses are equally labile under reinforcement, the masses are equal, and A12 describes one straight line.

The mass of a response is not necessarily associated with its physical ponderousness but with insensitivity to changes in the

force of reinforcement. Prepared responses will have relatively little mass, whereas contraprepared ones will have much greater mass. We should expect species-specific defense reactions, such as flight, to be easily conditioned to aversive stimulation and to have correspondingly small values of m in those contexts. In like manner, we should expect pigeon's keypecking to be associated with small values of m in contexts in which food is available. The data of Figure 12 are reproduced in Figure A2, along with the trajectories given by Equation A12, with $m_O/m_T = 1.9$.

(Manuscript received December 2, 1996;
revision accepted for publication November 17, 1997.)