5-2013

A Comparison of Resistance to Extinction Following Dynamic and Static Schedules of Reinforcement

Andrew R. Craig
Utah State University

Recommended Citation

Craig, Andrew R., "A Comparison of Resistance to Extinction Following Dynamic and Static Schedules of Reinforcement" (2013). *All Graduate Theses and Dissertations*. 2012.
https://digitalcommons.usu.edu/etd/2012

A COMPARISON OF RESISTANCE TO EXTINCTION FOLLOWING DYNAMIC AND STATIC SCHEDULES OF REINFORCEMENT

by

Andrew R. Craig

A thesis submitted in partial fulfillment
of the requirements for the degree

of

MASTER OF SCIENCE

in

Psychology

Approved:

Timothy A. Shahan
Major Professor

Amy L. Odum
Committee Member

Gregory J. Madden
Committee Member

Mark R. McLellan
Vice President for Research and
Dean of the School of Graduate Studies

UTAH STATE UNIVERSITY
Logan, Utah

2013

ABSTRACT

A Comparison of Resistance to Extinction Following Dynamic and Static Schedules of Reinforcement

by

Andrew R. Craig, Master of Science

Utah State University, 2013

Major Professor: Dr. Timothy A. Shahan
Department: Psychology

Resistance to extinction of single-schedule performance is negatively related to the reinforcer rate that an organism experienced in the pre-extinction context. This finding opposes the predications of behavioral momentum theory, which states that resistance to change, in general, is positively related to reinforcer rates. The quantitative model of extinction provided by behavioral momentum theory can describe resistance to extinction following single schedules in a post-hoc fashion, and only if the parameters of the model are allowed to vary considerably from those typically derived from multiple-schedule preparations. An application of the principles of Bayesian inference offers an alternative account of extinction performance following single schedules. According to the Bayesian change-detection algorithm, the temporal intervals of non-reinforcement that an organism experiences during extinction are compared to the temporal distribution of reinforcers that the organism experienced during baseline. A transition to extinction is more readily detectable when the previously collected distribution of reinforcers in time is populated with relatively short intervals (i.e., when more frequent reinforcement was experienced during baseline). The Bayesian change-detection algorithm also suggests that changes in reinforcer rates are more detectable when organisms have temporally proximal experience with frequently changing rates. The current experiment investigated this novel prediction. Pigeons pecked keys for food under schedules of reinforcement that arranged either relatively dynamic reinforcer rates or relatively static rates across conditions. Following each period of reinforcement, resistance to extinction was assessed. Persistence was greater following static contingencies than following dynamic contingencies for the majority of subjects. These data provide support for the Bayesian approach to understanding operant extinction and might serve to extend behavioral momentum theory by offering change detection as an additional mechanism through which extinction occurs.

(85 pages)

PUBLIC ABSTRACT

A Comparison of Resistance to Extinction Following
Dynamic and Static Schedules of Reinforcement

by

Andrew R. Craig, Master of Science

Utah State University, 2013

Major Professor: Timothy A. Shahan, Ph.D.
Department: Psychology

Behavioral momentum theory states that the Pavlovian stimulus-reinforcer relation governs resistance to extinction. Thus, higher baseline reinforcer rates should produce more persistent behavior than lower baseline reinforcer rates. Though this positive relation generally is observed when behavior is maintained and subsequently disrupted in multiple schedules, the opposite relation is observed when single schedules are used. An alternative framework of extinction performance based on Bayesian change detection may be applied intuitively to describe resistance to extinction in single schedules of reinforcement. This approach asserts that detection of changes in reinforcer rates during extinction should be easier following training with rich reinforcer rates than lean reinforcer rates. Further, the Bayesian-based approach to understanding operant extinction implies that experience with changing reinforcer rates during baseline should further facilitate change detection and thereby hasten extinction. The current experiment tested this prediction. Pigeons pecked keys for food according to both static (i.e., reinforcer rates were held constant) and dynamic (i.e., reinforcer rates changed both between and within session) variable-interval schedules across successive baseline conditions. Following each baseline condition, extinction was assessed. Proportion-of-baseline response rates in extinction following static training generally were higher than these rates following dynamic training. These data provide tentative support for the Bayesian-based approach to understanding operant extinction. They do not, however, challenge the underlying, quantitative framework of behavioral momentum theory if change detection is considered an additional mechanism through which extinction occurs.

ACKNOWLEDGMENTS

I would like to thank my colleagues in the Psychology Department's Behavior Analysis Seminar for their many insightful and energetic conversations about both change detection and Bayesian inference as they apply to operant behavior. Special thanks and deep gratitude are owed to my advisor, Dr. Timothy Shahan, for encouraging curiosity and the exploration of various theoretical accounts of operant behavior. I also thank my committee members, Drs. Gregory Madden and Amy Odum, for their thoughtful critiques, comments, and guidance.

I dedicate this thesis to my mother and stepfather, Sharon and Henry Rock, and to all of my friends and family who have offered interminable support, always.

Andrew R. Craig

CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Behavioral momentum theory states that the resistance to extinction of operant behavior is a direct function of the overall reinforcer rate that an organism experienced in the pre-extinction stimulus context. In multiple schedules, behavior maintained in a component associated with a relatively rich reinforcer rate is more resistant to extinction than behavior maintained in a component associated with a relatively lean rate (see Nevin, 2002, for a review). In the context of single schedules of reinforcement, however, the opposite relation typically is observed. Here, behavior maintained by a leaner schedule of reinforcement is relatively more resistant to extinction (e.g., Shull & Grimes, 2006). This relation is not well described by the augmented model of extinction based on behavioral momentum (see Nevin & Grace, 2000). It is, however, described by an alternative framework of extinction performance based on Bayesian statistical inference (Gallistel, 2012).

Bayesian change detection, as it applies to operant extinction, asserts that organisms compare previously collected information about reinforcer rate in a stimulus context to currently experienced temporal intervals between events to determine whether changes in reinforcer rate have occurred (see Gallistel, 2012). When comparing previously established distributions of reinforcers in time to the absence of reinforcers during extinction, transitions from rich schedules to extinction are more readily detectable than transitions from lean schedules to extinction. Changes in behavior are suggested to follow detection of changes in reinforcer rates.

The Bayesian change-detection algorithm also has been applied to describe the adaptability of organisms' choice behavior to changes in relative reinforcer rates in concurrent-schedule arrangements. When relative reinforcer rates transition from a period of prolonged stability to new relative rates, the behavioral allocation of rats (Gallistel, Mark, King, & Latham, 2001) and pigeons (Mazur, 1995, 1996) adapts slowly to the prevailing contingencies. Further, reversions to pre-change behavioral allocation occur early during the first few sessions following a change in relative reinforcer rates. When relative reinforcer rates change often, however, the behavioral allocation of rats rapidly adjusts to these changes with no reversions to pre-change allocation (Gallistel et al., 2001). These findings suggest that the ability of organisms to detect changes in reinforcer rates is positively related to the frequency with which these changes occur. Indeed, Bayesian inference considers the prior probability of changes in reinforcer rates as an additional determinant of change detection.

Experience with frequently changing reinforcer rates presumably alters the way an organism incorporates new temporal information into the distribution of reinforcers in time derived from pre-change experiences. Specifically, the rate at which new temporal information comes to dominate this distribution is positively related to the frequency with which changes have occurred (see Gallistel, 2012; Gallistel et al., 2001). Because the rate-detection component of Bayesian inference also predicts resistance to extinction following single schedules of reinforcement, it is reasonable to believe that frequent changes in reinforcer rates could affect resistance to extinction of single-schedule performance. This prediction, however, has not been investigated.

The purpose of the current experiment was to determine if frequent changes in reinforcer rates differentially impact resistance to extinction of single-schedule performance relative to extinction following stable reinforcer rates when overall, mean rates are held constant. Pigeons pecked keys for food according to either a Static-variable-interval (VI) or a Dynamic-VI schedule across conditions. All of the subjects first responded under a Dynamic-VI schedule in which reinforcer rates changed both between and within sessions (cf., Gallistel et al., 2001). Here, four-session blocks with the same programmed mean reinforcer rates were arranged to provide a non-arbitrary means of determining the reinforcer rates that were used in the following Static-VI condition. In the Static-VI condition, the VI value was determined by equating it to the mean reinforcer rate each subject obtained during the final eight sessions (i.e., two blocks) of the preceding Dynamic-VI condition to ensure that the reinforcer rates delivered by both schedules were comparable within subjects. Finally, each subject again responded under a Dynamic-VI schedule. Following each reinforcement condition, resistance to extinction was assessed.

CHAPTER II

BACKGROUND

Behavioral Momentum and Resistance to Change

Behavioral momentum theory is a conceptual framework that describes the effects of reinforcement on the persistence of operant behavior. According to behavioral momentum, there are two separable aspects of operant behavior that, in conjunction, constitute momentum: response rate and resistance to change (Nevin, 1992a; Nevin, Mandell, & Atak, 1983). Response rate is the frequency with which responding occurs in the experimental context and is driven by the association between responding and the delivery of reinforcers made contingent on the response. Resistance to change describes the persistence of behavior in the face of disruption and is governed by the Pavlovian relation between reinforcers and the stimulus context in which they are delivered.

Resistance to change generally is studied in multiple schedules in which two or more stimuli are presented sequentially, each associated with a separate schedule of reinforcement (see Nevin, 1974). In multiple schedules, different reinforcer rates are associated with different stimulus conditions (e.g., different response-key colors for pigeons). Using this paradigm, the effects of reinforcer rates on resistance to change of behavior can be studied within subjects. To assess resistance to change, disruptors of comparable magnitude are applied to all multiple schedule components simultaneously. Such disruptors as pre-session feeding, free inter-component-interval (ICI) food presentations, and extinction historically have been used to assess resistance to change (e.g., Harper, 1996; Nevin, 1974, 1992a, 2002; Nevin, Mandell, & Yarensky, 1981).

Nevin (1974, Experiment 1) conducted an experiment that is representative of the study of resistance to change in multiple schedules. Here, pigeons responded under a two-component multiple schedule in which one component was associated with a VI 60-s schedule of reinforcement and the other component was associated with a VI 180-s schedule. To test for resistance to change, Nevin delivered different frequencies of response-independent food on variable-time (VT) schedules during ICIs. Nevin observed that behavior maintained by the component correlated with the VI 60-s schedule was more resistant to the disruptive impact of free ICI food than was behavior maintained by the component correlated with the VI 180-s schedule. Further, there was a negative relation between the frequency with which VT food was presented and resistance to disruption. This finding is general to the investigation of resistance to change (for a review, see Nevin, 2002) and also has been observed in humans (e.g., Ahearn, Clark, Gardenier, Chung, & Dube, 2003; Cohen, 1996; Mace et al., 1990; Mace et al., 2010), rats (e.g., Blackman, 1968), and goldfish (Igaki & Sakagami, 2004).

Response rate also is positively associated with reinforcer rates when all reinforcers are delivered contingent on a single response (see Shull, 2005, for a review; see also Catania & Reynolds, 1968; Herrnstein, 1970). Evidence for the separable nature of response rate and resistance to change comes from studies in which additional reinforcers are made available either response independently or contingent on a separate, alternative response. Nevin, Tota, Torquato, and Shull (1990) conducted a series of experiments with pigeons that demonstrated differences in response rate and resistance to change under both of the above conditions. In their first experiment, pigeons responded under a two-component multiple schedule of food reinforcement. In one component, only response-dependent food was available and was delivered according to a VI 60-s schedule. In the second component, three types of conditions were arranged. In one condition, only response-dependent food was available according to a VI 60-s schedule. In the second type of condition, VT food was superimposed over the VI 60-s schedule and was delivered, on average, every 30 or 15 s. In the final type of condition, VI and VT foods concurrently were available such that the overall reinforcer rate in this component was equal to the reinforcer rate delivered by the VI 60-s schedule present in the other component (e.g., VI 180 s with VT 90 s; VI 300 s with VT 75 s). Resistance to change was examined by using both pre-feeding and extinction. Response rates were lower in the component with added VT food, but resistance to both pre-feeding and extinction was higher in this component when the overall rate of food presentations was higher than in the other VI-only component. In their second experiment, Nevin et al. (1990) arranged a three-component multiple-concurrent schedule of reinforcement. In the first component, a VI 240-s schedule operated on the target-response key, and a separate VI 80-s schedule operated on an alternative-response key. In the second component, a VI 240-s schedule again was available for target responding, but alternative responding had no consequences (i.e., extinction). Finally, in the third component, a VI 60-s schedule was available for target responding, and alternative responding had no consequences. Resistance to change was assessed by using both pre-feeding and extinction in different conditions. Again, target response rates were lower, but resistance to change was higher, in the component with additional food.

The finding that added alternative reinforcement differentially impacts response rate and resistance to change has been replicated in rats using qualitatively different reinforcers. For example, Shahan and Burke (2004) demonstrated that resistance to change of ethanol-maintained responding was higher when response-independent food was concurrently available. Further, response-independent deliveries of sweetened condensed milk have been demonstrated to increase resistance to change of food-maintained responding (Grimes & Shull, 2001). Thus, in terms of resistance to change in multiple schedules, it appears that neither the source of reinforcers nor the type of reinforcer matters: Resistance to change is a function of the overall reinforcer rate delivered in a stimulus context.

Modeling Resistance to Change

To isolate the contribution of reinforcer rate to resistance to change, it is important to ensure that the magnitudes of the disruptors applied to separate multiple-schedule components are either the same between components or that disruptors with different magnitudes be clearly and quantitatively specified. The disruptive impacts of pre-feeding or presenting free ICI food presumably are directly related to the magnitude of the disruptor (i.e., the amount of food presented). Nevin et al. (1983) proposed a model of resistance to change that assumed such constant magnitudes of disruption. In this model, decreases in responding given a disruptor were directly related to the magnitude of the disruptor and inversely related to behavioral mass (i.e., the strength of the response that is governed by reinforcer rate). The model is:

$$\log\left(\frac{B_x}{B_o}\right) = \frac{-f}{r^b}, \tag{1}$$

where the left side of the equation is log proportion-of-baseline response rates given a disruptor. On the right side of the equation, *f* is the impact of the disruptor, and the denominator is behavioral mass represented as baseline reinforcer rate (in reinforcers per hr) exponentiated by *b*, which represents sensitivity to reinforcer rate. Note that, in the original model, the denominator of the equation was simply a mass term, *m*, but subsequent parametric analyses have revealed that behavioral mass is well described by a power function of reinforcer rate (see Nevin, 2002).

The parameter *r* is set equal to obtained reinforcer rates (in reinforcer per hr) in the pre-disruption context. The parameters *f* and *b*, however, are free to vary and are set by fitting the predictions of the model to obtained data using least-squares regression. The value of the *f* parameter is positively related to the magnitude of the disruptor that is used and the sensitivity parameter, *b*, typically assumes a value of 0.5 (see Nevin, 2002). Because *f* and *b* are constant between multiple-schedule components and the value of *r* is separately determined for each component, the model predicts that behavior maintained by a component correlated with a higher reinforcer rate will be more resistant to disruption than behavior maintained by a component correlated with a lower rate.

Nevin, McLean, and Grace (2001) argued that the disruptive impact of extinction might differ from the disruptive impact of either pre-feeding or free ICI food, so a general disruptive term like *f* in Equation 1 might not be sufficient to model responding in extinction. Pre-feeding and presenting free ICI food may be classified as "external" disruptors because the underlying contingencies of reinforcement in the components of the multiple schedule remain intact when disruption is applied. One can assume that these disruptors are of equal force regardless of the schedule of reinforcement to which they are added. Extinction, however, is an "internal" disruptor. Because extinction involves the disruption of behavior by means of disrupting the underlying schedules of reinforcement, the magnitude of its disruptive impact might depend on the schedule it replaces.

Catania (1973) commented on two separate mechanisms through which extinction of operant behavior may occur: The delivery of reinforcers in the experimental context may be discontinued and/or the response-reinforcer contingency may be suspended (e.g., by replacing a VI schedule with a comparable VT schedule). Nevin and Grace (2000) proposed an augmented model of extinction based on behavioral momentum theory that accounts for the proposed differences between extinction and other forms of disruption. The model appears as follows:

$$\log\left(\frac{B_t}{B_o}\right) = \frac{-t(c + dr)}{r^b}. \tag{2}$$

Here, the ratio of B_t to B_o is proportion-of-baseline response rates at time t in extinction. The numerator of the right side of the equation represents the disruptive effects of extinction. The parameter t denotes time in extinction (measured in sessions), c is the disruptive impact of suspending the response-reinforcer contingency, r in the numerator indicates the change in reinforcer rates from baseline to extinction, and d scales the disruptive impact of this generalization decrement. The denominator of the equation is the strengthening term that describes the contribution of reinforcer rates to persistence of behavior in extinction. Here, behavioral mass is represented as it was in Equation 1.

The variable *t* increments by one with each session of extinction and the *r* variables in the numerator and the denominator are derived from the obtained reinforcer rate (in reinforcers per hr) delivered during baseline. The parameters *c*, *d*, and *b* are free to vary and are determined by fitting the model to obtained extinction data. When fit to extinction data from multiple schedules, these variables typically assume values of 1, 0.001, and 0.5, respectively (see Nevin & Grace, 2000). As the reinforcer rate delivered in a stimulus context increases, *r* increases in both the denominator (i.e., behavioral mass becomes greater) and the numerator (i.e., the disruptive impact of generalization decrement between baseline and extinction becomes greater). This feature of Nevin and Grace's model of extinction captures the different disruptive impacts present in transitions to extinction from schedules that differ in reinforcer rate.

To describe proportion-of-baseline response rates in standard units, both the left and right sides of Equation 2 may serve as the exponent of 10. This manipulation eliminates log transformation of proportion-of-baseline response rates on the left side of the equation. This equation appears as follows:

$$\frac{B_t}{B_o} = 10^{\frac{-t(c+dr)}{r^b}} \tag{3}$$

Here, all parameters are as they were in Equation 2. Figure 1 presents simulations of Equation 3 across a range of reinforcer rates using the typical free-parameter values described above.

As Catania (1973) suggested, the disruptive effects of suspending the response-reinforcer contingency (*c*) and of generalization decrement (*dr*) are assumed to be separate and independent processes in the augmented model. Extinction of operant responding may occur if the response-reinforcer contingency is terminated in the absence of a decrease in the rate of reinforcer presentations by delivering reinforcement independently of responding (see Boakes, 1973; Rescorla & Skucy, 1969). Cessation of responding tends to occur more rapidly when reinforcers are withheld entirely than when only the response-reinforcer contingency is suspended. In terms of the model, this is because both *c* and *dr* provide additive forms of disruption.

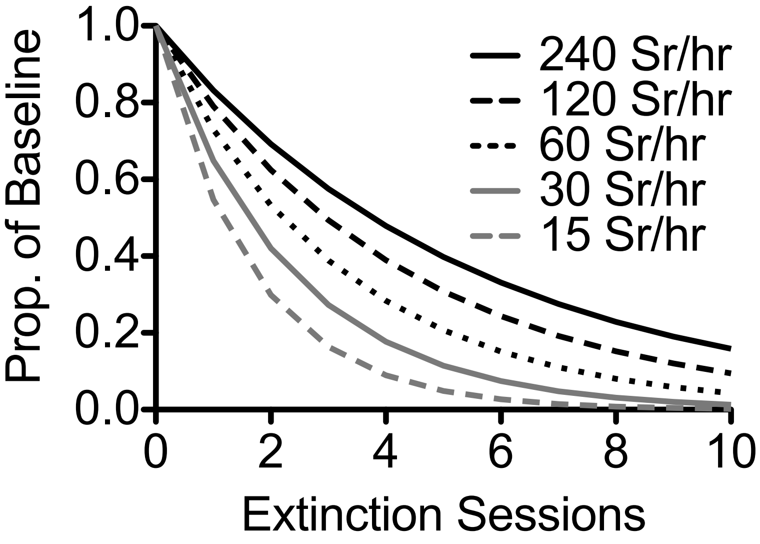

Figure 1. Simulations of Nevin and Grace's (2000) augmented model of extinction across various values of *r* (reinforcers per hr during baseline). Simulations were conducted using the following parameter values: $c = 1$, $d = 0.001$, and $b = 0.5$.

Single Schedules: A Caveat

The predictions of Equations 2 and 3 account for the general finding that resistance to extinction is greater in multiple-schedule components associated with higher reinforcer rates (see Figure 1). When single schedules of reinforcement are considered, however, the opposite effect typically is observed. A direct comparison of the differential effects of reinforcer rate on resistance to extinction in single schedules and multiple schedules was conducted by Cohen (1998). Here, one group of rats pressed levers for food in a multiple VI 30-s VI 120-s schedule. When reinforcement was suspended, responding in the component correlated with the VI 30-s schedule was more resistant to change than responding in the component correlated with the VI 120-s schedule. A second group of rats performed under single VI schedules with the same values as those used in the multiple schedule. When extinction was introduced, responding trained under the VI 120-s schedule was more resistant to extinction than responding trained under the VI 30-s schedule.

Shull and Grimes (2006, Experiment 1) subsequently replicated these results. Here, rats poked keys for food under separate VI 60-s and VI 480-s schedules. When reinforcement was suspended for a single, prolonged session, Shull and Grimes observed a negative relation between baseline reinforcer rates and resistance to extinction. They interpreted this result by suggesting that the transition from a relatively rich schedule of reinforcement to extinction was more discriminable than the transition from a relatively lean schedule to extinction.

To describe their data mathematically, Shull and Grimes (2006) used the augmented model of extinction (Equation 2). Fits of this model by least-squares regression to the data of individual subjects revealed larger than typical contributions of baseline reinforcer rates to response strength ($\bar{b}$ = 0.67), large disruptive impacts of generalization decrement ($\bar{d}$ = 0.23), and virtually no disruptive impact of suspending the response-reinforcer contingency (c = 0, in most cases). These parameter values differ substantially from those typically derived from fits of the augmented model to extinction data following multiple-schedule training (i.e., b = 0.5, d = 0.001, c = 1). It is theoretically unclear why generalization decrement would provide relatively little disruption in multiple schedules and a large amount of disruption in single schedules. Likewise, breaking of the response-reinforcer contingency occurs in extinction following both types of schedules. It is not intuitive to presuppose that the disruptive impact of doing so would be present only in multiple schedules. Further, these parameter values predict very steep extinction functions when carried out across sessions (see Figure 2 for model simulations using Shull and Grime's parameter values across the same range of reinforcer rates depicted in Figure 1, for ease of comparison). Given these theoretical considerations, the utility and generality of the augmented model of extinction, and thereby the description of extinction performance offered by behavioral momentum theory, might be restricted and the model's characterization of extinction performance might be incomplete.

Extinction from a Rationalist Perspective

Like Shull and Grimes (2006), Gallistel (2012) suggested that single schedules with lean reinforcer rates generate behavior that is more resistant to extinction than single schedules with rich rates because a transition to extinction following a lean schedule is relatively more difficult to detect than a transition to extinction following a rich schedule. Gallistel, however, took a different approach to describing and modeling this finding. According to Gallistel, organisms time the intervals between reinforcer deliveries in their environments to generate knowledge about the current state of reinforcer availability. That is, each reinforcer delivery or, more precisely, the intervals between reinforcer deliveries provide information about the distribution of reinforcers in time. During extinction, organisms use the information they previously gathered about the temporal

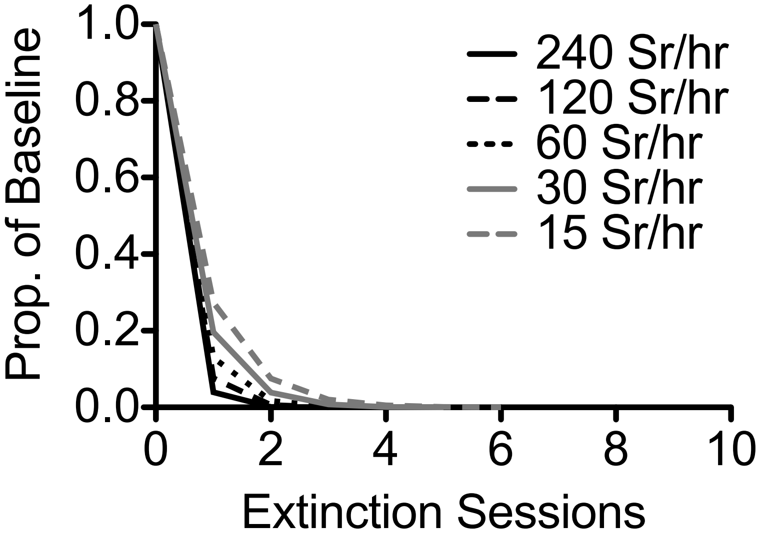

Figure 2. Prediction of Equation 3 across various reinforcer rates using the mean parameter values obtained from Shull and Grime's (2006) first experiment. These values were: $c = 0$, $d = 0.23$, and $b = 0.67$.

locations of reinforcers and the temporal information that they continue to extract from the environment to make decisions about the stability of reinforcer availability. If the time since a previous reinforcer becomes longer than anticipated (based on previously collected temporal information), organisms are able to detect that a change in reinforcer rates occurred.

Bayesian statistical inference provides both a means of modeling the way in which organisms detect these changes in reinforcer rate and an alternative approach to understanding the persistence of behavior during extinction of single-schedule performance. In the operant adaptation of Bayesian inference, an organism gathers two pieces of information that are crucial to detecting changes in the temporal distribution of reinforcers. The first piece of information is the relative probability distribution of inter-reinforcer interval (IRIs) and the second is a distribution of the probability of changes in reinforcer rates in the organism's environment. When new information is made available to the organism (e.g., an interval between reinforcer deliveries, the interval from the last reinforcer, or the interval from session onset if no reinforcers have been collected), it is compared to these distributions. If this temporal information differs sufficiently from the previously collected information, the organism may detect a change in reinforcer rates.

Prior experience with relatively lean reinforcer rates results in a distribution of IRIs that is populated with relatively long intervals. The opposite is true of experience with a relatively rich reinforcer rate. From an analytical standpoint, an organism should be able to detect a transition from a period of relatively frequent reinforcement to extinction more quickly than a transition from a period of relatively infrequent reinforcement for a simple reason: A long interval since reinforcement (such as those experienced during extinction) is more likely to belong to a distribution of long IRIs than to a distribution of short IRIs. The predictions of Bayesian inference, then, describe the results that are obtained when single schedules of reinforcement are replaced with extinction.

Effects of Frequent Change

The Bayesian change-detection mechanism Gallistel (2012) applied to the rates of the extinction of single-schedule performance also has been applied to describe choice in environments with static and dynamic relative reinforcer rates (see Gallistel et al., 2001). In this application, the probability with which changes in reinforcer rates occur in the organism's environment becomes important. Gallistel and colleagues suggested that choice is the product of detecting relative reinforcer rates in a concurrent-schedule context. Once relative rates accurately have been estimated, choice behavior is an innate, biologically selected policy that governs an organism's allocation of behavior according to the matching law (i.e., relative rates of responding across alternatives match the relative reinforcer rates associated with those alternatives; see also Gallistel et al., 2007).

The development of matching through change- and rate-detection mechanisms most clearly has been demonstrated by experiments investigating transitions from one choice situation to another with novel relative reinforcer rates. For example, Dreyfus (1991; Experiment 1) examined the effects of changing relative reinforcer rates both between and within sessions in pigeons while holding overall reinforcer rate constant across alternatives. In this experiment, Dreyfus used a switching-key, concurrent VI VI procedure (Findley, 1958). A single VI 60-s schedule determined the availability of a reinforcer, and each reinforcer was assigned to an alternative probabilistically to ensure overall reinforcer rates were controlled (cf., Stubbs & Pliskoff, 1969). The probability that a reinforcer would be assigned to one alternative was either .9 or .1; this probability of assignment changed between sessions and at regular intervals within sessions. When changes in relative reinforcer rates occurred, the pigeons' time allocation abruptly changed to match the new contingencies. Interpretation of these data in terms of Bayesian inference is complicated, however, by the regularity of schedule changes. That is, the ability of the pigeons to detect changes in relative reinforcer rates based on experienced reinforcer rates, alone, was confounded with the temporal information provided by the regularity with which changes in relative rates occurred.

The work of Mazur (1995, 1996) presented a clearer description of transitions from stable-state choice situations to novel choice situations. In Mazur's (1995) experiment, pigeons responded under concurrent schedules with stable relative reinforcer rates for a number of sessions. Reinforcers became available according to a single VI 30-s schedule, and the probability that a given reinforcer would be assigned to the left key initially was .1, .25, or .4. Then, relative reinforcer rates were switched such that left-key responding was reinforced with a probability of .9, .75, or .6 for a few sessions. At the beginning of the first four post-change sessions, Mazur observed 'spontaneous recovery' of pre-change behavioral allocation and that choice behavior of the pigeons adjusted to the new contingencies throughout the course of a session. More complete transitions occurred with continued exposure to post-change contingencies across sessions. Mazur (1996) used a similar procedure but presented data on shifts from a high probability of reinforcement on a key to a low probability of reinforcement on a key to investigate the possibility that the spontaneous recover of pre-change allocation observed in Mazur (1995) could have resulted from the pigeons' tendency to allocate behavior in a manner closer to indifference when changes in relative rates occurred. Again, Mazur (1996) observed both spontaneous recovery of pre-change allocation at the beginning of the first few post-change sessions and gradual adaptation to the new contingencies across sessions.

The spontaneous recovery Mazur (1995, 1996) observed may be characterized as a behavioral-momentum-like effect if one considers pre-change behavioral allocation as a behavioral unit and the change in relative reinforcer rates to be a disruptor. When this disruptor was applied, the behavior of the pigeons initially was resistant to changes in reinforcer rates. The spontaneous recovery of pre-change behavioral allocation decreased across sessions. This is similar to the gradual decrease in response rates across sessions observed when extinction is applied as a disruptor (see Figure 1; see also Nevin, 2002; Nevin & Grace, 2000). In this situation, too, it takes a number of sessions for behavior to fully adjust to the prevailing contingencies.

When relative reinforcer rates in choice situations change frequently, however, considerably different results are obtained. Gallistel et al. (2001) directly compared the way rats' choice behavior adapted to changes in relative reinforcer rates when these rates changed infrequently, as they did in Mazur's (1995, 1996) experiments, and when they changed rapidly. In this experiment, rats were presented first with stable relative rates of electronic brain stimulation in a two-lever choice situation. During the first phase, relative reinforcer rates between the two alternatives were held constant at either a 1:1 or 1:4 ratio delivered on the left and right alternatives, respectively. The overall reinforcer rate was held constant at 9.4 deliveries per min. After 32 sessions performing under these contingencies, relative reinforcer rates changed mid-session, but the overall rate was maintained. These new contingencies were held constant for 20 additional sessions following the session in which the change occurred. As Mazur observed in pigeons, behavior of the rats took a number of sessions to adjust to the post-change contingencies. Spontaneous recovery of pre-change behavioral allocation also was observed in the first few post-change sessions.

In the second phase of the experiment, Gallistel et al. (2001) arranged rapidly changing choice situations in which relative or absolute reinforcer rates changed both between and within sessions. Changes in both relative and absolute rates occurred across 20-session-long phases for each subject in a counterbalanced manner. Absolute rate changes randomly alternated between 2.1, 6, 9.4, and 18 reinforcers per min. When changes in relative rates occurred, the overall reinforcer rate was again held constant at 9.4 reinforcers per min, but the ratios of reinforcers delivered according to the left and right alternatives varied between 9:1, 3:1, 1:1, 1:3, and 1:9. When the contingencies changed often, rats' behavior adjusted to changes in relative reinforcer rates rapidly, within only a few visit cycles between the two alternative responses. Further, no reversions to pre-change behavioral allocation were observed.

Gallistel and colleagues (2001) interpreted these results from a Bayesian statistical perspective: If change was inherent in the prior distributions the rats extracted from the experimental context (as it was in the second phase of this experiment), comparisons of newly acquired information to the prior distribution were more sensitive to detection of changes in reinforcer rates. Once a change was detected by comparing these two pieces of information, the prior distribution was updated to reflect only the perceived, current contingencies. Information about overall reinforcer rates was not used to detect these changes in rate. This was not the case when changes were infrequent. In this situation, the prior distribution reflected the accumulation of temporal information across a larger window and did not completely update when a change was first detected. Here, the prior distribution more closely reflected overall reinforcer rates.

These results, coupled with those of investigations of resistance to extinction following single schedules of reinforcement (e.g., Cohen, 1998; Shull & Grimes, 2006), suggest detection of changes in reinforcer rates via the Bayesian change-detection algorithm can be affected most dramatically by two experimental parameters. The first parameter is the schedule of reinforcement that constitutes the prior distribution of reinforcers in time. In terms of extinction, this parameter is manifest in that behavior maintained by relatively lean single schedules is more resistant to extinction than behavior maintained by relatively rich single schedules. The second experimental parameter that putatively affects change detection is the frequency with which changes in reinforcer rates occur in the experimental context. The effects of this parameter on extinction performance is, as of yet, unknown.

CHAPTER III

PURPOSE

Extinction performance following single schedules of reinforcement is not well described by the augmented model of extinction based on behavioral momentum (Equations 2 and 3; see Cohen, 1998; Shull & Grimes, 2006), but it can be described using Bayesian statistical inference (see Gallistel, 2012). Here, behavior can be explained by assuming an organism gathers information about the temporal distribution of reinforces in its environment. In extinction, the organism compares currently experienced intervals to the previously generated distribution of reinforcers in time (i.e., IRIs) derived from pre-extinction experiences. When a currently experienced interval in extinction sufficiently exceeds the IRIs that comprise the prior distribution, the organism may detect a change in reinforcer rates and allocate its behavior accordingly.

An organism's ability to detect changes in reinforcer rates in choice situations is positively related to the frequency with which changes in rates occur (see Gallistel et al., 2001; see also Mazur, 1995, 1996). Again, this finding is well described by the Bayesian change-detection algorithm. From the standpoint of Bayesian statistical inference, frequent changes in reinforcer rates might also impact resistance to extinction of single-schedule performance, assuming these processes are mediated, at least in part, by the same mechanism. Frequent changes in reinforcer rates might alter the manner in which the organism tracks the temporal locations of reinforcers. In the case of choice behavior, currently perceived temporal intervals come to dominate the prior distribution of IRIs more quickly when change is highly probable. If this also is true of extinction performance following single schedules, it might result in a hastening of the extinction process.

The purpose of the current experiment was to investigate the extent to which experience with frequent changes in reinforcer rates impacts resistance to extinction relative to experience with infrequent changes when reinforcer rates were held constant across conditions. To determine if frequent changes in reinforcer rates affects resistance to extinction in single schedules, pigeons responded under Dynamic- and Static-VI schedules across conditions that alternated with extinction tests. All subjects first responded under a Dynamic-VI schedule in which reinforcer rates changed both between and within sessions and then under a Static-VI schedule in which rates were constant for the duration of the condition. Finally, all subjects again performed under a Dynamic-VI schedule.

All investigations of behavioral momentum theory historically have arranged stable reinforcer rates prior to extinction testing (see Nevin, 2002, for a review). It therefore is unclear if overall, mean reinforcer rates from baseline govern resistance to extinction or if response persistence is a function only of more recently experienced rates. To address this concern, four-session blocks with the same programmed mean reinforcer rate were arranged in the Dynamic-VI conditions. This arrangement ensured the mean reinforcer rates that each subject experienced in the final sessions preceding extinction testing were similar to the overall, mean reinforcer rate delivered in these conditions.

Because the individual VI schedules that comprised the Dynamic VI varied over a large range and were experienced for differing amounts of time within sessions, it was in principle possible for mean reinforcer rates during blocks of session in the Dynamic-VI conditions to vary between subjects. Accurate, a priori estimations of obtained reinforcer rates in these conditions therefore were not possible. The VI values used for the Static-VI condition were individually determined by equating the mean reinforcer rates delivered in this condition to the mean rate each subject obtained during the last two blocks of sessions in the preceding Dynamic-VI condition. This yoked design allowed for comparisons of resistance to extinction following Dynamic- and Static-VI conditions with similar reinforcer rates. Further, comparisons of resistance to extinction following both Dynamic-VI conditions helped to isolate any effects of repeated extinction tests on rate of extinction.

CHAPTER IV

METHOD

Design

A within-subjects ABCBAB design was used in this experiment. In phases labeled ‘A,’ subjects responded under a Dynamic-VI schedule in which reinforcer rates changed both between and within sessions. The ‘C’ phase was a Static-VI schedule. Here, the VI value for each subject was individually determined by averaging the reinforcer rates each subject obtained during the last eight sessions (i.e., two blocks) of the first Dynamic-VI condition. ‘B’ phases were extinction tests. The dependent variable of this experiment was resistance to extinction as measured by proportion-of-baseline response rates.

Subjects

Eight experimentally naïve homing pigeons (Double T Farm, Glenwood, IA) served. Each pigeon was housed separately in a colony room with a 12:12 hr light:dark cycle (lights on at 7:00 am) and had free access to water in its home cage. Sessions were conducted daily at approximately the same point during the light cycle. Pigeons were maintained at 80% of their free-feeding weights (± 15 g) by the use of supplementary post-session feedings when necessary. Reacquisition of key pecking during the Static-VI baseline condition was problematic for Pigeon 4277. This subject was maintained at the lower end of its 80% range for all conditions thereafter. Animal housing and care were conducted in accordance with the regulations of Utah State University's Institutional Animal Care and Use Committee (IACUC; protocol #1098).

Apparatus

Four sound-attenuating operant chambers for pigeons (dimensions approximately 29 cm long, 26 cm deep, and 29 cm high; Coulbourn Instruments, Whitehall, PA) were used. These chambers were constructed of clear plastic and aluminum, and each had an aluminum work panel on the front wall. Each aluminum work panel had two plastic response keys measuring 2.5 cm in diameter and located 16 cm from the floor of each chamber and 2 cm from either of the side walls. Only the right key was used during this experiment and was transilluminated white during pre-training and either orange or blue during Static- and Dynamic-VI conditions. The key colors that were associated with the Dynamic-VI and Static-VI schedules were counterbalanced across pigeons. A 28-v lamp located in the center of the work panel and 23 cm from the floor of the chamber provided general illumination. Both this lamp and the right key light were lighted at all times except during reinforcer deliveries and blackout periods, or when otherwise specified below. A food aperture measuring 6 cm wide by 5 cm high was located in the center of the work panel 5 cm from the floor of the chamber. A 28-v lamp illuminated this aperture during reinforcement, which consisted of 1.5 s of access to Purina Pigeon Checkers delivered by a solenoid-operated food hopper. White noise was present at all times to mask extraneous sound. All sessions were controlled by MedPC software (Med Associates, St. Albans, VT) by a PC computer.

Procedure

During all phases of the experiment, sessions began with a 60-s blackout period to allow the pigeons to acclimate to the operant chambers. All VI and VT schedules consisted of 10 intervals and were constructed using the constant-probability algorithm of Fleshler and Hoffman (1962). See Table 1 for a timeline of experimental conditions.

Hopper Training

Each pigeon first was trained to eat from the food hopper in the operant chamber. Direct observations of the subjects occurred regularly during this phase to ensure that each learned to eat when the hopper was raised. The food was delivered into the darkened chamber according to a VT 60-s schedule. Hopper-training sessions lasted 40 min and this phase terminated only when the pigeons ate reliably from the hoppers.

Table 1

Summary of Experimental Conditions

Condition	Sessions
Pre-Training	~ 30
Dynamic VI	36[a]
Extinction	14
Static VI	36
Extinction	15
Dynamic VI	36
Extinction	13

[a] This phase lasted 40 sessions for Pigeon 4277. See section "Dynamic-VI conditions" for more details.

Auto-Shaping/Hand Shaping

Following hopper training, each pigeon received a number of auto-shaping sessions to establish the key-peck response (see Brown & Jenkins, 1968). Each trial began with the illumination of the house light. After 45 s elapsed, the right key was illuminated for 5 s, after which both the house light and key light were extinguished and a reinforcer was delivered. If a peck to the lighted key occurred at any point during a trial, the trial terminated immediately and a reinforcer was delivered. After each reinforcer was delivered, the next trial began. Auto-shaping sessions consisted of 40 such trials. This phase last until each pigeon reliably pecked on at least 80% of the trials in a session.

Two pigeons failed to acquire key pecking after five sessions of auto-shaping. For these subjects, key pecking was shaped manually by delivering food contingently on successive approximations.

Ratio/Interval Training

After auto-shaping the pigeons moved to ratio training. In the initial session of ratio training, pigeons responded on the right key under a fixed-ratio (FR) 1 schedule. After 20 reinforcers were collected on FR 1, the ratio increased to FR 2 within the first session for an additional 20 reinforcers. During the second session, the FR increased from 2 to 4 after 20 reinforcers. The ratio continued to be doubled mid-session until each pigeon reliably pecked on an FR 16 schedule. If performance deteriorated at any point during this phase, the session and corresponding FR schedules were repeated until performance improved. This schedule then was replaced with a VI 10-s schedule. The

VI value increased across sessions in 5-s increments until each pigeon reliably pecked under a VI 60-s schedule. All sessions during this phase of training terminated after 40 min (excluding time for reinforcement) or after 40 reinforcers were delivered.

Dynamic-VI Conditions

In these conditions, responding was reinforced according to a VI schedule in which the VI value changed both between and within sessions. The VI values used for the Dynamic VI were 30, 70, 110, and 150 s and within-session changes to the VI schedule occurred 5, 15, or 25 min into a session. If an interval elapsed and a schedule change occurred before the pigeon collected the available reinforcer, that reinforcer was cancelled and the new schedule commenced immediately. Each subject experienced every possible combination of pre-change schedule, post-change schedule, and change point across sessions (i.e., a total of 36 schedule/change-point combinations were experienced across 36 sessions per condition). Four-session blocks of schedules were constructed such that each of the four possible schedules was experienced for the same amount of time in each block (see Table 2 for an exhaustive list of blocked sessions). This blocked arrangement introduced regularity in the mean reinforcer rates that each pigeon experienced across blocks to simplify determination of the reinforcer rates used in the Static VI condition (see below). The specific order of the schedules in each block and the order in which the pigeons experienced the blocks of sessions were randomized, but the subjects experienced the same progression of blocks during the first and second Dynamic-VI conditions to allow direct between-condition comparisons of extinction performance. Sessions lasted 30 min, excluding time for reinforcement.

Table 2

Schedules and Change Points for Blocks of Sessions Within the Dynamic-VI Conditions

Block	Change Point (Mins)	Pre-Change VI	Post-Change VI	Block	Change Point (Mins)	Pre-Change VI	Post-Change VI	Block	Change Point (Mins)	Pre-Change VI	Post-Change VI
1	5	30	70	4	5	70	110	7	15	30	150
	5	70	30		5	110	70		15	150	30
	15	110	150		25	30	150		25	70	110
	15	150	110		25	150	30		25	110	70
2	5	30	110	5	5	70	150	8	15	30	110
	5	110	30		5	150	70		15	110	30
	15	70	150		25	30	110		25	70	150
	15	150	70		25	110	30		25	150	70
3	5	30	150	6	5	110	150	9	15	30	70
	5	150	30		5	150	110		15	70	30
	15	70	110		25	30	70		25	110	150
	15	110	70		25	70	30		25	150	110

For Pigeon 4277, responding decreased inexplicably during the last block of the first Dynamic-VI condition. Accordingly, the first block of sessions this pigeon experienced was re-introduced prior to extinction testing. The second Dynamic-VI condition for this pigeon was truncated to 36 sessions by removing the first block of sessions from its progression.

Static-VI Condition

In the Static-VI condition, responding was reinforced according to a single VI schedule. The VI value for each subject was individually determined by equating it to the mean reinforcer rates obtained by that subject in the last eight sessions of the preceding Dynamic-VI condition. These values were rounded to the nearest 1,000th of a whole number (see Table 3 for a list of these values). The Static-VI condition included 36 sessions. Again, sessions terminated after 30 min, excluding time for reinforcement.

Table 3

Static-VI Values

Subject	Static VI (s)
4275	64.865
4276	67.290
4274	72.362
1499	67.290
4278	66.977
4188	68.246
4189	66.977
4277	70.244

Extinction Testing

Resistance to extinction was assessed for each subject following completion of each Dynamic-VI and Static-VI condition. The stimulus context in extinction was identical to that of the preceding baseline condition. Responding, however, had no programmed consequences. These phases lasted a minimum of 10 sessions and continued until responding for each subject was below 10% of the previous baseline condition. Each session of extinction lasted for 30 min.

Justification

Changes in reinforcer rates during the Dynamic-VI condition occurred both between and within sessions to reflect the procedure used by Gallistel et al. (2001). Reinforcer rates and session durations were constrained to avoid pigeons eating excessive amounts of food within sessions to help maintain constant levels of motivation throughout the experiment. Reinforcer rates were not equated within each session, however, to limit the information provided by the pre-change schedule about the point at which schedules would change, a problem encountered by Dreyfus (1991). This procedure also limited the information that a pre-change schedule provided about the upcoming post-change schedule.

Data Analysis

All statistics reported below were deemed significant at an α level of .05. For all analyses of variance (ANOVA), assumptions of sphericity were tested with Mauchly’s method. If this assumption was violated, Greenhous-Geisser corrections of degrees of freedom were used.

CHAPTER V

RESULTS

Reinforcer Rates

Mean obtained reinforcers per hr, standard error of the mean (SEM), and relative reinforcer rates (i.e., rates from either Dynamic-VI baseline condition divided by those from the Static-VI baseline condition) from the last eight sessions of each condition for all subjects are displayed in Table 4. Mean reinforcer rates, aggregated across subjects, also are displayed in Figure 3. Mean reinforcer rates for the first Dynamic-VI, Static-VI, and second Dynamic-VI baseline conditions were 52.97 (*SEM* = 3.30), 50.25 (*SEM* = 0.44), and 53.56 (*SEM* = 3.43) reinforcers per hr, respectively. For each pigeon, obtained reinforcer rates were lower in the Static-VI baseline condition than in either Dynamic-VI baseline conditions because programmed reinforcer rates for the Static-VI condition were yoked to obtained rates from the first Dynamic-VI condition. Under interval schedules of reinforcement, unless subjects earn reinforcers as soon as they are made available, obtained reinforcer rates necessarily underestimate programmed rates. These differences, however, were small: Relative reinforcer rates did not exceed 1.16 in any case (see Table 4). A one-way, repeated-measures ANOVA was used to examine differences in obtained reinforcer rates between conditions. The main effect was not significant, $F(1.05, 66.10) = 0.82$, $p = .38$.

Table 4

Mean Reinforcer Rates, SEM, and Relative Reinforcer Rates for the Last Eight Sessions of Each Baseline Condition

	Condition						Relative Reinforcer Rates	
	Dynamic 1		Static		Dynamic 2			
Subject	Reinforcer/Hr	SEM	Reinforcer/Hr	SEM	Reinforcer/Hr	SEM	D_1/S	D_2/S
4275	55.50	10.21	51.25	1.13	52.50	10.57	1.08	1.02
4276	53.50	7.90	50.75	0.84	55.00	8.13	1.05	1.08
4274	49.75	10.75	46.50	0.98	54.00	11.14	1.07	1.16
1499	53.50	11.13	51.50	1.35	52.50	11.26	1.04	1.02
4278	53.75	11.39	50.75	1.25	53.25	11.37	1.06	1.05
4188	52.75	9.14	51.00	1.13	54.25	8.75	1.03	1.06
4189	53.75	10.26	52.25	1.10	53.25	11.22	1.03	1.02
4277	51.25	7.57	48.00	1.13	53.75	9.18	1.07	1.12

Note. All values rounded to the nearest 100th of a whole number.

Baseline Response Rates

Mean rates of responding (plus SEM) for each subject during the last eight sessions of each baseline condition are included in Table 5. These data, aggregated across subjects, are shown in Figure 4. Mean response rates for the first Dynamic-VI, Static-VI, and second Dynamic-VI baseline conditions were 70.93 (*SEM* = 4.07), 59.54 (*SEM* = 2.25), and 60.92 (*SEM* = 2.14) responses per min, respectively. A one-way, repeated-measures ANOVA was used to examine differences in response rates between conditions. The main effect was significant, $F(1.47, 92.73) = 7.74$, $p < .01$. Bonferroni-corrected pairwise comparisons based on estimated marginal means revealed response

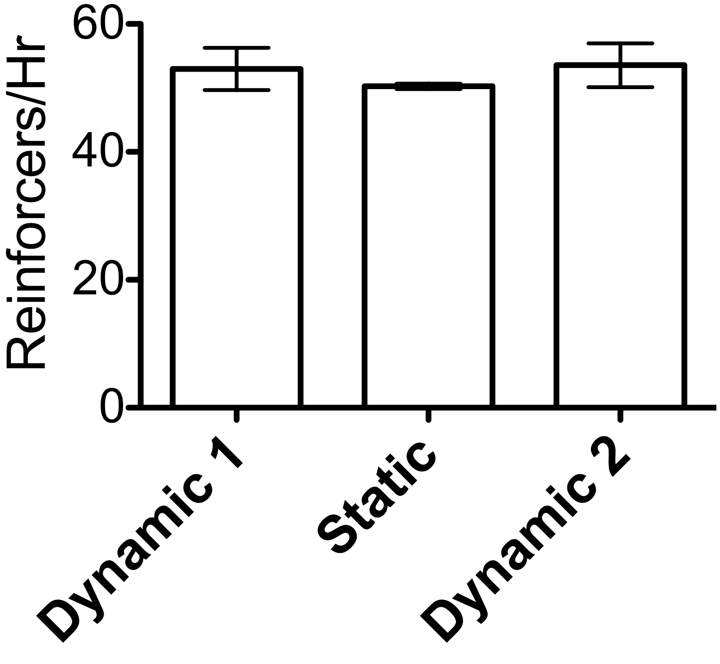

Figure 3. Mean group reinforcer rates during the last eight sessions of each condition. Error bars represent SEM.

Table 5

Mean Response Rates (and SEM) for the Last Eight Sessions of Each Baseline Condition

	Condition					
	Dynamic 1		Static		Dynamic 2	
Subject	Responses/Min	SEM	Responses/Min	SEM	Responses/Min	SEM
4275	81.54	2.15	46.27	2.75	78.93	3.70
4276	46.98	0.91	54.81	1.01	54.45	1.60
4274	32.75	3.29	32.67	0.59	35.12	2.63
1499	132.11	4.32	65.57	4.38	59.09	3.13
4278	65.86	1.82	47.73	1.08	52.35	1.53
4188	78.53	3.04	80.03	0.86	79.75	2.68
4189	93.86	3.56	86.94	2.05	78.80	2.80
4277	35.78	5.66	62.28	2.21	48.85	1.88

Note. All values were rounded to the nearest 100th of a whole number.

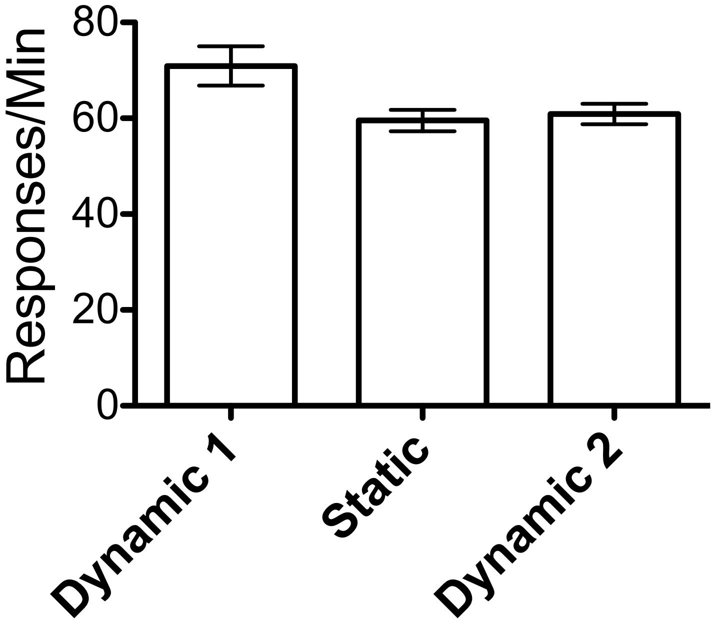

Figure 4. Mean group response rates during the last eight sessions of each condition. Error bars represent SEM.

rates from the first Dynamic-VI baseline condition were significantly higher than rates in both the Static-VI and second Dynamic-VI conditions. Response rates did not differ, however, between the latter two conditions.

Responding Under Dynamic-VI Schedules

Figures 5 and 6 show obtained reinforcer rates (triangles, plotted with respect to the right y-axis), and the rates of responding for each pigeon (circles, plotted with respect to the left y-axis), during the last eight sessions of the first and second Dynamic-VI baseline conditions, respectively. Rates of reinforcer delivery and responding for each session were divided into rates obtained prior to ("Before," filled data points) and following ("After," opened data points) within-session changes in reinforcer rates.

During the first Dynamic-VI baseline condition, rates of responding for five of eight subjects (Pigeons 4275, 4274, 4188, 4189, and 4277) were systematically related to the reinforcer rates currently arranged. That is, both across and within sessions, responding tended to occur more frequently in the presence of high reinforcer rates and less frequently in the presence of low reinforcer rates. For Pigeon 4276, responding tended to occur at relatively constant rates, regardless of reinforcer rate, and for Pigeons 1499 and 4278, responding was systematically higher early in a session (i.e., before mid-session changes in reinforcer rates) than it was later in a session (i.e., after mid-session changes). In the second Dynamic-VI baseline condition, response rates were systematically related to the current reinforcer rate for six of eight subjects (Pigeons 4275, 1499, 4278, 4188, 4189, and 4277). For Pigeons 4276 and 4274, responding was

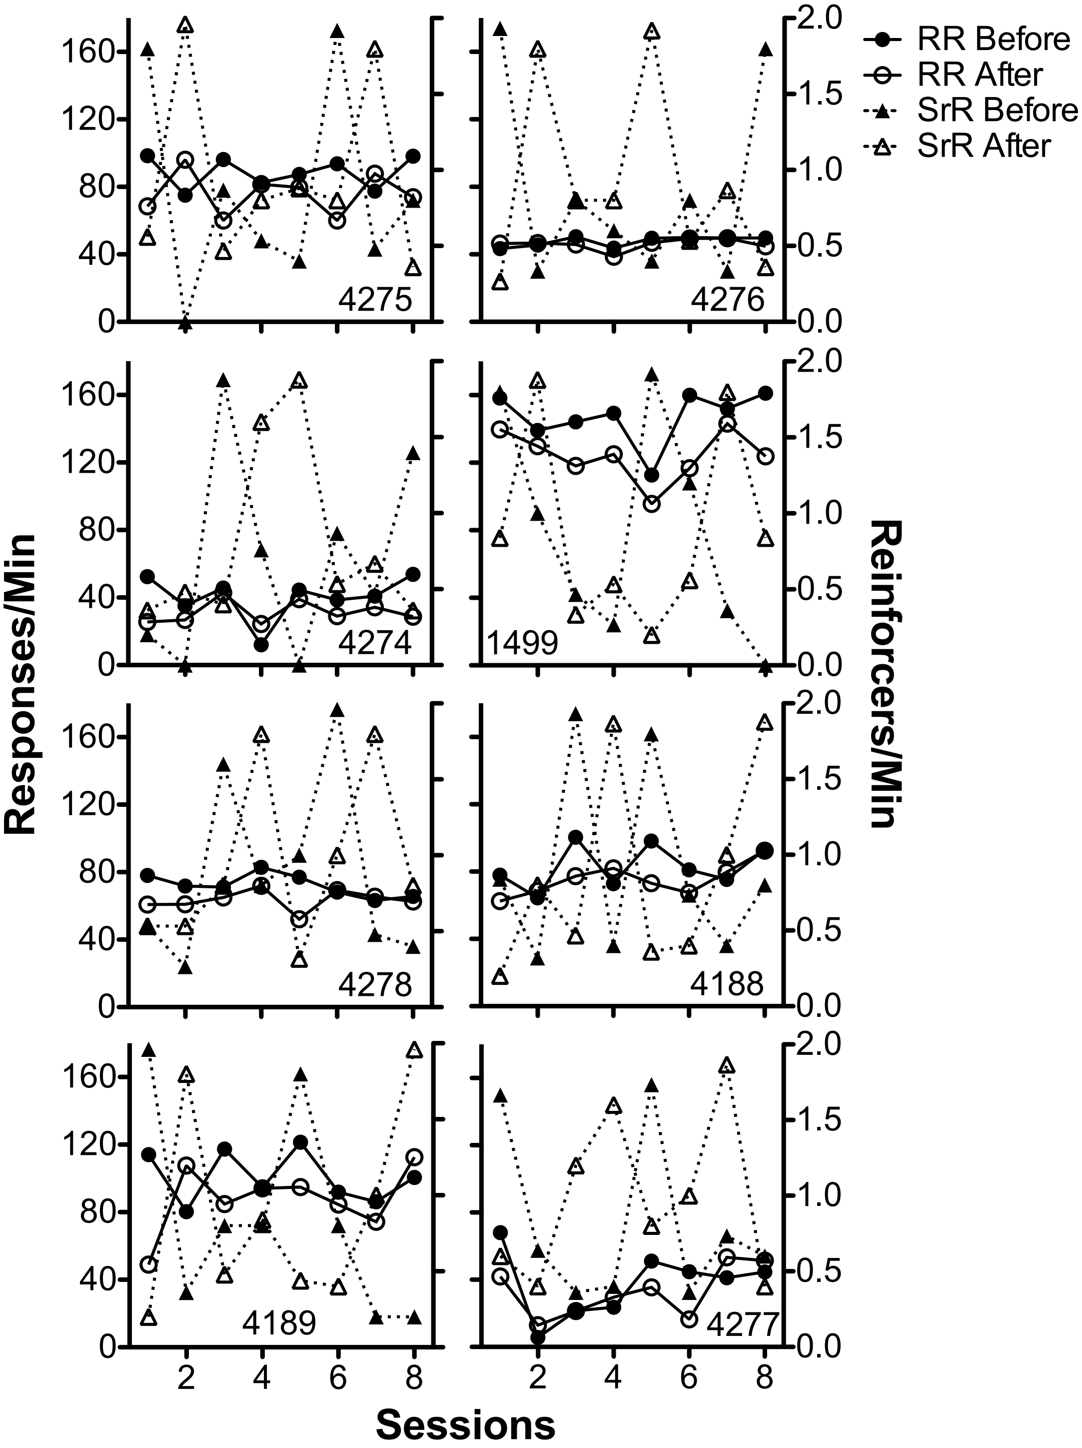

Figure 5. Response rates and reinforcer rates prior to ("Before") and following ("After") mid-sessions changes in contingencies during the last eight sessions of the first Dynamic-VI baseline condition for each subject.

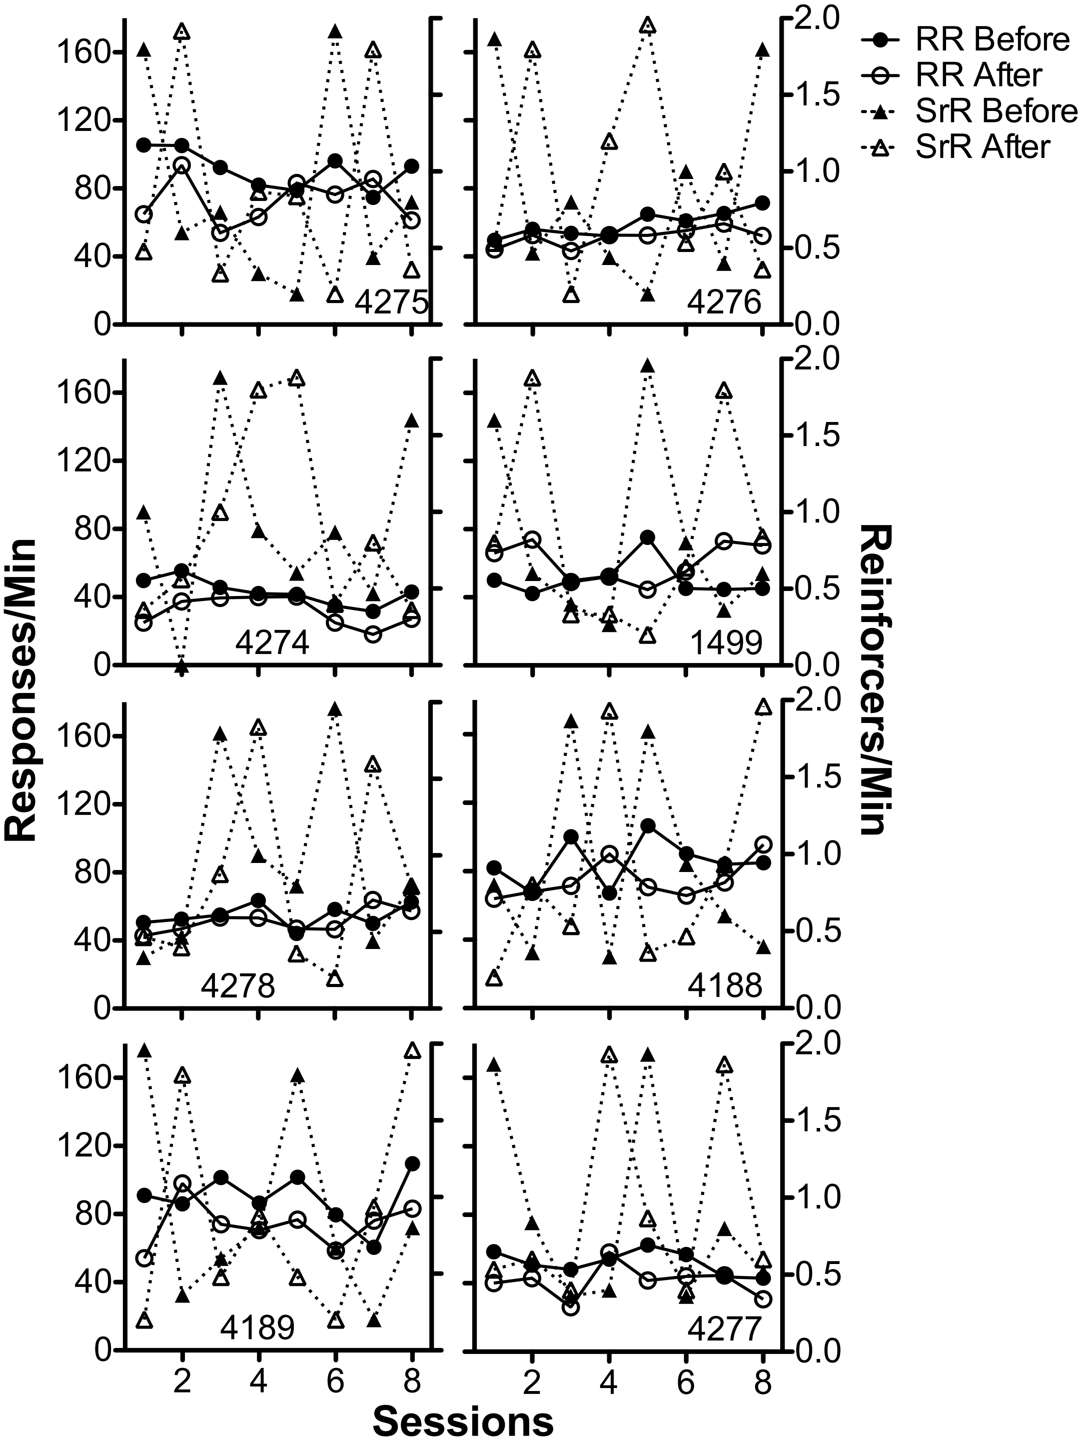

Figure 6. Response rates and reinforcer rates prior to ("Before") and following ("After") mid-sessions changes in contingencies during the last eight sessions of the second Dynamic-VI baseline condition for each subject.

systematically higher early in a session than it was later in a session.

Because response rates tended to co-vary with reinforcer rates for the majority of subjects in the final baseline sessions of both Dynamic-VI conditions, single-response matching analyses (one for each Dynamic-VI condition; see Herrnstein, 1970) were conducted to quantitatively describe these relations. For each subject, rates of responding and reinforcer rates in the presence of each individual VI (i.e., 30, 70, 110, and 150 s) schedule were averaged across the last eight sessions of each condition. Rates of responding (in responses per min) were plotted as a function of obtained reinforcer rates (in reinforcers per hr) under each schedule. The following equation then was fit to these data for each individual subject and for data aggregated across subjects using the non-linear curve-fitting function of Graphpad Prism:

$$B = \frac{kR}{R + R_e}. \tag{4}$$

Here, the parameters *B* and *R* are determined empirically and represent obtained rate of responding and reinforcer rate, respectively. The parameters *k* and R_e are free to vary and represent asymptotic response rate and rate of extraneous reinforcement (in units of experimentally derived reinforcers per hr).

Figures 7 and 8 show the results of these analyses for the first and second Dynamic-VI baseline conditions, respectively. Response- and reinforcer-rate data on which these analyses were based are included in Table 6. For the first and second Dynamic-VI baseline conditions, estimates of asymptotic response rates [*k*; *M* = 85.27 (*SEM* = 12.95) and 73.64 (*SEM* = 8.51), respectively], extraneous reinforcer rates [R_e; *M*

= 8.51 (*SEM* = 3.56) and 6.96 (*SEM* = 1.33), respectively], and R^2 [*M* = .57 (*SEM* = .11) and .76 (*SEM* = .11), respectively] varied between subjects. That is, response rates varied as a function of reinforcer rates more substantially for some subjects than for others.

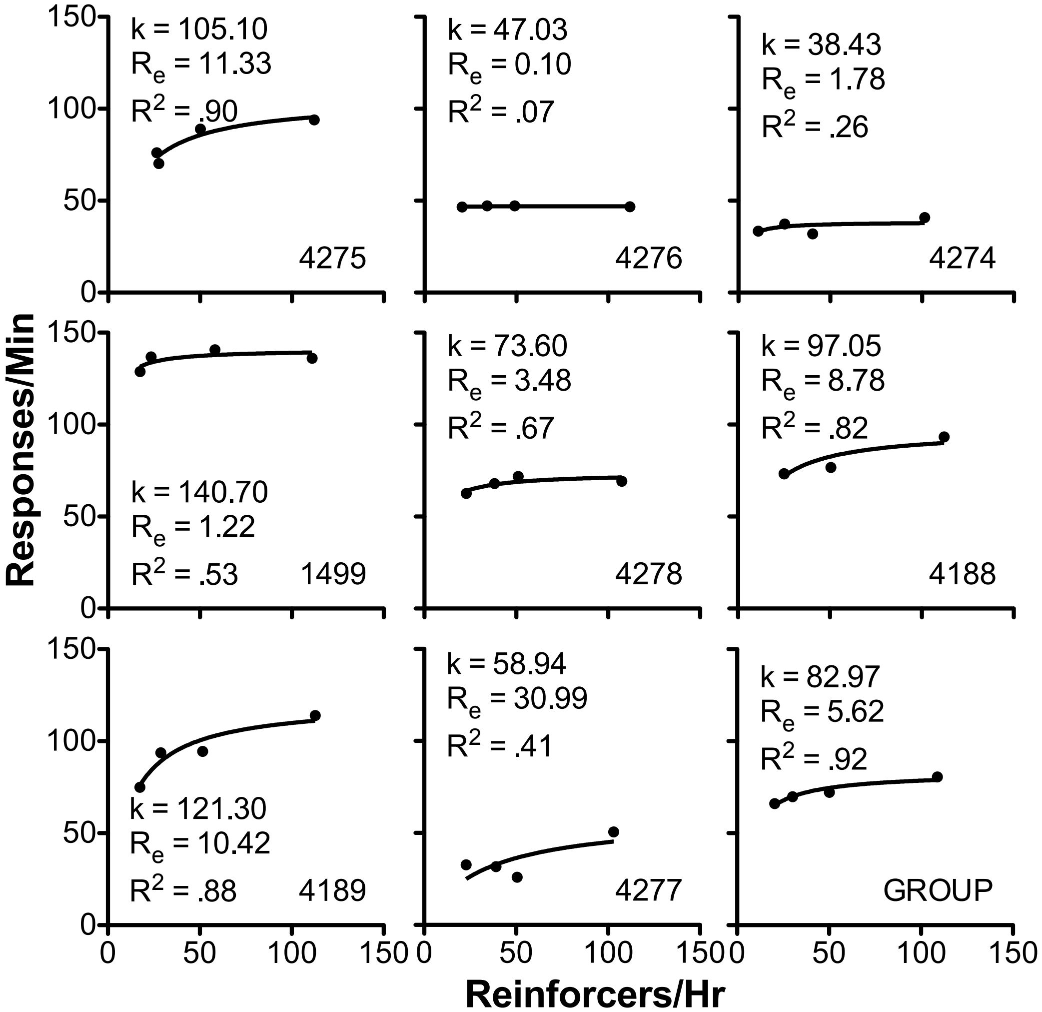

Figure 7. Response rates plotted as a function of obtained reinforcer rates for each VI schedule during the last eight sessions of the first Dynamic-VI baseline condition. Solid lines represent fits of Equation 4 to the data. Free-parameter and R^2 estimates are included on each panel.

Extinction

Calculating Proportion-of-Baseline Response Rates

As described above, response rates varied across sessions as a function of reinforcer rates during both Dynamic-VI baseline conditions for most subjects. To

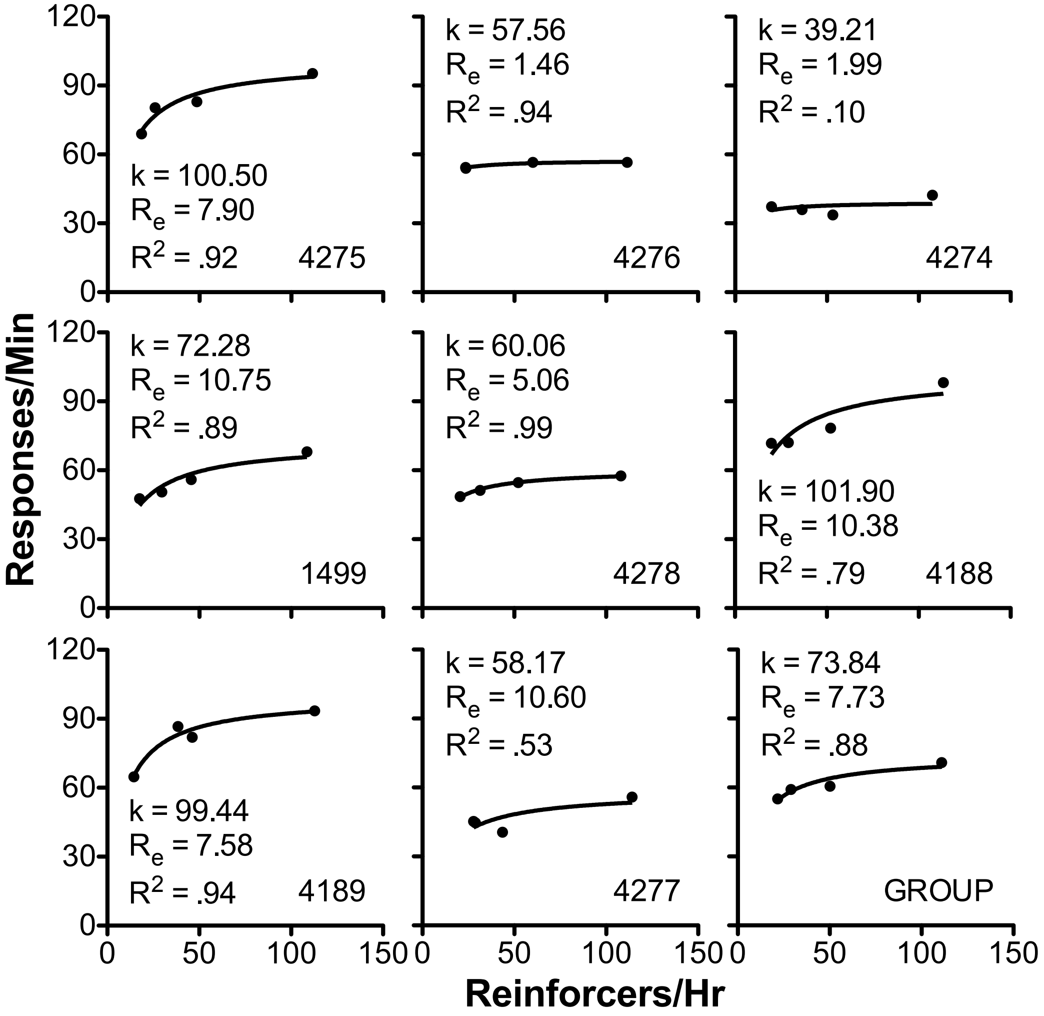

Figure 8. Response rates plotted as a function of obtained reinforcer rates for each VI schedule during the last eight sessions of the second Dynamic-VI baseline condition. Solid lines represent fits of Equation 4 to the data. Free-parameter and R^2 estimates are included on each panel.

Table 6

Response and Reinforcer Rates for Matching Analysis

		Responses/Min				Reinforcers/Hr			
Subject	D-VI BL	VI 30 s	VI 70 s	VI 110 s	VI 150 s	VI 30 s	VI 70 s	VI 110 s	VI 150 s
4275	1	93.98	88.90	70.21	76.01	112.20	50.20	27.60	26.40
	2	95.25	82.95	80.30	68.98	111.60	48.60	25.80	18.40
4276	1	46.68	47.18	47.15	46.58	111.80	49.00	34.00	20.40
	2	56.51	56.60	53.99	54.31	111.40	60.00	23.60	23.40
4274	1	40.75	31.95	37.30	33.49	101.40	40.40	25.20	10.80
	2	42.23	33.63	35.96	37.30	107.40	53.20	36.40	19.80
1499	1	136.03	140.67	136.81	128.75	111.00	58.20	23.40	17.40
	2	68.04	55.96	50.52	47.62	108.60	45.60	29.60	17.40
4278	1	69.15	71.85	67.92	62.59	107.40	51.00	38.20	22.80
	2	57.59	54.68	48.56	51.25	108.00	52.20	20.60	31.40
4188	1	93.25	76.74	73.25	73.24	112.20	50.60	25.00	25.20
	2	98.11	78.38	72.10	71.74	113.40	52.00	29.00	19.80
4189	1	113.91	94.43	93.73	74.99	112.80	51.60	28.80	17.40
	2	93.45	81.91	86.68	64.76	112.80	46.20	38.40	14.40
4277	1	50.53	25.92	31.64	32.66	103.00	50.60	39.00	22.80
	2	55.85	40.52	45.28	44.62	114.00	43.60	27.80	28.80
Group	1	80.54	72.20	69.75	66.04	108.98	50.20	30.15	20.40
	2	70.88	60.58	59.17	55.07	110.90	50.18	28.90	21.68

Note. All values were rounded to the nearest 100th of a whole number.

determine the extent to which these fluctuations in baseline response rates affected calculations of proportion-of-baseline response rates during extinction, proportion-of-baseline rates across the first 10 sessions of extinction were calculated for each subject using rates of responding from the last session of each baseline condition and the mean rates of responding from the last 8, 4, and 2 sessions of each condition. Three separate 4

X 11 (Method of Calculation X Session) repeated-measures ANOVA were used to compare proportion-of-baseline rates, one for each extinction condition. In each ANOVA, the main effect of Method was non-significant [for the first Dynamic-VI condition: *F*(3, 21) = 0.99, *p* = .42; for the Static-VI condition: *F*(1.05, 7.31) = 0.37, *p* = .57; and for the second Dynamic-VI condition: *F*(1.52, 10.70) = 0.33, *p* = .67]. Further, the Method X Day interaction from each ANOVA was non-significant [first Dynamic-VI condition: *F*(30, 210) = 1.29, *p* = .15; Static-VI condition: *F*(1.07, 7.54) = 0.64, *p* = .46; second Dynamic-VI condition: *F*(2.10, 14.68) = 0.44, *p* = .67]. Accordingly, proportion-of-baseline response rates across days of extinction were calculated using mean response rates from the last eight sessions of each condition.

Between-Condition Comparisons

Mean responses per min during the last eight sessions of each baseline condition and responses per min for each session of extinction for each subject are included in Figure 9. Mean proportion-of-baseline response rates from sessions 1-5 and 6-10 of extinction for each subject are included in Figure 10. In this figure, the first set of bars for each subject represent mean proportion-of-baseline response rates from the first five sessions of extinction, and the second set of bars represent these rates for the second five sessions. Absolute rates of responding and proportion-of-baseline response rates for the majority of subjects following the fifth session of extinction were reduced to low levels across conditions. A 3 X 2 (Condition X Session Block) repeated-measures ANOVA conducted on mean proportion-of-baseline response rates from the first and second five-session blocks of extinction revealed a non-significant main effect of Condition, *F*(1.2, 8.4) = 1.81, $p = .21$, a significant main effect of Session Block, $F(1, 7) = 56.98$, $p < .001$, and a significant Condition X Session Block interaction $F(2, 14) = 3.73$, $p < .05$. Consequently, only responding from days 0-5 of each extinction test is considered below.

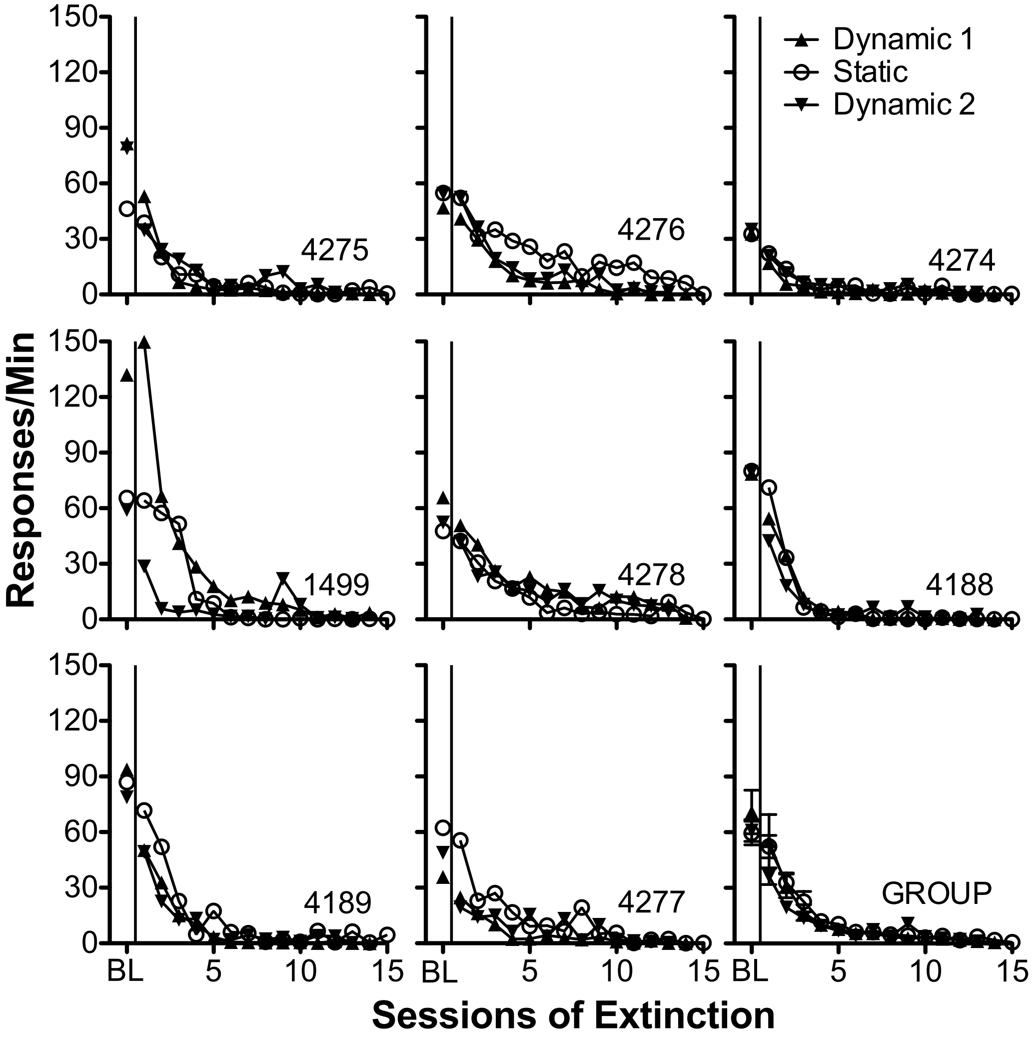

Figure 9. Mean rates of responding during the last eight sessions of baseline and absolute rates of responding from each session of extinction for each subject. Error bars for the group function represent SEM.

Proportion-of-baseline response rates across the first five sessions of extinction testing for each subject in each condition are displayed in Figure 11. With few exceptions, resistance to extinction initially was greater following the Static-VI baseline condition than following the first Dynamic-VI condition (for Pigeon 1499, responding was more persistent in the first session of extinction following the first Dynamic-VI condition than the Static-VI condition). This patterning persisted across the first five sessions of extinction for six of eight subjects. For Pigeons 4278 and 4188, proportion-of-baseline responding was largely undifferentiated between conditions. Extinction functions following the first and second Dynamic-VI conditions were not systematically different.

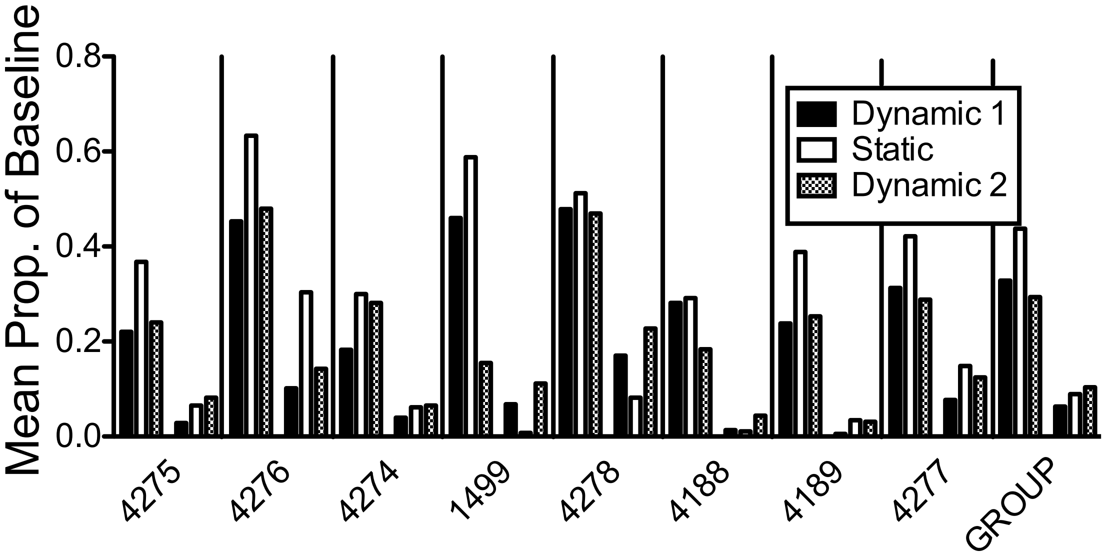

Figure 10. Mean proportion-of-baseline response rates from sessions 1-5 (left-most bars in each panel) and 6-10 (right-most bars in each panel) of extinction for each subject. Aggregated group data also are included.

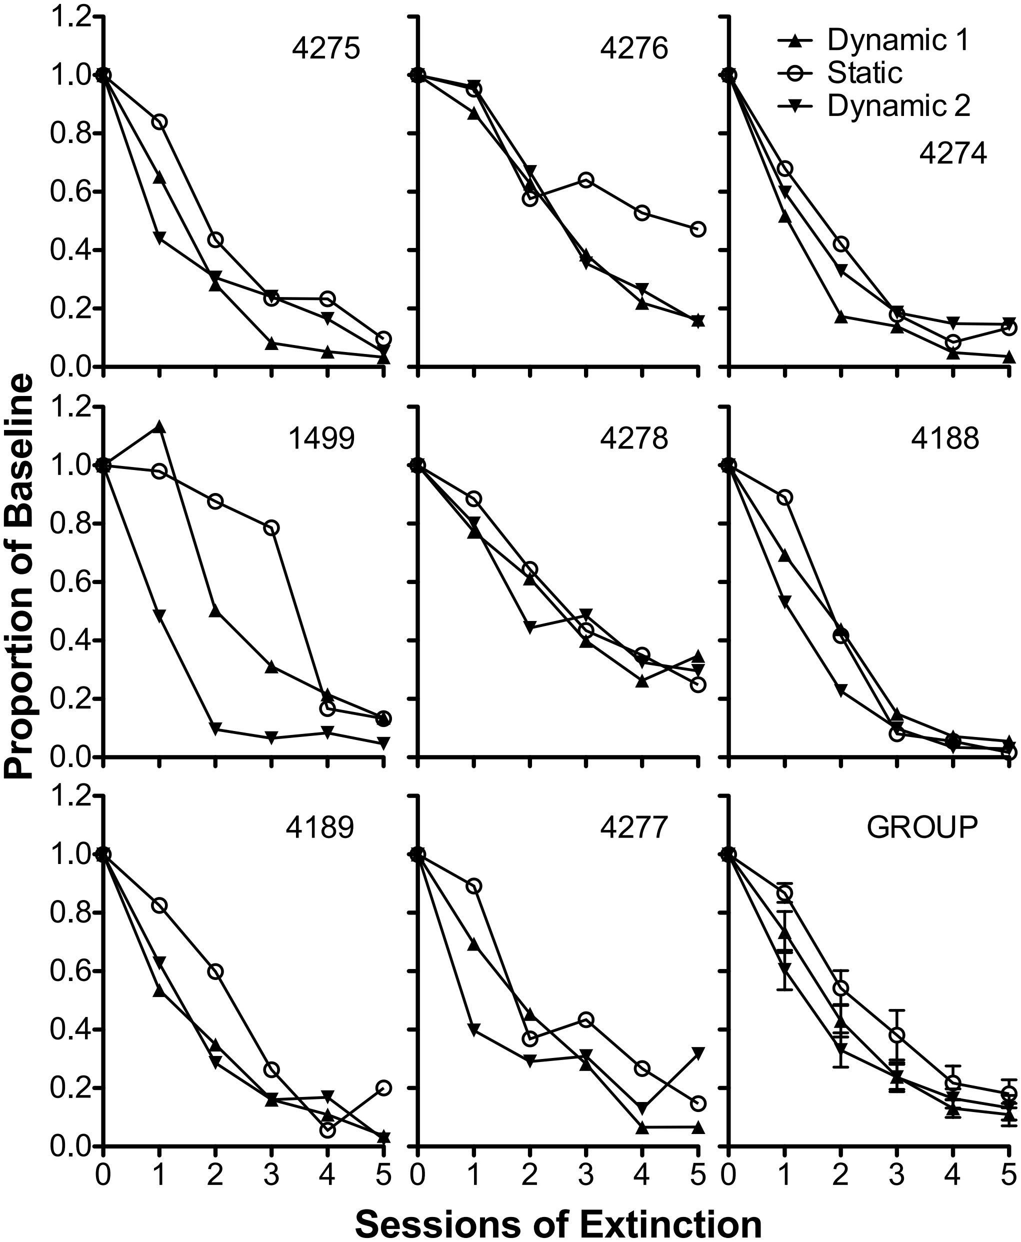

Figure 11. Proportion-of-baseline response rates across the first five sessions of extinction in each condition and for each subject. For the group function, error bars represent SEM.

A 3 X 6 (Condition X Day) repeated-measures ANOVA revealed a significant main effect of condition, $F(2, 14) = 7.89$, $p < .01$, a significant main effect of Day, $F(5, 35) = 248.09$, $p < .001$, and a significant Condition X Day interaction , $F(10, 70) = 2.40$, $p < .05$. Bonferroni-corrected pairwise comparisons based on estimated marginal means for Condition revealed that resistance to extinction following the Static-VI condition was significantly higher than resistance to extinction following both Dynamic-VI conditions, but resistance to extinction following the first and second Dynamic-VI conditions did not differ.

Relating Extinction to Other Training Variables

As described previously, it is unclear from the perspective of behavioral momentum theory whether resistance to change depends on overall reinforcer rates or only on more recently experienced rates. Because the Dynamic-VI baseline conditions arranged reinforcer rates that were regular across blocks of sessions but varied widely across sessions, correlation analyses were conducted to determine if relative resistance to extinction between the Dynamic- and Static-VI conditions was related to the relative reinforcer rates each pigeon experienced during the final session (where between-subject variability in rate was maximal in Dynamic-VI conditions), and the last eight sessions, of these conditions. Here, relative resistance to extinction was defined as the log ratio of mean proportion-of-baseline response rates from the first five sessions of Dynamic-VI and Static-VI extinction (cf., Grace & Nevin, 1997). Relative resistance to extinction was calculated separately for the first and second Dynamic-VI conditions. These values are plotted against log relative reinforcer rates (Dynamic/Static) in Figure 12. Relative resistance-to-extinction and reinforcer-rate values are included in Table 7.

The left panel of Figure 12 shows log mean proportion-of-baseline response rates plotted against log relative reinforcer rates from the last session of each baseline condition, and the right panel shows log mean proportion-of-baseline response rates plotted against log relative reinforcer rates using mean rates from the last eight sessions of each condition. Data points falling below the horizontal dashed line represent subjects for whom resistance to extinction was lower in the Dynamic-VI condition than in the Static-VI condition, and points that fall to the left of the vertical dashed line represent

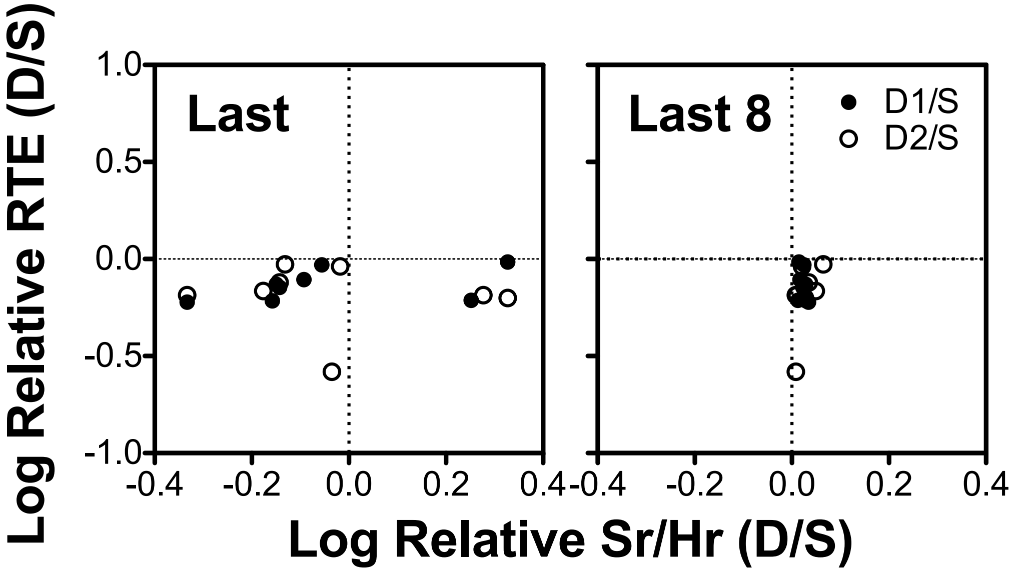

Figure 12. Relative resistance to extinction (Dynamic/Static) for each subject plotted as a function of log relative reinforcer rates (Dynamic/Static). Log relative resistance to extinction and reinforcer rates are displayed for both Dynamic-VI condition (D1/S and D2/S). The left panel shows log relative reinforcer rates from the last session of each baseline condition, and the right panel shows log relative reinforcer rates using mean rates from the last eight sessions of each condition. Dashed lines represent the point at which resistance to extinction (horizontal) or reinforcer rates (vertical) were equal between conditions.

Table 7

Log Relative Reinforcer-Rate and Log Relative Resistance-to-Extinction Values

	Log Relative Sr Rates		Log Relative Res. to EXT	
Subject	D_1/S	D_2/S	D_1/S	D_2/S
4275	-0.33	0.01	-0.22	-0.19
4276	-0.14	0.04	-0.15	-0.12
4274	-0.16	0.06	-0.21	-0.03
1499	-0.09	0.01	-0.11	-0.58
4278	-0.06	0.02	-0.03	-0.04
4188	0.33	0.03	-0.01	-0.20
4189	0.25	0.01	-0.21	-0.19
4277	-0.16	0.05	-0.13	-0.16

Note. All values were rounded to the nearest 100th of a whole number.

subjects for whom reinforcer rates were lower in the Dynamic-VI condition. In neither case were log mean proportion-of-baseline response rates significantly correlated with log relative reinforcer rates [for the left panel, $r(22) = .05$, $p = .85$; for the right panel, $r(22) = .35$, $p = .18$].

Response rates tended to be higher during the first Dynamic-VI baseline than in the other baseline conditions. Correlation analyses were conducted to examine the extent to which relative resistance to extinction between both Dynamic-VI and the Static-VI conditions was related to log relative baseline response rates from these conditions. Relative resistance to extinction was defined here as above. These values were plotted as a function of log relative response rates from the last eight sessions of baseline (Dynamic/Static; see Table 8 for relative response-rate values). Log relative response rates were calculated separately for the first and second Dynamic-VI conditions. These

Table 8

Log Relative Response-Rate Values

	Log Relative Resp./Min	
Subject	D_1/S	D_2/S
4275	0.25	0.23
4276	-0.07	-0.01
4274	0.01	0.03
1499	0.30	-0.05
4278	0.14	0.04
4188	-0.01	-0.01
4189	0.03	-0.04
4277	-0.24	-0.11

Note. All values were rounded to the nearest 100th of a whole number.

data are shown in Figure 13. Data points falling below the horizontal dashed line represent subjects for whom resistance to extinction was lower in the Dynamic-VI condition than in the Static-VI condition, and points that fall to the left of the vertical dashed line represent subjects for whom response rates were lower in the Dynamic-VI condition. No significant relation was present between log relative baseline response rates and log relative resistance to extinction, $r(6) = .12$, $p = .67$.

Model Fits

Equation 3 was fit to obtained proportion-of-baseline response rates for each subject, and to aggregated group proportion-of-baseline rates, across the first five sessions of extinction following each condition. Fits were conducted separately for each condition using Microsoft Excel Solver, a non-linear curve-fitting program. Free-parameter values (*c*, *d*, and *b*), reinforcer-rate parameters (*r*), and R^2 values for each fit may be found in Table 9.

As may be seen from the table, values of *c* (for the first Dynamic-VI, Static-VI, and second Dynamic-VI conditions, respectively, range: 0.98-1.00, 0.95-1.00, and 0.98-1.00) and *b* (range: 0.5-0.68, 0.41-0.68, and 0.48-0.68) were relatively constant across conditions. These values were similar to those values typically obtained from fits of Equation 3 to extinction data from multiple-schedule preparations (i.e., $c = 1$, $b = 0.5$; see Nevin & Grace, 2000). Values of *d*, however, tended to vary systematically between

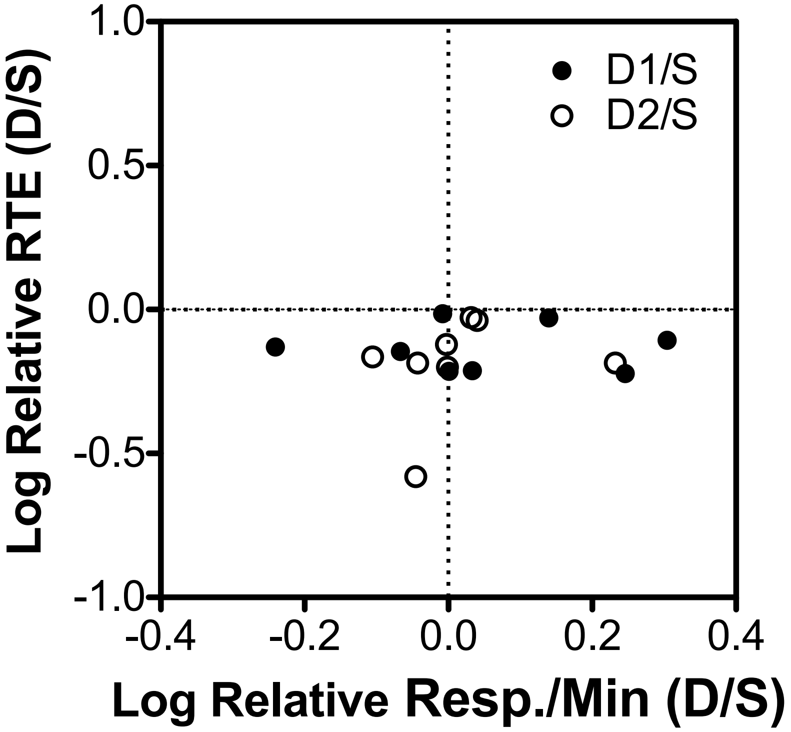

Figure 13. Log relative resistance to extinction (Dynamic/Static) plotted as a function of log relative response rates (Dynamic/Static). Log relative resistance to extinction and response rates are displayed for both Dynamic-VI condition (D1/S and D2/S). Dashed lines represent the point at which resistance to extinction (horizontal) or response rates (vertical) were equal between conditions.

Table 9

Parameter Values from Fits of Equation 3 to Extinction Data

Subject	Condition	c	d	b	r	R^2
4275	Dynamic 1	1.00	0.056	0.68	55.50	0.97
	Static	1.00	0.005	0.51	51.25	0.97
	Dynamic 2	1.00	0.041	0.63	52.50	0.97
4276	Dynamic 1	1.00	0.000	0.51	53.50	0.98
	Static	0.95	0.001	0.68	50.75	0.99
	Dynamic 2	0.99	0.001	0.53	55.00	0.96
4274	Dynamic 1	1.00	0.067	0.67	49.75	0.99
	Static	1.00	0.008	0.50	46.50	0.99
	Dynamic 2	1.00	0.028	0.60	54.00	0.99
1499	Dynamic 1	1.00	0.000	0.50	53.50	0.89
	Static	0.97	0.000	0.60	51.50	0.89
	Dynamic 2	1.00	0.061	0.61	52.50	0.97
4278	Dynamic 1	0.98	0.002	0.56	53.75	0.99
	Static	0.98	0.000	0.56	50.75	0.99
	Dynamic 2	0.98	0.004	0.58	53.25	0.99
4188	Dynamic 1	1.00	0.024	0.59	52.75	0.98
	Static	1.00	0.001	0.41	51.00	0.88
	Dynamic 2	1.00	0.068	0.68	54.25	0.99
4189	Dynamic 1	1.00	0.050	0.67	53.75	0.99
	Static	1.00	0.004	0.51	52.25	0.95
	Dynamic 2	1.00	0.035	0.62	53.25	0.99
4277	Dynamic 1	1.00	0.010	0.52	51.25	0.98
	Static	0.99	0.002	0.52	48.00	0.96
	Dynamic 2	1.00	0.011	0.48	53.75	0.88
Group	Dynamic 1	1.00	0.010	0.53	52.97	0.99
	Static	1.00	0.000	0.51	50.25	0.99
	Dynamic 2	1.00	0.008	0.48	53.56	0.99

Note. Values of *c* and *b* were rounded to the nearest 100th of a whole number. Values of *d* were rounded to the nearest 1,000th.

conditions. More specifically, for six of eight subjects, values of *d* were higher for fits to extinction functions from the first and second Dynamic-VI conditions (respective, range: 0-0.067 and 0.001-0.068) than for the Static-VI condition (range: 0-0.008). Values of *d* for the Dynamic-VI conditions, but not for the Static-VI condition, were substantially higher than those typically reported (i.e., $d = 0.001$).

CHAPTER VI

DISCUSSION

Operant extinction necessarily is, at least in part, a change-detection process. Presumably, changes in behavior follow changes in reinforcement contingencies if and only if (or when and only when) the organism discriminates that a change occurred. The Bayesian approach to understanding change-detection processes in the context of operant behavior states that organisms attend both to the rate at which reinforcers are delivered in their environment and to the frequency with which changes in rate occur (Gallistel, 2012; see also Gallistel et al., 2001). From this perspective, information about reinforcer rate is necessary to detect changes in rate, and information about the frequency with which changes in reinforcer rate occur governs the rapidity with which organisms adapt to changing contingencies. In essence, when environments are stable and change is not to be expected, the organism does not attend to those attributes of the environment that are necessary for rapid adaptation to change. If change is inherent, however, the organism might become attuned to detecting and adapting to change.

The Bayesian account of extinction performance may be contrasted to the quantitative account of extinction offered by behavioral momentum theory (Equations 2 and 3; see Nevin & Grace, 2000). This account states that resistance to extinction is a function of the Pavlovian stimulus-reinforcer relation present in an experimental context. In isolation of change-detection processes, this account would not predict a relation between variability in baseline reinforcer rates and resistance to extinction. Instead, momentum theory considers overall, mean baseline reinforcer rate to be the major determinant of resistance to extinction. Based on these insights, the purpose of the current experiment was to determine the extent to which recent experience with frequently changing (relative to experience with infrequently changing) reinforcer rates affects the speed of operant extinction when overall rates are held constant.

To this end, the Dynamic and Static-VI schedules used in the current experiment, which delivered equal overall reinforcer rates across the final sessions of baseline while introducing differential variability in local reinforcer rates within and between sessions, generated differential resistance to extinction. These data provide support for Gallistel's (2012) approach to understanding operant extinction from the perspective of Bayesian change detection. There are, however, a number of challenges inherent both in the procedures used in the current experiment and the data derived from this experiment that could complicate this interpretation. These challenges are discussed below.

Limitations, Alternative Explanations, and Rebuttals

First, baseline reinforcer rates were consistently lower in the Static-VI condition than in either Dynamic-VI conditions for each subject (see Table 4). In single schedules, resistance to extinction is negatively related to baseline reinforcer rates (see Cohen, 1998; Shull & Grimes, 2005). Thus, these consistent differences in baseline reinforcer rates could have contributed to the differences in resistance to extinction observed in the current experiment. Differences in reinforcer rates between conditions, however, were small relative to the three- or more-fold differences in reinforcer rates typically used in studies of resistance to change in multiple schedules (for review, see Nevin, 1992a; 2002) and single schedule (see Cohen, 1998; Shull & Grimes, 2006). It therefore is unlikely that these differences in reinforcer rate, alone, produced the differential resistance to extinction observed in the current experiment.

Baseline response rates also differed between conditions. Specifically, response rates were significantly higher in the first Dynamic-VI baseline condition than in the Static-VI baseline condition (see Figure 4). Some evidence suggests that, when reinforcer rates are held constant between separate stimulus contexts (as in the current experiment), relatively low response rates tend to be more persistent in the face of disruption than do relatively high response rates. For example, Lattal (1989) trained pigeons to peck keys for food in a two-component multiple schedule where one component was associated with a tandem VI FR schedule, and the other component was associated with a tandem VI differential-reinforcement-of-low-rate (DRL) schedule. These tandem schedules produced high- and low-rate responding, respectively, but reinforcer rates were held constant between components. When responding was disrupted by presenting free ICI food, behavior in the low-rate (tandem VI DRL) component was more persistent than responding in the high-rate (tandem VI FR) component. Nevin, Grace, Holland, and McLean (2001) conducted a similar experiment in which pigeons pecked keys for food in a two-component multiple schedule where one component was associated with a variable-ratio (VR) schedule and the other component was associated with a VI schedule. The VR component produced relatively high rates of responding, and the VI component produced relatively low rates of responding. Both schedules, however, arranged equal reinforcer rates. When behavior was disrupted by extinction, free ICI food, or both extinction and ICI food, behavior in the low-rate (VI) component was more persistent than behavior in the high-rate (VR) component.

Unlike Lattal (1989) and Nevin et al. (2001), no relation was present between relative baseline response rates and relative resistance to extinction in the present experiment (see Figure 13). In both the Lattal and Nevin et al. studies, however, differential rates of responding between the components of their multiple schedules were produced by schedules of reinforcement that placed different contingencies on inter-response times (IRTs): Under ratio schedules, IRTs are positively related to reinforcer rates, while under interval-based (VI or DRL) schedules, IRTs are either relatively unrelated to (VI) or negatively related to (DRL) reinforcer rates (see Nevin et al. for discussion). In contrast, only VI schedules were used in the current experiment, and differences in response rates between conditions presumably were produced by differential variability in local reinforcer rates. It might be the case that the negative relation between response rates and resistance to change demonstrated by Lattal and by Nevin et al. is limited to those situations where contingencies on long or short IRTs are programmed overtly.

Alternatively, it might be the case that baseline response rates were not sufficiently different between the conditions of the current experiment to produce statistically meaningful differences in resistance to extinction. On the level of the individual subject, however, the current data suggest that this was not the case. For example, two subjects (Pigeons 4276 and 4277) demonstrated substantially higher response rates in the Static-VI condition than in the first Dynamic-VI condition (see

Table 5). Extinction functions for these subjects were quantitatively more differentiated than those for roughly half of the other subjects (see Table 7 for relative resistance-to-extinction values), contrary to what one might expect if any meaningful relation existed between baseline response rates and resistance to extinction.

A third detail of the current experiment that might have limited these findings is the limited range of programmed reinforcer rates that comprised the Dynamic-VI schedules. This range was 24 (VI 150 s)-120 (VI 30 s) reinforcers per hr. These rates represent a five-fold difference between the richest and leanest schedules experienced. Gallistel et al. (2001), who demonstrated that frequent changes in relative reinforcer rates produce rapid adaptation of rats' relative response allocation to changing contingencies in choice situations, used a 10-fold range of reinforcer rates. Indeed, Gallistel et al. and Gallistel (2012) assert that detection of a change in reinforcer rates depends on the size of the change that occurs. That is, in principle it should be easier to detect a change in reinforcer rate if the change is large than if the change is small.

The range of reinforcer rates used for the Dynamic-VI baseline conditions in the current experiment was limited for practical reasons (i.e., to avoid the pigeons eating excessive amounts of food within sessions to maintain constant levels of motivation throughout each condition). Despite the restricted range in Dynamic-VI reinforcer rates, however, response rates for most subjects varied systematically with current reinforcer rates in these conditions (see Figures 7 and 8). This finding provides evidence that the pigeons detected changes in reinforcer rates during Dynamic-VI baseline conditions and that they adjusted their behavior within and between sessions in response to these changes. Larger degrees of variability in reinforcer rates might further have promoted change detection and, possibly, further hastened subsequent extinction. The extent to which this prediction is true is unknown and is a direction for future investigation.

Finally, this experiment used a fixed ABCBAB design to investigate the extent to which dynamic, relative to static, reinforcer rates affected resistance to extinction. Repeated extinction tests routinely result in fewer responses across exposures to extinction, especially early in successive extinction conditions (e.g., Anger & Anger, 1976; Bullock, 1960; Clark, 1964; Davenport, 1969; Jenkins, 1961). Further, Baum (2012) recently showed that temporally proximal contingencies other than those presently enforced can have substantial effects on responding during extinction. In Baum's experiment, pigeons pecked keys for food according to VI schedules where the value of the VI ranged from 0 s (i.e., FR 1) to 1200 s. Periods of food presentation and extinction tests alternated within sessions. The VI first increased across conditions from the richest schedule to the leanest schedule. Then, the schedules were experienced in the opposite order such that each condition was experienced twice. The number of pecks during extinction across schedules was substantially higher for the decreasing progression than for the initial increasing progression in most schedule pairs.

Proportion-of-baseline response rates did not differ statistically between the first and second Dynamic-VI conditions, where the pigeons experienced similar dynamic reinforcer rates. Visual inspection of individual subjects' extinction functions (see Figure 11) also reveals no systematic differences in extinction following these conditions. On the group level, however, resistance tended to be lower in the second Dynamic-VI condition. These data suggest that repeated extinction tests did not contribute to the higher resistance to extinction produced by static reinforcer rates. To the contrary, repeated extinction tests might have contributed to decreases in persistence across tests, thereby decreasing the differences in extinction functions between the first Dynamic-VI condition and the Static-VI condition.

All investigations of resistance to extinction from the perspective of momentum theory have assessed extinction following periods of stable reinforcer rates during baseline (see Nevin, 2012b; Nevin & Grace, 2000). To the contrary, the present experiment assessed resistance to extinction following periods of rapidly changing reinforcer rates. Thus, one alternative explanation for the current finding might be that, instead of overall reinforcer rates across a baseline condition, response strength and thereby resistance to extinction is governed by more recently experienced reinforcer rates. Indeed, the Devenports and their colleagues (e.g., Devenport & Devenport, 1993, 1994; Devenport, Hill, Wilson, & Ogden, 1997; Devenport, Patterson, & Devenport, 2005) have demonstrated that, in foraging situations, various species tend to rely most heavily on recent information when making choices between foraging options.

When relative reinforcer rates from the last session of each baseline condition (where maximal variation in rates between subjects occurred in Dynamic-VI conditions) were correlated with relative resistance to extinction (see Figure 12), no relation was observed. These data suggest that resistance to extinction in the present experiment did not depend on recently experienced reinforcer rates. It is possible, though, that this apparent lack of dependency could have resulted from a relatively small sample size, the relatively restricted range over which reinforcer rates varied in the final session of each Dynamic-VI baseline condition, or both. The method by which organisms incorporate information about past reinforcement into one assignment of value to a stimulus context warrants further investigation.

Implications for Momentum Theory

From the perspective of behavioral momentum theory, it is conceptually unclear how the effects of variability in local reinforcer rates should affect resistance to extinction. Equations 2 and 3 do not have dedicated mechanisms to account for these effects. When either of these equations is fit to extinction data from multiple-schedule preparations, the free parameters of the model are held constant across components. The only parameters that differ between components are the reinforcer-rate terms (r) present in the numerator and the denominator of the model (see Nevin & Grace, 2000; see also Nevin, 2012b). Holding these parameters constant assumes that, within subjects, the impact of suspending the response-reinforcer contingency, the scalar impact of generalization decrement, and sensitivity to baseline reinforcer rates are held constant between multiple-schedule components. When fit to extinction data from single schedules, similar assumptions are made (cf., Shull & Grimes, 2006). If the reinforcer-rate parameter, r, too is held constant, the augmented model of extinction cannot predict differential resistance to extinction. In the case that some treatment differed between schedule exposures or stimulus contexts, and to the extent that this treatment could, in principle, differentially impact some fundamental extinction process between components (as might be the case in the current experiment), it is justifiable to assume that at least one of the behaviorally relevant parameters in Equation 3 should differ between conditions. Accordingly, individual, condition-by-condition, model fits were used for the current data.

The parameter values derived from these model fits to the extinction data revealed consistent parametric effects of dynamic- versus static-VI contingencies. As may be seen in Table 9, the *c* parameter, which represents the impact of suspending the response-reinforcer contingency on responding during extinction, and *b* parameter, which represents sensitivity to baseline reinforcer rates, were relatively undifferentiated between conditions. The *d* parameter, which scales the impact of generalization decrement (i.e., the change in reinforcer rates between baseline and extinction), however, was systematically higher for six of eight subjects following the Dynamic-VI condition than following the Static-VI condition. Assuming change detection as an underlying mechanism of operant extinction, the effects of dynamic contingencies, relative to static contingencies, on the parameters of Equation 3 are intuitive. If experience with changing reinforcer rates promotes change detection as Gallistel (2012) and Gallistel et al. (2001) suggest, it is reasonable for this to be reflected in a greater impact on responding of changes in reinforcer rates to zero during extinction.

The findings of the current experiment, then, are not necessarily inconsistent with the theoretical underpinnings of behavioral momentum. Variation in those parameters posited to be behaviorally relevant by Equation 3 accounted for much of the variance present in the subjects' data. Instead, these results might serve to extend momentum theory by suggesting an additional mechanism through which these behaviorally relevant factors influence responding during extinction. The present model fits demonstrate that, from the perspective of behavioral momentum, change detection might act as a mediating mechanism (or one of various mediating mechanisms) through with generalization decrement exerts its effect on responding during extinction.

It is important to note that both the Bayesian and behavioral-momentum based approaches to understanding extinction performance, when considered alone, are limited. Behavioral momentum theory cannot easily account for resistance to extinction in single schedules, while the Bayesian approach to understanding extinction can. Likewise, the Bayesian approach cannot easily account for resistance to extinction in multiple schedules, while behavioral momentum theory can.

Gallistel (2012) suggested one possible explanation for the greater persistence occasioned by relatively high reinforcer rates in multiple schedules. According to Gallistel's hypothesis, in extinction it is in the interest of the organism to continue to sample from response options in the case that reinforcer availability is re-introduced. Further, sampling rate should be roughly proportional to previously experienced reinforcer rates. Thus, despite the enhanced detectability of transitions to extinction following relatively rich reinforcer rates, other behavioral processes might contribute to momentum-like effects. As Nevin (2012a) described, however, this interpretation necessarily is either incorrect or incomplete because it does not capture the negative relation between baseline reinforcer rates and resistance to extinction in single schedules.

That is, it is unclear why differential sampling in extinction should occur in single and multiple schedules.

It is possible that change-detection and response-strengthening processes both contribute to extinction (see Nevin, 2012a) and that the interplay between these processes differs depending on the preparation in which extinction is assessed. Based on this interpretation, response strength would be the major determinant of resistance to extinction when discriminated operant behavior (i.e., in a multiple schedule) is disrupted. When non-discriminated responding (i.e., in a single schedule) is disrupted, change-detection processes would have a greater influence over persistence. This interpretation is attractive in that it reconciles those discrepant findings produced when extinction is assessed in single- versus multiple-schedule contexts. If both of these proposed facets of operant extinction contribute to extinction performance, however, it is unclear how they interact and why one schedule context evokes one process more than the other.

Multiple schedules present complicated sources of information to the organism. It is possible that behavioral-momentum effects during extinction in multiple schedules are the manifestation of the organism using this information in equally complicated and perhaps unexpected ways to navigate the stochasticity of its environment. In multiple schedules, response persistence depends not only on reinforcer rates within a component but also on overall reinforcer rates between components (see Nevin, 1992b). Thus, the sources of information provided by reinforcer rates within multiple-schedule components likely are not considered as entirely separate sources of information by the organism. Instead, reinforcer rates in either component might contribute mutual information to the organism's overall assessment of reinforcer availability. If, for example, the organism detects that reinforcer presentations have ceased in one component during extinction, the strength of the evidence that presentations have ceased in the other component might not need to be as strong for the organism to conclude that extinction is in effect. Returning to Gallistel's (2012) argument, the organism might then sample frequently in the rich-schedule context because it also serves as a relatively rich source of information about the global state of reinforcer availability.

The data from the current experiment are not sufficient to discriminate between these theoretical possibilities. These results do, however, suggest that change detection is one potential mechanism that contributes to operant extinction. Further, these data suggest that the ability of organisms to detect changes in reinforcer rates can be experimentally manipulated by presenting periods during which reinforcer rates change rapidly (cf., Gallistel et al., 2001). Further work will be necessary to determine precisely how organisms navigate the temporal dynamics of reinforcer availability inherent in operant extinction.

CHAPTER VII

CONCLUSIONS

The results of the current study provide some tentative evidence that experience with changing reinforcer rates affect an organism's ability to detect, and subsequently adapt to, the changes in reinforcer rates associated with transitions to extinction. This evidence is consistent with the predictions of the Bayesian approach to understanding operant extinction in that exposure to rapidly changing reinforcer rates, relative to experience with static rates, hasted extinction. These data, however, might not contradict the assertions of behavioral momentum theory if change detection is considered a sieve through which the impact of generalization decrement affects behavior in extinction.

Though behavioral momentum theory cannot easily account for resistance to extinction following single schedules of reinforcement (cf., Cohen, 1998; Shull & Grimes, 2006), the Bayesian approach to understanding extinction performance also is limited. This approach states that change detection is dependent on the size of the change that occurs. Accordingly, the Bayesian change-detection algorithm, alone, cannot account for the momentum-like effects typically observed in multiple schedules. Future theoretical work will be necessary to reconcile the limitations and scopes of these two approaches to understanding the extinction process. It might be conceptually generative, however, to consider that these approaches to understanding extinction performance might not be mutually exclusive.

REFERENCES

Ahearn, W. H., Clark, K. M., Gardenier, N. C., Chung, B. I., & Dube, W. V. (2003) Persistence of stereotypic behavior: Examining the effects of external reinforcers. *Journal of Applied Behavior Analysis*, *36*, 439-448.

Anger, D., & Anger, K. (1976). Behavior changes during repeated eight-day extinctions. *Journal of the Experimental Analysis of Behavior*, *26*, 181-190.

Baum, W. M. (2012). Extinction as discrimination: The molar view. *Behavioural Processes*, *90*, 101-110.

Blackman, D. E. (1968). Response rate, reinforcement frequency, and conditioned suppression. *Journal of the Experimental Analysis of Behavior*, *11*, 503-516.

Boakes, R. A. (1973). Response decrements produced by extinction and by response-independent reinforcement. *Journal of the Experimental Analysis of Behavior*, *19*, 293-302.

Brown, P., & Jenkins, H. M. (1968). Auto-shaping of the pigeon's key peck. *Journal of the Experimental Analysis of Behavior*, *11*, 1-8.

Bullock, D. H. (1960). Repeated conditioning-extinction sessions as a function of the reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, *3*, 241-245.

Catania, A. C. (1973). The nature of learning. In J. A. Nevin & G. S. Reynolds (Eds.), *The study of behavior* (pp. 30-68). Glenview, IL: Scott Foresman.

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.

Clark, F. C. (1964). Effects of repeated VI reinforcement and extinction upon operant behavior. *Psychological Reports*, *15*, 943-955.

Cohen, S. L. (1996). Behavioral momentum of typing behavior in college students. *Journal of Behavior Analysis and Therapy*, *1*, 36-54.

Cohen, S. L. (1998). Behavioral momentum: The effects of the temporal separation of rates of reinforcement. *Journal of the Experimental Analysis of Behavior*, *69*, 29-47.

Davenport, J. W. (1969). Successive acquisition and extinctions of discrete bar-pressing in monkeys and rats. *Psychonomic Science*, *16*, 242-244.

Devenport, J. A., & Devenport, L. D. (1993). Time-dependent decisions in dogs (*Canis familiaris*). *Journal of Comparative Psychology*, *107*, 169-173.

Devenport, L. D., & Devenport, J. A. (1994). Time-dependent averaging of foraging information in least chipmunks and golden-mantled ground squirrels. *Animal Behavior*, *47*, 787-802.

Devenport, L. D., Hill, T., Wilson, M., & Ogden, E. (1997). Tracking and averaging in variable environments: A transition rule. *Journal of Experimental Psychology: Animal Behavior Processes*, *23*, 450-460.

Devenport, J. A., Patterson, M. R., & Devenport, L. D. (2005). Dynamic averaging and foraging decisions in horses (*Equus callabus*). *Journal of Comparative Psychology*, *119*, 352-358.

Dreyfus, L. R. (1991). Local shifts in relative reinforcement rate and time allocation on concurrent schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, *17*, 486-502.

Findley, J. D. (1958). Preference and switching under concurrent schedules. *Journal of the Experimental Analysis of Behavior*, *1*, 123-144.

Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, *5*, 529-530.

Gallistel, C. R. (2012). Extinction from a rationalist perspective. *Behavioural Processes*, *90*, 66-80.

Gallistel, C. R., King, A. P., Gottlieb, D., Balci, F., Papachristos, E. B., Szalecki, M., & Carbone, K. S. (2007). Is matching innate? *Journal of the Experimental Analysis of Behavior*, *87*, 161-199.

Gallistel, C. R., Mark, T. A., King, A. P., & Latham, P. E. (2001). The rat approximates an ideal detector of changes in rates of reward: Implications for the law of effect. *Journal of Experimental Psychology: Animal Behavioral Processes*, *27*, 354-372.

Grace, R. C., & Nevin, J. A. (1997). On the relation between preference and resistance to change. *Journal of the Experimental Analysis of Behavior*, *67*, 43-65.

Grimes, J. A., & Shull, R. L. (2001). Response-independent milk delivery enhances persistence of pellet-reinforced lever pressing by rats. *Journal of the Experimental Analysis of Behavior*, *76*, 179-194.

Harper, D. N. (1996). Response-independent food delivery and behavioral resistance to change. *Journal of the Experimental Analysis of Behavior*, *65*, 549-560.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, *13*, 243-266.

Igaki, T., & Sakagami, T. (2004). Resistance to change in goldfish. *Behavioural Processes*, *66*, 139-152.

Jenkins, H. M. (1961). The effect of discrimination training on extinction. *Journal of Experimental Psychology*, *61*, 111-121.

Lattal, K. A. (1989). Contingencies on response rate and resistance to change. *Learning and Motivation*, *20*, 191-203.

Mace, F. C., Lalli, J. S., Shea, M. C., Lalli, E. P., West, B. J., Roberts, M., & Nevin, J. A. (1990). The momentum of human behavior in a natural setting. *Journal of the Experimental Analysis of Behavior*, *54*, 163-172.

Mace, F. C., McComas, J. J., Mauro, B. C., Progar, P. R., Ervin, R., & Zangrillo, A. N. (2010). Differential reinforcement of alternative behavior increases resistance to extinction: Clinical demonstration, animal modeling, and clinical test of one solution. *Journal of the Experimental Analysis of Behavior*, *93*, 349-367.

Mazur, J. E. (1995). Development of preference and spontaneous recovery in choice behavior with concurrent variable-interval schedules. *Animal Learning & Behavior*, *23*, 93-103.

Mazur, J. E. (1996). Past experience, recency, and spontaneous recovery in choice behavior. *Animal Learning & Behavior*, *24*, 1-10.

Nevin, J. A. (1974). Response strength in multiple schedules. *Journal of the Experimental Analysis of Behavior*, *21*, 389-408.

Nevin, J. A. (1992a). An integrative model for the study of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, *57*, 301-316.

Nevin, J. A. (1992b). Behavioral contrast and behavioral momentum. *Journal of Experimental Psychology: Animal Behavior Processes*, *18*, 126-133.

Nevin, J. A. (2002). Measuring behavioral momentum. *Behavioural Processes*, *57*, 198-287.

Nevin, J. A. (2012a). Behavior theory and information theory: Some parallels. *Behavioural Processes*, *90*, 84-86.

Nevin, J. A. (2012b). Resistance to extinction and behavioral momentum. *Behavioral Processes*, *90*, 89-97.

Nevin, J. A., & Grace, R. C. (2000). Behavioral momentum and the law of effect. *Behavioral and Brain Sciences*, *23*, 73-130.

Nevin, J. A., Grace, R. C., Holland, S., & McLean, A. P. (2001). Variable-ratio versus variable-interval schedules: Response rate, resistance to change, and preference. *Journal of the Experimental Analysis of Behavior*, *76*, 43-74.

Nevin, J. A., Mandell, C., & Atak, J. R. (1983). The analysis of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, *39*, 49–59.

Nevin, J. A., Mandell, C., & Yarensky, P. (1981). Response rate and resistance to change in chained schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*, 278-294.

Nevin, J. A., McLean, A. P., & Grace, R. C. (2001). Resistance to extinction: Contingency termination and generalization decrement. *Animal Learning & Behavior*, *29*, 176-191.

Nevin, J. A., Tota, M. E., 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.

Rescorla, R. A., & Skucy, J. C. (1969). Effect of response-independent reinforcers during extinction. *Journal of Comparative and Physiological Psychology*, *67*, 381-389.

Shahan, T. A., & Burke, K. A. (2004). Ethanol-maintained responding of rats is more resistance to change in a context with added non-drug reinforcement. *Behavioral Pharmacology*, *15*, 279-285.

Shull, R. L. (2005). The sensitivity of response rate to the rate of variable-interval reinforcement for pigeons and rats: A review. *Journal of the Experimental Analysis of Behavior*, *84*, 99-110.

Shull, R. L., & Grimes, J. A. (2006). Resistance to extinction following variable-interval reinforcement: Reinforcer rate and amount. *Journal of the Experimental Analysis of Behavior*, *85*, 23-39.

Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*, 887-895.