Evaluation of Brain Stimulation Reward in Rats: Heuristics, Exemplars, and Context-Dependency.

Bonnie Sonnenschein

A Thesis

in

The Department

of

Psychology

Presented in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy at
Concordia University
Montreal, Quebec, Canada

February 2004

National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisisitons et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence
ISBN: 0-612-90404-0
Our file Notre référence
ISBN: 0-612-90404-0

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou aturement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this dissertation.

While these forms may be included in the document page count, their removal does not represent any loss of content from the dissertation.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de ce manuscrit.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

Canadä

CONCORDIA UNIVERSITY
SCHOOL OF GRADUATE STUDIES

This is to certify that the thesis prepared

By: **Bonnie Sonnenschein**

Entitled: **Evaluation of Brain Stimulation reward in Rats: Heuristics, Exemplars, and Context-Dependency**

and submitted in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY (Psychology)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

______________________ Chair
Dr. D. Frost

______________________ External Examiner
Dr. R. Gallistel

______________________ External to Program
Dr. G. Saad

______________________ Examiner
Dr. A. Arvanitogiannis

______________________ Examiner
Dr. M. von Grunau

______________________ Thesis Supervisor
Dr. P. Shizgal

Approved by ______________________
Chair of Department or Graduate Program Director

March 4 2004

____
Dean of Faculty

Abstract

Evaluation of Brain Stimulation Reward in Rats: Heuristics, Exemplars, and Context-Dependency

Bonnie H. Sonnenschein, Ph.D.
Concordia University, 2004

The current study examined the cognitive heuristics employed by rats evaluating rewarding brain stimulation. Rats implanted with electrodes aimed at the lateral hypothalamus were trained to self-administer brain stimulation. In Experiment 1, one type of stimulation train was presented at a time: some trains were set at a constant frequency, whereas others were composite-frequency trains in which the frequency at the end of the train was set to be lower than in the preceding component. If the "peak" and "end" reward values were used as exemplars in evaluating the stimulation, the weaker end should degrade the overall reward value of the composite trains. However, the performance curves for all trains were overlapping, suggesting that all trains were equally rewarding. A model in which the "peak" exemplar alone is employed, is sufficient to account for these findings. In Experiment 2, constant- and composite-frequency trains were presented simultaneously. Adding a weaker end had no effect on preference, again suggesting that the simple "peak" alone model was sufficient. Experiment 3 was similar to the first experiment, in that constant- and composite-frequency trains were presented one at a time. However, the duration of the weaker, terminal portion of the composite trains was now greater. The performance curves for all trains were again generally overlapping, indicating no consistent preference for any train type. This result is again consistent with the "peak-only" model. The aim of Experiment 4 was to assess whether the "beginning" exemplar played a role in the evaluation of rewarding stimulation. Preliminary results did not support this hypothesis. Lastly, Experiment 5 examined whether duration could serve as an exemplar when the evaluation context was changed. In an earlier experiment involving trains presented one at a time, rats were indifferent to train duration increases once the stimulation exceeded 2-4 s. In Experiment 5, when comparable train durations were presented simultaneously, the longer train was more effective than the shorter trains in 4 out of 7 subjects. The implications of the study findings and future directions are discussed.

Acknowledgements

This thesis would not have been possible without the input and help of many individuals.

I would like to extend my thanks to computer guru Steve Cabilio for his tireless efforts to convince PREF2 that yes, we DO want it to run, upload (or is it download? I can never remember) data, etc. etc. Thanks are also due to hardware guru Dave Munro for his tireless efforts to convince the PREF2 operant boxes that yes, we do want them to run the experiments and behave properly.

I wish also to thank my committee – Dr. Michael von Grünau, Dr. Andreas Arvanitogiannis, Dr. Gad Saad, and Dr. Randy Gallistel – for their insightful comments and their efforts to make this thesis as clear and understandable as possible.

Kent Conover is also owed a debt of thanks, for all the work he did to help develop the software programs to run the experiments and analyze the data, and also for the numerous times he served as a sounding board. His willingness to discuss therapy-related, and metaphysical issues regarding the nature of the universe, was also appreciated.

Thanks are also due to a group of people whom, while they did not directly contribute to the thesis, are still due many thanks for helping me to develop more fully the clinical side of my nature. Thus, I wish to extend heartfelt thanks to the faculty of the

Montreal General Hospital: Drs. Tinaz Chinoy, Marie-France Boudreault, Richard Karmel, and Allen Surkis. Thanks are particularly due to Dr. Surkis for introducing me to his particular brand of psychodynamic therapy, which has greatly advanced my functioning as a therapist and my understanding of patients and the way they function as people.

I wish also to thank both the Fonds FCAR of Quebec, and the Natural Sciences and Engineering Research Council of Canada (NSERC), for ensuring that I was always able to afford the *gourmet* bread and water over the course of my studies.

Most importantly, I would like to thank my research supervisor, Dr. Peter Shizgal. Not only has he put up with me and my idiosyncrasies for the last 6.5 years of my studies, but he has always been willing to: a) help me work through my Mathematics Disorder (see the Diagnostic and Statistical Manual of Mental Disorders - IV), b) help me work through my logic in understanding how to interpret what the rats are telling us, and c) meeting with me whenever I needed it. He deserves special thanks for all the efforts he has made to help me develop more fully as a researcher, and I look forward to future collaborations (until he gets thoroughly sick of me).

Another group of people deserve thanks. Again, not for contributing directly to the thesis, but for contributing to my continued mental health (such as it is). Keeping me out of 4-East at the Montreal General was a great help in allowing me to write this thesis – being kept in restraints makes it damn hard to type, after all. In this vein, I must thank my parents, Sam and Bev Sonnenschein, and particularly my dear friends Kathy and Harvey Auerback.

But the most important person of the "indirect contributors" whom I must thank is my sweetie, François Verpaelst. He has not only been there for me in innumerable ways over the last 4 ½ years, but he has helped me to turn off my "Internal Editor" (superego, to you clinical types) and give my creativity free rein in the areas of music and artistic expression. The ability to not only play Celtic fiddle (and recently, guitar) but to learn by ear, as well as to be able to teach others how to play, has added a whole new facet to who I am as a person, and Frank has been instrumental (no pun intended) in helping me to help myself achieve these skills. He is the best friend anyone could ever ask for, and I love him.

This thesis is thus dedicated to the three people who have made the most important contributions to my personal development over the last while, in those three critical areas of scientist, practitioner, and artist – Peter Shizgal, Allen Surkis, and François Verpaelst.

Table of Contents

Page

LIST OF TABLES xi

LIST OF FIGURES xii

GENERAL INTRODUCTION 1
- The Discovery of Brain Stimulation Reward (BSR), and the Relationship Between BSR and Natural Reinforcers 1
- Characteristics of the Neural Substrate of BSR 5
- Calculating the Subjective Intensity of Reward, and Turning it into Behaviour 17
- Using the TL Model to Resolve the Discrepancy 22
- The Temporal Compression Rule: Peak-Only versus Peak-and-End 38
- The Current Study 46

EXPERIMENT 1 48
- Introduction 48
- Method 51
  - Subjects 51
  - Surgery 51
  - Apparatus 52
    - Screening and training 52
    - Experiments 53
  - Procedure 54
    - Screening and training 54
    - Baseline condition 56
    - Single-operant peak-and-end test 57
- Results 59
- Discussion 76

EXPERIMENT 2 78
- Introduction 78
- Method 80
  - Subjects 80
  - Surgery 80
  - Apparatus 80
  - Procedure 80
    - Screening and training 80
    - Baseline condition 81
    - Dual-operant peak-and-end experiment 81
- Results 87
- Discussion 102

EXPERIMENT 3 104
- Introduction 104
- Method 105
  - Subjects 105
  - Surgery 105
  - Apparatus 105
  - Procedure 105
    - Screening and training 106
    - Baseline condition 106
    - Single-operant "2+6" peak-and-end test 106
- Results 108
- Discussion 123

EXPERIMENT 4 129
- Introduction 129
- Method 132
  - Subjects 132
  - Surgery 132
  - Apparatus 132
  - Procedure 132
    - Baseline condition 132
    - Dual-operant beginnings-and-peaks test 132
- Results 135
- Discussion 144

EXPERIMENT 5 146
- Introduction 146
- Method 148
  - Subjects 148
  - Surgery 148
  - Apparatus 148
  - Procedure 148
    - Baseline condition 148
    - Dual-operant preference-reversal test 148
  - Histology 149
- Results 153
- Discussion 167

GENERAL DISCUSSION 170
- Peak-and-End Tests 171
- Beginnings-and-Peaks versus Temporal Summation 175
- Preference Reversal, Evaluation Context, and Duration Neglect 176
- Summary and Conclusions 180

REFERENCES 183
APPENDIX 191

List of Tables

Table Page

1 Results of fitting the half-maximal performance data points with a two-dimensional rectangular hyperbola. 29

2 Results of fitting the data from both experiments with the Triple-Logistic Shizgal Model. 37

3 Differences in F-values in the single-operant peak-and-end test, depending on whether a "Common" or "Separate" slope model was used. 68

4 Differences in F-values in the single-operant "2+6" peak-and-end test, depending on whether a "Common" or "Separate" slope model was used. 114

List of Figures

Figure Page

1 Illustrative hyperbolic function, showing the relationship between required stimulation "strength" and train duration. 10

2 Graphical representation of temporal summation in the integrator, adapted from Fouriezos (1995), and based on Gallistel's (1978) model 13

3 Graphical representation of the different elements comprising the Shizgal Triple-Logistic (TL) Model. 20

4 Rescaled data results from the single-operant frequency sweeps test. 25

5 The strength-duration fits arising from the single-operant frequency sweeps test. 28

6 Rescaled data results from the duration sweeps test. 32

7 Results of fitting the rescaled data from both the frequency and duration sweeps experiments using the Shizgal TL model. 35

8 Raw data results from the single-operant peak-and-end test. 61

9 A sample single-operant graph, showing the broken-line fit to a performance curve. 63

10 Rescaled means from the single-operant peak-and-end test. 66

11 Additional frequency sweeps data. 72

12 Additional frequency sweeps data for the remaining subjects. 74

13 A sample curve showing the "shoulder" region of a single-operant performance curve. 84

14 A sample dual operant graph showing the broken-line fits and the estimated cross point (CP). 89

15 Sample dual-operant graph, showing the comparison between the estimated cross point (CP) and the frequency of the standard lever (vertical magenta line). 92

16 Results from the first condition of the dual-operant peak-and-end experiment. 95

17 Results from the duration neglect condition of the dual-operant peak-and-end experiment. 98

18 Results of the dual-operant peak-and-end test. 100

19 Raw data results from the single-operant 2+6 peak-and-end test. 110

20 Rescaled means from the single-operant 2+6 peak-and-end test. 113

21 Graphs comparing performance (rescaled data) for 2 s and 8 s trains in both the frequency sweeps and 2+6 experiments. 120

22 Graphs comparing performance (rescaled data) for 2 s and 8 s trains in both the frequency sweeps and 2+6 experiments, for the remaining subjects. 122

23 Preliminary results from the dual-operant 2+6 peak-and-end test. 126

24 Results of the delayed Std train vs. the 2 s Alt train condition of the dual-operant beginnings-and-peaks test. 137

25 Results of the "6w + 2 s" Std train vs. the 2 s Alt train condition of the dual-operant beginnings-and-peaks test. 139

26 Results of the modified dual-operant beginnings-and-peaks test, for 1 rat. 142

27 The locations of the electrode tips. 152

28 Results of the 4 s Std train vs. 4 s Alt train condition of the dual-operant preference reversal experiment. 155

29 Results from the dual-operant preference reversal test, for subjects tested with 4 s Alt trains. 158

30 Results from the dual-operant preference reversal test, for 1 subject tested with a 2 s Alt train. 160

31 Results from the dual-operant preference reversal test, plotted with the adjusted Std value (vertical dot-dashed line), for subjects tested with 4 s Alt trains. 163

32 Results from the dual-operant preference reversal test, plotted with the adjusted Std value (vertical dot-dashed line), for 1 subject tested with a 2 s Alt train. 165

The Discovery of Brain Stimulation Reward (BSR), and the Relationship Between BSR and Natural Reinforcers

The study of the phenomenon of brain stimulation reward (BSR) began with a serendipitous discovery by Olds and Milner in 1953 (Milner, 1989; Olds, 1973; Olds & Milner, 1954). These two researchers were attempting to electrically stimulate the reticular formation of one of their rats. They found that the animal would move forward and display exploratory behaviour whenever the stimulation was turned on, whereas when the stimulation was unavailable, the rat would cease these behaviours. They also discovered that it was possible to direct this particular animal to move in any particular direction they wanted, simply by giving the animal a short burst of stimulation every time it turned in the appropriate direction (Milner, 1989). Indeed, this procedure, as simple as it might sound, may have some practical real-world applications today. Talwar et al. (2002) implanted rats with both stimulation electrodes and electrodes aimed at the right and left "whisker" sensory areas. The goal was to train the rats to navigate in certain directions selected by the experimenters, with the "whisker" area stimulation used to direct the rats to turn either right or left, and stimulation of the medial forebrain bundle (MFB) to reward the rats for moving in the selected direction. If such rats were given "backpacks" containing all the necessary equipment, such as remote-controlled microstimulators and batteries, it could be possible to one day use such "guided" rats for search and rescue operations in cases of building collapse and other disasters.

The experimentation of Olds and Milner did not stop at directing their rat to move about. One of their alternative hypotheses was that the electrical stimulation was "forcing" the animal to move forward. To eliminate this possibility they set up a more complex self-stimulation situation that would require the rat to press a lever to stimulate itself (and with no experimenters nearby to inadvertently give the rat cues). Using this procedure, it was found that this particular animal not only learned to press the lever to stimulate itself, but would do so persistently for long periods of time (Milner, 1989).

Olds and Milner attempted to replicate this finding with other subjects with electrodes aimed at the reticular formation, but were mystified when the new animals did not self-stimulate. As they were reluctant to sacrifice the one animal they possessed which displayed this behaviour, they eventually decided to anesthetize this particular rat and x-ray its head, and upon doing so they discovered that the electrode had missed its target and was located in the septal area. With a new target to aim for, Milner and Olds soon discovered that other rats implanted with stimulating electrodes in the same area would reliably self-stimulate (Milner, 1989).

Olds and Milner (1954) then went on to implant electrodes in other brain regions, and again attempted to train them to bar-press for stimulation. The acquisition of lever-pressing, as well as the extinction of the operant response after the stimulation was turned off, was then correlated with the site of stimulation to determine whether stimulation at that site was “rewarding” , “neutral”, or “punishing”. The highest responding for the brain stimulation was again seen in rats with electrodes aimed at the septal area, while other areas, such as the mammillothalamic-tract and cingulate cortex, were also shown to support robust self-stimulation. Stimulation of other sites, such as the caudate nucleus and corpus callosum, did not lead to much change in behaviour; these sites were therefore labeled "neutral". Lastly, stimulation of areas such as the medial geniculate was described as "punishing", as lever-pressing was more likely when the stimulation was unavailable.

Having discovered that BSR was not unique to the septal area, Milner and Olds suggested that a network of several brain regions may underlie this phenomenon, and they further suggested that the term "reinforcing structures" should be used "as a general name for the septal area and other structures which produce the reward phenomenon" (Milner & Olds, 1954, p. 425). Indeed, there are a number of such "reinforcing structures" in various areas of the brain (see Wise (1996), or Shizgal & Murray (1989). One researcher has even asserted that "self-stimulation can be obtained in electrode sites in roughly one-fifth of the total volume of the rat brain" (Yeomans, 1988, p. 227). The lateral hypothalamic area (LHA) is the site most often studied, because acquisition and maintenance of self-stimulation tends to be very reliable.

For what reason would researchers be interested in studying so "artificial" a phenomenon, where electrical stimulation is "injected" into a subject's brain? One reason is that it is believed that the neural system(s) activated by the electrical stimulation overlaps those subserving other motivated behaviours, such as feeding, drinking, sexual behaviour, and drug seeking and consumption (Wise, 1980). The work of many early researchers suggested that BSR is related to the rewarding effects of natural goal-objects (Balagura & Hoebel, 1967; Hoebel, 1968; Hoebel & Teitelbaum, 1962; Hoebel & Thompson, 1969; Margules & Olds, 1961; Morgan & Mogenson, 1966).

In more recent work, it has been argued that the rewarding stimulation mimics some but not all of the properties of natural goal-objects. BSR mimics enough properties of natural goal-objects that it can compete and summate with them. As an example of such competition, when rats were presented with a choice between BSR and intraoral sucrose, the sucrose was preferred when the frequency of the BSR was low, whereas preference switched to BSR when its frequency was high (Conover & Shizgal, 1994a). In the case of summation, when BSR was presented in a "compound" with sucrose and rats were offered a choice between this compound and BSR alone, the compound stimulus was preferred (Conover & Shizgal, 1994a, 1994b). Similar results have also been obtained with saline rewards (Conover, Woodside, & Shizgal, 1994). Taken together, these results suggest that the rewarding effects of BSR and sucrose/saline were evaluated using a common "currency", and thus a "common" unit of measurement was used. However, manipulating the physiological state of the rats through sodium depletion or intragastric cannulae (Conover & Shizgal, 1994b) affects natural rewards, but not BSR. This suggests that the rats are not "hallucinating" a gustatory reward while receiving BSR. Furthermore, these results suggest that the reward signal from the BSR electrodes is "injected" downstream in the neural pathway from where sensory information from natural rewards are weighted by physiological feedback.

Thus it appears that both natural rewards and BSR are evaluated using a common "currency". BSR, however, has some additional advantages over natural rewards. First, it can be "aimed" directly at particular brain sites of interest. Second, it can be rigorously controlled by the experimenter. Last, performance for BSR is stable over time (i.e., performance during experimental sessions does not "satiate"). Thus, while BSR could be seen as an "artificial" reward, it can uniquely contribute to the search for how all rewards, including natural rewards, are evaluated by the brain.

Characteristics of the Neural Substrate of BSR

The actual identity of the neurons directly stimulated by the electrodes, hereafter referred to as the "first-stage" neurons, is not currently known. Nonetheless, many of the characteristics of these first-stage neurons have been described. For example, these neurons are myelinated, have absolute refractory periods ranging from 0.5 to 1.2 msec, and have conduction velocities of 2-8 m/s (for a review, see Shizgal and Murray (1989), or Shizgal (1997)). These findings indicate that the MFB first-stage neurons cannot belong to any of the known catecholamine systems, because the neurons that generally subserve the latter systems are known to be unmyelinated, have refractory periods greater than 1.5 msec, and their conduction velocities are less than 1.0 m/s. For example, Yeomans, Maidment and Bunney (1988), in their extracellular recordings of stimulated dopamine (DA) neurons in the substantia nigra, zona compacta, and ventral tegmental area (VTA), found that the refractory periods of these neurons ranged from 1.2 to 2.5 msec. Based on this and other findings, they concluded that DA neurons are easiest to stimulate when high currents, long pulse durations, or small electrode tips are employed, neither of which is typically the case when BSR researchers aim at MFB stimulation sites.

There are various models that describe how the rewarding stimulation is thought to influence the directly-stimulated neurons. Typically, the stimulation is generally made up of trains of short-duration, low current, electrical pulses. When the pulse duration is held constant, the strength of a train is determined by both the pulse amplitude (current) and the frequency (in pulses per second; pps)[1]. At the level of the neural tissue, each individual electrical pulse produces a volley of action potentials (nerve impulses) along the axons of the first-stage neurons. These action potentials are then believed to pass over a synapse or set of synapses into a neural network, usually termed the "neural integrator" (Gallistel, Shizgal, & Yeomans, 1981), which is believed to integrate the post-synaptic effects of the volley over space and time.

Gallistel and colleagues (Gallistel, 1978; Gallistel et al., 1981) initially put forth a "counter model" that attempted to describe how the reward signal (output) of the neural integrator is determined by the aggregate rate of firing (input) coming from the first-stage neurons. Their model suggests that the input is summated over space and time in such a way that it is not important what combination of frequency and stimulation current was delivered. Put another way, the frequency and current intensity are reciprocally related such that provided the aggregate rate of impulse flow stays the same, it does not matter whether the impulse flow was produced by firing many neurons a few times (high current, low frequency) or a few neurons many times (low current, high frequency) (Gallistel, 1976). In both cases, according to

---

[1] Stimulation "strength" refers to the set of parameters that determine the number of elicited action potentials per second in the first-stage (directly stimulated) neurons. This set consists of the current, pulse duration, and frequency. Stimulation train duration, which is a temporally-extended parameter of the stimulation, is usually not considered to contribute to the "strength" of a stimulation train, under this definition.

the counter model, the reward signal from the integrator will be the same, always providing that the aggregate rate of impulse flow stays the same.

Shizgal and Matthews (1977) demonstrated that the rewarding effect of BSR grows over time, leveling off once the stimulation exceeds some critical value. They conducted an experiment in which rats were implanted with LHA stimulating electrodes, and trained to work in a dual-operant paradigm in which pressing on one lever turned the stimulation on, and pressing the other lever turned the stimulation off. Stimulation current intensity was traded-off against the duration of the "bursts" of stimulation, and the results indicated that when the burst width exceeded 1-2 seconds, the current intensity required to maintain criterion performance leveled off.

The results of a subsequent study (Gallistel, 1978), are consistent with those of Shizgal and Matthews (1977). In his experiment, rats were trained to run back and forth between two levers separated by a passageway, and were required to respond on one lever first, before they were allowed to run down the passageway to the reward lever. The site of the stimulation was the LHA, and the behavioural performance criterion selected for all analyses was 70% of maximum running speed in the passageway. The stimulation trains were identical on both levers, with the exception that the duration presented on each lever varied from session to session. The results suggested that stimulation trains of short duration needed to be very strong in order to hold the output of the neural integrator constant, but as the duration of the stimulation increased, the strength of the stimulation required to maintain the subject's behaviour decreased. This relationship was shown to follow a rectangular hyperbolic function, as shown in Figure 1. This function is often referred to as a "strength-duration function". As the trains get longer, the strength of the stimulation required to maintain performance decreases, but this decrement becomes more gradual as the train duration lengthens. Eventually, as the trains become longer and longer, the required strength levels off at an asymptote, which Gallistel (1978) termed the "rheobase". It should also be noted, however, that while short trains would need to be strong in order to hold the output of the integrator constant, longer trains would require more firings than short trains, simply because with longer stimulation trains there is more time for the excitation to decay while the train is delivered. The form of the strength-duration (hyperbolic) function is assumed to reflect the rate of "leakiness" (decay characteristics) of the integrator, and this "curviness" of the hyperbolic function is described by the "chronaxie" of the hyperbola (Gallistel, 1978). When this chronaxie is large, the shape of the hyperbolic fit will have less of a curve to it and will appear more like a straight line, whereas when the chronaxie is small, the shape of the hyperbolic fit will be more deeply curved. The chronaxie is defined as the duration at which the required strength is twice the rheobase (Gallistel, Shizgal, & Yeomans, 1981).

Given that the rewarding effect gradually increases and then levels off, Gallistel (1978) initially proposed a "leaky bucket" model of how the aggregate rate of impulse flow is integrated over time. He suggested that there is a single neural integrator involved in summating the reward signal (output) arising from stimulation of the first-stage neurons, and that this integrator reacts to the stimulation as if it were a "leaky bucket being filled by a spurting hose" (p. 978). If one pours water into a leaky bucket at a steady rate, initially the water level will be low and the water will leak out at a slow rate. As the water level rises,

Figure 1. An illustrative hyperbolic function, displaying the relationship between required stimulation "strength" and train duration. The green line traces a rectangular hyperbola. Train duration in seconds is plotted on the x-axis, and the frequency required to obtain half-maximal (50%) performance is expressed in logarithmic units (pulses per second) and plotted as the y-axis variable.

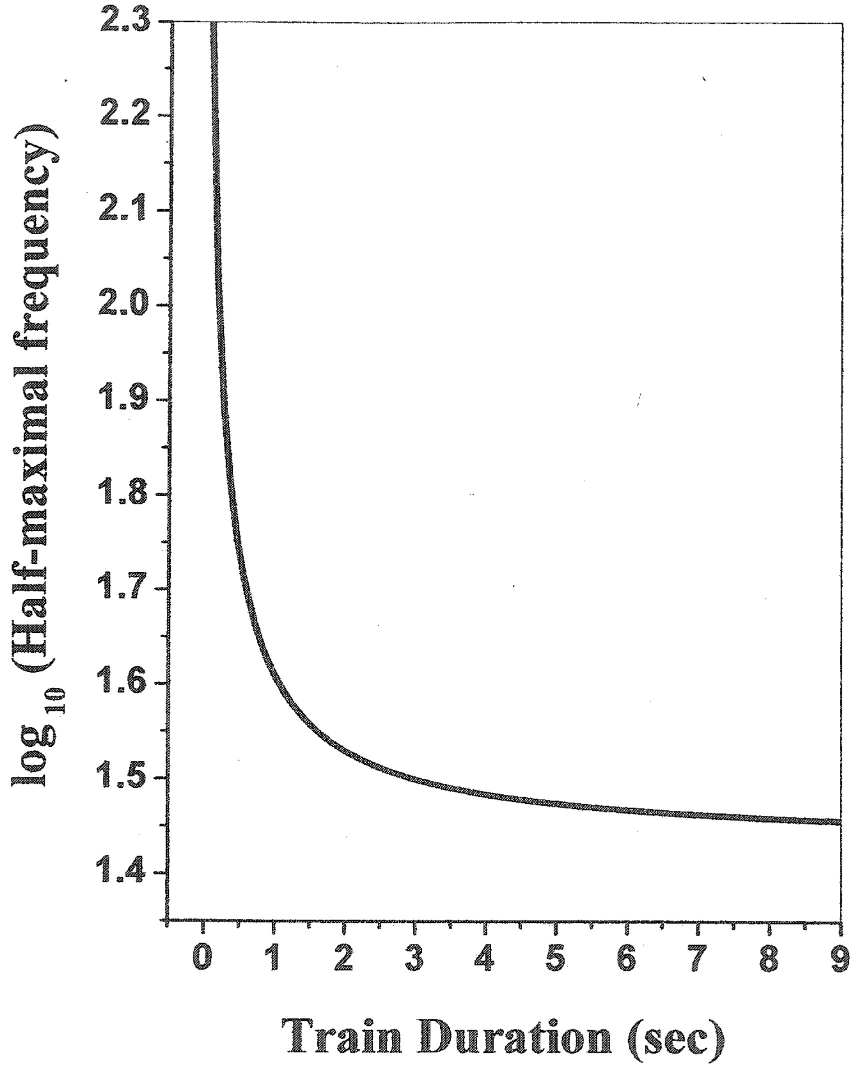

Figure 1. Illustrative hyperbolic function, showing the relationship between required stimulation strength and train duration

however, the water will leak out more quickly, and eventually the rate of water being poured into the bucket will be the same as that leaking out, and the water level will thus have leveled off. If the neural integrator did behave in such a way, the number of firings of the first-stage neurons required to produce a given reward output would depend on the duration of the input (the train of stimulation), providing everything else is held constant. If the input to the integrator is interrupted, the excitation in the integrator is thought to decay as a function of time. Figure 2 (adapted from Fouriezos, 1995) shows a graphical representation of how this decay rate could influence the output. However, this "leaky bucket" explanation of how reward is integrated over time does not in fact capture how the neural integrator performs its function, as will be described later.

Other researchers, however, have not been satisfied with the notion that only one neural integrator is responsible for computing the reward output. Mason and Milner (1986a), implanted rats with electrodes aimed at the LHA, and then gave the rats a choice between two rewards in a Y-maze. In the first phase of the experiment, the train duration of the "standard" lever reward was held constant, while the train duration of the "test" lever reward was varied. On both levers the frequencies were also held constant, at either a low (70 Hz) or a high (200 Hz) frequency. A second phase of the experiment employed a similar paradigm, except that three different frequencies were used: 70, 200, or 400 Hz. The results suggested that when the lower frequency reward (70 Hz) was delivered, it took longer than 5 seconds for reward value to approach asymptote, whereas if the higher frequency (200 Hz) train was delivered, reward value saturated at train durations of approximately 2 seconds. In addition, when the very high frequency (400 Hz) was used, reward value approached asymptote at train durations as short

Figure 2. Graphical representation of temporal summation in the integrator, adapted from Fouriezos (1995), and based on Gallistel's (1978) model. The vertical dimension of the graph represents excitation in the neural network underlying the reward effect (the integrator), and the horizontal dimension of the graph represents time. In part a, the trains are of equal duration. As the input rate (stimulation frequency) is increased, the summated excitation approaches, and eventually surpasses, the threshold (represented by the dotted line). Between pulses, excitation "leaks" away (decays). In part b, inserting a pause in the middle of the train lowers the maximum level attained. (The frequency is the same input rate as the middle trace in part a).

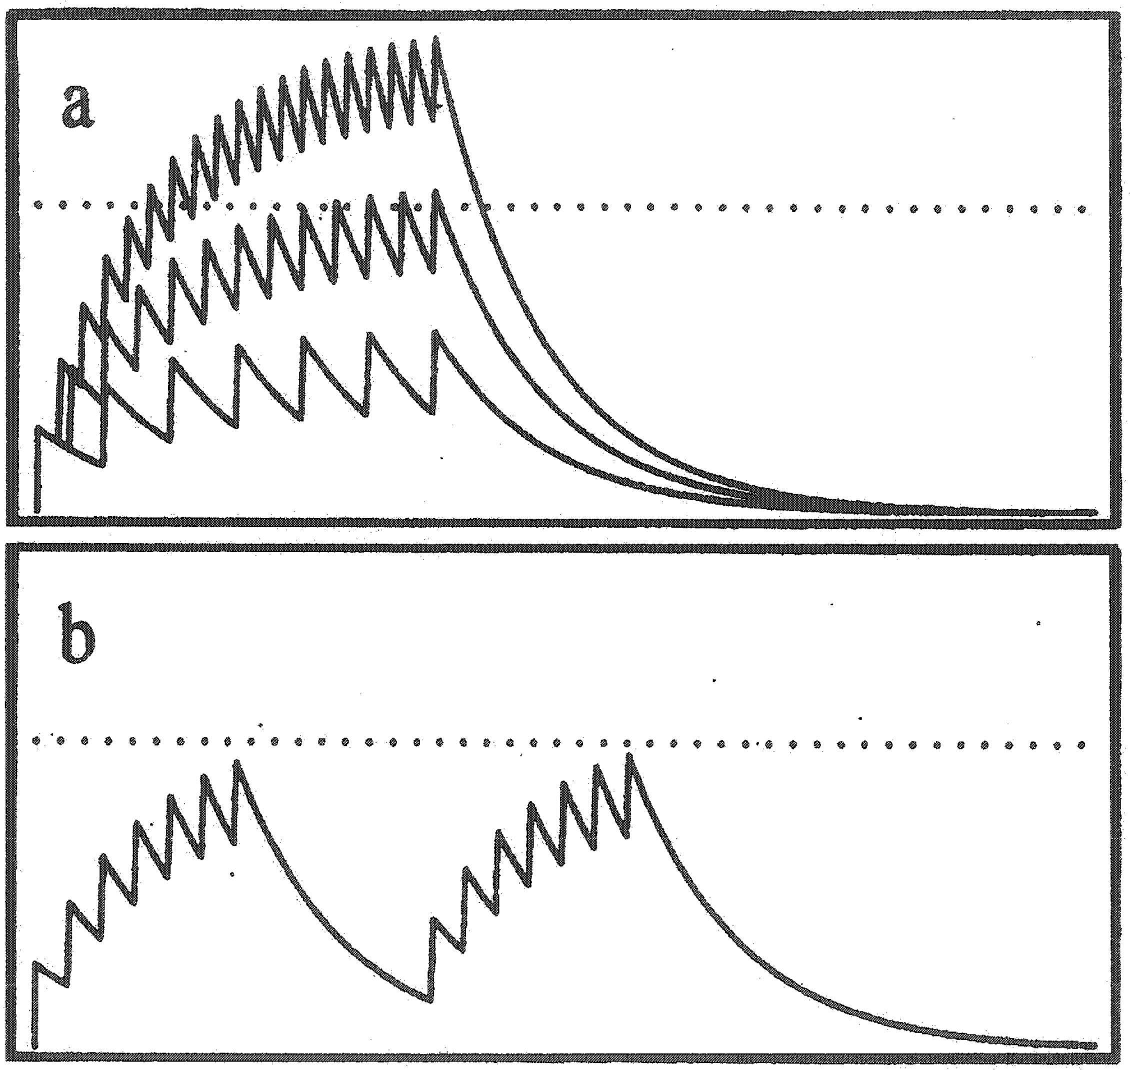

Figure 2. Graphical representation of temporal summation in the integrator, adapted from Fouriezos (1995), and based on Gallistel's (1978) model

as 0.5 to 1 s. Taken together, these results suggested that the train duration at which "fatigue" (duration neglect) set in was dependent upon the frequency used.

Mason and Milner (1986b) also conducted a second study, using a similar Y-maze paradigm, but in this case the duration of the reward on both the "standard" and "test" levers was always the same. Another modification in this experiment was that the test reinforcement was composed of two portions – part A, the first 2.5 seconds of the test reward, was always set to the same stimulation characteristics as the standard reward, and part B, the remaining portion of the test reward, was sometimes varied in frequency from the previous portion of the train. In phase one of the study, "part A" of the test lever reward and the standard reward were both set at a frequency of 100 Hz, and part B of the test reward was set at either 50, 100, or 200 Hz. The second phase of the experiment was similar, except in this case the standard reward and part A of the test reward were both held constant at 250 Hz, and part B was set at either 125, 250, or 500 Hz. Similar to the results of Mason and Milner's earlier (1986a) study, when both the standard and part A were set at the lower (100 Hz) frequency, the rats preferred the test rewards in which part B was set at 200 Hz (a 100% increase in the frequency), over those in which part B was either 100 (unchanged) or 50 Hz (decreased). This pattern did not hold when the standard and part A of the test reward were delivered at a higher frequency (250 Hz). They interpreted their results to suggest that when low frequency stimulation trains are delivered, the neural integrators do not fatigue and reward value has not yet approached asymptote, and thus further increases in stimulation strength were preferred by the rats. When high frequency trains are delivered, however, the neural integrators become fatigued more quickly and reward value approaches asymptote, and therefore the rats become indifferent to whether stimulation strength is further increased.

In contrast to Gallistel's (1978) study, Mason and Milner suggested that a dual-integrator model was needed to account for the results of their two studies (Mason & Milner, 1986a, 1986b, and see also Mason & Milner, 1985). They asserted not only that a second neural integrator had to be involved, but also that this second integrator had to have a longer time constant of "leakiness" (decay). This integrator modulates the output of the first integrator. They claim that the dual-integrator model yielded a good approximation of their data. The time constants they derived from their data were 450 msec for the first integrator and approximately 6.5 s for the second.

In summary, there is both empirical and theoretical disagreement between Gallistel (1978) and Mason and Milner (1985, 1986a, 1986b). Empirically, the data from the former research suggests that there is a characteristic time course for reward integration. In contrast the data from the latter camp suggests that this time course is dependent on stimulation strength. In terms of theoretical disagreement, Gallistel (1978) contends that a single neural integrator is responsible for carrying out reward integration, whereas Mason and Milner (1985, 1986a, 1986b) assert that two neural integrators with different time constants are responsible.

It is also interesting to note that the empirical evidence obtained by Mason and Milner (1986a, 1986b) demonstrates that "leaky bucket" integration, as described earlier (Gallistel, 1978) cannot be the operation used by the neural integrator. If an exponential function did govern this process, the different profiles seen between high and low frequencies would not have been obtained. This suggests that a new model of reward integration over time is needed, one that captures not only the effects seen by both Gallistel (1978) and Mason and Milner (1986a, 1986b), but also suggests a new form for the integration of reward over time.

Gallistel (1978) and Mason and Milner (1986a, 1986b) do agree that the integration level rises quickly at first, and then levels off, which will hereafter be referred to as "progressive insensitivity to duration". Recall that Gallistel (1978) observed that as the train duration lengthens, there are smaller and smaller changes in the frequency of the stimulation required to maintain half-maximal performance. This progressive insensitivity to duration follows a hyperbolic function, as noted above. Once the stimulation exceeds a particular duration, further increases in duration no longer lead to measurable changes in the required frequency (as shown by the "flattening out" of the hyperbolic strength-duration function at the longer train durations; see Figure 1).

Progressive insensitivity to duration is compatible with the findings of Shizgal and Matthews (1977), described earlier, as well as with another study conducted by Mark and Gallistel (1993), using a somewhat different methodology. In the latter experiment, rats with stimulating electrodes in various MFB regions were trained to respond on two levers, although in this case each lever had a different variable-interval schedule of reinforcement available. The stimulation trains on both the "standard" and "alternate" levers were held at the same frequency, and the train durations and variable intervals were traded-off against each other. The results of this study demonstrated that the subjective reward increased as a function of train duration, with subjective reward magnitude reaching a peak at durations of approximately 1 s, and then declining. Thus, the findings of all 3 studies (Gallistel, 1978; Mark & Gallistel, 1993; Shizgal & Matthews, 1977), although using different paradigms, suggested that once stimulation exceeds 1-2 seconds, duration neglect is seen, as represented by a lack of measurable changes in performance. In other words, the subjects displayed progressive insensitivity to duration, culminating in a point beyond which further increases in train duration no longer lead to measurable changes in required frequency.

Calculating the Subjective Intensity of Reward, and Turning it into Behaviour

In order to resolve the apparent discrepancy between the positions of Gallistel (1978) and of Mason and Milner (1986a, 1986b), while retaining the idea of progressive insensitivity to duration, Shizgal developed a model, referred to as the "triple-logistic" or TL model. This approach integrates the findings of both camps within a single neural-integrator model of BSR. This model incorporates three elements: a) the trade-off function between train duration and the frequency required to maintain half-maximal reward value approximates a rectangular hyperbola, thus encompassing the findings of Gallistel (1978), b) at a given stimulation train duration, the intensity of reward grows as a sigmoidal (s-shaped) function of the frequency (Gallistel & Leon, 1991), thus incorporating the influence of stimulation frequency/strength on reward value into the model and also encompassing the ideas of Mason and Milner (1986a, 1986b) and c) operant response performance grows as a sigmoidal function of the subjective reward intensity.

The key features of the TL model are represented graphically in Figure 3. The first element of the model is that the relationship between train duration and the stimulation frequency required to maintain half-maximal reward value approximates a hyperbola. As shown in Figure 3a, when these required frequencies (in pulses per second) are plotted on the y-axis and the train duration is plotted on the x-axis, this is indeed the case: the decrease in the required frequency initially occurs rapidly, but then begins to decrease more gradually, and eventually levels off at the rheobase. Thus, at shorter train durations, higher frequencies are needed to maintain half-maximal reward, and as the stimulation trains increase in duration, the required frequency gradually decreases and eventually levels off at the rheobase, while at train durations beyond the rheobasic frequency, further increases in stimulation duration have no further measurable effect on performance. This lack of measurable effect of further increases in duration will be called "duration neglect" in the remainder of this document.

The second element of the TL model holds that at a given train duration, subjective reward intensity grows as a sigmoidal function of the frequency. This is shown graphically in Figure 3b, where the common logarithms of the frequencies (in pulses per second) are represented on the x-axis, and the "Subjective Intensity of Reward" is represented on the y-axis. In such s-shaped curves, as can be seen from the figure, at both very high and very low values of the frequency, there is very little change in the reward value, resulting in the "flat" sections of the sigmoid function. In contrast, when the frequencies used are in a moderate range, large changes in the reward value (rising slope of the sigmoid function) are observed in response to changes in stimulation frequency. Stated another way, lower

Figure 3. Graphical representation of the different elements comprising Shizgal's Triple-Logistic (TL) Model. Part 3a (upper left quadrant) represents the first part of the model, which relates train duration to the frequency required to maintain half-maximal reward. As in Figure 1, this relationship is represented as a rectangular hyperbola (green curve). Part 3b (upper right quadrant) represents the second component of the TL model, which describes the relationship between the frequency of the stimulation and the "subjective intensity of reward" as a sigmoid function (green curve). Part 3c (lower left quadrant) represents the third component of the model, which relates subjective intensity of reward to performance (proportion of the available rewards harvested), also in sigmoidal fashion (green curve). Lastly, the form of the complete TL model is represented by the black grid "mountain" in the 3-D graph (lower right quadrant, 3d). The model thus is able to account for how train duration and frequency interact to produce performance (Proportion of Rewards Harvested).

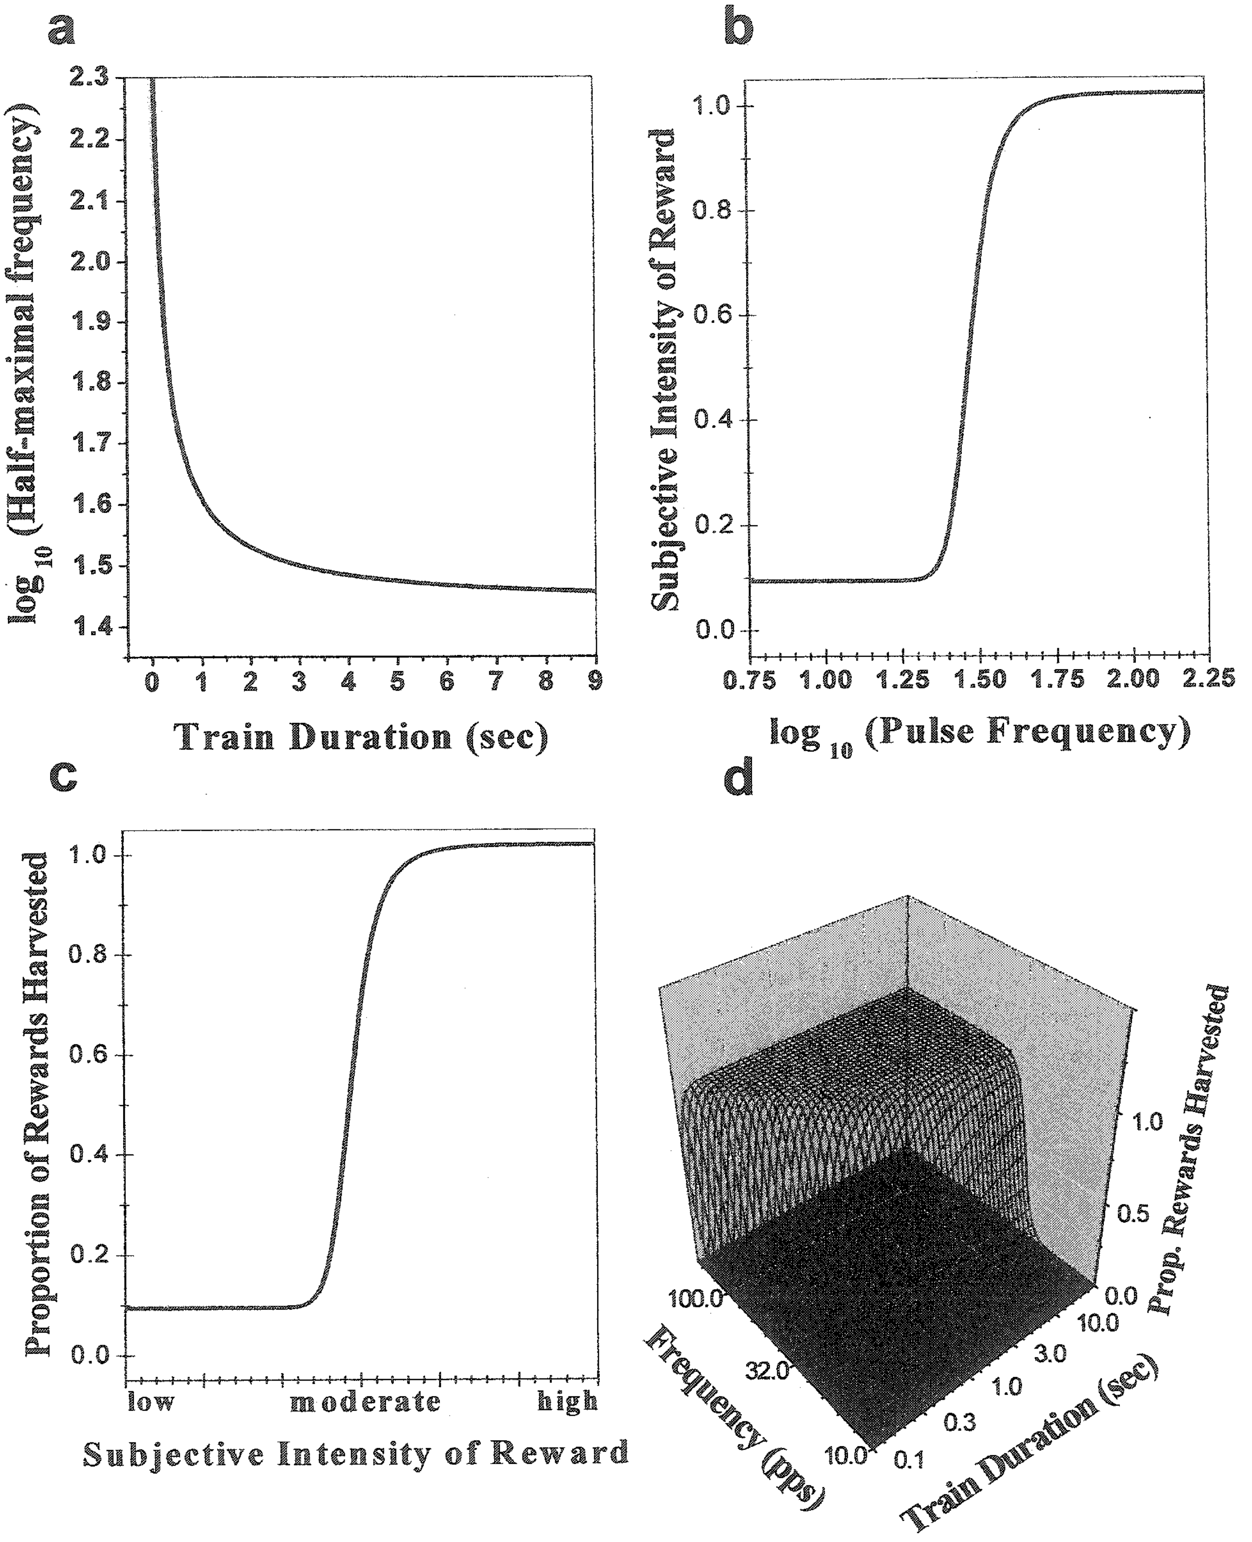

Figure 3. Graphical representation of the different elements comprising the Shizgal Triple-Logistic (TL) Model.

frequencies are too weak to support much change in reward value, and the same occurs at very high frequencies because the neural integrator is approaching saturation, whereas at moderate frequencies there are no such "limitations" on reward growth, and thus changes in stimulation frequency in this range can lead to large changes in reward value.

The third element of the TL model holds that operant responses from the subjects grow as a sigmoid function of the reward intensity. The reason this element is included in the model is because while the earlier two components of the model deal with subjective intensity of reward, the operant performance of the subjects is as yet the only means by which an experimenter can evaluate the subjective reward value experienced by experimental animals, and thus a third element was required to translate the inner experience of subjective reward value into the observable output of performance. This element of the TL model is an extension of Herrnstein's (1970) single-operant matching law, to a context in which a continuous reinforcement schedule is used. This relationship is displayed graphically in Figure 3c, with subjective intensity of reward plotted on the x-axis, and Proportion of Rewards Harvested (performance) plotted on the y-axis. Much like the sigmoid just described, very little change happens in the function when reward value is either very low or very high, whereas at moderate reward intensities, lever pressing changes sharply in response to relatively small changes in reward value. Stated another way, when at very low intensity the reward value is too weak to motivate the subject to engage in the self-stimulating behaviour, and thus there is little or no change in performance. When the reward value is very high, however, there is also very little or no change in performance because the rat will already be at a behavioural ceiling, harvesting all available rewards. When reward intensities are more moderate, however, such "limitations" are not a factor, and thus changes in reward value do lead to large changes in performance. For the interested reader, further details on the TL model are included in the Appendix.

If all these elements are put together and plotted in a three-dimensional space, a "mountain" graph is obtained, as shown in Figure 3d. The x-axis represents the train duration in seconds, the y-axis represents the frequency in pulses per second, and the z-axis represents the Proportion of Rewards Harvested (PRH). The black grid "mountain" thus represents the predicted harvest for any given combination of stimulation frequency and train duration.

Using the TL Model to Resolve the Discrepancy

Recall the apparent discrepancy between the findings of Gallistel (1978) and colleagues (Mark & Gallistel, 1993; Shizgal & Matthews, 1977), and those of Mason and Milner (1986a, 1986b). The work of the former researcher had suggested that there is a characteristic time course for reward integration, whereas the work of the latter researchers suggested that the time course is dependent on the strength of the stimulation.

Two experiments, conducted as part of the author's Master's research (Sonnenschein, Conover, & Shizgal, 2003), attempted to resolve the apparent discrepancy and to test the TL model. The first study was designed to be a replication of Gallistel's (1978) strength-duration experiment. In this first experiment, 6 rats were implanted with electrodes aimed at the LHA, and trained to bar-press for rewarding brain stimulation. The experiment was a trade-off experiment in which the train duration was held constant throughout the session, while the frequency was "swept" in a descending series. All other characteristics of the stimulation, such as pulse duration and current intensity, were held constant. The rats were exposed to several different train durations – 8, 4, 2, 1, 0.5, 0.25, and 0.125 s in some subjects – and each rat had to reach a stability criterion with one train duration before they were exposed to a subsequent train duration.

The results of this "frequency sweeps" experiment are shown in Figure 4. The y-axis represents the Proportion of Rewards Harvested (out of a maximum of 20 rewards available to the rat on each trial), and the x-axis represents the frequency in common logarithmic units on the lower x-axis, and in untransformed units on the upper x-axis (in pulses per second (pps)). All curves shown represent data averaged over 6 sweeps (descending frequency sweeps). From examination of the figure, it becomes evident that for all subjects tested, the performance curves lie approximately parallel to each other, and that as train duration increases, the distance between each curve becomes progressively smaller, a pattern which is most evident for rats B9, B15, and B26. This progressive insensitivity to duration is consistent with past findings (Gallistel, 1978; Mason & Milner, 1986a, 1986b).

From examination of Figure 4, for rats B9 and B28, the 4 s (magenta) and 8 s (orange) curves are lying close together, suggesting that duration neglect was displayed once the trains exceeded 4 s. For two of the remaining subjects, rats B15 and B26, the 4 s, 8 s, and 2 s (cyan)

Figure 4. Results of the Sonnenschein, Conover, and Shizgal (2003) frequency-sweeps experiment, for 6 rats (identified by the labels in the lower right of each graph). The frequency and its common logarithm are represented on the upper and lower x-axes, respectively. Plotted on the y-axis is the rescaled proportion of rewards harvested, out of a maximum of 20. The rescaling procedure sets the lower asymptote of each curve to zero and the upper asymptote to one. Each data point is the mean of six observations; error bars represent the standard error of the mean. Rats were exposed to several different train durations: 0.125 s (purple), 0.25 s (red), 0.5 s (green), 1 s (blue), 2 s (cyan), 4 s, (magenta), and 8 s (orange).

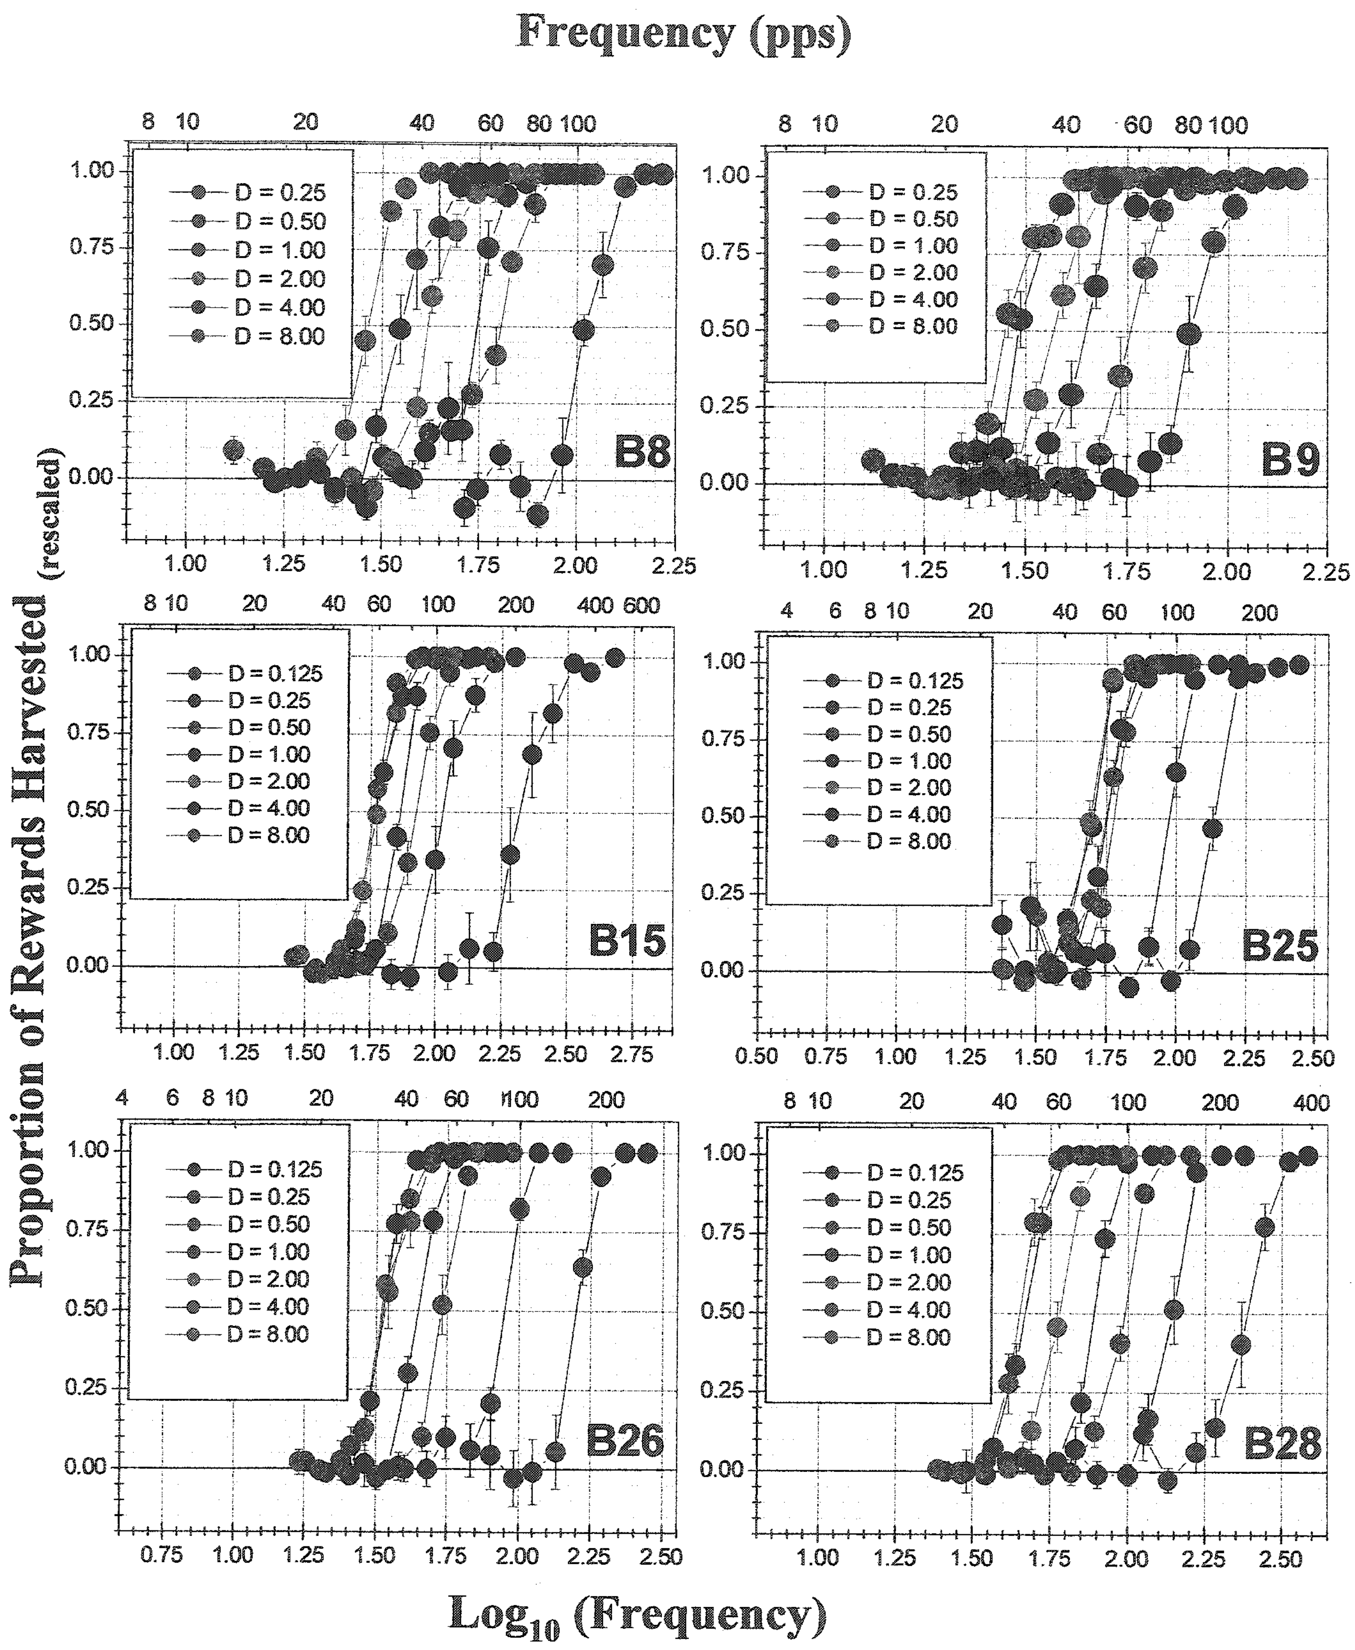

Figure 4. Rescaled data results from the single-operant frequency sweeps test.

curves all overlap, suggesting that for these rats, duration neglect was seen once the trains exceeded 2 s in duration. Thus, for these 4 subjects, once the train duration became longer than 2-4 seconds, further increases in train duration had no further effect on performance. Of the remaining 2 rats, rat B8 had apparently not yet reached the point at which duration neglect was displayed, as there is still a sizeable distance between the 4 s and 8 s curves, and rat B25 the effect of train duration did not appear to be monotonic, in contrast to the other subjects, as the 4 s and 8 s curves fall to the right of the 1 s (blue) and 2 s curves.

The next step in the frequency sweep experiment was to assess whether the results would follow a rectangular hyperbola, thus replicating Gallistel's (1978) findings. The 50% performance (half-maximal) criterion was selected, and the common logarithm of the frequency (in pulses per second) required to maintain this criterion performance was plotted against the train duration, as shown in Figure 5. The rectangular hyperbolic function (green curve) fit to the data (open red circles) is indeed a good approximation of the data, particularly for rats B9, B15, B26, and B28; the 6 open circles at each train duration are from each of the 6 sweeps (frequency sweeps) making up the performance curves. The remaining 2 subjects displayed systematic yet small deviations. As can be seen from the figure, the fit of the hyperbolic function deviates from the 4 s and 8 s curves for rat B8, and from the 1 s, 4 s, and 8 s, curves for rat B25. In all cases the coefficients of determination are quite high, ranging from 0.884 (B25) to 0.988 (B9); a complete list of the coefficients of determination is shown in Table 1.

The results from the frequency sweeps study thus replicate Gallistel's (1978) results: the relationship between required frequency and train duration approximated a rectangular

Figure 5. The strength-duration fits arising from the single-operant frequency sweeps test (Sonnenschein, Conover, & Shizgal, 2003), for 6 subjects (identified by the labels in the lower right of each graph). The red open circles represent the common logarithms of the frequencies ("required frequencies") at which the rescaled reward harvest was 50% of the maximum available rewards. The train duration (in seconds), is plotted on the x-axis. The green curve represents a rectangular hyperbola fitted to the required frequencies. Due to the heteroscedasticity evident in Figure 4, each point was weighted by the inverse of the variance for that train duration, with a maximum weight of 10. Copyright © 2003 by the American Psychological Association. Reprinted with permission.

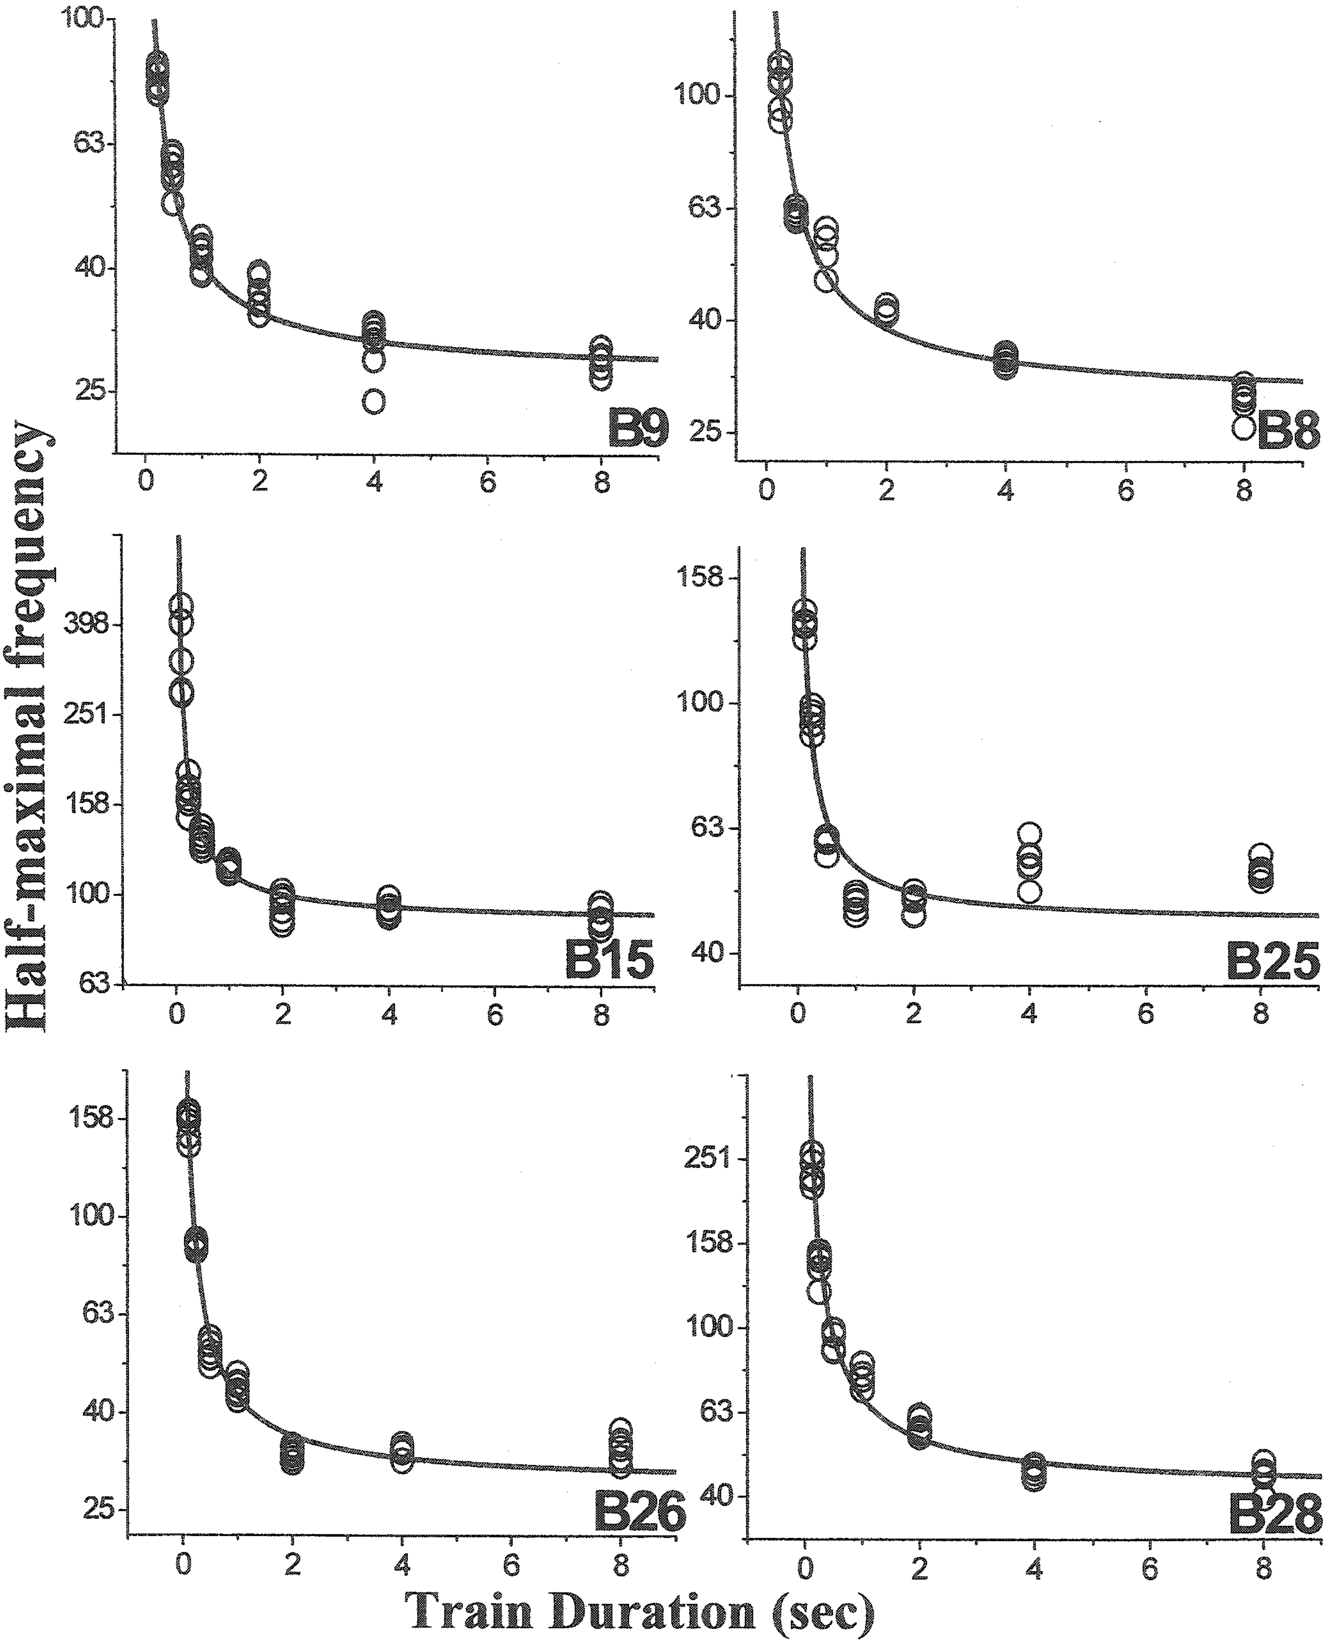

Figure 5. The strength-duration fits arising from the single-operant frequency sweeps test.

Table 1.

Results of Fitting the Half-Maximal Performance Data Points with a Two-Dimensional Rectangular Hyperbola

Rat	Coefficient of Determination	Rheobase	Chronaxie
B8	0.951	32.82	0.453
B9	0.988	27.06	0.503
B15	0.928	55.77	0.268
B25	0.884	44.83	0.224
B26	0.963	28.71	0.493
B28	0.973	41.59	0.640

hyperbola. In addition, although duration neglect appeared to take somewhat longer (2-4 s) to be displayed in the frequency sweeps experiment, the range overlaps those reported by Gallistel and colleagues (Gallistel, 1978; Mark & Gallistel, 1993; Shizgal & Matthews, 1977) (1-2 s).

A subsequent study was designed to replicate a key feature of Mason and Milner's (1986a, 1986b) experiments. The same 6 rats were used and in this case during the experimental sessions the frequency was now held constant while train duration was varied in a descending series. Once again, all other characteristics of the stimulation were held constant. Furthermore, in this "duration sweeps" experiment, two frequencies were employed: a "high" frequency and a "low" one, with the exact values of the frequencies for each individual subject determined by their prior performance in the frequency sweeps experiment.

The results of the duration sweeps experiment are shown in Figure 6. The y-axis once again represents the Proportion of Rewards Harvested (out of a maximum of 20), and the x-axis now represents the train duration. In the case of the high frequency (black) curve, performance approaches asymptote rapidly, with subjects B8, B9, B15, and B25 performing at maximum even when the train durations are merely 0.75 s to 1 s in duration. In the remaining 2 rats, performance approaches asymptote at 3 s (B28) and 5 s (B26). In the case of the low frequency (gold curve), performance increased much more slowly. For half the animals (B8, B25, and B28), performance approached asymptote at approximately 8 s, while in the remaining subjects (B9, B15, and B26), performance still had not approached maximum even at the 8 s duration, the longest duration tested. Thus, if examined from the point of view of duration neglect, when the high frequency trains were available, duration neglect was observed

Figure 6. Results of the duration-sweeps experiment (Sonnenschein, Conover, & Shizgal, 2003), for 6 rats (identified by the labels in the lower right of each graph). The x-axis represents the train duration (in seconds). Plotted on the y-axis is the mean rescaled proportion of rewards harvested, out of a maximum of 20. Error bars represent the standard error of the mean. The lower asymptote of the curves was set to zero by the rescaling procedure; no upper asymptote was assumed, and the maximum proportions were free to vary. The black curves represent the high frequency, and the gold curves represent the lower frequency. The high and low frequencies for each rat were determined individually, based on their performance in the prior frequency-sweeps experiment. Each data point represents the mean of six observations. Copyright © 2003 by the American Psychological Association. Adapted with permission.

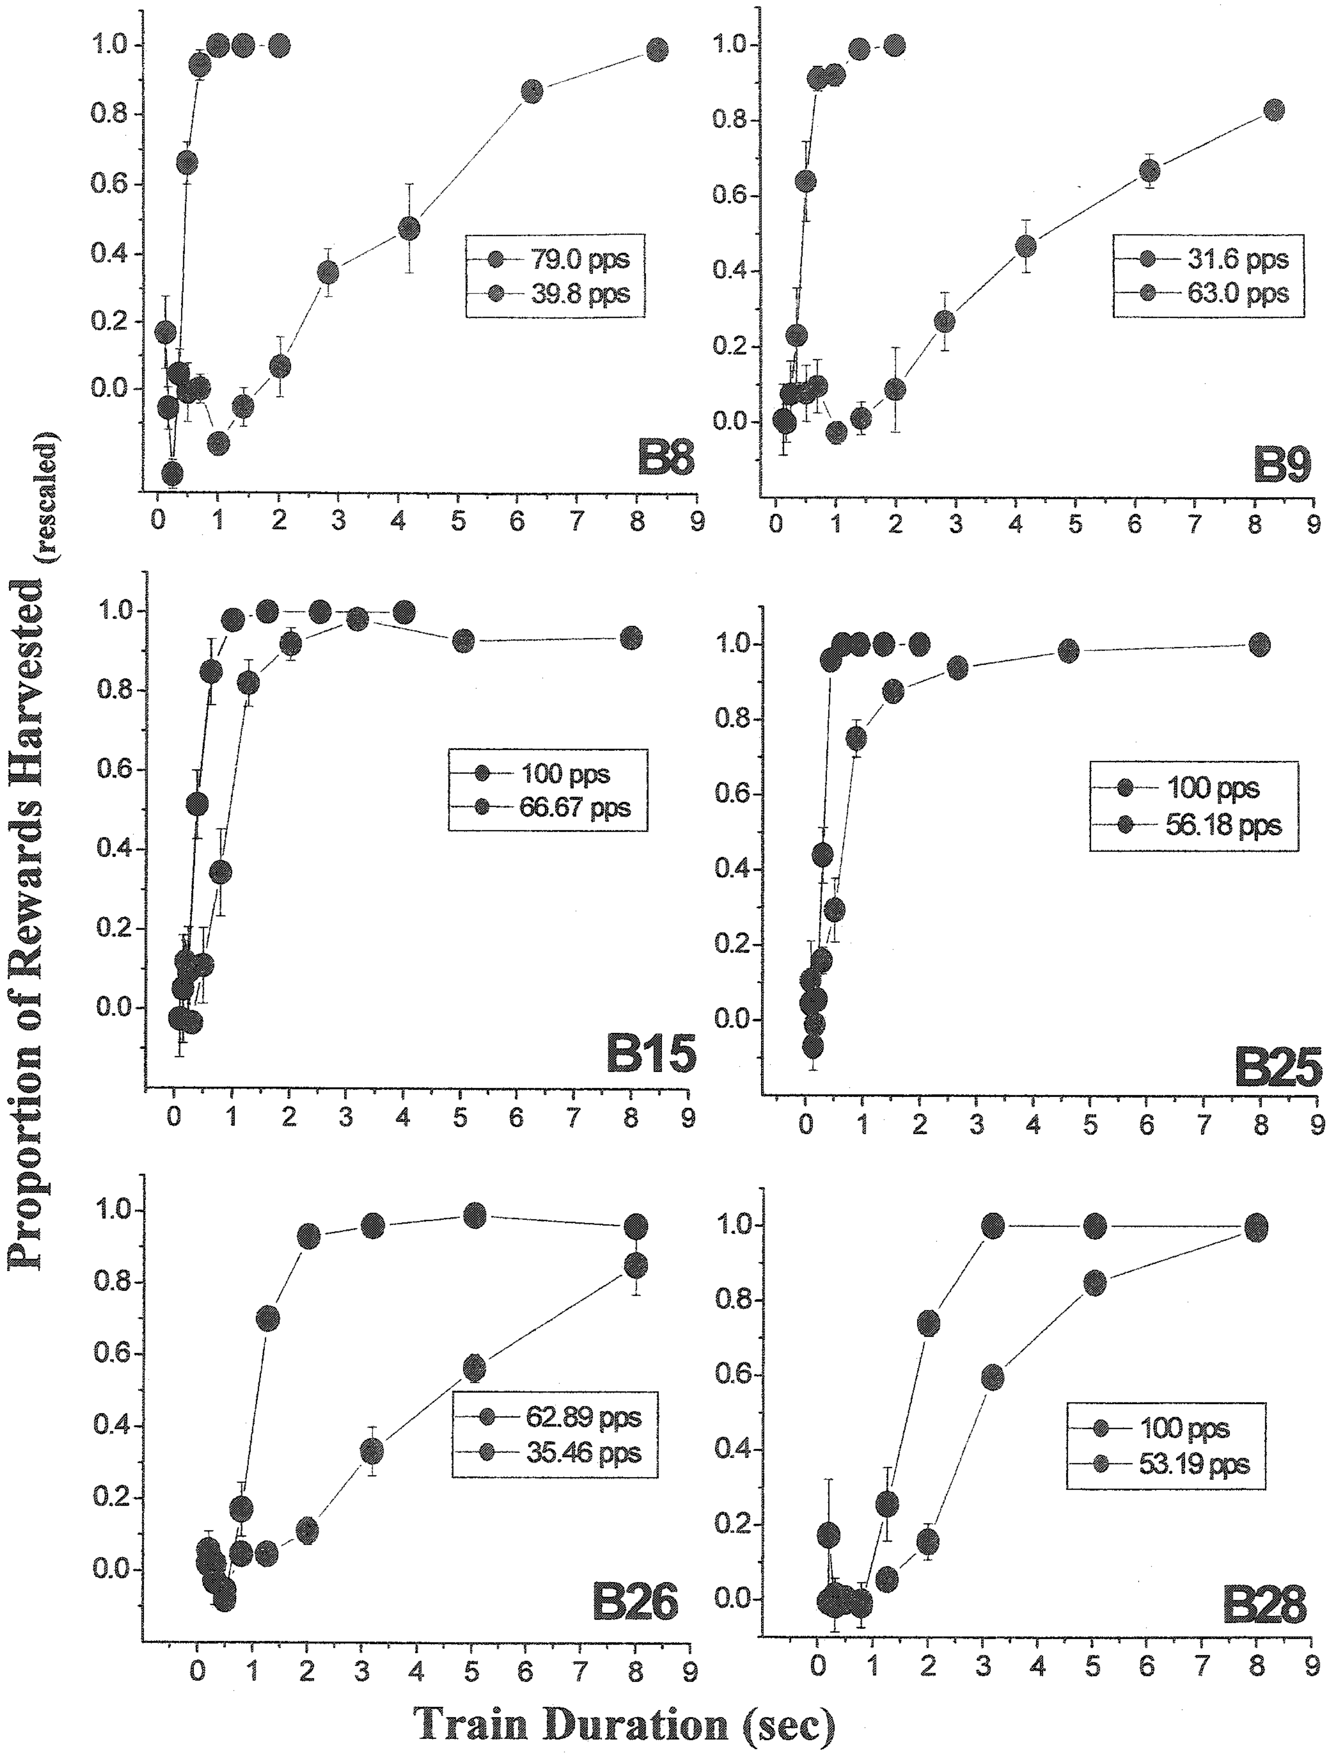

Figure 6. Rescaled data results from the duration sweeps test.

once the train durations exceeded 0.75-1 s, whereas when a lower frequency was available, duration neglect was observed only when the trains reached a duration of 8 s, if indeed duration neglect was demonstrated at all at the low frequency.

The results from the duration sweeps experiment are in line with those of Mason and Milner (1986a, 1986b): at higher frequencies, duration neglect appears at shorter train durations than at lower frequencies. Therefore, the apparently discrepant results of both camps (Gallistel, 1978; Mark & Gallistel, 1993; Mason & Milner, 1986a, 1986b; Shizgal & Matthews, 1977) were replicated.

The final step in analyzing the results of both the frequency and duration sweeps was to see whether the Shizgal Triple-Logistic (TL) Model could account for the results. Recall that Shizgal developed a model which attempted to integrate the findings of both Gallistel (1978) and Mason and Milner (1986a, 1986b) within a single neural-integrator model of BSR. This model attempted to not only integrate Gallistel's (1978) hyperbolic function with Mason and Milner's (1986a, 1986b) observations, but also to link the subjective experience of the reward (which cannot be directly measured) to observable performance.

The results from both the frequency sweeps and duration sweeps experiments were plotted in the same space, and the TL model was fit to the data, yielding the results shown in Figure 7. The x-axis variable is the common logarithm of the train duration in seconds, the y-axis variable is the common logarithm of the frequency in pulses per second, and the z-axis variable is the Proportion of Rewards Harvested (PRH). The fit of the TL model is

Figure 7. Results of fitting the rescaled data from both the frequency- and duration- sweeps experiments (Sonnenschein, Conover, & Shizgal, 2003) with the Shizgal TL model, for 6 subjects (identified by the labels in the lower right of each graph). The x-axis represents the common logarithm of the train durations (in seconds), the y-axis represents the common logarithms of the frequency (in pulses per second), and the z-axis represents the mean rescaled proportion of rewards harvested, out of a maximum of 20. Train duration (in seconds) is represented by coloured spheres: 0.125 s (purple), 0.25 s (red), 0.5 s (green), 1 s (blue), 2 s (cyan), 4 s (magenta), and 8 s (orange). In addition, the high-frequency data from the duration-sweep experiment are represented by black spheres, and the low frequency data are represented by gold spheres. The fit of the Shizgal TL model to the data is represented by the black grid. Each data point is the mean of six observations. The inverse of the variance was used to weight the surface fits. Copyright © 2003 by the American Psychological Association. Reprinted with permission.

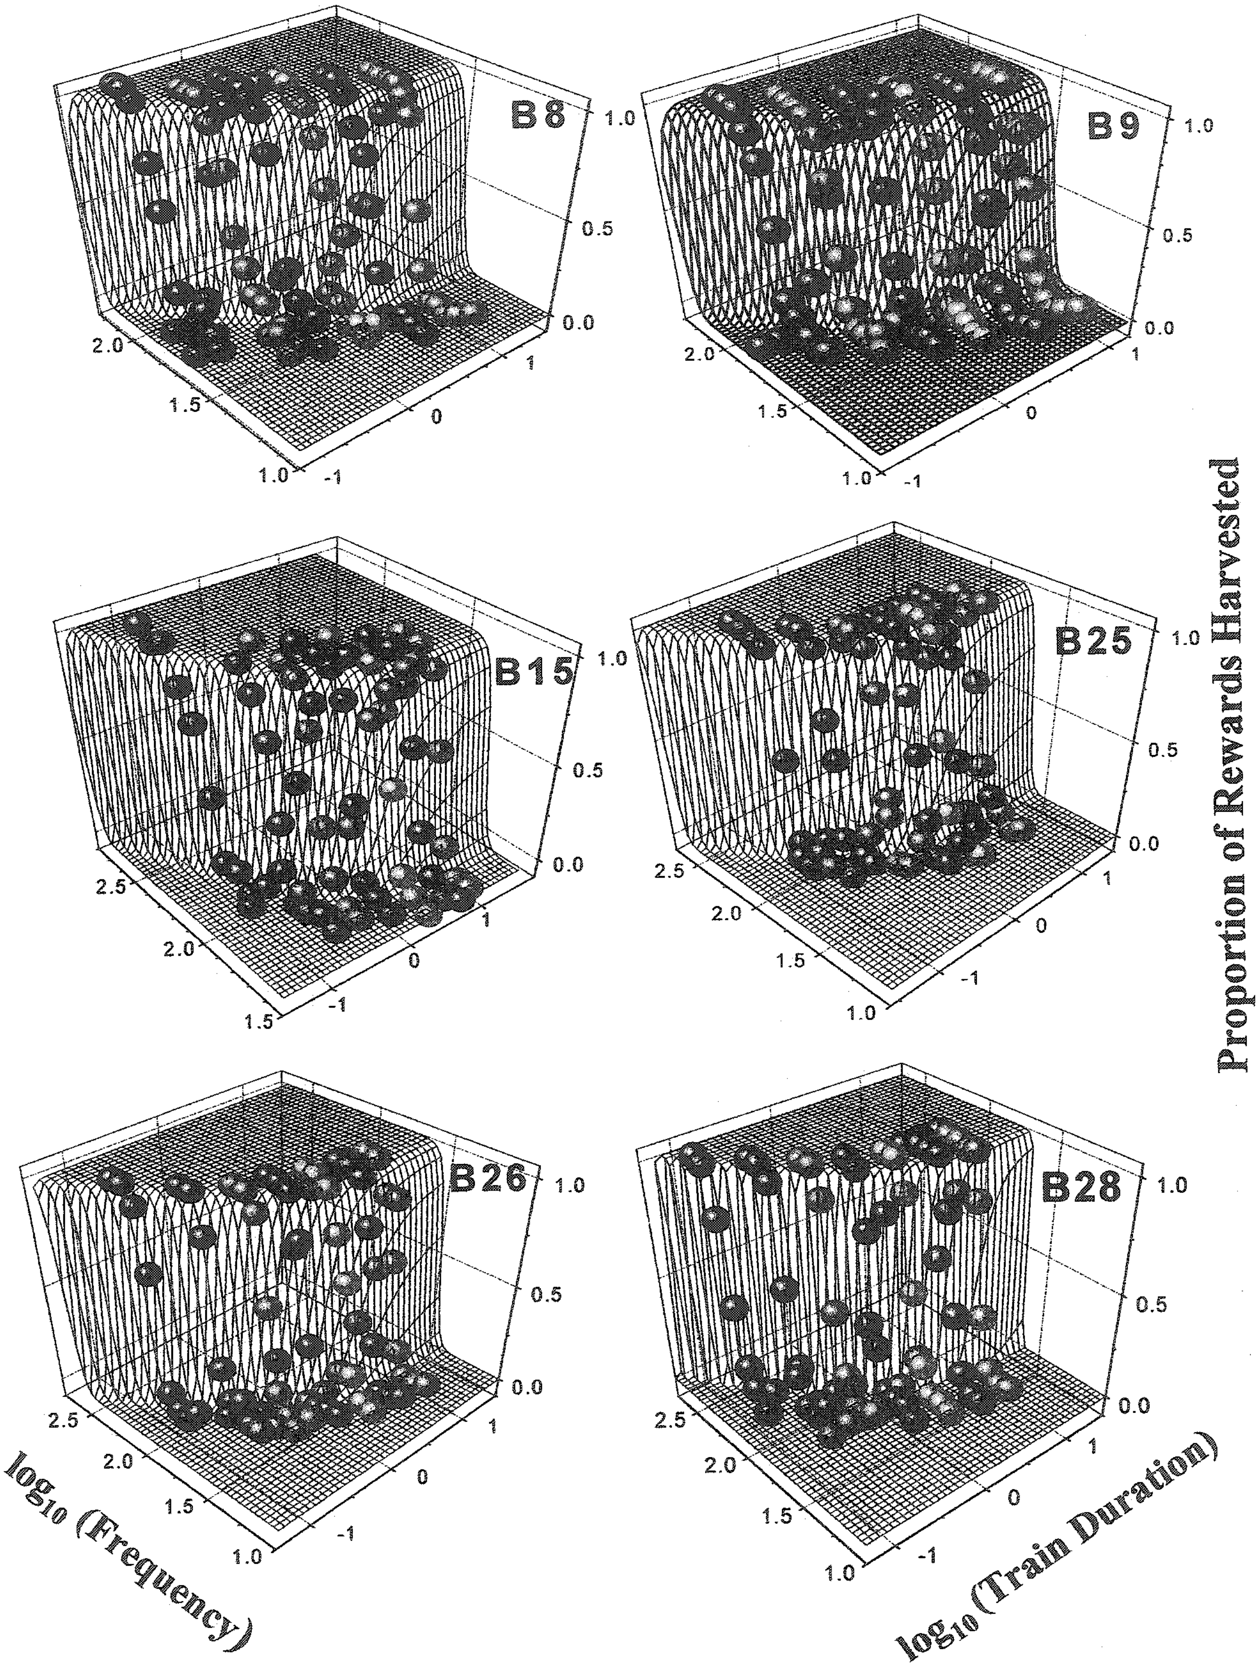

Figure 7. Results of fitting the rescaled data from both the pulse frequency and duration sweeps experiments using the Shizgal TL model.

once again represented by the black grid. There was quite a high correspondence between the data points and the fitted surface of the TL function, which yielded coefficients of determination ranging from 0.958 (B26) to 0.985 (B15) (the complete list of coefficients of determination is shown in Table 2). Furthermore, the close approximation of the fits to the data suggests that regardless of whether the frequency was held constant and the train duration was swept, or vice-versa, a given combination of train duration and frequency led to similar levels of operant responding.

Thus, it appears that the discrepancy between the two camps of researchers is not so much a discrepancy, but an illusion created by the fact that each camp of researchers was effectively looking at the 3-D "mountain" from only one or more 2-D perspectives. Basically, one needs to look at the 3-D structure in its entirety in order to fully appreciate the contributions of both train duration and frequency to the display of duration neglect. These two experiments also helped to resolve the issue of a single neural integrator (Gallistel, 1978) versus two neural integrators (Mason & Milner, 1985). Given that the discrepancy between the two camps on the issue of duration neglect is an illusion, and thus the results of Mason and Milner (1985) are really no different from those of Gallistel (1978), this suggests that Gallistel's (1978) simpler single-integrator theory is sufficient to account for both results.

To summarize, the TL model is able to account for the findings of both Gallistel and colleagues and Mason and Milner. Thus, both progressive insensitivity to duration, and the fact that this insensitivity occurs more rapidly when a high stimulation strength is

Table 2.

Results of Fitting the Data from Both Experiments with the Triple-Logistic Shizgal Model

Rat	Coefficient of Determination	Rheobase	RRR	g	p
B8	0.976	35.60	0.377	4.406	7.487
B9	0.984	34.32	0.342	3.025	7.432
B15	0.985	73.53	0.010	15.138	1.214
B25	0.973	61.48	0.020	12.476	1.531
B26	0.958	43.37	0.073	7.718	2.988
B28	0.984	43.90	0.554	15.065	9.418

used, were not only replicated empirically, but could also now be accounted for within a single model.

The Temporal Compression Rule: Peak-Only versus Peak-and-End

Although the TL model clarified the relative contributions of train duration and stimulation strength to the time course of reward integration, there is another aspect of evaluating reward which is equally important. When integrating reward value over time, a "temporal compression" rule of some sort is needed to boil down the set of instantaneous utilities associated with an experience (such as a BSR train) into a single value.

Gallistel (1978) argued that a peak-only model is sufficient to account for how organisms evaluate appetitive experiences. BSR models to date, including the TL model, have generally assumed that it is only the "peak" reward value of the stimulation which is recorded into memory, in line with Gallistel's view.

In contrast, Daniel Kahneman and colleagues (Fredrickson & Kahneman, 1993; Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993; Redelmeier & Kahneman, 1996) have argued that the rule an organism uses to remember a past experience is a "peak-and-end" rule. They have argued that although an experience, be it aversive or appetitive, is experienced in a moment-to-moment fashion (instantaneous utility), the overall goodness or badness of the experience is not recalled in the same moment-to-moment manner. Instead, they assert that a two-step process is in operation. First, organisms take two exemplars from the experience, the “peak” pain or reward (depending on the context) of the experience, and the amount of pain or reward at the “end” of the experience. Second, an intermediate value between these two exemplars is selected and recorded into memory (remembered utility) and used in later decision-making and behaviour. Such models are usually referred to as “representation by exemplar” models, and the “peak” model of Gallistel (1978) can also be conceived of as a “representation by exemplar” model. These models are important because when organisms are evaluating a past experience, prodigious mnemonic resources and considerable time would be required to “replay” the entire experience in a real-time fashion. It is much more efficient to “boil down” a temporally-extended experience into a few key exemplars, which the neural system can then use to compute a rating of the overall “goodness” or “badness” of the experience. In this vein, Gallistel’s (1978) peak-detection model and Kahneman’s (Kahneman et al., 1993) peak-and-end model are both aimed at assessing which of the candidate segments of an experience are most likely to be used as exemplars, when an organism is performing evaluations.

Kahneman’s group has performed a number of experiments that appear to support their contention that the two important exemplars are indeed the “peak” and “end” values of an experience. In one study (Redelmeier & Kahneman, 1996) human patients undergoing colonoscopies and lithotripsies, both uncomfortable and often painful medical procedures, were instructed to make real-time evaluations of the intensity of pain during the procedure, as well as a retrospective evaluation after the procedure was concluded. The results demonstrated that the patients’ retrospective evaluations were highly correlated with both the "peak" pain, and pain experiences at the closing moments of the procedure ("end").

Another aspect of the Kahneman "peak-and-end" model, which was also evident in this study, was the minimal role duration played in forming retrospective evaluations. The results of the Redelmeier study revealed that there was no significant correlation (r=0.03 for colonoscopies, r=0.08 for lithotripsies) between duration of the medical procedure and the retrospective pain ratings, despite the fact that the procedures the patients underwent varied widely in duration, with some patients undergoing procedures as brief as 4 minutes, to as long as 67 minutes. Kahneman and colleagues termed this indifference to the duration of an experience as "duration neglect" (analogous to the duration neglect in BSR studies described earlier). This neglect is an inevitable consequence of the "peak-and-end" cognitive heuristic, which boils down the experience to a small set of exemplars and does not include a representation of the duration.

Another, more rigorously-controlled study by Kahneman, Frederickson, Schreiber, and Redelmeier (1993), revealed similar support for the idea of "peak" and "end" as the key exemplars, and the neglect of duration information as a central feature. In this experiment, human subjects placed their hands in cold water during two types of trials. In one trial ("short" trial) the subjects were required to keep their hands in the water for 60 s while the temperature of the water was held constant, whereas in the second type of trial ("long" trial) the subjects had to keep their hands in the cold water for 90 s. In this "long trial", the first 60 s of the trial were identical to the "short" trials in terms of water temperature, but during the additional 30 s the water was warmed slightly. The subjects were then led to believe that they would have to experience a third cold water trial, and were given the opportunity to choose whether to repeat the short trial or the long trial. Most of the subjects elected to repeat the long trial, which is in accordance with the predictions of the peak-and-end model. In the short trials, the "end" of the procedure was likely just as painful as the "peak" moment of pain, and thus this trial would have been remembered as more aversive, whereas in the case of the long trial, the "end, with its slightly less cold water, would have led to a lower overall remembered pain rating. In addition, this experiment also demonstrated the phenomenon of "duration neglect", as it seems counter-intuitive that the subjects would want to suffer through the longer trials, yet this is what they elected to do (although they were never actually required to experience a third trial). In fact, Kahneman et al. (1993) found that the duration of the trials accounted for a mere 2% of the variance in subjects' overall retrospective evaluations. Lastly, it is also interesting to note that even though some subjects indicated there was no lessening of discomfort in the long trial, this was not a problem for the peak-and-end model. It predicts in that case that the long and short trials should be approximately equally aversive, a finding which was indeed confirmed in this subset of subjects.

A third study by Frederickson and Kahneman (1993) involved the use of pleasant and aversive film clips to see if the peak-and-end model (and duration neglect) would also hold in this situation. Subjects were presented with brief, plotless, and emotionally evocative films designed to be either pleasant (e.g. penguins diving off a glacier) or aversive (e.g. a medical film of an amputation procedure), and were presented in either a short or long form (the short versions were one third as long as the longer versions). Subjects made moment-to-moment reports of their affective reaction to the clips, and were also asked to make global evaluations of the pleasure or discomfort they experienced while watching each clip (retrospective evaluation); these ratings were obtained immediately after each clip finished. The results demonstrated that the reported "peak" and "end" affects accounted for 82.8% of the variance in the global evaluation of the aversive films, and for 47.6% of the variance in evaluation of the pleasant films.

In summary, according to the Kahneman (Kahneman et al., 1993) model of peak-and-end evaluation, an organism tends to take two exemplars from their experiences, the "peak" moment of pain or pleasure, and the amount of pleasure/pain at the "end" of the experience, and select an intermediate value between the two exemplars. Information about the duration of the experience, however, appears to play a very minor or non-existent role when forming these evaluations. According to the model, because of this tendency to neglect duration information and retain only peak and end information, in situations in which the painful experience becomes somewhat less aversive towards the end, this experience would be remembered as less aversive than another experience in which the peak and end moments of pain are more similar. As alluded to above, a counter-intuitive prediction is suggested by this model: if two aversive experiences were delivered, and the one with the less painful end was actually longer in temporal duration than the experience with the more aversive end, organisms would still prefer the longer experience, even if in reality they are experiencing *more* pain, so long as the "end" of the experience is less aversive than the "peak" moment of pain. If this is indeed the case, one practical application of this in the domain of medical treatment would be to add a less painful "end" to procedures in which patient compliance is a problem. For example, in the case of colonoscopies, perhaps leaving the instrument motionless inside the patient (still uncomfortable, yet not as painful as the rest of the procedure) for a few moments after the procedure is completed will lead the patients to form a more favourable retrospective evaluation of the procedure (even if they are suffering longer than they need to). This has been demonstrated by Redelmeier, Katz, and Kahneman (2003).

If the peak-and-end model is also applied to appetitive experiences, such as BSR, a similar counter-intuitive prediction is obtained: If the organism is presented with two rewarding experiences, a short experience in which the "peak" and "end" moments of the experience are approximately equally rewarding, and a longer experience in which the "end" is less rewarding than the rest of the train, this model would predict that the shorter train would be preferred. Even though the organism is passing up extra seconds of stimulation, given that duration information is "neglected", this organism is selecting the experience with the higher average of "peak" and "end" reward values. Put another way, adding a weaker (less rewarding) "end" to a stimulation train would degrade the overall reward value of the train, in comparison to a shorter train with a more constant reward value.

Most BSR models have assumed Gallistel's (1978) view, that is, that the "peak" exemplar is the only one selected from the experience of the rewarding train and encoded into memory. However, past BSR studies have not been designed in such a way as to distinguish between the "peak"-only and peak-and-end (Kahneman et al., 1993) models. BSR studies have tended to use stimulation trains in which the frequency is held constant throughout, termed “rectangular” or “constant-frequency” trains. Thus, it is as yet unknown whether adding a weak “end” (more stimulation but at a weaker frequency) would impact the overall evaluation of the trains. Also, it is impossible to separate the two models using constant-frequency stimulation, as both models make identical predictions in that case.

The key test would be to compare performance for constant-frequency trains (analogous to the “short” trains described in the Kahneman group studies) against that for composite frequency trains, in which the frequency (and reward value) would be higher in one section of the train (the “peak”) than in a different section (the “end”). If indeed such a train were to be constructed, one could argue that the course of the subjective intensity of reward would be different for such a train than that of the rectangular train.

The TL model has clarified how reward grows when rectangular (constant-frequency) stimulation trains are employed. Reward value rises as the train progresses, and it reaches a maximum (“peak”) during train delivery. Once this train ends, the reward value begins to decay to a minimal level. .Although a composite train can be very simply made up of two rectangular trains, one at a high stimulation strength and a second at a lower strength, the TL model cannot account precisely for how the reward signal decays during the weaker segment.

As yet, no BSR researchers have employed composite trains of the type just described. However, a few researchers have used stimulation trains with interruptions (“holes”) in the middle of the train. Fouriezos (1995) conducted a study in which stimulation trains were separated into 2-burst pairs, separated by an inter-burst-interval (IBI). The first burst of each 2-burst pair was set at a subthreshold frequency, and the IBI was increased between the two pairs. The frequency of the second burst was manipulated, in order to determine what frequency of the stimulation was required to maintain responding by the rats, for various IBI durations. The results demonstrated that an increasing number of pulses was required in the second burst to offset the effects of increasing IBIs. With longer IBIs, as there would be more time for the rewarding effects of the first burst of the pair to decay, more pulses would be needed in the second burst to compensate for this decrease in reward value. Fouriezos went on to determine the time constants of this decay for various subjects, and concluded that the decay rate was on the order of tenths of a second.

It should be noted that Sax & Gallistel (1984) also utilized stimulation trains with "holes" in them. Inserting a hole with a duration up to as much as 9 seconds had no effect on the degree of summation, which was estimated from running speeds recorded in a straight alley. However, given that Fouriezos' (1995) data is more extensive and yet he failed to replicate the results of Sax and Gallistel (1984), the latter results will not be considered further here.

The results of Fouriezos' (1995) work suggest that if one wanted to create trains in which the reward value at the "end" of the train was noticeably different from the reward value of the "peak", a particularly long weak ending segment would not be needed. Given the average time-constant estimate (0.44 s), only a second or so would be required for the reward signal to decay to a low level by the end of the train (~5% after 1.32 s). Thus, if one were to construct a composite train in the fashion described above, made up of a strong rectangular train followed by a weak rectangular train, one would expect that the reward value would rise during the delivery of the first rectangular train, eventually reaching maximal reward value. Then, once the strong train terminated and the weaker one began, reward value would decrease rapidly. This reward signal would not decay down to a minimal level, however, as the weak train would still have sufficient strength to maintain reward, albeit at a lower level than the first component of the train. Thus, the decay of reward from the first train portion might well stop once the reward reached the level it would have reached, had the weak train been delivered alone.

The Current Study

Given that the issue of "duration neglect" had now been settled with respect to BSR through the use of the Shizgal TL model, the focus of the current study was as follows: Experiment 1 involved testing whether the Kahneman (Kahneman et al., 1993) peak-and-end model would hold when rats were presented with constant- and composite-frequency trains concurrently in a single-operant paradigm. It was hypothesized that if the addition of a weaker end indeed reduces the remembered value of the train, the curve relating performance to stimulation strength would be shifted to the right. The aim of Experiment 2 was to assess whether the Kahneman peak-and-end model would hold when rats were presented with both constant- and composite-frequency trains simultaneously in a dual-operant paradigm. Again, it was hypothesized that if the peak-and-end model was the heuristic in operation, a constant-frequency train of a given strength would be equivalent to a stronger composite train. Experiment 3 was a single-operant study aimed at assessing whether delivering composite-frequency trains with longer "ends" would make the effect of adding a weaker end more evident. In Experiment 4 the aim was to determine whether a different combination of exemplars, such as the "beginnings" and "peaks" of the stimulation trains, would better account for performance than the peak-alone model. It was hypothesized that if "beginnings" of positive experiences are indeed salient, stimulation trains with strong "beginnings" should be preferred over trains with either "weak" or "zero" beginnings. Lastly, the purpose of Experiment 5 was to assess the concordance of measures obtained in single- and dual-operant paradigms. In the frequency sweeps experiment described earlier, duration neglect set in by 2-4 s. Experiments in human subjects show inconsistencies between the evaluation of options presented singly and the same options presented in pairs. The experiment assessed whether a similar inconsistency would be observed in the case of self-stimulating rats.

Experiment 1

The aim of the first experiment was to determine whether Kahneman's (Kahneman et al., 1993) peak-and-end model would hold for self-stimulating rats.

The peak-and-end model has two aspects. The first is that the "peak" and the "end" are the two exemplars extracted from an experience, and which are later averaged to yield a retrospective evaluation of the overall value. Second, information about the duration of the experience is "neglected". In contrast, Gallistel's (1978) model suggests that the "peak" reward value is the only exemplar selected from the BSR train.

Past BSR research has not utilized trains that varied in their "peaks" and "ends", and thus a comparison between the two models has not been possible. The use of composite-frequency trains would be a better test, preferably using trains long enough to be within the range where duration neglect is displayed. The frequency and duration sweeps experiments described earlier demonstrated that duration neglect was observed once trains exceeded 2-4 s, given a sufficiently high frequency.

To create appropriate composite-frequency trains, the first part of the train would have to be at a sufficiently long duration so as to fall within the neglected range, as well as within the range of frequencies that would rapidly approach asymptote. This would constitute the "peak" part of the train. For the weak "end" of the train, it would be optimal to pick a train that falls far short of the summit of the 3-D "mountain" described by the TL model, and thus a train at a sufficiently lower frequency would be required. In addition, as the work of Fouriezos (1995) suggests, the weak train segment should exceed one second minimum, in order to allow sufficient time for the effects of the strong train component to decay.

If a weak train were appended to the end of the strong train, one might expect that the subjective intensity of reward would gradually increase as the strong train progressed, and then once it ended and the weak train segment began, the subjective intensity of reward should decrease (as it is now driven by a weaker input). The reward value might then level off at the maximal reward value that the weaker train is capable of inducing.

The results of the frequency and duration sweeps allowed for a clarification of the relationship between train duration and frequency. This understanding was then used to select appropriate frequencies for the composite-frequency train segments, such that each segment would have differing reward intensities. At the same time, the previous experiments also indicated which frequency/train duration combinations for the composite trains were most likely to fall within the range where duration neglect would be observed. Use of the Shizgal TL model therefore allowed for the identification of appropriate frequencies and train durations for each rat, to generate differing "peaks" and "ends".

The single-operant peak-and-end test (Experiment 1) was designed to be analogous to the "short" and "long" trial types described in the Kahneman, Frederickson, Schreiber, and Redelmeier (1996) cold water study described earlier. Specifically, the constant-frequency trains, with their equally rewarding "peaks" and "ends", were intended to be analogous to the short trials, and the composite-frequency trains with their weaker “ends”, were intended to be analogous to the long trials.

The aim of the present experiment was to determine which model, peak-alone or peak-and-end, constitutes the most parsimonious explanation of how retrospective evaluations of BSR were performed by the rats responding for BSR. If the peak-and-end model held, it was predicted that the psychometric functions (performance curves) for one or all of the composite-frequency trains would be rightward-shifted along the frequency axis, compared to the constant-frequency trains. Such rightward shifts of the composite-frequency trains would indicate a reduction in their reward value. If the psychometric functions for all train types overlapped, this would suggest that the peak-only model was sufficient.

Method

Subjects

The subjects were 6 male Long-Evans rats from the Charles River breeding farms (St-Constant, Quebec). One of the rats (B9) had served in two previous experiments (Sonnenschein, Conover, & Shizgal, 2003), and the remaining subjects had run shortened versions of the frequency sweeps experiment to ensure that they were behaving as expected in terms of duration neglect. Rats weighed between 300-500 g at the time of surgery, and were housed individually in plastic cages with water available *ad libitum*. The colony room was maintained on a 12-hr dark/12-hr light cycle. All testing was conducted during the dark part of the cycle. To prevent the rats' weight from affecting their responding for BSR (Abrahamsen, Berman, & Carr, 1995; Fulton, Woodside, & Shizgal, 2000), once subjects' weights exceeded the cut-off[2] (650 g), food intake was restricted in order to maintain them at that body weight.

Surgery

Rats were given an injection of atropine sulfate (0.05 mg/kg) subcutaneously, to reduce mucus secretions. Twenty minutes later, sodium pentobarbital (Somnotol, 65 mg/kg) was

---

[2] As there is little data available on the "normal" weights of Long-Evans strain rats beyond 15 weeks of age (approximately 450 g; see the Charles River website at http://www.criver.com/03CAT/rm/rats/longevansRats.html), the 650 g cut-off was decided upon in consultation with S. Fulton (personal communication). Restricting the subjects' weight was intended to avoid the possibility that rats with electrodes in food-restriction-sensitive sites would show BSR threshold increases as they aged and gained weight (Abrahamsen, Berman, & Carr, 1999; Fulton, Woodside, and Shizgal, 2000). Given the fact that restriction-sensitive areas lie mere fractions of a millimeter from the LHA coordinates used in the present study, this precaution was deemed appropriate.

administered intraperitoneally to induce anesthesia. Supplements were administered during the surgery as required. Chronic stimulating electrodes were implanted bilaterally and aimed at the lateral hypothalamic area (LHA) according to the Paxinos and Watson (1998) coordinates: 2.8 mm posterior to bregma, 1.7 mm lateral to the mid-sagittal sinus, and 7.8 mm below the dura. Electrodes were constructed from 0.25 mm stainless-steel insect pins that were insulated with Formvar. The bottom 0.5 mm of the electrodes were bared of insulation. The return for the stimulation current was made up of two stainless steel jeweler's screws fixed in the skull, around which the return wire (the ground) was wrapped. The electrodes and return wire were secured to the skull with dental acrylic with the male Amphenol pins (Arrow Electronics) attached to electrodes and the ground inserted into a McIntyre miniature connector (Science Technology Centre, Carleton University, Ottawa) that was also secured to the skull with dental acrylic. For all behavioural testing the caps were attached by a threaded ring to a matching connector mounted at the end of the stimulation cable. Subjects were allowed to recover for a minimum of three days before training began.

Apparatus

Screening and training. Single-lever operant boxes were used for initial screening and testing for self-stimulation. These boxes were constructed out of wood, with dimensions 27.4 cm wide by 26 cm deep by 64 cm high, with Plexiglas front panels and one non-retractable lever (made by Concordia University's Science Technical Centre), positioned 6 cm above the wire mesh flooring. There was also a small keylight positioned 4 cm above the lever, which was illuminated whenever stimulation was available. The stimulation cable was attached to the

McIntyre miniature connector on the rats' heads, with the other end of the cable attached to the stimulator by a slip-ring assembly (Airflyte Electronics, New Jersey) mounted in the center of the ceiling of the operant chamber. All temporal parameters of the stimulation were set by hand-set digital pulse generators, and the stimulation current was regulated by constant-current amplifiers (Mundl, 1980). Build-up of charge at the interface of brain and electrode was prevented by a circuit which shunted the stimulator outputs through a 1-kΩ resistor when no pulses were delivered. The stimulation current was monitored using a Metermaster MM200 oscilloscope, by measuring the voltage drop across a 1-kΩ resistor (1% precision) in series with the subject. Stimulation consisted of trains of rectangular cathodal pulses, with the pulse duration always set at 0.1 ms, and the intensity set at 400 μA.

Experiments. Dual-lever operant boxes were used, which were constructed of welded aluminum frames with gray PVC plastic panels mounted on the sides and back, and clear PVC plastic panels mounted on the front of the chamber. The dimensions of these boxes were 33.0 cm wide by 23.5 cm deep by 60.5 cm high. The boxes were equipped with two retractable levers (MED Associates Inc., ENV-112B) positioned across from each other, and 10 cm above the mesh flooring. Keylights (approximately 1 cm in diameter) were positioned 6 cm above each lever, and a houselight (10 cm wide by 4.5 cm high) was positioned on the rear wall, 38 cm above the floor. The keylights above each lever were illuminated whenever the stimulation was available, and the houselight flashed on and off throughout the inter-trial-interval (ITI). The stimulation cable was attached to the rats' McIntyre miniature connectors. To allow the subject to move freely within the cage without tangling or twisting of the lead, phone handset swivel jacks (Archer "Untangler", model no. 279-299) were used as slip rings.

Stimulation current was monitored using Tektronix 2205 oscilloscopes, by reading the voltage drop across a 1-kΩ resistor in series with the rat. The intensity of the stimulation was regulated and monitored as described above. All temporal characteristics of the stimulation were set using custom-developed software programs written by CSBN System Manager Steve Cabilio.

Procedure

Screening and training. Rats were initially screened for the presence of aversive or disruptive stimulation-induced movements, by stimulating them with 0.5 s trains of 40 cathodal pulses, set at an intensity of 400 μA. If no disruptive movements or aversive effects were observed, the stimulation frequency was gradually increased to the highest value the rats could tolerate, and the rats were then trained using conventional shaping procedures to press the lever to obtain brain stimulation; trains were delivered to the rat upon successively closer approximations to pressing the lever (facing the lever, approaching the lever, putting their paw on the lever), until the animals were self-stimulating reliably. Only unilateral stimulation was used in the present study, and the electrode selected was the one with the least motoric or aversive effects, and/or the electrode which supported the most robust responding. The stimulation consisted of 0.5 s trains set at 400 μA, and generally in the range of 60 to 40 cathodal pulses (frequency range 118 to 78 Hz), with the final stimulation frequency used dependent upon the individual rat. The pulse duration was always set at 0.1 milliseconds.

On the subsequent 2-5 days, more shaping was conducted if the rats were not bar-pressing reliably. In addition, the animals were habituated to "blackout delays" after each bar press. During the initial screening and training described earlier, after each bar-press the stimulation became available again after only a 0.6 s delay (during which the keylight was extinguished). In the experiment proper, blackout delays of 16 s would be used, and thus initial blackout training was employed to habituate the rats to longer delays between one bar-press and the next available reward delivery. The rats were exposed to several increasing blackout delays (first 1 s, then 15 s, and lastly 30 s), and once they reliably returned to the lever after several presentations of the longest blackout delay, the next phase of training began.

In the next phase of training, the frequency of the trains was "swept" from trial to trial; the parameters of the stimulation at the beginning of each session were set to the maximum tolerable value used during initial screening (i.e., stimulation for which the rat was responding vigorously and there were no strong aversive effects or disruptive movements), and on subsequent trials the frequency was decreased by 0.05 $\log_{10}$ unit steps, until the stimulation delivered proved insufficient to support steady responding. Three of these sweeps were run each training session, and training continued for 2-4 days, until the rat was not appreciably disrupted by the period of low reward, and would re-institute responding at the beginning of the next sweep.

The last phase of training was a final screening to confirm that the rats could withstand stimulation trains up to 8 seconds in duration. Stimulation similar to that used in prior training was used once again, except now the train duration was gradually increased. The train durations sampled in this phase were 1, 2, 4, and 8 seconds. In the case of all trains, the frequency for each duration was calculated in accordance with predictions based on Gallistel's (1978) strength-duration experiments. As noted earlier, the stimulation intensity was set to 400 μA.

Baseline condition. This condition was run daily for 18 minutes before the appropriate experimental session, in order to track any changes in the rats' performance over time (caused by moving of the stimulating electrode, loosening of the headcap assembly, etc.). These tests took place in the dual-lever chambers described earlier. Upon introduction to the dual-lever boxes, rats were trained to lever-press for 0.5 s descending frequency sweep sets, identical to those described earlier. Once responding had stabilized (i.e., operationally defined as less than 0.05 $\log_{10}$ unit shifts observed in the performance curves over repeated sessions), a 5 s "blackout delay" (BD) was instituted: once the lever was pressed, it retracted into the wall of the chamber, and only became available again once the BD had timed out. For the duration of the BD, the keylight above the lever was extinguished. The frequency was once again "swept" in descending order during the experimental "sweep"; on each successive trial the frequency was reduced by approximately 0.05 $\log_{10}$ units for rat B9, and by approximately 0.075 $\log_{10}$ units for all remaining subjects. All other stimulation parameters were held constant, as noted earlier.

A single "priming" train was delivered during each inter-trial-interval (ITI), 15 s after the start of the respective ITI. The baseline condition began on an ITI, and the ITI was 20 s long. As each trial was 100 s in duration, rats could harvest a maximum of 20 rewards per trial.

When the steps were 0.05 $\log_{10}$ units (rat B9), 12 trials were run each sweep to yield a complete performance curve; for all other subjects (0.075 $\log_{10}$ unit steps), 9 trials were required to yield the complete curve. The priming train characteristics (i.e., stimulation current, frequency, and train duration) were always identical to those of the stimulation trains that were available during the trial that immediately followed, with the exception that priming train delivery was not contingent upon a lever-press by the subject. In fact, levers were not available at all during the ITI.

Single-operant peak-and-end test. Although the present experiment took place in operant boxes equipped with two levers, only one lever was available during the sessions, while the other was kept retracted and its accompanying keylight left extinguished over the course of the trials. As in the baseline condition, the houselight was flashed during the 20 s ITIs, levers were retracted into the walls, and one priming train, identical in stimulation parameters to the reward that would be available during the subsequent trial, was delivered noncontingently. However, the priming train began 5 s after the start of the ITI. Each ITI began immediately after the termination of the previous trial, and all trials were 320s in duration. Either 13 or 9 trials were run each sweep, depending upon whether the frequencies decreased in 0.05 $\log_{10}$ unit steps (rat B9) or 0.075 $\log_{10}$ unit steps (all other subjects), respectively. Daily experimental sessions were made up of three sweeps.

A "blackout delay" (BD) of 16 s was employed. This duration was selected in order to allow enough time for the decrease of any aversive build-up, particularly in the case of the longest train duration used (8 s), yet without making the test session unnecessarily long. The use of a BD also held the maximum rate of reinforcement constant; as with the trial time set to 20 times the BD (320 s), the maximum number of rewards that could be earned was fixed at 20 rewards per trial.

Five types of stimulation trains were presented to the rats over the course of the experiment: a constant-frequency 8 s train, a constant-frequency 6 s train, and three composite-frequency trains. In the latter trains, the first 6 s were set at a constant frequency (identical to that used in the constant-frequency 6 s train, at a comparable "step" in the descending sweep), and the last 2 s were at a lower constant frequency. The ratio of frequencies of the first 6 seconds of the composite trains to the last 2 seconds were 4:1, 2:1, and 1.4:1 for the first, second, and third composite trains, respectively. Each experimental sweep consisted of a single descending frequency sweep (made up of the appropriate constant or composite set of frequencies), with all other characteristics of the stimulation held constant. The subjects were exposed to only one of the five stimulation trains during an experimental session, and were required to reach a stability criterion with each train type before they were presented with another train type.

Results

The first trial in each sweep was considered as a warm-up trial, and performance during this trial was not used in subsequent statistical analyses. All analyses were conducted using data averaged over six sweeps (i.e., six individual descending frequency sweeps), to minimize some of the effects of normal daily variation between individual sessions.

The raw data from the single-operant peak-and-end test are shown in Figure 8. The y-axis variable is the proportion of rewards harvested (PRH), and the x-axis variable is the common logarithm of the frequencies in pulses per second (pps). This figure reveals that while the upper asymptote (20/20 rewards harvested by the rat) was nearly always reached in each condition, the lower asymptotes had a tendency to be more variable. In addition, the variability of the data points tended to be greater on the rising segments (slopes) of each performance curve, compared to data points lying on the upper asymptotes of each curve.

In the work of Sonnenschein, Conover, and Shizgal (2003), similar issues of variability came up and were dealt with by rescaling and weighting each curve to remove some of the effects of this variability. The first step was to estimate a new lower asymptote using a 3-segment "broken-line function". This function consists of one horizontal line fit to the data points scattered at the lower end of the performance curve, one horizontal line fit to the data points scattered at the upper end of the curve, and one rising segment or "riser" fit to the data points scattered along the "slope" of the performance curve, thus linking both of the horizontal segments. A sample broken-line fit to illustrative data is shown in Figure 9. The next step was

Figure 8. Graphs of the raw means from the "6+2" variant of the peak-and-end test, for 6 rats (individual subjects are identified by the label in the lower right of each graph). The independent variable is the common logarithm of the pulse frequency, in pulses per second (the represented frequency corresponds to the value used for the 6 s portion of the composite train). The dependent variable is the proportion of rewards harvested, out of a maximum of 20. Error bars represent the standard error of the mean. Each performance curve represents a different train duration/type: 8 s (red) or 6 s (yellow) constant-frequency trains, and three composite "6+2" trains. For each composite, the first 6 s portion of the train is identical to the 6 s constant-frequency train, and the last 2 s portion is set at a lower frequency. The ratio of the frequencies in the first component of the train to the second component of the train was set to either 2:1 (green), 4:1 (magenta) or 1.4:1 (blue). Each performance curve represents responding averaged over six descending frequency "sweeps" at that particular train duration/type.

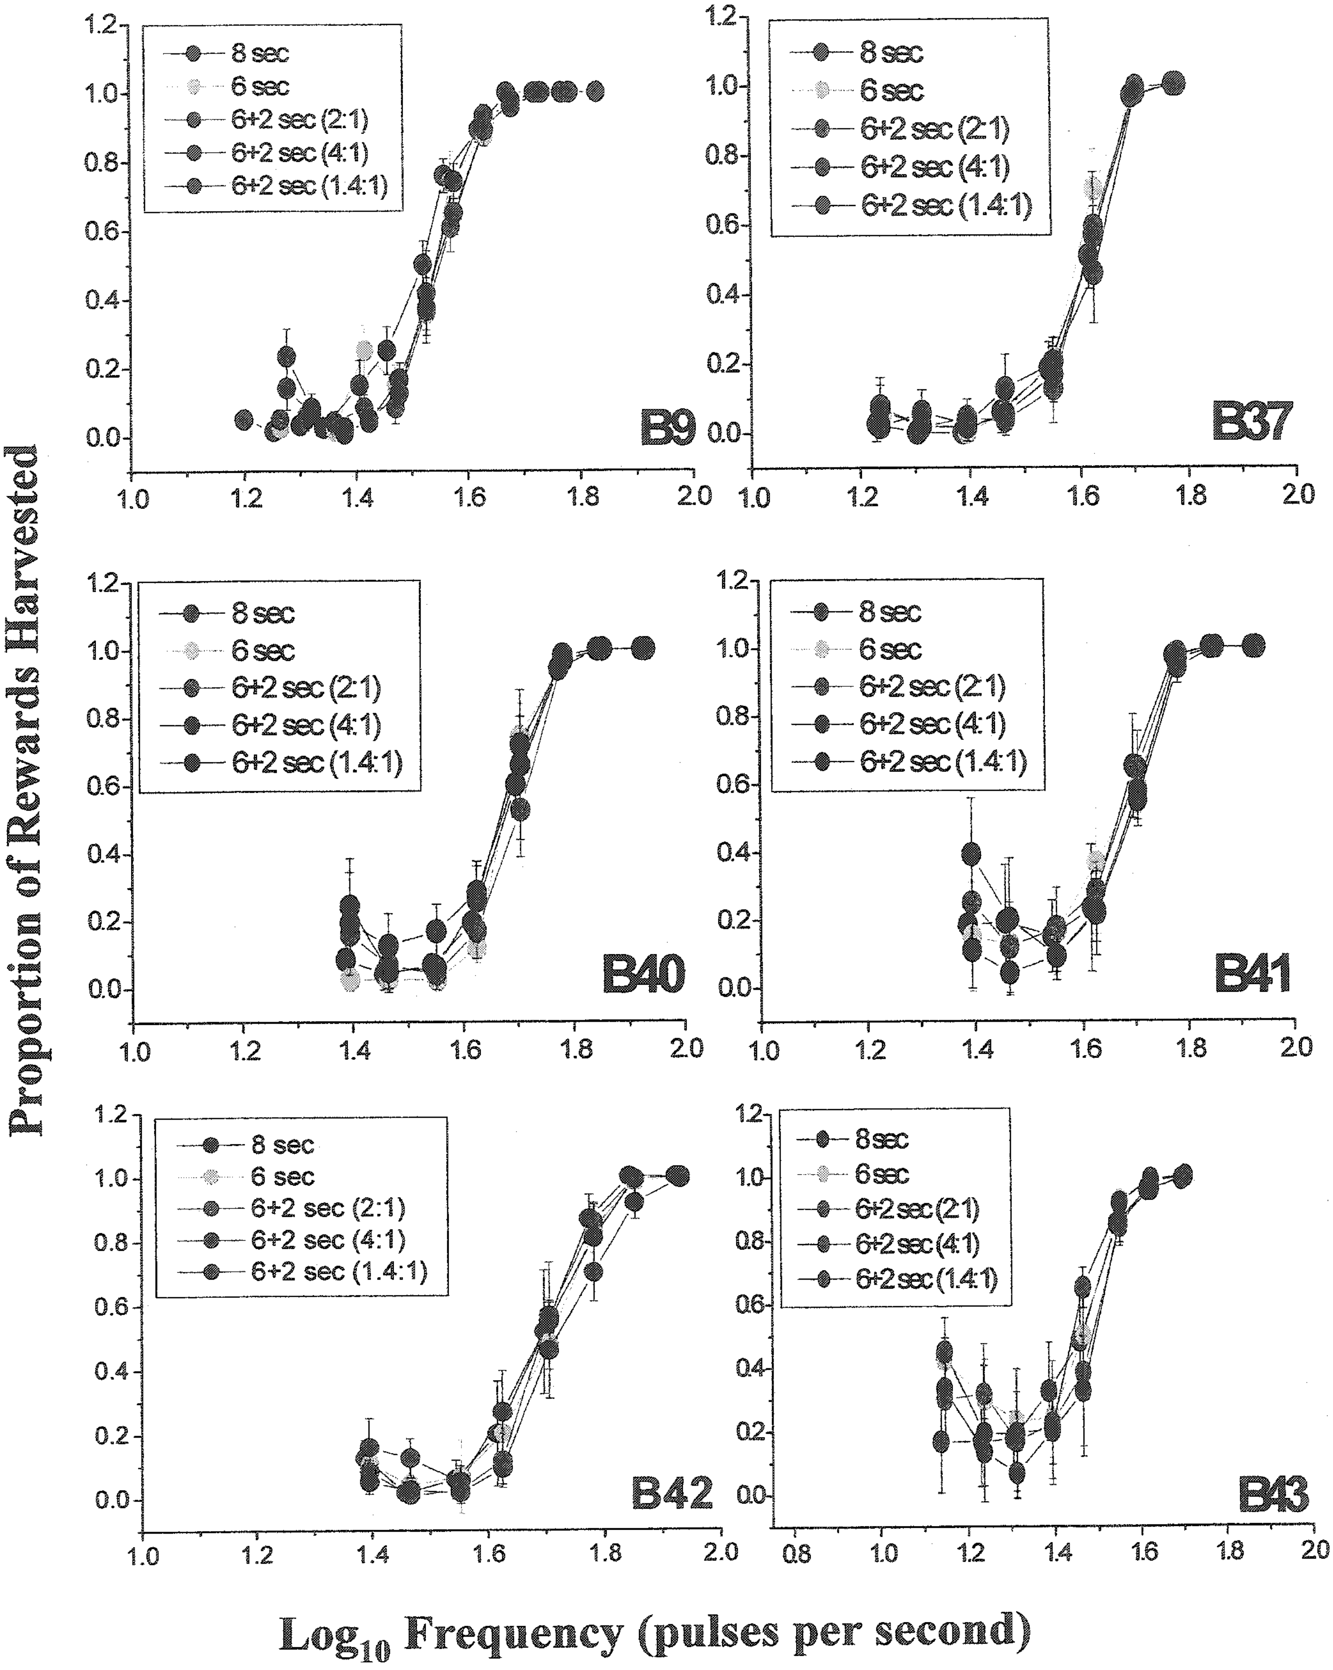

Figure 8. Raw data results from the single-operant peak-and-end test.

Figure 9. A representative broken-line fit to illustrative single-operant data. The green line represents the function fit to the data points (open green circles). The pulse frequency (in pulses per second), expressed in logarithmic units, is the x-axis variable, and the raw harvest is the y-axis variable. A broken-line fit consists of a horizontal line fit to the data points falling along the upper asymptote of the performance curve, a horizontal line fit to the data points falling along the lower asymptote of the curve, and a diagonal line segment or "riser" that connects these estimated asymptotes.

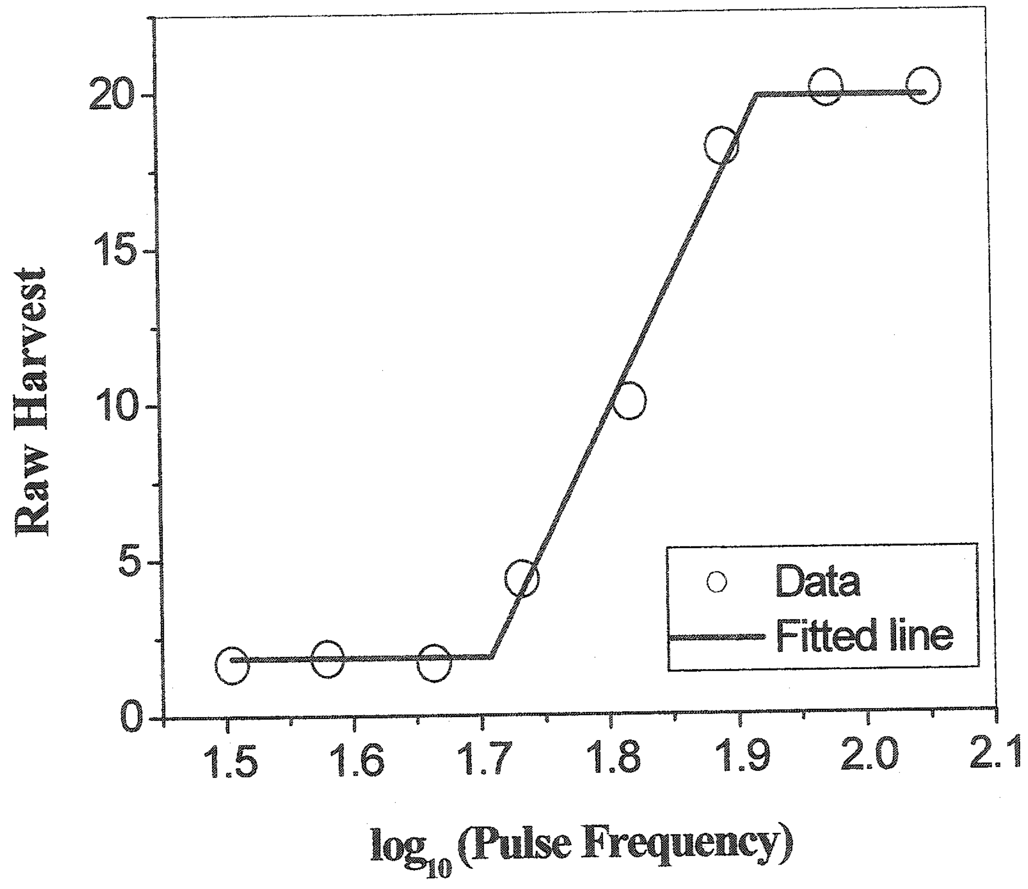

Figure 9. A sample single-operant graph, showing the broken-line fit to a performance curve.

to set the new estimated lower asymptote to zero (no rewards harvested), and to rescale the remaining data points in the psychometric functions so that their means ranged between 0 and 1.

Finally, to compensate for the heteroscedasticity (unequal variance within different segments of the performance curves) noted earlier, an attempt was made to give more weight to the relatively stable points on the upper asymptotes, and less weight to the less stable data points on the lower asymptotes and rising segments. To accomplish this, each individual observation making up each curve was weighted, and the weights used were the inverse of each observation's variance. The maximum weight was set to be 10.

The rescaled and weighted data are shown in Figure 10. Axes are identical to those shown in Figure 8. To assess the statistical reliability of shifts between curves, an Analysis of Covariance (ANCOVA) was performed, with the frequency set as the covariate, the train type as the predictor variable, and the rescaled proportion of rewards harvested as the dependent variable. The alpha level was set to 0.05 for all statistical tests. ANCOVAs were carried out on only the rising segments of each curve, as the upper and lower asymptotes as a general rule do not yield information about horizontal shifts between performance curves.

An expert statistical system was used to analyze the data (RS/1 Version 6.1, Brooks Automation, Inc., MA). In some cases, the system suggested running an ANCOVA model which fit all conditions (i.e., all train duration performance curves) with lines with a common slope, whereas for other rats the system suggested it would be best to fit separate slopes to

Figure 10. Graphs of the rescaled means from the “6+2” variant of the single-operant peak-and-end test, for 6 rats (individual subjects are identified by the label in the lower right of each graph). The x-axis variable is the frequency (in pulses per second (pps)), of the stimulation delivered during the 6 s portion of the composite train, expressed in logarithmic units. The y-axis variable is the proportion of rewards harvested, out of a maximum of 20. Error bars represent the standard error of the mean. Each curve represents a different train duration/type tested: 8 s (red) and 6 s (yellow) constant-frequency trains, and three composite “6+2” trains. For each “6+2” composite train, the first 6 s is equal to the 6 s constant-frequency train, and the last 2 s are at a lower frequency. The ratio of the frequencies of the 6 s portion to the 2 s portion of these trains are either 2:1 (green), 4:1 (magenta) or 1.4:1 (blue). Each of the five curves represents performance over 6 descending frequency “sweeps” at each train duration/type.

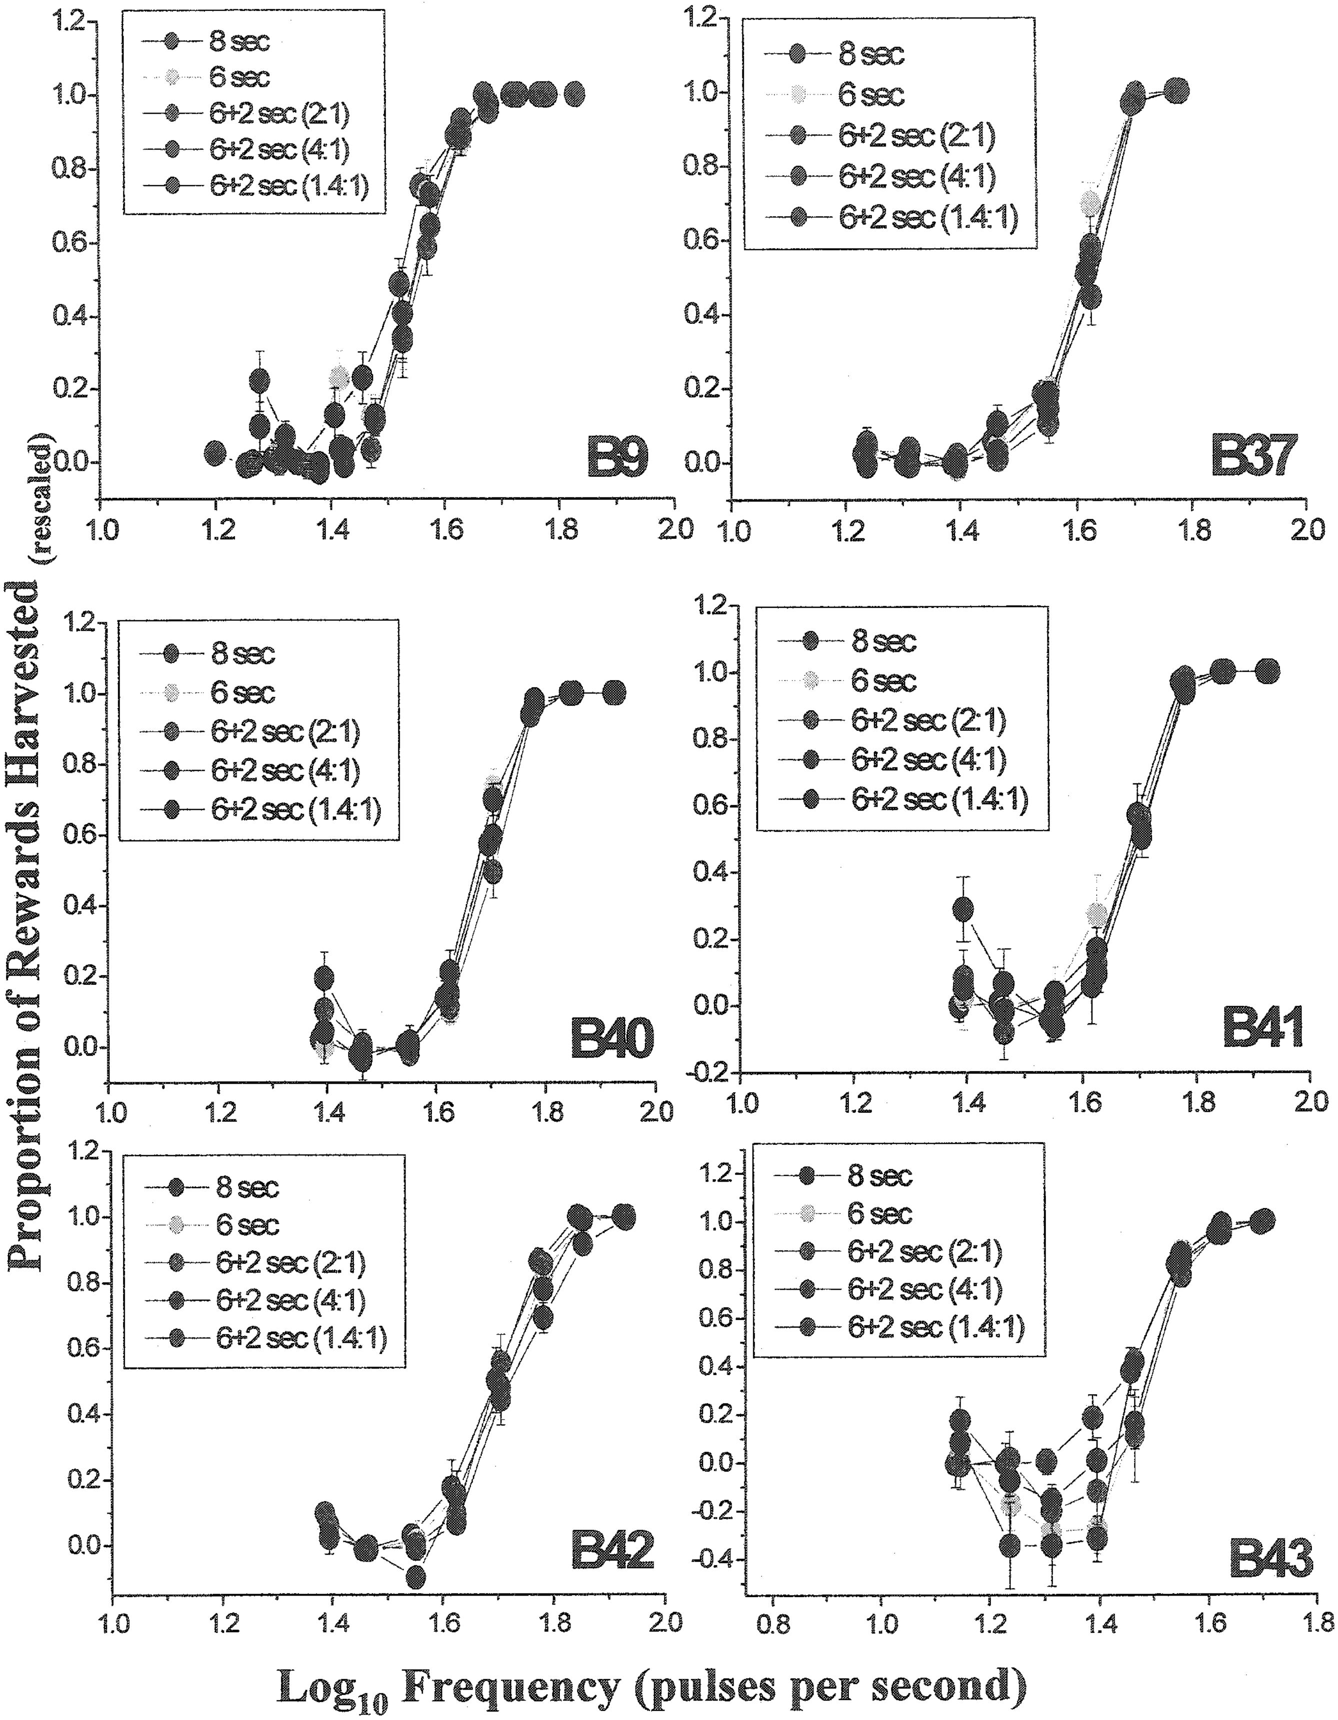

Figure 10. Rescaled means from the single-operant peak-and-end test.

each performance curve. In all, the system initially suggested the use of the Common slope option for rats B37, B40, and B42, and the Separate slope model for the remaining subjects, rats B9, B41 and B43. In the case of a separate slopes analysis, however, it is much more difficult to make pairwise comparisons between the different performance curves, as one must specify a particular value on the x-axis for each comparison of interest. Also, the use of a separate slopes analysis did not appreciably change the p-values obtained, as shown in Table 3. It was decided, to “force” the system to run a common slopes ANCOVA, even for those subjects in which a separate slope model was suggested.

Another, less widespread problem in the ANCOVA issue was differential spread of residuals in some performance curves. For one rat (B43), the system reported that some performance curves had larger spreads than others, and it recommended that weighted computations be performed such that the less reliable performance curves were down-weighted. For all other subjects this was not the case, and unweighted computations were performed for all remaining subjects.

Lastly, the system recommended that for some rats a “Robust” rather than a “Least Squares” method should be used to estimate the slopes and intercepts of the curves, as the former method would minimize the influence of outliers. This Robust method was recommended (and followed) for rats B37, B40, B41 and B42. For the remaining 2 subjects, B9 and B43, the usual Least Squares estimates were used.

Table 3.

Differences in F-values in the Single-Operant Peak-and-End Test

Rat	Common Slope Model	Separate Slope Model
B9	F(4,138) = 5.004, $p<0.001$	F(4,134) = 5.43, $p = 0.0004$
B37	F(1,84) = 1.498, $p>0.05$	F(4,80) = 1.44, $p = 0.2271$
B40	F(4,84) = 1.990, $p>0.05$	F(4,80) = 1.91, $p = 0.1169$
B41	F(4,90) = 0.606, $p>0.05$	F(4,86) = 0.962, $p = 0.4324$
B42	F(1,108) = 5.281, $p<0.001$	F(4,104) = 5.09, $p = 0.0009$
B43	F(4,90) = 6.761, $p<0.001$	F(4,86) = 7.45, $p = 0$

Based on Kahneman's (Kahneman et al., 1993) peak-and-end model, as well as the experiments of Sonnenschein, Conover, and Shizgal (2003), it was expected that the 8 s and 6 s psychometric functions would lie in close proximity to each other, as both should fall within the range of duration neglect. The peak-and-end model predicts that all or some of the "6+2" composite trains curves should fall to the right of the 8 s and 6 s curves, as the "weaker" end was expected to degrade the overall reward value of the train.

From examination of Figure 10, it becomes apparent that this is not in fact the case. For rat B37, there is no significant difference between any of the performance curves, $F(1,84) = 1.498$, $p>0.05$. A similar result holds for rat B40, with no significant differences between any of the curves, $F(4,84) = 1.990$, $p>0.05$, as well as for rat B41, $F(4,90) = 0.606$, $p>0.05$.

For the remaining subjects, there are some significant differences, but these do not meet the predictions based on the peak-and-end model. For rat B9, there is a significant difference between curves, $F(4,138) = 5.004$, $p<0.001$, with tests based on simultaneous 95% confidence intervals revealing that this significant difference lies between the 8 s (red) curve and both the 1.4:1 (blue curve) and 2:1 (green curve) composites. However, the 6 s curve does not deviate from the 6+2 composite curves. For rat B42, there is also a significant difference between performance curves, $F(1,108) = 5.281$, $p<0.001$, with the tests based on simultaneous 95% confidence intervals showing that this difference is again between the 8 s (red) curve (but not the 6 s curve) and some of the composites, the 4:1 (magenta) and 1.4:1 (blue) in this case. Lastly, for rat B43, there is a significant difference between several of the psychometric functions, $F(4,90) = 6.761$, $p<0.001$, with tests based on simultaneous 95% confidence intervals revealing that this difference is between the 8 s (red) curve and the 6 s (yellow), 2:1 (green), and 4:1 (magenta), respectively. There is also a statistically significant difference between the 6 s (yellow curve) and one of the composites, the 1.4:1 (blue) curve.

To summarize, it appears that for most subjects, there is no significant difference between responding for the 6 s regular train and responding for a "6+2" composite train. Even in the case of the one data set in which there was a statistically significant difference between the 6 s and one of the composites (rat B43, 6 s vs. 1.4:1), this result appears spurious, given that the composite in question is not significantly different from the 8 s curve. In addition, the 6 s train and the other two, "weaker" composites, are overlapping and do not significantly differ from each other.

While a number of subjects did show statistically significant differences between the 8 s curve and some of the composites, this is not a problem for the determination of "peak-alone" versus "peak-and-end", given that the key comparison is between the 6 s train and the 6+2 composites, with the 8 s train merely there as a control. In addition, even if some of the results are *statistically* significant, it does not necessarily follow that these results are *behaviourally* significant. The "statistically" significant shifts seen in the current experiment are very small compared to shifts displayed by the same rats in the frequency sweeps experiment (rat B9: Figure 4) or in pilot frequency sweeps tests (all other rats, Figures 11 and 12). For example, the distance between the curves obtained at the two shortest durations tested for each rat (B9: 0.25 s and 0.50s, all others rats 1 s and 2 s) is much greater than any "statistically" significant shift seen in the present experiment. The small size of these shifts in relation to the frequency

Figure 11. Rescaled data from the pilot (short-version) frequency-sweeps tests for 6 rats (identified by the labels in the lower right of each graph). The pulse frequency and its common logarithm are represented on the upper and lower x-axes, respectively. Plotted on the y-axis is the rescaled proportion of rewards harvested, out of a maximum of 20. The rescaling procedure sets the lower asymptote of each curve to zero and the upper asymptote to one. Each data point is the mean of six observations; error bars represent the standard error of the mean. Rats were exposed to four different train durations: 1 s (blue), 2 s (cyan), 4 s, (magenta), and 8 s (orange).

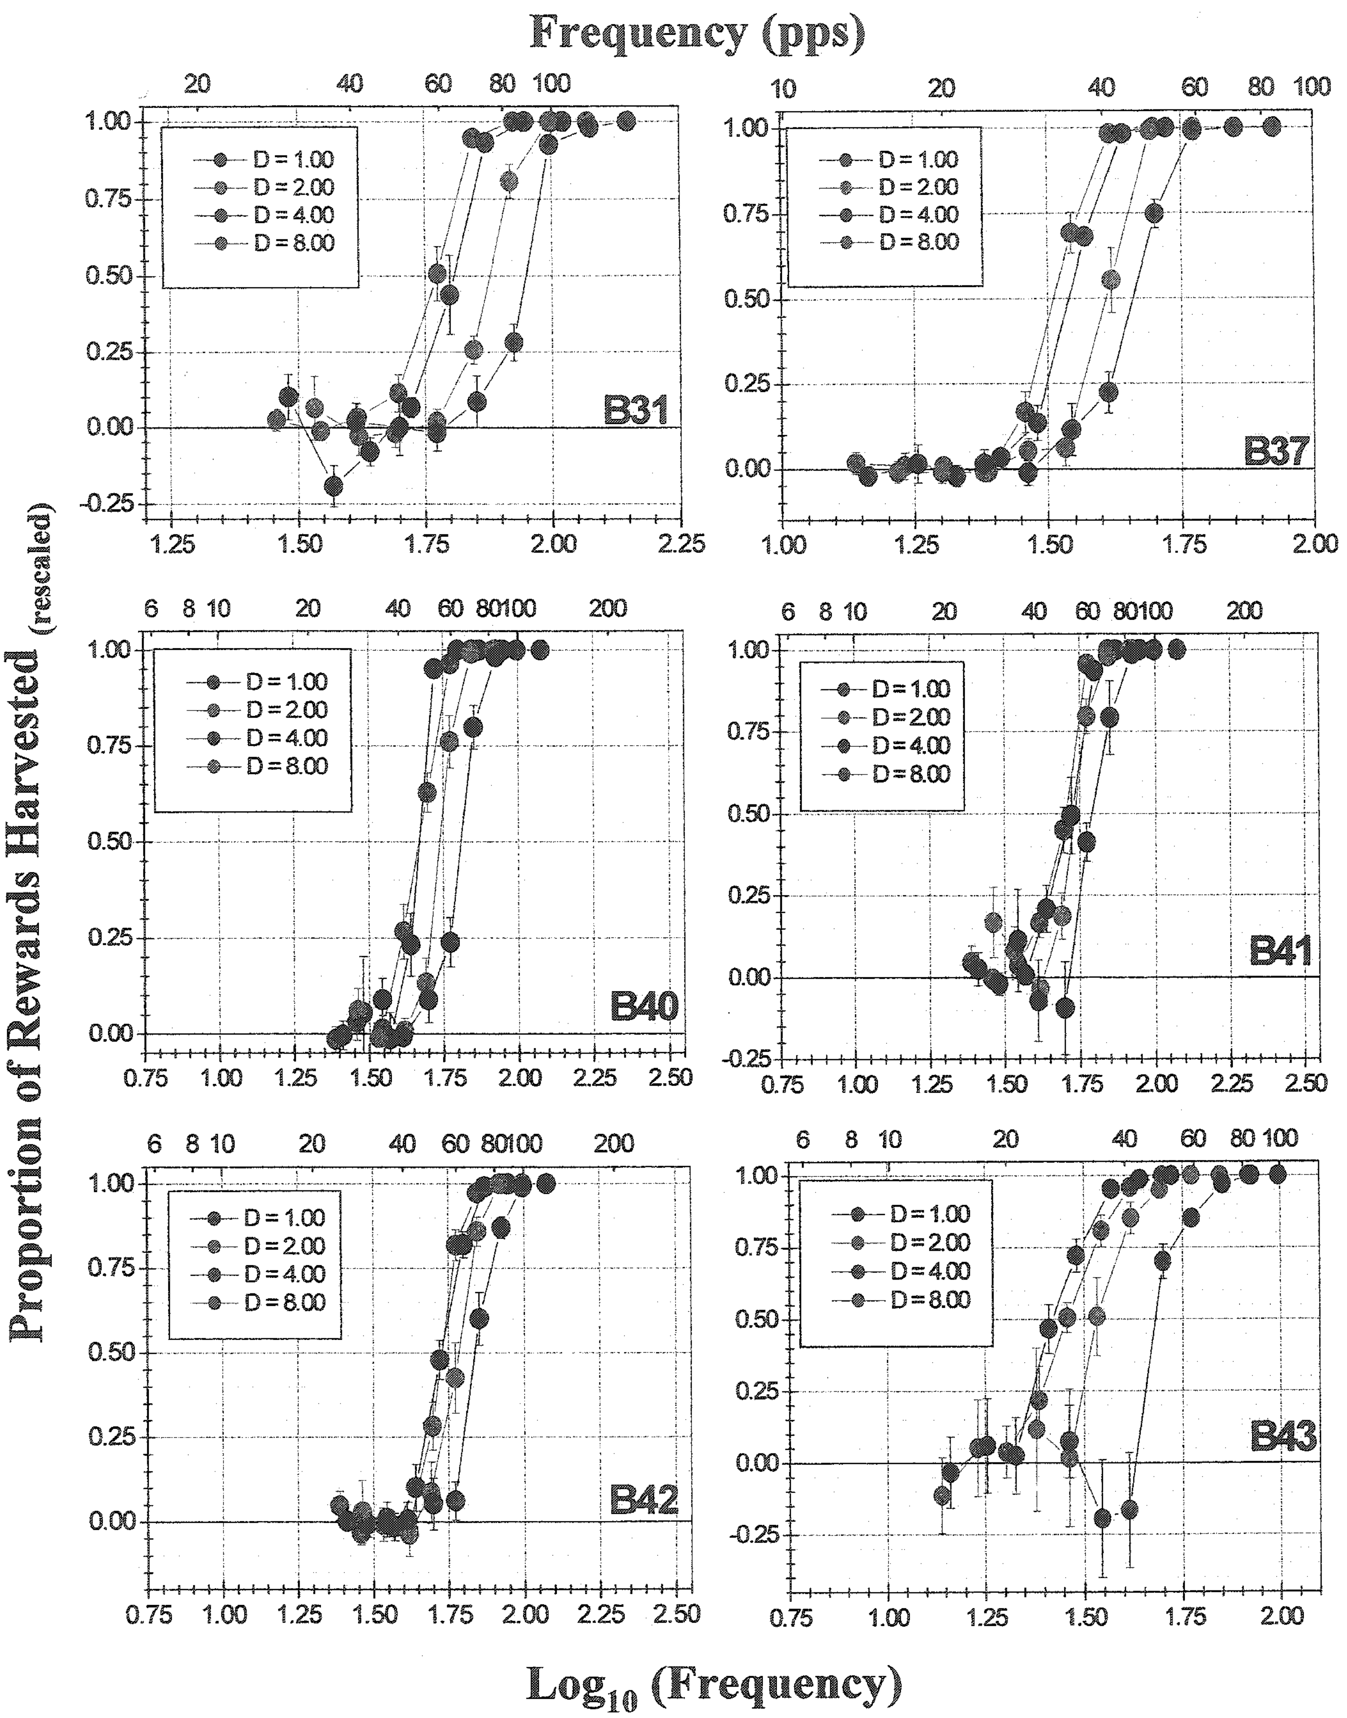

Figure 11. Additional frequency sweeps data.

Figure 12. Rescaled data from the pilot (short-version) frequency-sweeps tests for an additional 3 rats. (For details, see caption for Figure 11.)

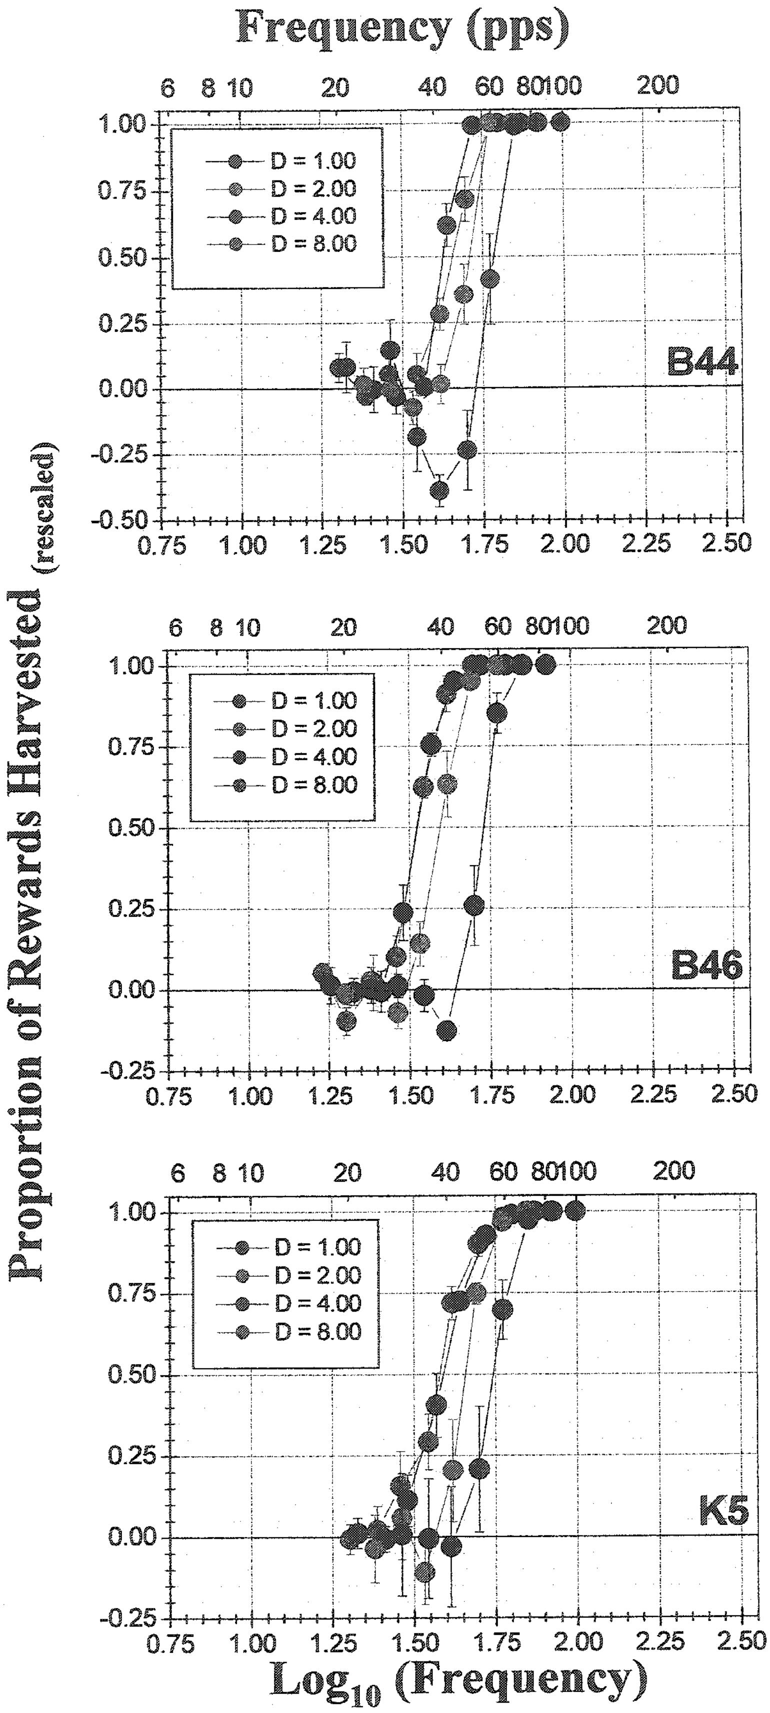

Figure 12. Additional frequency sweeps data for the remaining subjects.

sweeps thus suggests that the shifts in the present experiment, where they occur, are not truly all that behaviourally significant, even if statistically significant.

Discussion

The results of the single-operant peak-and-end experiment (Experiment 1) suggest that the peak-only model of Gallistel (1978) is sufficient. It appears that it is only the "best" or "most rewarding" moment of the BSR train that is remembered, rather than some combination of "best" moment and "ending" moment reward value. Adding 2 s of weaker stimulation to a 6 s train did not degrade the overall subjective intensity of reward for that train, as performance for these trains as opposed to performance for the shorter constant-frequency 6 s train was roughly the same. The only statistically significant difference between a 6 s train and a 6+2 s train was displayed by rat B43, and that effect was likely spurious.

The results also provide further evidence for the display of duration neglect: once duration exceeds a critical value, further increases in train duration no longer produce curve shifts. All rats except 1 (B43) displayed no significant differences between the 8 s and 6 s constant-frequency curves, and this is consistent with the duration neglect findings from Sonnenschein, Conover, and Shizgal (2003). Findings from the latter study suggested that trains exceeding 2-4 s lead to the display of duration neglect, given a sufficiently strong frequency. This finding is therefore generally consistent with the "duration neglect" displayed by the human subjects in the Kahneman studies (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996). The results are also consistent with those of Gallistel (1978) and colleagues (Mark & Gallistel, 1993; Shizgal & Matthews, 1977), and of Mason and Milner (1986a, 1986b).

The findings of Experiment 1 do not support the extension of the Kahneman group's peak-and-end model to BSR in rats. Their finding of duration neglect, however, was replicated in the present experiment. One possible reason why the weak "ends" may not have proved salient enough to devalue the overall reward value of the trains may have to do with stimulus presentation. In a single-operant paradigm, only one train is presented for several days, and thus the subject is unable to make a direct comparison between constant-frequency and composite-frequency trains. In addition, it is not clear whether the single-operant paradigm itself is analogous to how the Kahneman group experiments were carried out. Is the rat in fact performing a global rating – i.e., comparing its memory of one of the past train types to the present train – or is it merely responding to its memory of the "goodness" of the train experienced most recently? Perhaps if the constant-frequency and composite-frequency trains were presented simultaneously to the rat, the weaker "end" would prove more salient in such a context. Experiment 2 was aimed at determining whether this would in fact be the case.

Experiment 2

The Kahneman (Kahneman et al., 1993) peak-and-end model does not appear to account for performance for BSR trains in a single-operant paradigm. If two trains are presented simultaneously, however, the rat is then forced to make a decision in real-time as to which stimulation train is most rewarding. Perhaps in such an evaluation context, a train with a weaker "end" would lead to a lower associated reward value, than in a situation where only one train is evaluated at a time, as in Experiment 1.

Hsee and colleagues (Hsee, Blount, Loewenstein, & Bazerman, 1999) have argued that when options are presented one at a time (single evaluation; SE), their characteristics may be weighted much differently than when multiple options are presented (joint evaluation; JE). Given that the single-operant paradigm can be thought of as analogous to the SE context, and the dual-operant paradigm to the JE context, perhaps in the SE context the weak end was less relevant to the rats when the train "options" were only presented singly. If two or more trains are presented at the same time, perhaps information about the "ends" of the stimulation would become more relevant.

Thus, the aims of Experiment 2 were: 1) to determine if the peak model of Gallistel (1978) would again be sufficient to account for the results in the new evaluation context, and 2) to determine if "duration neglect" between 6 s and 8 s trains would be replicated in a dual-lever paradigm. If Gallistel's (1978) peak-only model held, rats would display equipreference between a constant-frequency and a composite-frequency train presented simultaneously.

However, if the Kahneman (Kahneman et al., 1993) peak-and-end model held, adding a weaker "end" would decrease the value of the composite train, and thus the rats would show a preference for the shorter constant-frequency train.

Method

Subjects

The 6 subjects in the dual-operant peak-and-end experiment were the same as those utilized in Experiment 1, with the exception of rats B28 and B31. All rats were housed and food restricted in the same manner as described in Experiment 1.

Surgery

"New" subjects had been surgically implanted with stimulating electrodes in the same manner as those subjects described earlier in Experiment 1.

Apparatus

All operant boxes employed were identical to those described in Experiment 1. The only exception was that the houselight no longer flashed during the inter-trial-intervals. Instead, the houselight flashed during the "switchover", described below.

Procedure

Screening and training. All "new" animals had been screened for aversive effects and disruptive stimulation-induced movements, as well as trained to self-stimulate reliably. For one "new" animal, B28, 400 μA proved too aversive, but an intensity of 200 μA was effective, and thus this current was used in all experiments with this subject.

Baseline condition. This condition continued for all rats, including the two added subjects, and was identical to the daily baseline test described previously.

Dual-operant peak-and-end experiment. All testing took place in the dual-lever operant boxes described earlier. Each trial began with a 20 s ITI. In this case, however, during ITIs both levers were unavailable, the houselight was left extinguished, and no priming train was delivered. Trials lasted 320 s, and there were 11 trials per sweep for rat B9 (0.05 log unit steps) and 9 trials per sweep for all remaining rats (0.075 log unit steps). In total, there were 33 trials per session (3 sweeps) for rat B9, and 27 trials per session (3 sweeps) for all remaining rats.

Both levers were available to the rats during the trials. Similar to Experiment 1, a 16 s "blackout delay" (BD) was employed, however in this case once either lever was pressed, both levers retracted into the chamber wall, and both extended into the chamber again only after termination of the BD. The keylights over both levers were also extinguished during the BD. The use of a 16 s BD served the same purpose as described in Experiment 1.

The trains available on each lever were set as follows. One lever, termed the "standard" lever (Std) delivered a reward which was always set to the same moderately-rewarding frequency. The frequency used for each rat was determined by their performance in previous experiments. This value was usually selected to be a "shoulder value"; if a rat's performance for a descending frequency sweep at a particular train duration is plotted, this value is the portion of the curve at which the rat is obtaining most of the available rewards, but has not reached a behavioural ceiling (i.e., "harvesting" all available rewards). A representative psychometric function and its shoulder are shown in Figure 13. The frequency of the Std train was held to this shoulder value throughout all trials, and did not vary from trial-to-trial or from sweep to sweep within an experimental condition.

On the remaining lever, termed the "alternate" lever (Alt), the set of frequencies used was initially determined from the descending frequency sweeps used during the peak-and-end single-operant test, such that the steps were either 0.05 (B9) or 0.075 (all other subjects) $\log_{10}$ units apart. This set of frequencies was then varied in the following manner: the set of frequencies was split into two halves (a "high frequency" set and a "low frequency" set), and the order of each set was mixed randomly. The "high" and "low" sets were then combined into one series for the "sweep", such that on every second trial, the frequency of the Alt train would be higher than that of the Std train, while on the remaining trials, the opposite was the case. The only exception was the warm-up trial, which was identical to Trial 1, and both were always selected from the "high" frequency set.

A "switchover" was also included in each trial, to prevent the development of a "side bias". If one lever had been always set to trigger the Alt train and the other the Std train, the rats may have developed a preference for one side of the operant box, regardless of the rewards available on each lever. The switchover was used to guard against this, and was set to

Figure 13. Sample data indicating the "shoulder" region of a single-operant performance curve. Illustrative data are represented by the connected open green circles. The common logarithm of the frequency (pulses per second) is the x-axis variable, and the Raw Harvest (out of a maximum of 20) is plotted on the y-axis. The "shoulder region" of the curve, the portion of the psychometric function where the rat is obtaining most of the available rewards but has not reached behavioural ceiling, is indicated by the section of the curve bounded by the dotted line. Values from this region were used to select the initial pulse frequencies for the Std trains in the present dual-operant studies.

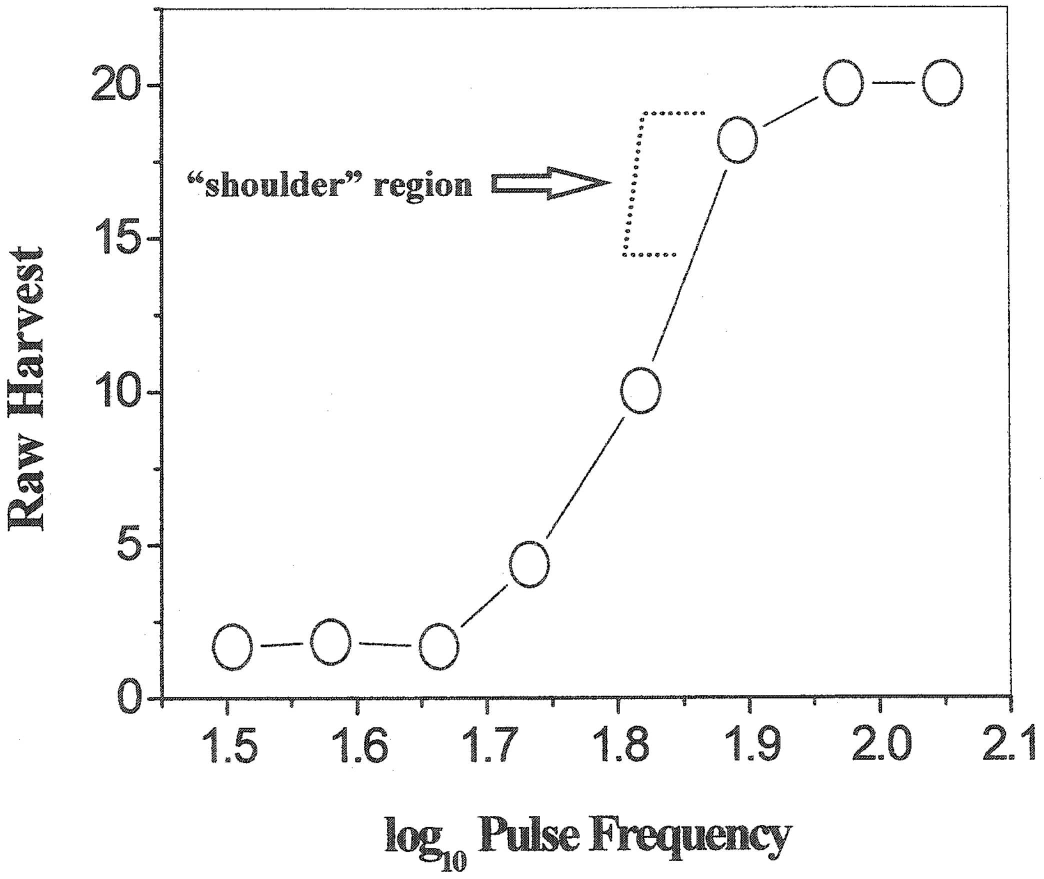

Figure 13. A sample curve showing the "shoulder" region of a single-operant performance curve.

occur at the half-way point in each trial (i.e., 160 s into the trial); the houselight flashed for 5 s and the mapping between levers and trains was reversed, such that the lever which originally delivered the Alt train was now delivering the Std train, and vice-versa.

The first condition of the present experiment involved a check to ensure that the rat was discriminating appropriately between the Alt and Std trains. This involved setting the duration of both the Alt and Std trains to 6 s (with the exception of rat B9, for whom both trains were set to 4 s). All other stimulation parameters, such as stimulation current, and frequency, were set as described previously: the frequency of the Std train was always set to the same moderately-rewarding value, and the Alt frequency varied from trial-to-trial, such that on some trials, the Alt train was more rewarding than the Std lever reward, and on other trials, less rewarding.

The second condition was a test for duration neglect. The frequency sweeps experiment (Sonnenschein, Conover, & Shizgal, 2003) had demonstrated that in a single-operant condition, most subjects demonstrated duration neglect once the stimulation trains exceeded 2-4 seconds, but a similar phenomenon had yet to be tested in a dual-lever situation. While Mark and Gallistel (1993) carried out dual-operant experiments in which subjects were required to choose between stimulation trains that differed in duration, these experiments involved variable-interval (VI) rates of reinforcement rather than continuous reinforcement. The use of the latter reinforcement schedule would be more comparable to the reinforcement schedule used in Sonnenschein, Conover, and Shizgal (2003) and Experiment 1. In the second condition of the present experiment, the duration of the Std lever reward was set to 8 s, and the duration of the Alt lever trains were set to 6 s, with all other stimulation parameters set as described earlier.

In the third and final condition, the Alt lever triggered a 6 s constant-frequency train, and the Std lever triggered a composite train. Specifically, the Std was a "6+2" train, in which the first 6 s of the train was set to a constant frequency, and the last 2 s of the train were set to a lower frequency. In Experiment 1, three different composite trains were tested, in which the ratios of the frequencies of the first 6 s to the last 2 s of the trains were either 4:1, 2:1, or 1.4:1, but as no consistent differences in performance were displayed between the three "degrees" of weak endings used in Experiment 1, only the 2:1 ratio difference between the "6" and "2" segments was employed in the present study.

Results

The first trial in each sweep was again considered a "warm-up" trial, and responding for both the standard and alternate levers during this trial was not used in subsequent statistical analyses. All analyses were conducted using data averaged over 12 sweeps.

In the present experiment, the measure of interest was where the subject would "cross over" from responding preferentially on one lever, to responding on the other lever. This point of cross-over is usually referred to as the point of "equipreference", or the "cross-point" (CP), and it is expressed as the stimulation frequency at which the reward value of both trains is inferred to be equal. The CP could only be inferred because the rats were never presented with trials in which the reward value on both trains was equal.

The calculation of the CP was carried out in the following manner. First, a "broken-line function" was fit to each individual Alt and Std response curve, in a manner analogous to the procedure employed in Experiment 1. Briefly, horizontal lines were fit to the upper asymptote, the rising segment, and the lower asymptote, for both the Std and Alt performance curves. A sample broken-line fit to data from a set of dual-operant sessions is shown in Figure 14. For the purposes of the current experiment, for each condition (for each rat) 12 individual Alt lever response curves and 12 individual Std lever response curves were fit with this function, yielding 24 separate broken-line fits for each graph.

Figure 14. An illustrative broken-line fit to a dual-operant data set. The open red circles represent rewards obtained by pressing the Alt lever, and the open green circles represent rewards obtained by pressing the Std lever; the green and red lines are the respective broken-line fits. The common logarithm of the frequency (in pulses per second) is the x-axis variable, and the proportion of rewards harvested (out of 20) is plotted on the y-axis. The cross-point (CP), represented as a blue open circle, is the estimated point at which the subject "crosses-over" from responding preferentially on one lever to responding preferentially on the other. The confidence band around the CP is the 95% confidence interval.

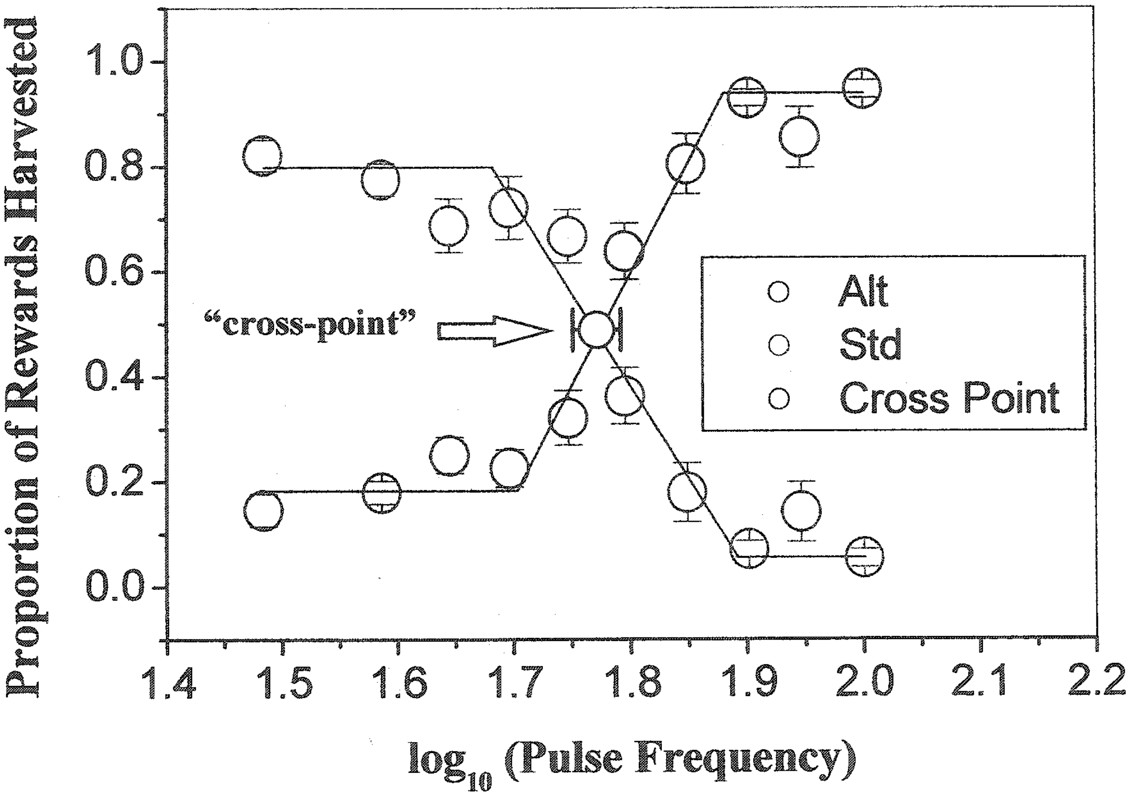

Figure 14. A sample dual operant graph showing the broken-line fits and the estimated cross point (CP).

The frequency of the Alt train at which the rising portions (slopes) of each broken-line fit intersect is the cross-point (CP; see Figure 14). In practice, not every sweep yielded a good CP; sometimes the intersection of the Std and Alt curves did not lie on the diagonal rising segments. Only individual sweeps where a well-behaved CP was obtained (i.e., the two rising segments intersect) were retained for statistical and graphing purposes. As all analyses were carried out on data that had been averaged over 12 sweeps, rats had to obtain a minimum of 12 sweeps with good CPs before the next condition could be run.

Once the broken-line fits on the "well-behaved" CPs were performed, the x- and y-values of each of the 12 CPs were taken from each sweep, and the set of 12 x-values was averaged together, with a similar procedure carried out separately for the set of 12 y-values. These two values, the averaged x-value and the averaged y-value, were then used to graph a mean CP for the full set of 12 sweeps. Lastly, a 95% confidence band was calculated for the mean CP of the 12 sweeps, positioned around the x-values (horizontal band).

The CP is only meaningful in relation to the frequency of the standard lever. A sample dual-operant graph is plotted in Figure 15 for the purpose of explanation. In this graph, the x-axis variable is the logarithm of the frequency on the Alt lever, and the y-axis variable is the proportion of rewards harvested (PRH). Responses on the Std lever (a 6 s train, in this case) are represented by the green data points, whereas responses on the Alt lever (another 6 s train) are represented by the red points. The blue point on the graph is the CP that was calculated from the individual CPs of each of the 12 sweeps. Lastly, a vertical magenta line is plotted on the graph, which indicates the frequency of the stimulation on the Std lever.

Figure 15. A sample dual-operant graph. Lines represent broken-line fits to illustrative data points, with rewards obtained by pressing the Alt lever represented by open red circles, and rewards obtained by pressing the Std lever represented by open green circles. The green and red lines are the respective broken-line fits. The common logarithm of the frequency (in pulses per second; pps) is plotted on the x-axis, and the proportion of rewards harvested (out of 20) is plotted on the y-axis. The cross-point (CP) is represented as an open blue circle, with a 95% confidence band. The frequency of the Std train is represented as a vertical magenta line. Both the Std and the Alt lever trains are 6 s in duration.

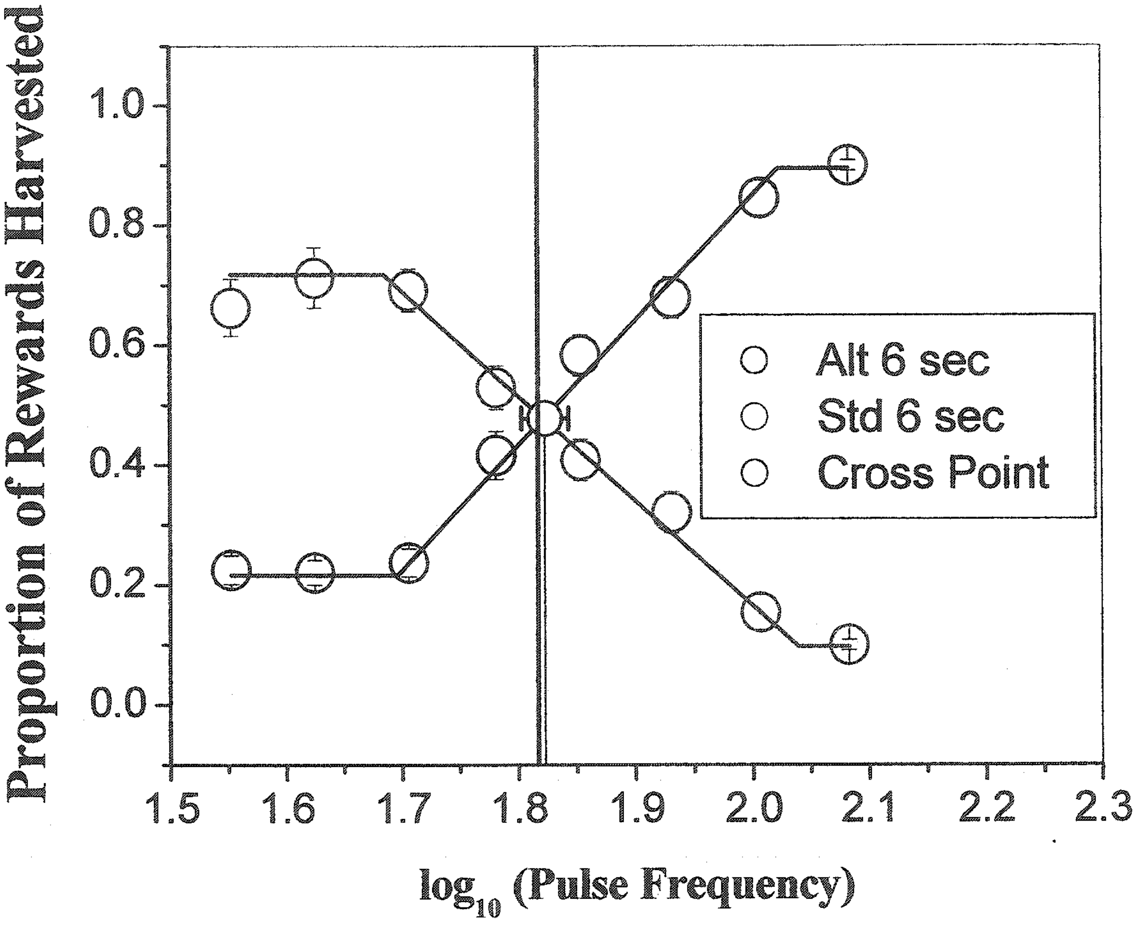

Figure 15. Sample dual-operant graph, showing the comparison between the estimated cross-point (CP) and the frequency of the standard lever (vertical magenta line).

To interpret any dual-operant graph, one must keep in mind that on the area of the graph to the *left* of the magenta line (the frequency of the Std train), the frequency of the Alt lever is less than that of the Std lever, and thus the rats *should* respond more for the Std train than the Alt train. On the other half of the graph, to the *right* of the magenta line, the frequency of the available Alt trains is higher than that of the Std train, and thus the rats *should* respond more for the Alt train than the Std train.

As shown in Figure 16, the rats behaved as expected when the Alt and Std trains were identical. The frequency of the Alt lever train is graphed in logarithmic units on the x-axis, and the magenta line represents the frequency of the Std lever train, also in logarithmic units. The y-axis variable is the proportion of rewards harvested (PRH). As can be seen from the figure, for all subjects except B42 (bottom right), whenever the frequency of the Std train exceeds the frequency of the Alt train (left half of the graph), the rats do respond more for the Std lever reward (green curve falls above the red curve), and the opposite holds when the Alt lever train is more rewarding (right half of the graph). In addition, if the durations of the two rewards are the same, as they are in the present condition, the CP and magenta line should overlap. In other words, the rats should "cross-over" from pressing preferentially on one lever to the other, at the frequency of the Std lever train. For all but one of the subjects, this did in fact occur. In the case of B42, there is a statistically significant but small difference between the CP (and its confidence band) and the frequency of the standard, with the size of the behavioural shift equaling 0.04 log units. Thus, with the exception of the latter rat, all the subjects discriminated between the two rewards as expected.

Figure 16. Graphs illustrating the reliability of performance in the dual operant peak-and-end test, for 6 rats (individual subjects are identified by the label in the lower right of each graph). The x-axis value is the common logarithm of the pulse frequency of the Alt lever train, in pulses per second, with the vertical magenta line representing the common logarithm of the frequency of the Std train, also in pulses per second. The y-axis value is the proportion of rewards harvested, out of a maximum value of 20. The red circles represent rewards obtained by pressing the Alt lever, and green circles rewards obtained by pressing the Std lever, with error bars around these data points representing the standard error of the mean. The green and red lines are the respective broken-line fits. Both data curves represent performance over 12 sweeps. The blue circle represents the estimated "cross-over" between the Std lever and the Alt lever, and its error bar is the 95% confidence interval. Both the Std and Alt lever trains were 6 s in duration (except for rat B9, who had both trains set to 4 s).

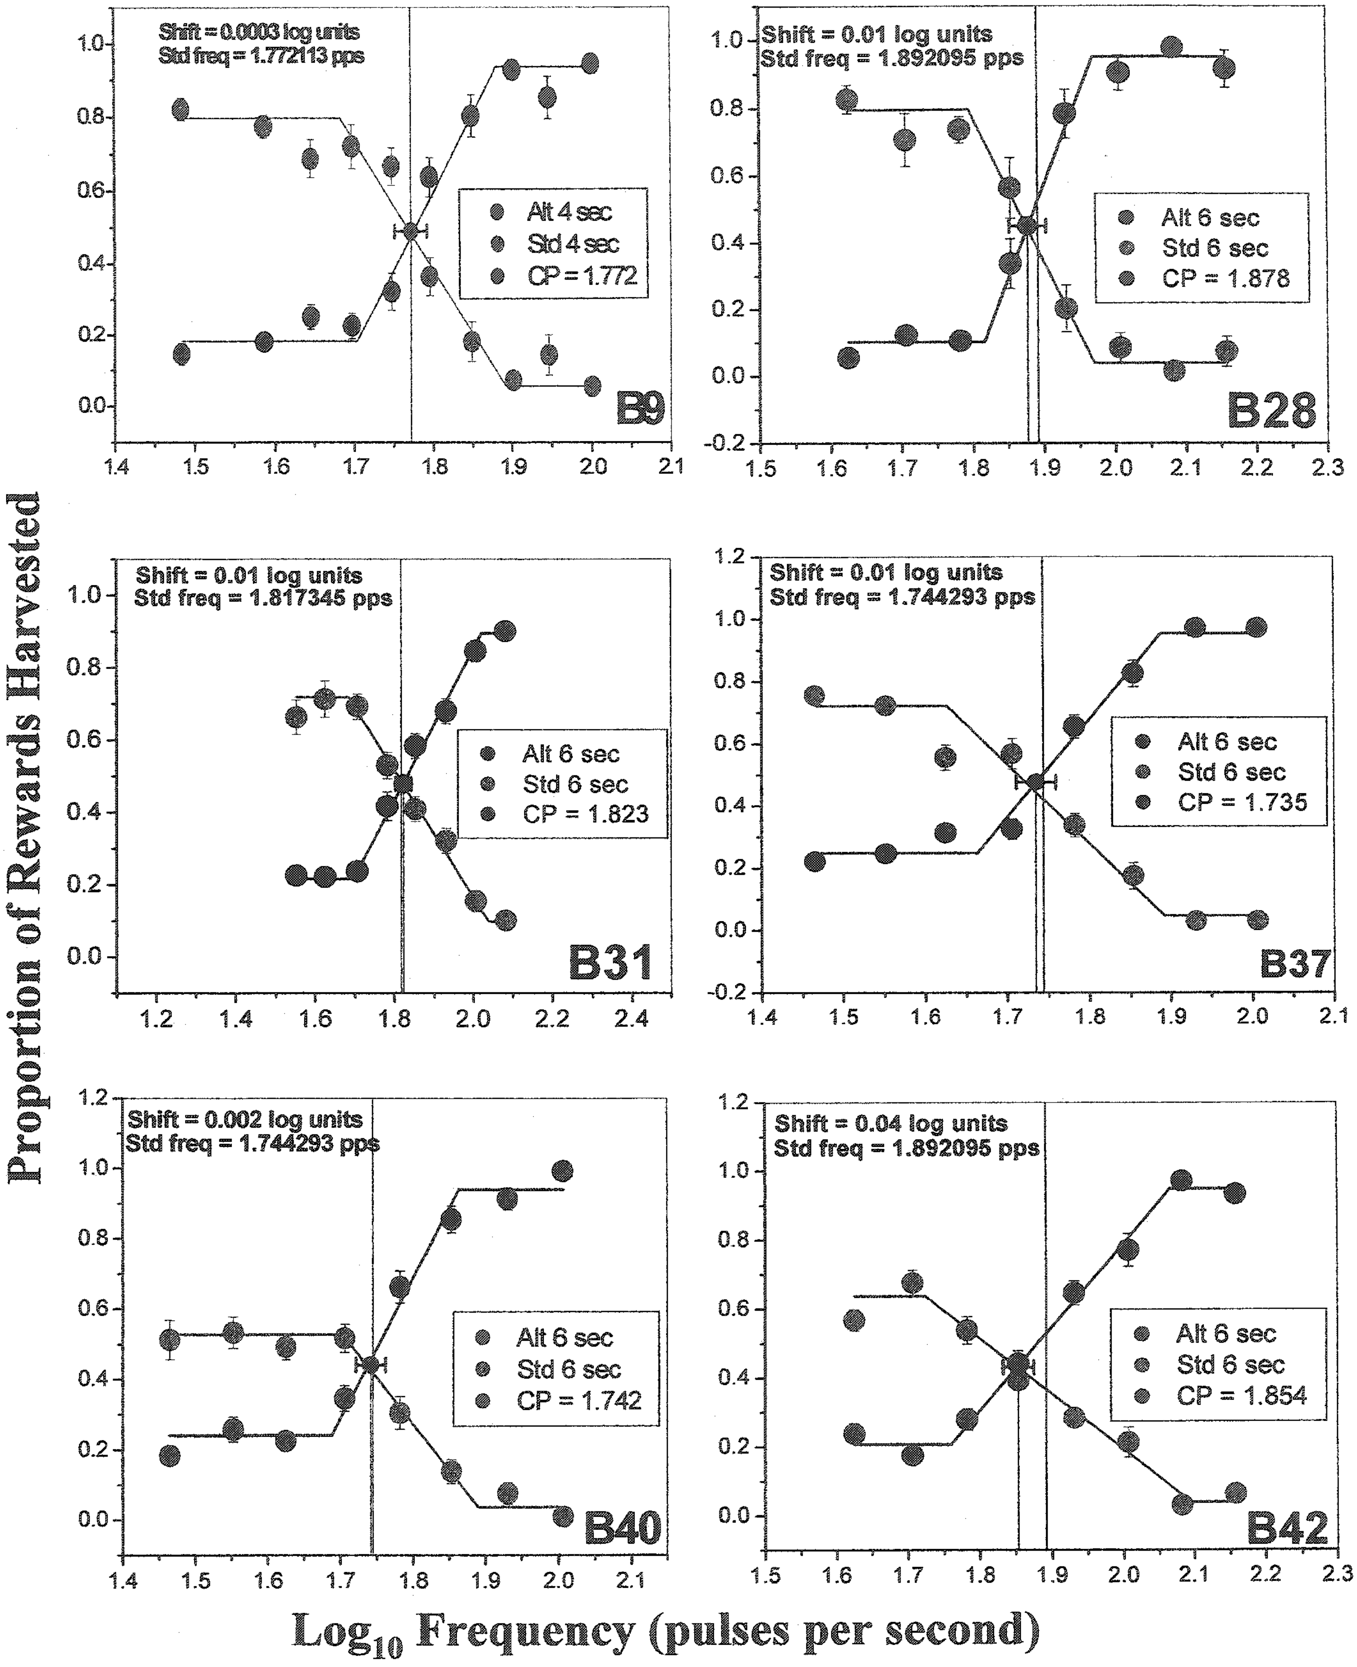

Figure 16. Results from the first condition of the dual-operant peak-and-end experiment.

The second condition was a test for duration neglect (B9 did not run this condition, as he was a pilot rat run many months before the remaining subjects, and at that time, a test for duration neglect was not carried out). Here, all other parameters were identical to the first condition above, except the Std train was now 8 s in duration, and the Alt lever delivered a 6 s train. The results of this condition are shown in Figure 17. All axes are identical to the previous Figure. For all subjects except B31, there is no statistically significant difference between the CP (blue point) and the frequency of the Std train (magenta line). It is also notable that although B42 proved unreliable in the 6 s vs. 6 s condition, he performed as expected in the current condition. Another rat, B31, did show a *statistically* significant preference for the Alt lever reward, as indicated by the fact that the CP and its 95% confidence band lie to the left of the magenta line. However, this is not a highly significant shift *behaviourally*, as the difference is a small 0.02 log unit difference (a mere 4.7% increase in the frequency of the Alt train was needed to motivate this rat to "cross over" to that lever). Thus, once again the majority of the subjects displayed equipreference between the two rewards.

The final condition of the present experiment was the peak-and-end test. The Std lever reward was a "6+2" train, in which the first 6 s were at a moderately-rewarding frequency, and the last 2 s of the train were set at a frequency half of that of the 6 s part. The Alt train was set to be 6 s in duration. The results of this condition are shown in Figure 18, where the axes have the same meaning as prior graphs. It is evident that there are no statistically (or behaviourally) significant shifts between the CP and the magenta line for any subject, implying that both rewards were equivalent to the rats.

Figure 17. Graphs of the means from the dual-operant duration-neglect test, for 5 rats (individual subjects are identified by the label in the lower right of each graph). The x-axis value is the frequency, expressed in logarithmic units, of the Alt train (in pulses per second). The vertical magenta line represents the common logarithm of the Std-train frequency. The dependent variable is the proportion of rewards harvested. The red data points represent rewards obtained by pressing the Alt lever, and green data points rewards obtained by pressing the Std lever, both with error bars representing the standard error of the mean. The green and red lines are the respective broken-line fits. Both curves represent performance over 12 individual sweeps. The blue point represents the estimated "cross-over" between the Std and the Alt, with the 95% confidence band plotted around it. The Std train was 8 s in duration, and the Alt train was 6 s in duration.

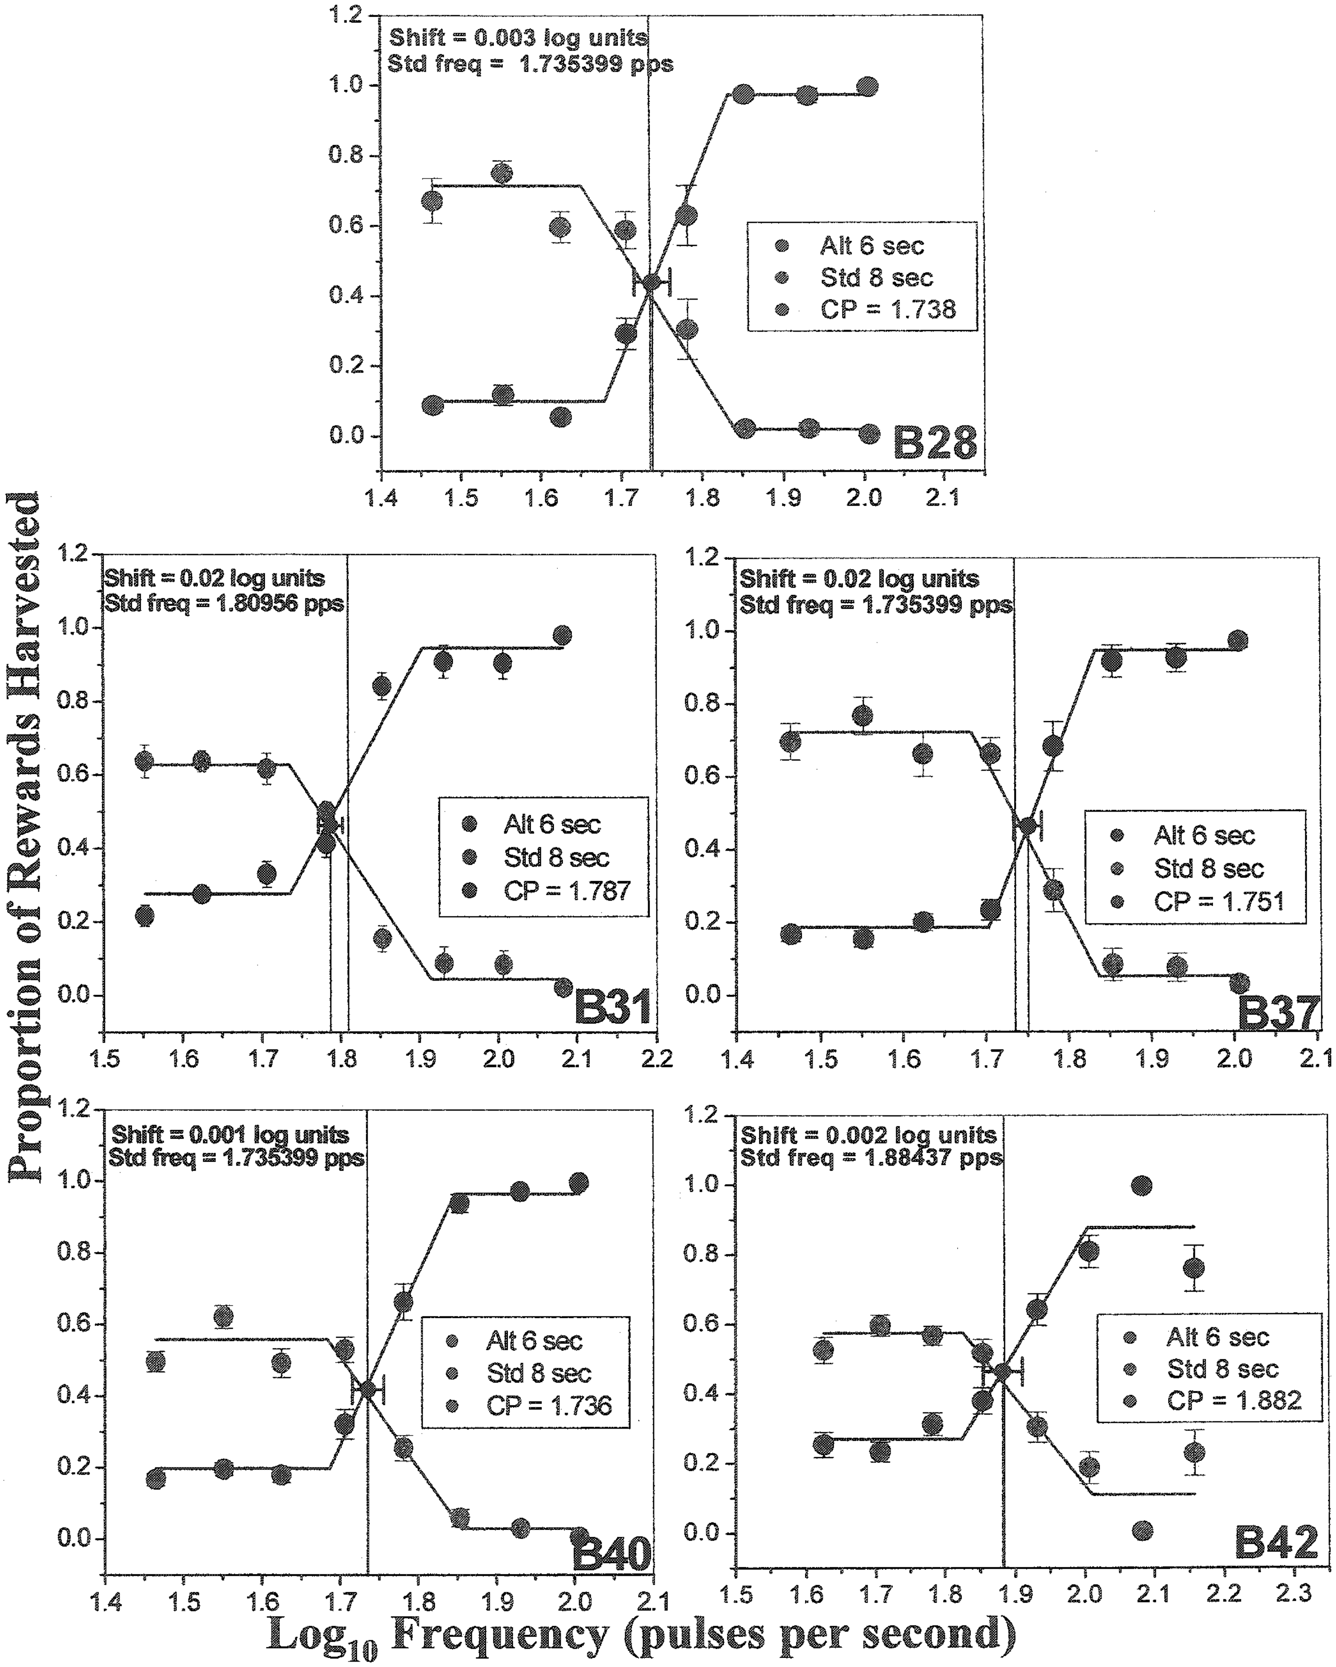

Figure 17. Results from the duration neglect condition of the dual-operant peak-and-end experiment.

Figure 18. Graphs of the means from the dual-operant peak and end test, for 6 subjects. The x-axis variable is the common logarithm of the frequency of the Alt train, in pulses per second. The vertical magenta line denotes the common logarithm of the frequency of the Std train (the frequency of the 6 s segment of the train), also in pulses per second. The y-axis variable is the proportion of rewards harvested. The red points represent rewards obtained by pressing the Alt lever, and green points rewards obtained by pressing the Std lever, with error bars representing the standard error of the mean. The green and red lines are the respective broken-line fits. Data from both curves are averaged from 12 individual sweeps. The blue point represents the estimated "cross-over" between the Std lever and the Alt lever, with a 95% confidence band. The Std was a composite frequency "6+2" train, in which the first 6 s were set at a high frequency, and the last 2 s at a lower pulse frequency. The Alt train was 6 s in duration.

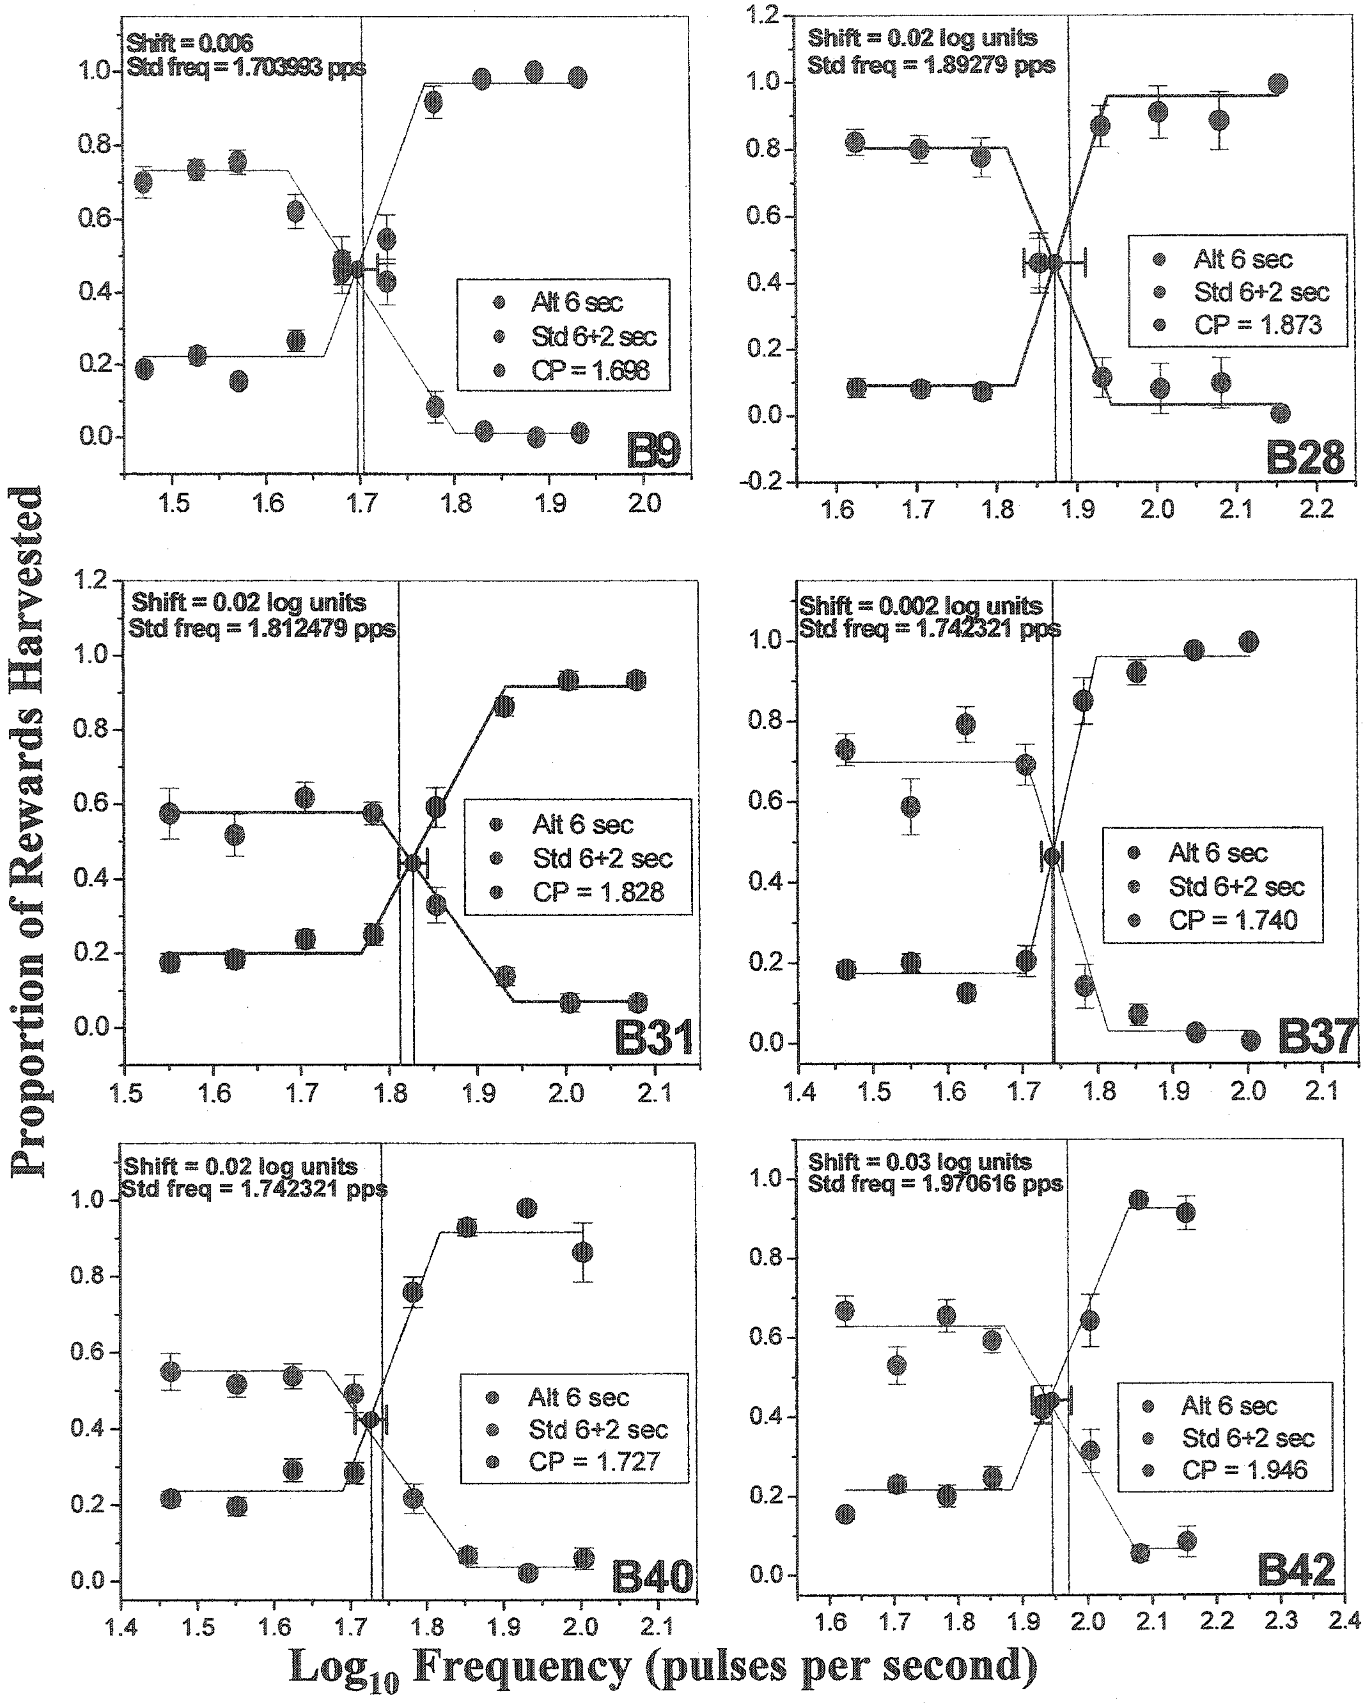

Figure 18. Results of the dual-operant peak-and-end test.

To summarize, the data suggest that adding a "weaker" end did not degrade the overall reward value of a 6 s train, and duration neglect between a 6 s and an 8 s train was replicated.

Discussion

The dual-operant peak-and-end experiment (Experiment 2) yielded results similar to those of the single-operant peak-and-end test (Experiment 1). Again, duration neglect was demonstrated, as the point of equipreference (CP) was indistinguishable from the frequency of the standard lever when the comparison was between an 8 s Std train and a 6 s Alt train. When the 6+2 s composite Std train and a constant-frequency 6 s Alt train were presented simultaneously, the point of equipreference was again indistinguishable from the Std train frequency. Both the constant- and composite-frequency trains were again seen as essentially equivalent by the subjects, and thus adding a weaker end appeared to have no measurable effect on the overall reward value. These results are in line with the predictions of Gallistel's (1978) peak-alone model, rather than the peak-and-end model (Kahneman et al., 1993).

Kahneman's (Kahneman et al., 1993) model was therefore not supported in either the single- or dual-operant paradigm. This lack of an effect may have to do with the stimulation trains employed. Although it was assumed that a brief 2 s "weak" train would be sufficiently salient to degrade the overall reward value of the entire composite train, as suggested by the work of Fouriezos (1995), perhaps this was not the case.

It is assumed that once a stimulation train terminates, the reward effect begins to decay, eventually returning to a minimal level. In the case of the 6+2 composite trains, it is assumed that once the strong 6 s portion of the train ends, reward value declines, toward the level that would have been asymptotic had the entire train been delivered at the lower frequency. The work of Fouriezos (1995) suggests that on average a train longer than 1.32 s should be sufficient for the effects of the first train portion to decay to 5% of the maximum achieved during the strong portion of the train (the duration 1.32 s is 3 times the average time constant estimated by Fouriezos). However, there was considerable variance in the time-constant estimates reported by Fouriezos. Thus, it is possible that the amount of decay during the 2 s weak component was insufficient. Experiment 3 was thus designed to address this issue.

Experiment 3

One possible reason why the "peak" appeared to be the only exemplar taken into account in previous experiments, may be because the weak 2 s train segment may not have been sufficiently long.

In an attempt to rule out this possibility, Experiment 3 made use of a novel composite train, in which the durations of the "strong" and "weak" portions were reversed: composite stimulation now consisted of "2+6" trains in which the first 2 s were fixed at a particular frequency, and were immediately followed by a 6 s train at a lower frequency.

Although a "6+6" train could have been employed, in which both train portions would have been identical in duration, this would have increased the total train duration to 12 s, and given that the maximum duration tested to date had only been 8 s, it was decided that constructing a 2+6 (8 s total) train was the best option. Additionally, had a 6+6 train been used, a longer blackout delay would have been needed to offset the effects of the long train.

Thus, the aim of Experiment 3 was to determine whether the peak-and-end model (Kahneman et al., 1993) would be better able to account for performance than the peak-only model (Gallistel, 1978), now that the "weak endings" were longer and thus possibly more salient. If the peak-and-end model held in this situation, the psychometric functions for some or all of the 2+6 composites would be rightward-shifted relative to the 2 s constant-frequency trains, due to the addition of the weaker end.

Method

Subjects

Three of the 6 subjects (B37, B40, and B43) had participated in previous experiments, and the remaining rats were experimentally naive. All rats were housed and fed in the same manner as described in the previous experiments.

Surgery

Experimentally naive subjects were surgically implanted with bilateral electrodes in the same manner as previous subjects, with the exception that K5's implantation coordinates were 2.8 mm posterior to bregma, 1.7 mm lateral to the mid-sagittal sinus, and 8.3 mm below the dura (Paxinos & Watson, 1998). (This rat had been prepared originally for another study.)

Apparatus

All operant boxes were identical to those described in Experiments 1 and 2.

Procedure

Screening and training. All "new" subjects had been screened for disruptive stimulation-induced movements and/or aversive effects, and trained to self-stimulate reliably, in the same manner as in earlier experiments.

Baseline condition. This condition continued for all rats from Experiments 1 and 2, and was instituted for all "new" subjects. All procedures were conducted in the same manner as described earlier.

Single-operant "2+6" peak-and-end test. The present experiment was essentially identical to the single-operant peak-and-end experiment (Experiment 1). Only one lever was available during the trials. During the 20 s ITIs the houselight was flashed, and one priming train, identical in characteristics to the reward that would be available during the subsequent trial, was delivered noncontingently. All trials were once again 320 s in duration, with a 16 s blackout delay, and each sweep was made up of 9 trials.

Five types of stimulation trains were available over the course of the experiment: a constant-frequency 8 s train, a constant-frequency 2 s train, and three composite "2+6" trains. In the latter train type, the first 2 seconds of the stimulation were set at a constant frequency identical to that of the constant-frequency 2 s train, at a comparable "step" in the descending sweep, and the last 6 seconds were set at a lower constant frequency. As in Experiment 1, the ratio of the frequencies of the 2 s component of the composite train to the 6 s component of the train was 2:1, 4:1, and 1.4:1 for the three composite trains tested. Each experimental "sweep" consisted of a descending series of pulse frequencies. Subjects were exposed to only one of the five stimulation trains during an experimental session, and were required to reach a stability criterion with each train type before the next train type was presented.

The frequencies for the 2 s and 6 s portions of the composite trains were selected for each rat by examining their performance in the pilot frequency sweeps experiments. Specifically, the aim was to select frequencies for the 2 s portion that were strong enough to maintain robust responding by the rat, whereas the 6 s train portion required frequencies sufficiently weak that the rat was unlikely to work reliably for it in isolation. This was particularly important in the current experiment because with an "end" that is so extended in duration, had the 6 s portion not been "weak" enough, there may have been enough time for the reward signal during the 6 s portion to rise above the level achieved during the 2 s portion.

Results

The first trial of each sweep was not included in subsequent analyses, and all analyses were performed on means that had been averaged over six individual sweeps for each condition.

The raw data results of the single-operant "2+6" peak-and-end test are shown in Figure 19. The dependent variable is the proportion of rewards harvested (PRH) out of a maximum of 20, and the independent variable is the logarithm of the frequency in pulses per second (pps). From examining this figure, it appears that once again the performance curves are quite variable: while the upper asymptote (20/20 rewards harvested by the rat) was nearly always reached in each condition (with the exception of the 8 s condition (red) in the case of rat B46), the lower asymptotes had a tendency to be more variable, as was seen in Experiment 1. In addition, the distributions of the data points at the different frequencies were once again heteroscedastic, with the variability of the data points greater on the rising segments (slopes) of each performance curve, compared to those lying on the upper asymptotes of each curve.

These problems were again dealt with by rescaling and weighting each curve to remove some of the effects of this variability, as was described in Experiment 1. Briefly, a new lower asymptote was estimated using a broken-line function, and then this new estimated lower asymptote was set to zero (no rewards harvested), and the remaining data points in the curves were rescaled in order that their means would range between 0 and 1. Lastly, to remove some of the effects of heteroscedasticity, each observation making up each individual performance

Figure 19. Graphs of the raw means from the "2+6" variant of the peak-and-end test, for 6 rats (individual subjects are identified by the label in the lower right of each graph). Plotted on the x-axis are the common logarithms of the pulse frequency (taken from the 2 s portion of the train), in pulses per second. The y-axis variable is the proportion of rewards harvested, out of a maximum of 20. Error bars represent the standard error of the mean. Each curve denotes a different train duration/type tested: 8 s (red) or 2 s (yellow) constant-frequency trains, and three composite "2+6" trains. For each of these composites, the first 2 s of the train are the same frequency as the 2 s constant-frequency train, and the last 6 s are set at a lower frequency. The ratio of the frequencies in the first component of the train to the second component of the train was set to either 2:1 (green), 4:1 (cyan) or 1.4:1 (blue). Each performance curve represents responding over six descending pulse frequency "sweeps" at that particular train duration/type.

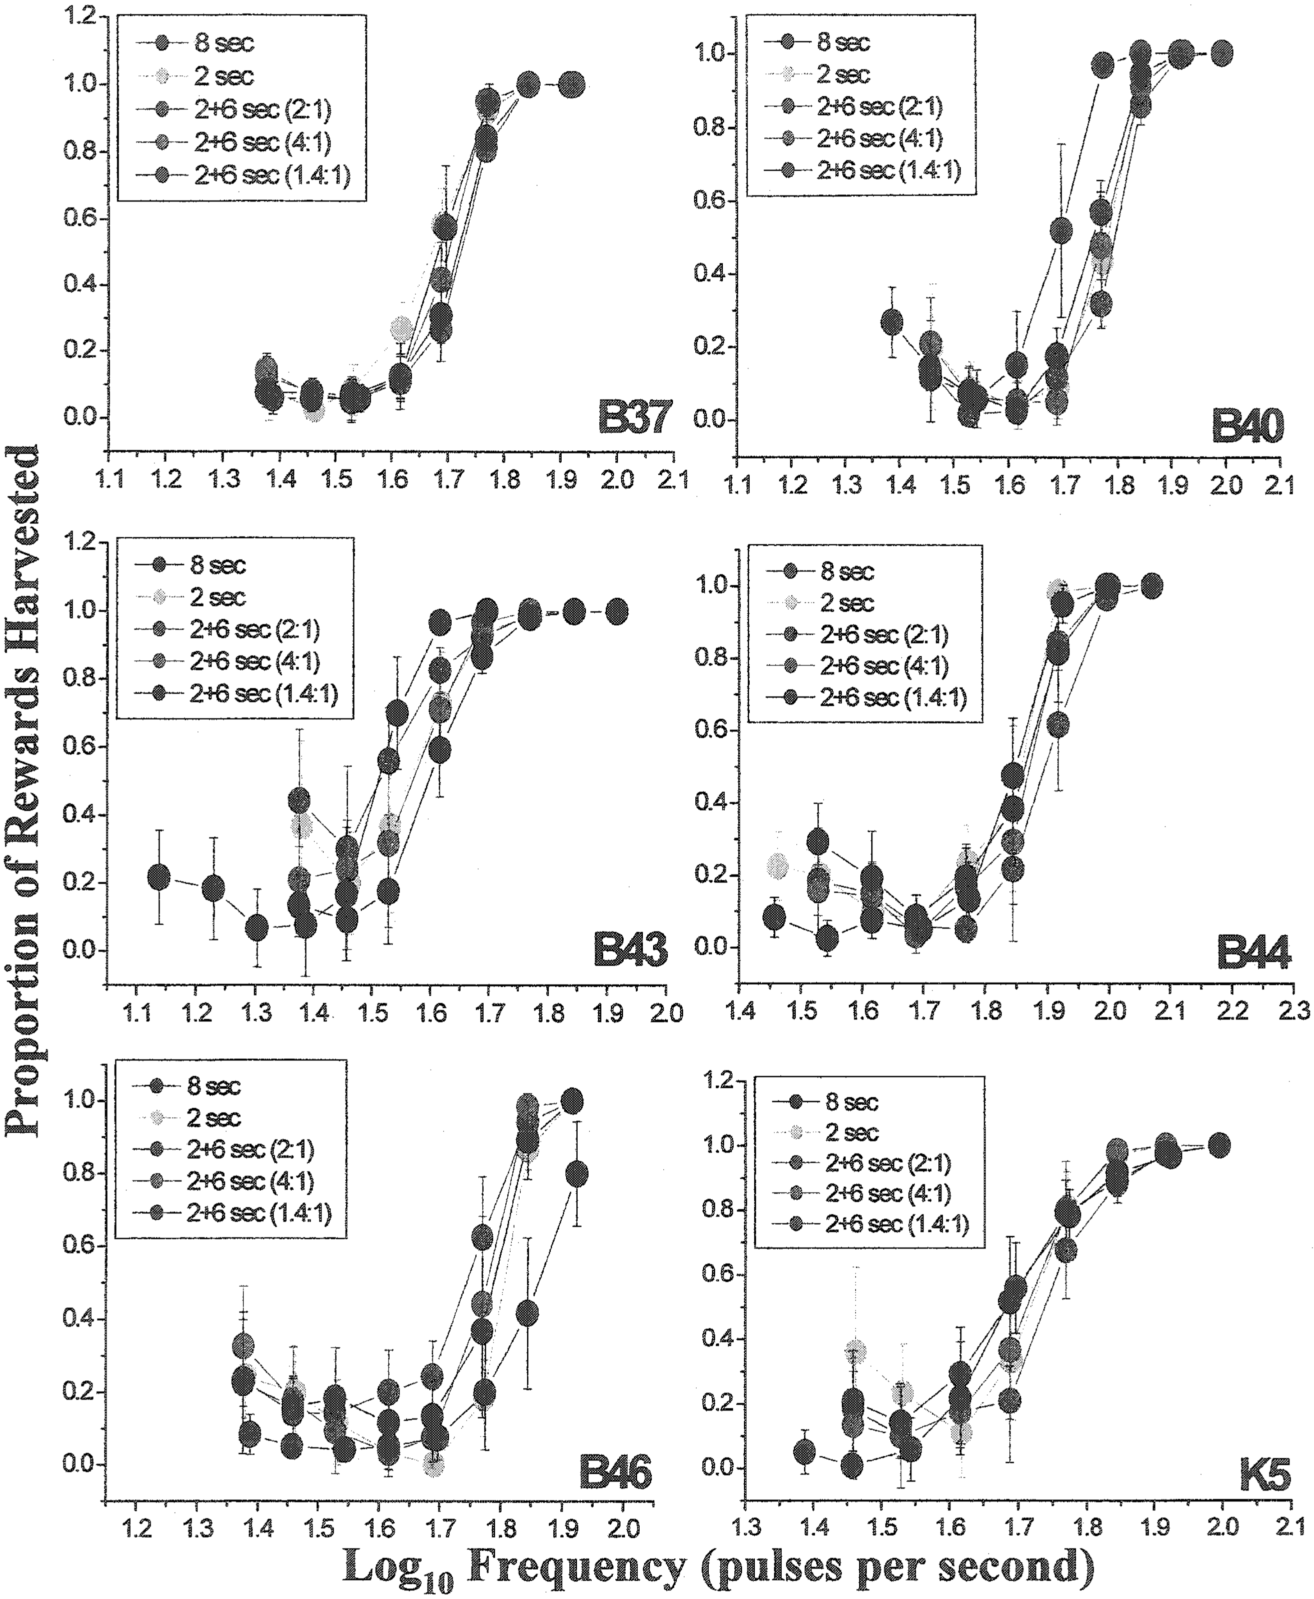

Figure 19. Raw data results from the single-operant 2+6 peak-and-end test.

curve was weighted, with the inverse of each observation's variance used as its weight, and the maximum weight set at 10.

The rescaled and weighted data are shown in Figure 20. X- and y-axes are once again identical to those in the previous figure. To aid in the interpretation of statistically significant shifts between curves, an Analysis of Covariance (ANCOVA) was performed, with the frequency set as the covariate, the train type set as the predictor variable, and the rescaled proportion of rewards harvested (PRH) as the dependent variable. The alpha level was set to 0.05 for all statistical tests. The ANCOVAs were again carried out on data points on the rising segments of each psychometric function.

As was the case in Experiment 1, some of the data, despite rescaling, were still prone to outliers and differences in the spread of residuals in some conditions and some rats. The expert statistical system (RS/1 Version 6.1, Brooks Automation, Inc., MA) suggested running an ANCOVA model based on fitting each performance curve with a separate slope, as opposed to fitting a common slope to all curves. For all rats except one (B43), the system suggested the use of the Separate Slopes model. Once again, running the Separate Slopes version did not appreciably change the p-values compared to the Common Slope version, as shown in Table 4. For this reason, and also because using a separate slopes model precludes making pairwise comparisons between performance curves, the decision was made to "force" the analyses of all rats to run as a Common slope model.

Figure 20. Graphs of the rescaled means from the “2+6” variant of the peak-and-end test, for 6 rats (individual subjects are identified by the label in the lower right of each graph). The independent variable is the frequency (expressed in logarithmic units) in pulses per second (pps) (taken from the 2 s portion of the train), and the dependent variable is the proportion of rewards harvested, out of a maximum of 20. Error bars denote the standard error of the mean. Each curve represents a different train duration/type tested: 8 s (red) or 2 s (yellow) constant-frequency trains, and three composite “2+6” trains. For each individual “2+6” composite train, the first 2-s portion is set to the same frequency as the 2 s constant-frequency train, and the last 6-s portion is at a lower frequency. The ratio of the frequencies of the “2” portion to the “6” portion of the trains are either 2:1 (green), 4:1 (cyan) or1.4:1 (blue). Each curve represents performance over 6 descending frequency “sweeps” at each train duration/type.

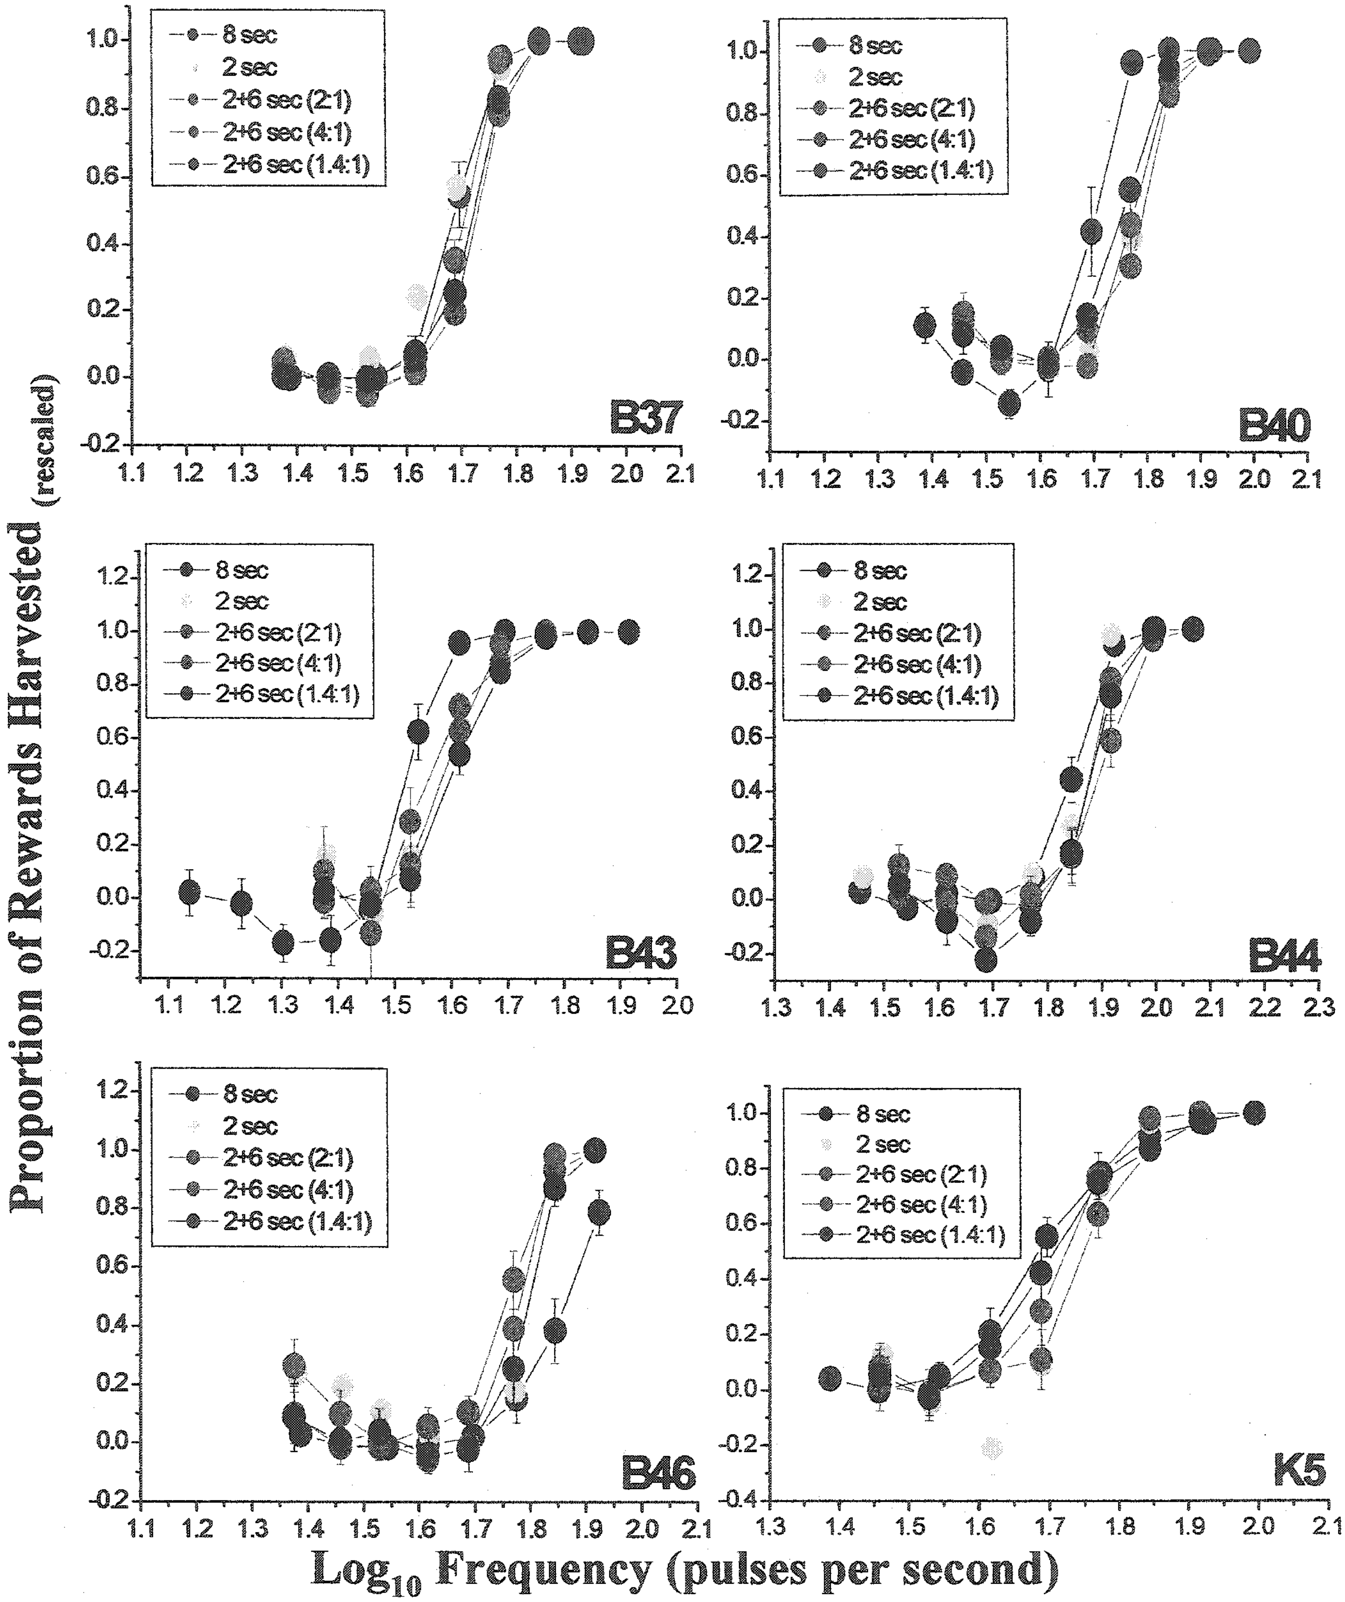

Figure 20. Rescaled means from the single-operant 2+6 peak-and-end test.

Table 4.

Differences in F-values in the Single-Operant "2+6" Peak-and-End Test

Rat	Common Slope Model	Separate Slope Model
B37	$F(4,78) = 12.241, p<0.05$	$F(4,74) = 18.2, p = 0$
B40	$F(4,90) = 17.86, p<0.05$	$F(4,86) = 31.5, p = 0$
B43	$F(4,84) = 12.108, p<0.05$	$F(4,80) = 11.8, p = 0$
B44	$F(4,90) = 20.963, p<0.05$	$F(4,86) = 23.9, p = 0$
B46	$F(4,78) = 4.902, p<0.05$	$F(4,74) = 7.53, p = 0$
K5	$F(4,120) = 5.799, p<0.05$	$F(4,116) = 6.18, p = 0.0002$

Additionally, for rats B37, B40, B44, and K5, the expert statistical system reported that some performance curves had larger spreads than others, and it recommended that weighted computations be performed such that the data points from the less reliable performance curves would be down-weighted. For the remaining 2 subjects, B43 and B46, the spread of each performance curve was roughly equal, and thus the system recommended unweighted computations for these subjects. In all cases, the recommendations of the expert system were followed.

Lastly, the system once again recommended (for rats B37, B40, B43, and B44) that a "Robust" rather than a "Least Squares" method be used to estimate the slopes and intercepts of the curves, as the Robust method is designed to minimize the effects of outliers. For the remaining rats, B46 and K5, the Least Squares estimates were sufficient. Again, the recommendations of the expert system were followed.

For rat B40, there is a significant difference between curves, $F(4,90) = 17.86$, $p<0.05$, with tests based on simultaneous 95% confidence intervals revealing that the significant difference lies between the 8 s (red) curve and each of the other curves. Only one other comparison is significant, which is the 2:1 composite (green) curve versus the 1.4:1 composite (blue) curve, but the weakest composite, the 4:1 composite (cyan) curve, falls between the other two curves. Given that none of the composites differ significantly from 2 s (the key comparison) this finding, though statistically significant, is likely spurious. In the case of rat B43, the difference between the curves was statistically reliable, $F(4,84) = 12.108$, $p<0.05$, due to a difference between the 8 s (red) curve and all the other curves, which do not differ significantly from each other.

In the case of rat B46, there is a significant difference between performance curves, F(4,78) = 4.902, $p<0.05$, and the tests based on simultaneous 95% confidence intervals revealed that this significant difference is between the 8 s (red) curve and the 2:1 composite (green) curve. There is also a significant difference between the 8 s curve and the 4:1 composite (cyan) curve. Rat K5 displays a significant difference, F(4,120) = 5.799, $p<0.05$, between the 8 s curve (red) and the 2:1 composite (green) curve. For both of these subjects, therefore, there are no statistically significant differences between the 2 s curve and the "2+6" curves, the comparison of key interest, but the expected difference between the 8 s curves and *all* the other curves did not hold. Indeed, in the case of subject B46, the 8 s curve is displaced to the right rather than the left. This suggests that the 8 s train was *less* rewarding than the shorter 2 s trains, contradicting this subject's own performance in pilot testing (see Figure 12).

The final two subjects show even less orderly results. Rat B37 displays a significant difference between curves, F(4,78) = 12.241, $p<0.05$, with comparisons based on simultaneous 95% confidence intervals revealing that there is a significant difference between the 2:1 composite (green) curve and each of the following curves: 8 s (red), 2 s (yellow), and the 4:1 composite (cyan). In addition, there is a significant difference between the 2 s (yellow) curve and the 1.4:1 composite (blue) curve. For rat B44, there is once again a statistically significant difference, F(4,90) = 20.963, $p<0.05$, with tests based on simultaneous 95% confidence intervals revealing six significant pairwise differences: 2:1 vs. all 4 of the other curves, and

1.4:1 vs. both 8 s (red) and 2 s (yellow). Thus, in the case of both rats, there is a significant difference between the 2 s curve and two of the composites (1.4:1 and 2:1). It should be noted, however, that for both animals the *weakest* of the three composites, the 4:1 (cyan), is *not* significantly different from 2 s, and lies to the left of the other two composites. In terms of the expected difference between the 8 s curve and all other curves, this again did not hold.

The majority of the subjects in the current experiment displayed statistically significant shifts between several psychometric functions. However, as noted in the Discussion of Experiment 1, this does not necessarily mean that these shifts were systematic and meaningful. The shifts are quite small compared to those demonstrated by the same rats in the pilot frequency-sweeps tests (Figures 11 and 12). The small magnitude of these shifts, coupled with their non-systematic nature, argues against treating them as behaviourally significant.

Given the findings of Sonnenschein, Conover, and Shizgal (2003), which suggested that duration neglect is observed once train durations exceed 2-4 s (given a sufficiently high frequency), one might have expected that there would be a significant difference between the 2 s and 8 s psychometric functions in the present experiment. The 8 s train falls well within the previously established range of duration neglect, whereas the 2 s train might not be sufficiently long to show neglect in the case of most rats. While a significant difference between 2 s and 8 s curves was not seen consistently in the present experiment, the shifts that were obtained are not inconsistent with previous findings. As noted earlier, all rats serving in the present experiment had previously trained in a shortened version of the frequency sweeps experiment. The results for the 2 s and 8 s conditions from both the current experiment and the frequency sweeps experiment are plotted in Figures 21 and 22. In the majority of cases (B40, B43, B44, and K5), the difference between the 2 s and 8 s conditions was comparable in both experiments.

To summarize, it appears that for 4 out of 6 subjects, there is no behaviourally significant difference between responding for the 2 s regular train and responding for a "2+6" composite train. Although larger shifts were expected between the 2 s and 8 s curves, the shifts obtained in some rats were consistent with their performance for the same trains in the frequency sweeps pilot studies. This is not a problem for making an evaluation of the "peak-alone" versus the "peak-and-end" model, in any case, as the key comparison is between the 2 s vs. 2+6 s composite trains.

Figure 21. Graphs comparing performance in the pilot frequency-sweeps test to performance in the 2+6 variant of the peak-and-end test. Only the performance curves for the 8 s and 2 s trains are shown in each graph. Rats are identified by the labels in the lower right of each graph set. For both graphs, the common logarithm of the frequency is represented on the lower x-axis. Plotted on the y-axis is the rescaled proportion of rewards harvested, out of a maximum of 20. The rescaling procedure sets the lower asymptote of each curve to zero and the upper asymptote to one. Each data point is the mean of six observations; error bars represent the standard error of the mean. In the frequency-sweeps graphs, the 2 s curve is shown in cyan, and the 8 s curve in orange. In the 2+6 graphs, the 2 s curve is shown in yellow, and the 8 s curve in red.

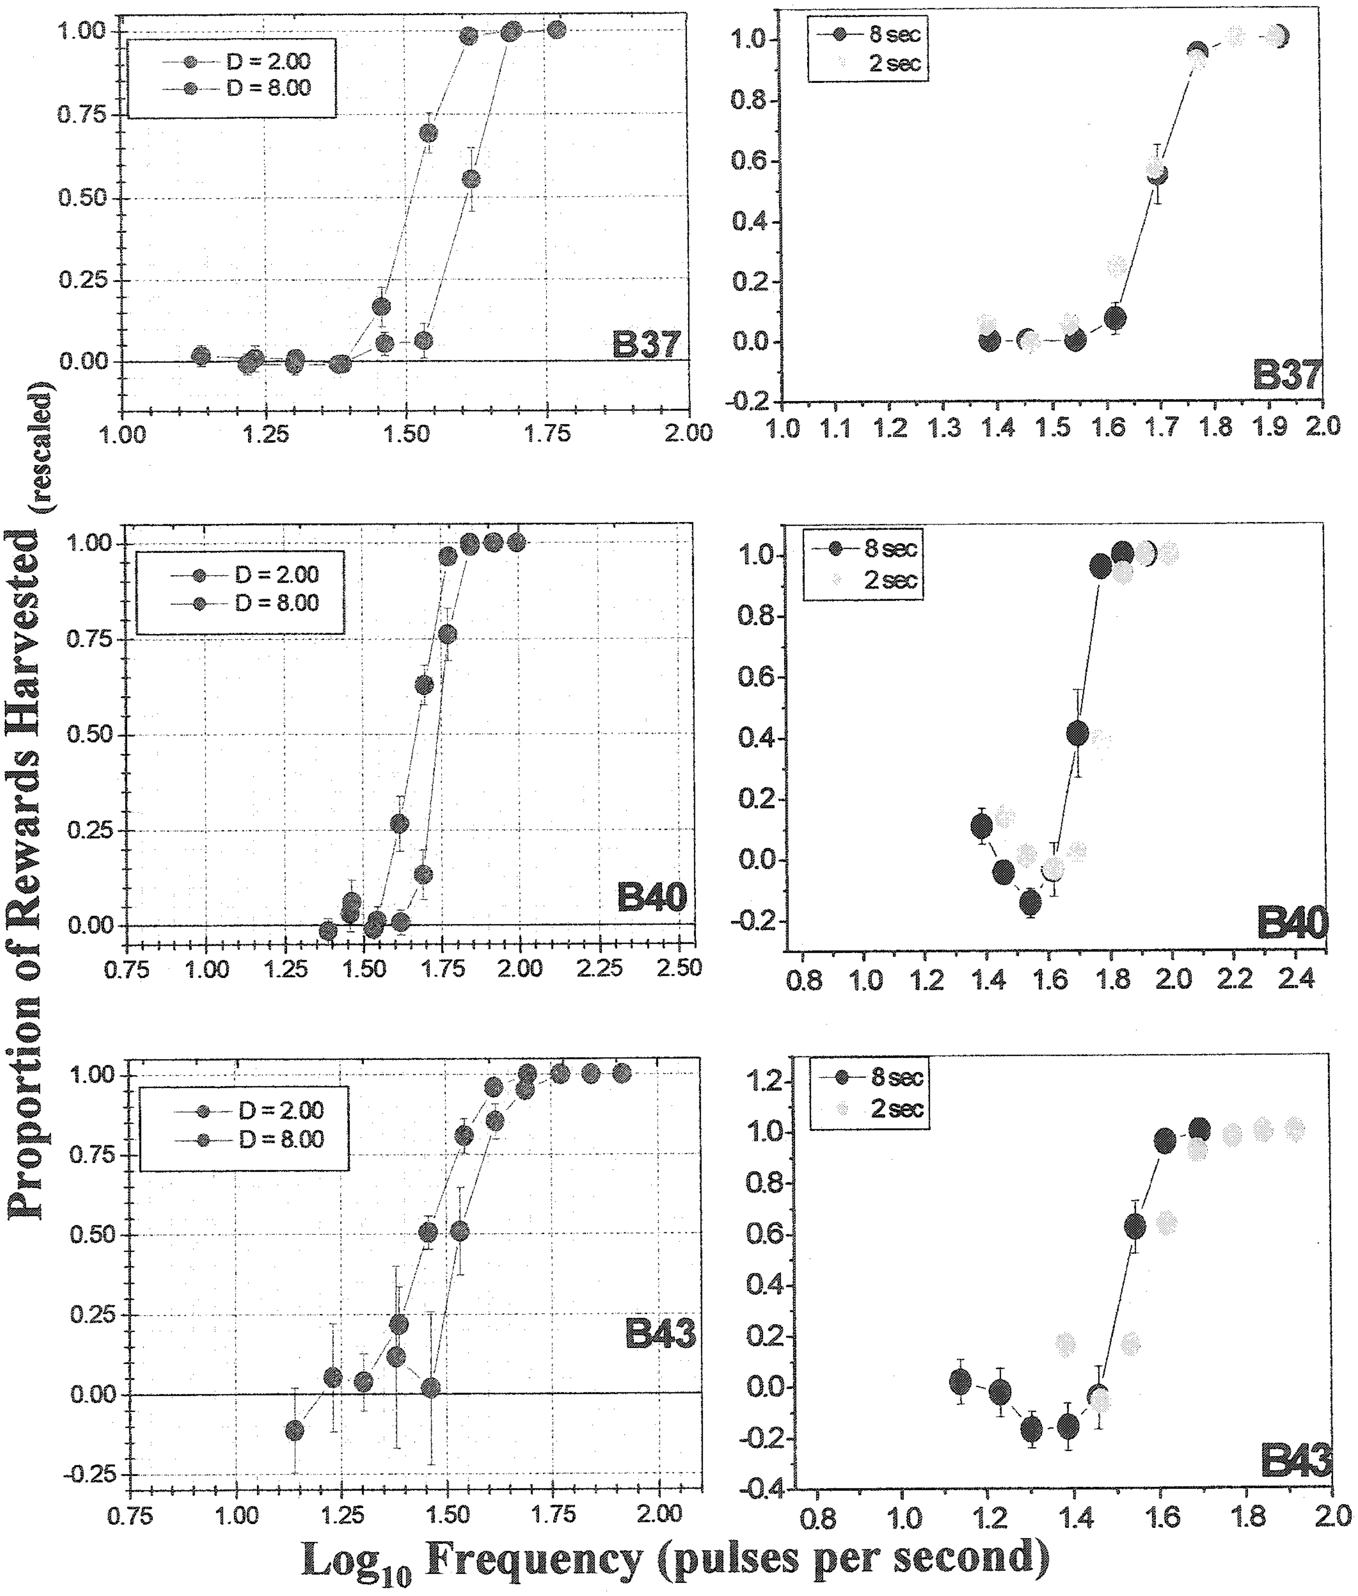

Figure 21. Graphs comparing performance (rescaled data) for 2 s and 8 s rewards in both the frequency sweeps and 2+6 experiments.

Figure 22. Additional comparisons of performance in the pilot frequency-sweeps test to performance in the 2+6 variant of the peak-and-end test. (For details, see caption for Figure 21.)

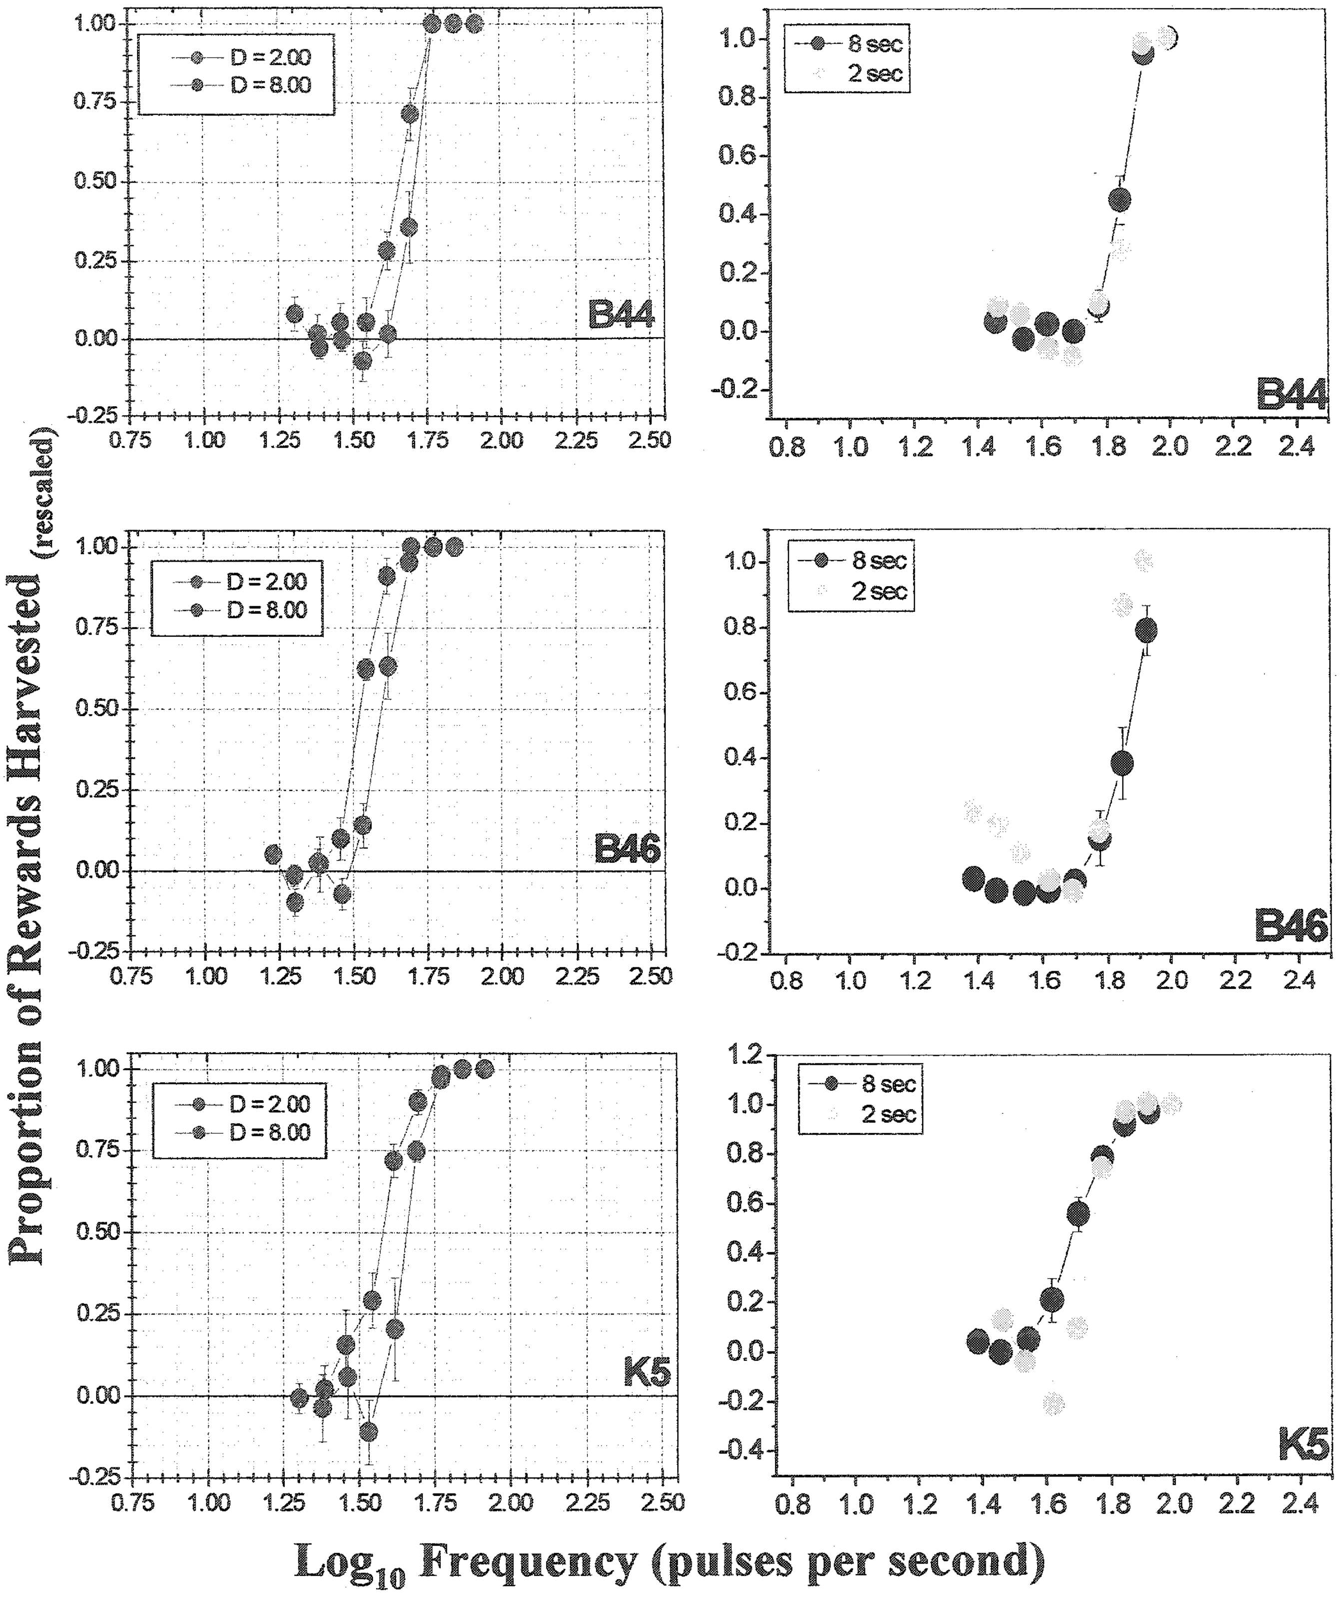

Figure 22. Graphs comparing performance (rescaled data) for 2 s and 8 s rewards in both the frequency sweeps and 2+6 experiments, for the remaining subjects.

Discussion

Although the results of Experiment 3 are less clear than those of Experiments 1 and 2, they are not seriously inconsistent. For the majority of the subjects (4 out of 6), there were no significant differences between responding for the 2 s train and responding for a "2+6" composite train. Once again the peak-alone model (Gallistel, 1978) is sufficient to account for the results, and there is little support for the peak-and-end model (Kahneman et al., 1993). Even if one were prepared to argue that the small statistically-significant shifts indicate the presence of a peak-and-end effect (particularly for rats B37 and B44), the effect of the "end" even in those cases is quite small, particularly in comparison to shifts seen in the earlier frequency sweeps experiments. Also, the shifts between the 2 s and 8 s curves were often consistent with the psychometric functions for 2 s and 8 s trains in previous experiments, suggesting that while large shifts between the two curves were generally not observed in the present experiment, the findings are in line with previous demonstrations of progressive insensitivity to duration. This was not an issue in the "6+2" experiment, as the two control trains, 8 s and 6 s, both fell within the range of durations where rats were likely to show duration neglect. In the current experiment, rats were expected to be more sensitive to the difference between the 2 s and 8 s, due to the fact that most of the rats in Sonnenschein, Conover, and Shizgal (2003) only displayed duration neglect once trains exceeded 4 s. In fact, although the distance between these curves for many rats in Experiment 3 were not that large, many rats nonetheless still displayed shifts consistent with those obtained in their earlier frequency sweeps (Figures 21 and 22).

In light of the less-than-orderly data of some of the subjects, it was decided to undertake a brief one-condition dual-operant study. In this case, the Std train was a "2+6" composite, with a 2:1 ratio between the frequency of the first portion and the frequency of the second portion, and the Alt train was a constant-frequency 2 s train. All other parameters of the experiment were set as in previous dual-operant experiments. This served as a test of the peak-and-end model using a composite with a longer "weak" ending, but within a paradigm that would allow for simultaneous choice between both types of trains.

Preliminary results from this experiment, available at the time of this writing in two rats, are shown in Figure 23. The dependent variable is the proportion of rewards harvested, and the independent variable is the common logarithm of the frequency of the Alt train, with the frequency of the Std train, also in common logarithmic units, shown as a vertical magenta line. There is no statistically significant difference between the cross-point (CP) and the magenta line (frequency of the Std train). Also, the behavioural shift is small, a mere 0.01 to 0.02 log units. Recall that rat B37, in the single-operant 2+6 test, displayed a statistically significant difference between the 2:1 composite train and the 2 s constant-frequency train, with the former train rightward-shifted. The failure to observe a leftward shift in the CP in the present dual-operant test thus suggests that the small yet statistically significant difference between the same trains in the single-operant context may indeed have been due to chance. In addition, as rat K5 did not display any difference between the performance curves for the 2+6 and 2 s curves in the single-operant context, the results of the dual-operant 2+6 test are consistent. The results of the present brief dual-operant experiment are thus in line with those of the previous experiments.

Figure 23. Results from the dual-operant variant of the “2+6” experiment, for 2 subjects (individual subjects are identified by the label in the lower right of each graph). The Std was a composite “2+6” train, and the Alt was a 2 s train. The 6 s portion of the “2+6 s” train was set to a low pulse frequency, which was half that of the 2 s portion of the train. The x-axis variable is the common logarithm of the frequency of the Alt train, in pulses per second (pps). The common logarithm of the frequency (in pps) of the Std train (taken from the 2 s portion of the composite train) is shown as a vertical magenta line. The proportion of rewards harvested (out of a maximum of 20) is the y-axis variable. Red circles represent rewards obtained by pressing the Alt lever, and green circles represent rewards obtained by pressing the Std lever. The green and red lines are the respective broken-line fits. Error bars represent the standard error of the mean, and all curves are made up of data that has been averaged over 12 experimental sweeps. The blue circle represents the “cross-over” between responding preferentially on one lever to responding preferentially on the other lever, with a 95% confidence band.

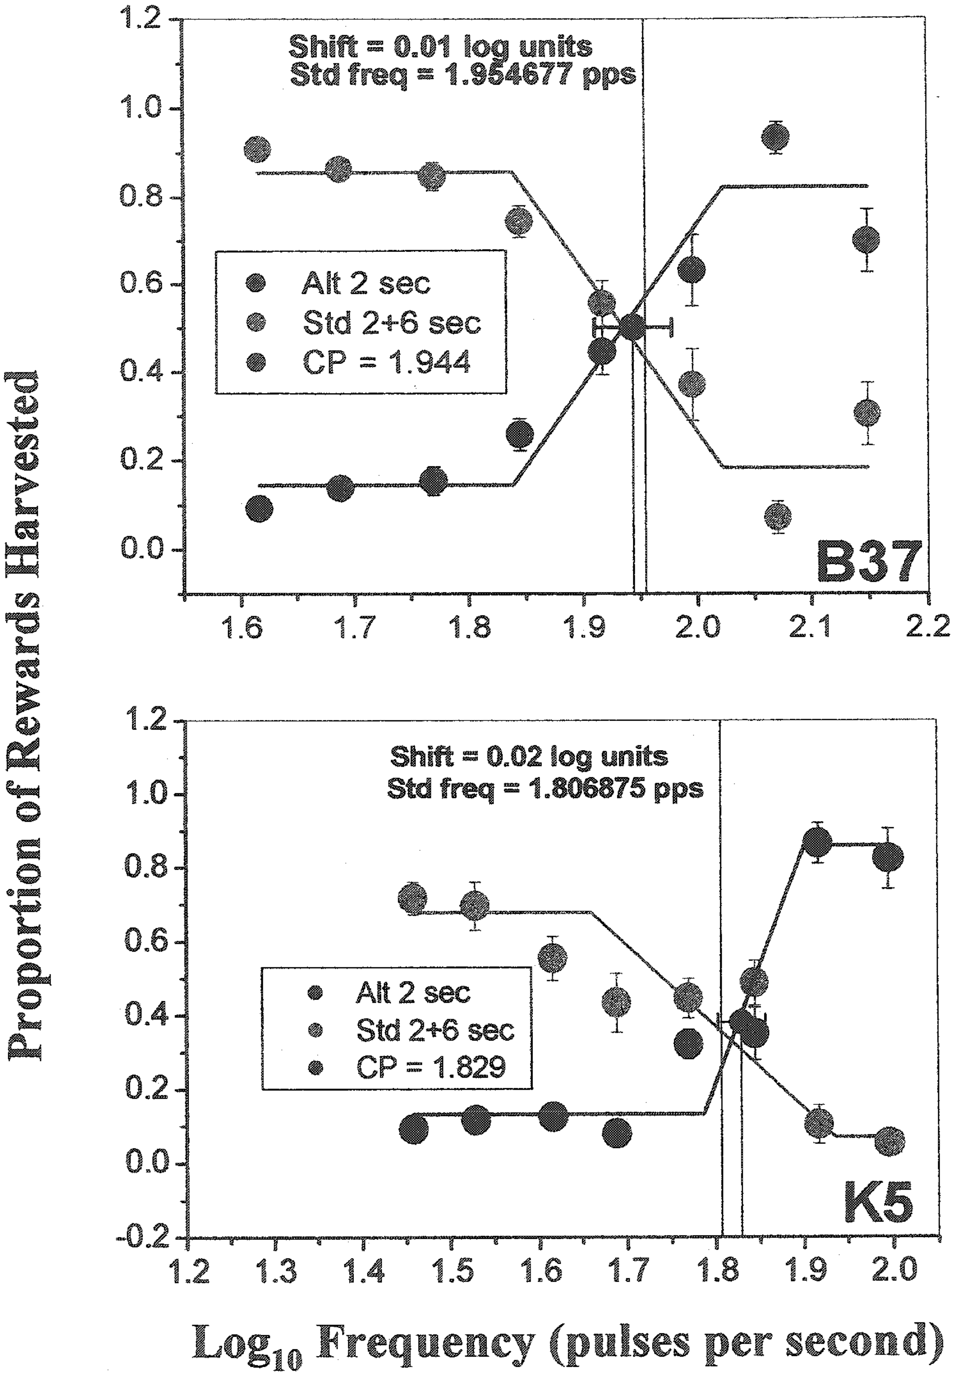

Figure 23. Preliminary results from the dual-operant 2+6 peak-and-end test

Taken together, the single- and (minimal) dual-operant data suggest that even when a longer weak ending is added to a strong train, and regardless of whether the experimental paradigm is conducted in either single- or dual-operant mode, appending this weak end did not exert a large influence on the overall reward value of the train.

Given that single-operant paradigms are especially vulnerable to changes in reward value caused by factors outside of the experimental condition, such as the electrode drift, and given that the results of some rats (B46) appear grossly inconsistent with the results obtained in the frequency sweeps experiments, one might be tempted to conclude that perhaps a general shift did occur in the rats' response to stimulation. However, as mentioned earlier, to ensure that such a shift did not go unnoticed, the baseline condition was run continuously throughout the experiment, and the results did not show any comparable shifts in the baseline data, as one might expect if some reward-altering process was taking place.

To summarize, it appears that irrespective of whether a single- or dual- operant paradigm is used, or whether the weak "ending" of the composite trains is short or long in duration, the simpler peak-alone model of Gallistel (1978) is sufficient to account for the behaviour of the rats. There are several differences between the Kahneman group's studies and the current set of experiments that may offer some clues as to why a peak-and-end effect was not seen in the present studies. Most of these issues will be discussed more fully in the General Discussion, but to briefly summarize: 1) the Kahneman studies employed human subjects whereas the current studies employed nonhuman animals, 2) Kahneman's group used aversive stimuli whereas the current experiments made use of appetitive stimuli, 3) in the Kahneman studies the humans were generally the passive recipients, whereas in the present study the rats had to be actively responding in order to obtain the stimulation, and 4) the Kahneman experiments used stimuli ranging from one minute to an hour, whereas in the present experiments the longest "experience" possible was 8 seconds.

The results of Experiments 1-3 suggest that the simpler peak-only model (Gallistel, 1978) is sufficient to account for the rats' behaviour. However, one limitation of the current study was that only the ends of the stimulation trains were varied. If the beginnings of the trains were varied, this might not only influence the reward value of the train, but also unmask another exemplar that could be taken into account when evaluating appetitive experiences. The aim of Experiment 4 was to examine this question.

Experiment 4

Kahneman and colleagues (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996) suggest that when people evaluate *aversive* experiences, the two exemplars taken from the experience while performing retrospective evaluations are the "peak" moment of pain, and the amount of pain at the "end" of the experience.

Whether the same moments of the experience are selected as exemplars when the experience is *positive*, however, has not yet been assessed. Everyone has had the experience of going to a restaurant and eating a rich dessert. The first bite of the dessert is certainly a salient moment, as the patron gets a first taste of the "rewarding" flavour of the dessert. This suggests that when an appetitive experience is evaluated, the reward level at the "beginning" of the experience is a salient exemplar, perhaps more so than the reward at the "end" of the experience. In effect, perhaps different heuristics are employed for different types of experiences: peak-and-end computations for aversive experiences, and beginning-and-peak (or peak-alone) for appetitive experiences. It seems logical that in the case of an aversive experience, an organism would be interested in having the experience *terminate* as quickly as possible, and thus the "ending" may have particular significance. In the case of an appetitive experience, however, it could be argued that an organism would be most interested in having the experience *start* as soon as possible, and thus the "beginning" may have particular salience.

The aim of Experiment 4 was to determine whether a new combination of exemplars, the "beginning" and the "peak", would better account for the rats' behaviour than the "peak" alone model (Gallistel, 1978). Although Experiments 1-3 supported the peak-only model, in none of these experiments was the frequency at the beginning of the train lower than the frequency at the end of the train. The present experiment therefore utilized a novel type of composite train, in which the first 6 s of the train were set to the weak frequency, and the second, 2 s portion of the train, was set to the stronger frequency.

The key comparison pitted a train consisting of 6 s of weak stimulation followed by 2 s of stronger stimulation (6w+2 s) against a constant-frequency 2 s train. This comparison directly assesses the effect of adding a "weak" beginning. However, there is a potential confound. The composite train not only has a "weak" beginning, but the "strong" portion of the composite train is delayed by 6 s. The rat presses the lever, but then is not immediately rewarded with a strong stimulation train. This delay of reward could have quite an impact on the evaluation of the stimulation. Studies have shown that delaying a reinforcer can slow the acquisition of self-stimulation behaviour in rats (Black, Belluzzi, & Stein, 1985), although neither training nor immediate reinforcement is necessary to establish a new behaviour, and acquisition of behaviour is still possible even with delayed reinforcers (Lattal & Gleeson, 1990). It has also been shown that delaying reinforcement causes increases in BSR thresholds (Fouriezos & Randall, 1997), and in order for a delayed train to be preferred to an immediate one, its strength must be increased (Mazur, Stellar, & Waraczynski, 1987).

Delay discounting could have implications for how the composite train is evaluated. One possibility is that the effect of delay, as just described, will combine with the effect of adding a "weak beginning", and this will lead the composite train to be the most devalued. The "zero beginning" train, with just a delayed peak, will also be devalued relative to the 2 s constant-frequency train, but less so than the 6w+2 s composite. This argument is analogous to the predictions of the peaks-and-ends model (Kahneman et al., 1993).

Another possibility is that the "weak" portion of the train could "bridge" the gap between the operant response and the beginning of the more rewarding train segment. Put another way, the weak train undermines the delay discounting effect by mitigating the delay period, thus lessening the effect of the delay. Indeed, "filling" the delay period with some kind of signal, such as illuminating the operant box with a light, has been shown to produce higher rates of responding for a natural reinforcer than when the delay is not accompanied by such a "signal" (Reed & Reilly, 1990). If so, adding a "weak" beginning could have two opposite effects: the value of the 2 s component would be less undermined by the delay, but the average of the weak signal at the beginning of the train and the strong signal at the end would be less than the 2 s component alone.

In order to assess the contribution of delay discounting, the composite train was compared both to a 2 s train, and to a 2 s train delayed by 6 s.

Method

Subjects

All rats in the present experiment had served in several of the previous experiments. All were housed and food restricted as described earlier.

Surgery

All 3 subjects had been implanted with stimulating electrodes as described previously.

Apparatus

All operant boxes used were identical to those described in Experiments 1-3.

Procedure

Baseline condition. This condition continued throughout the beginnings-and-peaks experiment.

Dual-operant beginnings-and-peaks test. The experimental procedure was essentially identical to the previous dual-operant experiment (Experiment 2). The initial condition comprised a delayed 2 s train, available on the Std lever, and a regular 2 s train available on the

Alt lever. In the case of the Std train, delivery of the stimulation proceeded as follows: after the lever was pressed by the rat, the lever retracted into the wall, but nothing was delivered for 6 s. Once this 6 s period had elapsed, the 2 s train was delivered. As the reliability of performance had already been assessed in the context of another recent experiment, a 2 s Std vs. a 2 s Alt condition was not run in the present study. Although an attempt was made to train the rats to press for delayed stimulation trains on both the Alt and Std levers, the rats proved unable to work reliably. It was thus decided to utilize an undelayed Alt train for all conditions in the present experiment. In addition, as the Std lever reward was initially set at only a moderately-rewarding frequency, there would be plenty of opportunity to raise its frequency to offset any devaluing of stimulation due to delay discounting, in order to compensate for the effect of the 6 s delay.

In the second condition, the Std lever reward was a "6w+2 s" train, in which the first 6 s of the train was the low frequency portion of the train, and the last 2 s were at the higher frequency (ratio of frequencies 1:2). The Alt train was a constant-frequency 2 s train, and all other characteristics of the trains and experimental sessions were set as described earlier.

It was also decided to run another version of the dual-operant test. In this case, an attempt was made to raise the frequency of the Std train sufficiently, so that the reward value of its strong 2 s component would approach5 an asymptotic level. For the first condition, this "raised" Std train was set at 2 s and run against an Alt lever reward that was also 2 s in duration. For the second condition the "raised" Std train was delayed by 6 s ("zero beginning") and run against a 2 s Alt train. Lastly, the "weak beginning" condition was run, using the new “raised” 6w+2 s Std train, with the 2 s at the new higher value, and the 6w segment set to the same value used in the previous “unraised” version of the current experiment.

Results

The results from the first condition of the dual-operant beginnings-and-peaks test are shown in Figure 24. This test pitted a 2 s Std train delayed by 6 s, against a 2 s Alt train. X- and y-axis variables are the same as those described in Experiment 2, and the distance between the vertical magenta line (frequency of the Std lever reward) and the cross-point (CP; blue circle) is again the main feature of interest. As can be seen from the graphs, there is a statistically significant difference between the CP and the Std lever frequency, as indicated by the magenta line falling far outside the 95% confidence interval around the CP. Also, the CP is shifted left relative to the magenta line, indicating the rats preferred the undelayed Alt train to the delayed Std train. Lastly, these shifts are very large, as the obtained log unit shifts of 0.24 to 0.33 represent a 73.8-113.8% increase in the frequency required to motivate the rats to switch from the Alt lever to the Std lever.

The second condition involved a composite "6w+2 s" train, available on the Std lever, and a 2 s Alt train. Recall that in the case of the 6w+2 s train, the first 6 s were set to a low frequency, and the last 2 s to a high frequency (ratio 1:2). The results are shown in Figure 25, where it can be seen that adding a weak beginning to the "2 s" portion of the Std train led to a decrease in the distance between the magenta line and the CP. In other words, the effectiveness of the Std was increased by adding the weak beginning.

For rat B37, while there was still a statistically significant shift between CP and the frequency of the Std lever, it was much smaller than that of the first condition. Only a 0.07 log

Figure 24. Results from the delay condition of the dual-operant beginnings-and-peaks test, for 3 subjects (individual subjects are identified by the label in the lower right of each graph). The Std was a rectangular 2 s train delayed by 6 s, and the Alt was a rectangular 2 s train. The frequency (expressed in logarithmic units) of the Alt train, in pulses per second (pps), is the independent variable, and the frequency (expressed in logarithmic units) of the Std train is denoted by a vertical magenta line. The proportion of rewards harvested (out of 20) is the dependent variable. Red curves represent rewards obtained by pressing the Alt lever, and green curves rewards obtained by pressing the Std lever. The green and red lines are the respective broken-line fits. Error bars denote the standard error of the mean, with all data points drawn from data collected and averaged over 12 experimental sweeps. The blue point represents the estimated "cross-over" between responding preferentially on the Std lever to responding preferentially on the Alt lever, with a 95% confidence band.

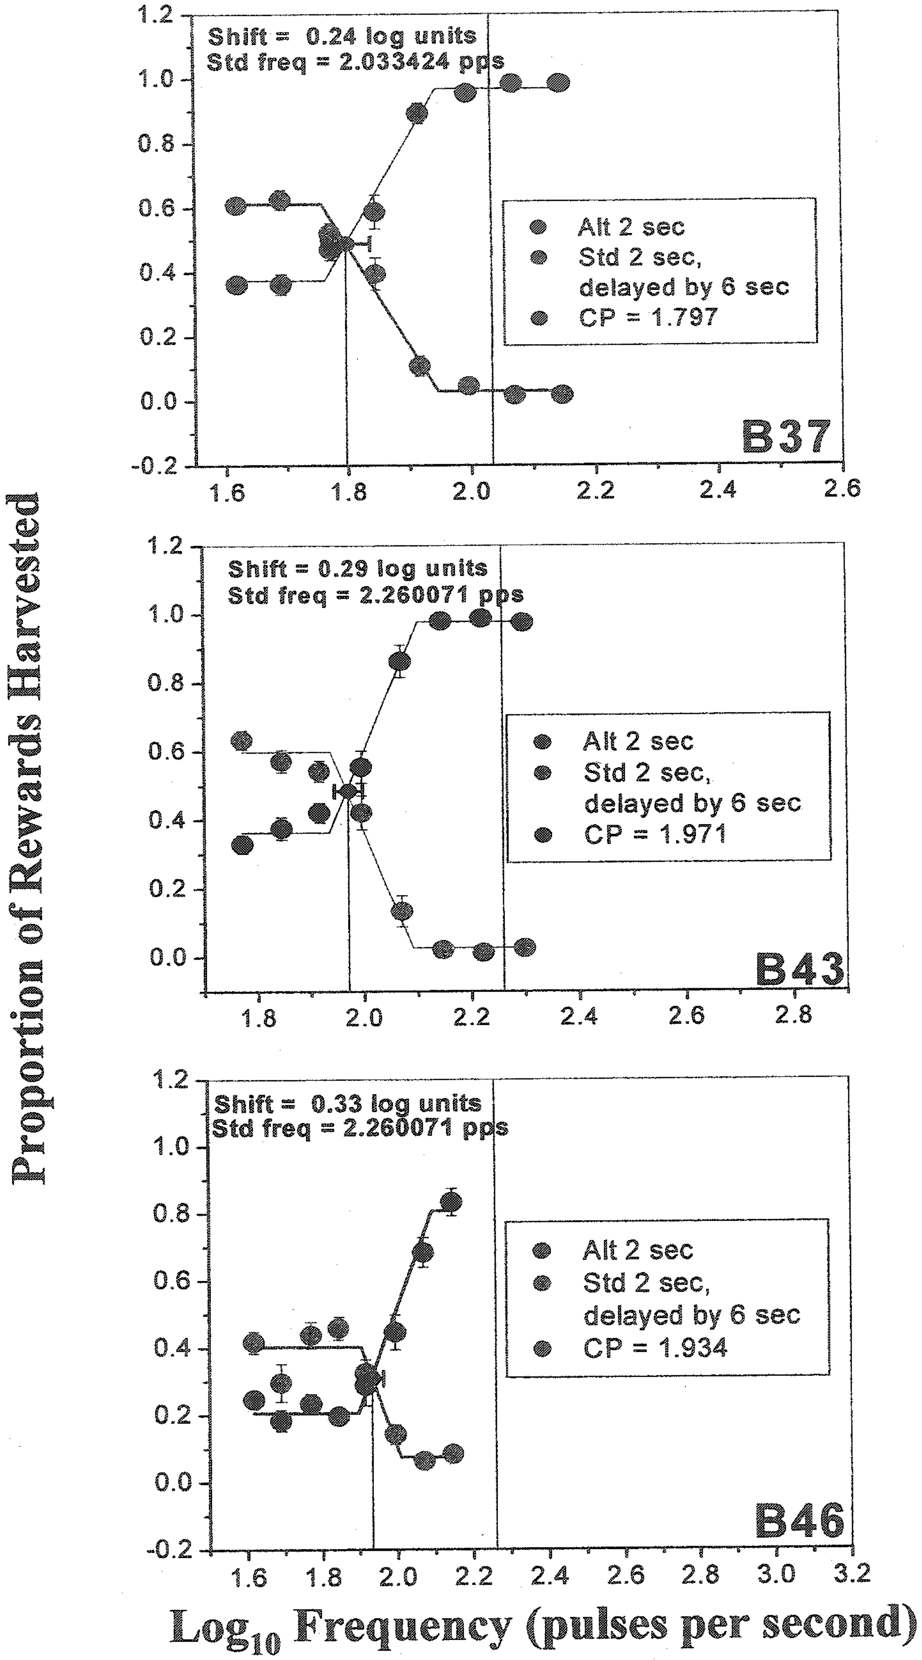

Figure 24. Results of the delayed Std train vs. the 2 s Alt train condition of the dual-operant beginnings-and-peaks test.

Figure 25. Results from the dual-operant beginnings-and-peaks test, for 3 subjects (individual subjects are identified by the label in the lower right of each graph). The Std was a composite "6w+2 s" train, and the Alt was a 2 s train. The 6 s portion of the "6w+2 s" train was set to a low pulse frequency, and the 2 s portion to a high frequency, with the ratio of frequencies in the 2 train components set to 1:2). The common logarithm of the frequency of the Alt train, in pulses per second (pps), is the x-axis variable, and the common logarithm of the frequency (in pps) of the Std train (taken from the frequency of the 2 s segment of the composite train) is denoted by a vertical magenta line. The proportion of rewards harvested (out of 20) is the y-axis variable. Red curves represent rewards obtained by pressing on the Alt lever, and green curves represent rewards obtained by pressing on the Std lever, with error bars representing the standard error of the mean. The green and red lines are the respective broken-line fits. All curves are made up of data averaged over 12 sweeps. The blue point represents the "cross-over" between responding on one lever to responding preferentially on the other lever, with a 95% confidence band plotted around it.

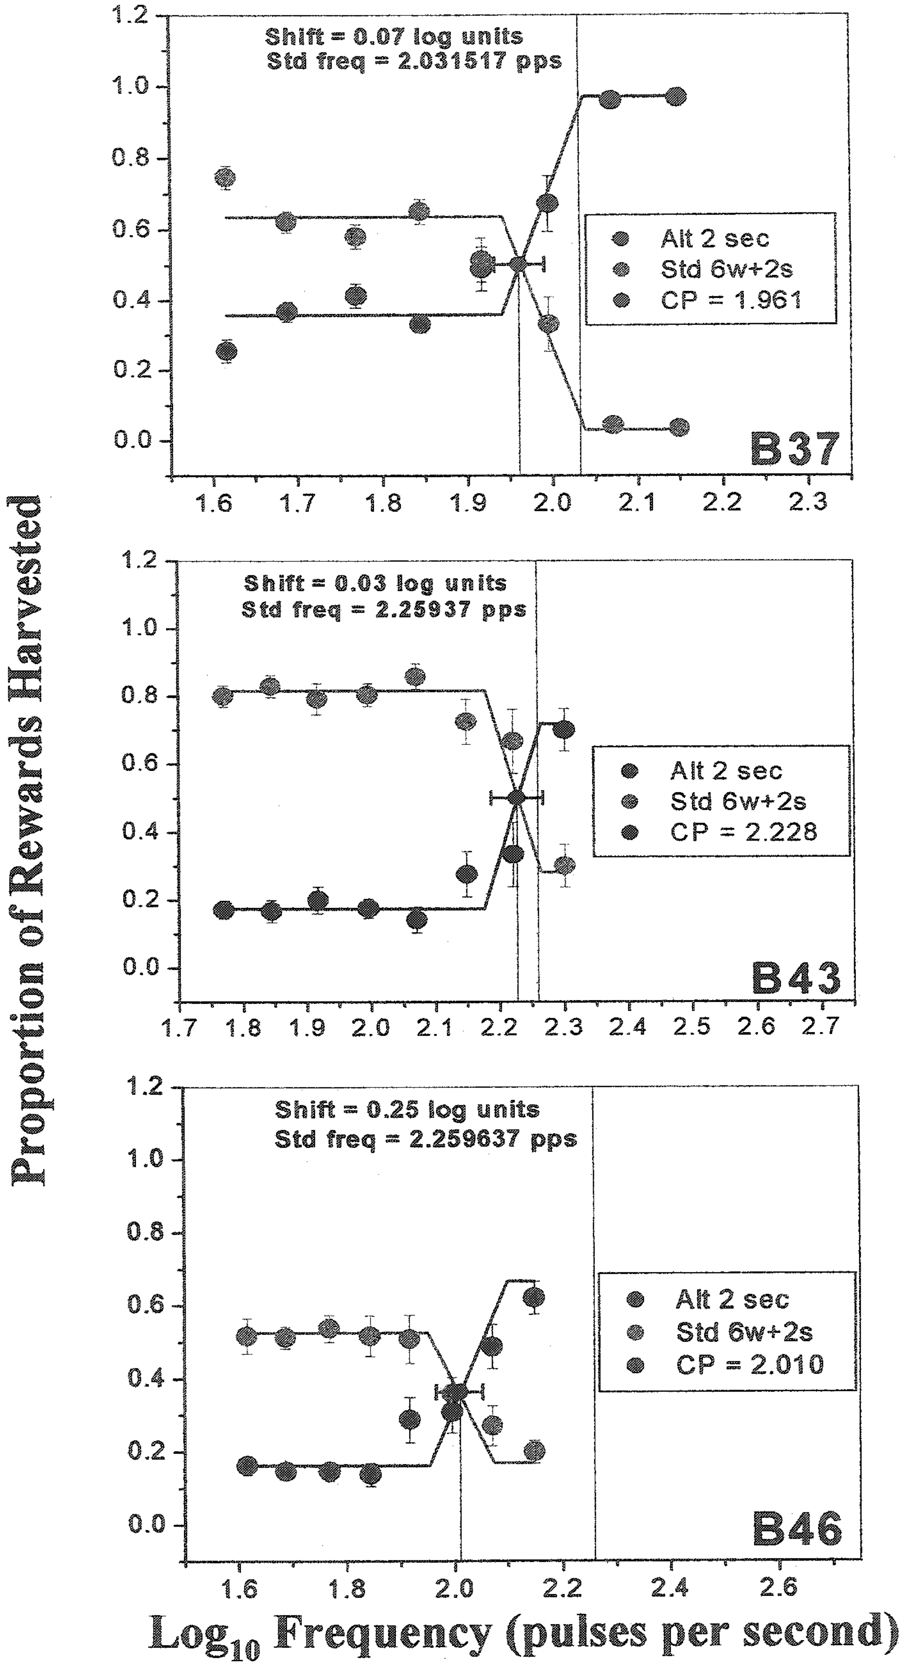

Figure 25. Results of the "6w +2 s" Std train vs. the 2 s Alt train condition of the dual-operant beginnings-and-peaks test.

unit shift was displayed when the Std included a weak beginning, as compared to a 0.24 log unit shift (a 17.5% increase required to switch, as opposed to a 73.8% increase) when the Std consisted of a delayed 2 s train. For rat B43, the difference between conditions was even more dramatic, the difference between CP and Std frequency disappeared altogether. There was a large and statistically significant shift (0.29 log units) when the delayed Std train was pitted against the 2 s Alt train in the first condition, but when the weaker beginning was added to the Std train, the statistical significance disappeared, and the CP was shifted by only 0.03 log units (a 95% increase versus a 7.2% increase, respectively). Lastly, for B46 there was still a statistically significant shift, which was 0.08 log units smaller than in the previous condition: a 0.33 log unit shift when the Std was a delayed 2 s train, and a 0.25 log unit shift when the Std was the 6w+2 s composite train (a 113.8% increase versus a 77.8% increase, respectively).

One rat[3] was run in a modified version of the experiment. The objective was to increase the frequency of the Std 2 s train until the rat was indifferent between the Alt and Std trains, even at the highest Alt frequencies. As shown in Figure 26, this objective was not attained convincingly. Although there was no clear preference between the Alt and Std levers at the highest Alt frequency, it is possible that a preference for the Alt would have been observed had even higher frequencies been tested. Unfortunately, the side-effects of the high-frequency stimulation were too strong to make such a test feasible.

---

[3] This experiment was initially run with more subjects. However, due to various factors such as illness, old age, and headcap problems, no other subjects were able to complete the experiment.

Figure 26. Results from the modified version of the dual-operant beginnings-and-peaks test, for 1 subject. In the first condition, the frequency of the Std train was increased until the rat was nearly indifferent between rewards at the highest Alt train frequency. Both the Std and Alt trains were 2 s in duration. The second panel depicts the results from the delayed-Std train condition. The Std was a rectangular 2 s train identical in strength to that used in the previous condition; the onset of this train was delayed by six seconds. The Alt was a rectangular 2 s train (undelayed). The third panel depicts the results from the composite-Std condition. The "2 s" portion of the "6w+2 s" Std train was identical to that used in the previous two conditions, and the "6w" portion was set to a lower frequency, identical to the "6w" portion used in the previous version of the experiment (Figure 25). The Alt was a 2 s train. For all conditions, the common logarithms of the frequency of the train triggered by the Alt lever, in pulses per second (pps), is the independent variable, and the common logarithm of the frequency (in pps) of the train delivered by the Std lever (pulse frequency taken from the 2 s portion of the composite train), in the third condition is denoted by a vertical magenta line. The proportion of rewards harvested is the dependent variable. Red data points indicate rewards obtained by pressing the Alt lever, and green data points indicate rewards obtained by pressing the Std lever. The green and red lines are the respective broken-line fits. Error bars denote the standard error of the mean. The data has been averaged over 12 sweeps. The blue circle represents the estimated "cross-over" between responding preferentially on the Std lever to responding preferentially on the Alt lever, with a 95% confidence band.

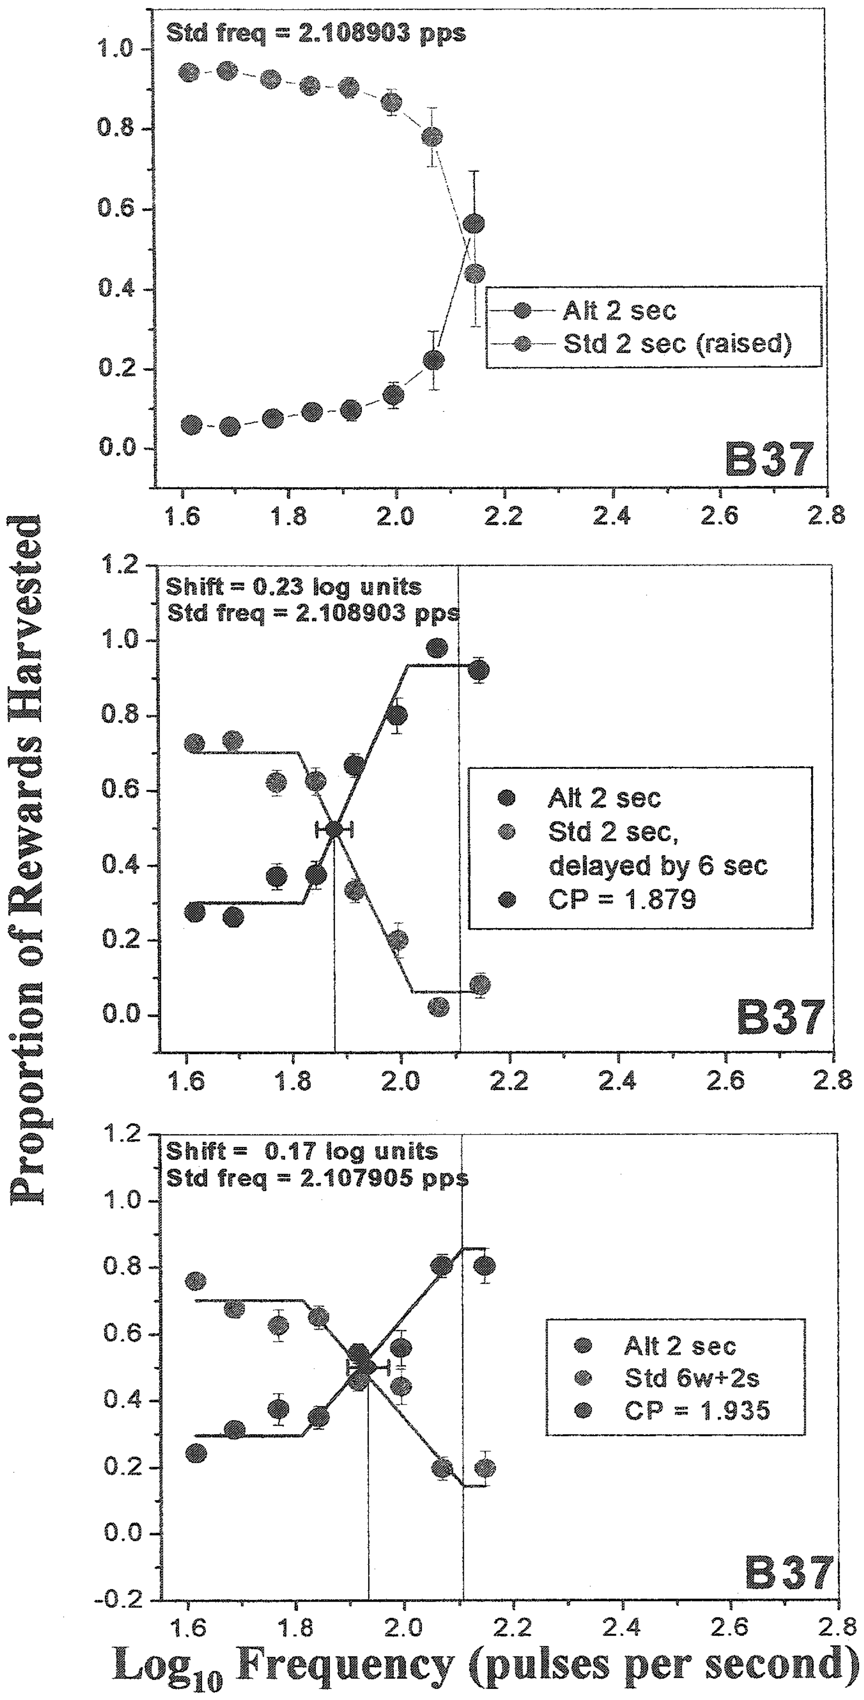

Figure 26. Results of the modified dual-operant beginnings-and-peaks test, for 1 rat.

In the second condition of the modified experiment, the (raised) Std train was again delayed by 6 s, and the Alt lever delivered a 2 s train. The results are graphed in the middle of Figure 26, and show that there was a large statistically significant shift between the CP and the vertical magenta line representing the Std frequency. In this case, the shift was on the order of 0.23 log units, indicating that a 69.8% increase was required for the rat to switch from the Alt lever to the Std lever.

The last condition of the modified experiment pitted a "6w+2 s" composite train against a 2 s Alt; the frequency of the 2 s segment of the composite train was set to the new high value, and the frequency of the 6 s component was the same as in the initial version of this experiment. The results are shown in the bottom panel of Figure 26. Note that the distance between the CP and the Std frequency (vertical magenta line) is almost as large as in the delay discount test in the middle panel. The difference between the two shifts (between conditions) is 0.06 log units, much smaller than the 0.17 log unit difference observed previously. This suggests that once the reward value of the Std train is increased to a maximal level, the effect of adding a weak beginning to such a train is almost nullified.

Discussion

Several findings emerged from the dual-operant beginnings-and-peaks experiments. First, equipreference was shown between a weak Alt train and a much stronger delayed Std train. Second, when the delayed Std was preceded by 6 s of weak stimulation, the strength of the Alt had to be increased in order to achieve equipreference. Lastly (although demonstrated in only one rat thus far), when the strength of the Std train was increased substantially, adding the 6 s of weak stimulation no longer pulled the crosspoint very far to the right.

That delay substantially devalues a train of stimulation is consistent with the literature on delay discounting (Black et al., 1985; Fouriezos & Randall, 1997; Mazur et al., 1987). Also of interest is Reed and Reilly's (1990) demonstration that "filling" the delay period with a signal (for example, a light) resulted in higher rates of responding than when the delay was not bridged by such a signal. Thus, the bridging stimulus appears to lessen the degree of delay discounting. Such a view is consistent with the results in Figure 25, showing that the addition of the 6 s of weak stimulation appeared to largely restore the value of the Std train. However, the failure of the 6 s bridge to mitigate the effects of delay when the 2 s end component was very strong, is not consistent with the bridging hypothesis. Keep in mind that the 6 s component was identical in the two versions of the experiment, and thus should have been equally salient.

Although data from more subjects is required, the preliminary findings in Figure 26 suggest that temporal summation rather than bridging accounts for the effect of the weak beginning. During the 6 s of weak stimulation, the reward signal will likely undergo temporal summation. If so, the reward signal would be above zero at the onset of the 2 s strong component. Thus, the reward signal is likely to climb to a higher peak during the 2 s strong component following 6 s of weak stimulation, than following an "empty" 6 s delay. This hypothesis can be tested, by raising the strength of the 2 s component so that it produces a maximal reward. If the reward is maximal, the preceding baseline is irrelevant. That the 6 s of weak stimulation largely failed to mitigate the effect of delay when the 2 s end was very strong is consistent with this hypothesis.

The temporal summation hypothesis is a simple extension of the notion that the only exemplar important to the evaluation of the train is the peak (Gallistel, 1978). Adding a weaker beginning enhanced the rewarding effect, rather than undermining it, as one might have expected on the basis of peak-and-end averaging. As in the case of the peak-and-end experiments (Experiments 1-3) there was no evidence that multiple exemplars contributed to the evaluation. The peak alone appears to suffice.

Experiment 5

Hsee and colleagues (Hsee et al., 1999) have argued that when a person evaluates various options, different characteristics of these options may become salient depending upon the evaluation context. When only one option is evaluated at a time (single evaluation), characteristics A, B and C may carry the greatest weight, whereas when two or more options are evaluated at one time (joint evaluation), characteristics D, E, and F may be weighted more heavily. Suppose that a consumer is interested in purchasing a new car. If they visit a dealership and only one car of the type they want is available (single-evaluation context), they may place more importance on the price of the car, its mileage, whether it has air conditioning, etc., than on other characteristics of the car, such as its colour. If this consumer visited another dealership with two or more desirable cars (joint-evaluation context), then paint colour or other characteristics may carry more weight. This change in the salience of the various characteristics may even lead the consumer to select a different car in joint-evaluation than in single-evaluation; perhaps the consumer selected Car A once all characteristics were weighed in the single-evaluation context, but would select Car B rather than Car A if both were evaluated simultaneously in a joint-evaluation context. Such a change in preference between evaluation contexts is termed a "preference reversal".

The two paradigms employed in Experiments 1-4, could be thought of as analogous to the single- and joint-evaluation contexts just described. The single-operant paradigm could be thought of as analogous to single-evaluation, and the dual-operant paradigm as joint-evaluation. Sonnenschein, Conover, and Shizgal (2003) had demonstrated duration neglect once trains exceeded 2-4 s, in a single-operant (single-evaluation) paradigm. Given that different characteristics of the stimulation, including duration information, may become relevant if the evaluation contexts were changed, it was of interest to determine whether duration neglect would again be observed once trains exceeded 2-4 s, if trains were presented simultaneously in a joint-evaluation (dual-operant) context.

No apparent differences were seen in terms of peak-and-end/peak-only evaluation (Experiments 1-3) regardless of whether single- and dual-operant paradigms were used. Duration neglect, however, was only briefly assessed in one dual-operant paradigm, the 6 s Alt vs. 8 s Std condition in Experiment 2. However, this constituted a weak test of duration neglect, as the change in duration is relatively small. In the current experiment, a larger range of durations was explored (2 or 4 s vs. 8 s).

It was predicted that if duration information became salient in the dual-operant (joint-evaluation) context, rats would show a preference for the longer of two trains, even though these trains appeared equivalent in a single-evaluation context (Sonnenschein, Conover, & Shizgal, 2003).

Method

Subjects

All 7 rats in the present experiment had served in previous studies. Housing and feeding was identical to that of prior experiments.

Surgery

The stimulating electrodes were implanted as described previously.

Apparatus

The same operant boxes were employed as in the previous experiments.

Procedure

Baseline condition. This condition continued for all rats, and was identical to the daily baseline test described in Experiments 1-4.

Dual-operant preference-reversal test. Many of the characteristics of the present experiment were essentially identical to the dual-operant experiments described earlier. This included procedures for determining the frequencies of Alt and Std trains on each trial, the structure of the trials and ITIs, and the duration of the BDs.

The initial condition of the preference reversal test pitted a 4 s Std train against a 4 s Alt train. Once again, such a condition was employed to ensure that the rats were behaving reliably. For rat B9, 33 trials were run per session (0.05 $\log_{10}$ unit steps; 11 trials per sweep), and for all remaining subjects 27 trials were run per session (0.075 $\log_{10}$ unit steps; 9 trials per sweep).

The last condition, the preference reversal test, involved setting the duration of the Std to 8 s for all rats. For the Alt lever train, the duration used depended on each rat's performance during the single-operant frequency-sweeps experiment (taken either from Sonnenschein, Conover, and Shizgal (2003), or from shortened pilot versions of the same experiment). As the majority of the rats began to display duration neglect in the single-evaluation context once train durations exceeded 4 s (see Figure 4 for rats B9 and B28, and Figure 11 for rats B31, B37, B40 and B42), their Alt train duration was set to 4 s. The remaining subject, B26, had displayed duration neglect once trains exceeded 2 s (see Figure 4), and thus the Alt train was set to 2 s for that subject.

Histology

Upon completion of the study, rats were sacrificed with an overdose of sodium pentobarbital (Somnotol) and perfused through the heart with physiological (0.9%) saline, followed by a 10% formalin solution. The brain was then extracted by removing the ventral surface of the skull and carefully sliding the brain off the electrodes. Brains were stored in 10% formalin for at least one week, and then frozen and sliced in 30 μm sections using a cryostat. Alternate sections were mounted onto glass slides, allowed to dry for a minimum of 2-3 days, and then stained with formol thionine to identify the site of the electrode tips. Slides were examined under a microscope and electrode tip verifications were carried out with reference to the stereotaxic atlas of Paxinos and Watson (1998). The locations of the electrode tips are shown in Figure 27 (data for one rat, K5, is missing due to loss of the critical sections).

Figure 27. The locations of the electrode tips for all subjects. The coordinate of the coronal section (distance from bregma) is indicated at the lower right. MFB refers to the medial forebrain bundle, and LH refers to the lateral hypothalamus. Electrode placements for individual rats are indicated by coloured stars; the legend appears in the top right.

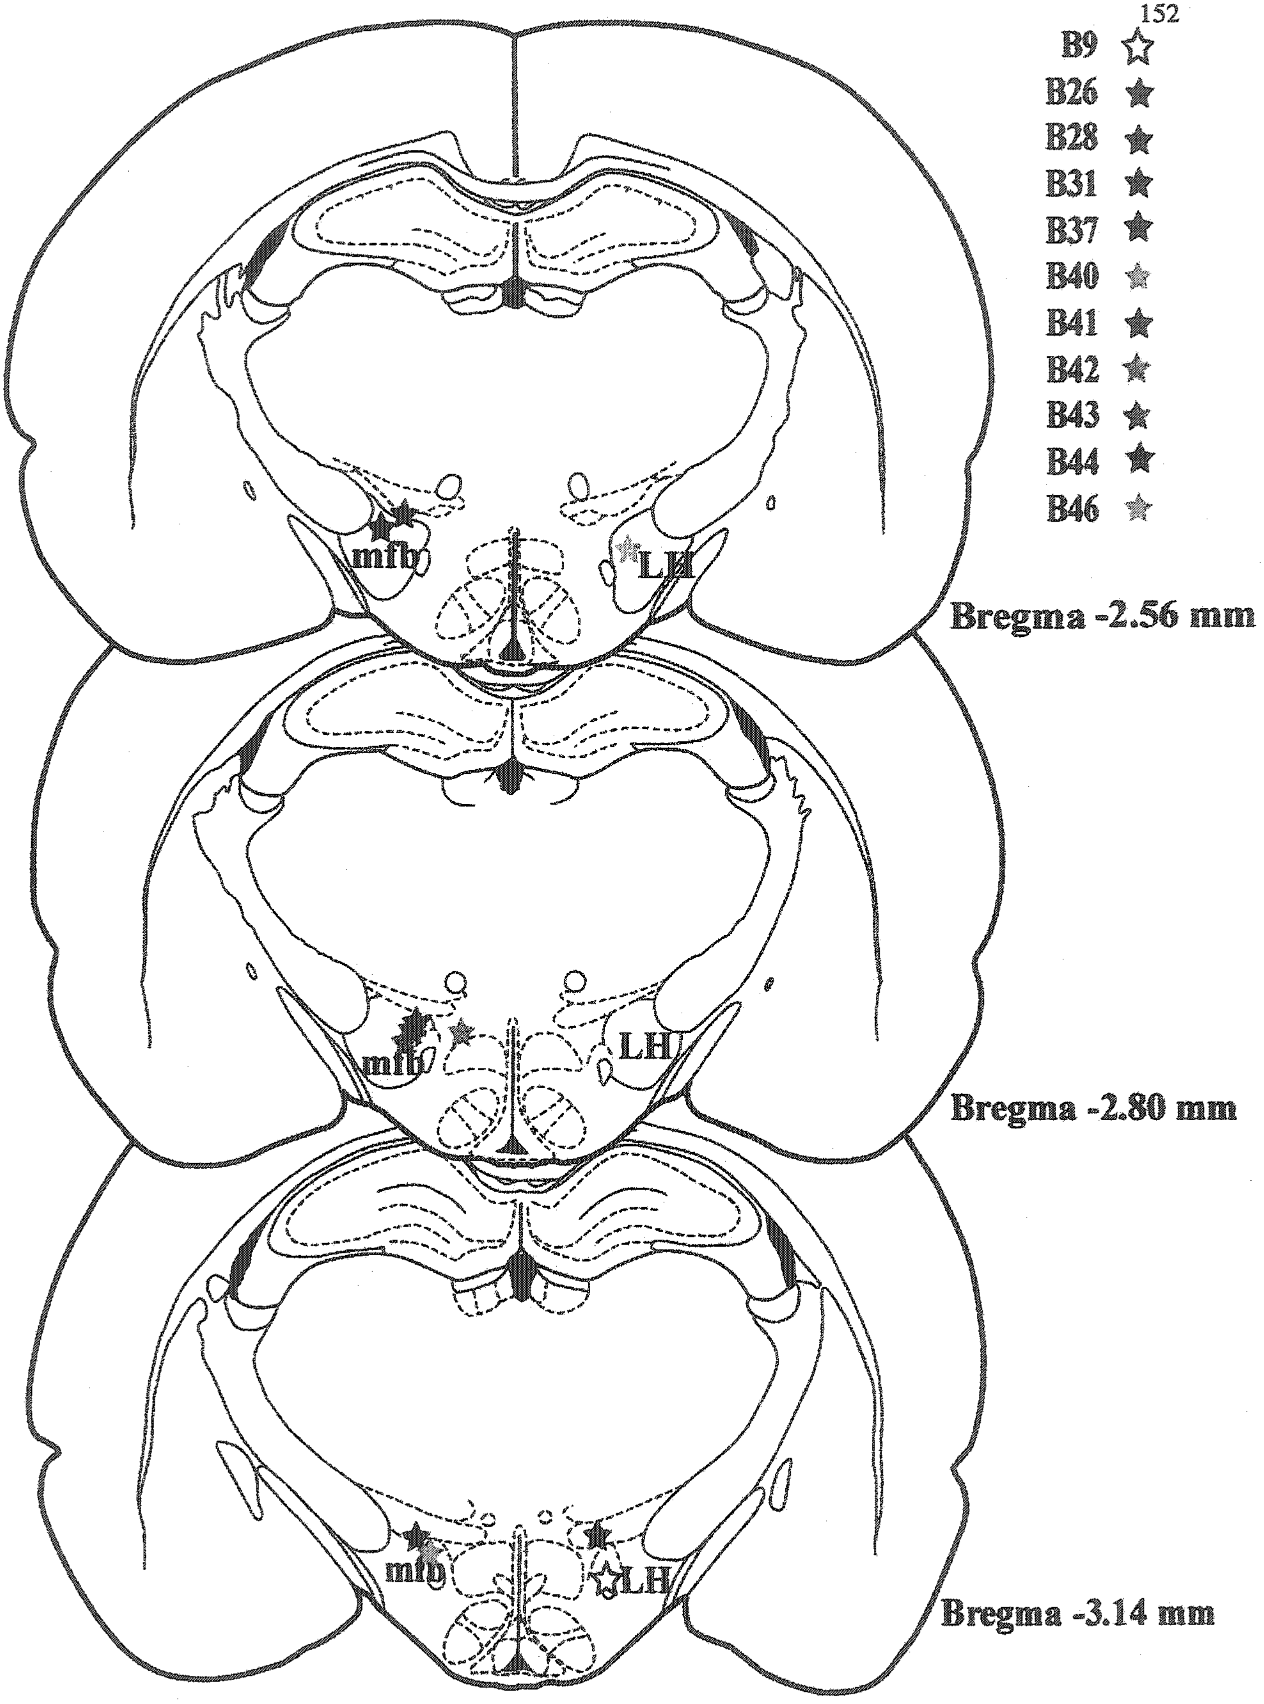

Figure 27. The locations of the electrode tips.

Results

As was the case in the previous dual-operant experiments, the initial warm-up trial in each sweep was not included in the statistical analyses. All analyses were conducted on means that had been averaged over 12 such sweeps.

The cross-point (CP) and where it falls in relation to the frequency of the Std train was once again the key feature of interest. CPs were calculated in the same manner as described in previous dual-operant experiments. Briefly, broken-line functions were fit to each Alt and Std response curve for each of the 12 sweeps. Only those sweeps where a well-behaved CP was obtained (i.e., the two diagonal segments intersected) were retained for analysis. The x- and y-values of the intersections arising from each of the 12 broken-line-fits were averaged to yield a mean x-axis value for the CP and a mean y-axis value. From these values, the mean CP for the full set of 12 sweeps was obtained, and a 95% confidence band was calculated around the mean x-value.

In the first condition of the preference reversal study, both the Std and Alt trains were set to the same duration. The results are shown in Figure 28, where the x-axis represents the common logarithm of the frequency of the Alt train, and the y-axis is the proportion of rewards harvested (PRH) by the subjects. The vertical magenta line represents the frequency of the Std train. The rule of thumb employed to assess statistical reliability was overlap of the confidence interval with the magenta line. According to this criterion, there are no statistically significant differences between the estimated CP and the magenta line, for all 6 subjects tested.

Figure 28. Results from the 4 s versus 4 s condition of the dual-operant preference-reversal test, for 6 subjects (individual subjects are identified by the label in the lower right of each graph). The common logarithm of the pulse frequency of the Alt train in pulses per second (pps) is plotted on the x-axis, with the vertical magenta line denoting the common logarithm of the frequency (pps) of the Std train. The dependent variable is the proportion of rewards harvested (out of a maximum of 20). Red data points represent rewards obtained by pressing the Alt lever, and green data points rewards obtained by pressing the Std lever, with error bars representing the standard error of the mean. The green and red lines are the respective broken-line fits. All means are drawn from data that has been averaged over 12 individual sweeps. The blue point represents the estimated "cross-over" between responding on the Std lever to responding on the Alt lever, with a 95% confidence band. Both the Alt and Std were rectangular 4 s trains.

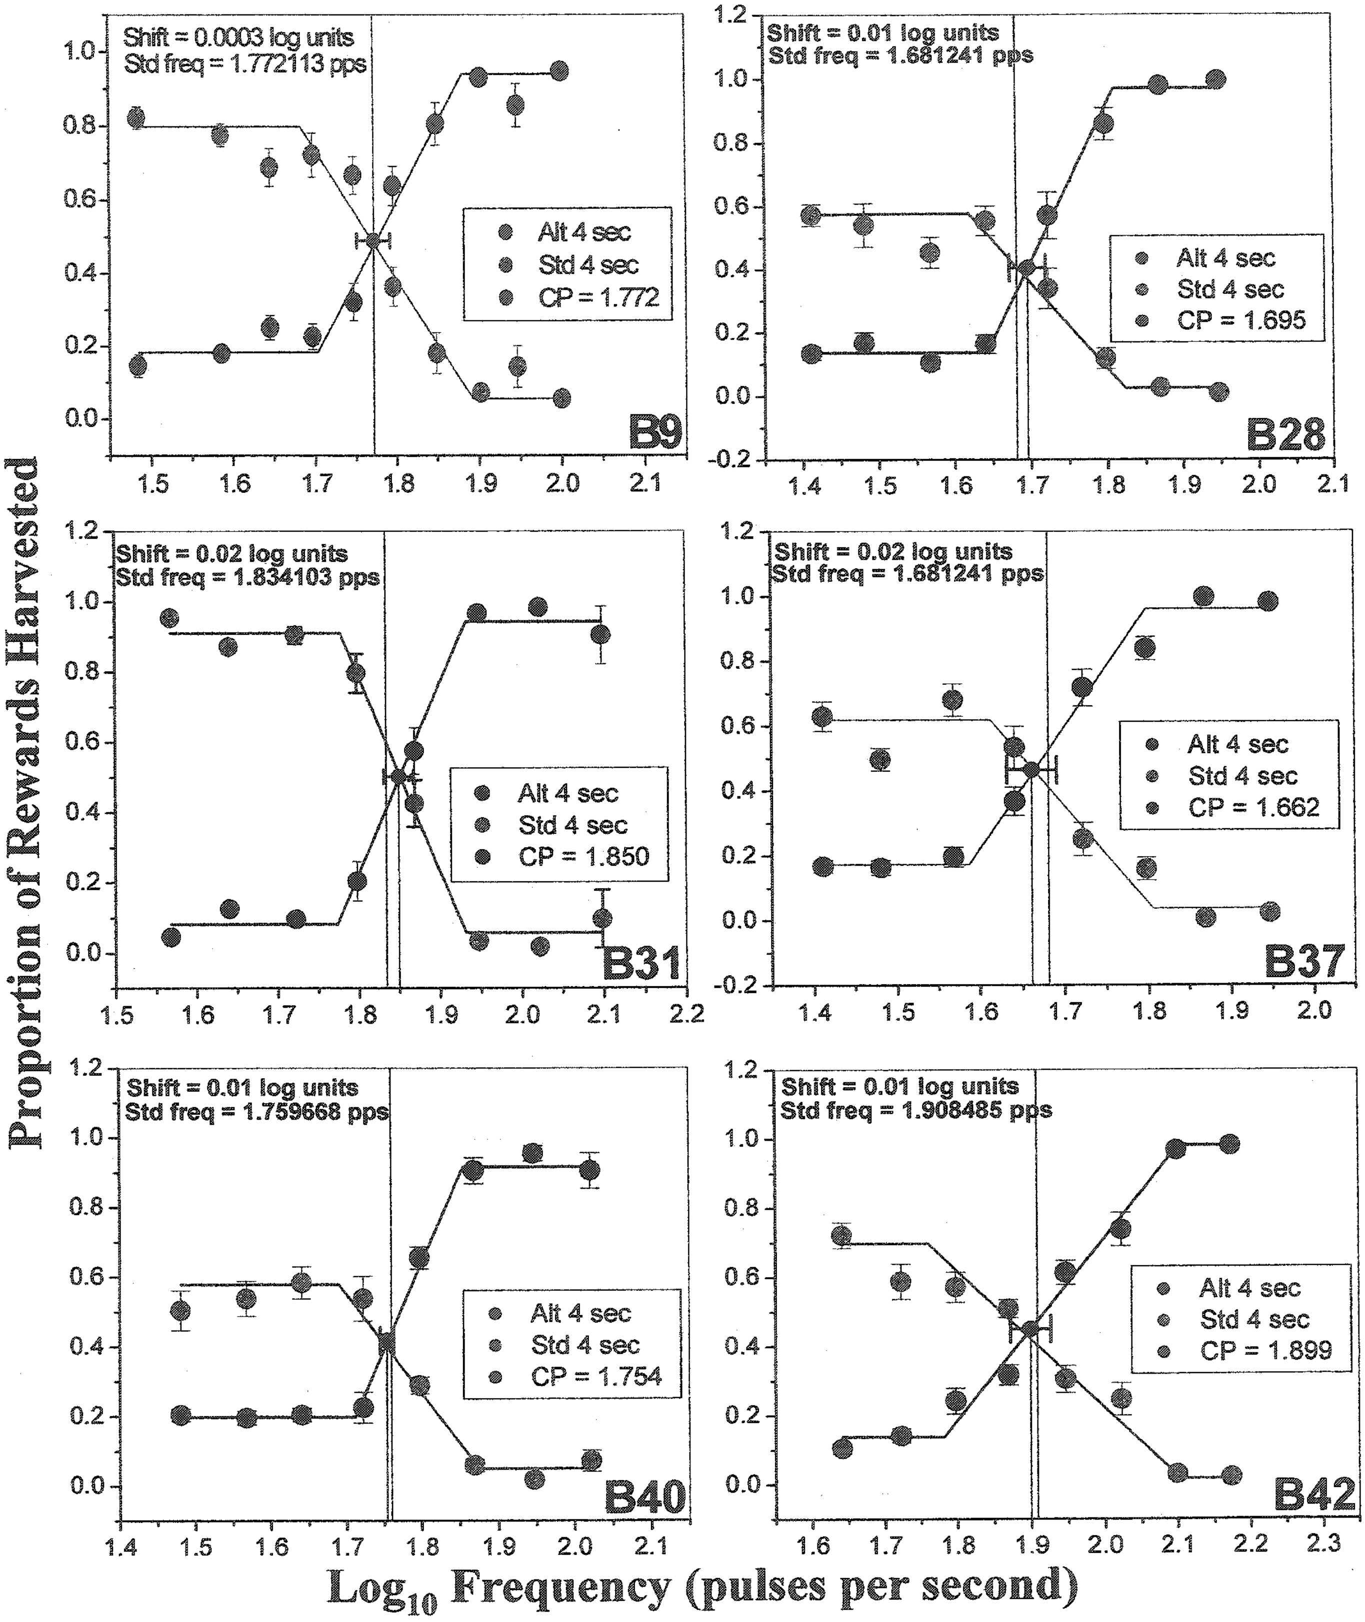

Figure 28. Results of the 4 s Std train vs. 4 s Alt train condition of the dual-operant preference reversal experiment.

The second condition of the current experiment was the preference-reversal test. All parameters were the same as in the condition just described, except that the Std train was now set to 8 s, and the Alt lever train to 4 s (for 6 out of the 7 rats). The independent and dependent axes are identical to those described earlier. As can be seen in Figure 29, all of the subjects demonstrated a statistically significant shift between the CP and the frequency of the Std train, as indicated by the vertical magenta line falling outside the 95% confidence interval around the CP. The size of this shift varies among the rats, with four of them (B28, B37, B40 and B42) showing a moderately large 0.04 log unit shift (a 9.6% increase in the frequency of the Alt train was required to motivate the rats to "cross over" to that lever). In the remaining rats (B9, B31) shifts of 0.05 to 0.10 log units were observed (a 12-26% increase in the Alt train was required). In all cases, the crosspoint is shifted rightward, relative to the vertical magenta line, indicating that the 8 s train was preferred to the 4 s train.

This pattern also held for B26, who was tested with different train durations. This particular subject was tested with 2 s Alt and Std trains in the first condition, and an 8 s Std train pitted against a 2 s Alt train in the preference reversal test. The results are shown in Figure 30, and are comparable to those of subjects tested with 4 and 8 s trains. In the first condition, shown in the top panel of Figure 30, there is no statistically significant difference between the CP and the frequency of the Std train. In the preference reversal condition, shown in the bottom panel of the figure, there was again a statistically significant difference between the CP and the Std lever train frequency. The shift was 0.05 log units, indicating that a 12% increase in the frequency of the Alt lever train was needed for this rat to switch responding from the Std

Figure 29. Results from the dual-operant preference-reversal test, for 6 subjects (individual subjects are identified by the label in the lower right of each graph). A 4 s Alt train and an 8 s Std train were used. The independent variable is the frequency of the Alt train in pulses per second (pps), expressed in logarithmic units. The vertical magenta line represents the common logarithm of the frequency (in pps) of the Std train. The dependent variable is the proportion of rewards harvested, out of a maximum value of 20. Red circles represent rewards obtained by pressing the Alt lever, and green circles rewards obtained by pressing the Std lever. The green and red lines are the respective broken-line fits. Error bars represent the standard error of the mean, and all means are drawn from data that has been averaged over 12 separate sweeps. The blue circle denotes the estimated "cross-over" between responding preferentially on the Std lever to responding preferentially on the Alt lever, with a 95% confidence band.

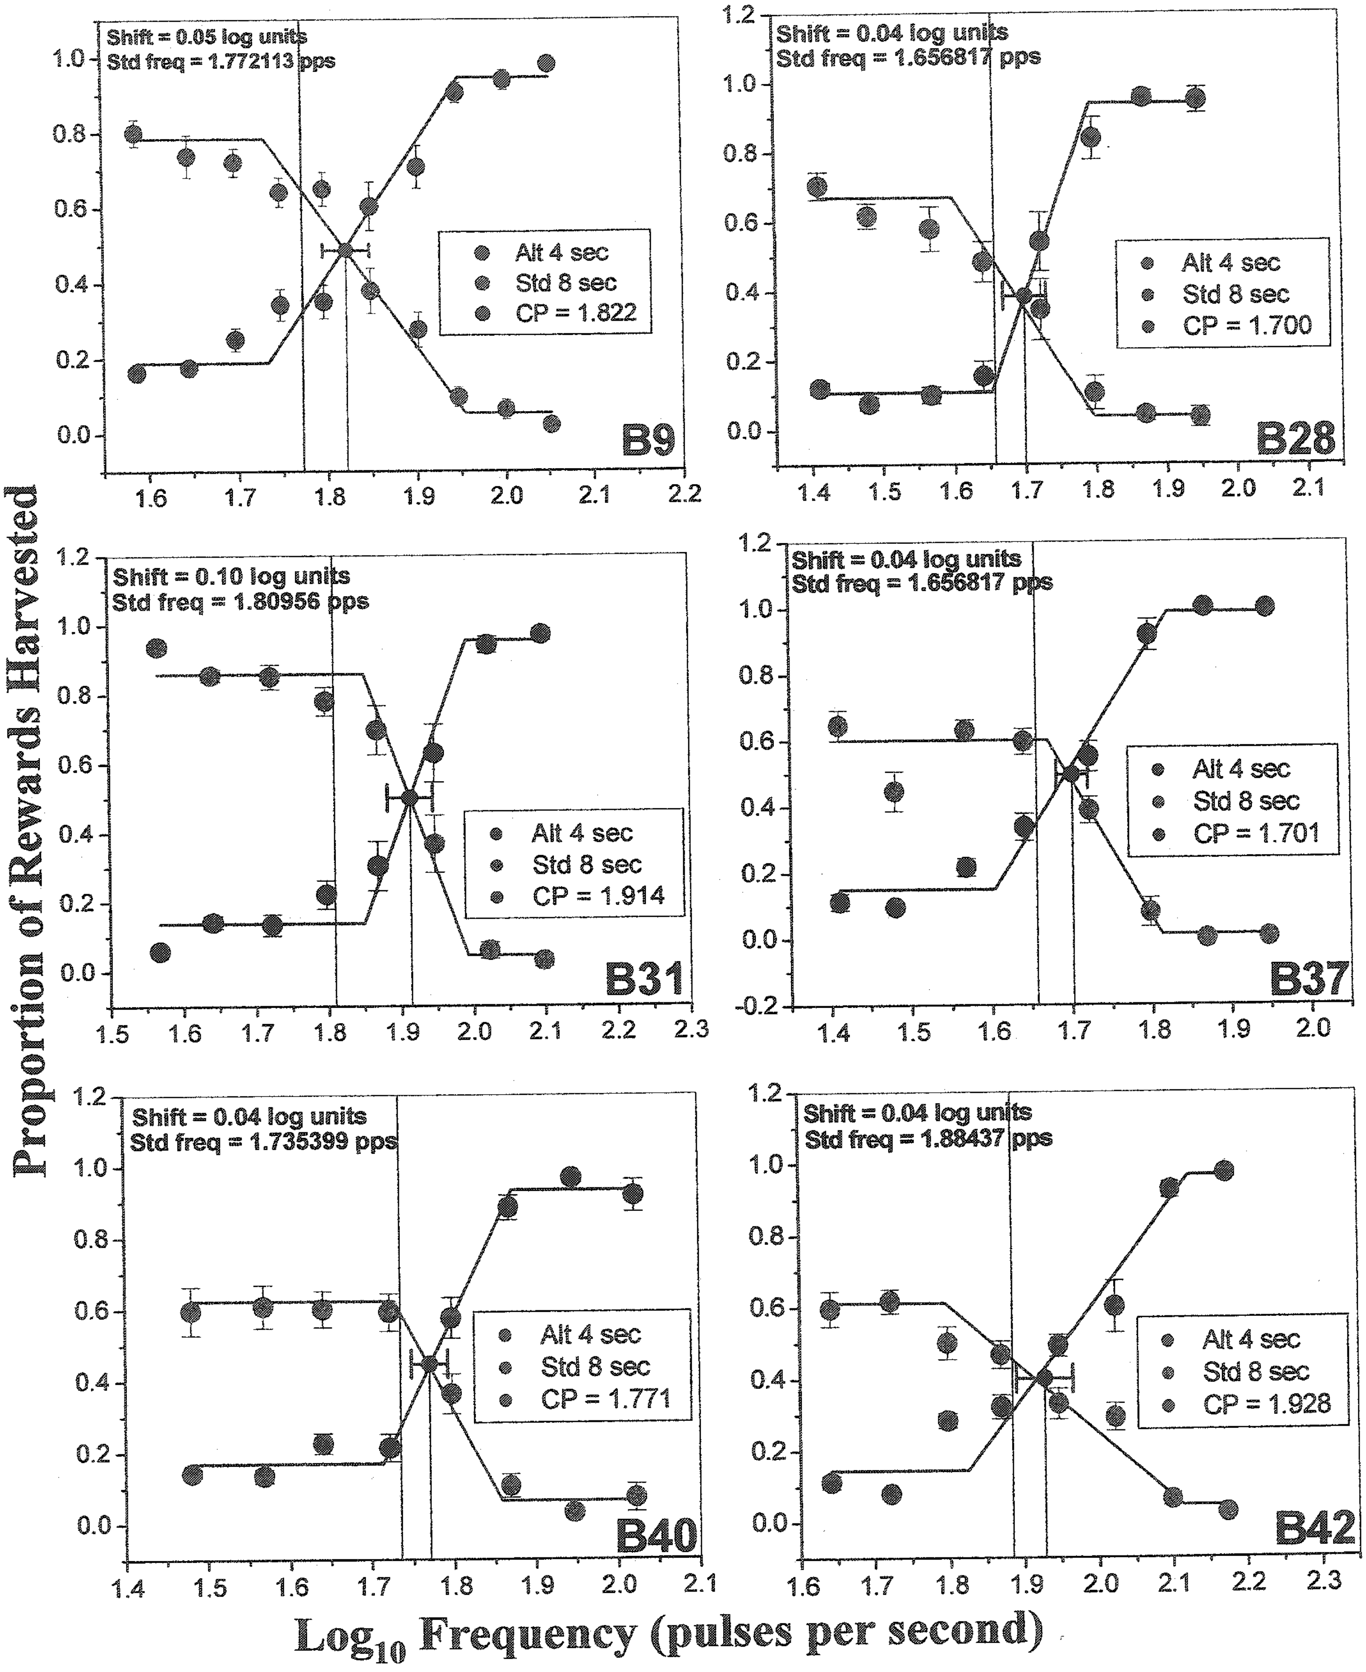

Figure 29. Results from the dual-operant preference reversal test, for subjects tested with 4 s Alt trains.

Figure 30. Results from both conditions of the dual operant preference reversal test, for 1 subject tested with a 2 s Alt train and an 8 s Std train. The independent variable is the frequency of the Alt train in pulses per second (pps), expressed in logarithmic units. The vertical magenta line represents the pulse frequency (expressed in logarithmic units) of the Std train. The dependent variable is the proportion of rewards harvested, out of a maximum value of 20. Red circles represent rewards obtained by pressing on the Alt lever, and green circles rewards obtained by pressing on the Std lever. The green and red lines are the respective broken-line fits. Error bars represent the standard error of the mean, and all means are drawn from data that has been averaged over 12 sweeps. The blue circle denotes the estimated "cross-over" between responding preferentially on the Std lever to responding preferentially on the Alt lever, with a 95% confidence band.

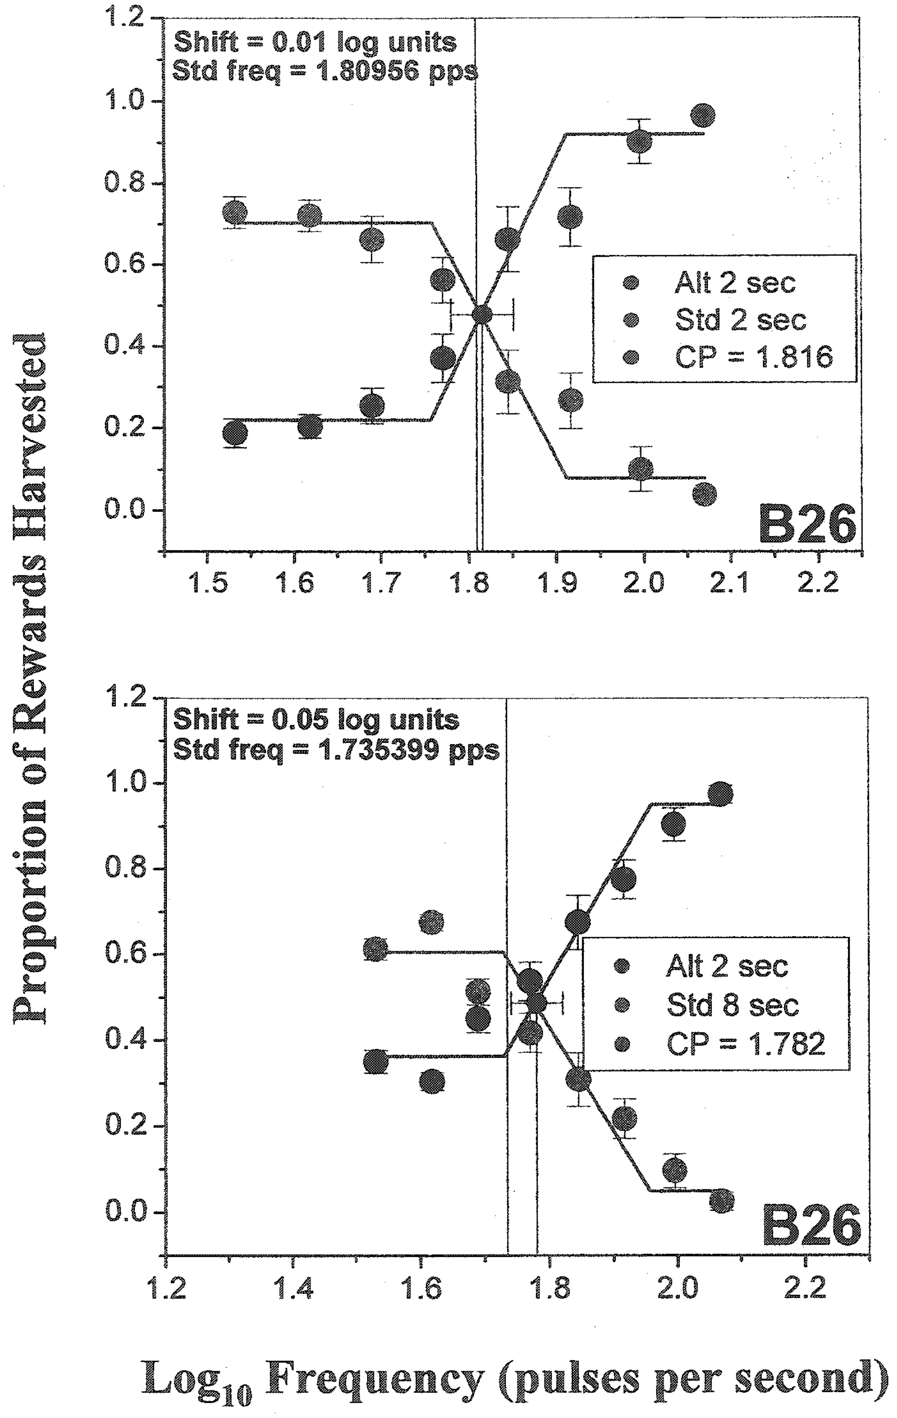

Figure 30. Results from the dual-operant preference reversal test, for 1 subject tested with a 2 s Alt train.

lever to the Alt lever. Again, the crosspoint was shifted rightward relative to the frequency of the Std train frequency, indicating a preference for the 8 s train.

There is a problem, however, with interpreting the results of the preference reversal experiment in the manner just described. From examining the frequency sweeps data (Figure 4 for rats B9, B26, and B28, and Figure 11 for rats B31, B37, B40, and B42), it is evident that for several of the subjects, the 4 s and 8 s curves did not always overlap completely. In other words, duration neglect was approached, but not always achieved. Therefore, the most rigorous and appropriate baseline for measuring the preference reversal effect is not the frequency of the 8 s Std train, but rather a value that includes any small shifts displayed between the 8 s and 4 s curves in the frequency sweeps experiment.

These small shifts were accommodated by calculating the frequency required for the rat to harvest half of the available rewards (the mid-point of the psychometric function). This value will be called the "required frequency". This calculation was performed for each of the six broken-line fits to the 4 s (2 s in the case of B26) and 8 s curves. The required frequency estimates for the curves obtained at the 4 s duration were averaged, as were the required frequency estimates for the curves obtained at the 8 s duration. For each rat the difference between these estimates was calculated and added to the frequency of the Std train. The sum of the frequency of the Std train and the shift is represented in Figures 31 and 32 as a black dot-dashed line. Moving the vertical baseline to the right takes into account any demonstrated advantage of the 8 s train over the 4 s train under single-operant conditions.

Figure 31. Results from the preference-reversal test, with the results from the prior frequency sweeps experiments taken into account, for 6 subjects (individual subjects are identified by the label in the lower right of each graph). A 4 s Alt train was pitted against an 8 s Std train. The independent variable is the frequency of the Alt train in pulses per second (pps), expressed in logarithmic units. The vertical magenta line represents the common logarithm of the pulse frequency (in pps) of the Std train. The vertical black dot-dashed line represents the correction made to the vertical magenta line, after the distance between the 4 s and 8 s performance curves in the frequency sweeps experiments is taken into account. The dependent variable is the proportion of rewards harvested, out of a maximum value of 20. Red circles represent rewards obtained by pressing the Alt lever, and green circles rewards obtained by pressing the Std lever. The green and red lines are the respective broken-line fits. Error bars represent the standard error of the mean, and all means are drawn from data that has been averaged over 12 separate sweeps. The blue circle denotes the estimated "cross-over" between responding preferentially on the Std lever to responding preferentially on the Alt lever, with a 95% confidence band.

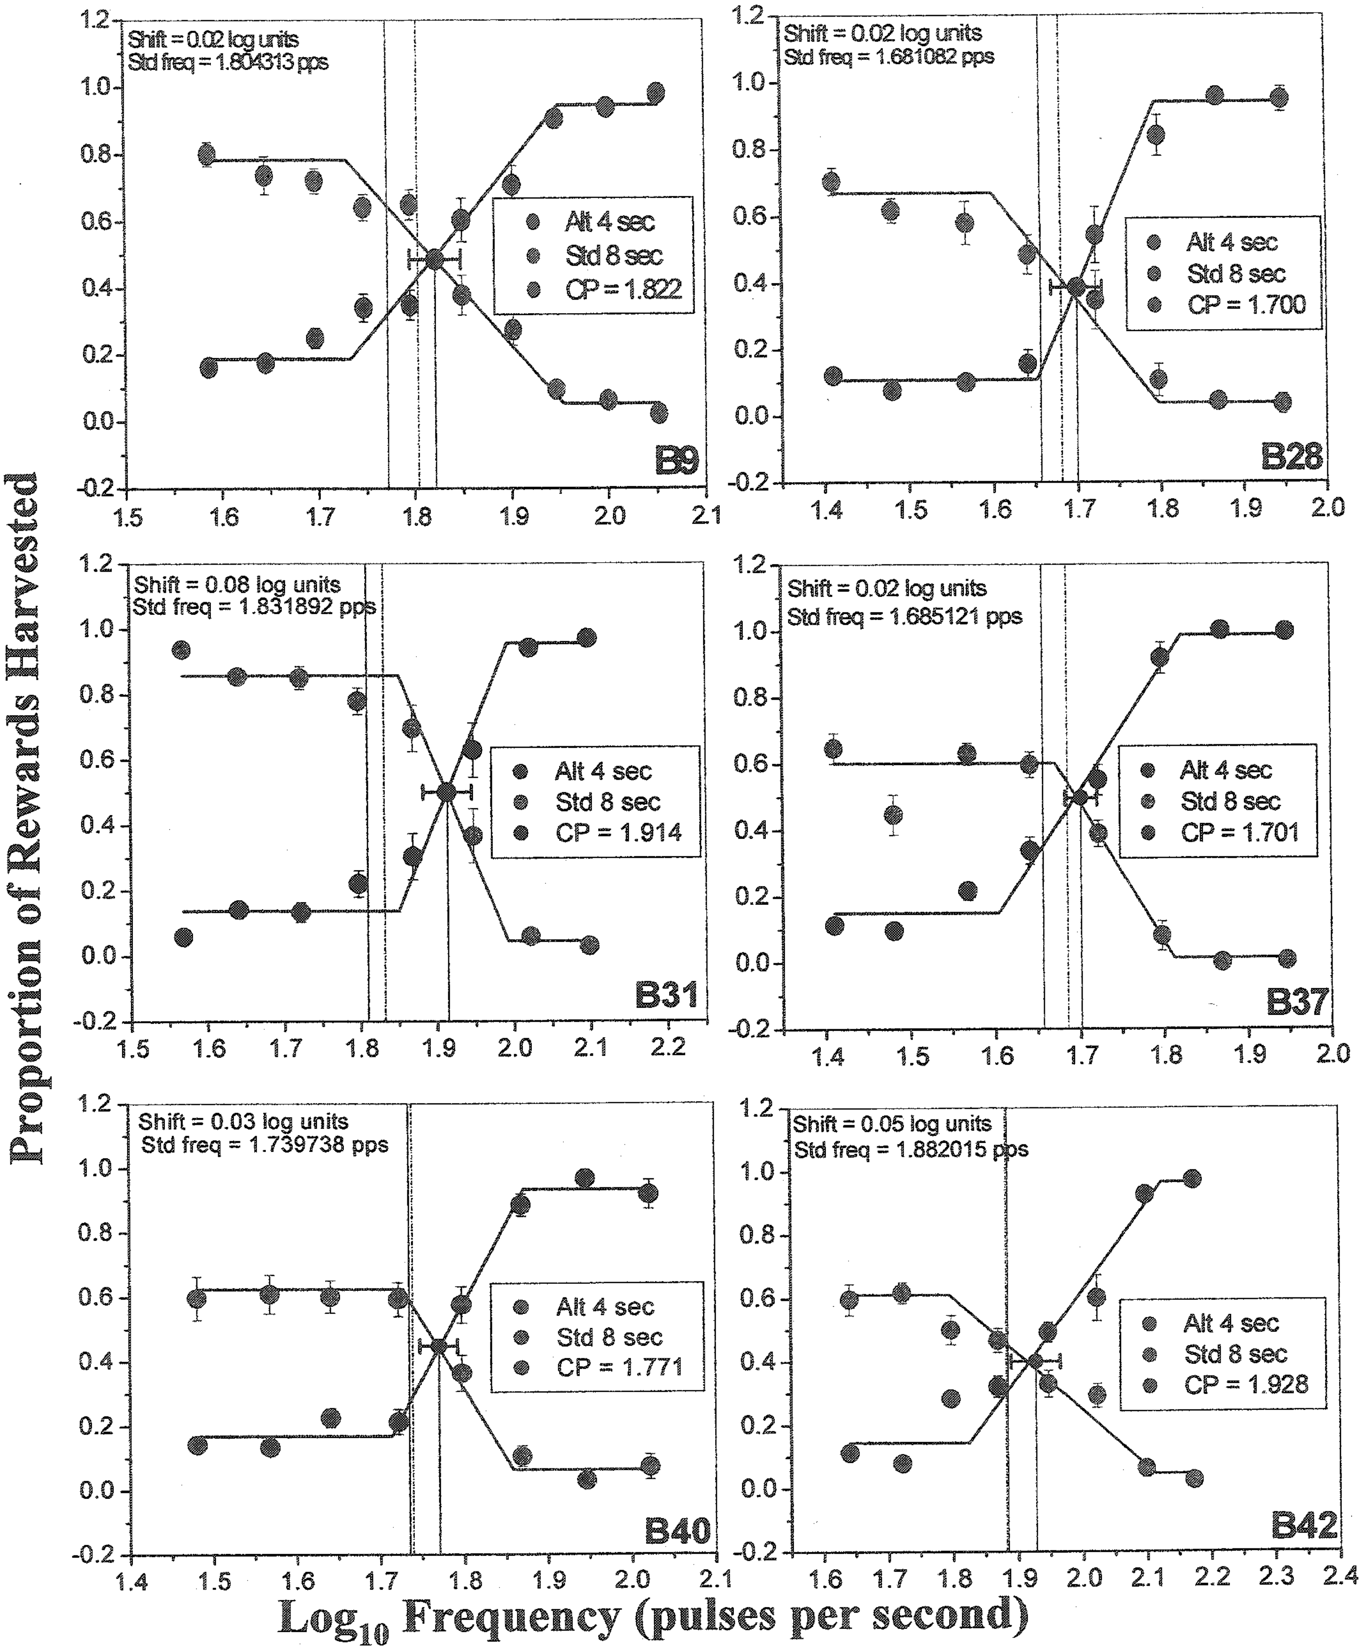

Figure 31. Results from the dual-operant preference reversal test, plotted with the adjusted Std value (vertical dot-dashed line), for subjects tested with 4 s Alt trains.

Figure 32. Results from both conditions of the preference-reversal test, with the results from the prior frequency sweeps experiments taken into account, for 1 subject. A 2 s Alt train was pitted against an 8 s Std train. The independent variable is the frequency of the Alt train in pulses per second (pps), expressed in logarithmic units. The vertical magenta line represents the pulse frequency, in logarithmic units, of the Std train. The vertical black dot-dashed line represents the correction made to the vertical magenta line, after the distance between the rat's 2 s and 8 s performance curves in the frequency sweeps experiment is taken into account. The dependent variable is the proportion of rewards harvested, out of a maximum value of 20. Red circles represent rewards obtained by pressing on the Alt lever, and green circles rewards obtained by pressing on the Std lever. The green and red lines are the respective broken-line fits. Error bars represent the standard error of the mean, and all means are drawn from data that has been averaged over 12 sweeps. The blue circle denotes the estimated "cross-over" between responding preferentially on the Std lever to responding preferentially on the Alt lever, with a 95% confidence band.

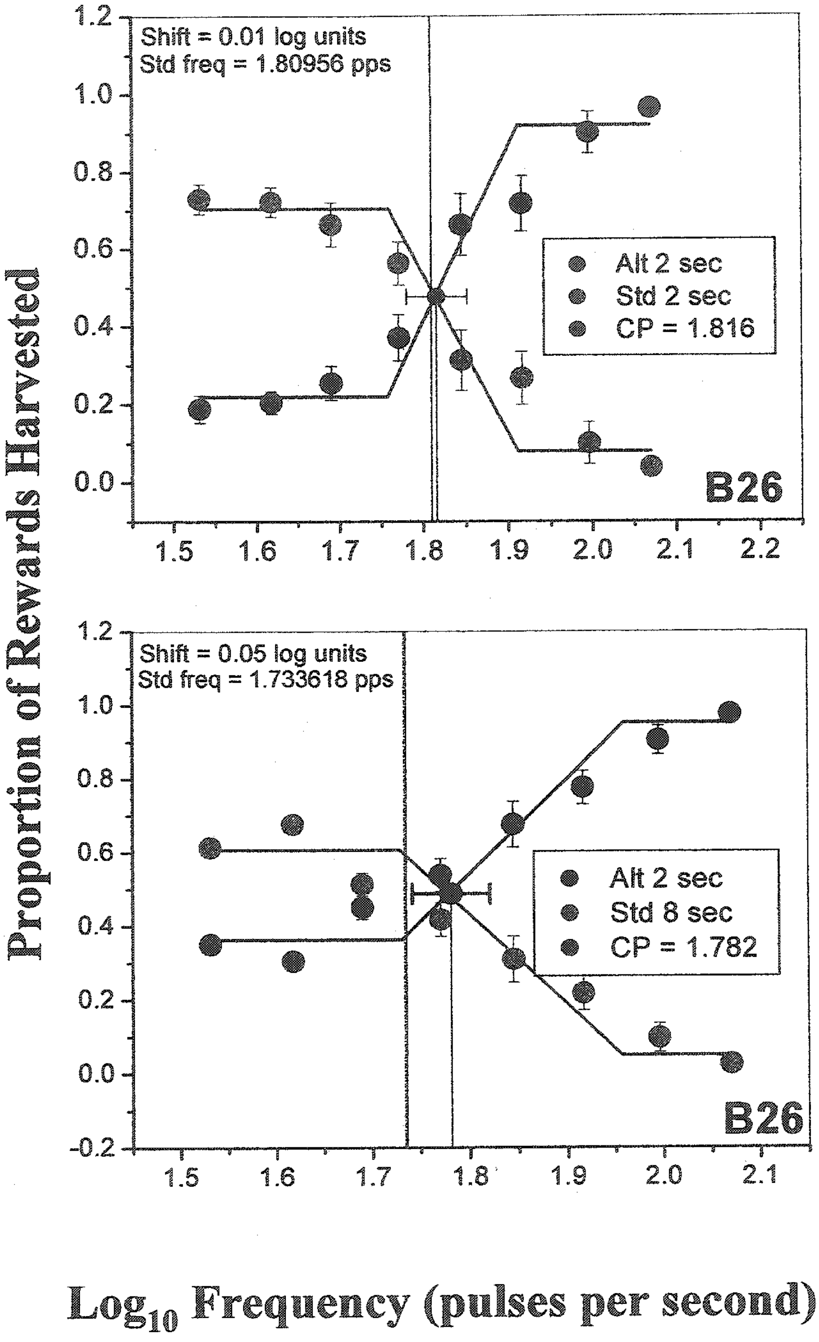

Figure 32. Results from the dual-operant preference reversal test, plotted with the adjusted Std value (vertical dot-dashed line), for 1 subject tested with a 2 s Alt train.

The results of this analysis are shown in Figures 31 and 32. All axes have the same meaning as in prior graphs. In this case, both the original magenta line and the new comparison line (plotted as a black dot-dashed line) are shown. As can be seen from the figures, once the new line is used as the means of comparison, the preference reversal effect no longer meets the criterion for significance for rats B9, B28, or B37, with the black dot-dashed line falling within the 95% confidence band around the CP. For all three subjects, the new log unit shift is 0.02, (a mere 4.7% increase in the Alt train frequency was required to motivate the rats to switch to that lever). However, for the other four subjects, the preference reversal effect remains, as the black dot-dashed line and the confidence band do not overlap. The shifts for these rats were in the range of 0.03 to 0.08 log units (a 7.2 to 20% increase in the Alt lever frequency required to motivate the rats to press on that lever). Regardless of whether a significant preference reversal effect was observed, all results were still in the same direction, with the CP always rightward-shifted regardless of whether the magenta line or black line was the comparison measure.

Discussion

In 4 out of the 7 rats in the present experiment, the 8 s train was more effective than 4 s (and 2 s, in 1 subject) train, once these two trains were presented simultaneously in a joint-evaluation (dual-operant) paradigm. The results for those 4 subjects are in contrast to their performance in the single-evaluation (single-operant) context, where the effectiveness of the two trains was very similar or indistinguishable. Thus, a preference reversal was displayed by these particular rats, consistent with the ideas of Hsee et al. (1999). Although the remaining 3 rats did not display a significant preference in favour of the 8 s train, their results were still in the same direction (CP rightward-shifted relative to the comparison line). Taken together, the results from all 7 rats suggest that at least in some cases, information about the duration of the experience, while not salient in the single-operant paradigm, became more salient in the dual-operant paradigm. This is not surprising, as individual differences could be expected to play a role in the evaluation of choices, over and above any contributions to the decision process brought about by the evaluation context.

The results of the present study therefore suggest that (at least in some subjects) the expression of duration neglect, rather than being "absolute", may differ depending on the context. Ariely and Loewenstein (2000) have argued that the duration neglect seen in the Kahneman studies (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996) may not be as absolute as the latter group has argued. The work of Ariely and Loewenstein (2000) has suggested that the way subjects are asked to make their retrospective evaluation will influence whether duration information will be taken into account or not. For example, they argue that asking subjects to make an *overall* evaluation of the experience, as was typically done in the Kahneman studies, leads subjects to neglect duration information. In contrast, they suggest that asking subjects to make a different kind of rating, such as comparing one experience to a "standard" experience (analogous to a dual-operant condition) leads to a retrospective evaluation which does include information about the duration of the experience (Ariely & Loewenstein, 2000). Perhaps in the context of the present studies, running rats in single- and dual-operant paradigms constitutes "different ways of asking the question", which could then be expected to influence how duration information is used in evaluation and decision-making.

It should be noted that duration could influence the reward signal in at least two ways. Prior to the onset of duration neglect, increasing the train duration allows more time for the reward signal to build up. Thus, increasing the duration increases the peak value of the signal. Even if the peak were the only exemplar employed in the evaluation of the reward signal, duration would contribute indirectly via its effect on the peak. Pushing the train duration into the range of duration neglect makes it possible to test whether duration contributes to the evaluation in a manner independent of the peak. Once duration neglect has been attained, the reward signal has approached asymptote and hence, its peak value is no longer changing. If further increases in duration alter the evaluation of the reward signal, then these can be attributed to the effect of duration per se.

The psychometric functions obtained at the longest durations in the abbreviated frequency sweeps test (Figures 11 and 12) and in the experiment of Sonnenschein, Conover, and Shizgal (Figure 4) are virtually superimposable in many cases (B15, B26, B28, B40, B41, B42, B46, K5); in most of the remaining cases, duration neglect is approached at the longest durations tested. Duration neglect has therefore been achieved or approached in the single-operant tests. The current experiment investigated whether duration became more salient in a dual-operant paradigm and contributed to the evaluation of the reward signal, even in cases where the duration was increased beyond the values that yielded neglect in the single-operant paradigm. As noted above, in 4 out of 7 cases (Figures 31 and 32), evidence that duration per se contributes to the evaluation of the reward signals was indeed obtained in the dual-operant paradigm.

General Discussion

The main aim of the present study was to elucidate the cognitive heuristics that rats may use when evaluating positive experiences.

Experiment 1 presented self-stimulating rats the opportunity to bar-press for constant-frequency (8 s and 6 s) and composite-frequency trains (6+2), presented in a single-operant paradigm. This experiment tested the peak-and-end hypothesis in the single-operant paradigm. No systematic evidence for peak-and-end averaging was obtained. Experiment 2 was aimed at determining whether the lack of a peak-and-end effect was paradigm-specific, i.e., would a peak-and-end effect be unmasked if constant- and composite-frequency trains were presented simultaneously in a dual-operant context. Experiment 3 was aimed at determining whether extending the weak ending of the composite train would lead to the end exemplar becoming more salient than had been the case in prior experiments. The goal of Experiment 4 was to assess whether the "beginning" may serve as an exemplar in addition to the "peak". Lastly, Experiment 5 was aimed at determining whether information about the duration of the stimulation would become more salient if the trains were presented concurrently (joint-evaluation), than had been the case when the same trains had been presented singly (single evaluation).

The following topics will be discussed below: the results of the peak-and-end tests and their implications, the results of the beginnings-and-peaks test, and the findings of the preference reversal test and how duration neglect may be affected by the evaluation context. Finally, the main conclusions arising from the present study will be discussed.

Peak-and-End Tests

In Experiments 1-3, the peak-alone model of Gallistel (1978) was consistently supported by the data, rather than the peak-and-end model of Kahneman (Kahneman et al., 1993). Regardless of whether the composite- and constant-frequency trains were presented in either a single-operant paradigm, a dual-operant paradigm, or with longer weak "endings" (2+6 s trains), the composite trains were not generally found to be devalued relative to the shorter constant-frequency trains. This indicates that the reward signal of the "end" segment of the train did not act as an exemplar in the evaluation of the reward signal. The peak-only model of Gallistel (1978) was sufficient to account for the rats' behaviour in Experiments 1-3.

As alluded to earlier on, there are many differences between the Kahneman group's experiments (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996) and the present study which may have accounted for the lack of a peak-and-end effect.

The first obvious difference is that the Kahneman studies employed human subjects, whereas the present study employed non-human animals. While it is true that rats and humans differ substantially in terms of their cortical structures, most neuroscience research involving nonhuman animals assumes that the *subcortical* structures of humans and other mammals, are more similar. Given that the MFB reward system appears to be largely subcortical, its structure may be very similar in rats and humans. Another possible species difference may be that rats react differently than humans to positive and negative experiences. Given that the clearest results with respect to peak-and-end evaluation have occurred in human subjects undergoing aversive experiences, rats could be exposed to negative stimuli (e.g., exposure to an inescapable bright light, to a heat lamp, etc.) to determine whether they too display a peak-and-end effect in such situations.

The possible difference between evaluating positive and negative experiences is also relevant given that the second difference between the Kahneman studies and the present study was the type of stimulus employed. Kahneman and colleagues have focused mostly on aversive experiences, whereas the present study used only positive stimuli. One might imagine that the importance of different exemplars may vary depending on the experience. Frederickson and Kahneman (1993) conducted a study in which humans evaluated six positively-valenced film clips and six negatively-valenced film clips. The effect of duration was small and significant only in the case of the aversive films, and evaluations of the aversive film clips were more consistent than those of the positive films. However, it is not clear whether these results reflect fundamental differences in how positive and negative experiences are evaluated, or are simply due to the specific characteristics of the six film clips in each category. Even if future work in human subjects reveals that positive and negative stimuli are evaluated differently, it would still be important to determine whether this is so in laboratory animals.

There is yet another form of stimulus difference between the present study and the Kahneman studies, however, which could have influenced the results. In the Kahneman studies, all the stimuli were meaningful in the sensory and perceptual domains. The subjects knew what the stimuli were, as well as how they felt about them. In the present study, in contrast, the use of BSR meant that the rats were receiving a stimulus thought to lack many "perceptual" stimulus characteristics (Shizgal, 1999). The stimulation is unlikely to produce an interpretable percept, such as a smell, sound, or visual image (i.e., the rat does not hallucinate a large piece of cheese or other natural goal object while receiving the stimulation). Perhaps this impoverished sensory nature of the BSR played a role in the lack of a peak-and-end effect. One possible way to get around this problem would be to assess whether the "peak-and-end" model will hold when rats are responding for more naturalistic types of appetitive experiences. For example, one could employ a variant of the intraoral/intragastric cannulae experiments described in the introduction (Conover & Shizgal, 1994a, 1994b; Conover et al., 1994; Grill, Spector, Schwartz, Kaplan, & Flynn, 1987). If the concentration of the sucrose solution could be changed during delivery of the sucrose to the rat, in order to create a "weaker" sweet taste at the end, this could represent a situation analogous to the types of BSR trains employed in the current study. The sucrose would, however, have more meaningful stimulus characteristics (such as taste) than trains of electrical brain stimulation.

A third difference between the two sets of studies is that in the Kahneman studies, the human subjects were generally the passive recipients of the negative stimuli, whereas in the present studies the rats were required to make an active response (bar-pressing) to obtain the rewarding stimulus. Perhaps some exemplars will be weighted differently or ignored altogether if some effort from the organism is required to contact the stimulus. The peak-and-end model might hold when an organism is in a position where it is "forced" to receive a stimulus, whereas the simpler peak-alone model holds when the organism must make a decision whether to approach and perform in order to obtain a stimulus.

Future studies could determine whether the active/passive distinction is an important one, by turning rats into passive recipients of BSR, and then observing whether this leads to a peak-and-end effect. For example, one could use an open-field conditioned place-preference (CPP) paradigm (Vezina & Stewart, 1987a, 1987b), in which tactile cues such as different flooring types (grid versus bar) are used. Trains could be delivered noncontingently to the rats on different floorings. If on the test day the rats spend significantly more time on floorings that had been previously paired with constant-frequency trains, as opposed to floorings that had been paired with composite-frequency trains, this would support a peak-and-end effect. Still another possible study, this time with human subjects, could address this issue by enabling human subjects to be more active "seekers" of their experiences.

Lastly, a fourth difference between the two studies was that the Kahneman studies generally made use of longer stimuli than the present study. The majority of the studies by Kahneman and co-workers used stimuli ranging from 21 s to an hour or more (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996). In one experiment where a shorter stimulus (8 s), comparable to the present study, was employed, the authors found a peak-and-end effect, but also effects of duration and other variables (Schreiber & Kahneman, 2000). Given that the present study generally used stimuli much shorter than those used in the majority of the Kahneman studies, perhaps there is a difference between the way the neural substrate evaluates "shorter" temporally-extended experiences as opposed to longer temporally-extended experiences. The rats might behave much differently if the stimulation lasted minutes or even hours. One problem with delivering longer constant-frequency trains, however, is that the aversive effects tend to build up as train duration increases (Shizgal & Matthews, 1977), and thus longer trains will tend to get more aversive over time. Thus, a careful selection of the subjects may be necessary to get around this problem.

Beginnings-and-Peaks versus Temporal Summation

Initially, it was hypothesized that two possible processes would govern how the rats would respond to the delayed and weak-beginning trains. Recall that the first possibility was that the weak beginning would *worsen* the overall reward value of the 6w+2 s train, which is in line with a beginnings-and-peaks hypothesis. A second possibility was that the weak beginning would "bridge" the 6 s delay period, and thus would mitigate the effect of the delay. The original version of the experiment appeared to be in line with the bridging hypothesis. When a delayed Std was preceded by 6 s of weak stimulation, the strength of the Alt train had to be increased in order to achieve equipreference. When the Std was a 2 s train delayed by 6 s, only a weak Alt was sufficient for the rats to display equipreference between the two rewards.

Temporal summation offers a third explanation for the mitigating effect of the weak beginning. The weak train drives the reward signal up to a certain level, and once the strong train is delivered, its reward signal is summated with that of the weak train, reaching some higher level. The modified version of the beginnings-and-peaks experiment permitted a test of whether temporal summation was indeed in operation. By raising the Std train to a frequency sufficiently high to produce a maximal rewarding effect, temporal summation should have been irrelevant. Regardless of the signal level attained during the first 6 s, a maximum level would be attained during the 2 s component. Indeed, increasing the strength of the 2 s component dramatically reduced the contribution of the weak 6 s component. Thus, the evidence is stronger for temporal summation rather than "bridging". However, only one rat to date has completed the modified experiment, and thus more subjects need to be included.

Taken together, the results from Experiments 1-4 suggest the "best" moment of the experience (the peak) is the most important exemplar that is recorded into memory (Gallistel, 1978). These experiments provide no evidence that the signal level at the beginning or end is used as a separate exemplar in evaluating trains of rewarding brain stimulation.

Preference Reversal, Evaluation Context, and Duration Neglect

Based on the results of Experiment 5, it appears that, at least for some subjects, when evaluations are made in a joint-evaluation context such as a dual-operant paradigm, duration information per se may serve as an exemplar. In 4 of 7 rats, the performance curves for 2-4 s and 8 s trains were generally overlapping in the single-operant paradigm, but the strength of the shorter duration train had to be increased in the dual-operant paradigm, in order for it to compete effectively with the 8 s train. On the basis of the single-operant results, one might surmise that these particular rats were indifferent between the 2-4 s and 8 s trains. If so, changing the evaluation context from single to joint, produced a preference reversal. This suggests that duration neglect, at least for some individuals, is not absolute but rather to a certain extent is dependent upon the evaluation context (single-evaluation or joint-evaluation) in which the stimulation trains are presented.

Given that duration neglect is relevant to the peak-and-end model (Kahneman et al., 1993), how do the results of the peak-and-end tests (Experiments 1-3) fit with those of the preference reversal experiment? In Experiment 1, all stimulation trains were presented in a single-operant context, and in 5 out of 6 subjects, there were no significant differences between the 6 s and 8 s trains. These rats demonstrated duration neglect, consistent with the findings of Sonnenschein, Conover, and Shizgal (2003), which were also obtained in a single-evaluation context. In Experiment 2, a joint-evaluation context just like the preference reversal test, 4 out of the 5 rats showed equipreference between the 6 s and 8 s trains. This is not surprising given the relative similarity in duration between the two trains tested (6 s and 8 s), in contrast to the more disparate durations used in the preference reversal test (4 s and 8 s). Even in those subjects where duration information became more salient in joint evaluation, it seems likely that if two trains very similar in duration (6 s and 8 s) were to be presented to these rats, duration information would again prove unlikely to have a great impact.

As noted earlier, some researchers (Ariely, 1998; Ariely & Loewenstein, 2000) have argued that duration neglect is not an “absolute finding”. For example, the study of Ariely (1998) demonstrated that when an aversive stimulus (heat applied to the arm, or pressure applied to a finger) was of a constant pain intensity, duration information was not taken into account in the retrospective evaluations. Another set of stimuli were used as well, in which the profile of the aversive stimulus was manipulated. For example, in the "Up" presentations, the stimulus began at a low intensity, and then was increased over the course of the stimulus. In "Down" presentations, in contrast, the stimulus began at a high intensity, and then was gradually decreased over the course of the presentation. More complex stimulus presentations were also employed, such as mixed "Up & Down" and "Down & Up" stimuli. When these "patterned" pain intensity stimuli were used, duration information was taken into account by the subjects, such that increases in duration increased the remembered pain intensity.

Another example of the lack of absolutism of duration neglect can be found in the experiments of Ariely and Loewenstein (2000). In four experiments, human subjects were presented with various aversive sounds and asked to give different evaluations of the aversiveness of these sounds. In Experiment 1, they were asked to make an overall evaluation of the aversiveness of the sound, on a scale from 1-100. In Experiment 2, they were asked about their willingness to repeat some of the sound sequences, in exchange for money. In Experiment 3, they were given a fixed standard sound, and asked to rate the sounds they heard in reference to the standard. Finally, in Experiment 4, subjects were required to make choices between sound sequences. The purpose of these experiments was to determine if in fact the subjects in the Kahneman studies (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996) had been aware of duration information, but did not include this information in their retrospective evaluations. For example, if someone asks you to rate how your vacation went, when contemplating this rating you may consider the weather, the local events, and so on, but leave duration of your trip out of the picture (Ariely & Loewenstein, 2000).

Consistent with the predictions of Ariely and Loewenstein (2000), duration neglect was only seen in the first experiment, where the subjects were asked to make global evaluations. The authors argued these ratings were analogous to the types of ratings subjects had been asked to make in the Kahneman studies. In the other three experiments, however, duration information *was* taken into account by the subjects, lending support to the authors' assertions that subjects do encode information about stimulus duration, but it is only included in their retrospective evaluations when they believe this information is relevant to "the question".

Although one cannot verbally "change the question" with nonhuman animal subjects, perhaps running the animals in both single- and dual-operant paradigms constitutes "two ways of asking the question". In the case of the dual-operant paradigm, one could argue that such paradigms are a type of discrimination task in which the rats are performing a "rating" between the two trains presented, before deciding which train to respond for. Such a mode of computation, according to Ariely and Loewenstein (2000), should lead duration information to be included in the evaluation. The issue of what "question" is asked of the rats is less clear when one is using a single-operant paradigm. Are the rats indeed asked to make a kind of global "overall" evaluation of the current day's train? Are they calculating some kind of rating of the currently available train against their memory of other trains presented singly in previous days or weeks? In accounts of single-operant responding by Herrnstein (1970, 1974), as well as by behavioural economists (Kagel, Battalio, & Green, 1995) single-operant responding is viewed as another instance of choice, one where the subject chooses between working for the experimenter or engaging in other ("leisure") behaviours in the operant chamber. Thus, it may be more accurate to think of the single-operant paradigm as a "willingness to pay" (with "labour"; bar-pressing) paradigm. Ariely and Loewenstein's (2000) willingness-to-pay experiment (Experiment 2), however, suggested that duration information was salient to their subjects, and thus is in contrast to the findings of the single-operant Sonnenschein, Conover, and Shizgal (2003) experiment. In the latter experiment, it appeared that duration information was not heavily weighted when forming evaluations of the stimulation, at least when trains 2-4 s and longer were delivered.

Future BSR studies could be aimed at determining more precisely at what point duration neglect is likely to be displayed in a dual-operant paradigm. This would entail determining the chronaxie of the strength-duration function for trains in a dual-operant paradigm. The results of Experiment 5 suggest that the chronaxie would be longer, at least for some rats, in the dual-operant test than in the single-operant test.

Summary and Conclusions

To summarize, the "peak" is the only exemplar supported by the available evidence concerning the evaluation of BSR in single-operant paradigms. Duration may also act as an exemplar in some cases, in the dual-operant paradigm. The "peak" is the key exemplar regardless of whether the composite trains are designed with weaker "beginnings" or weaker "ends". When trains that varied in the strength of their "beginnings" were presented in a dual-operant paradigm, preliminary findings suggested that it was temporal summation and peak detection, rather than a bridging or beginnings-and-peaks heuristic, which accounted for the results. Lastly, in terms of the preference reversal data, 4 out of 7 rats demonstrated an influence of duration information when the 4 s and 8 s trains were presented in joint-evaluation, in comparison to their performance for the same trains in a single-evaluation context. In addition, duration neglect, as seen in Sonnenschein, Conover, and Shizgal (2003), was demonstrated reliably in all single-operant experiments in the present study.

Given that the "peak" model appears to be a robust finding across all four of the composite train experiments, it may be of interest to assess whether the "peak" model applies to other, more naturalistic types of appetitive experiences. Although BSR has some advantages over natural reinforcers, it would be possible to design an experiment that could avoid many of the problems associated with using natural reinforcers. For example, one could use the intraoral/intragastric cannulae procedures (Conover & Shizgal, 1994a, 1994b; Conover et al., 1994; Grill et al., 1987) described earlier. A sucrose solution delivered intraorally to the rat would not require a consummatory response separate from the instrumental response (and thus would be more similar to BSR delivery), and an intragastric cannula would prevent satiety. As mentioned earlier, if one could find a way to weaken the concentration of the sucrose solution, *during* sucrose delivery, perhaps this would suffice as a test of the peak-and-end model, this time using a more natural goal object.

Given that the rewarding effects of MFB stimulation have also been linked to drug addiction (Wise, 1987), it may also be of interest to determine whether a "peak" heuristic or similar exemplar is taken into account when rats are receiving either drugs or BSR that mimics drug effects (Lepore & Franklin, 1992). BSR trains designed to mimic drug effects could be very useful in this regard, as they would allow rigorous experimental control over the stimulation, yet are arguably more "naturalistic" than constant-frequency stimulation trains.

In conclusion, the present study indicates that in comparison to multiple-exemplar models, the "peak"-alone model of Gallistel (1978) seems best able to account for the behaviour of rats responding for rewarding brain stimulation. The present work has also suggested several fruitful avenues for future research, in terms of the dependence of duration neglect on evaluation context.

BSR experiments are germane to examining the neural computations involved in evaluating hedonic experiences. Such investigations may eventually allow the identification and elucidation of the underlying neural systems. Given the importance of these neural systems in motivated behaviours such as feeding, mating, and drug use, the identification of the candidate neurons would represent an important advance for neuroscience, psychology, and decision research.

References

Abrahamsen, G. C., Berman, Y., & Carr, K. D. (1995). Curve-shift analysis of self-stimulation in food-restricted rats: Relationship between daily meal, plasma corticosterone and reward sensitization. *Brain Research, 695*(2), 186-194.

Ariely, D. (1998). Combining experiences over time: The effects of duration, intensity changes and on-line measurements on retrospective pain evaluations. *Journal of Behavioral Decision Making, 11*, 19-45.

Ariely, D., & Loewenstein, G. (2000). When does duration matter in judgment and decision making? *Journal of Experimental Psychology: General, 129*(4), 508-523.

Balagura, S., & Hoebel, B. G. (1967). Self-stimulation of the lateral hypothalamus modified by insulin and glucagon. *Physiology and Behavior, 2*(4), 337-340.

Black, J., Belluzzi, J. D., & Stein, L. (1985). Reinforcement delay of one second severely impairs acquisition of brain self-stimulation. *Brain Research, 359*(1-2), 113-119.

Conover, K. L., & Shizgal, P. (1994a). Competition and summation between rewarding effects of sucrose and lateral hypothalamic stimulation in the rat. *Behavioral Neuroscience, 108*(3), 537-548.

Conover, K. L., & Shizgal, P. (1994b). Differential effects of postingestive feedback on the reward value of sucrose and lateral hypothalamic stimulation in rats. *Behavioral Neuroscience, 108*(3), 559-572.

Conover, K. L., Woodside, B., & Shizgal, P. (1994). Effects of sodium depletion on competition and summation between rewarding effects of salt and lateral hypothalamic stimulation in the rat. *Behavioral Neuroscience, 108*(3), 549-558.

Fouriezos, G. (1995). Temporal integration in self-stimulation: A paradox lost? *Behavioral Neuroscience, 109*(5), 965-971.

Fouriezos, G., & Randall, D. (1997). The cost of delaying rewarding brain stimulation. *Behavioral Brain Research, 87*(1), 111-113.

Fredrickson, B. L., & Kahneman, D. (1993). Duration neglect in retrospective evaluations of affective episodes. *Journal of Personality and Social Psychology, 65*(1), 45-55.

Fulton, S., Woodside, B., & Shizgal, P. (2000). Modulation of brain reward circuitry by leptin. *Science, 287*(5450), 125-128.

Gallistel, C. R. (1976). Spatial and temporal summation in the neural circuit subserving brain-stimulation reward. In A. Waquier & E. T. Rolls (Eds.), *Brain Stimulation Reward* (pp. 97-99). New York: Elsevier.

Gallistel, C. R. (1978). Self-stimulation in the rat: Quantitative characteristics of the reward pathway. *Journal of Comparative and Physiological Psychology, 92*(6), 977-998.

Gallistel, C. R., & Leon, M. (1991). Measuring the subjective magnitude of brain stimulation reward by titration with rate of reward. *Behavioral Neuroscience, 105*(6), 913-925.

Gallistel, C. R., Shizgal, P., & Yeomans, J. S. (1981). A portrait of the substrate for self-stimulation. *Psychological Review, 88*(3), 228-273.

Grill, H. J., Spector, A. C., Schwartz, G. J., Kaplan, J. M., & Flynn, F. W. (1987). Evaluating taste effects on ingestive behavior. In F. M. Toates & N. Rowland (Eds.), *Feeding and drinking* (pp. 151-188). Amsterdam: Elsevier.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*(2), 243-266.

Herrnstein, R. J. (1974). Formal properties of the matching law. *Journal of the Experimental Analysis of Behavior, 21*(1), 159-164.

Hoebel, B. G. (1968). Inhibition and disinhibition of self-stimulation and feeding: Hypothalamic control and postingestional factors. *Journal of Comparative and Physiological Psychology, 66*(1), 89-100.

Hoebel, B. G., & Teitelbaum, P. (1962). Hypothalamic control of feeding and self-stimulation. *Science, 135*, 375-377.

Hoebel, B. G., & Thompson, R. D. (1969). Aversion to lateral hypothalamic stimulation caused by intragastric feeding or obesity. *Journal of Comparative and Physiological Psychology, 68*(4), 536-543.

Hsee, C. K., Blount, S., Loewenstein, G. F., & Bazerman, M. H. (1999). Preference reversals between joint and separate evaluations of options: A review and theoretical analysis. *Psychological Bulletin, 125*(5), 576-590.

Kagel, J. K., Battalio, R. C., & Green, L. (1995). *Economic choice theory: An experimental model of animal behavior*. Cambridge: Cambridge University Press.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science, 4*(6), 401-405.

Lattal, K. A., & Gleeson, S. (1990). Response acquisition with delayed reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes, 16*(1), 27-39.

Lepore, M., & Franklin, K. B. (1992). Modelling drug kinetics with brain stimulation: Dopamine antagonists increase self-stimulation. *Pharmacology, Biochemistry, and Behavior, 41*(3), 489-496.

Margules, D. L., & Olds, J. (1961). Identical "feeding" and "rewarding" systems in the lateral hypothalamus of rats. *Science, 135*, 374-375.

Mark, T. A., & Gallistel, C. R. (1993). Subjective reward magnitude of medial forebrain stimulation as a function of train duration and pulse frequency. *Behavioral Neuroscience, 107*(2), 389-401.

Mason, P. A., & Milner, P. M. (1985). Short- and long-term summation characteristics of electrical self-stimulation reward. *Behavioral Brain Research, 18*(3), 223-231.

Mason, P. A., & Milner, P. M. (1986a). Temporal characteristics of electrical self-stimulation reward: Fatigue rather than adaptation. *Physiology and Behavior, 36*(5), 857-860.

Mason, P. A., & Milner, P. M. (1986b). Further evidence against adaptation of prolonged electrical self-stimulation reward. *Physiology and Behavior, 36*(5), 861-865.

Mazur, J. E., Stellar, J. R., & Waraczynski, M. (1987). Self-control choice with electrical stimulation of the brain as a reinforcer. *Behavioural Processes, 15*(2-3), 143-153.

Milner, P. M. (1989). The discovery of self-stimulation and other stories. *Neuroscience and Biobehavioral Reviews, 13*(2-3), 61-67.

Morgan, C. W., & Mogenson, G. J. (1966). Preference of water-deprived rats for stimulation of the lateral hypothalamus rather than water. *Psychonomic Science, 6*(7), 337-338.

Mundl, W. J. (1980). A constant-current stimulator. *Physiology and Behavior, 24*, 991-993.

Olds, J. (1973). The discovery of reward systems in the brain. In E. S. Valenstein (Ed.), *Brain stimulation and motivation: Research and commentary.* (pp. 81-99). Glenview, Ill: Scott Foresman.

Olds, J., & Milner, P. (1954). Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of Comparative and Physiological Psychology, 47*, 419-427.

Paxinos, G., & Watson, C. (1998). *The rat brain in stereotaxic coordinates* (4 ed.). San Diego, California: Academic Press.

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain, 66*(1), 3-8.

Redelmeier, D. A., Katz, J., & Kahneman, D. (2003). Memories of colonoscopy: A randomized trial. *Pain, 104*, 187-194.

Reed, P., & Reilly, S. (1990). Context extinction following conditioning with delayed reward enhances subsequent instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes, 16*(1), 48-55.

Sax, L. D., & Gallistel, C. R. (1984). Temporal integration in self-stimulation: A paradox. *Behavioral Neuroscience, 98*(3), 467-468.

Schreiber, C. A., & Kahneman, D. (2000). Determinants of the remembered utility of aversive sounds. *Journal of Experimental Psychology: General, 129*(1), 27-42.

Shizgal, P. (1997). Neural basis of utility estimation. *Current Opinions in Neurobiology, 7*(2), 198-208.

Shizgal, P. (1999). On the neural computation of utility: Implications from studies of brain stimulation reward. In D. Kahneman (Ed.), *Well being: The foundations of hedonic psychology* (pp. 500-524). New York, NY, US: Russell Sage Foundation.

Shizgal, P., & Matthews, G. (1977). Electrical stimulation of the rat diencephalon: Differential effects of interrupted stimulation on on- and off-responding. *Brain Research, 129*(2), 319-333.

Shizgal, P., & Murray, B. (1989). Neuronal basis of intracranial self-stimulation. In J. Liebman & S. Cooper (Eds.), *The neuropharmacological basis of reward. Topics in experimental psychopharmacology.* (Vol. 1, pp. 106-163). Oxford: Clarendon Press/Oxford University Press.

Sonnenschein, B. H., Conover, K. L., & Shizgal, P. (2003). Growth of brain stimulation reward as a function of duration and stimulation strength. *Behavioral Neuroscience, 117*(5), 978-994.

Talwar, S. K., Xu, S., Hawley, E. S., Weiss, S. A., Moxon, K. A., & Chapin, J. K. (2002). Rat navigation guided by remote control. *Nature, 417*(6884), 37-38.

Vezina, P., & Stewart, J. (1987a). Morphine conditioned place preference and locomotion: The effect of confinement during training. *Psychopharmacology, 93*(2), 257-260.

Vezina, P., & Stewart, J. (1987b). Conditioned locomotion and place preference elicited by tactile cues paired exclusively with morphine in an open field. *Psychopharmacology, 91*, 375-380.

Wise, R. A. (1980). Action of drugs of abuse on brain reward systems. *Pharmacology, Biochemistry, and Behavior, 13 Suppl 1*, 213-223.

Wise, R. A. (1987). The role of reward pathways in the development of drug dependence. *Pharmacology and Therapeutics, 35*(1-2), 227-263.

Wise, R. A. (1996). Addictive drugs and brain stimulation reward. *Annual Review of Neuroscience, 19*, 319-340.

Yeomans, J. (1988). Mechanisms of brain-stimulation reward. In J. M. Sprague & A. N. Epstein (Eds.), *Progress in psychobiology and physiological psychology* (Vol. 13, pp. 227-266). New York: Academic Press.

Yeomans, J. S., Maidment, N. T., & Bunney, B. S. (1988). Excitability properties of medial forebrain bundle axons of A9 and A10 dopamine cells. *Brain Research, 450*(1-2), 86-93.

Appendix

(Note: this section closely follows the Appendix in Sonnenschein, Conover, and Shizgal (2003), which was originally written by P. Shizgal).

Derivation of the triple-logistic model

Dr. Shizgal's model of temporal integration is represented by the black mesh in Figure 3d (as well as in Figure 7). This "triple-logistic" (TL) model was based on three previously established relationships. The first relationship is the hyperbolic strength-duration function proposed by Gallistel (1978), and shown in Figure 3a. If such a relationship holds not only for performance (i.e., bar-pressing), but also for the subjective intensity of the reward, then the frequency required to produce a reward of half-maximal rreward intensity will decrease hyperbolically as the train duration is increased.

This is captured in the equation:

$$F_{hm} = F_R \times \left(1 + \frac{C}{D}\right) \tag{1}$$

where C = chronaxie of strength - duration function for trains (the duration at which F_{hm} equals twice F_R)

D = train duration

F_{hm} = pulse frequency required to produce a reward of half - maximal intensity

F_R = frequency at which intensity is half - maximal at an infinite train duration

The second equation in the TL model encompasses the idea that reward growth increases as a sigmoid function. This function was derived based on matching experiments conducted by Gallistel and Leon (1991). Using self-stimulation sites similar to those employed in the present study, they demonstrated that the subjective intensity of reward grows initially as a power function of the frequency, and then levels off. A logistic is used to capture this relationship in the following equation (and shown in Figure 3b):

$$I = \frac{I_{\max}}{1 + \left[\frac{F_{hm}}{F}\right]^{g}} \qquad (2)$$

where F = pulse frequency
F_{hm} = frequency at which the intensity of the BSR is half - maximal
g = exponent of intensity growth
I = intensity of BSR
$I_{\max}$ = maximum intensity of BSR

If Equation 1 is substituted for F_{hm} in Equation 2, we obtain the following equation, which combines both the influence of frequency and train duration into a single expression. This equation is thus the central feature of the TL model and is necessary to its ability to account for the apparently discrepant results described by Gallistel (1978) on the one hand, and by Mason and Milner (1986a, 1986b) on the other.

$$I = \frac{I_{\max}}{1 + \left[\frac{\left\{ F_R \times \left(1 + \frac{C}{D} \right) \right\}}{F} \right]^g} \tag{3}$$

where C = chronaxie of strength - duration function for trains
D = train duration
F = stimulation frequency
F_R = frequency at which intensity of BSR is half - maximal at an infinite train duration
g = exponent of intensity growth
I = intensity of BSR
$I_{\max}$ = maximum intensity of BSR

The third and final component of the model is the equation relating performance (bar-pressing) to the subjective intensity of the reward signal. Another sigmoid function captures this relationship (and is shown in Figure 3c):

$$PRH_{RS} = \frac{1}{1 + \left[\frac{U_e}{R \times I} \right]^p} \tag{4}$$

where I = intensity of BSR
p = performance exponent
PRH_{RS} = proportion of rewards harvested (rescaled)
R = rate of reinforcement
and U_e = utility of "everything else"

When $p = 1$, Equation 4 reduces to Herrnstein's (1970) single-operant matching law, an expression which was formulated to account for performance on variable-interval schedules. Given that a continuous reinforcement schedule was in effect throughout the present study, it was expected that performance would grow much more steeply as a function of payoff, than would be the case under a variable-interval schedule. Thus, the exponent p is expected to be greater than unity. The "everything else" in Equation 4 refers to activities other than working for BSR which the rat may engage in, such as exploration, grooming, chewing, and so on. U_e, the utility of "everything else", can be defined as the product of the maximal subjective intensity of BSR (I_{max}) and the rate of reinforcement (R_e). R_e can be defined as the rate at which the stimulation that produces the subjective intensity of I_{max} must be delivered to produce a payoff equal to U_e.

If Equation 3 is substituted for I in Equation 4, and keeping in mind that U_e is the product of I_{max} and R_e, the complete TL model is obtained:

$$PRH_{RS} = \frac{1}{\left\{1+\left[\frac{R_e}{R}\times\left(1+\left\{\frac{F_R}{F}\times\left(1+\frac{C}{D}\right)\right\}^{g}\right)\right]^{p}\right\}} \tag{5}$$

where C = chronaxie
D = train duration
F = frequency
F_R = rheobasic frequency
g = exponent of intensity growth
p = performance exponent
PRH_{RS}= proportion of rewards harvested (rescaled)
R = nominal rate of reinforcement
R_e = rate at which utility of a maximal BSR equals the utility of "everything else"

This equation, represented by the black mesh surface in Figure 3d (as well as in Figure 7), predicts performance (expressed as the proportion of rewards harvested) as a function of the two independent variables, frequency (*F*) and train duration (*D*).

The rate of reinforcement is considered to be nominal, because the continuous reinforcement schedule employed in the frequency and duration sweeps experiments makes it impossible to measure this variable precisely. Thus a single parameter, the RRR ("relative rate of reinforcement"), was substituted for Re/R when the model was fit to the data. This RRR value can be thought of as a scaling factor.