INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA

UMI®
800-521-0600

Temporal Integration in the Neural Substrate for Brain Stimulation Reward:

Duration Neglect and the Peak-and-End Model

Bonnie Sonnenschein

A Thesis

in

The Department

of

Psychology

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

August 1999

0-612-43623-3

ABSTRACT

Temporal Integration in the Neural Substrate for Brain Stimulation Reward: Duration Neglect and the Peak-and-End Model

Bonnie H. Sonnenschein

There are two competing theories of how organisms evaluate temporally-extended appetitive experiences. The "peak" model states that the most rewarding moment of the experience is recorded into memory, while the "peak-and-end" model asserts that remembered reward value is an averaging of both the "peak" reward and the reward value experienced at the end of the event. Both models also suggest that information about the temporal duration of the experience plays little role in retrospective evaluations ("duration neglect"). The current brain stimulation reward (BSR) studies: a) tested a mathematical model designed to predict how lever-pressing performance in self-stimulating rats is affected by varying the train duration and pulse frequency, b) examined how train duration and pulse frequency affect "duration neglect", and c) compared the "peak" and "peak-and-end" models. Two male rats were chronically implanted with stimulating electrodes aimed at the lateral hypothalamic area. In the first experiment, train duration was held constant while the pulse frequency was varied. In the second experiment, frequency was held constant while the train duration was varied. In the third experiment, constant-frequency trains were compared to trains in which the frequency at the "end" of the train was less rewarding than that of the "peak". The results were consistent with the mathematical model, and allowed for a better understanding of "duration neglect" in BSR. Furthermore, while one subject displayed a lack of a peak-and-end effect, the other rat's results appeared to support the peak-and-end model, although instabilities in performance in the latter rat may have led to a spurious result.

Acknowledgements

No thesis would be complete without an acknowledgement of all those (besides the author) who helped to create it.

First of all, I would like to thank my supervisor, Peter Shizgal, for the countless hours of support and guidance, both academic and social in nature, that he has provided to me over the course of this work. His explanations have helped to turn an area of technical and mathematical mystery into something that even a "Mathematics Disorder" sufferer such as myself can understand and explain to others, and for that I am most grateful. I am also grateful to him for being understanding during my battles to embrace both the Clinical and Experimental sides of my nature, something that not all research supervisors are willing to do.

Secondly, I would like to thank the two members of my thesis committee, Pierre-Paul Rompré and Norman Segalowitz, for their helpful and timely comments during the "construction" of this work, and for helping to refine the final stages.

I am also very much indebted to those who laboured long and hard on the technical and technological aspects of this thesis. Many thanks go to the Shizgal lab Research Associate, Kent Conover, for the many hours he has put in on software development, and for the many interesting discussions/debates he and I have engaged in on behavioral economics, the philosophy of the mind, mental disorder, and operant

psychology, among other topics. I am also indebted to Steve Cabilio, computer "guru", for his generous help during periods of software, hardware, and neural breakdown. Furthermore, I am indebted to Dave Munro, who spent much time and effort in the construction and debugging of the "Generation II PREF" operant chambers.

I would also like to thank Fonds FCAR of Quebec, and the Natural Sciences and Engineering Research Council of Canada (NSERC) for their funding of my personal feeding and housing over the last two years, without which conducting my research would have been difficult indeed!

I am also indebted to my parents, Sam and Beverly Sonnenschein, and my brother Harley, for taking my bouts of "thesis-related" insanity in stride for quite some time, and for all their love and support of my endeavours over the course of my existence. Thanks are also due to my herd of friends who have been kind enough to deal with my foibles and idiosyncrasies, in some cases for nearly a decade: Heather McCauley, Nathalie Benchitrit, Catherine Filion, Langdon Roberts, Nicole Boucher, Devin Mueller, Eric Hargreaves, Tansy Evely, Paula Bennett, Tracy Davey, Tara Conroy, Helen Filipopoulos, Christy Atwood, Bettina Grassman, and Kathy De Sousa.

Lastly, I would like to thank my significant other, Frank Verpaelst, for all the love and support that he has given me, especially over the last few months. The completion of this thesis would have been much more difficult without his encouragement and pep-talks, and the "reality-checks" and wider perspective that he is always willing to offer.

Table of Contents

List of Figures

Past studies have demonstrated that when forming evaluations of past experiences (retrospective evaluations), humans do not recall the experience in a "minute-by-minute" fashion. Rather, they appear to take certain characteristics of the experience and record those into memory, to be used in later evaluations and choice behaviour. In other words, rather than "taking a film" of the experience, the organism just takes a few "snapshots" (Ariely & Carmon, in press; Fredrickson & Kahneman, 1993). These "snapshots" are usually referred to as "exemplars" (Fredrickson & Kahneman, 1993) or "gestalt characteristics" (Ariely & Carmon, in press), and constitute a major component of the "representation by exemplar" model, which can be thought of as a type of cognitive heuristic.

One group of researchers has proposed that humans may use just such a type of "representation by exemplar" model, called the "peak-and-end" model (Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993) when evaluating aversive experiences and forming memories of such experiences. One study suggested that patients undergoing painful medical procedures, such as colonoscopies, seem to take the "peak" of the painful experience (i.e. the "worst" moment of pain) and the "end" of the experience (i.e. the intensity of pain right before the end of the procedure), and select an intermediate value between these two exemplars which is then encoded into memory (Redelmeier & Kahneman, 1996). The actual temporal length of the procedure seems to play a minimal role in how patients actually remember the experience, and this neglect of temporal information is called "duration neglect" (Kahneman et al., 1993).

These ideas are illustrated in a series of studies (Fredrickson & Kahneman, 1993;

Kahneman et al., 1993; Redelmeier & Kahneman, 1996). In one study (Kahneman et al., 1993), human subjects were asked to place their hand in cold water. In one set of trials ("short trials") subjects kept their hands in cold water for 60 seconds while the temperature of the water remained constant. In another set of trials ("long trials") the subjects kept their hands in the cold water for 60 seconds, and then for an additional 30 seconds while the water temperature was raised slightly. When asked which trial they preferred to repeat, most subjects selected the long trials. Since it is counter-intuitive that the subjects would wish to suffer for a longer period of time, these results suggested that the duration of the aversive experience played a minimal role in their formation of retrospective evaluations. In fact, the duration of the experience accounted for only 2% of the variance in the subjects' retrospective evaluations.

The study of Redelmeier et al. (1996) revealed similar results. In this case, patients undergoing painful medical procedures were asked to make both real-time and retrospective evaluations about the intensity of pain during the procedure. Results revealed that even though the medical procedures varied widely in total duration (from 4 minutes to 67 minutes), there was no significant correlation between the duration of the procedure and the patients' retrospective evaluations. Thus, the patients showed "duration neglect". In contrast, the retrospective evaluations were highly correlated with both the "peak" moment of pain and the pain that was experienced at the "end" of the procedure. This pattern was found for both delayed and immediate retrospective evaluations.

Duration neglect was also observed in a study of retrospective evaluations of emotional experiences (Fredrickson & Kahneman, 1993). In this study subjects had to make retrospective evaluations of their affective responses to various pleasant and aversive film clips, which varied in duration. Results showed that even though the subjects were able to distinguish between long and short film clips, the actual duration of the film had little impact on their retrospective evaluations, as in the studies mentioned above.

To summarize, it appears that humans take a few exemplars of an experience, such as the "peak" and the "end", and use them to form evaluations, while information about the temporal length of the experience is neglected. This means that, according to the model, people should prefer a *longer* aversive experience (pain spread out over a longer period of time) as long as the pain lessens towards the end of the experience (Redelmeier & Kahneman, 1996), which initially appears to be a counter-intuitive prediction.

In the domain of appetitive stimuli, one might expect the peak-and-end model to lead to another counter-intuitive prediction. Consider the situation in which a human or animal subject must choose between a "short" pleasant experience which has a constant reward value, and a "long" pleasant experience where the reward value drops off somewhat at the end (analogous to the cold water study by Kahneman et al. (1993) mentioned earlier). The peak-and-end model predicts, in effect, that if the actual duration of the experience is not taken into account, the subject should prefer the shorter appetitive experience, with its constant reward value, over the longer experience with its "weaker" end.

Thus, it was of interest in the current series of studies to examine whether the peak-and-end model could account for the evaluation of appetitive experiences. The appetitive experience used was brain stimulation reward (BSR). While it is generally presumed in BSR research that when discrete stimulation (such as short bursts of pulses) is used the reward value of the stimulation builds up over a brief amount of time, eventually reaching an asymptote, it is unclear how an organism performs temporal integration and evaluation of BSR over *prolonged* appetitive stimulation in which the rewarding effect waxes and wanes.

It should be mentioned at this point that there is a difference between BSR stimulation trains as they are usually administered, and the aversive experiences that the subjects in the Kahneman experiments typically underwent. In most of the aversive experiences used by Kahneman et al., such as the cold water study mentioned earlier (Kahneman, 1993), the peak moment of pain usually occurs near the middle of the experience, whereas in BSR, it is assumed that because reward value rises as the duration of the stimulation increases, up to an asymptote, the "peak" of a BSR train is assumed to occur at the "end" of the train. Thus, one might believe that the "peaks" and "ends" of BSR stimulation trains are confounded. However, according to Kahneman's theory, the location of the "peak" within the experience is irrelevant. If the "peak" and the "end" do occur at the same point in the profile of the experience, then the organism should still take both exemplars into account and perform an averaging on them. For example, suppose that in a pain experiment, the highest pain rating possible is 10. If a painful experience in which the peak is rated as a 10 and occurs in the middle of the experience, and in which the end is also rated as a 10 (experience "A"), is compared to a painful experience in which the worst moment of pain occurs at the end of the experience and is also rated as a 10 (experience

"B"), according to the theory both will yield the same overall pain evaluation. In experience "A", 10 (the peak) plus 10 (the end) summed and averaged is still 10, while in experience "B", the peak and the end are both 10 (they occur at the same moment and are thus equal according to the theory), and summing and averaging them will yield the same value as in experience A. In addition, if both experience A and B are compared to an experience "C", in which the peak occurs in the middle of the experience and is rated as a 10, and the end of the experience is rated as a 5, both A and B will be remembered as being more aversive than experience C, as averaging 10 (the peak) and 5 (the end) of experience C will yield 7.5, a lower overall pain rating than the other two pain profiles. The same should also be true of appetitive experiences, according to the model. If all three experiences are pleasant ones, with a "10" being the rating of highest "pleasantness", then evaluations of the positiveness of experiences A and B should be equal, and both should be remembered as being more rewarding than experience C (as the lower value (7.5) now corresponds to a lower rating of pleasantness). Thus, even though the "peaks" and "ends" of BSR trains are confounded, this does not pose a problem for the testing of Kahneman's theory.

However, past BSR research has not directly addressed the Kahneman theory, that is, the question of how a weaker "end" added to a stimulation train may degrade the overall subjective reward value of the train, because most BSR researchers usually use stimulation trains in which the pulse frequency is held constant throughout the train, usually referred to as "rectangular" trains. It is assumed in cases such as constant-frequency trains, as described earlier, that the reward value of the train increases as the duration of the input increases, until an asymptote is reached, with both the peak and end thus being the same (as in experience "B"). However, none of the past BSR research has been conducted with any stimulation trains analogous to experience "C"; no studies have compared constant-frequency trains with composite trains in which the pulse frequency (and thus the reward value) is greater in one portion of the train than another. It is assumed in the Shizgal lab that if one were to construct such a stimulation train, in which the first portion of the train was made up of a high pulse frequency, and the second portion of the train of a lower and less rewarding pulse frequency, that the "peak" of the subjective intensity of reward would be reached sometime near the end of the "strong" portion of the train. Then, as the less effective stimulation in the second half of the train continues, it is hypothesized that the overall reward value of the train would degrade, compared to a standard constant-frequency train equal in frequency to the "strong" portion of the composite train (experience B vs. experience C, as described above).

Based on experimentation with the typical constant-frequency rectangular BSR trains, one research group (Norman & Gallistel, 1978) has suggested that it is only the "peak" reward value of the train that is recorded into memory. However, as such a theory has not been tested with composite trains in which the "peak" and "end" of the train differ in pulse frequency and reward value, the current studies attempted to pit the "peak" model against the "peak-and-end" model using composite-frequency trains, to assess which model constitutes the most parsimonious explanation of how organisms perform temporal integration of the reward value of extended experiences.

It therefore appears that BSR can be used as an appetitive experience for the purposes of testing the predictions of Kahneman's peak-and-end model. The phenomenon of BSR was first discovered by James Olds and Peter Milner in 1953 at the McGill University psychology department (Milner, 1989). Olds had been attempting to implant an electrode into the tegmental area of one of his rats, but during surgery the electrode had been pushed forward into the septal area (although this was not discovered until later). Olds discovered a few days after the surgery that when this particular rat was placed on a tabletop and the stimulation was turned on, it would advance, sniff, and perform other exploratory behaviours, and that it would stop or turn back when the stimulation was terminated. Furthermore, by giving the rat a brief burst of stimulation every time it turned in a certain direction it was possible to guide the rat to any region of the table that the experimenter desired, and the shaping of other behaviours was also fairly easy to accomplish. This included training the rat to bar-press for stimulation. On the basis that the phenomenon would be more believable if the rat could be trained to initiate the stimulation itself in the absence of possible cues from the experimenter, the rat was placed in a Skinner box constructed for the occasion, where it rapidly learned to bar-press for stimulation. Once the reliability of this behaviour was established, Milner tried to replicate the results by implanting other rats in the tegmental area, but as these rats failed to self-stimulate, it was suspected that the electrode was in fact in a different location than had been originally intended, and after x-raying the head of Olds' rat (as he was obviously unwilling to sacrifice it at that point), the location of the electrode was discovered to be near the septal area. Subsequently, it was confirmed that rats implanted with stimulating electrodes in the septal area would perform self-stimulation.

Other areas of the brain were also later found to support self-stimulation (Olds & Milner, 1954). Rats were chronically implanted with stimulating electrodes in various regions of the forebrain and midbrain, and again operant boxes were used in order to train the rats to bar-press for the delivery of the stimulation. The strongest operant responding was seen in subjects that had electrodes implanted in the septal area, as was discovered earlier. For example, even the lowest-responding septal rat responded at a rate of 285 bar-presses per hour. However, rats with electrodes implanted in other areas of the forebrain, such as the cingulate cortex and the mammillothalamic tract, were also shown to engage in vigorous responding, while rats stimulated in other brain areas did not display such robust lever-pressing. Milner and Olds went on to conclude that since rewarding effects were found to arise from stimulation of several brain regions, this suggested that there exists within the brain a "reward system", which they referred to as a collection of "reinforcing structures" (p. 425) (Olds & Milner, 1954).

More information about the characteristics of this neural "reward system" was revealed in later studies by many different researchers. The work of all these groups and individuals is well summarized in a review by Gallistel, Shizgal, and Yeomans (1981), who report that the results of these various studies have shown that: a) the neurons that tend to support self-stimulation when directly stimulated by electrode form part of a bundle of axons that run through the "medial forebrain bundle" (MFB) and the midbrain tegmentum (thus, the reward system is often referred to as the MFB reward system), b) they have absolute refractory periods that range from 0.5 to 1.2 milliseconds, c) their diameter is approximately 0.3 to 1.5 micrometers (µm), d) the axons of these neurons are myelinated, e) the directly-stimulated, "first-stage" neurons, conduct action potentials at velocities ranging from 2 to 8 meters per second, and f) these directly stimulated neurons cannot be dopaminergic, noradrenergic, nor any other type of catecholamine neuron, because their properties are different from those associated with catecholamine neurons; catecholamine neurons are unmyelinated, conduct impulses at a velocity less than 1 meter per second, and have refractory periods that range from 1.2 to 2.5 milliseconds (for example, see Yeomans, Maidment and Bunney (1988) and their study of dopamine neurons).

Even though there remains much to be learned about the physiological properties of the "first-stage" neurons, even less is known about the next step in the link between the stimulated neurons and the operant behaviour that the self-stimulating subject performs. It is believed that the volley of action potentials that arises from the electrical stimulation courses down the axons of the first-stage neurons, and then passes across a synapse or several synapses into a neural network that is usually referred to as the neural "integrator" (Gallistel et al., 1981), which integrates the number of action potentials in the volley over space and time. Research has demonstrated that the output (the reward signal) of this integrator is determined by the aggregate rate of firing at its input. The reward value of the stimulation is summated over space and time in such a way that it does not matter whether ten directly stimulated neurons fire ten times each within a certain length of time or whether fifty neurons fire only two times each, for example. The integrator's output will be the same in both cases, provided that the aggregate impulse flow is constant. This idea is referred to as the "counter model" (Gallistel, 1978; Gallistel, Shizgal & Yeomans, 1981).

However, this counter model theory was not the last word on temporal integration in the neural substrate of reward. For example, a study conducted by Sax and Gallistel (1984) revealed that if a stimulation train is separated into two bursts (short trains) separated by an interval (referred to as the inter-burst-interval or IBI), there was no significant effect of the IBI for intervals of up to 2 seconds. At IBIs of more than 2 seconds, the pair of bursts was less rewarding to the rats than any of those with IBIs of 2 seconds or less. That is, at IBIs of 2 seconds or shorter, the rats treated the two bursts as if they were only one train, while at longer IBIs, they appeared to treat each burst as a separate (and thus less rewarding) small train. According to Sax and Gallistel, these results implied that the neural integrators were conducting "perfect integration of that portion of a reward signal that exceeds some threshold" (p. 473). In other words, if the interval between the two bursts was short enough, the neural integrator recognised both bursts as being part of the same train and summated their reward values. For example, it was found that two 0.2-second bursts separated by an IBI of 2 seconds was as rewarding to the rats as a single 0.4-second train with all the same characteristics except for the presence of an IBI. However, Fouriezos (1995) later discovered that such a "perfect integration above a threshold" was not necessary to account for the results, and instead Gallistel's (1987) older "leaky integrator" model (described below) can account for cases in which bursts of stimulation are separated by IBIs of 2 seconds or less. In Fouriezos' study, the first burst of the pair was set at a subthreshold frequency, and then the second burst was used to see what pulse frequency was required to sustain responding, with increasing IBIs between the two bursts. The results showed that more and more pulses in the second train were necessary to offset the effects of longer and longer IBIs. The longer the IBI, the longer the period for the rewarding effect of the first burst to decline, and thus more pulses were necessary in the second burst to compensate for this decay in reward value. Fouriezos further concluded that the rate of decay of reward value in his study was predicted by Gallistel's (1978) "leaky integrator" model.

Gallistel's (1978) leaky integrator model of temporal integration of BSR is as follows: he believes that there is one integrator involved in generating the reward signal, and he argues that this single integrator behaves much like a "leaky bucket being filled by a spurting hose" (p. 978). In such a system, the number of firings of the stimulated neurons required to produce a given output (a behavioural response, such as lever-pressing) would depend on the duration of the input (the stimulation train), everything else being held constant. Also, as noted above, if the input to the integrator is stopped or delayed, the excitation in the integrator decays at an exponential rate, as a function of time. In addition, a short train duration would have to be strong in order to support stable behaviour from the subject, but as the train duration increases, the required strength would decrease (Gallistel, 1978). This latter relationship is illustrated in Figure 1, where the duration of the train is plotted on the x-axis and the pulse frequency (pulses per second) required to maintain half-maximal performance from the subject is plotted on the y-axis. As shown in the figure, the decrease in the required strength of the input initially occurs rapidly, but then becomes more gradual, until it levels off at an asymptote, which Gallistel (1978) refers to as the "rheobase". From studying the figure, it appears that the relationship between the frequency required to maintain half-maximal performance and the duration of the stimulation train used follows a

Figure 1. **Representative hyperbolic strength-duration function, based on Gallistel's (1978) equation. The green curve represents the hypothetical hyperbolic function. Plotted on the x-axis is the stimulation train duration, in seconds, and on the y-axis, the common logarithm of the pulse frequency required to maintain half-maximal (50%) of the maximum possible subjective intensity of reward.**

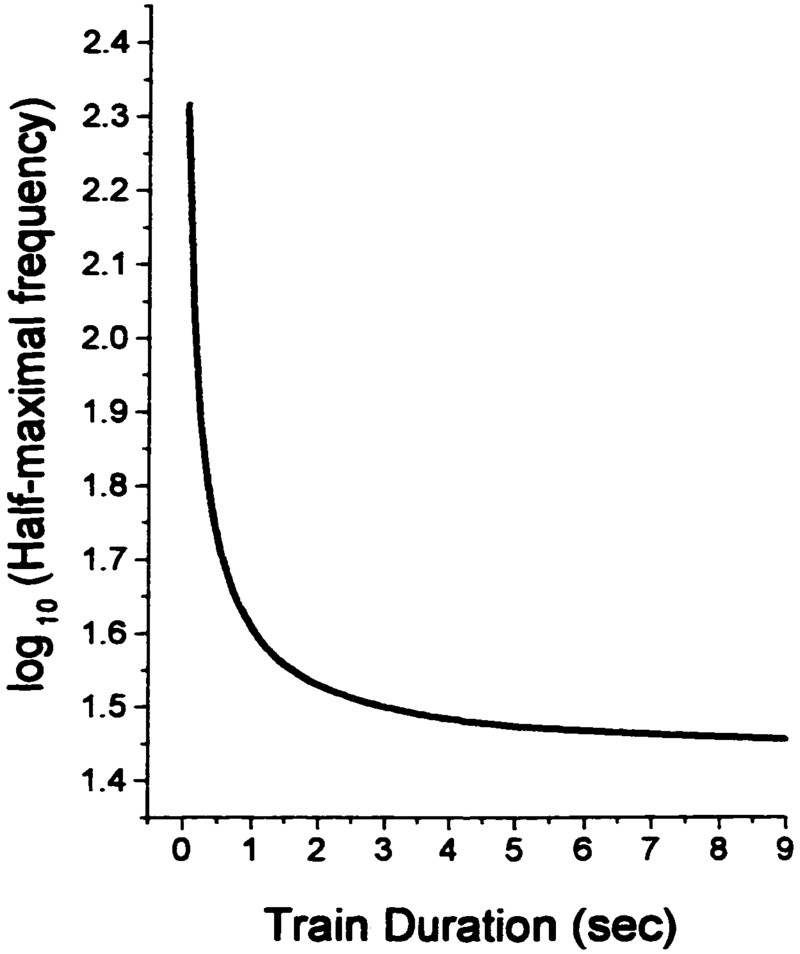

Figure 1. Typical hyperbolic strength-duration function.

rectangular hyperbolic function (which is usually referred to as the "strength-duration function"). This model of the reward integrator suggests that more firings are required when long-duration rather than short-duration trains are used because there is more time for the decay of excitation ("leaking") during the longer stimulation trains. Thus, the rectangular hyperbolic function is assumed to reflect the decay characteristics (the "leakiness") of this neural integrator.

The idea that the relationship between the duration of the stimulation train and the pulse frequency required to maintain a certain performance criterion is a rectangular hyperbola has been supported by some studies. In one study (Gallistel, 1978), rats that had been chronically implanted with electrodes aimed at the lateral hypothalamic area (LHA) were trained to shuttle back and forth between a priming lever and a reward lever, separated by a runway. The stimulation presented on the reward lever was identical to the stimulation available on the priming lever, and the duration of the stimulation trains presented on the two levers was varied from session to session, with the stimulation pulse frequency required to produce 70% of maximum performance was selected for analysis. The results showed that when train duration was plotted against the charge (in microcoulombs) required to maintain 70% of maximum running speed (the "required charge") in the runway between the two levers, there was a strong linear relationship between train duration and required charge for the majority of the subjects. In order to obtain the hyperbolic function mentioned earlier, one must first take the equation that describes the straight-line relationship between required charge and duration (Gallistel et al., 1981; Mason & Milner, 1985):

$$Q = RD + Q_0$$

where

Q = the required charge (the charge required to maintain criterion performance)

D = the stimulation train duration

R = slope of the straight line

Q_0 = the charge required to produce the criterion performance when the train duration goes to zero; the y-intercept.

Since the charge (Q) is the product of the stimulation strength and the train duration (D), one can derive the hyperbolic strength-duration function by dividing the linear equation given above by the train duration (D), to obtain (Gallistel et al., 1981):

$$Q' = R + Q_0/D$$

where

Q' = stimulation strength (equivalent to Q divided by D)

This equation can also be written in another form (Gallistel et al., 1981):

$$Q' = R(1 + C/D)$$

where

Q' = stimulation strength (in microcoulombs per second)

D = the train duration

R = the rheobase of the strength-duration function

C = the chronaxie (equivalent to Q_0/R), an index of the "curviness" of the rectangular hyperbola

This is the equation for the hyperbolic function (Gallistel, 1978; Mason & Milner, 1985) in Figure 1, mentioned earlier.

However, as stated earlier, not all researchers agree with Gallistel's (1974, 1978) conception of the neural integrator and how it performs temporal integration. Another group of researchers (Mason & Milner, 1985) attempted to fit Gallistel's (1978) single "leaky integrator" model to their data. In this study, rats implanted with electrodes in their LHAs were allowed to choose between rewards presented in a "Y-Maze" apparatus. In this paradigm, rats obtain rewarding stimulation by pressing once on a lever in either of two goal-arms. In one goal-arm the train duration was fixed, while in the other the train durations varied. Charge-duration curves at each behavioural criterion were obtained for each rat. Mason and Milner (1985) first plotted the duration of the test reinforcement against the charge required to make the rats choose the test reinforcement over the standard reinforcement on 50% of the trials. They discovered not only that at the longer train durations, the function departs from linearity (and thus, a hyperbolic function cannot be derived from it), but also that functions based on Gallistel's (1978) single leaky integrator did not fit their results. Instead, they suggested that there could be a second leaky integrator involved, one that has a longer time constant of decay ("leakiness"), and which acts on the first leaky integrator (Mason & Milner, 1985).

To summarize, there are two competing theories of the nature of the neural integrator that performs the aggregation of action potentials arising from BSR, as well as competing ideas of how this (these) integrator(s) performs temporal integration. One group (Gallistel, 1978) believes that only one "leaky" neural integrator is sufficient, and that the relationship between input strength and train duration required to maintain a certain criterion performance follows a hyperbolic function, while a second group (Mason & Milner, 1985) suggests that two "leaky" integrators with different "rates of leakiness" are necessary, and that the relationship between input strength and train duration required to maintain a particular criterion performance does not conform to a rectangular hyperbola.

These two research groups also differ on their conceptions of when "duration neglect" occurs in BSR, which is of interest in the current studies. Recall that when Dr. Kahneman tested his peak-and-end model in humans, he discovered that the subjects tended to neglect temporal information about the experience (i.e. showed duration neglect) (Kahneman et al., 1993). Thus, it would be especially interesting to test his model using BSR in a comparable situation, in which the rats would also pay more attention to the exemplars than the duration of the train. Does duration neglect exist in BSR? Returning to Gallistel's (1978) hyperbolic strength-duration function for a moment (Figure 1), one can see that at durations at and beyond the rheobase, one cannot tell the difference between the required strength needed for one train duration compared to other durations, as further increases in the train duration have no effect on the required strength, and thus duration can be said to be "neglected". Alternatively, it can be said that beyond the rheobase, further increases in train duration have no effect on the subjective reward value of the train (i.e. at that particular train duration and beyond, the rat performs the same regardless of changes in the strength of the stimulation).

At what point does the rat start to treat all train durations as if they were the same? Several experiments (Mark & Gallistel, 1993; Shizgal & Mathews, 1977) have demonstrated that once the duration of the stimulation trains used exceeded 1-2 seconds, further prolongation of the train did not appear to increase reward value. In one study (Shizgal & Mathews, 1977), rats that had been implanted with electrodes in the LHA were trained in a standard operant chamber to press one lever to turn stimulation on, and a second lever to turn stimulation off. In this experiment, the duration of "bursts" of stimulation was traded-off against the stimulation current intensity. Results showed that when the "burst width" became longer than 1 or 2 seconds, the current intensity required to maintain criterion performance levelled off, analogous to the rheobase described earlier in Gallistel's (1978) strength-duration function. In another study (Mark & Gallistel, 1993), rats with electrodes aimed at various points of the MFB were trained to press two levers, each of which had a different variable-interval (VI) schedule of reinforcement. One lever was associated with stimulation consisting of a 1-second train whose reward value was kept constant, and with a VI that varied between sessions, while the stimulation on the other lever had similar characteristics but varied in train duration on each trial and its VI schedule was fixed. Thus, it was possible to examine at what point "equipreference" (the point at which the difference in VI schedules between the 2 levers was offset by the differences in the train durations presented) was demonstrated in each trial. The results showed that subjective reward magnitude grows as a function of stimulation train duration, up to train durations of about 1 second. Beyond train durations of 1 second, however, subjective reward value levels off. Thus, it would appear that duration neglect sets in once the stimulation trains exceed 1-2 seconds in duration.

The Mason and Milner (1986a, 1986b) research group have a different conception of where duration neglect occurs in BSR. Recall that they claim, as mentioned earlier, that there are two neural integrators involved in generating the reward signal. However, they also assert that the reward signal should "saturate" ("fatigue"). This is similar to the conception of "duration neglect", in that once the neural integrators are saturated or fatigued further increases in the train duration should not result in an increase in reward value. The exact particulars of how these neural integrators fatigue was examined in one study (Mason & Milner, 1986a) which used a Y-maze paradigm similar to the one described earlier. Rats were implanted with electrodes in the LHA, and asked to choose between a standard reinforcement and a test reinforcement. In the first part of the experiment, the duration of the standard reinforcement was fixed, and the duration of the test reinforcement was varied. The pulse frequencies on both levers were also fixed, at either a high (200 Hz) or a low (70 Hz) pulse frequency. The second part of the experiment was of similar design, but three pulse frequencies were used: 70, 200 or 400 Hz. The results of the experiments demonstrated that at a high pulse frequency (200 Hz), the reward value reached asymptote or "saturated" at about 2 seconds, while at a lower pulse frequency (70 Hz), it took longer than 5 seconds for the reward value to approach asymptote. When a very high (400 Hz) pulse frequency was used, it took train durations of only 0.5 to 1 seconds for the reward value to approach asymptote. A second experiment (Mason & Milner, 1986b) also appeared to confirm the dependency of "saturation" on pulse frequency. They used a similar paradigm, but the duration of the "standard" and "alternate" trains was always the same. Also, the first 2.5 seconds of the alternate reinforcement ("part A") always had the same stimulation characteristics as the standard reinforcement, while the remainder of the alternate reinforcement ("part B") sometimes varied in pulse frequency from part A. In the first part of the study, the pulse frequency of both alternate "part A" and the standard train was 100 Hz, while alternate part B was set at either 50, 100, or 200 Hz. In the second part of the study, the standard reinforcement and alternate part A were both set at a pulse frequency of 250 Hz, and the pulse frequency of alternate part B was set at either 125, 250, or 500 Hz. The results revealed that at a low pulse frequency (100 HZ), the rats chose a 100% increase in pulse frequency during part B of the train over either no change or a decrease in the frequency. They did not show this tendency when delivered a high pulse frequency (250 Hz) for part A. Thus, when stimulation is initially delivered at a low pulse frequency, the neural integrators are not saturated, and reward value has not yet reached its maximum, and thus the rats choose to increase the pulse frequency further. On the other hand, when the initial pulse frequency is high, the neural integrators are saturated quickly and reward value reaches its maximum, and thus it makes no difference to the rats whether further stimulation pulse frequencies increase, decrease, or remain the same.

To summarize, it appears that there are discrepant results when examining the issue of "duration neglect" in BSR. Gallistel and colleagues (Gallistel, 1978; Mark & Gallistel, 1993; Shizgal & Mathews, 1977) claim that reward value approaches asymptote between approximately 1 to 2 seconds, while Mason and Milner (1986a, 1986b) claim that the rate at which reward value levels off depends entirely upon how much stimulation is "poured" into the integrators. This posed a problem for the current studies. Recall that the Kahneman studies in humans mentioned earlier were carried out under conditions in which the subjects appeared to show neglect of duration. Thus, in order to test the peak-and-end model, not only would one require pulse frequencies that differ in reward value, so that one can create stimulation trains in which the "peak" and "end" differ strongly in reward value, but it would also be optimal to have trains which fall in the range of duration neglect.

The first of the current studies attempted to resolve this controversy. Shizgal recently developed a mathematical model in which a single neural integrator would be able to account for Mason and Milner's (1986a, 1986b) findings, as well as other effects of manipulating train duration and pulse frequency. This model, called the "triple-logistic model", has three elements: 1) the relationship between the pulse frequency required to produce half-maximal reward value and train duration is a rectangular hyperbola, 2) the subjective intensity of the reward value grows as a sigmoid (s-shaped) function of the pulse frequency used, and 3) performance increases as a sigmoid function of the subjective intensity of the reward.

The first part of the model, shown in Figure 2, is Gallistel's (1978) strength-duration

Figure 2. Description of the equation describing Gallistel's (1978) strength-duration function. The C represents the chronaxie of the function for stimulation trains, in milliseconds (ms), and is an index of the "curviness" of the resulting hyperbola. The D represents the duration of the stimulation train, in seconds. The F_{hm} represents the stimulation frequency required to maintain half the maximum (50%) possible subjective reward value. The F_R represents the frequency at which the subjective reward value is half the maximum (50%) possible, at an infinitely long train duration. In other words, it is the pulse frequency at which further increases in the train duration have no effect on the subjective intensity of reward.

Strength-duration function

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

where C = chronaxie of strength - duration function for trains

D = train duration

F_{hm} = stimulation frequency required to produce half - maximal subjective intensity

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

Figure 2. Equation describing Gallistel's (1978) hyperbolic strength-duration function.

function. This is the first logistic equation in the Shizgal model, as a hyperbolic equation is a logistic with an exponent of "1". As mentioned earlier, the "C" or chronaxie term shown in the equation in Figure 2 is an index of the "curviness" of this hyperbola (over a given domain, the smaller the value of "C", the closer the hyperbola's shape is to a right-angled curve; the larger the value, the closer it is to a straight line). The chronaxie is the duration at which the required strength is twice the rheobase (Gallistel et al., 1981). The "D" in the equation is the train duration used, and the "F_R" is the lowest pulse frequency at which subjective reward value is at half-maximal intensity (i.e. half of the maximum value that the reward can possibly attain) at any train duration, and is also referred to as the "rheobasic frequency". Lastly, the "F_{hm}" refers to the stimulation frequency required to produce a reward value of half-maximal intensity at a given train duration. Thus, this first part of the triple-logistic model produces a curve similar to the one shown in Figure 1. At short train durations, a higher pulse frequency is required to maintain a given reward intensity, and as longer durations are used, this required pulse frequency decreases and levels off at an asymptote, the rheobase. Beyond this asymptote, further increases in train duration have no effect on the subjective intensity of the reward.

The second equation built into the Shizgal model is shown in Figure 3. This logistic, called the "reward-growth function", describes the relationship between pulse frequency and subjective intensity of reward, and was derived by Shizgal from experiments done by Gallistel, such as the "matching" (dual-lever) study described earlier (Mark & Gallistel, 1993). In this equation, the "F" refers to the pulse frequency used, the "F_{hm}" is the same parameter as that described in the strength-duration equation above, "g" is an exponent

Figure 3. Equation defining the relationship between pulse frequency and the subjective reward value, a sigmoidal function referred to as the Reward-Growth Function. The F represents the train stimulation pulse frequency. The F_{hm} represents the stimulation frequency required to maintain half the maximum (50%) possible subjective reward value. The g is an exponent that represents the growth of reward value. The I represents the subjective reward value, and the I_{max} represents the maximum possible subjective intensity of the reward.

Reward-Growth Function

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

F = stimulation frequency

F_{hm} = frequency at which subjective intensity is half-maximal

g = exponent of intensity growth

I = subjective intensity of reinforcement

$I_{\max}$ = maximum intensity of reinforcement

Figure 3. Equation defining the sigmoidal relationship between pulse frequency and the subjective intensity of reward.

which represents the rate at which reward intensity grows, "I" represents the reward intensity of the stimulation, and the "I_{max}" represents the maximum reward intensity of the stimulation. A graphical representation of what this equation describes is shown in Figure 4, and it can be seen to be a sigmoid or s-shaped function. In such functions, at very low or at very high values of the x-axis variable (the lower and upper asymptotes of the function), very little change in the value of the y-axis variable occurs, while at moderate values of the x-variable, there are large changes in the value of the y-axis variable (the rising slope of the function). In the sigmoid (logistic) function described here, the relationship described is between pulse frequency and subjective reward intensity; at very low pulse frequencies, since they are too low to support responding, there is little change in reward intensity, and the same occurs at very high pulse frequencies, since the neural integrator is likely approaching saturation. At moderate pulse frequencies, however, there are changes in the reward intensity, as shown by the rising slope of the curve. In fact, it is the exponent in this equation (g) which determines the "steepness" of the rising portion of the sigmoid; in other words, the steepness at which the subjective intensity of reward grows.

When the F_{hm} in the reward-growth function (Equation 2) is replaced by Gallistel's strength-duration function (Equation 1), a combined equation is obtained, which is shown in Figure 5. This combined equation, called the "generalized reward-growth equation", represents the growth of reward intensity of BSR as a function of both the train duration and the pulse frequency. This combined equation is the most important element of the triple-logistic logistic model, and since it takes both train duration and pulse frequency into

Figure 4. Representative sigmoidal relationship. The green curve represents a hypothetical sigmoid. Plotted on the x-axis is the common logarithm of the simulation pulse frequency (pulses per second), and plotted on the y-axis is the proportion of rewards harvested (maximum harvest = 20 reinforcements).

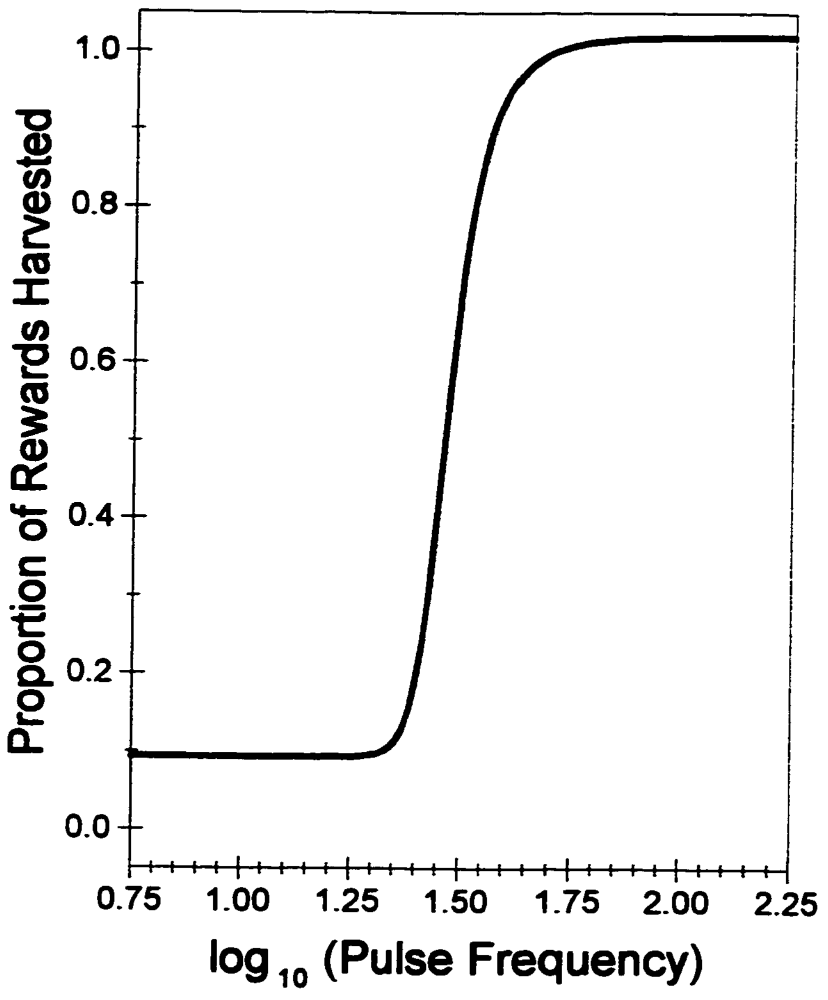

Figure 4. Sigmoidal relationship between pulse frequency and reward value.

Figure 5. Equation describing the growth in the subjective intensity of reward as a function of both the stimulation train duration and the pulse frequency, referred to as the Generalized Reward-Growth Function. The C represents the chronaxie of the function for stimulation trains, in milliseconds (ms). The D represents the stimulation train duration, in seconds. The F represents the train stimulation pulse frequency. The F_R represents the frequency at which the subjective reward value is half the maximum (50%) possible, at an infinitely long train duration; the pulse frequency at which further increases in the train duration have no effect on the subjective intensity of reward. The g is an exponent that represents the growth of reward value. The I represents the subjective reward value, and the I_{max} represents the maximum possible subjective intensity of the reward.

Generalized Reward-Growth Equation

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

where $C =$ chronaxie of strength - duration function for trains

$D =$ train duration

$F =$ stimulation frequency

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

$g =$ exponent of intensity growth

$I =$ subjective intensity of reinforcement

$I_{max} =$ maximum intensity of reinforcement

Figure 5. Equation describing the growth of the subjective reward intensity as a function of both the pulse frequency and the train duration.

account, it was hypothesized that this model could account for the discrepancy between Gallistel's (1978) and Mason and Milner's (1986a, 1986b) findings.

One final element of the triple-logistic model remains. Equations 1 and 2 of this model deal with subjective intensity of reward, i.e. processes in the brain of the subject, but they do not speak to the issue of performance. Since the goal of the current studies was to fit the triple-logistic model to the behaviour of the subjects, a third equation was required in order to bridge the gap between subjective reward value in the brain of the animal and the performance (lever-pressing) that the animal would display. This equation is shown in Figure 6, and relates the subjective intensity of the reward to performance. This equation, called the "performance function", is a modified version of Herrnstein's matching law (Herrnstein, 1970), which was originally designed to relate rates of responding for a reinforcer to the rate or magnitude of the reinforcement. In the equation, "I" represents the subjective intensity of reward, "p" is an exponent that refers to the rate at which performance increases, "R" is the rate of reinforcement (i.e. the rate at which the rat "harvests" the reinforcements), "U_e" is the utility of everything else in the environment (i.e. the utility to the rat of engaging in other behaviours unrelated to lever-pressing, such as grooming, sniffing, exploring the experimental chamber, and so on), and "PRH" is the proportion of rewards harvested (i.e. the amount of rewards actually harvested by the rat, divided by the maximum number of rewards that could have been harvested). The performance function, like the reward-growth function, is also a sigmoid (Figure 4.). At very low reward intensities, the reward value is too weak to motivate the rat to bar-press, and thus there is little, if any, change in performance. At very high reward intensities, the rat reaches

Figure 6. The equation describing the sigmoidal relationship between observed performance (lever-pressing behaviour) and the subjective reward intensity, referred to as the Performance Model. The I represents the subjective intensity of the reward, the p is an exponent representing the increase in performance, and the PRH_{RS} represents the proportion of rewards harvested (out of a maximum of 20), rescaled to remove heteroscedasticity. The R represents the rate of reinforcement (reinforcement schedule), and the U_e represents the "utility of everything else", i.e. the likelihood that the rat will engage in behaviour other than lever-pressing, such as grooming, scratching, exploring, sleeping, and so on.

Performance model

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

where I = intensity of reinforcement

p = performance exponent

PRH_{RS} = proportion of rewards harvested (rescaled)

R = rate of reinforcement

and U_e = utility of "everything else"

Figure 6. Equation describing the sigmoidal relationship between performance and subjective reward intensity.

a behavioural ceiling and harvests all the available rewards and therefore again there is no change in performance. At moderate reward intensities, on the other hand, lever-pressing performance does change as the reward intensity of the stimulation is increased. As with Equation 2 described earlier, the exponent (p) determines the steepness of the sigmoid's rising portion. In this equation, however, it refers to the steepness at which performance grows.

The full single-integrator model is obtained by substituting the "I" in the performance function (Equation 4) with the generalized reward-growth equation (Equation 3). The complete model is shown in Figure 7. Most of the terms have the same meaning as in previous equations, but one new term has been introduced: "R_e" represents the reinforcement rate at which the utility of a maximal reward equals the utility of everything else (i.e. the reinforcement rate at which the rat divides his time equally between bar-pressing and other (non-bar-pressing) activities), replacing the "U_e" described earlier. If this model is plotted on a three-dimensional graph, as shown in Figure 8, a "mountain" is obtained. This mountain shows the predicted proportions of rewards the rat will harvest for any given combination of train duration and pulse frequency.

Thus, one aim of the first two studies was to test the triple-logistic model. A secondary aim of these studies was also to use this model, if validated, to resolve the duration neglect controversy described earlier. The first study was designed to replicate Gallistel's (1978) findings, by holding the train duration constant during each experimental session, and varying the pulse frequencies of the trains within each session. This procedure

Figure 7. The equation describing the 3-dimensional relationship between pulse frequency, train duration, and observed performance (lever pressing), referred to as the Full Temporal Integration Model or the Shizgal Triple-Logistic Model (the hyperbola and two sigmoids embedded within the equation are all logistic functions, a hyperbola being a sigmoid that is raised to an exponent of 1). In the full equation, the C represents the chronaxie of the function for stimulation trains, in milliseconds (ms), the D represents the train duration in seconds, the F represents the pulse frequency, and the F_R represents the frequency at which further increases in the train duration have no effect on the subjective intensity of reward. The g is an exponent that represents the growth of reward value, while the p is an exponent that represents the growth of performance. the PRH_{RS} represents the proportion of rewards harvested (out of a maximum of 20), rescaled to remove heteroscedasticity. The R represents the rate of reinforcement (reinforcement schedule), and a new term, R_e, represents the reinforcement rate at which the utility of a maximally-rewarding BSR train equals the utility of engaging in behaviours other than self-stimulation.

Full temporal integration model

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

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"

Figure 7. Equation describing the 3-dimensional relationship between pulse frequency, train duration, and observed performance.

Figure 8. Representative 3-dimensional prediction (black grid) of the Shizgal triple-logistic model. Plotted on the x-axis is the stimulation train duration, in seconds. Plotted on the y-axis is the pulse frequency (pulses per second). Lastly, plotted on the z-axis is the proportion of rewards harvested, out of a maximum of 20.

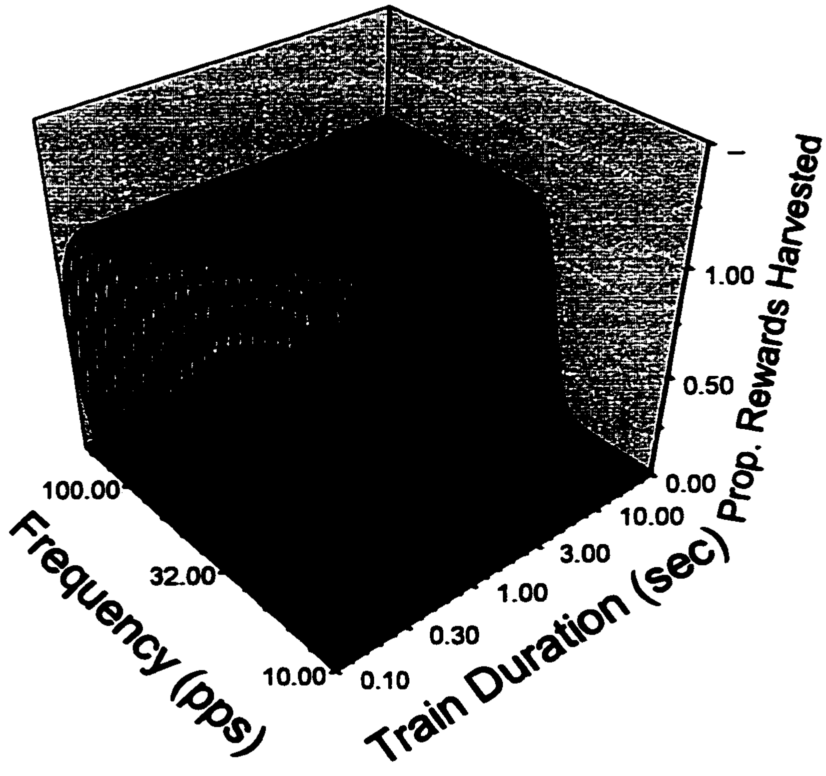

Figure 8. The 3D model of the relationship between pulse frequency, train duration, and performance.

allowed us to take the pulse frequencies required to reach the performance criterion (the point at which 50% of the rewards have been harvested) for each train duration, and plot them against the train duration, similar to the strength-duration function described earlier (Gallistel, 1978). The second study was designed to assess Mason and Milner's (1986a, 1986b) contentions, by holding the pulse frequency constant during each experimental session, and varying the train durations used within each session. This allowed us to examine whether or not duration neglect ("saturation") sets in sooner with high-frequency rather than low-frequency trains. The results from these two studies were then used to select particular stimulation frequencies and train durations for the third study, which was designed to pit Kahneman's (1993) peak-and-end model against the peak model (Norman & Gallistel, 1978). More specifically, the results of the first two studies were used to select stimulation frequencies which were perceived as having differing reward values by the rats, yet which also fell within the range of frequencies demonstrated to show duration neglect.

The results of the first two studies were important for the testing of the peak-and-end model in the third study, because the experiments conducted by the Kahneman group have typically been carried out under conditions in which the subjects appeared to show neglect for the duration of the aversive experience. Thus, in the current studies, it was of interest to see if parallel conditions could be obtained in a BSR study, in which the subjects would also show duration neglect, and the Shizgal triple-logistic model was used in an attempt to ascertain where duration neglect would set in. Also, as there had to be a difference in reward value between the "peaks" and "ends" of the stimulation trains in order to test the Kahneman model, the triple-logistic model also allowed for an identification of the appropriate frequencies to use to generate such "peaks" and "ends".

In the peak-and-end study, the third and last of the present studies, two types of stimulation trains were used. These trains can be thought of as being analogous to the "short" and "long" trial types described in the Kahneman (1993) experiment. One type of stimulation, the constant-frequency trains, never varied in pulse frequency, and thus the "peak" and "end" reward values occurred at the end of the train and were equal. In the second type of stimulation train, composite trains, the pulse frequency was decreased during the last seconds of the train, allowing the "end" to be less rewarding than the train's "peak", and lowering the overall reward value of the train. This allowed for a test of the earlier "peak" model against the "peak-and-end" model, in order to assess which model can best account for the evaluation of reward.

To summarize, the current series of experiments was designed to examine how temporally extended rewarding experiences are evaluated, and how such evaluations affect performance. The first two studies were intended to resolve a controversy in the literature relating to the point at which "duration neglect" sets in, as well as a validation of the Shizgal triple-logistic model. The last study was designed to pit the "peak" model (Norman & Gallistel, 1978) against the "peak-and-end" model (Kahneman et al., 1993).

Method

Subjects

The subjects were 2 experimentally-naïve male Long-Evans rats from the Charles River breeding farms (St-Constant, Quebec). They weighed approximately 500g at the start of the experiment. They were housed individually in plastic cages, in a room with a 12-12 dark/light cycle (lights off between 7:00 a.m. and 7:00 p.m.). Food and water were available *ad lib*.

Surgery

Under sodium pentobarbital anaesthesia (Somnotol, 65 mg/kg), electrodes were implanted bilaterally in the lateral hypothalamic area (LHA) using the following 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. The electrodes were made from 0.25 mm stainless-steel insect pins insulated with Formvar. The bottom 0.5 mm of the tips of the electrodes were bared of insulation. The return for the current consisted of two stainless steel screws fixed in the skull, around which the return wire (the ground) was wrapped. The electrodes were permanently anchored to the cranium with dental cement and jeweller's screws. Subjects were allowed to recover for 3 days before training commenced.

Apparatus

Screening and training. Single-lever operant boxes were used in the initial screening and testing of the rats. The boxes were constructed out of wood (27.4 cm by 26 cm by 64 cm) with Plexiglas fronts and one non-retractible lever (made for the lab by Concordia technicians) positioned 6 cm above the wire mesh flooring. In addition, a keylight was positioned 4 cm above the lever. This keylight was illuminated whenever stimulation was available. All temporal characteristics of the stimulation were set by digital pulse generators, and the intensity (μamps) was regulated by constant-current amplifiers. In order to prevent a build-up of charge at the interface of brain and electrode, when no pulse was present the charge was shunted through a 1 k Ω resistor. The stimulation current was monitored using a Metermaster MM200 oscilloscope, by reading the voltage drop across a 1 k Ω resistor (1% precision) in series with the electrode. Stimulation consisted of trains of cathodal pulses, with the pulse duration set at 0.1 ms, and the intensity set at 400 μamps. To allow the rat to move freely within the cage without tangling or twisting the lead, phone handset swivel jacks (Archer "Untangler", model no. 279-299) were used as "plugs" for the leads at the top of the operant chamber.

Experiments. Dual-lever operant boxes were used in all stages of the experiments. The boxes 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 (33 cm by 23.5 cm by 60.5 cm) with two retractible levers (MED Associates Inc., ENV-112B), positioned 10 cm above the mesh floor. The two levers were positioned across from each other, and only one lever was used in the course of the experiments; the other lever was left retracted. There were keylights positioned 6 cm above each lever (although only the light above the active lever was used in the course of the experiments), and a houselight was positioned on the rear wall, 38 cm above the floor. The keylight above the active lever was illuminated whenever the stimulation was available, and the houselight was flashed on and off throughout the 20-second inter-trial-intervals (ITIs), to signal that the next trial was imminent.

Stimulation currents were monitored using a Tektronix 2205 oscilloscope, by reading the voltage drop across a 1 k Ω resistor in series with the electrode. As with the screening boxes described above, the intensity of the stimulation was regulated by constant-current amplifiers, and in order to prevent the build-up of charge at the interface of brain and electrode, when no stimulation pulses were present, the charge was shunted through a 1 k Ω resistor. All temporal characteristics of the stimulation were set using a custom-developed software program written by Steve Cabilio.

Procedure

Screening and training. Before the beginning of the current studies, several rats were screened for the presence of aversive or motor effects, by stimulating them with 0.5-second trains of 40 cathodal pulses, set at an intensity of 400 μamps. If no motor or aversive effects were observed, the rats (two remained after initial screening) were next shaped to press the lever to obtain brain stimulation; trains were delivered to the rat upon successively closer approximations to pressing the lever (for example, facing the lever, approaching the lever, putting their paw on the lever), until the animals were self-stimulating reliably. The stimulation consisted of 0.5 second trains set at 400 μamps, and generally in the range of 60 to 40 cathodal pulses (frequency range 118 to 78 Hz), but the stimulation intensity and pulse number used were dependent upon the individual rat. The pulse duration was always set at 0.1 milliseconds.

Next, the number of pulses in each train was "swept"; the parameters of the stimulation at the beginnings of the sessions were set to values used during initial screening (i.e. stimulation at which the rat was responding vigorously), and then on each subsequent trial the number of pulses in the train was decreased by approximately 0.05 $\log_{10}$ units, until the stimulation was insufficient to support responding. Training in this phase continued for several days.

Lastly, the rats were screened to ensure that they were able to tolerate stimulation trains as long as 8 seconds in duration. The stimulation was initially set to the same parameters used during training and pulse number sweeps, but over successive trials the duration of the train was made progressively longer (train durations sampled were 1, 2, 4, and 8 seconds), with the appropriate pulse frequency for each train duration adjusted in accordance with predictions based on Gallistel's (1978) strength-duration experiments, and the stimulation intensity held constant.

Baseline condition. During this phase of the experiment, rats were moved to the dual-lever chambers described above. The rats were trained to lever-press for 0.5-second trains identical to the ones used in the pulse sweeps described above, to encourage them to generalize their lever-pressing behaviour to the novel operant box. Once their bar-pressing had stabilized (i.e. less than 0.05 $\log_{10}$ unit shifts in the performance curves from session to session), a 5-second "blackout delay" (BD) was instituted: once the lever was pressed, it would retract into the wall of the operant chamber and would only become available again once this BD had elapsed. The pulse number (frequency) was swept during the experimental session; on each subsequent trial the pulse frequency was reduced by approximately 0.05 log units. All other stimulation parameters were held constant. One "priming" stimulation train was delivered during the inter-trial-interval (ITI), ten seconds after the start of the ITI. The ITI was always 20 seconds in duration, and commenced immediately after the end of the 100-second trial. The delivery of the priming train was not contingent upon an operant response by the subject. The intensity, pulse frequency, and duration of the priming stimulation were always identical to that of the stimulation delivered during the trial that immediately followed. In all, twelve trials were run during each session (1 "sweep"). This condition was used throughout the experiment, to monitor possible large-scale shifts or instabilities in the rats' performance over time, and was run each day before the appropriate experimental session was run.

Pulse frequency sweeps experiment. As with the baseline condition, each individual session consisted of a series of pulse frequency sweeps, in which each successive trial contained approximately 0.05 log units fewer pulses than the preceding trial. Again, all other stimulation parameters were held constant, and one priming train was delivered to the subject during the 20-second ITI, which commenced immediately after the end of each trial, the trial duration being 320 seconds. However, in this experiment the blackout delay (BD) was extended to 16 seconds; once the lever was triggered, it would retract into the wall of the operant chamber and would only become available again after 16 seconds had elapsed. The BD served two purposes: first, it prevented the rapid build-up of aversive effects during long stimulation trains by slowing down the rate at which the rats could obtain the stimulation, and second, it controlled for reinforcement rate effects. If the BD had not been used, and the lever had become available again immediately after the stimulation train ended, the rats would have been able to obtain more rewards during short train duration trials than during long train duration trials. Therefore, through the use of a 16-second BD, the rats were limited to the same maximum number of rewards on every trial, regardless of the train duration used. This eliminated the confounding effects of differential reinforcement rates due to differences in stimulation train duration length and the availability of the lever. The duration of 16 seconds was selected for the BD in order to allow for some time for the decrease of aversive build-up at even the longest train duration (8 seconds), without making the experimental sessions unnecessarily lengthy.

In addition, in this phase of the experiment the rats were exposed to six different train durations. The rats were initially exposed to a 0.5-second train to acclimate them to this experimental condition, and then the different train durations were presented in descending order (i.e. 8, 4, 2, 1, 0.5, and 0.25 seconds). The rats were required to meet a stability criterion with one train duration before being exposed to a subsequent train duration. Each complete "sweep" consisted of thirteen trials and three sweeps were run during each experimental session, thus each subject underwent 39 trials in one session.

Train duration sweeps experiment. Experimental sessions in this phase were similar to those in the preceding experiment, except that each session consisted of a series of train duration sweeps at a constant pulse frequency. All other stimulation parameters were also held constant within each session. As with the previous experiment, the BD was set at 16 seconds. The pulse frequencies were set throughout the experimental sessions at either a "high" or a "low" pulse frequency, with the actual values of these frequencies determined by the individual rats' performance in the previous experiment. The rats were required to meet a stability criterion with the first pulse frequency (the "high" frequency) before being exposed to the second pulse frequency (the "low" frequency). During each experimental session, the train durations were swept in descending order (i.e. 8, 4, 2, 1, 0.5 and 0.25 seconds). During each experimental session there were ten trials in each sweep, and three sweeps were run (30 trials in total). Also, a single priming train was delivered during the ITI, and as in the previous experiment the ITI commenced immediately following each trial, with each trial lasting 320 seconds.

Peak-and-end test. This phase of the experiment was similar to the pulse frequency and train duration sweeps experiments described earlier, except the rats were exposed to five stimulation trains: a constant-frequency 8 second train, a constant-frequency 6 second train, and three composite trains in which the first 6 seconds were set at a constant pulse frequency, and the last 2 seconds were at a lower constant pulse 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. In addition, as with the previous experiments, the BD was set at 16 seconds, and each session consisted of a series of pulse sweeps (made up of the appropriate constant or composite pulse frequencies), with all other stimulation parameters held constant. As before, the rats were required to reach a stability criterion with each of the 5 train types before being exposed to the next train type. Again, one priming train was delivered to the rat during the ITI. As before, the ITI began immediately after the termination of the 320-second trial. The first train type presented to the rats was the 8-second train, followed by the 6-second train, and then by the 4:1, 2:1, and 1.4:1 composite trains, respectively. Each session consisted of thirteen trials per sweep, with three sweeps being run during each experimental session, for a total of 39 trials per session.

Results

For each of the three studies described, the first trial in each sweep was considered to be a "warm-up" trial, and was not included in subsequent analyses. Also, in order to obtain as accurate a picture as possible of the subjects' behaviour given normal variation between individual sessions, all analyses were carried out on data averaged over six "sweeps" (two experimental sessions, with three sweeps per session).

Frequency sweeps experiment

The raw data for animals B8 and B9 are graphed in semi-logarithmic space in Figure 9. The dependent measure on the y-axis is the proportion of rewards harvested (PRH), which has a maximum value of 1 (20 out of 20 rewards harvested). The pulse frequency on the x-axis has been transformed to logarithmic units. Upon examination of these performance curves, it is evident that while the upper asymptote (all available rewards harvested) was nearly always reached, the lower asymptote varied with the train duration. Also, there tended to be a greater amount of variability in the data points on the rising segment (slope) of all performance curves and at their lower asymptotes, compared to the data points at the upper asymptotes of the curves.

The non-zero value of the lower asymptotes and the deviation of these asymptotes with the train duration is beyond the scope of the Shizgal triple-logistic model. These features are not key elements of interest and are also likely to compromise curve-fitting, thus

Figure 9. Graphs of the raw means from the pulse frequency sweeps experiment, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the bottom x-axis is the common logarithm of the pulse frequency (pulses per second), and the top x-axis of each graph also displays the nonlogarithmic pulse frequency values (pulses per second). Plotted on the y-axis 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 tested (0.25, 0.5, 1, 2, 4, or 8 seconds). Each curve represents performance over 6 descending pulse "sweeps" at that particular train duration.

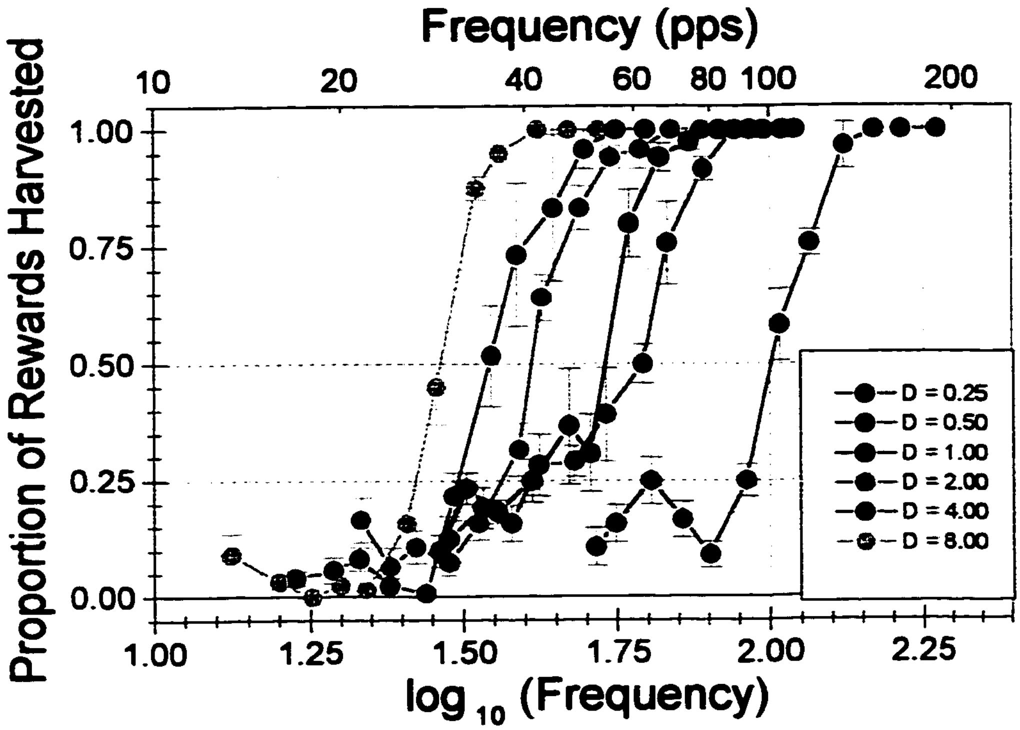

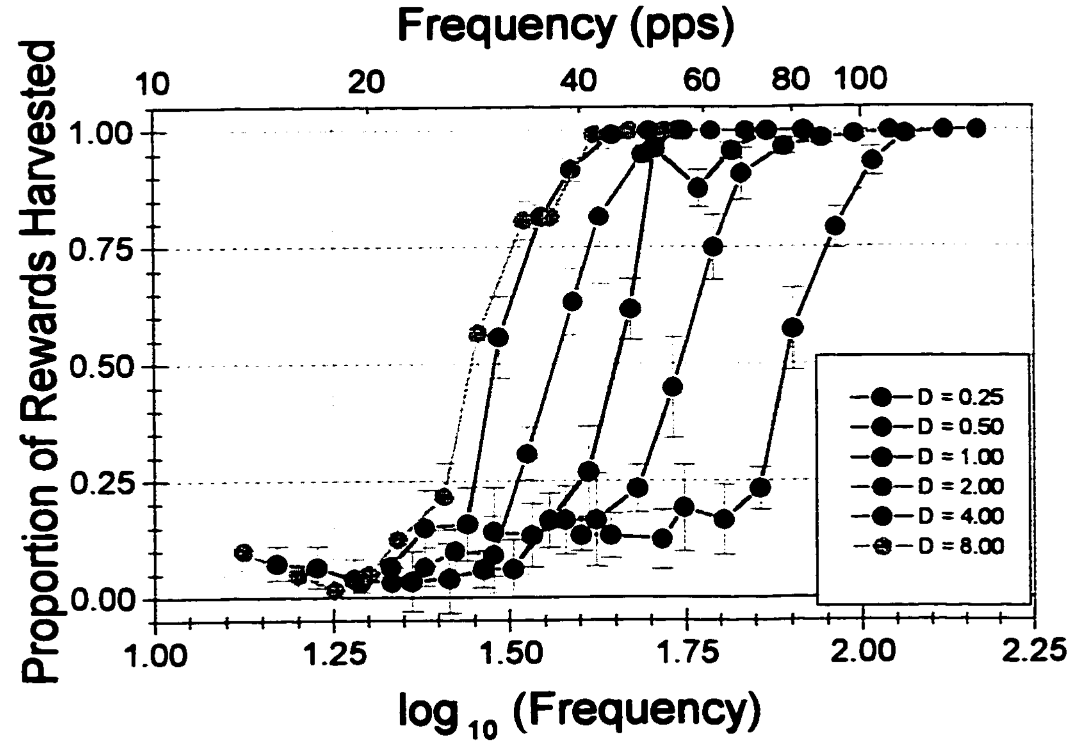

Figure 9. Raw data of frequency sweeps for rat B8 (top graph) and rat B9 (bottom graph).

each curve was rescaled to remove this variability. The first step was to estimate a new lower asymptote using a "broken-line function". In this function, the best-fitting line is plotted in three line segments: one lower horizontal line, one upper horizontal line, and one rising segment linking the two horizontal segments, and this function gives rise to the type of fit shown in Figure 10. Next, the estimated lower asymptote was set to zero (no rewards harvested) and the data points were rescaled so that the means ranged between 0 and 1.

Also, to compensate for the heteroscedasticity of the performance curves, and thus give more weight to the comparatively stable points on top and less weight to the more unstable points on the risers and lower asymptotes, each individual observation (on each of the six pulse frequency sweeps) was weighted. The weights used were the inverse of each observation's variance, with the maximum weight set at 10.

The rescaled data for rats B8 and B9 are shown in Figure 11. For both subjects the six performance curves lie approximately parallel to each other, and the distance between each curve appears to decrease as the train is increased, although this is more evident for rat B9 than rat B8. Also, the 4-second (magenta) and 8-second (orange) curves lie very close to each other, suggesting that duration neglect is setting in within this range, although this is again more evident for rat B9.

In order to assess whether or not the performance curves fit Gallistel's (1978) strength-duration function, the first step in the analysis was to isolate the portion of the curve corresponding to half of maximum performance (0.5 of rewards harvested); this was the

Figure 10. A representative broken-line fit. The blue curve represents a hypothetical fit to the data points (open orange circles). The common logarithm of the pulse frequency (pulses per second) is plotted on the x-axis, and the "Raw Harvest" (out of 20) is plotted on the y-axis. As shown in the graph, a broken line function consists of a horizontal line fitted to the upper asymptote of the data points, a horizontal line fitted to the lower asymptote of the data points, and a "rising segment" (slope) that connects these two estimated asymptotes.

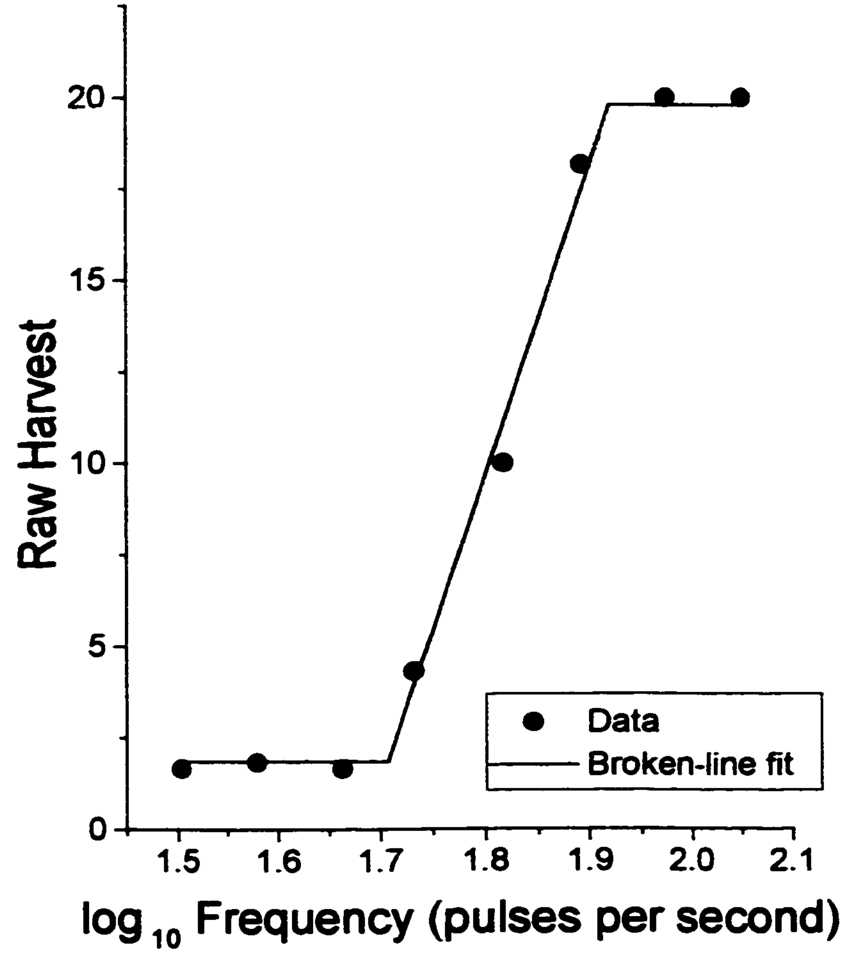

Figure 10. Example of a typical broken-line fit.

Figure 11. Graphs of the rescaled means from the pulse frequency sweeps experiment, for rat B8 (top graph) and rat B9 (bottom graph). The data were transformed to deal with problems with heteroscedasticity and the variation of the lower asymptote of each curve with the train duration. Plotted on the bottom x-axis is the common logarithm of the pulse frequency (pulses per second), and the top x-axis of each graph also displays the nonlogarithmic pulse frequency values (pulses per second). Plotted on the y-axis is the proportion of rewards harvested, out of a maximum of 20 (rescaled values). Error bars represent the standard error of the mean. Each curve represents a different train duration tested (0.25, 0.5, 1, 2, 4, or 8 seconds). Each curve represents performance over 6 descending pulse "sweeps" at that particular train duration.

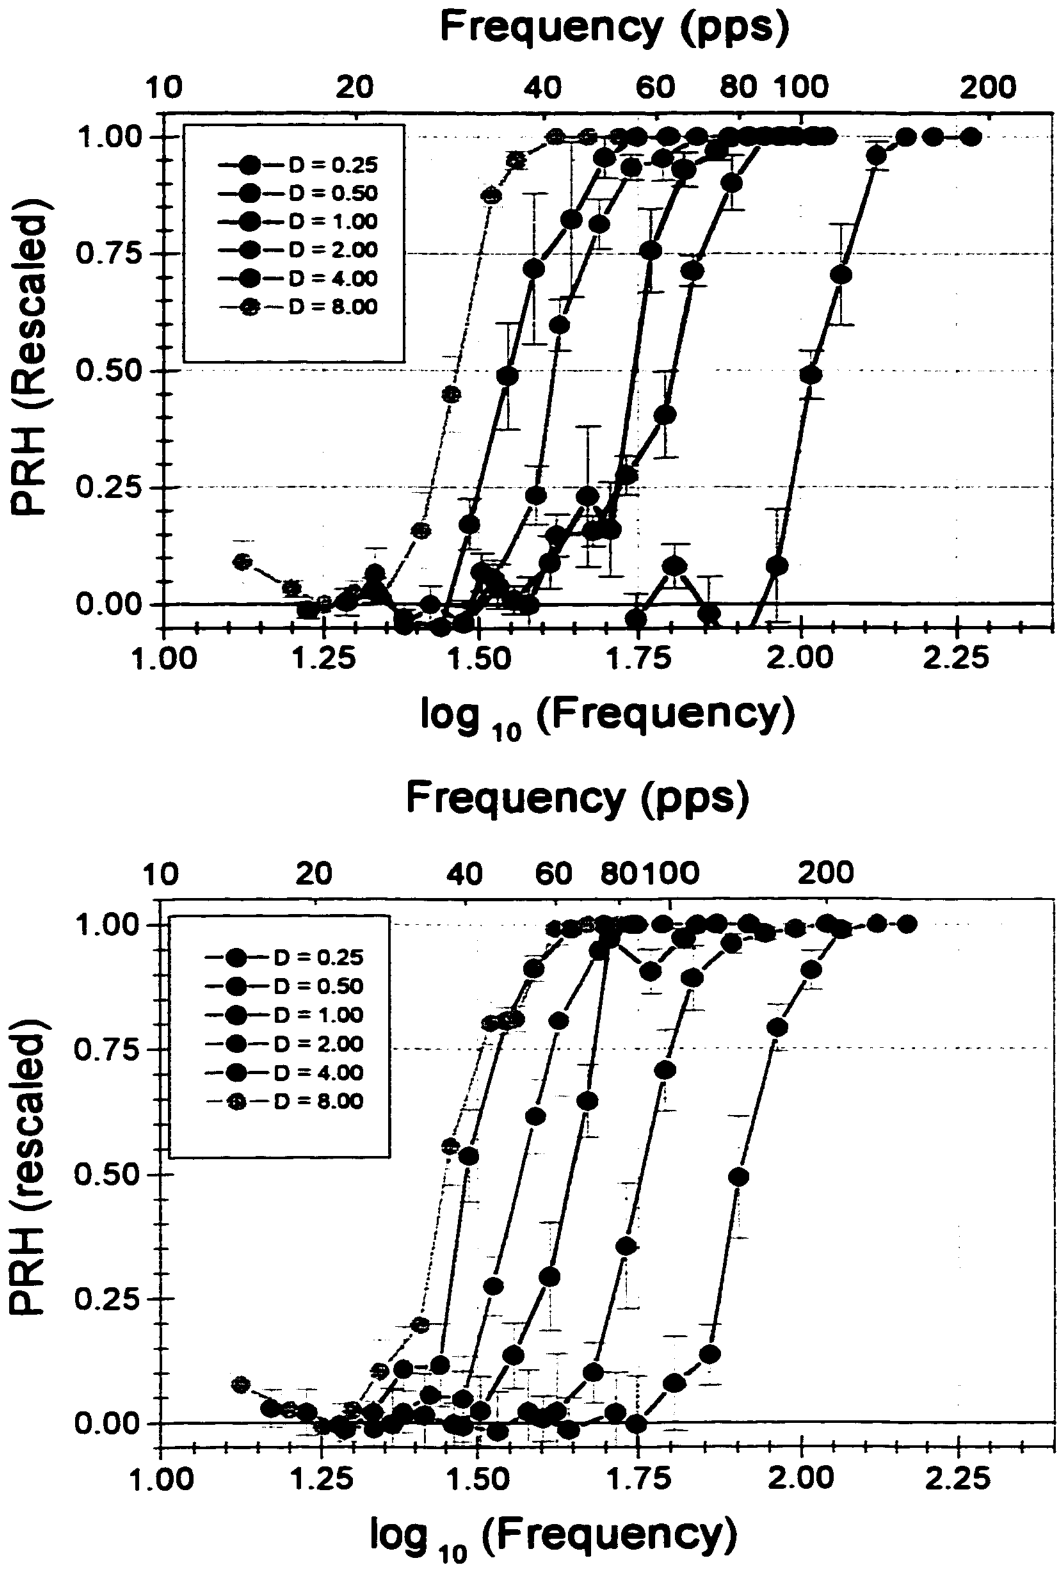

Figure 11. Rescaled data of frequency sweeps for rat B8 (top) and rat B9 (bottom).

criterion selected. Then the 2-dimensional (2-D) relationship between the train duration used and the pulse frequency required to obtain this 50% performance criterion ("required frequency") was examined by fitting a rectangular hyperbola to the common logarithms of these required frequencies. During this procedure, the curve-fitting procedure reported that the required frequency values were themselves heteroscedastic, particularly for rat B8, and thus these values were weighted using a procedure identical to that described above; the weights were the inverse of the variance, with a maximum weight of 10. Also, it was noted that for one of the animals (B8), two of the data points deviated significantly from the normal distribution of the remaining data points. These outliers were removed and not used in subsequent analyses.

The fit of Gallistel's (1978) hyperbolic strength-duration function to the required frequencies are shown in Figure 12 for rats B8 and B9. The red open circles show the actual required frequencies (six means for each train duration, one from each sweep), while the green line is the hyperbolic function derived from the rescaled data points. As can be seen from the graphs, the function does provide a good approximation of the data for rat B9, but deviates from the two longest durations (4 and 8 seconds) for rat B8. The coefficient of determination for rat B8 was 0.951, and 0.988 for rat B9.

Other variables of interest arising from this analysis were the estimated values of the rheobase and the chronaxie. Recall that the "curviness" of a rectangular hyperbola can be described by the x-axis value (chronaxie) at which the value of the function is twice the horizontal asymptote (the rheobase). The values of the rheobase were 32.82 and 27.06 for

Figure 12. Graph of the strength-duration function fitted to the rescaled common logarithms of the pulse frequency (pulses per second) required to maintain 50% of maximum lever-pressing performance, for rat B8 (top graph) and rat B9 (bottom graph). The stimulation train duration, in seconds, is plotted on the x-axis, and the common logarithms of the required pulse frequencies are plotted on the y-axis. The green curve represents the fitted function, and the red open circles are the rescaled means. Each circle represents performance during 1of 6 descending pulse "sweeps" at that particular train duration.

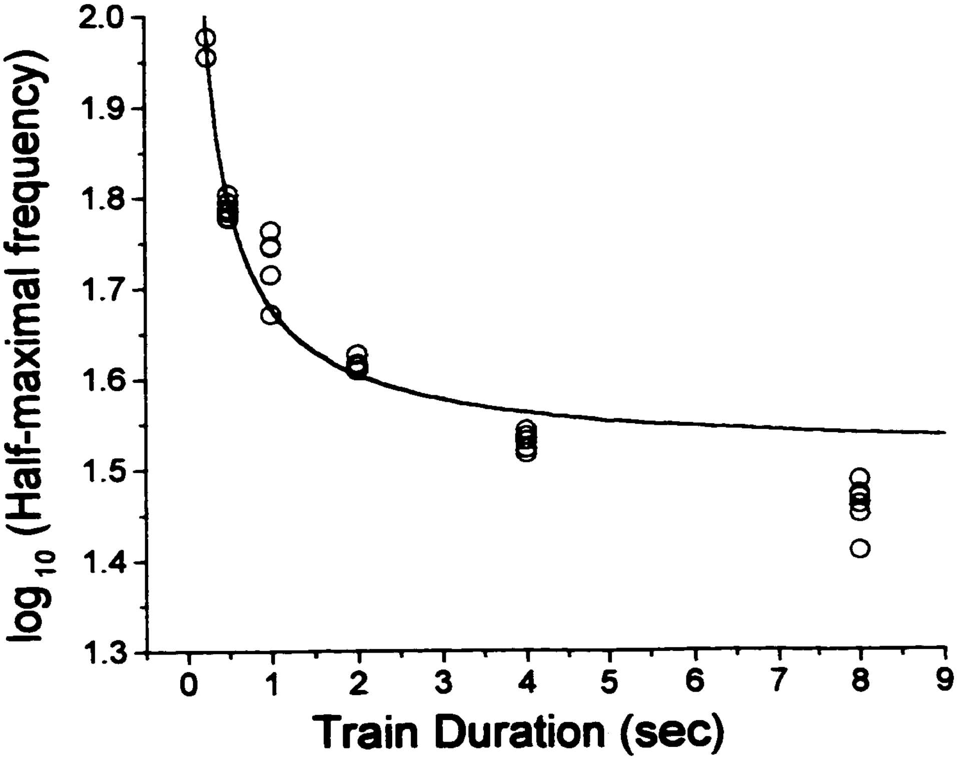

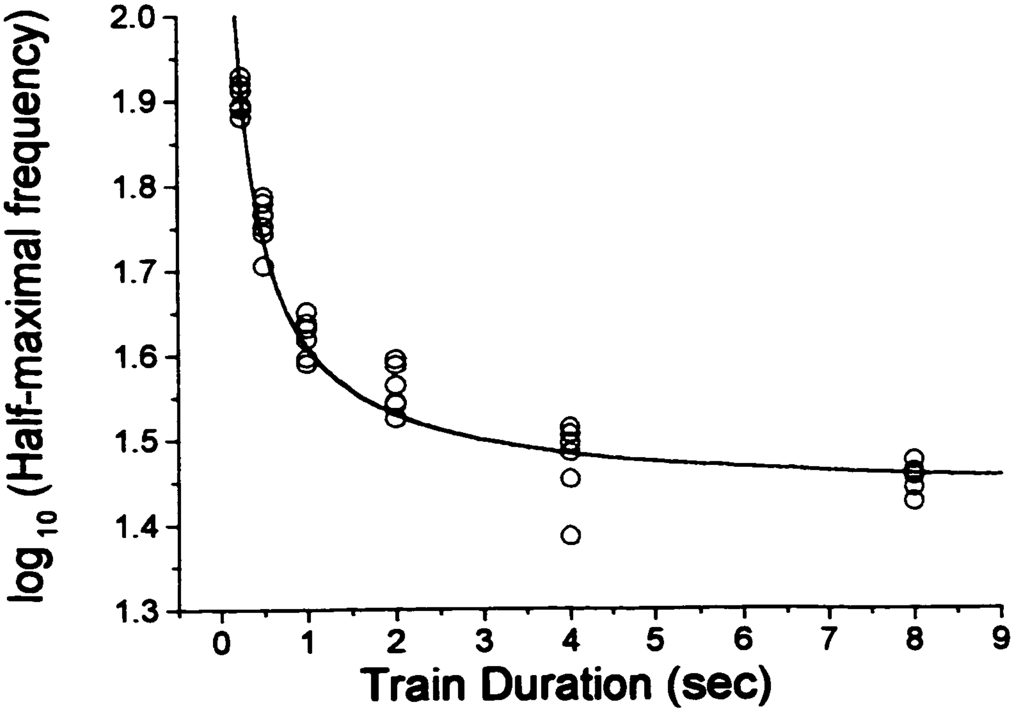

Figure 12. Fitted strength-duration function for rat B8 (top) and rat B9 (bottom).

rats B8 and B9, respectively, while the values of the chronaxie were 0.4533 and 0.503. The estimated chronaxies were used as constants in the 3-dimensional (3-D) fit that follows.

The last step in the analysis of the pulse frequency sweep experiment was to fit the Shizgal triple-logistic model to the transformed data. The logarithms of the frequencies and train durations served as the independent variables and the proportion of rewards harvested (PRH) served as the dependent variable. The chronaxies estimated using the 2-D strength-duration fits described earlier were set as constants in the 3-D fit to the Shizgal model. The estimated rheobase was not used as a constant, however, because while the rheobase is supposed to represent the frequency at which the reward value is half-maximal, the rheobase obtained from the 2-D fits were the frequencies at which *performance* was half-maximal. Thus, it was decided to leave the rheobasic parameter free to vary during the 3-D analysis. The weights used to remove heteroscedasticity when rescaling the data were also included as part of the 3-D fit. Furthermore, when using an iterative nonlinear curve-fitting procedure, one must supply an initial "guess" of each parameter's values. The initial values used in the analysis were 50.0 for the rheobase, 0.5 for RRR, and 5.0 for both g and p. It should be mentioned here that the "RRR" parameter replaces the ratio "R_e/R" used in the full temporal integration model described in the introduction. This was done because the "R" value, the rate of reinforcement, cannot be estimated in the current studies due to the fact that a continuous reinforcement paradigm was used, and also due to the use of the BD to control this effect as was mentioned earlier. Also, the "R_e" is a hidden parameter that cannot be directly assessed. Thus, the ratio of R_e/R was combined in one parameter in the current curve-fitting procedure, to yield a parameter called RRR, which represents the relative rate of reinforcement.

The results of fitting the 3-D Shizgal model to the rescaled data for rats B8 and B9 is shown in Figure 13. The common logarithms of the frequencies are shown on the y-axis, the common logarithms of the train durations are shown on the x-axis, and the PRH is plotted on the z-axis. These graphs show the function derived from the 3-D fit, shown by the black grid, and the actual performance curves, which are identical to the rescaled performance curves shown in 2-D space in Figure 11, are plotted on the same graph to show their comparison to the fitted function. The fitted surface is a good approximation of the data points, with the coefficient of determination 0.984 for rat B8, and 0.986 for rat B9.

Other variables of interest arising from the 3-D analysis were the estimated values of the relative rate of reinforcement (RRR), the growth exponent (g), the performance exponent (p), and the rheobase. Recall that in the Shizgal triple-logistic model "g" is an exponent which represents the rate at which reward intensity grows and "p" is an exponent that refers to the rate at which performance grows. Also, the "RRR" parameter used in the current analysis replaces the "R_o/R", as mentioned earlier. The rheobase here is identical in meaning to the rheobase estimated during the earlier 2-D fits. For B8, the values of R, g, p, and the rheobase were 0.327, 3.962, 7.088, and 37.014, respectively. For B9, the values of R, g, p, and the rheobase were 0.010, 6.089, 2.128, and 59.351, respectively.

However, it should be noted that the results of the 3-D fits for RRR, g, p and the rheobase were heavily influenced by the initial values used to start the fitting procedure.

Figure 13. Graph of the rescaled data points and prediction of the Shizgal triple-logistic model for the pulse frequency sweeps experiment for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis are the common logarithms of the stimulation train durations, in seconds. Plotted on the y-axis are the common logarithms of the pulse frequency (pulses per second). Plotted on the z-axis is the proportion of rewards harvested, out of a maximum of 20. The coloured spheres represent the pulse frequency sweeps (rescaled data points) at each train duration, with the following legend: red, 0.25 sec duration, green, 0.5 sec, blue, 1 sec, cyan, 2 sec, magenta, 4 sec, and orange, 8 sec. Each curve represents performance over 6 descending pulse "sweeps" at that particular train duration. The black grid represents the surface predicted by the Shizgal triple-logistic model after the chronaxie estimated from the 2-dimensional strength-duration fit is entered into the model as a constant, while all other parameters are allowed to vary.

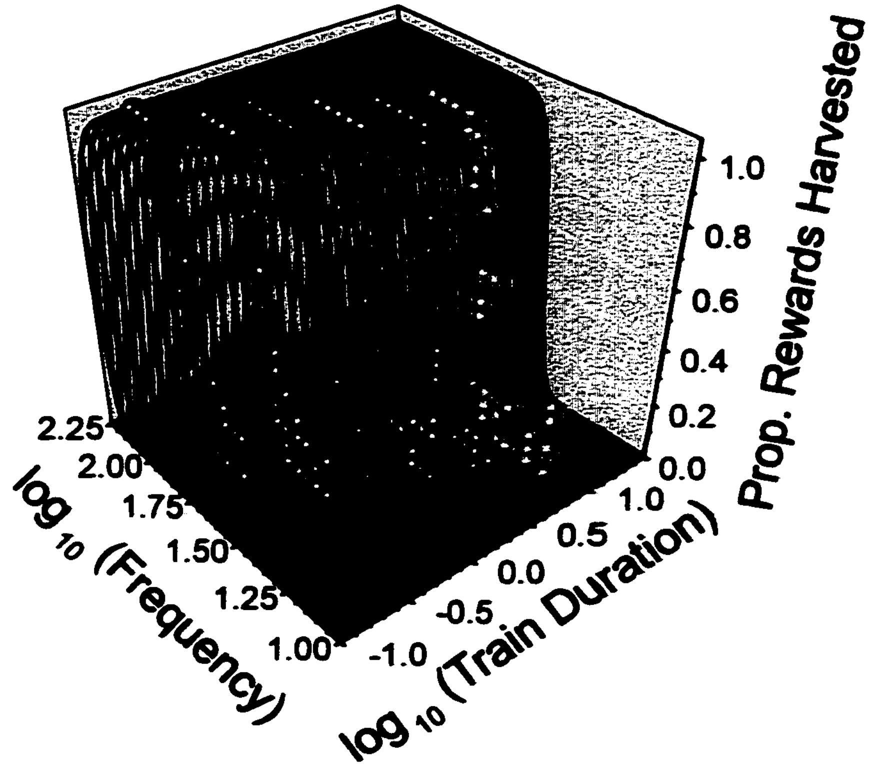

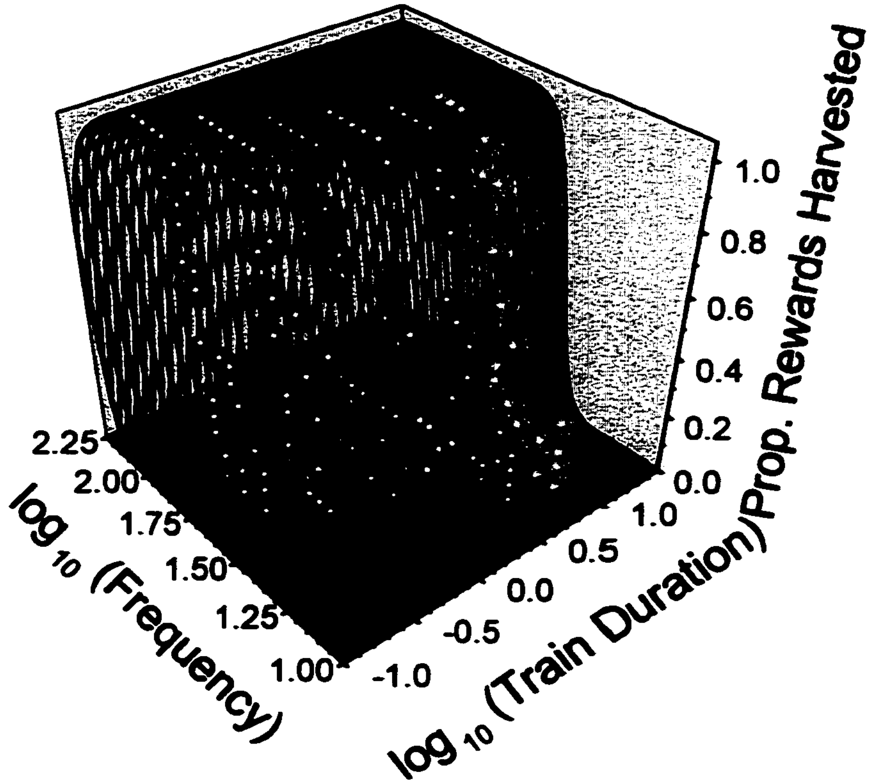

Figure 13. 3-D fit of rescaled frequency sweeps for rat B8 (top) and rat B9 (bottom).

This is a problem for the Shizgal model, although not for the accuracy of the results of the fitted surface (grid), and it will be discussed in later sections.

Train duration sweeps experiment

The raw data for rats B8 and B9 is graphed in Figure 14. The dependent measure on the y-axis is the proportion of rewards harvested (PRH), and train duration is plotted on the x-axis. As mentioned earlier, the actual values of the high and low pulse frequencies for each rat were determined by their individual results in the previous experiment. As with the frequency sweeps described above, the data was weighted and rescaled to deal with differential variation at different portions of the duration sweep, as well as the fact that the rats' performance never decreased to zero rewards harvested. However, since in some cases (rat B9) the performance curves never reached 1 (all available rewards harvested), and such "shortfalls" are predicted by the triple-logistic model (follow the line in the grid corresponding to a logarithmic frequency of 1.50, which is 31.6 Hz in non-logarithmic units on rat B9's 3-D fit in Figure 13), in such cases the rescaling procedure was modified to fix the upper asymptote where it was, and then stretch the bottom of the curve down to zero, in a manner analogous to the rescaling described above.

The rescaled data for rats B8 and B9 is shown in Figure 15. As can be seen from the graphs, at high pulse frequencies (black curve) performance approaches asymptote rapidly (by approximately 1 second for rat B8, and by 1.5 seconds for rat B9), while at the lower pulse frequency (gold curve) performance rises much more slowly, and in some cases (rat

Figure 14. Graphs of the raw means from the train duration sweeps experiment, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis is the train duration (in seconds). Plotted on the y-axis is the proportion of rewards harvested, out of a maximum of 20. Error bars represent the standard error of the mean. The black curves represent the high pulse frequencies tested (value determined individually for each rat), and the gold curves represent the lower pulse frequency tested (also determined individually for each rat). Each curve represents performance over 6 descending train duration "sweeps" at that particular pulse frequency.

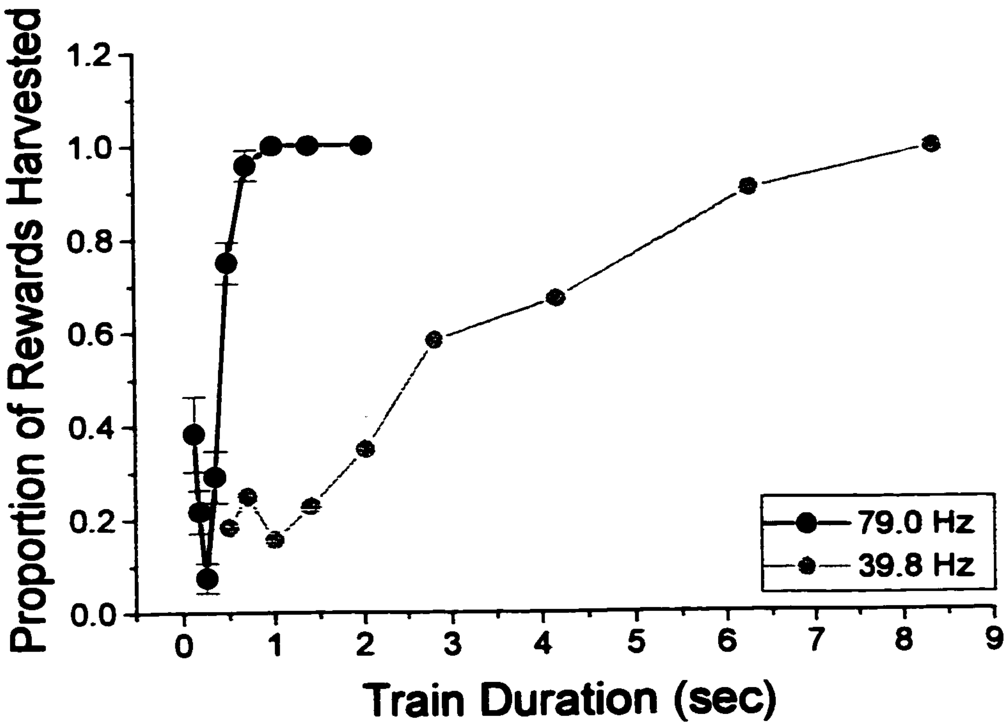

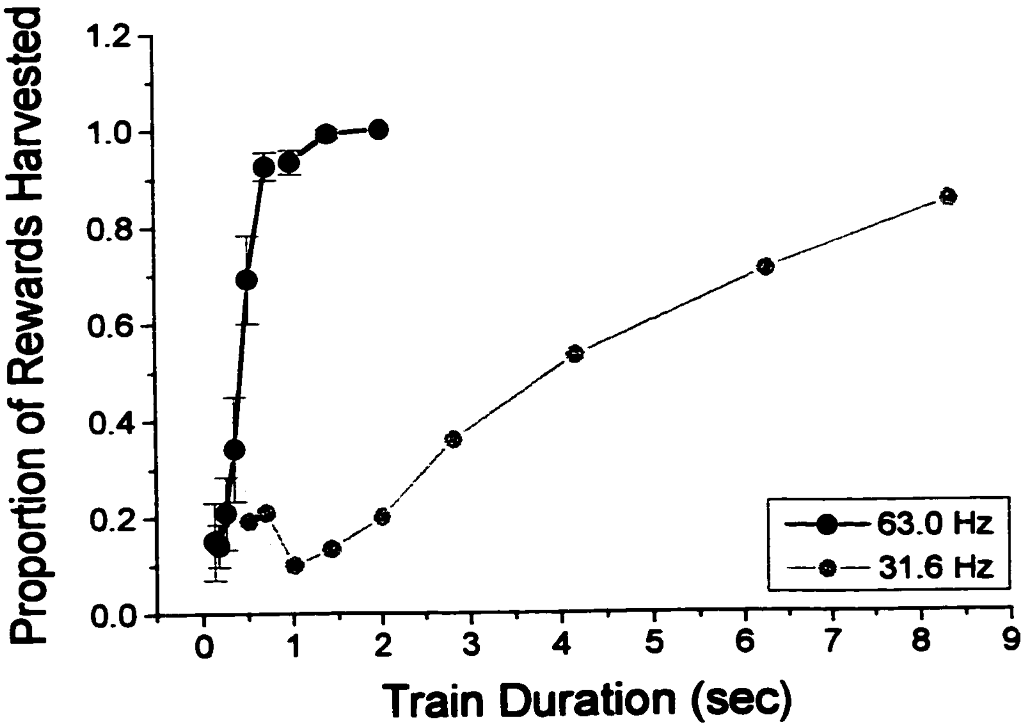

Figure 14. Raw data of train duration sweeps for rat B8 (top) and rat B9 (bottom).

Figure 15. Graphs of the rescaled means from the train duration sweeps experiment, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis are the common logarithms of the train durations (in seconds). Plotted on the y-axis is the proportion of rewards harvested, out of a maximum of 20. Error bars represent the standard error of the mean. The black curves represent the high pulse frequencies tested (value determined individually for each rat), and the gold curves represent the lower pulse frequency tested (also determined individually for each rat). Each curve represents performance over 6 descending train duration "sweeps" at that particular pulse frequency.

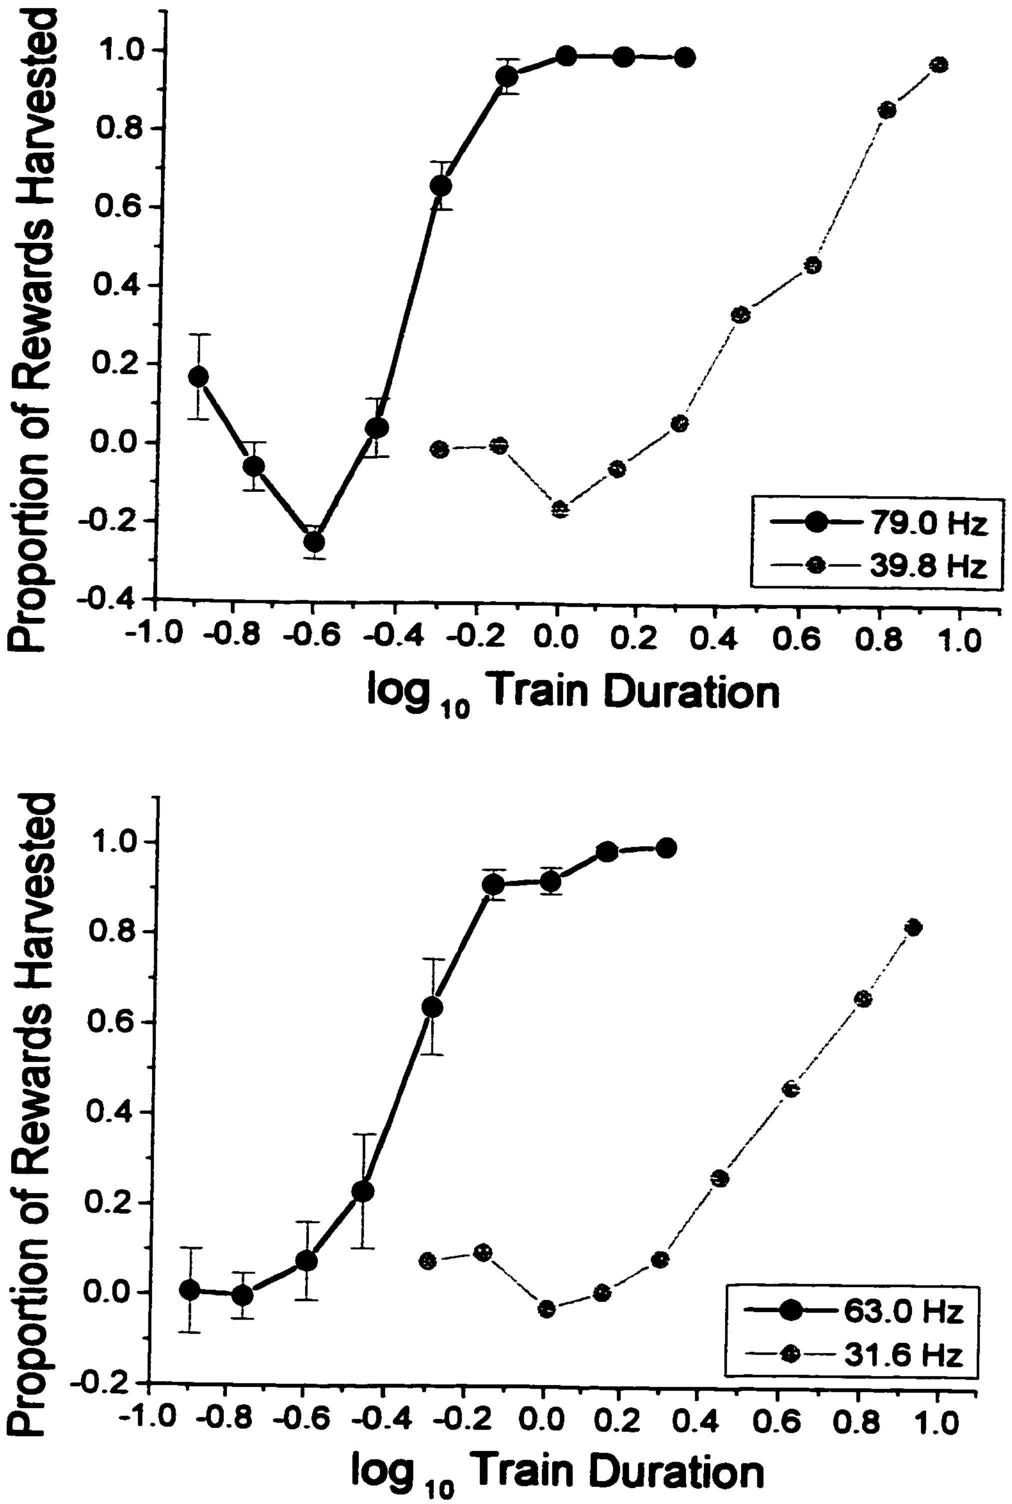

Figure 15. Rescaled data of train duration sweeps for rat B8 (top) and rat B9 (bottom).

B9) does not reach upper asymptote at the longest duration tested (8 sec).

The results of the duration sweeps were also assessed to determine whether they were consistent with the triple-logistic model. The logarithms of the frequencies and train durations again served as the independent variables and the proportion of rewards harvested (PRH) served as the dependent variable. The chronaxies estimated using the 2-D strength-duration fits were again set as constants in the 3-D fit, while the rheobase was again left free to vary. Again, the weights used to remove heteroscedasticity when rescaling the data were included in the fitting procedure.

The data sets of both the frequency and duration sweeps experiments are plotted in Figure 16 for rats B8 and B9. Again, the logarithms of the frequencies are shown on the y-axis, the logarithms of the train durations are shown on the x-axis, and the PRH is plotted on the z-axis. These graphs display the function derived from the new 3-D fit, shown by the black grid. Also, the rescaled data sets from both the frequency and duration sweeps (Figures 11 and 15) have been plotted in the same space for comparison. The correspondence between the fitted surface and the data points is again quite good, with the coefficient of determination 0.976 for rat B8, and 0.984 for rat B9. Thus, it appears that regardless of whether frequency sweeps or duration sweeps were conducted, a given combination of frequency and duration yields similar performance.

Figure 16. Graph of the rescaled data points and prediction of the Shizgal triple-logistic model for the pulse frequency sweeps and train duration sweeps experiments plotted in the same space, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis are the common logarithms of the stimulation train durations, in seconds. Plotted on the y-axis are the common logarithms of the pulse frequency (pulses per second). Plotted on the z-axis is the proportion of rewards harvested, out of a maximum of 20. The following coloured spheres represent the pulse frequency sweeps (rescaled data points) at each train duration: red, 0.25 sec duration, green, 0.5 sec, blue, 1 sec, cyan, 2 sec, magenta, 4 sec, and orange, 8 sec, while the following coloured spheres represent the train duration sweeps (rescaled data points) at each pulse frequency: black, the high pulse frequency used, and gold, the low pulse frequency used. Each curve represents performance over 6 descending "sweeps". The black grid represents the surface predicted by the Shizgal triple-logistic model after the chronaxie estimated from the 2-dimensional strength-duration fit is entered into the model as a constant, and all other parameters are allowed to vary.

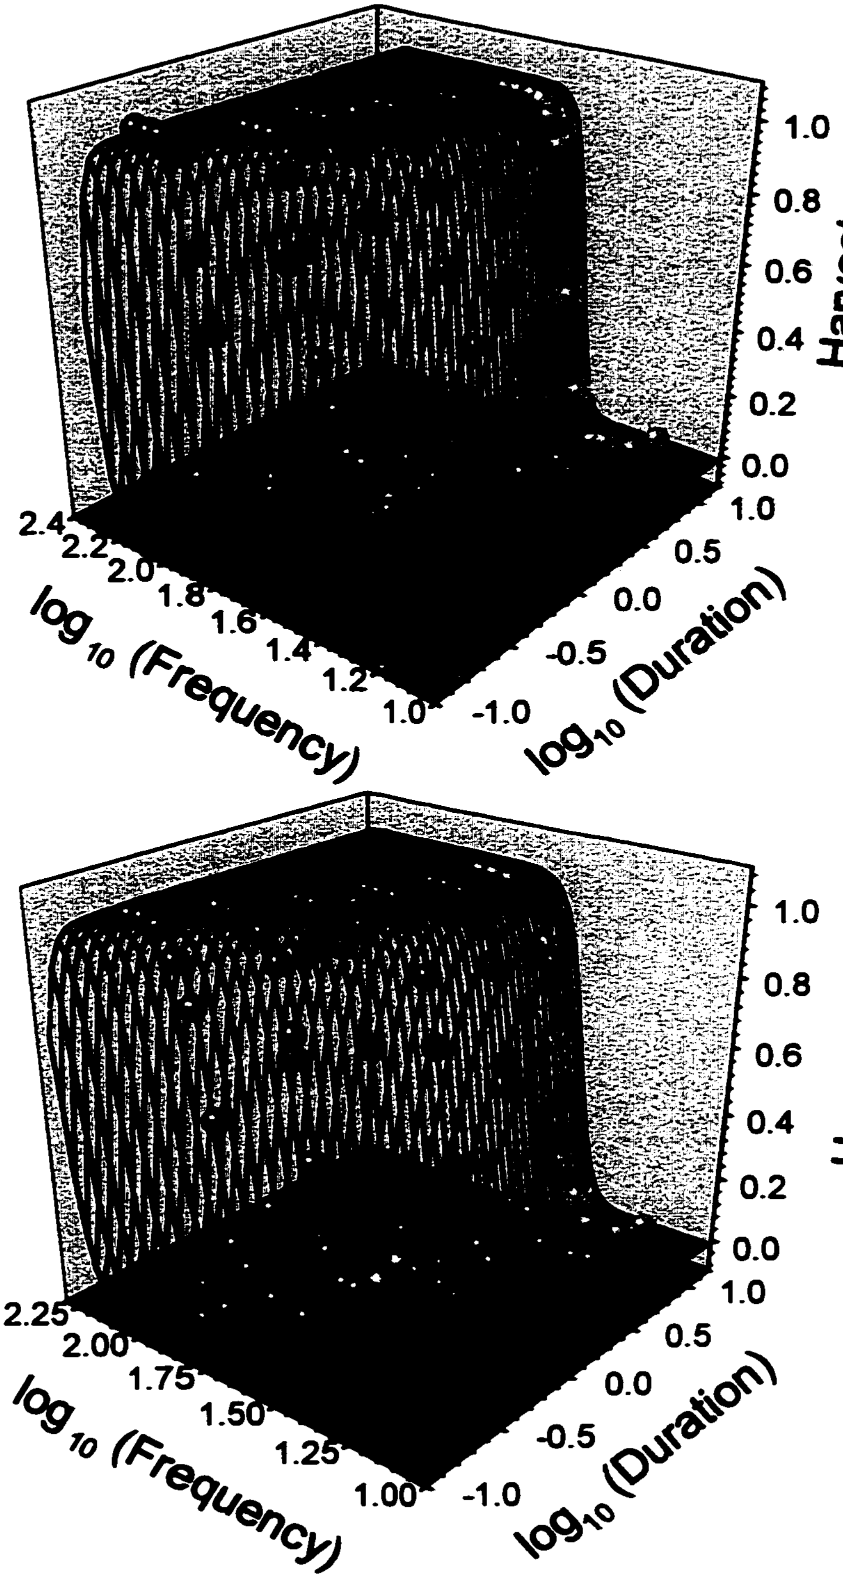

Figure 16. 3-D mountain fit to data from both the frequency and duration sweeps experiments, for rat B8 (top) and B9 (bottom).

As with the 3-D fits described above, other parameters of interest arising from the 3-D analysis were the estimated values of the relative rate of reinforcement (RRR), the growth exponent (g), the performance exponent (p), and the rheobase. For B8, the values of RRR, g, p, and the rheobase were 0.377, 4.406, 7.487, and 35.603, respectively, while for B9, the values of RRR, g, p, and the rheobase were 0.342, 3.025, 7.432, and 34.319, respectively. However, the results of the new 3-D fits for RRR, the rheobase, and g and p were again heavily influenced by the initial values used to start the fitting procedure. The initial values used in the analysis were the same as those used in the 3-D fits described in the pulse frequency sweeps section.

Peak-and-end test

The raw data results for the peak-and-end tests for rats B8 and B9 are shown in Figure 17. The y-axis variable is again the Proportion of Rewards Harvested (PRH), while the x-axis variable is the common logarithms of the pulse frequencies. The key comparisons of interest are whether the 8-sec (red) and 6-sec (yellow) constant-frequency trains are different, as a lack of difference would indicate whether duration neglect had been reached. The other comparison of interest is whether any of the composite frequency "6+2" trains (green, magenta, and blue) differ from the 6-sec constant train, which is a test of the "peak-and-end" model versus the "peak" model. For rat B9, it appears that all the curves are roughly on top of each other, while for rat B8, it appears that both the constant-frequency trains may differ from the composite-frequency trains.

Figure 17. Graphs of the raw means from the peak-and-end test, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis is the common logarithm of the pulse frequency (pulses per second). Plotted on the y-axis 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 or 6 sec constant-frequency train, and three composite "6+2" trains. For each "6+2" train, the first 6 seconds of the train is equal to the 6-second constant-frequency train, and the last 2 seconds is at a lower pulse frequency, with the ratio of the frequencies in the first portion of the train to the second portion of the train being either 2:1, 4:1 or 1.4:1. Each curve represents performance over 6 descending pulse "sweeps" at that particular train duration/type.

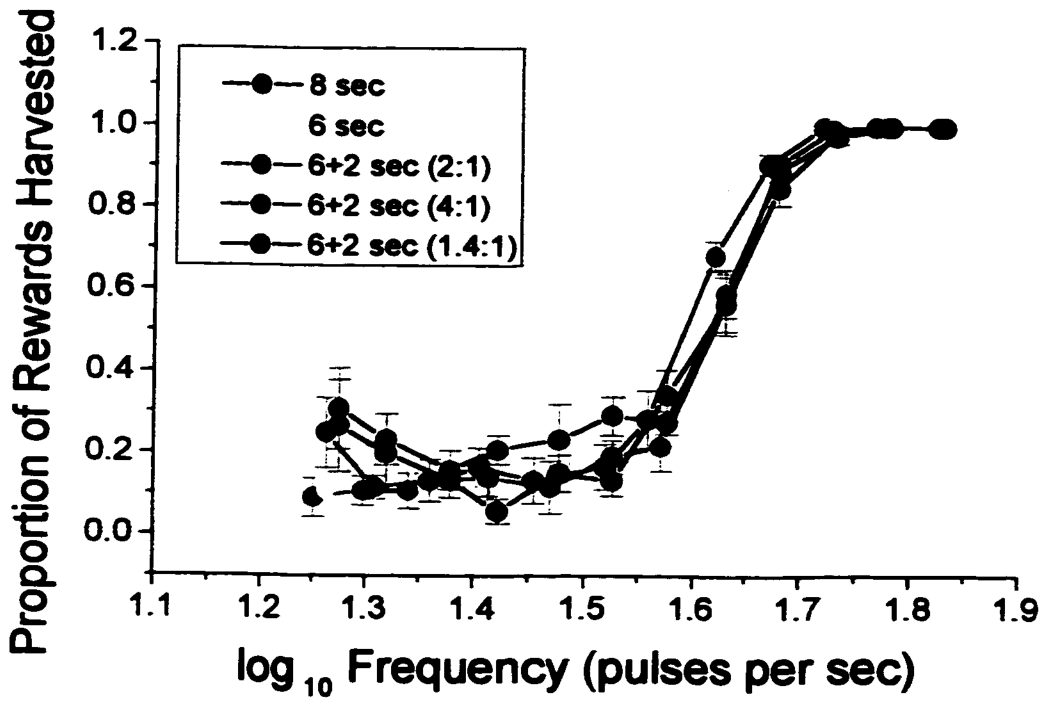

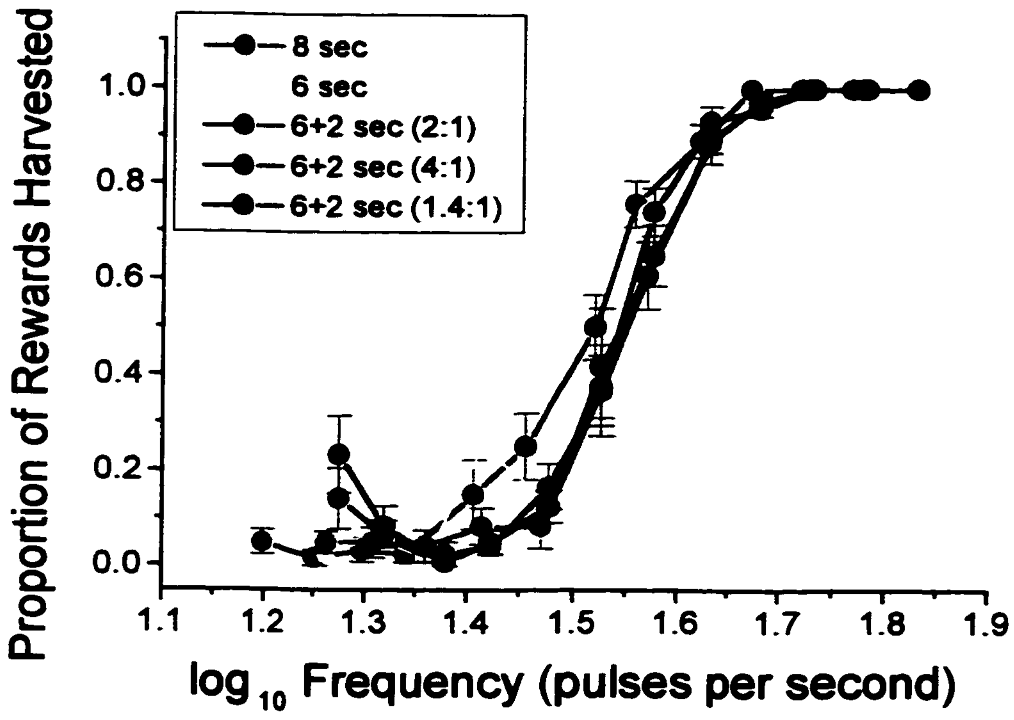

Figure 17. Raw data results of peak-and-end test for rat B8 (top) and rat B9 (bottom).

In order to assess whether any of the performance curves differed significantly from the other curves, an Analysis of Covariance (ANCOVA) was performed, with the alpha level set at 0.05. As the performance curves again suffered from the same problems with heteroscedasticity and with "noisy" lower asymptotes, weighting and rescaling of the curves was first carried out, in the same fashion as with the pulse frequency sweeps. The rescaled data results for B8 and B9 for the peak-and-end test are shown in Figure 18. An ANCOVA was then carried out on the rising segments of the curves, to determine if there were any differences between the constant-frequency 6- and 8-second curves and the three composite "6+2"-second performance curves. Only the rising segments were used in the ANCOVA because the upper and lower asymptotes do not yield much information about differences between the curves, as the upper and lower asymptotes are generally the same regardless of the train type used. In this analysis, the dependent measure was the PRH, the covariate was the pulse frequencies used, and the predictor variable was the train type.

The data points left after the removal of the upper and lower asymptotes, as well as the best-fitting lines through the individual rising segments predicted by the ANCOVA analysis are shown for rats B8 and B9 in Figure 19. The results of the ANCOVA show that for rat B9, there is no significant difference between the 6-second and 8-second constant-frequency trains [$F(1,63) = 3.359$, $p>0.05$]. Also, there were no significant differences between performance for the 6-second constant-frequency train compared to all three 6+2 composite-frequency trains [$F(3,109) = 1.803$, $p>0.05$]. Rat B8's results revealed that while there is no significant difference between the 8-second and 6-second curves [$F(1,45) =$

Figure 18. Graphs of the rescaled means from the peak-and-end test, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis is the common logarithm of the pulse frequency (pulses per second). Plotted on the y-axis 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 or 6 sec constant-frequency train, and three composite "6+2" trains. For each "6+2" train, the first 6 seconds of the train is equal to the 6-second constant-frequency train, and the last 2 seconds is at a lower pulse frequency, with the ratio of the frequencies in the first portion of the train to the second portion of the train being either 2:1, 4:1 or 1.4:1. Each curve represents performance over 6 descending pulse "sweeps" at that particular train duration/type.

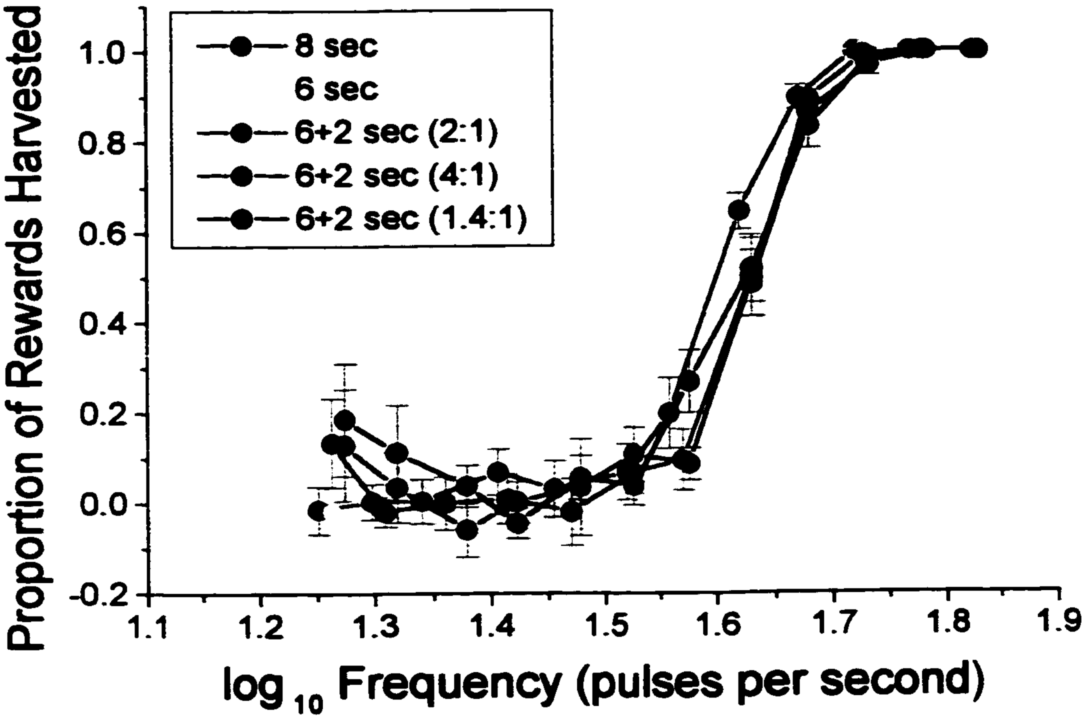

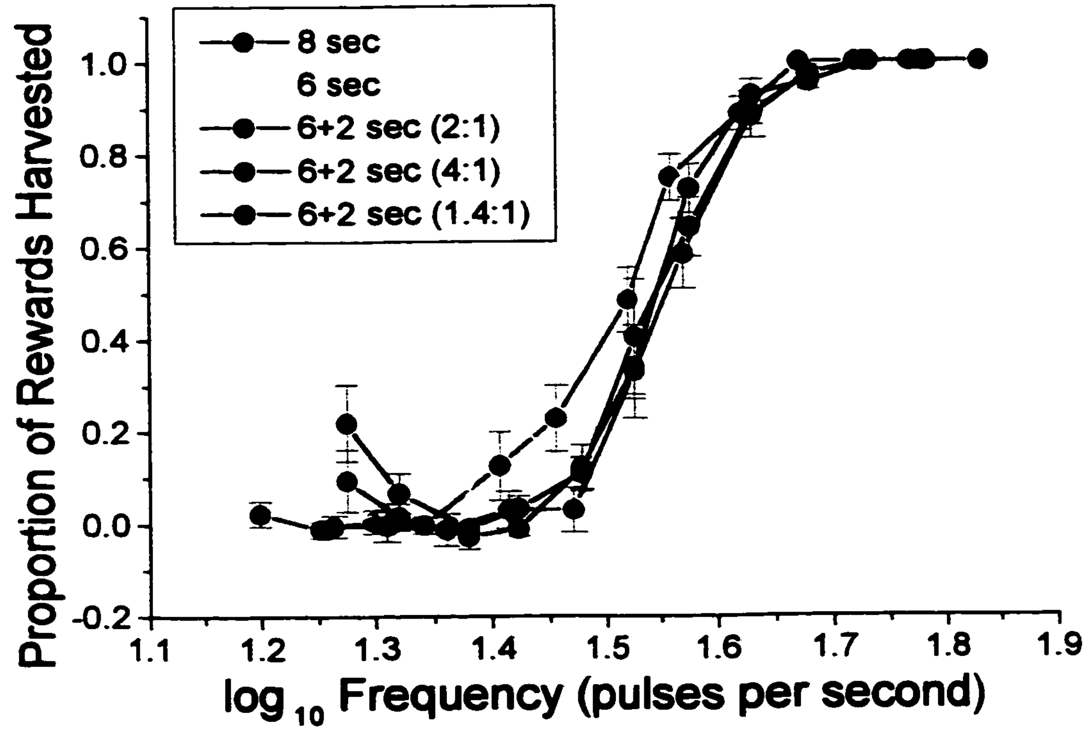

Figure 18. Rescaled data results of peak-and-end test for rat B8 (top) and rat B9 (bottom).

Figure 19. Graphs of the rescaled means and rising segments from the Analysis of Covariance (ANCOVA) carried out on the peak-and-end test, for rat B8 (top graph) and rat B9 (bottom graph). Plotted on the x-axis are the common logarithms of the pulse frequencies (pulses per second). Plotted on the y-axis is the proportion of rewards harvested, out of a maximum of 20. Each rising segment (and matching colour of data point) represents a different train duration/type tested: 8 or 6 sec constant-frequency train, and three composite "6+2" trains. For each "6+2" train, the first 6 seconds of the train is equal to the 6-second constant-frequency train, and the last 2 seconds is at a lower pulse frequency, with the ratio of the frequencies in the first portion of the train to the second portion of the train being either 2:1, 4:1 or 1.4:1. Each curve represents performance over 6 descending pulse "sweeps" at that particular train duration/type.

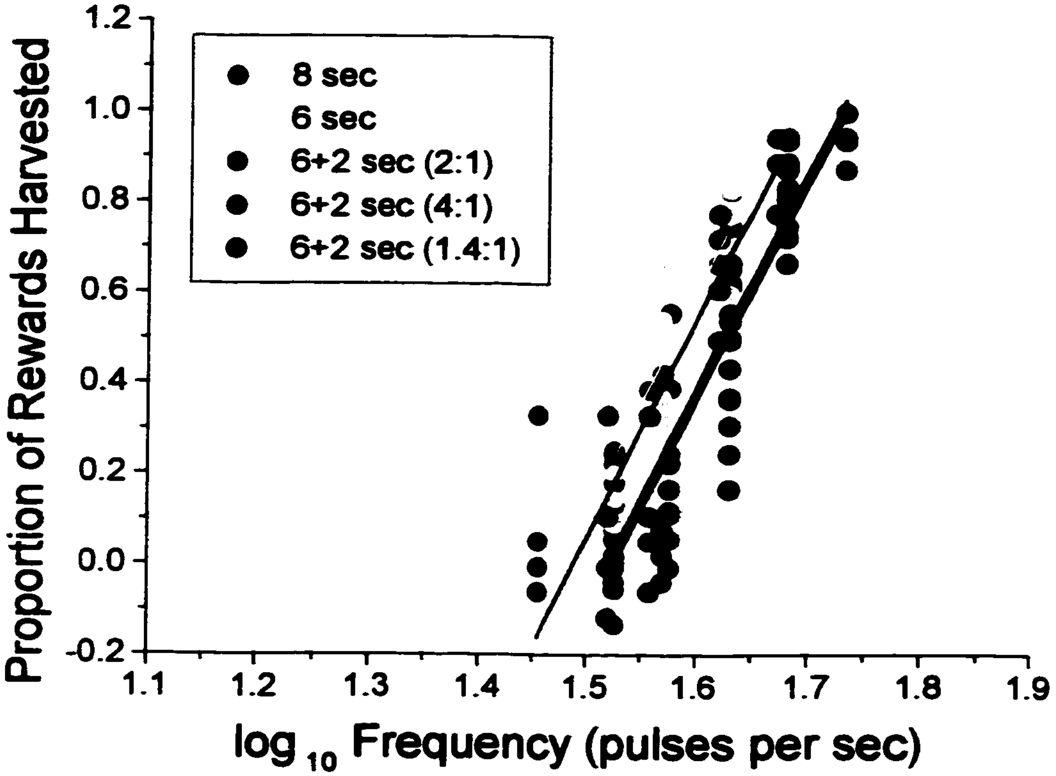

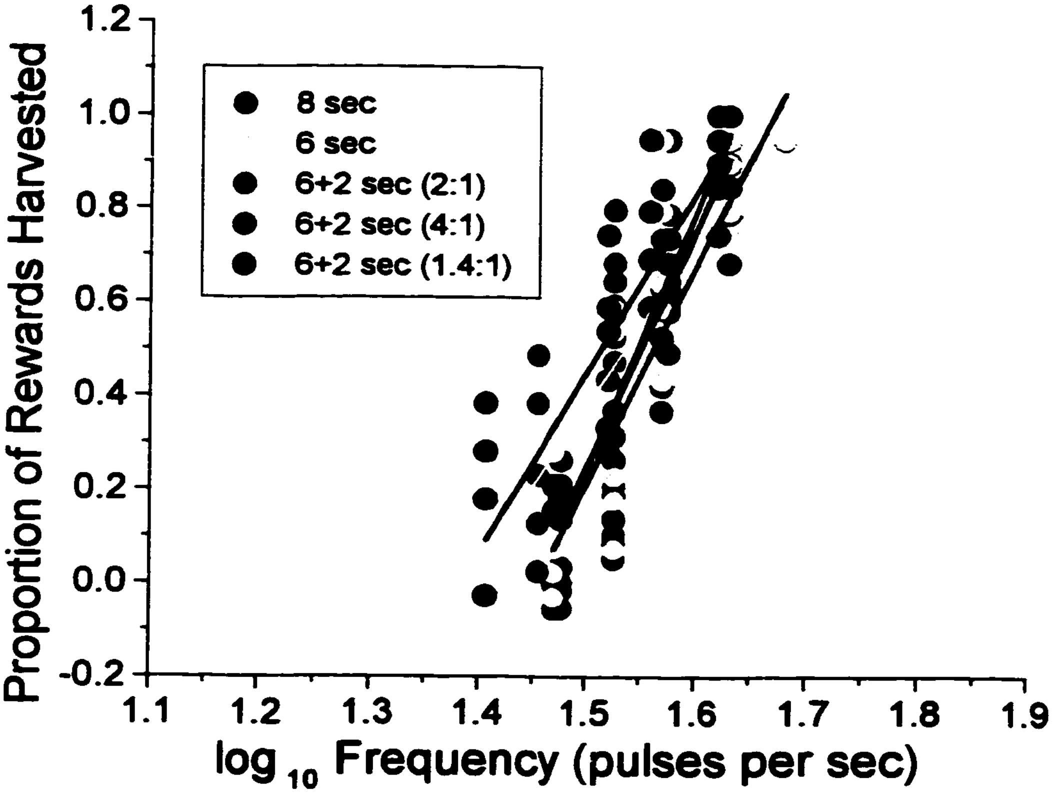

Figure 19. Rising segments and rescaled data points from peak-and-end test for rat B8 (top) and rat B9 (bottom).

1.377, $p>0.05$], there is a highly significant difference between the 6-second train and the 6+2 trains [$F(3, 97) = 11.473$, $p<0.001$]. Post-hoc comparisons revealed that for rat B8 the 6-second constant-frequency train differed from each of the composite trains ($p>0.05$), and that none of the composite trains differed from each other ($p<0.05$).

Thus, one rat (B8) appeared to show a "peak-and-end" effect, while the other rat (B9) only showed a "peak" effect. Thus, it is unclear whether adding a weaker "end" in the composite trains weakens their overall reward value compared to trains in which the "peak" and "end" are of equal reward value. However, as there were no statistically significant differences in performance between the two constant-frequency (6- and 8-second) trains for either subject, this suggests that "duration neglect" was reached; i.e. further increases in train duration beyond 6 seconds does not increase the reward value. These results are thus consistent with those of the earlier frequency and train duration sweeps studies.

Discussion

The current series of experiments focussed on two areas of interest. One aim was to resolve the discrepancy between Gallistel's (1978) assertion that duration neglect in BSR occurs after 1-2 seconds of stimulation, versus Mason and Milner's (1986a, 1986b) claim that duration neglect is dependent upon the pulse frequency used. "Duration neglect" for the purposes of the current studies was defined as the point at which further increases in train duration cease to add to the reward value. The second aim was to test Kahneman's (1993) "peak-and-end" model by pitting it against the earlier "peak" model (Norman & Gallistel, 1978). The "peak-and-end" model suggests that organisms form evaluations of experiences by taking two exemplars of the experience, the "peak" ("best" moment for appetitive experiences) and the "end", and performing some sort of averaging of the two, while the "peak" model asserts that it is only the "best" moment of the appetitive experience that is remembered.

In the first study, where pulse frequency sweeps were conducted while the train duration was held constant within the session, it was shown that the rats' performance curves were fit quite well by Gallistel's (1978) strength-duration function, although the results of B9's curve fit were better than those of B8, as his results at the two longest train durations deviated from the fit. The strength-duration function is a rectangular hyperbolic curve that relates the pulse frequency required to attain a particular criterion performance to the train duration used. One difference between Gallistel's (1978) results and those of the current study was that his rats appeared to reach the rheobase (duration neglect) after 2 seconds, whereas the most reliable subject in the present study (B9) demonstrated duration neglect after approximately 4 seconds. Overall, however, the results of both rats did appear to follow a hyperbolic function when the pulse frequency required to reach the performance criterion was plotted against train duration. Thus, the results of the current study appeared to support Gallistel's (1978) claims that the neural "integrator" that is responsible for integrating action potentials over space and time follows a hyperbolic function.

In contrast, in the second study, where train duration sweeps were conducted while the pulse frequency was held constant within the session, it was shown that the rats' performance curves appeared to conform to the ideas put forth by Mason and Milner (1986a, 1986b). At higher pulse frequencies, performance hit asymptote much sooner than when low pulse frequencies were used, for both subjects.

Thus, aspects of the formulations of both Gallistel (1978) and Mason and Milner (1986a, 1986b) appear to be correct, although their ideas are essentially opposing. One camp appears to claim that duration neglect always sets in around the same train duration, regardless of the pulse frequency used (Gallistel, 1978), while another claims that duration neglect is dependent upon the pulse frequency used [Mason, 1986a #56; 1986b #55]. How can this discrepancy be resolved?

The Shizgal triple-logistic model was a very good fit to the data for *both* experiments, and also resolves the discrepancy. Recall that the triple-logistic model has three elements. The first element is that the relationship between train duration and the pulse frequency needed to produce half-maximal reward value follows a rectangular hyperbola, just as in the Gallistel (1978) strength-duration function, and this allows the model to take his ideas into account. The second element of the model is that the subjective value of the reward follows a sigmoidal function of the pulse frequency used, and it is the combination of this second logistic equation with the first logistic equation of the model (the hyperbolic function) which takes into account both train duration and pulse frequency, thus integrating the ideas of Mason and Milner (1986a, 1986b) and Gallistel (1978) within the same mathematical model. The third and last element of the model is an equation that relates the non-observable, subjective intensity of reward to observable performance (bar-pressing), and this relationship is also a sigmoid.

Thus, it appears that the Shizgal model integrates the ideas of both camps of theorists. How does it succeed in reconciling these two seemingly irreconcilable ideas? This question can be answered by examining closely the 3-D "mountain" graph that is produced by the Shizgal model (see Figure 16). Recall that the pulse frequency, in log units, is plotted on the x-axis, the train duration, also logged, is plotted on the y-axis, and the PRH (proportion of rewards harvested) is plotted on the z-axis. One might imagine that if one were to stand on one side of the "mountain", one might see a very different picture than if one stood on the other side of it. For example, if one were to stand on the pulse frequency axis, one would see six curves lying nearly parallel to each other, with the curves getting closer together as they approach duration neglect. However, one need not look at just the train durations run during the present experiments; one could pick any combination of lines in the grid (function predicted by the Shizgal model) and follow those, and see a similar effect. Such effects are what Gallistel (1978) observed, and what is embodied in his hyperbolic strength-duration function. If, on the other hand, one were to stand along the train duration axis of the mountain graph, one would see two curves, one which is at a high pulse frequency and reaches behavioural asymptote very rapidly, and a second which is at a lower pulse frequency and reaches asymptote more slowly (if at all). Again, one need not look at only those pulse frequencies selected for use in the current studies; one would see the same effect by selecting any other high and low pulse frequency curves shown in the grid predicted by the Shizgal model. Such effects are what Mason and Milner (1986a, 1986b) observed, and this is embodied in their claims that duration neglect is dependent upon the pulse frequency used.

Thus, it appears that both camps of researchers were correct, but they were each only looking at half of what is essentially a 3-dimensional structure. When one examines each side of the mountain (in 2 dimensions), the claims of both camps appear discrepant, and it is only when one takes the entire 3-dimensional structure into account that one can resolve the seemingly discrepant findings in past research, as well as reliably predict at what point duration neglect sets in.

However, one could come to the same conclusions from careful examination of the 2-D graphs as well. For simplicity, the graph in Figure 20 is presented with only the predicted curves plotted, without any actual data points. If one examines more closely these predicted performance curves, one can see that taking a horizontal plane through the graph from any point of the y-axis yields the expected hyperbolic relationship; as the train duration

Figure 20. Graph of hypothetical pulse frequency sweep fits. Plotted on the bottom x-axis are the common logarithms of the pulse frequencies (pulses per second), and the top x-axis of each graph also displays the nonlogarithmic pulse frequency values (pulses per second). Plotted on the y-axis is the proportion of rewards harvested, out of a maximum of 20. Each curve represents a different train duration tested (0.25, 0.5, 1, 2, 4, or 8 seconds), and the black dashed vertical line represents the high pulse frequency train duration "sweep", and the gold dashed line the low pulse frequency train duration sweep. If one looks carefully at the graph, it becomes evident that while Gallistel's (1978) hyperbolic function can easily be seen on this graph, the effect that Mason and Milner (1986a, 1986b) predicted can also be shown. By following the vertical dashed black line (high frequency) from the bottom of the graph to the top, one can see that the hypothetical rat reaches asymptote (1.00 rewards harvested, or 20/20 rewards) at about 1 second (blue curve), whereas by following the gold dashed line (low frequency), one can see that the rat never reaches maximum reward harvest, even at the 8-second condition (orange curve). Thus, Mason and Milner's (1986a, 1986b) prediction can be shown in this 2-dimensional graph: at a high pulse frequency (black line), the reward value reaches asymptote quickly (in one second), and at a lower pulse frequency, the reward value rises more slowly (after at least 8 seconds).

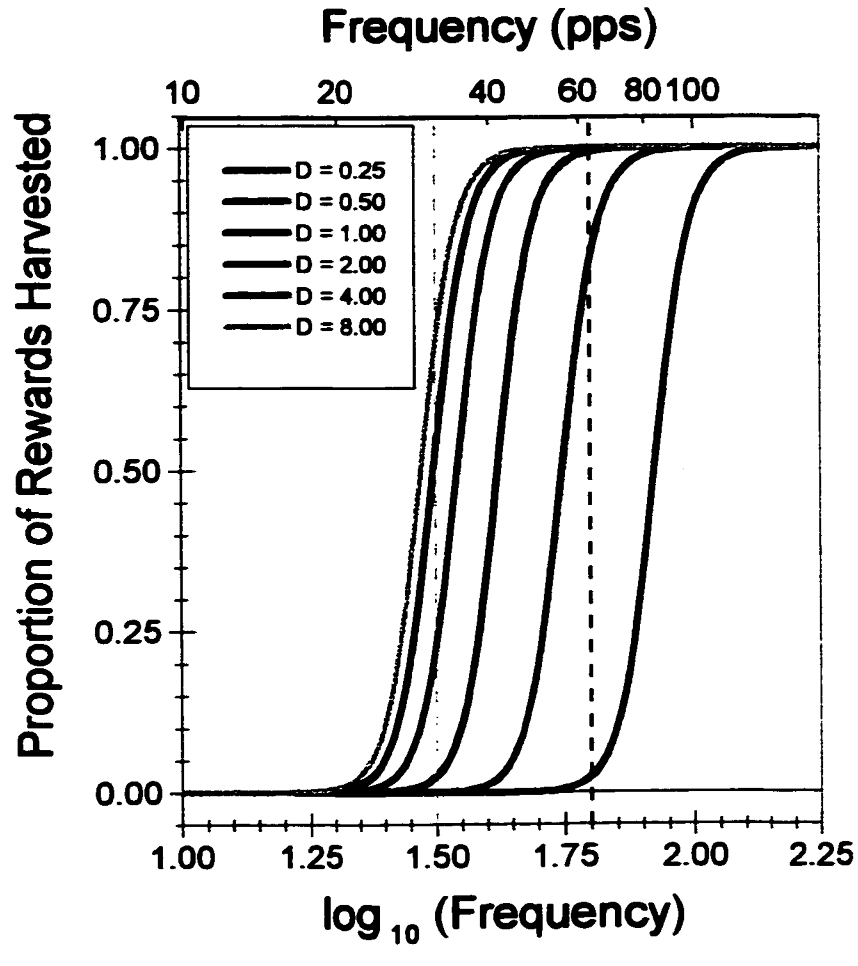

Figure 20. Predicted results of frequency sweeps.

increases, the spaces between each performance curve (or each point along the horizontal cut, in this case) decrease, until duration neglect is reached. The spacing between each train duration's point along the x-axis is the approximately same regardless of where one takes the horizontal cut used to derive Gallistel's (1978) hyperbolic strength-duration function. In fact, one can think of the horizontal cut as being equivalent to the criterion performance selected, which is then used to derive the function.

However, if one instead takes a *vertical* plane through the graph from any point on the x-axis, the effect that Mason and Milner (1986a, 1986b) described becomes evident, even in this 2-D display. One can think of such vertical cuts through the 2-D graph as a duration sweep, similar to those conducted in the present set of studies. In contrast to the horizontal cuts, in which the curves were parallel to each other, when one takes a vertical section through the graph, the curves are not parallel, due to the fact they are not linear. Thus, the profile one obtains with one's vertical cut is highly dependent upon where that cut is taken. If, for example, one takes the vertical cut (duration sweep) at a high pulse frequency (the black dashed line), the two shortest train durations (red: 0.25 seconds, green: 0.5 seconds) are spaced very far apart, while the remaining train durations are clustered together at the top of the graph, at maximal performance. In other words, if one were to move along this vertical cut, across the y-axis from the bottom to the top, one would see that at a train duration of 0.25 seconds, few rewards are being harvested, at 0.5 seconds, more than 75% of rewards are being harvested, and at the train durations of 1, 2, 4 and 8 seconds, all available rewards are being obtained. Thus, at this high pulse frequency, reward value rises quickly, with the rat obtaining all available rewards at train durations of 1 second and longer. In contrast, when a low pulse frequency is used (gold dashed line), a different profile is obtained. Here, the curves with the shortest durations are at nearly zero performance, and the curves representing longer durations are the ones spaced apart. Put another way, if one again were to move along this vertical cut, across the y-axis from the bottom to the top, one would see that at train durations of 0.25, 0.5, and even 1 second, bar-pressing for self-stimulation is minimal, at 2 seconds, the rat is obtaining 25% of the rewards, at 4 seconds, 50% of the rewards, and at 8 seconds, merely 75% of the rewards. Thus, in contrast to the high pulse frequency (black line), reward value rises slowly, with not even the longest train duration used yielding maximal performance.

Thus, it appears that the discrepant results between Gallistel's (1978) and Mason and Milner's (1986a, 1986b) studies is merely an illusion caused by taking different 2-D perspectives of the same 3-D structure, and that the Shizgal triple-logistic model is able to account of both sets of data, as well as providing a good fit to the performance curves obtained from the frequency and duration sweeps conducted in the present series of studies.

A second aim of testing the Shizgal model was to use its predictions in pitting Kahneman's (1993) "peak-and-end" model of retrospective evaluation against the older "peak" model (Norman & Gallistel, 1978). Since the humans in the peak-and-end studies displayed duration neglect, the aim of using the Shizgal triple-logistic model was to identify train durations and pulse frequencies that would be within the range of duration neglect, yet at which frequencies selected to serve as "peaks" and "ends" in the composite stimulation trains would have different subjective reward values to the rats.

The results of the peak-and-end tests revealed that when evaluating the reward value of BSR, one rat (B9) did not appear to take both the "peak" and "end" exemplars of the experience into account, while the other one (B8) appeared to use both exemplars. The fact that there were no significant differences between the constant-frequency 6-second train and all three composite-frequency trains for rat B9 suggests that this subject was only taking the "peak" into account, and that adding a weaker "end" did not degrade the overall reward value of the train. In contrast, rat B8 did show a significant difference between the constant-frequency train and the three composite-frequency trains, with the constant frequency train being more rewarding. This suggests that adding a weaker "end" in the composite trains did degrade the overall reward value of those trains. Thus, the "peak" (alone) model (Norman & Gallistel, 1978) can most parsimoniously account for the behaviour of rat B9, but the "peak-and-end" model (Kahneman, 1993) best accounts for the results of rat B8. However, for both subjects, there were no differences between the two constant frequency (6- and 8-second) curves, suggesting that both rats were displaying duration neglect, as did the human subjects in the Kahneman (1993) study.

How is it possible that the peak-and-end effect was obtained with one subject but not the other? While it could be that the peak-and-end effect is a real finding for rat B8, this animal's baseline data (not shown) became very unstable a few weeks after the peak-and-end test was run, suggesting that the electrode was moving. Tests with 4-second trains run on rat B8 around this time also revealed an instability in his data, as shown in the top half of

Figure 21. Graph displaying the vast shifts in B8's performance over the course of the experiments. Plotted on the x-axis are the common logarithms of the pulse frequencies (pulses per second), and the rewards earned (out of 20) are plotted on the y-axis. The top graph displays the results of two pulse frequency sweeps at a train duration of 4-seconds, the first run (magenta) taken 3 weeks after termination of the peak-and-end-test, and the second run (green) conducted a week later. The bottom graph shows a comparison between two pulse frequency sweeps taken at a train duration of 8-seconds. The initial 8-sec train (yellow) is the same curve that was run during the first experiment, plotted here for the purposes of comparison, while the second 8-second curve (blue) was run during the same period as the two 4-second trains in the top graph. Error bars represent the standard error of the mean. Each curve represents performance over 6 descending pulse frequency "sweeps" at that particular train duration.

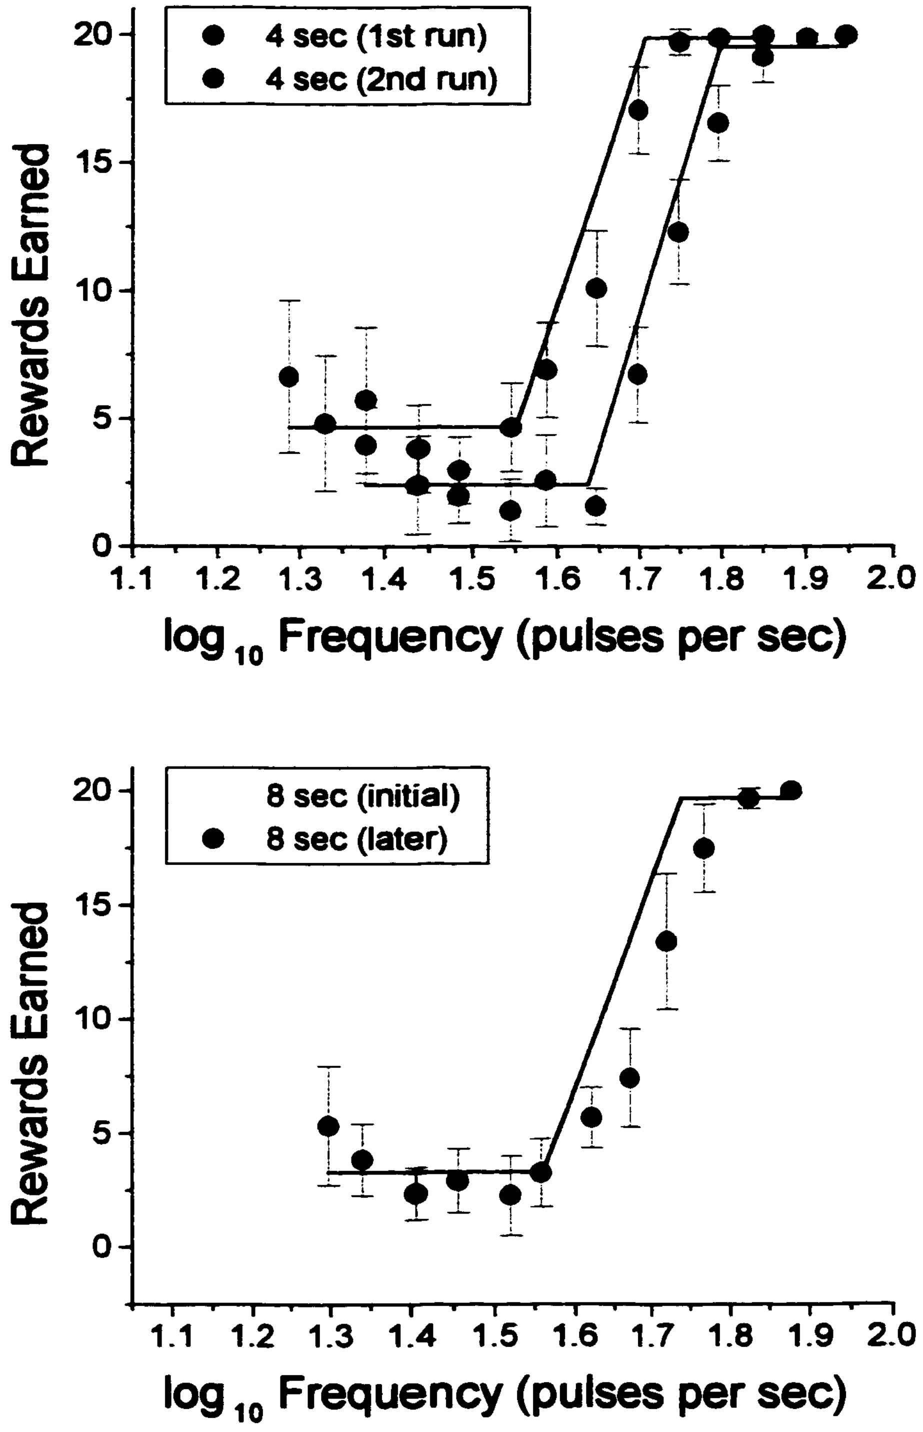

Figure 21. Shifts in performance over the course of the experiments by rat B8.

21. As can be seen from studying the graph, there is a large difference between B8's performance for the first 4-second curve and the second 4-second curve, run only a week later (recall that throughout the experiment the animals were run for several days to ensure stable performance at the current train duration/pulse frequency condition before being moved to the next condition. A similar amount of time allotted in the current 4-second versus 4-second tests to ensure stability in the individual performance curves in this case as well) . However, although one might argue that since this instability occurred after the peak-and-end tests were conducted, it likely had no effect on the results, the bottom half of Figure 21 shows that this may not be the case. The first 8-second train plotted in the graph is identical to the 8-second train plotted in the graph of the raw data results of the pulse frequency sweeps (Figure 9), and the second 8-second train is one that was taken around the same time as the two 4-second trains mentioned above (again, as with performance curves in the previous experiments, performance curves were run for several days to ensure that stable performance had been attained). Again, there is a large difference between these two curves, suggesting that B8's behaviour may have been unstable over the course of the experiments. Unfortunately, such effects did not show up in the baseline measures taken throughout the experiment, and are not easily detectable in the first two studies, because of the way that they were conducted. Because the different train durations were presented in decreasing order (i.e. 8, 4, 2, 1, 0.5, 0.25 seconds), and one would expect the later curves to be shifted rightward compared to the earlier curves (as shown in Figure 11) merely as a consequence of train duration, perhaps the order of presentation of the train durations masked a decrease in the rewarding effectiveness of the stimulation. Such an explanation could also account for the results of B8's peak-and-end test, as all three composite trains

were run after the 6- and 8-second constant frequency trains were conducted. In fact, the 8-second train was run after the 6-second train, and even though there were no significant differences between these two trains, the 8-second curve does lie on the "wrong" side (i.e. to the right) of the 6-second curve.

Therefore, the appearance of possible instability in B8's performance over the course of the three studies calls into doubt the apparent support of the peak-and-end model in the third study, and perhaps the results of this rat in the first two studies as well. However, as the very reliable and stable performance of B9 shows all the same effects as B8's performance except in the peak-and-end test, it seems possible that B8's results in the frequency and duration sweeps experiment are still relatively accurate.

It is possible that B8's display of a peak-and-end effect is not spurious, and that perhaps the rightward shifting of the three composite trains was a real effect. It seems prudent, therefore, to run more animals in the three experiments, in order to better assess whether the "peak" alone or "peak-and-end" model holds, and to verify whether B8's (and B9's) results on the first two experiments also constitute real findings. In fact, four animals are currently running in the frequency sweeps experiment as of this writing, and thus it will soon be possible to examine their performance and compare it to the results already mentioned here.

However, assuming that B9's results, which support the peak-alone model, are true and that the peak-and-end model is not appropriate to describe evaluations of reward in

BSR, there are several possible reasons why this could be the case. There are several differences between the present studies of BSR in rats, and the studies of retrospective evaluation of aversive experiences in humans conducted by Kahneman's group (Kahneman et al., 1993; Redelmeier & Kahneman, 1996).

First of all, one difference between the two sets of studies is that the stimulation Kahneman's group (Kahneman et al., 1993; Redelmeier & Kahneman, 1996) used was aversive, whereas the stimulation employed in the current series of studies was appetitive. One might speculate that different laws govern the evaluation of different types of experiences. For example, while "peak" and "end" exemplars may be valuable in the evaluation of negative experiences, perhaps the "end" of a positive experience is not as important, and it is rather the "beginning" and "peak" exemplars that are recorded in memory and used in later evaluation. Such questions need to be addressed in future research, perhaps by using composite BSR trains in which the "ends" of all trains are of equal reward intensity, and it is the "beginning" segments of the composite trains that are made to be "weaker".

A second difference between the studies is that the human subjects in the Kahneman studies (Kahneman et al., 1993; Redelmeier & Kahneman, 1996) were usually *passive* recipients of the aversive experience, whereas the rats in the current BSR studies were *active* recipients of the appetitive stimulation; the rat had to initiate delivery of the stimulation by pressing the lever. Perhaps evaluations such as "peak-and-end" averaging only hold when the organism is forced to passively receive the experience, whereas another type of combination or process holds when the organism must make active efforts to bring the experience upon itself. Future experiments in the area of BSR could examine the possible contributions of this active-passive difference, by forcing the rats to passively receive the rewarding stimulation. For example, one might place the rats in an open-field conditioned place preference paradigm (Vezina & Stewart, 1987a; Vezina & Stewart, 1987b), in which rats are trained over several days to associate one type of floor with the delivery of a constant-frequency train, and another type of flooring with the delivery of a composite train in which the "end" is weaker than the "peak". If on the test day, the rats spend significantly more time on the flooring that had previously been associated with the constant-frequency train, this would demonstrate that the weaker "end" of the composite train had degraded the overall reward value of those trains (and thus the reward associated with that floor type), and would thus constitute a demonstration of the peak-and-end effect, in a case where the rats are passive recipients of the experience.

A third difference between the present studies and the Kahneman *et al.* studies is the fact that in the Kahneman studies, the aversive experience took place over an extended length of time, whereas the stimulation trains used in the current studies were discrete experiences which were never longer than 8 seconds. Since it is possible that there could be a vast difference in how the brain evaluates experiences that are 8 seconds in length compared to those that are 60 seconds or even an hour or more in length, it would likely be of interest in future studies to attempt a peak-and-end test using longer stimulation trains. An experiment currently being carried out in our lab has shown that rats are capable of receiving constant-frequency stimulation trains as long as 60 seconds in duration, and therefore a test of the peak-and-end model using such trains is certainly feasible. However, since aversive effects tend to build up as the duration of the stimulation train lengthens (Shizgal & Mathews, 1977), if trains of several minutes or longer are to be tried, it may be best to use stimulation trains that change in pulse frequency over time, as in the experiment of Lepore and Franklin (1992). They devised stimulation trains in which the pulse frequency was modulated over time, and for the purposes of testing a peak-and-end effect with extended stimulation trains, the utility of using long-duration modulated-frequency trains to prevent rapid build-up of aversive effects would be very useful.

However, other groups researching the general concept of "representation by exemplar" in retrospective evaluation contend that exemplars other than just "peak" and "end" play a role in retrospective evaluations (Ariely, 1998; Ariely & Carmon, in press; Ariely & Loewenstein, in press). Although this group of researchers agrees with the Kahneman group that "peak" and "end" are important exemplars, they contend that the "trend" of the experience is also a valuable piece of information to the organism. They contend that experiences in which the pain gets less aversive (or the pleasure increases) over time are preferred to those in which the pain increases (or the pleasure diminishes) (Ariely, 1998). For example, in a study by Ariely (1998), humans were exposed to two different pain modalities in two experiments: the first experiment used the application of heat to the subjects' arms, and the second used mechanical pressure applied to the subjects' finger. The stimuli used were in many ways analogous to the constant-frequency and composite-frequency stimulation trains used here: some of the aversive stimuli involved holding the temperature or pressure constant throughout the duration of the presentation, while others involved various changes in temperature or pressure over time. For example, in the pressure experiment, the composite ("patterned designs") trains consisted of "Up" (pressure starts low, and then rises until it hits the peak at the end of the presentation), "Up & Down" (pressure starts low, until it hits the peak pressure at the midpoint of the presentation, and then the pressure decreases, until it is again at a low point at the end of the presentation), "Down" (pressure starts high and decreases until it hits a low at the end of the presentation), and "Down & Up" (pressure starts high and decreases until it hits a low point at the midpoint of the presentation, and then the pressure rises again until it again hits a peak at the end of the presentation). Thus, some of the "patterned" (composite) pressure experiences had improving trends (Down), worsening trends (Up), and those which were some combination of improvement and worsening (Down & Up, Up & Down).

Ariely's (1998) results showed, for both heat and mechanical pressure, that the retrospective evaluations were influenced primarily by both the trend of the last half of the aversive presentation and the final pain intensity (the "end"). Thus, it appears that the trend of the experience, particularly the experience's improvement or worsening in the last moments of the experience may also be important when evaluating experiences. However, Ariely (1998) acknowledges that in his experiments the trend of the stimulus presentations was closely linked with the "end" of the presentation, in that a highly painful "end" was preceded by a trend in which the pain levels increased, and "end"s that were less painful were preceded by a trend in which the pain levels decreased, making it difficult to tease apart the two exemplars.

Since stimulation trains can be constructed to vary in peak, end, and trend, it may be of interest in future studies to attempt to separate trend of the reward from the end intensity of reward. Perhaps this can be done while holding the "peak" and "end" reward values constant, but varying the modulation of the pulse frequency (again, similar to the trains used in (Lepore & Franklin, 1992)) in an attempt to create different trends of reward value between the peak and the end values.

However, it is also interesting to note that Ariely *et al.* (Ariely, 1998; Ariely & Carmon, in press; Ariely & Loewenstein, in press) have also shown that duration neglect is not as absolute as the Kahneman group's work has suggested. Returning to the studies of Ariely (1998) mentioned earlier, he discovered that with both of the types of noxious stimulation used (heat and pressure), while the duration of the experience had little impact on retrospective evaluations when the stimuli of constant pain intensity (analogous to the constant-frequency trains in the present studies), when patterned (analogous to the composite-frequency trains) stimuli were used, suddenly the duration of the experience did play a role. Specifically, the results indicated that an increase in duration in the patterned stimuli increased the perceived pain intensity of the presentation.

Another study (Ariely & Loewenstein, in press) suggested that the observation of duration neglect in the studies conducted by the Kahneman group was in fact a result of the type of reporting the subjects were asked to give about the painful experience. In the Ariely and Loewenstein (in press) study, subjects in different experiments were asked to give different reports of how "annoying" they found various aversive sounds. In the first experiment, subjects were asked to rate "overall how annoying was the sound" on a scale from 0 to 100, which is analogous to the global rating methods used in the studies conducted by the Kahneman group. In the second experiment, they were asked their willingness to repeat certain sound sequences in exchange for payment, in the third subjects were asked to rate the sounds in reference to a fixed standard sound, and in the fourth experiment, they were asked to make a choice between sound sequences. The purpose of using these different methods of rating the aversiveness of the annoying sounds was to see if it was possible in the Kahneman studies that in fact the subjects had been aware of the duration of the stimuli, but thought that such information was not relevant to the global rating of aversiveness that they were being asked to make, and thus did not take it into account when reporting their rating. Consistent with this idea, the results of the Ariely and Loewenstein (in press) study demonstrated that while some signs of duration neglect were seen in the first study, which was analogous to the types of ratings usually conducted by the Kahneman group in their studies, duration was an important concern to the subjects in the other three experiments, in which they had been asked to go about their ratings differently. They thus suggested that subjects do encode the duration of the experience, but they only include it when forming their retrospective evaluations when they feel that it is relevant to the type of rating they are being asked to make.

Although the results of the Ariely group's (Ariely, 1998; Ariely & Loewenstein, in press) studies are relevant to the Kahneman studies (Fredrickson & Kahneman, 1993; Kahneman et al., 1993; Redelmeier & Kahneman, 1996), it is unclear how their results apply to the current studies. Whether rats in a single-operant paradigm are being asked to make global ratings of the reward value of the stimulation train, as opposed to a different type of rating in which duration information may be more important, is open to debate. However, considering that the rats in the current single-operant studies displayed robust duration neglect, and that in choice experiments conducted later on with the same rats (not reported here), in which they were forced to make a choice between two levers (a "standard" lever which never varied in reward value, and an "alternative" lever on which the reward value varied), showed robust evidence of duration neglect in both circumstance, perhaps the concerns of the Ariely group are not as germane to BSR research.

Although an examination of the effects of trend on retrospective evaluation of BSR may be very informative, other future studies also suggest themselves. For example, as noted earlier, the results of the estimated parameters resulting from the 3-D model-fitting were found to be heavily dependent upon the starting values entered into the model. One possible way to lessen the model's reliance on the initial values is to set more constant parameters when running the fit. The chronaxie was derived from the 2-D strength-duration fits and set as a constant in all later model fits. Another parameter that could likely be derived from additional experiments and then set as a constant in the 3-D analysis is the parameter "R", the reinforcement rate. In principle, by setting the R with a certain schedule of reinforcement (such as either a fixed interval, fixed ratio, variable interval, or variable ratio), one could independently estimate R_e, g and p. One could do this using a "price" study, in which the rats would have to "pay" (perform) a certain number of bar-presses for stimulation trains that vary in certain characteristics (such as train duration or pulse frequency). By requiring the rats to "pay" bar-presses on a particular schedule, it will be possible to enter the "R" into the iterative curve-fitting procedure as a constant, along with the chronaxie, and will allow for a fitting procedure that is less influenced by the initial values. However, it should be noted here that the final form of the surface of the 3-D grid predicted by the curve-fit is itself not much influenced by the initial values, since as one parameter's value increases or decreases, the curve-fitting procedure decreases or increases the values of the other parameters to compensate for this. Thus, it is not the form (prediction) of the 3-D mountain that is at issue, but rather the precision of the individual estimated parameter values (RRR, g, p, and the rheobase) themselves that is the problem to be resolved by efforts such as "price" experiments. As the current study was more concerned with the actual prediction (form) of the model rather than the exact values of the estimated parameters themselves, the lack of precision in the final estimated values do not constitute a serious problem.

Lastly, if further tests of the peak-and-end model also support the simpler "peak" alone model, it will be of interest to study this model in more detail. For example, does it matter if the "peak" reward value comes at the beginning, the middle, or the end of the train? Also, it will be of interest to see if the "peak" model holds for other types of appetitive experiences, such as the delivery of a sucrose solution, or some other "natural" reinforcer.

In conclusion, the results of the present series of studies demonstrated that the Shizgal triple-logistic model allows for the most complete picture of the relationship between train duration and pulse frequency, and its effects on performance and "duration neglect". Specifically, it was seen that the discrepancy between Gallistel's (1978) and

Mason and Milner's (1986a, 1986b) claims in regards to the point at which duration neglect occurs was due to each group taking different 2-dimensional perspectives on what is essentially a 3-dimensional structure. The current set of studies also demonstrated that the Kahneman (1993) "peak-and-end" model cannot account for the behaviour of one of the subjects exposed to both constant-frequency and composite-frequency trains, and that the simpler "peak" model (Norman & Gallistel, 1978) is able to account for the results with this rat, although the other subject did show a peak-and-end effect. Therefore, more research with more animals needs to be conducted in the future, and even if the peak-and-end model is not supported in the long run, a larger investigation of the "peak-and-end" and "peak" models is warranted.

The broader implications of studies of retrospective evaluation of reward using BSR is that they may lead to the identification and description of the neural bases of psychological mechanisms of evaluation and choice. The identification of the neural underpinnings of the processes of evaluation and decision-making behaviour are of great importance to many in the fields of neuroscience, psychology, and economics, and may eventually allow us to better understand these and other cognitive processes.

References

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., & Carmon, Z. (in press). Evaluations of hedonic experiences: Combining sequences of pleasure and pain.

Ariely, D., & Loewenstein, G. (in press). The importance of duration in ratings of, and choices between, sequences of outcomes.

Fouriezos, G. (1995). Temporal integration in self-stimulation: a paradox lost? Behavioral Neuroscience, 109(5), 965-71.

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

Gallistel, C. R. (1978). Self-stimulation in the rat: quantitative characteristics of the reward pathway. Journal of Comparative and Physiological Psychology, 92(6), 977-98.

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

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

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.

Lepore, M., & Franklin, K. B. (1992). Modelling drug kinetics with brain stimulation: dopamine antagonists increase self-stimulation. Pharmacology, Biochemistry, and Behavior, 41(3), 489-96.

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. Behavioural Brain Research, 18(3), 223-31.

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.

Milner, P. M. (1989). The discovery of self-stimulation and other stories. Neuroscience and Biobehavioral Reviews, 13(2-3), 61-7.

Norman, M. F., & Gallistel, C. R. (1978). What can one learn from a strength-duration experiment? Journal of Mathematical Psychology, 18, 1-24.

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.

Sax, L. D., & Gallistel, C. R. (1984). Temporal integration in self-stimulation: a paradox. Behavioral Neuroscience, 98(3), 467-8.

Shizgal, P., & Mathews, G. (1977). Electrical stimulation of the rat diencephalon: differential effects of interrupted stimulation on on- and off-responding. Brain Research, 129(2), 319-33.

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

Vezina, P., & Stewart, J. (1987b). Morphine conditioned place preference and locomotion: the effect of confinement during training. Psychopharmacology, 93(2), 257-60.

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.