NOTE TO USERS

LINKING CHANGES IN PERFORMANCE FOR BRAIN STIMULATION REWARD TO STAGES OF NEURAL PROCESSING

Ada Mullett

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

April 2005

Your file Votre référence
ISBN: 0-494-04322-9
Our file Notre référence
ISBN: 0-494-04322-9

Abstract

Linking Changes in Performance for Brain Stimulation Reward to Stages of Neural Processing

Ada Mullett

The electrically evoked activity responsible for brain stimulation reward (BSR) is believed to mimic aspects of the neural activation caused by natural rewards, such as food or copulation. There is considerable evidence to suggest that this activation plays an important role in goal selection. Consequently, tracing the circuit responsible for the rewarding effect and working out its operating principles is an important step toward understanding goal-directed behaviour. This thesis tested a model of how brain reward circuitry operates. The proportion of a subject's time that it is willing to dedicate to performance of an operant task (time allocation) depends on the strength and cost of the reward. In the experiment described in this thesis, reward strength was determined by the pulse frequency of the electrical stimulation. The price of the rewards was manipulated by varying the average time that the subject had to hold down a lever in order to earn a train of rewarding stimulation. Plotting time allocation as a function of pulse frequency and price yields a three-dimensional structure called the "mountain." According to the model tested by the experiment, the mountain reflects processing of information about reward strength and price in a multi-stage network. An initial stage translates impulse flow in the directly activated neurons into a neural signal representing reward intensity. A later stage combines this intensity signal with the subjective mapping of reward price to yield a net payoff, and the final stage translates payoffs into behaviour. To understand how manipulations that alter operant performance produce their effects, it is necessary to determine at which stage(s) of processing they act. The primary goal of this experiment was to test the multistage model by establishing whether a manipulation that acts on the initial stage shifts the mountain along only one of its axes, as predicted. Reducing the duration of a stimulation train reduces the time during which stimulation-induced post-synaptic excitation can build up in the reward substrate. Increasing the pulse frequency can compensate for the reduction in train duration and thus, the mountain is predicted to shift along the frequency axis but not along the price axis. In all five subjects, reducing

the train duration from 1 second to 0.25 seconds resulted in a marked rightward shift along the frequency axis but little shift along the price axis. The results support the principal prediction of the mountain model. The demonstration that the mountain can shift along one axis without shifting along the other supports the notion that the integrative model on which the mountain is based can distinguish manipulations that affect the reward system at different stages. It should be possible to determine, in future experiments, at what stages of the circuit drugs, lesions, and physiological variables alter the rewarding effect of electrical brain stimulation.

Acknowledgements

I would like to thank all of the many wonderful people who played an important role in my life academically and emotionally while I completed the work required for this thesis. My lab mates have been very encouraging. I have had the pleasure of working closely with a generous group of people and I feel lucky to have been a part of the Shizgal Lab. Thanks to those who assisted in the data collection: Yuri Tekada, Stacey Noble, Rebecca Solomon, and Yannick Breton. The other members of the lab who have inspired and encouraged me include: Stephanie Fulton, Keiji Oda, Nafissa Ismail, Bonnie Sonnenschein, Selma Hamdani, Giovanni Hernandez, Renu Chitra and the rats. The work conducted here would not have been possible without the help and ingenuity of the technical wizards at the CSBN: Steve Cabilio and Dave Munro.

Special thanks needs to be extended to Dr. Kent Conover who has an amazing approach to statistics and problem solving. Kent taught me all of the hands-on techniques I needed. My academic supervisor, Dr. Peter Shizgal has been of immense inspiration. I am unlikely to find someone as brilliant or generous to work with again. I am both spoiled and lucky. Thank you.

My family, as always has been supportive and wonderful throughout this process. My mother, Dr. Sheila Mason, is my hero!

The research conducted for this thesis was supported by a grant to P. Shizgal from the Canadian Institutes of Health Research.

Table of Contents

Glossary

a	Behavioral allocation exponent ▪ it is the slope of the price function ▪ it determines how spread out the contour lines are on the price side of the mountain ▪ determines the price-sensitivity of the behavioural-allocation function
behavioural-allocation function	results from the translation of the reward growth function into observable behaviour
BSR	Brain stimulation reward ▪ refers to the phenomenon that animals will work to receive electrical stimulation through electrodes implanted in specific brain regions
fhm	frequency producing half-maximal reward
FVI	Free-Running Variable-Interval schedule (FVI). The lever is armed after a time interval (the price), sampled from an exponential distribution, has elapsed. In order to harvest the reward the rat has to be holding the lever down when the interval times out
g	Growth Exponent ▪ it contributes to the determination of the slope of the frequency side of the mountain ▪ determines the rate of reward growth for BSR
Half-maximal reward	the point on the reward growth function that corresponds to the position half way up the function (TA_{max} / 2)
Integrator	Spatio-temporal integrator
Leisure	includes all activities the rat might engage in while in the operant chamber such as grooming, investigating, or sleeping
$Log_{10} P$	Log of the Price
Matching Law	Herrnstein's single operant matching law (1970) which describes the way in which animals apportion their time to different choices
MFB	Medial forebrain bundle ▪ A bundle of axons that connect the forebrain to the hindbrain ▪ heavily implicated in reward perception
Mountain	3D psychometric function ▪ quantifies reward value based on how a subject spends its time in an operant chamber
Multi-stage model	A model that describes the multiple stages that occur between the time reward neurons are directly stimulated by the BSR electrode and the final behavioral output of the rat

Objective price	The price set by the experimenter
P_E	the price at which time allocation is 0.5 the subjective price at which the payoff from a maximal BSR equals the payoff from everything else
P_{min}	the minimum subjective price
price	Price is defined as the average time the rat has to hold down the lever in order to earn a reward (set by the experimenter)
Reward-growth function	represents the assessment of rewards through a variety of strengths in such a way that the rat can have a mental representation of reward
Subjective price	subject's interpretation of the objective price
Substitutability	describes the similarity of competing goods
TA	Time allocation
TA_{max}	maximum time allocation
TD1A	1st baseline condition (train duration of 1 sec)
TD1B	2nd baseline condition (train duration of 1 sec)
TDP25	Experimental condition (train duration of 0.25 sec)
Time allocation	Determined based on how a rat chooses to allocate its time between work and leisure activities
TA_{min}	minimum time allocation
TP	Transition parameter ▪ it determines how sharp the curve between the price and the frequency axis is
VI	Variable interval schedule of reinforcement
Work	refers specifically to pressing the lever for stimulation

Linking Changes in Performance for Brain Stimulation Reward to Stages of Neural Processing

The phenomenon of brain stimulation reward

A male rat is calm, sitting in an operant chamber, grooming. It has become familiar with the operant box and has been trained to press a bar to obtain a reward. The rat knows that a reward will be available soon. Suddenly a lever appears in the operant chamber. The rat leaps on the lever and presses frantically. It is working for neither food, nor water, nor access to an estrus female. Yet the rat is willing to work very hard, and when the reward is strong the rat will spend almost all of its time pressing the lever. Its payoff is a train of electrical current pulses delivered through the tip of an electrode implanted in its brain.

The effect of the electrical stimulation that the rat seeks to reinstate is called "brain stimulation reward" (BSR). This rewarding effect is so compelling that rats will cross electrified grids, climb obstacles, and tolerate discomfort in order to access the stimulation (Frank, Preshaw & Stutz 1982; Olds, 1958). The electrically evoked activity responsible for BSR is believed to mimic the neural activation caused by natural rewards, such as food or copulation (Hoebel, 1969; Shizgal, 1999). Consequently, tracing the circuit responsible for the rewarding effect and working out its operating principles is an important step toward understanding goal-directed behaviour.

Why is the problem important?

In a series of experiments conducted by Conover, Woodside and Shizgal (1994), rats evaluated brain stimulation reward and a sucrose solution. These are both very

compelling rewards. Conover et al. found that these drastically different rewards could be combined to result in a reward greater than either alone; that is, these rewards could be summated. Furthermore, when the two rewards were compared, the rat chose the larger of the two. The results of these experiments suggest that BSR and natural rewards are evaluated along a common dimension. If so, the neurons underlying BSR appear to participate in computing a common currency for evaluating natural goal objects.

This thesis will test a model of how brain reward circuitry operates. If validated, this model could play an important role in finding the reward-related neurons, a key step toward understanding how natural goal objects are evaluated and selected (Shizgal, 1997).

Malfunction of the circuitry underlying the evaluation and selection of goal objects would likely have disastrous consequences. For instance, drugs that hijack this circuit could overshadow the influence of goal objects crucial for survival. Malfunction of this circuitry might also contribute to other impulse-control disorders, such as gambling disorders and obesity (Shizgal, 1999). Thus, identifying and understanding the circuitry underlying BSR might contribute to the development of improved treatments for disorders that undermine the well-being of individuals, families, and societies.

An integrative model of BSR

To build an adequate model of the BSR phenomenon, we have to fill in the blanks between the stimulation electrode and the behaviour observed. This phenomenon involves multiple stages of processing. The goal is to describe, control, and find the behavioural neural signatures of each stage. In the following sections, the multi-stage

model is developed and a strategy for distinguishing the behavioural signatures of its different components is described.

Stages of processing underlying intracranial self-stimulation

As Figure 1 shows, the first stage of the circuit underlying BSR consists of the directly activated neurons whose outputs are ultimately translated into the rewarding effect. Each stimulation pulse triggers an action potential in the axons that are close to the electrode tip. The next stage of the model integrates the post-synaptic impact of the firings elicited in the first-stage neurons. This integrator assesses the impact of action potentials over a specific region of space (the area occupied by the directly stimulated neurons) and over a certain amount of time (the duration of the train of pulses). This spatio-temporal integration reflects the sum of the neural firings over space and time (Gallistel, 1978; Gallistel, Shizgal & Yeomans, 1981). This sum is translated into a reward-intensity value (Gallistel & Leon, 1991). That is, the output of the integrator is translated into an enduring record of reward (Figure 2). The function that maps the input to the integrator into its output is called the "reward-growth function". The reward intensity (the output) first increases with the firing rate (the input) but eventually levels off (Simmons & Gallistel, 1994). This stage of the model is called the "inner" function, since its output is "hidden" inside the brain and is not manifested directly in behaviour.

Translation of the stored reward-intensity value into behaviour requires a second function: a "behaviour-allocation function." To define this function, it is necessary to categorize behaviour as either "work" or "leisure." Work consists of activities (e.g., pressing a lever) that earn BSR (the rat's "wage"); leisure refers to everything else the rat does in the box, such as grooming, resting, or exploring. The behaviour-allocation

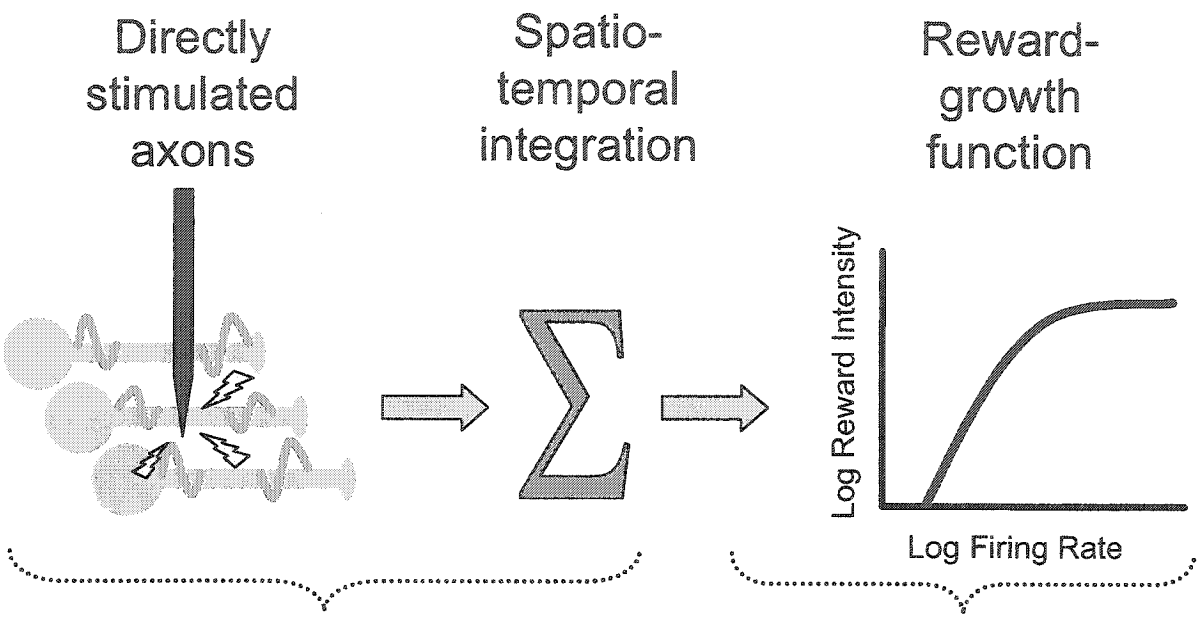

Figure 1 a

Figure 1 b

Figure 1a. Artificially stimulating a set of neurons, with a BSR electrode, will result in a corresponding increase in neural firing. The corresponding increase in neural firing is interpreted by the Spatio-temporal integrator. The integrator takes into account the total number of firings over space and time.

Figure 1b. The output of the integrator results in a reward-growth function.

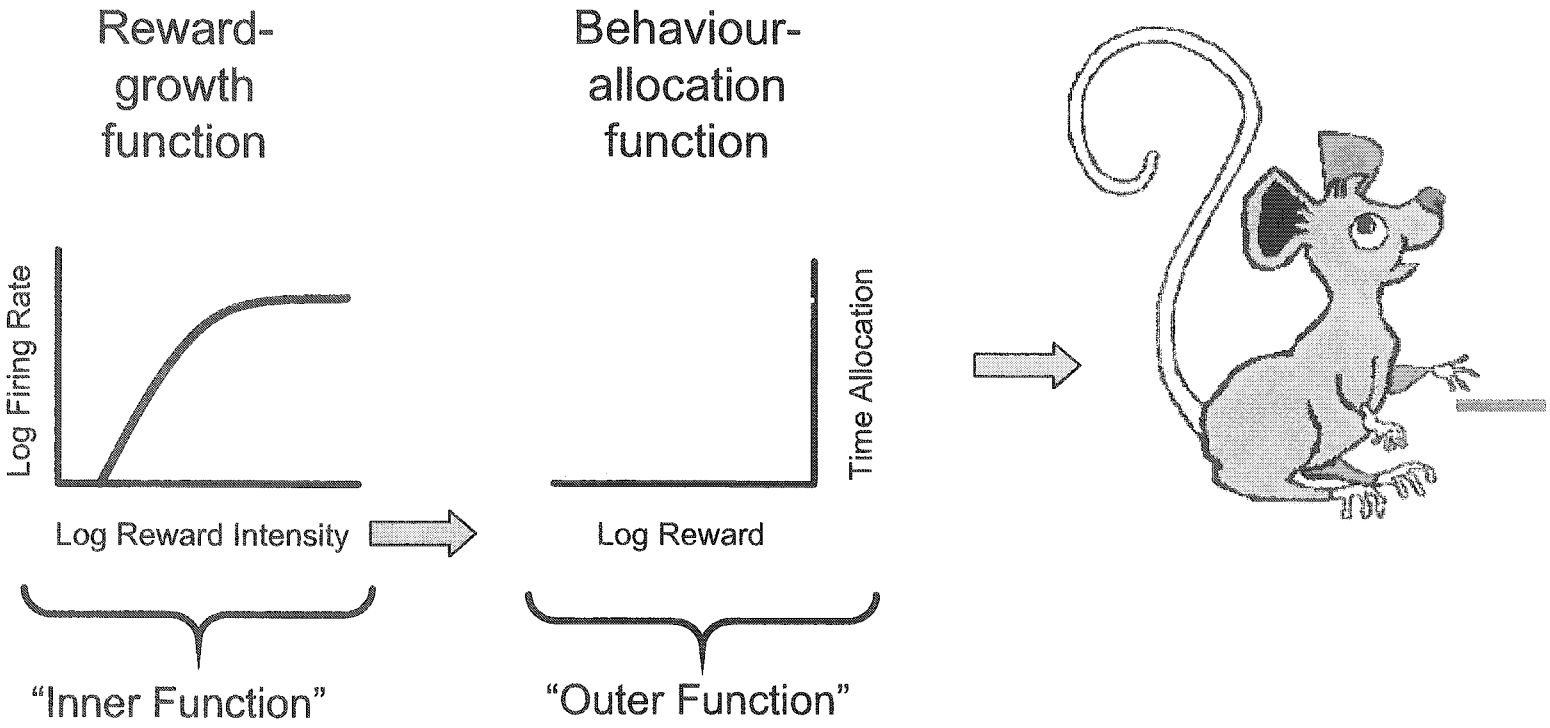

Figure 2. The reward-growth function is translated into a corresponding behavior: the behaviour-allocation function. The behaviour-allocation function translates a given reward intensity into a given partition of the rat's time between work and leisure.

function translates a given reward intensity into a given partition of the rat's time between work and leisure (Figure 2). When the reward is weak, the rat spends little time working and most of its time in alternate activities. When the reward is strong, the rat will dedicate all or most of its time to earning stimulation trains. The behaviour-allocation function is called the "outer function" since its output is directly observable (Figure 3) (Gallistel, 1978; Gallistel et al., 1981).

A psychometric function relates an independent stimulus variable controlled by the experimenter to an observable dependent variable, a behavioural output. Linking the reward-growth and behaviour-allocation functions produces a psychometric function (Figure 4) (Gallistel, 1978; Gallistel et al., 1981). This function relates the strength of the electrical stimulation (the physical variables that determine the number of elicited neural firings) to the rat's allocation of time between work and leisure (Shizgal, 2004). The model developed here provides a means of decomposing psychometric functions for BSR into their component parts, thus rendering a hidden signal, reward intensity, amenable to measurement. Moreover, as the next section shows, the model also provides a way to link manipulations that alter BSR to specific stages of processing (the inner or outer function).

Figures 5a and 5b present two commonly used psychometric functions. The first function, Figure 5a, illustrates how time allocation increases with the stimulation frequency. (Recall that reward strength increases with frequency.) This function is the familiar "reward-summation function" that has been used extensively to measure BSR (Edmonds & Gallistel, 1974; Gallistel et al., 1981; Miliaressis, Rompre, Laviolette, Philippe, & Coulombe, 1986; Shizgal & Murray, 1989). In Figure 5b, time allocation

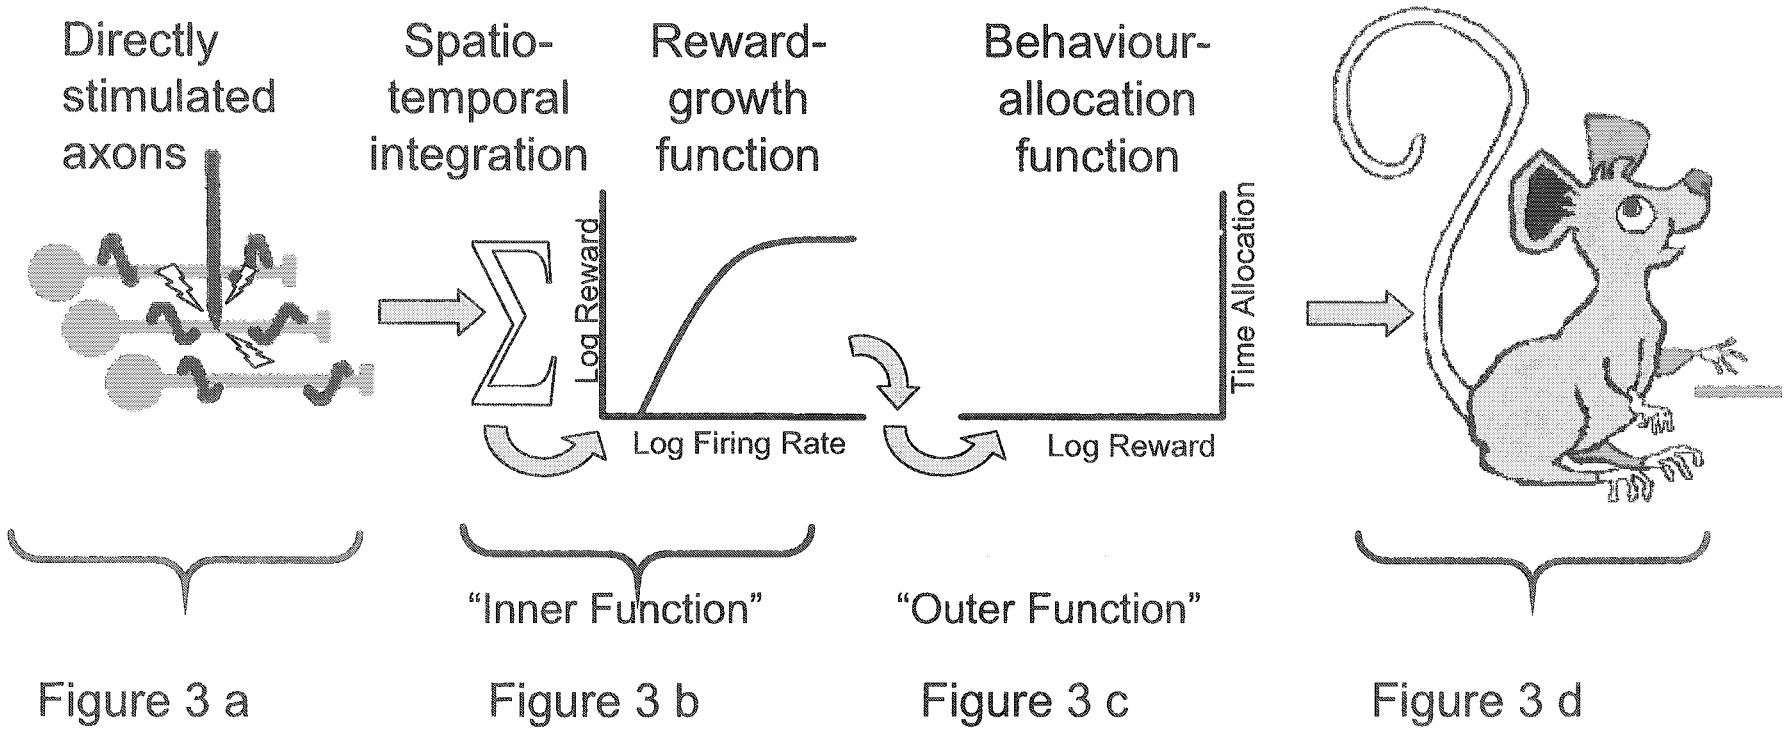

Figure 3. The multi-stage model that represents the relationship between neural firing and the output of the integrator in the reward-growth function and the behavior allocation function.

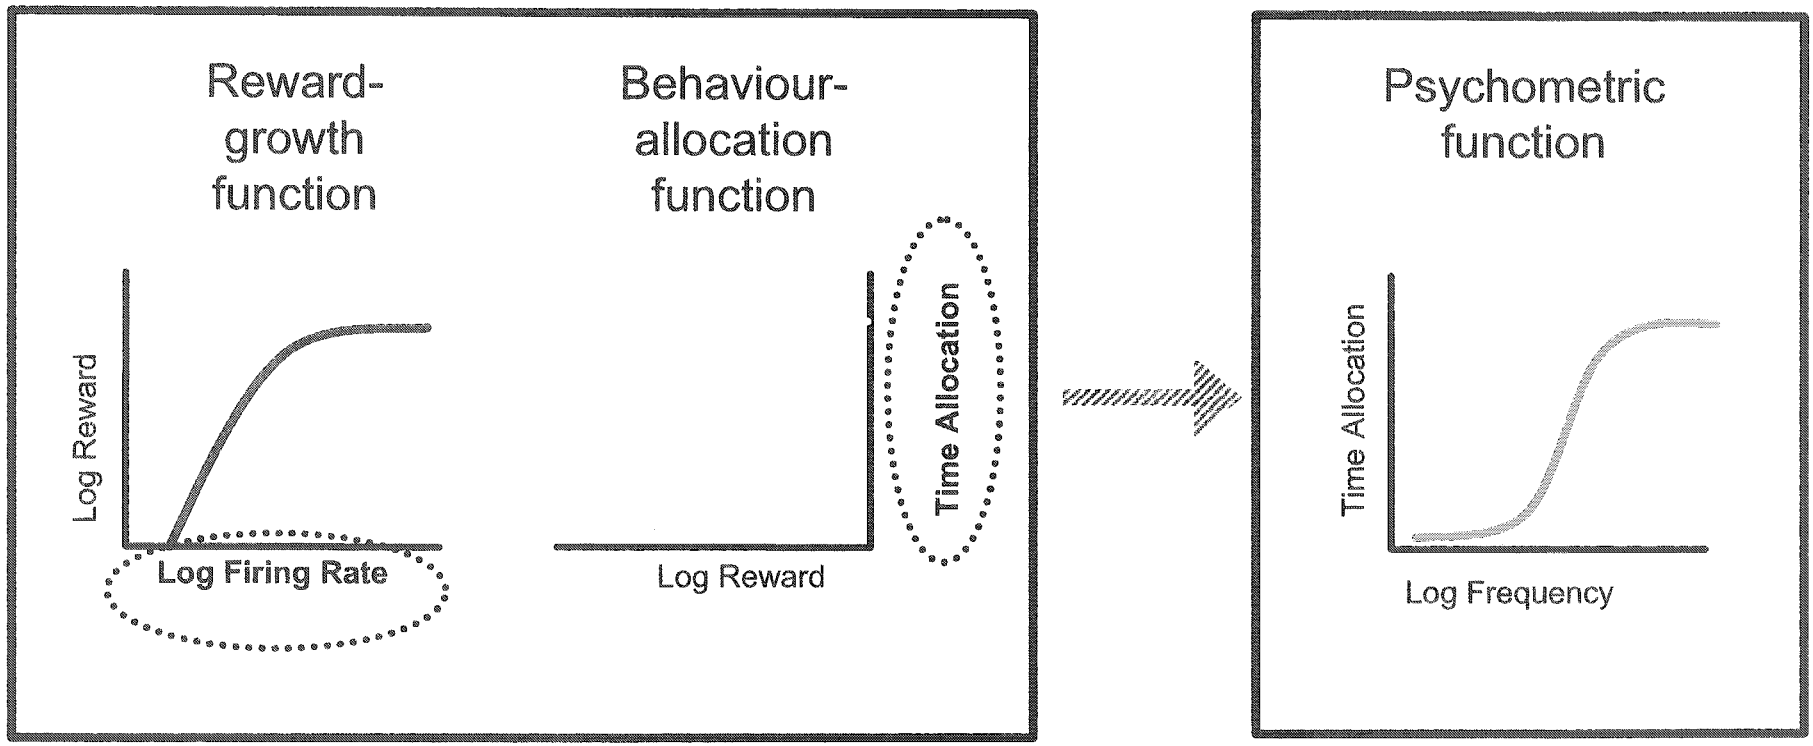

Figure 4. By combining components of the reward-growth function and the behaviour-allocation function a psychometric function is created.

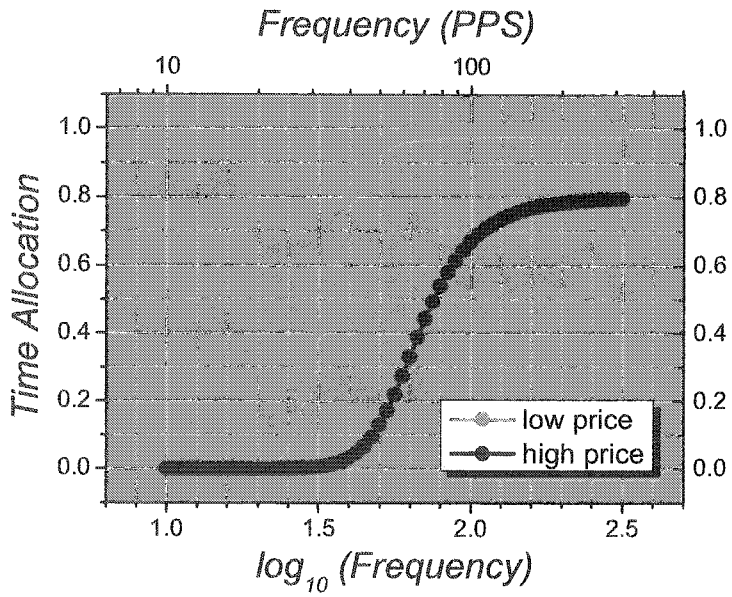

a

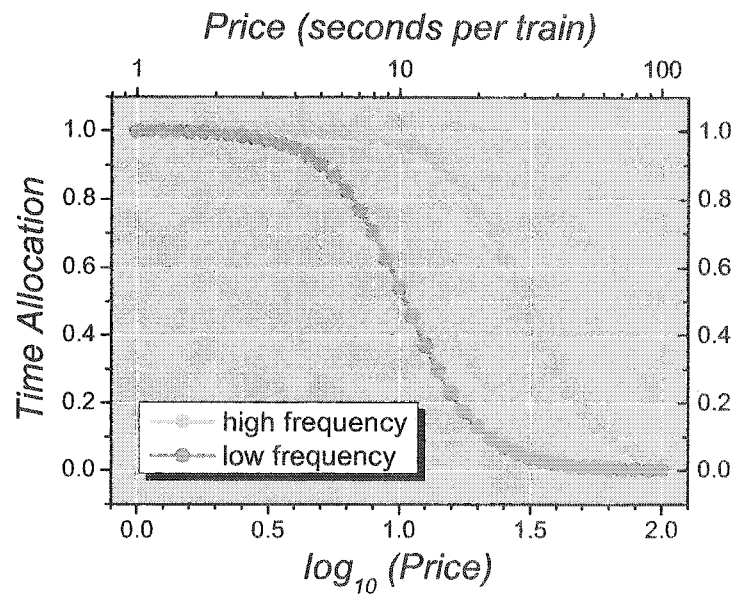

b

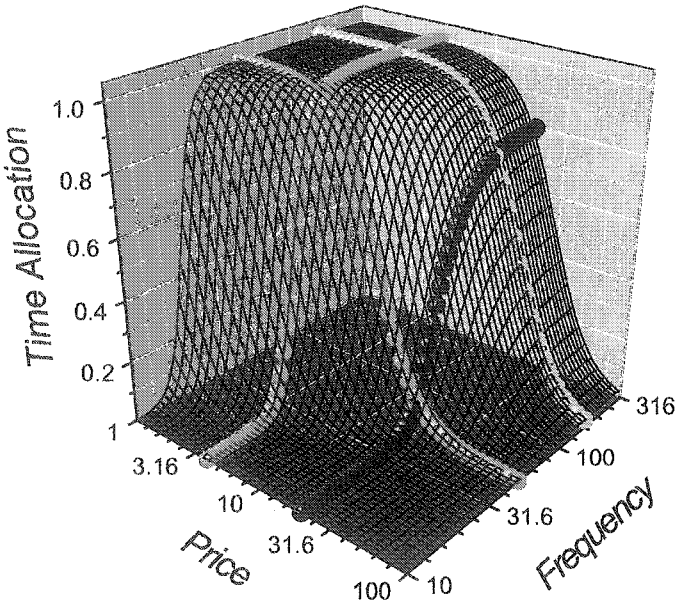

c

Figures 5a & b. Two commonly used psychometric functions. *Figure 5a* is the "reward-summation function" whereas the function presented in *Figure 5b* is analogous to the one derived from progressive ratio measurements of BSR.

Figure 5c. The nature of the interaction between price and frequency is revealed by plotting both sets of 2-dimentional (2D) psychometric functions in the same 3-dimentional (3D) space.

decreases as the time required to earn each reward ("price") increases. This function is analogous to the one derived from progressive ratio measurements of BSR. Note that the two reward-summation functions in Figure 5a differ because they were obtained at different prices, and the two functions in Figure 5b differ because they were obtained using different frequencies. Thus, frequency and price interact in determining time allocation. The nature of this interaction is revealed by plotting both sets of two-dimensional (2D) psychometric functions in the same three-dimensional (3D) space (Figure 5c). Obtaining a high time allocation requires both a low price and high reward strength are required. If the frequency is very low or the price is very high, then the time allocation will be low.

Figure 6 shows how the 3D representation can reveal the "hidden" inner (reward-growth) function. Time allocation has been color-coded with each decile represented by a unique color. The boundaries between the colors (white lines) are contour lines; they consist of pairs of frequencies and prices that generate the same time allocation. Projecting the white contour lines of the mountain upward onto a surface generates a contour map that provides a 2D representation of the 3D structure. This contour map is a concise summary both of the shape of the mountain and of its location along the frequency and price axes. In this thesis, this representation will be used to describe the effects of a variable that moves the mountain.

Figures 5 and 6 develop the mountain graphically. The following sections describe the equations that correspond to the figures and define key terms. The balance of the introduction further discusses the advantages of 3D representation relative to 2D representation.

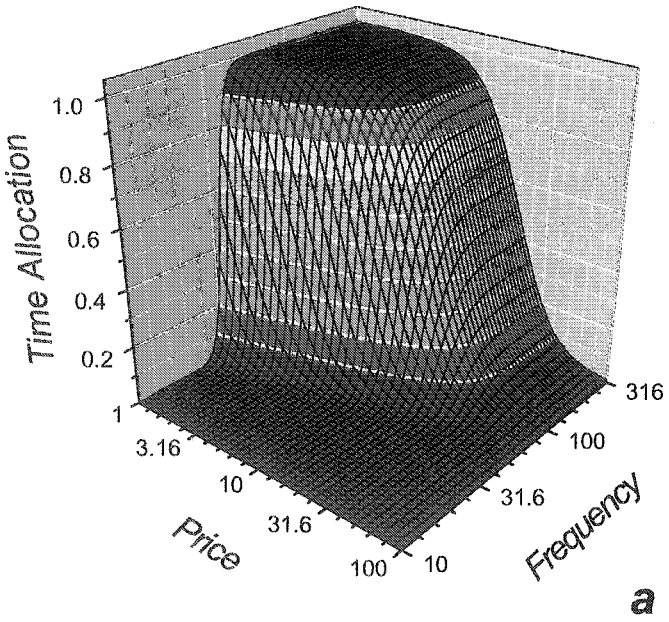

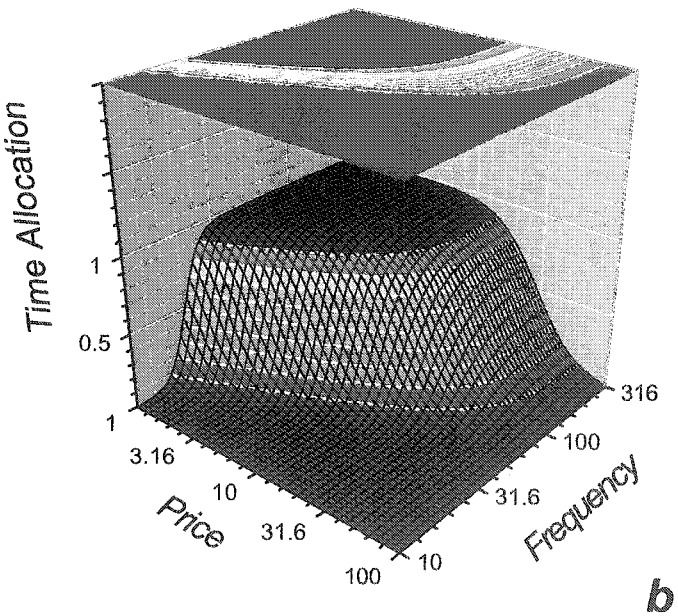

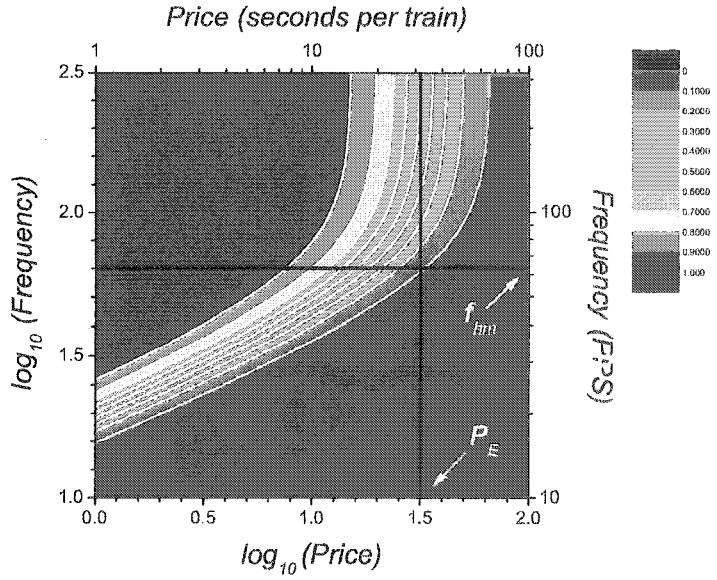

Figure 6a. The wire frame structure of the mountain can be color coded to represent the position of the structure in space.

Figure 6b. Projecting the contours of the mountain upwards onto a surface results in the contour graph presented in *Figure 6c*.

Derivation of the model

One can model the reward-growth function presented in Figure 1 as a logistic:

$$I = I_{max} \times \frac{f^g}{f^g + f_{hm}{}^g}, \quad (1)$$

where f is the frequency of the electrical stimulation,
f_{hm} is the frequency that produces a reward of half-maximal intensity,
g is the parameter that determines the steepness of reward-growth,
I is the reward intensity,
and I_{max} is the maximal reward intensity.
(Shizgal, 2004, is the source for all of the equations presented in this thesis).

The crucial features of this function are that reward intensity initially grows steeply (at a rate determined by g) as frequency is increased but then levels off (at I_{max}) as frequency increases further. The location of the reward-growth function is determined by f_{hm} (frequency that produces a reward of half-maximal intensity).

Figure 2 portrays the behaviour-allocation function with respect to reward strength alone, but in Figure 5 presents it as comprising both the reward strength and the price the rat faces in order to obtain the reward. Essentially, the payoff from BSR is compared with all the payoffs available in the environment.

One can model the behaviour-allocation function in Figure 2 as a second logistic:

$$TA = \frac{U_B{}^a}{U_B{}^a \times U_E{}^a}, \quad (2)$$

where a determines the rate at which time allocation increases as a function of the payoff from BSR,
TA is the time allocation,
U_B is the payoff from BSR,
and U_E is the payoff from everything else (the alternative activities available to the rat such as grooming, resting, or exploring).

This function is the same as Herrnstein's single operant matching law (1970, 1974) with one exception: the addition of the exponent, a. Including this exponent allows us to forgo the controversial assumption that the experimenter-controlled reward is a perfect economic substitute for "everything else" (Rachlin, 1980).

Equations 1 and 2 correspond to the reward-growth and behaviour-allocation functions depicted in Figure 4. In order to connect these two functions, the payoffs in equation 2 must be redefined in terms of the independent variables and the parameters of the reward-growth function.

The model treats payoffs as ratios of a benefit and a set of costs. The benefit is the intensity of the BSR. Costs include both the time required to earn the reward (price) and the perceived effort involved. Thus,

$$U_B = \frac{I}{\xi \times P} \quad , \qquad (3)$$

where I is the reward intensity,
P is the price the rat has to pay (amount of time spent working at the lever) to obtain a train of BSR,
U_B is the payoff from BSR,
and ξ is the perceived effort per unit time expended in holding down the lever.

Can one find a corresponding definition for the value of "everything else?" It follows from Equation 2 that when the time allocation is 50 percent that the value of the BSR is equal to that of "everything else." This equivalence allows one to define the payoff of "everything else" in terms of the payoff from BSR. To do this, one introduces a parameter, P_E, that determines the position of the mountain along the price axis, just as f_{hm} determines the position of the mountain along the frequency axis. Formally, P_E is the price at which the rat allocates 50 percent of its time to working for a maximally intense BSR (I_{max}). Thus,

$$U_E = \frac{I_{max}}{\xi \times P_E} \quad , \tag{4}$$

where I_{max} is the maximal reward intensity,
P_E is the price at which the rat spends 50 percent of its time working for a maximally intense BSR,
U_E is the payoff from "everything else,"
and ξ is the perceived effort per unit time required to obtain the train of BSR.

One can now rewrite Equation 2, substituting from equations 1, 3, and 4 and simplifying:

$$TA = \frac{\left(\frac{f^g}{f^g + f_{hm}{}^g}\right)^a}{\left(\frac{f^g}{f^g + f_{hm}{}^g}\right)^a + \left(\frac{P}{P_E}\right)^a} \quad , \tag{5}$$

where a determines the rate at which time allocation increases as a function of the payoff from BSR,
f is the frequency of the electrical stimulation,
f_{hm} is the frequency that produces a reward of half-maximal intensity,
g is the parameter that determines the steepness of reward-growth,
P is the price of the BSR (average time required to earn a train),
P_E is the price at which the rat spends 50 percent of its time working for a maximally intense BSR,
and TA is the time allocation.

Figure 6 illustrates how time allocation is a function of two independent variables: price and frequency. Equation 5 specifies this function in terms of four parameters. Two of these parameters, a and g, affect the rate at which time allocation increases as a function of the independent variables; the slope along the price axis is determined by a alone, whereas the slope along the frequency axis is determined by both a and g. The two remaining parameters, f_{hm} and P_E, specify the position of the mountain along each axis; measurement of these location parameters is at the core of this experiment.

The mountain models reward processing at its different stages

The primary goal of this experiment is to establish whether the behaviour-allocation function is indeed dissociable from the reward-growth function. If so, then it should prove possible to determine which of these two functions (and their corresponding neural underpinnings) are affected by manipulations that alter performance for BSR. The following section explains the effects on performance for BSR of decreasing the current or increasing the effort required to hold down the lever. These examples illustrate how variables such as different drug treatments, different physiological changes (such as food or water deprivation), or different lesions can produce similar changes in time allocation when acting at different stages of processing. In order to work out the neural circuitry that links the electrical stimulation to behaviour it is necessary to determine the stage of

processing at which each of these manipulations acts. The model allows one to do just that. As the following examples show, the location parameters (f_{hm} and P_E) provide the information required to associate the effect of a manipulation that changes performance for BSR with a given stage of processing.

Lowering the current decreases the number of neurons fired by the stimulation electrode. This action, at the very first stage of processing, results in a decreased rate of firing, and thus in a lower reward intensity (Figure 3b). This decrease in reward intensity constitutes a smaller input to the behaviour-allocation function (Figure 3c), and thus time allocation falls.

Assume that the original current is sufficient to drive the output of the reward-growth function to half its maximal value; following the decrease in current, reward intensity will be less than the half maximal value. Given that the firing rate is the product of the number of neurons stimulated and the rate at which they fire, one could restore the reward intensity to the half-maximal level by increasing the frequency, thus compensating for the decrease in current. This example shows that decreasing the current increases f_{hm}, the frequency that corresponds to the half-maximal reward value. The decrease in current is an example of manipulations that affect the early stages of processing, before the output of the integrator. These manipulations alter the value of f_{hm}.

In contrast to decreasing the current, increasing the effort required to hold the lever down (ξ) affects the later stages of the circuitry. One can see the consequences of this change in required effort by rearranging Equation 4 as follows:

$$P_E = \frac{I_{max}}{\xi \times U_E} \quad . \tag{6}$$

Increasing the effort required to hold the lever down (ξ) makes the denominator larger, thus decreasing the value of P_E. Equation 5 shows that a decrease in P_E decreases time allocation. The change in time allocation results from a change in a parameter of the behaviour-allocation function at a late stage of processing. It is possible to generalize the relation between the perceived effort and time allocation to all manipulations that result in a change in the parameters of the behavioural-allocation function.

The model allows one to determine at what stage various manipulations act by decomposing changes in time allocation into alterations of the location parameters of the mountain: f_{hm} and P_E. This thesis tests the model by assessing the independence of the reward-growth and behaviour-allocation functions and by determining whether it is possible to change f_{hm} without changing P_E. The next section shows why it is not possible to infer changes in f_{hm} and P_E unambiguously from traditional 2D representations.

Why the 2D representations are ambiguous.

In Figure 7a, the silhouette of the mountain along the darker grey wall (time allocation versus frequency) is drawn in yellow. This wall is positioned at a very low price (1 second). Thus, the yellow curve corresponds roughly to a conventional 2D psychometric function obtained under conditions of continuous reinforcement. A manipulation, such as decreasing the current that affects the circuitry before the output of

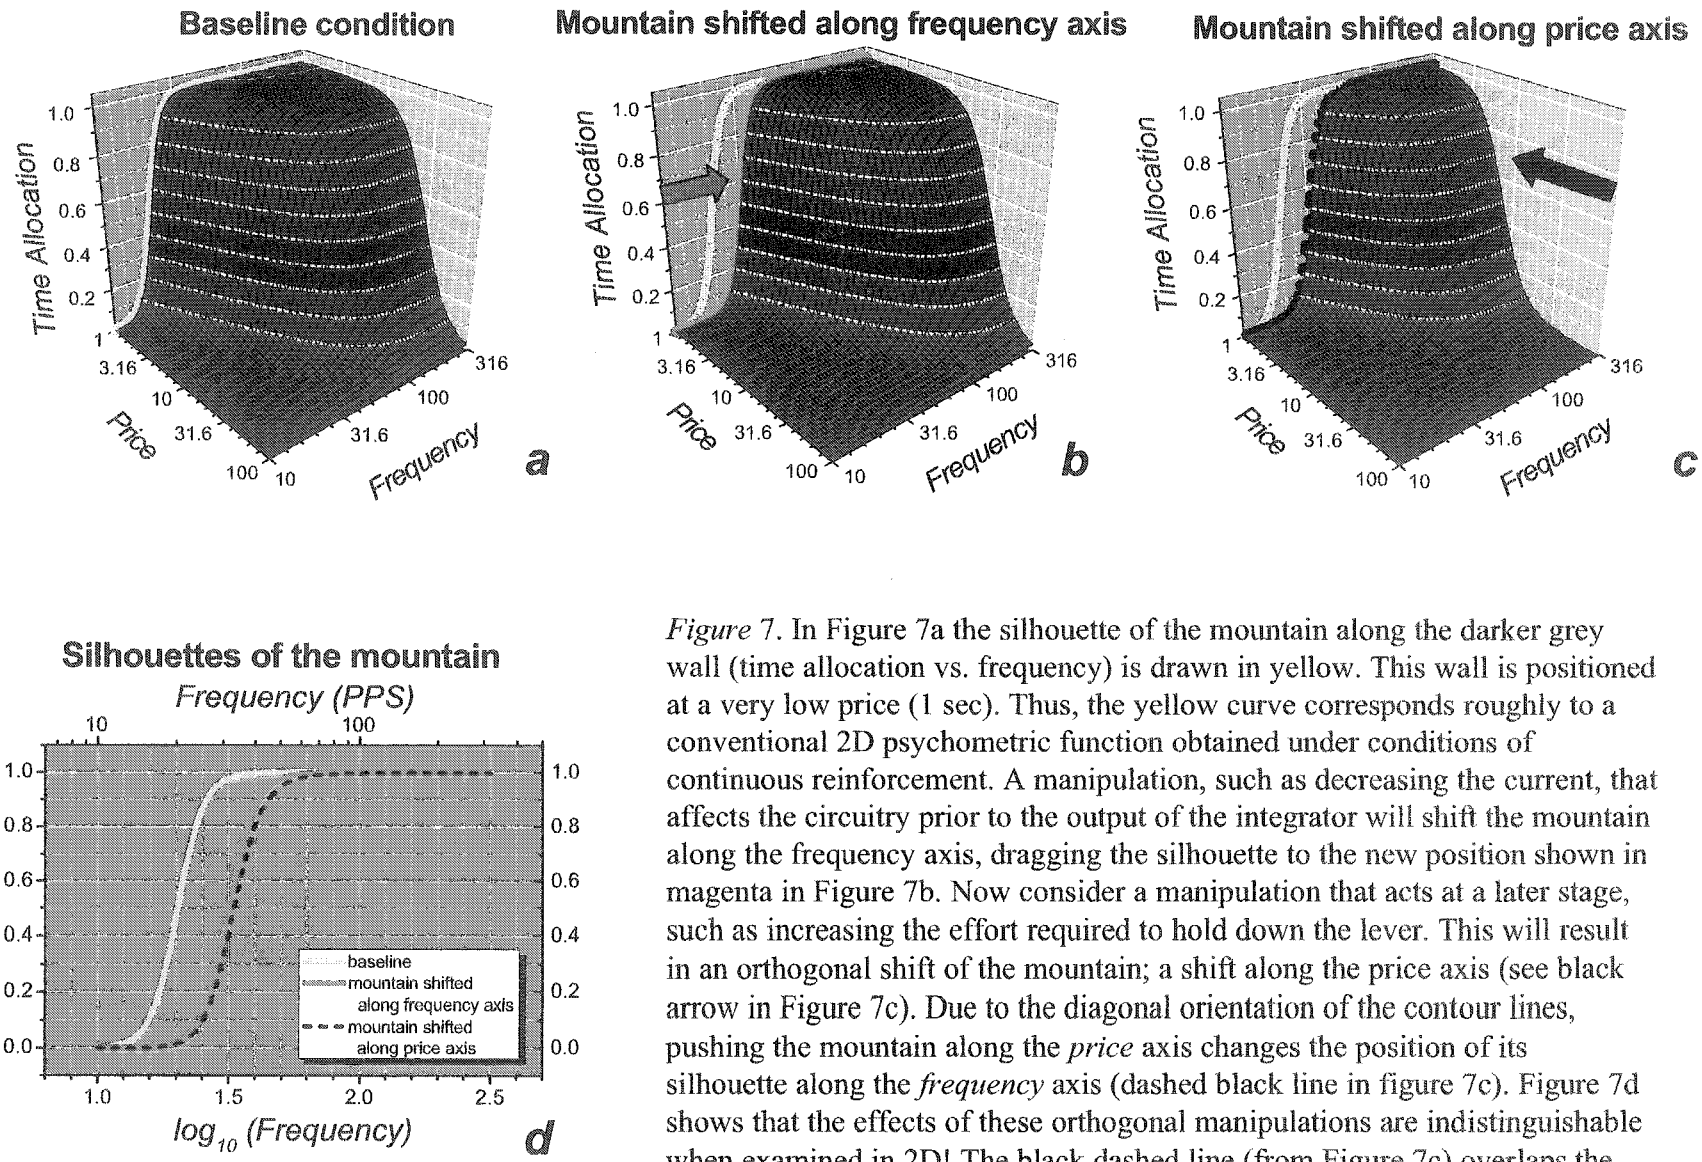

Figure 7. In Figure 7a the silhouette of the mountain along the darker grey wall (time allocation vs. frequency) is drawn in yellow. This wall is positioned at a very low price (1 sec). Thus, the yellow curve corresponds roughly to a conventional 2D psychometric function obtained under conditions of continuous reinforcement. A manipulation, such as decreasing the current, that affects the circuitry prior to the output of the integrator will shift the mountain along the frequency axis, dragging the silhouette to the new position shown in magenta in Figure 7b. Now consider a manipulation that acts at a later stage, such as increasing the effort required to hold down the lever. This will result in an orthogonal shift of the mountain; a shift along the price axis (see black arrow in Figure 7c). Due to the diagonal orientation of the contour lines, pushing the mountain along the *price* axis changes the position of its silhouette along the *frequency* axis (dashed black line in figure 7c). Figure 7d shows that the effects of these orthogonal manipulations are indistinguishable when examined in 2D! The black dashed line (from Figure 7c) overlaps the magenta line (from 7b).

the integrator, will shift the mountain along the frequency axis, dragging the silhouette to the new position shown in magenta in Figure 7b. Now consider a manipulation that acts at a later stage, such as increasing the effort required to hold down the lever. The later manipulation will result in an orthogonal shift of the mountain: a shift along the price axis (see black arrow in Figure 7c). Owing to the diagonal orientation of the contour lines, pushing the mountain along the *price* axis changes the position of its silhouette along the *frequency* axis (black line in figure 7c). Figure 7d shows that the effects of these orthogonal manipulations are indistinguishable when examined in 2D! The black dashed line (from Figure 7c) overlaps the magenta line (from 7b). This example shows that conventional 2D representations of performance for BSR fail to distinguish the effects of manipulations that act at different stages of the underlying circuitry.

Figures 8 and 9 show that the shifts that are confounded in the 2D representation in Figure 7d are readily distinguished in the 3D representations. Figures 8b and 8c and Figures 9b and 9c correspond to the yellow silhouette in Figure 7, Figure 8d corresponds to the magenta silhouette in Figure 7, and Figure 9d corresponds to the dashed black silhouette. Comparison of the lines denoting f_{hm} and P_E shows unambiguously that the mountain has shifted along the *frequency* axis in Figure 8 and along the *price* axis in Figure 9. The shifts that were indistinguishable in the 2D representation in Figure 7d are now clear.

Testing the model by varying the train duration

If the reward-growth and behaviour-allocation functions differ from one another, then it should be possible to alter one without affecting the other. Figure 8 shows the effect of an early-stage manipulation that alters the location parameter of the reward-

Figure 8. shows the effect of an early-stage manipulation that alters the location parameter of the reward-growth function, f_{hm}, leaving unchanged the location parameter of the behavioural-allocation function, P_E. The shifts that are confounded in the 2D representation in Figure 7d are readily distinguished in the 3D representations. To facilitate comparison with the maps of the shifted mountain, the contour graph corresponding to the yellow silhouette in Figure 7 is shown twice, in the upper right and lower left panels of Figures 8. The contour graph in the lower right panel of Figure 8 corresponds to the magenta silhouette in Figure 7. Comparison of the lines denoting *fhm* and P_E show unambiguously that the mountain has been shifted along the *frequency* axis in Figure 8. The shifts that were indistinguishable in the 2D representation in figure 7d are now abundantly clear.

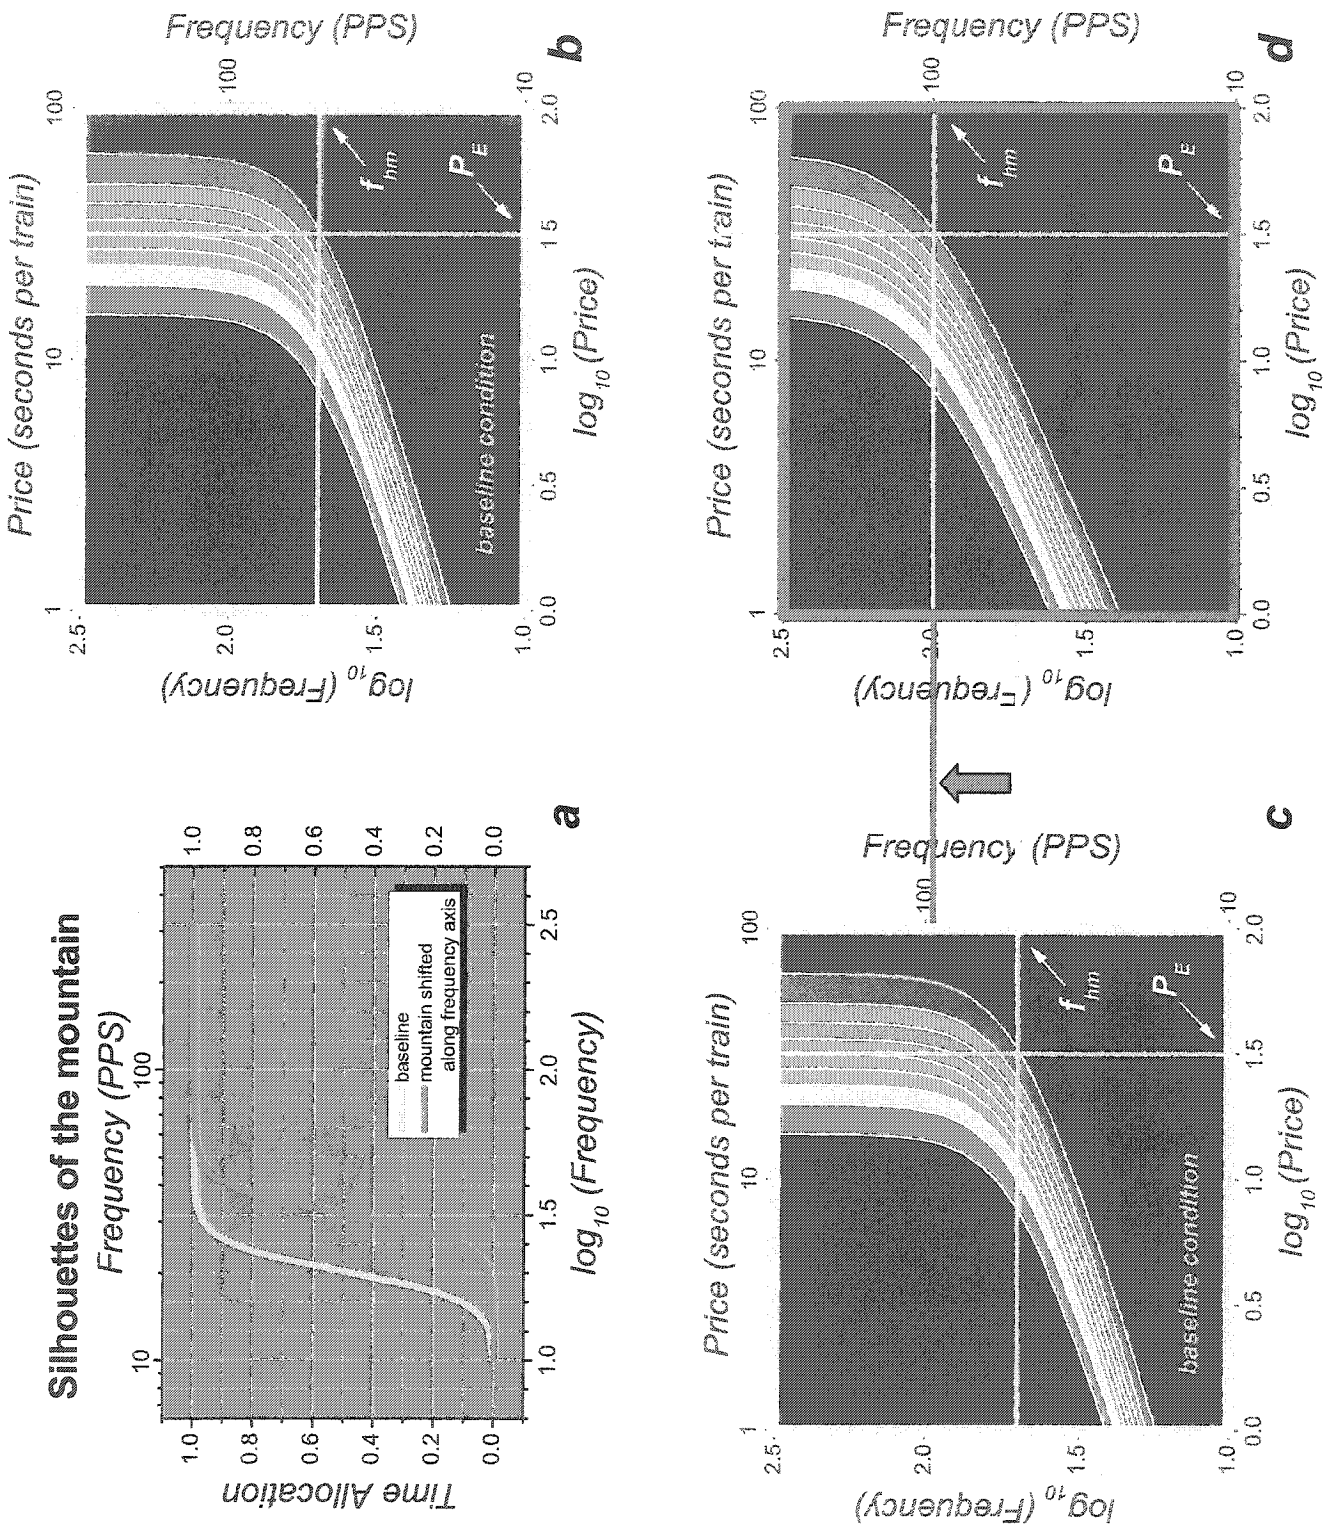
Silhouettes of the mountain
Frequency (PPS)
Time Allocation
baseline
mountain shifted along frequency axis
log10 (Frequency)
a
Price (seconds per train)
Frequency (PPS)
log10 (Frequency)
log10 (Price)
baseline condition
fhm
PE
b
c
d

*Figure 9. T*he shifts that are confounded in the 2D representation in Figure 7d are readily distinguished in the 3D representations. To facilitate comparison with the maps of the shifted mountain, the contour graph corresponding to the yellow silhouette in Figure 7 is shown twice, in the upper right and lower left panels of Figures 9. The contour graph in the lower right panel of Figure 9 corresponds to the dashed black silhouette. Comparison of the lines denoting *fhm* and P_E show unambiguously that the mountain has been shifted along the *price* axis in figure 9. This corresponds to a later stage manipulation that alters the location parameter of the behavioural-allocation function, P_E, leaving unchanged *fhm*, the location parameter of the reward-growth function. The shifts that were indistinguishable in the 2D representation in figure 7d are now abundantly clear.

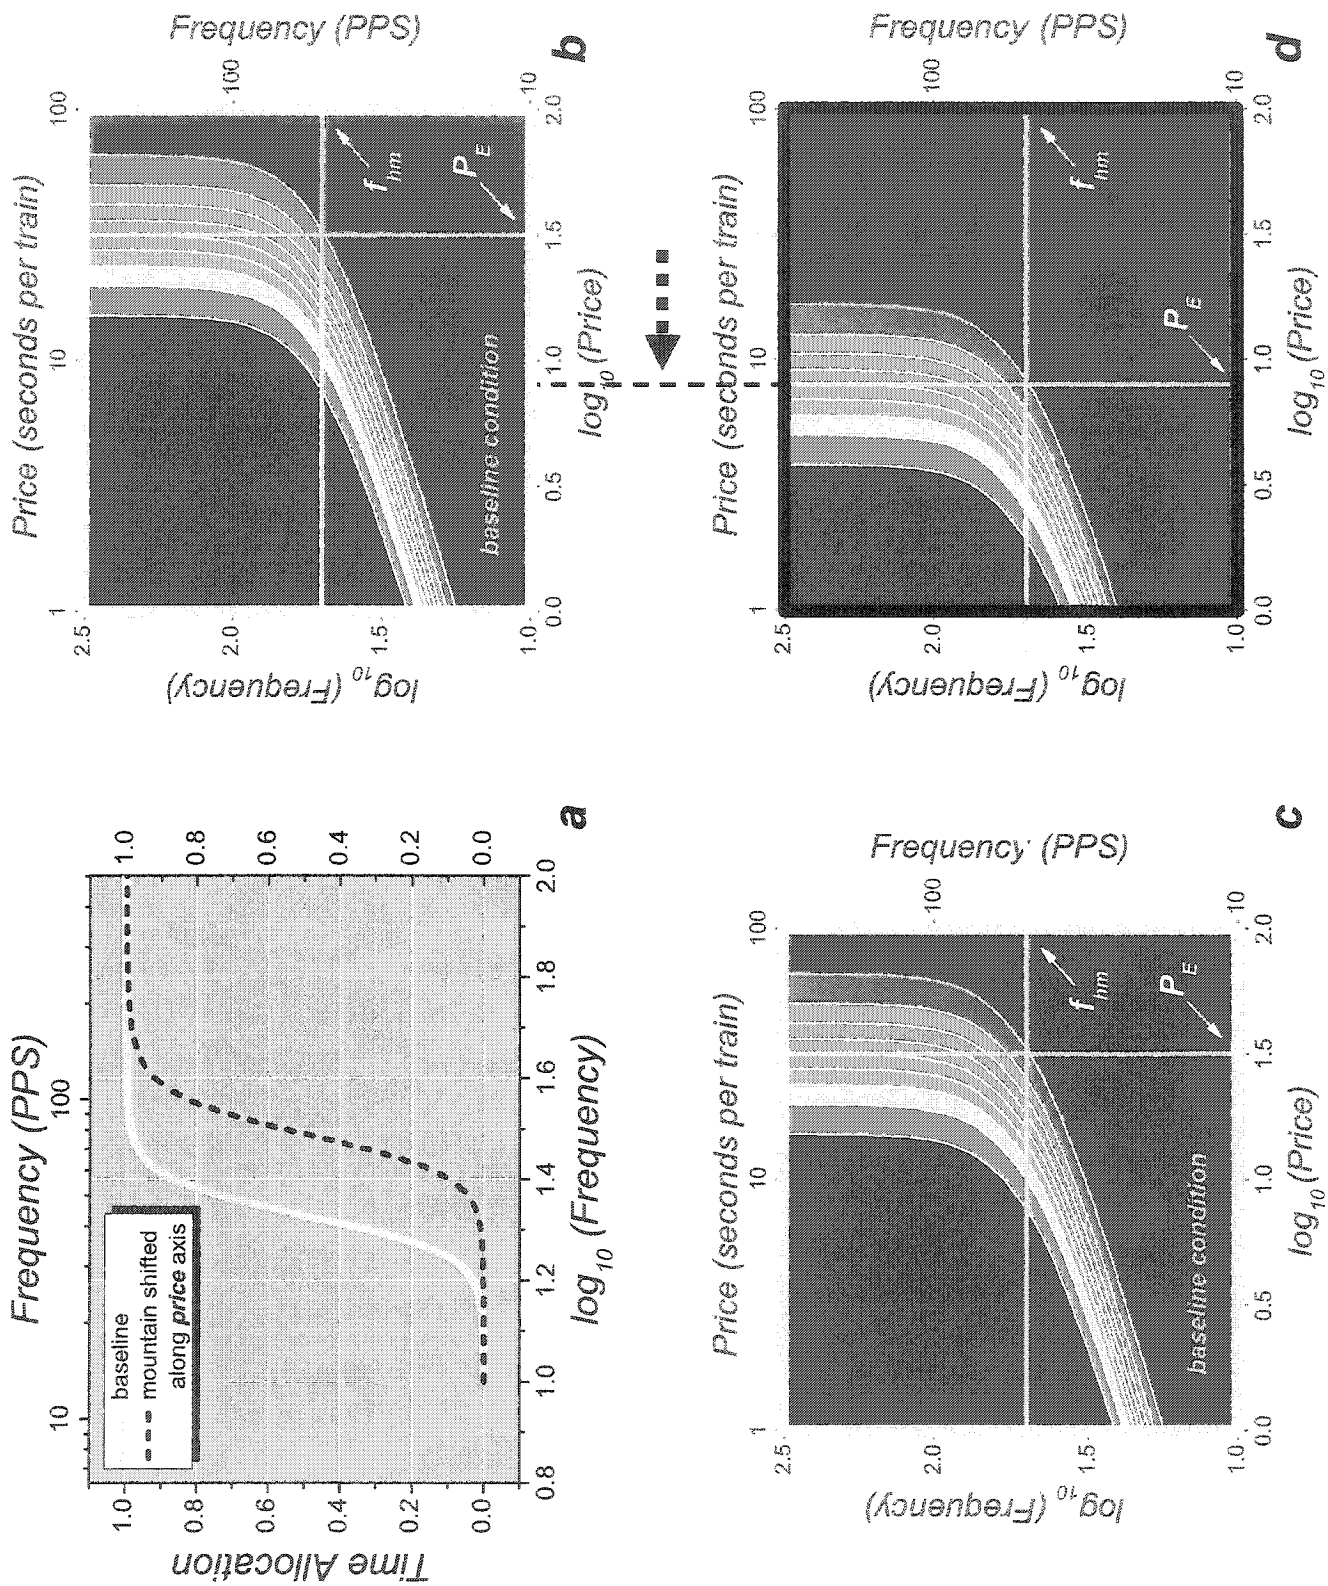
Frequency (PPS)
Price (seconds per train)
log10 (Frequency)
log10 (Price)
baseline condition
fhm
PE
a
b
c
d
Time Allocation
baseline
mountain shifted
along price axis

growth function, f_{hm}, but leaves unchanged the location parameter of the behaviour-allocation function, P_E. Figure 9 shows the effect of a later-stage manipulation that alters the location parameter of the behaviour-allocation function, P_E, but leaves unchanged f_{hm}, the location parameter of the reward-growth function. If it is possible to dissociate the reward-growth and behaviour-allocation functions, then it should be possible to produce changes such as those shown in both Figures 8 and 9. The experiment described below uses a manipulation that was expected to produce effects such as those seen in Figure 8-that is, a change in f_{hm} and not P_E.

A decrease in the duration of the train of stimulation pulses decreases the time during which excitation in the reward substrate can build up. To achieve the same reward intensity as a decrease in the train duration, it is necessary to increase stimulation frequency; the result is a shift of the position of the mountain along the frequency axis. Since the increase in frequency has compensated for the decrease in train duration, there is no change in the behaviour-allocation function. Arvanitogiannis (1997) carried out a similar experiment in which a traditional VI schedule was employed, and as such, the price of the BSR could not be strictly controlled.

Equation 6 defines P_E, the location parameter of the behaviour-allocation function, in terms of three variables: *Imax* (maximal reward intensity), ξ (perceived effort), and U_E (payoff from alternate activities such as grooming or resting). A decrease in the train duration should not change *Imax*, since the increase in frequency compensates for the reduced integration time. The duration of the train should not affect either the effort required to hold down the lever or the payoff from alternative activities. Therefore, decreasing the train duration should not change P_E and, as Figures 8b and 8d show, the

mountain should not shift along the price axis. The experiment described below tests this hypothesis.

A practical issue: How do we control price?

Recall that price is defined as the time that the rat has to spend in holding down the lever in order to earn a reward. Conover and Shizgal (2004) describe a schedule of reinforcement that provides precise control of the average price. They call this the "free-running variable interval schedule" (FVI). In order to reap a reward, the rat has to be holding the lever down at the moment in which an unsignaled interval sampled from an exponential distribution ends. The FVI schedule forces the rat to decide at any point in time whether to work at the lever or not. Under a FVI schedule, unlike a traditional variable interval (VI) schedule, the reward does not wait for the rat: the lever does not remain armed when the interval has elapsed. On a VI schedule, the cumulative probability that the lever is armed increases over time. The longer the rat waits, the higher is the likelihood that it can earn a reward without having to sacrifice leisure time. By contrast, the rat's average "earnings" on the FVI schedule are proportional to its time allocation.

Notion of subjective and objective price

The reward-growth function translates an objective variable, stimulation frequency, into a subjective variable, reward intensity. Could an analogous function operate in the time domain? Whereas the distinction between a price of 10 seconds and a price of 20 seconds is likely quite significant to the rat, a distinction between a price of 0.0010 seconds and a price of 0.0020 seconds is likely to be irrelevant. However, as the

contour representation shows (Figure 6), changing the price from 0.0010 to 0.0020 seconds would shift a 2D psychometric function (orange curve in Figure 4) by about the same amount as a change in price from 10 to 20 seconds. Clearly this is unreasonable: the rat's ability to discriminate changes in price must break down at some point. Thus a function is needed to translate objective prices into their subjective equivalents. Shizgal, (personal communication) has suggested such a function:

$$P_S = a + b \times \ln\left(1 + e^{\frac{P_O - a}{b}}\right), \qquad (7)$$

where a equals minimum subjective price,
b determines the sharpness of the transition between the horizontal and diagonal components (Figure 10),
P_O is the objective price,
and P_S is the subjective price.

According to this equation, represented in Figure 10, the subjective and objective prices correspond when the objective price exceeds a critical value, but the subjective price approaches an asymptote (a) as the objective price gets smaller and smaller. Equation 8 incorporates into Equation 5 the concept of a subjective price. In Equation 8, time allocation is now a function of the subjective price, P_S:

$$TA = \frac{\left(\frac{f^g}{f^g + f_{hm}{}^g}\right)^a}{\left(\frac{f^g}{f^g + f_{hm}{}^g}\right)^a + \left(\frac{P_S}{P_E}\right)^a}. \qquad (8)$$

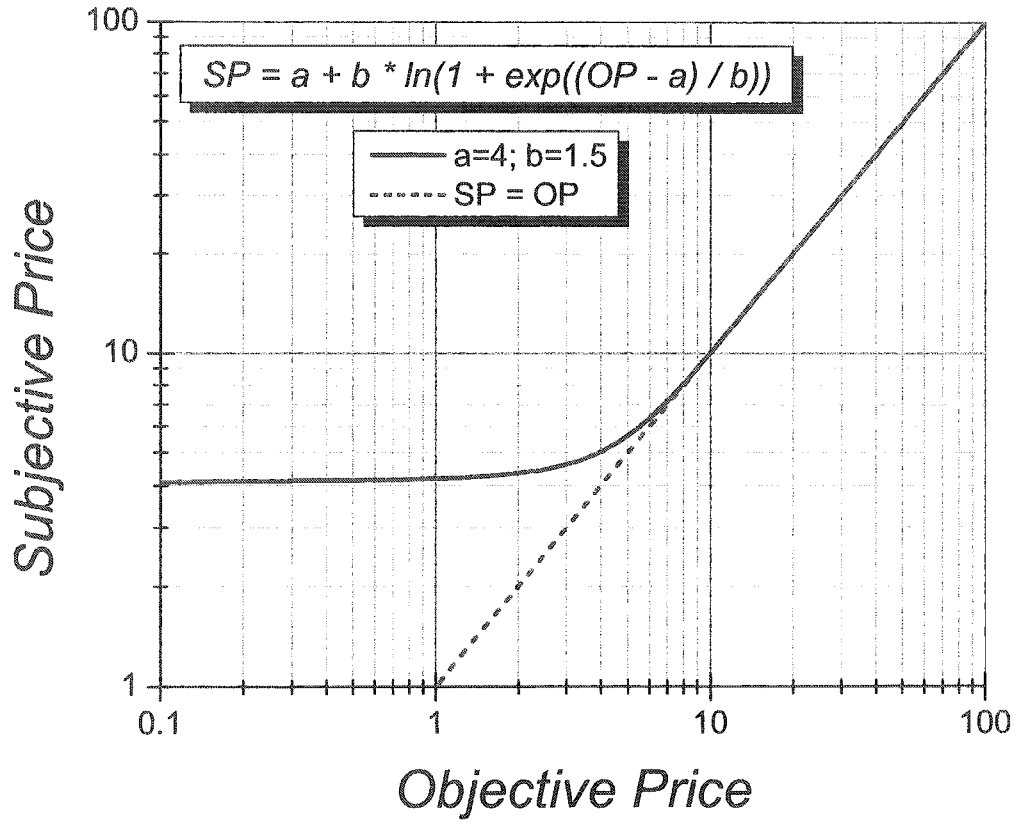

Figure 10. Comparison of the objective price to the subjective price. At low prices the effect of altering price has no subjective effect to the rat and corresponds to a leveling off of the subjective price function.

The experiment

The primary goal of this experiment is to establish whether it is possible to distinguish changes in the behaviour-allocation function from changes in the reward-growth function -- that is, whether it is possible to shift the mountain along only one of its axes. If it is possible, then it should also be possible to determine, in future experiments, at what stages of the circuit drugs, lesions, and physiological variables will alter the rewarding effect of electrical brain stimulation. The experiment determines whether a decrease in the train duration will, as predicted, shift the mountain along the frequency axis but not along the price axis.

Method

Subjects

The subjects were five male Long-Evans rats from Charles River Breeding Farms (St-Constant, Quebec). They were housed individually on a 12h/12h reverse-light cycle. The rats had *ad libitum* access to food and water while in their home cages. At the time of surgery, the subjects weighed between 350 and 500g. Testing and surgery were conducted during the dark phase of the light cycle.

Materials

Electrodes were made from 000 insect pins, coated with Formvar to within 0.5 mm of the tip. A male Amphenol pin was soldered to the blunt end of the insect pin and inserted into a 9-pin cylindrical connector (Ginder Scientific, Ottawa, Ontario) before surgery.

Surgery

The bilateral monopolar electrodes were implanted into the medial forebrain bundle (MFB) with the following stereotaxic coordinates: 2.8 mm caudal to Bregma, 1.7 mm lateral to the midline, and 8.3 mm ventral to the dura mater for the left electrode and 8.4 mm for the right electrode. Surgery was performed under sodium pentobarbital anesthesia (Somnotol, 60 mg/kg, IP). In order to reduce bronchial secretions, atropine sulfate (1 mg, SC) was administered 15 minutes before the Somnotol. Anesthesia was maintained with 2 percent Isofluorane. Dental acrylic was applied to anchor the connector assembly to the skull and to five jeweler's screws threaded into the frontal and parietal bones. The return wire was wound around the three most rostral screws. A

subcutaneous injection of buprenorphine (0.05 ml) was administered post-operatively to minimize discomfort.

Behavioural testing

Training.

A minimum of five days after surgery, the rats were taught to lever press for electrical brain stimulation, using standard shaping techniques. The duration of the stimulation train was 1 second, and the train consisted of 0.1 millisecond cathodal, rectangular pulses. The frequencies and current were adjusted to determine the working range for each rat. Initially, all stimulation sites were screened with a 300 μA current and a frequency of 19 Hz. If the subject showed no signs of forced movement or aversion, both the current and the frequency were increased until the rat responded vigorously.

Apparatus

The operant chambers (23 cm deep, 34 cm wide, and 60 cm high) were made of grey polyvinyl chloride (PVC) and had a wire-mesh floor. Each chamber had a clear Plexiglas front panel with a hinged door. Two retractable rodent levers (Med Associates Inc., GA) were installed approximately 5 cm above the floor, one on the right wall and the other on the left wall. For the purpose of this experiment only one of the two levers was employed. Above each lever was a yellow "jewel" light approximately 1.5 cm in diameter. This light was lit during the experiment whenever the retractable lever was extended. If the rat harvested the reward, the lever was retracted and the computer timed an 8-second "black-out delay" before the lever was extended again. An orange house

light, located on the back wall, 25 cm from the top of the box, flashed for 10 seconds to signal the beginning of the session and the inter-trial interval.

Procedure

The experiment used an A-B-A design (Figure 11). Each phase of the experiment consisted of repeated sweeps through values of one stimulation parameter (either the stimulation frequency or price) while the remaining parameters were held constant. The dataset for each phase included a set of frequency sweeps at a low-price (4 seconds), a set of price sweeps (at the highest stimulation frequency), and a second set of frequency sweeps, at a high price (ranging between 12.6 and 18 seconds, depending on the rat) (Figure 11). Sets of such frequency and price sweeps were collected first at a train duration of 1 second (A: first baseline), then at a train duration of 0.25 seconds (B: experimental phase), and finally at a train duration of 1 second (A': second baseline). Bracketing the experimental phase of the experiment permitted an assessment of baseline reliability.

In the case of the frequency sweeps, the rat was given the opportunity to obtain an average of 20 rewards at each frequency, whereupon the frequency was lowered. The initial spacing between frequencies was 0.05 $\log_{10}$ units (~12.2 percent); this value was adjusted throughout the experiment in order to sample the rate-frequency curves optimally for each rat and condition. A 2D plot of the data from a frequency sweep yields a sigmoidal curve. To sample a sigmoidal curve adequately, data are required from the rising portion, the upper asymptote, and the lower asymptote. Thus, the sequence of frequencies must be spaced to provide information about the position of all

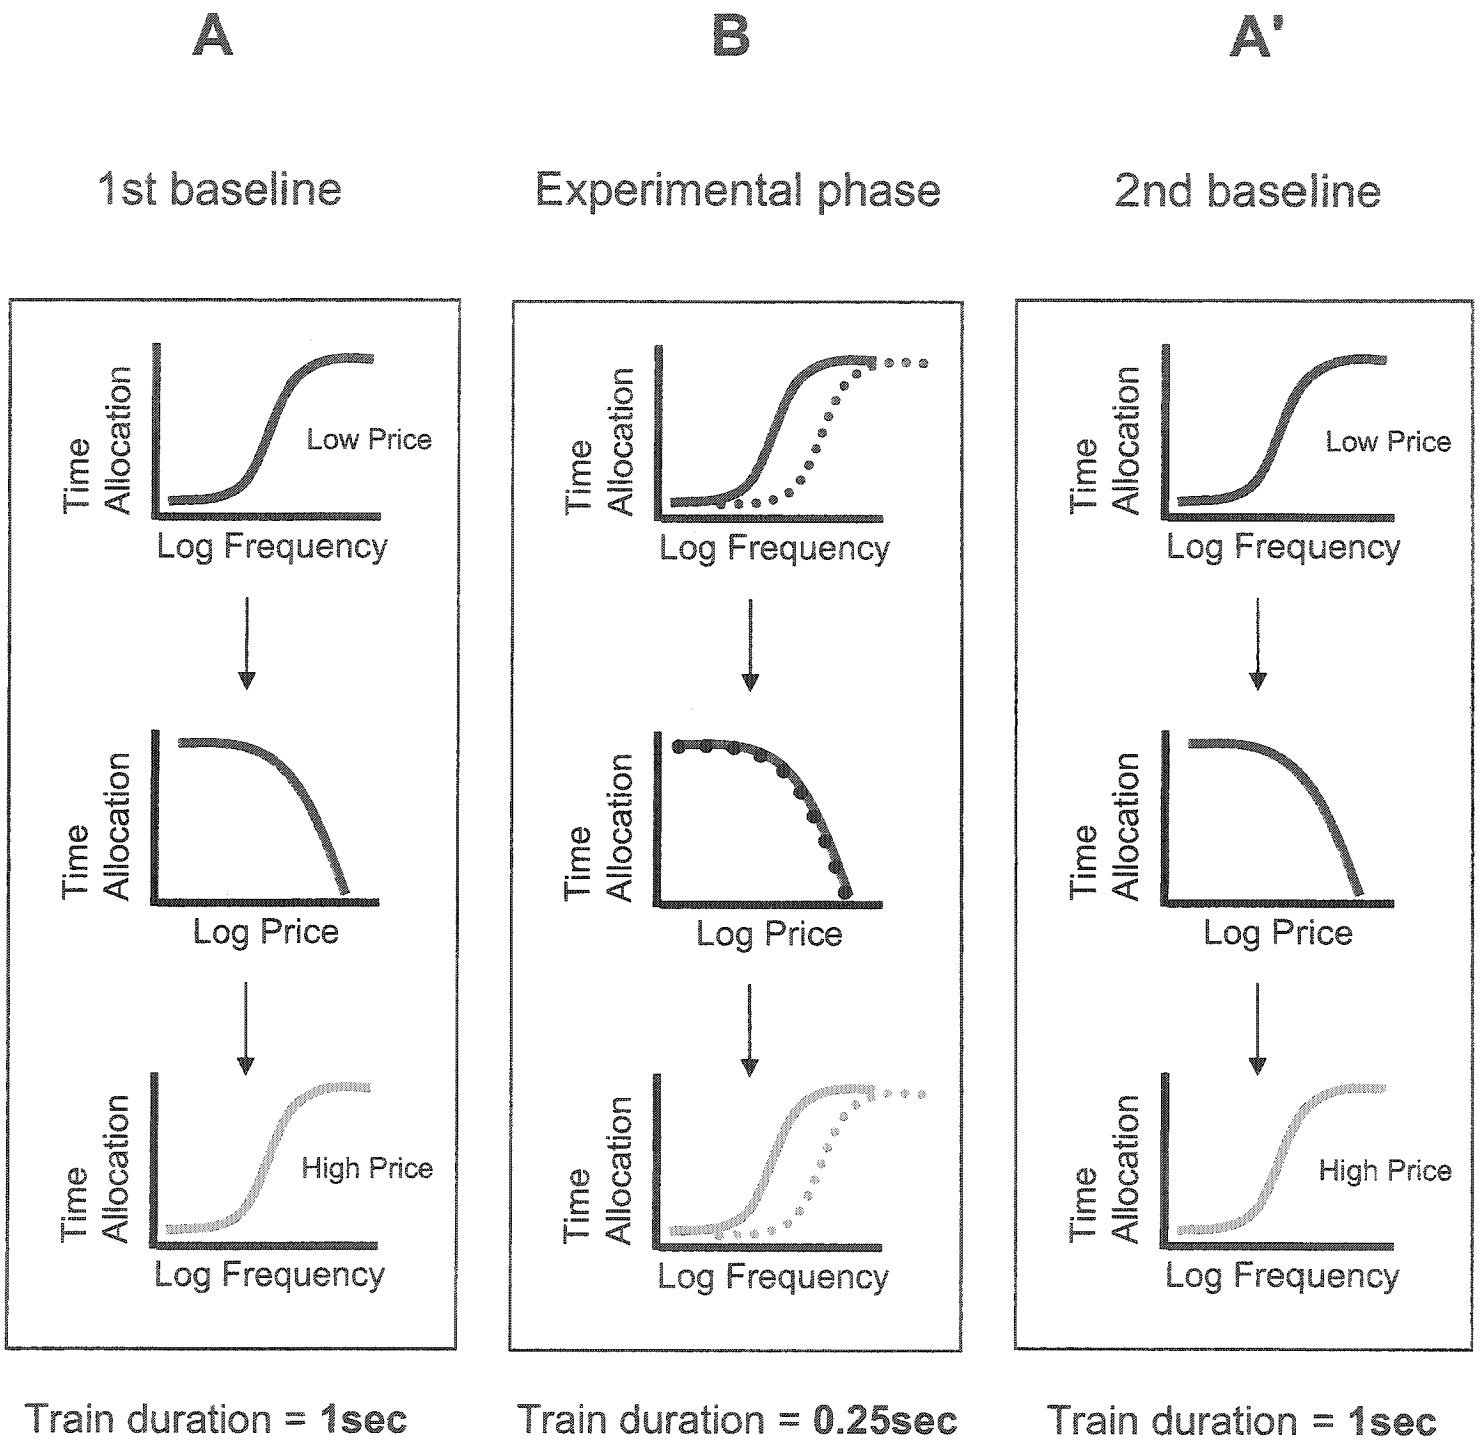

Figure 11. A B A' design and predictions. Each phase of the experiment is represented above. The 1st baseline condition was conducted first, followed by the experimental phase and then the 2nd baseline. The dotted lines in the experimental phase represent our expected result whereas the solid lines represent the expected results from the baseline conditions.

three parts of the sigmoidal curve: the steeper is the rising segment, the closer is the spacing of the frequencies (Figure 12).

Sessions devoted to low-price frequency sweeps consisted of seven sweeps, each comprising ten trials. The first and second trial of each sweep were identical; the first trial was treated as a warm-up, and the data from this trial were discarded. Each of the last eight trials tested a successively lower frequency. Figure 13b shows the average of the time-allocation values obtained at each frequency in the last six sweeps per session.

In the case of the price sweep, two or three sweeps were acquired per session and the data from the necessary number of sessions were then combined to yield six sets of sweeps. Figure 14c shows the average result from six price sweeps.

The high-price frequency sweeps were run in sets of four, and the final six curves to be kept were drawn from sweeps two through four from two different sessions (Figure 15c). As in the case of the low-price frequency sweeps, the first high-price frequency sweep of the session was discarded as a warm-up.

There were a few exceptions to the procedures detailed above. Some of the initial data were collected in sets not easily combined into six. Rats M11 and M16 experienced five sets of sweeps per trial for the low-price frequency sweep for the A phase of the experiment. In order to come up with the final set of six sweeps, the first two sweeps in each session were discarded.

Two sessions were run per day, and the rats were given a one-hour break between sessions. During the break, each rat was placed alone in a shoebox cage with free access to food and water. The room lights were extinguished during the test sessions and the break.

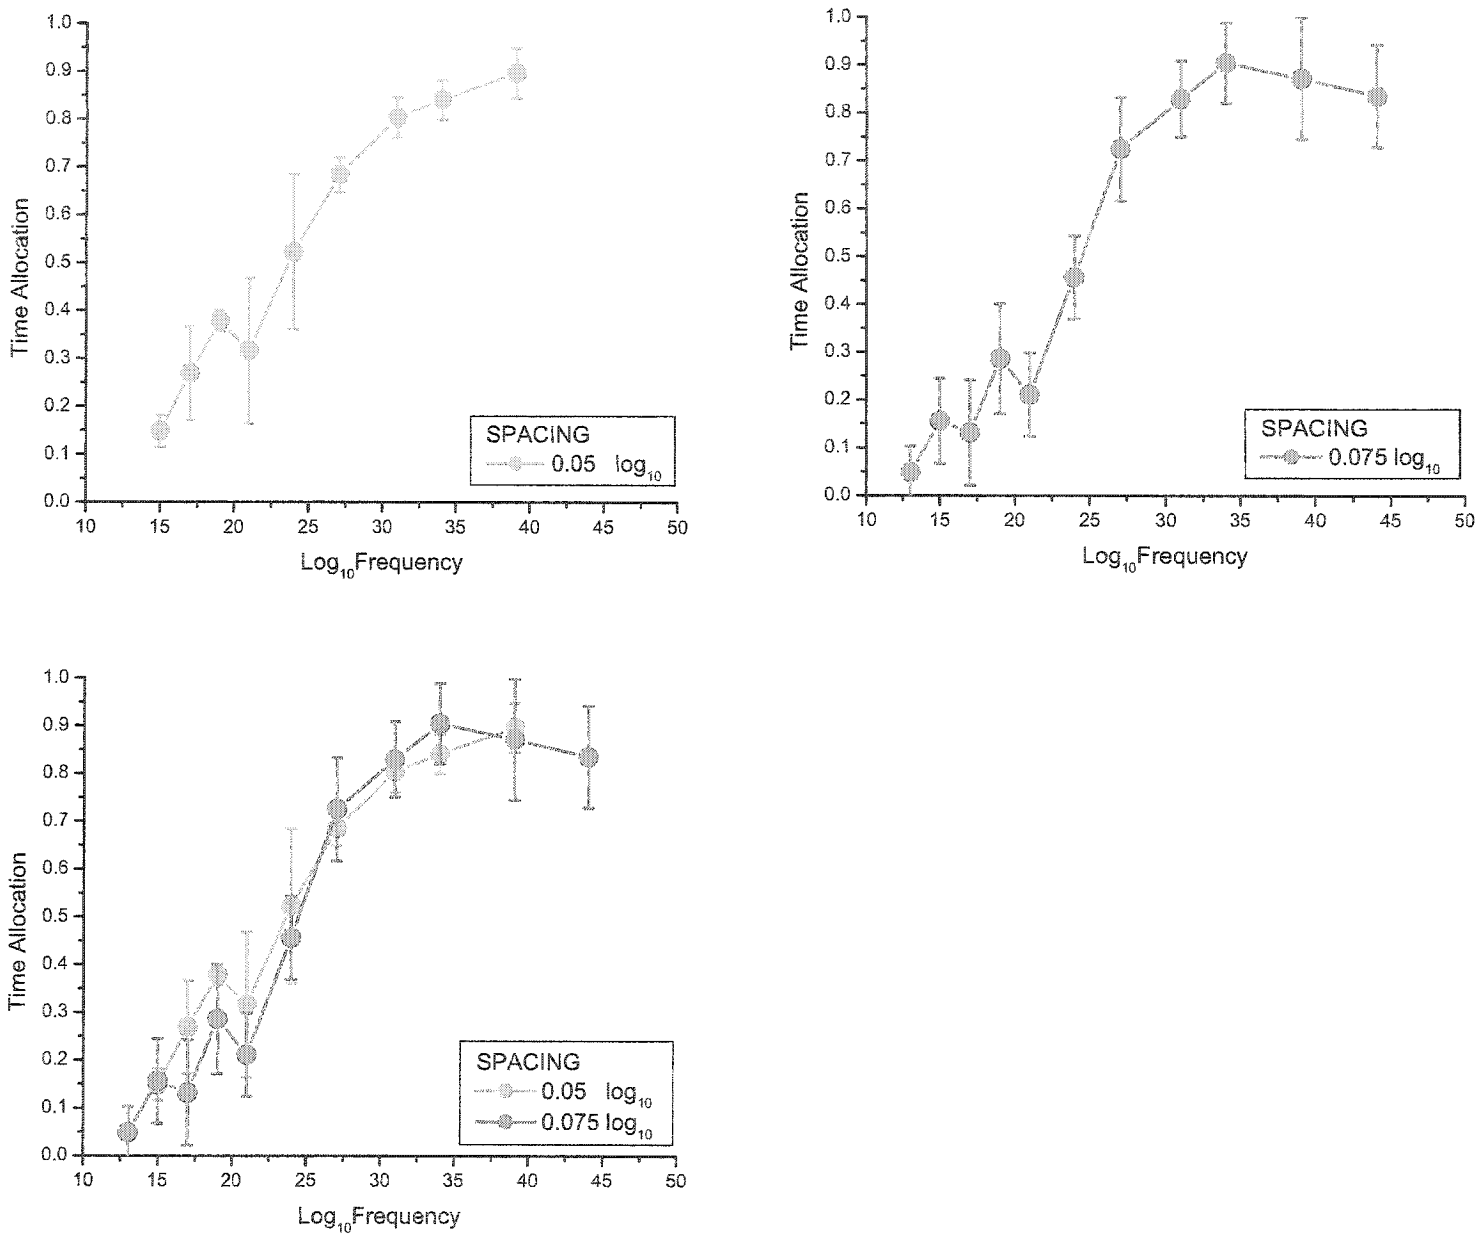

Figure 12. Training data is presented for subject M14. By increasing the spacing between stimulation frequencies from 0.05 log to 0.075 log, the sigmoidal curve is more thoroughly represented. The upper left panel was obtained with a spacing of 0.05 log units between frequencies whereas the upper right panel the spacing was 0.075 log apart. The lower panel presents both sets together in one figure.

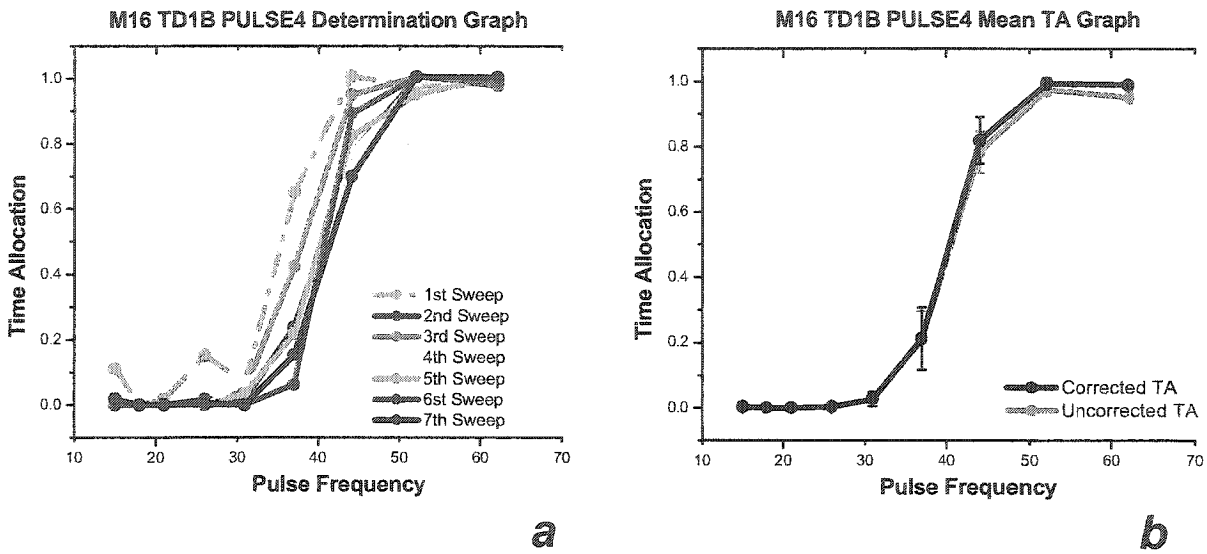

Figure 13a. One set of frequency sweeps from subject M16. From the second baseline condition at a 4 sec price. The first sweep is discarded as a warm-up and shaded grey in this figure.

Figure 13b. The mean time-allocation graphs that correspond to the sweeps for subject M16 in figure M3a. The corrected data are presented in black and the uncorrected in grey.

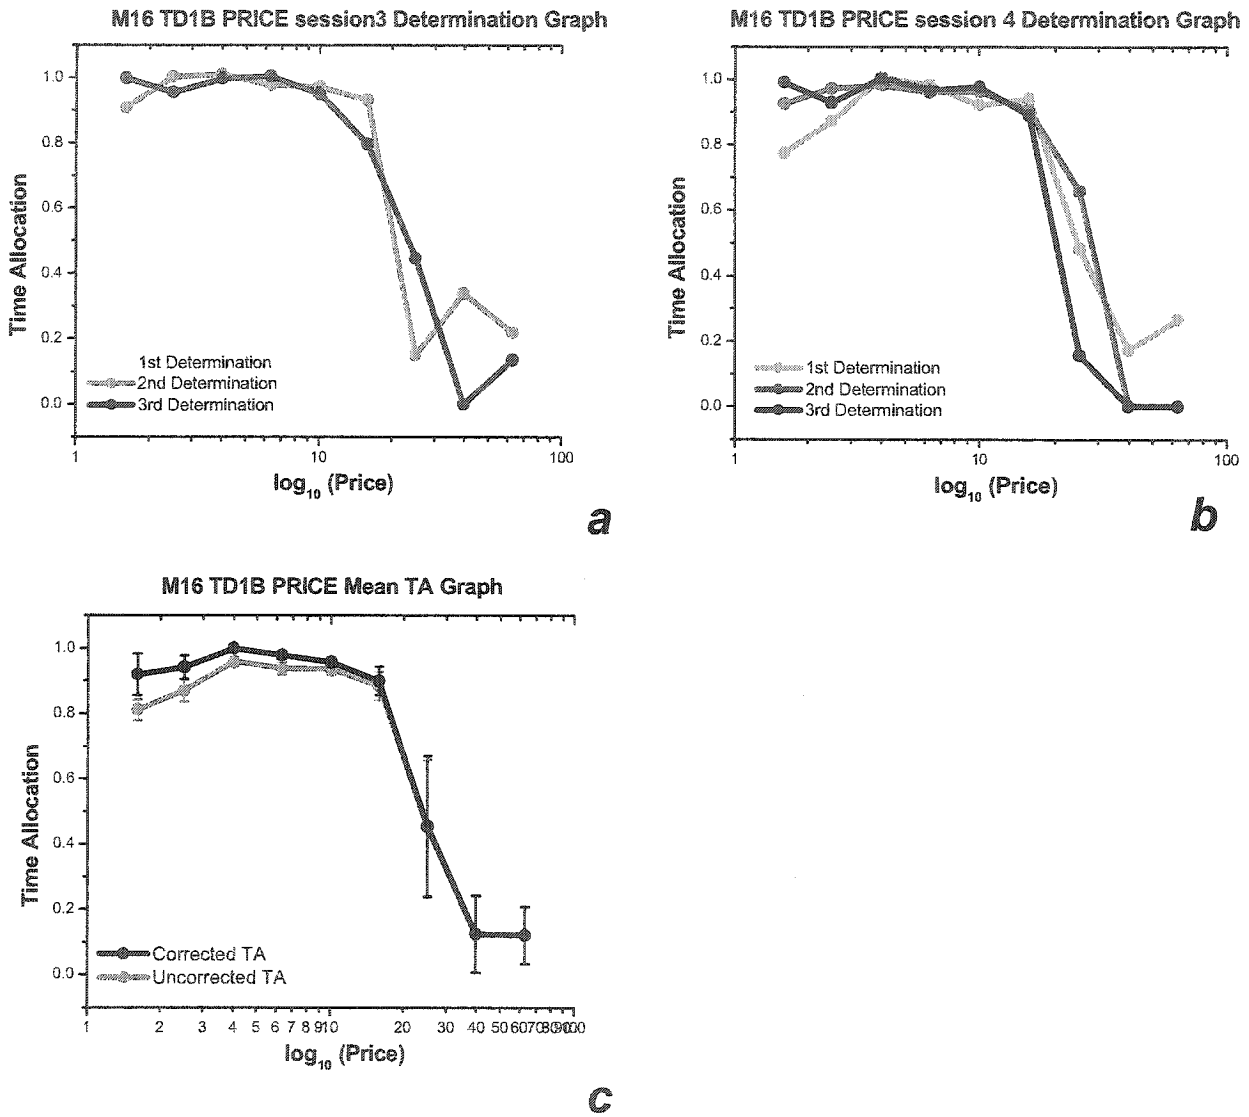

Figure 14a & b. Two different sessions were combined for the price sweep data. There was no warm-up condition for the price sweeps.

Figure 14c. This is the mean time-allocation graph. It is the mean of the raw data from the two sessions above.

The data was corrected for a bias, releases of the lever that were less than 1 sec in duration were re-classified as work time. The black line on the graph represents the time allocated to working at the lever for BSR after the correction whereas the grey line is uncorrected.

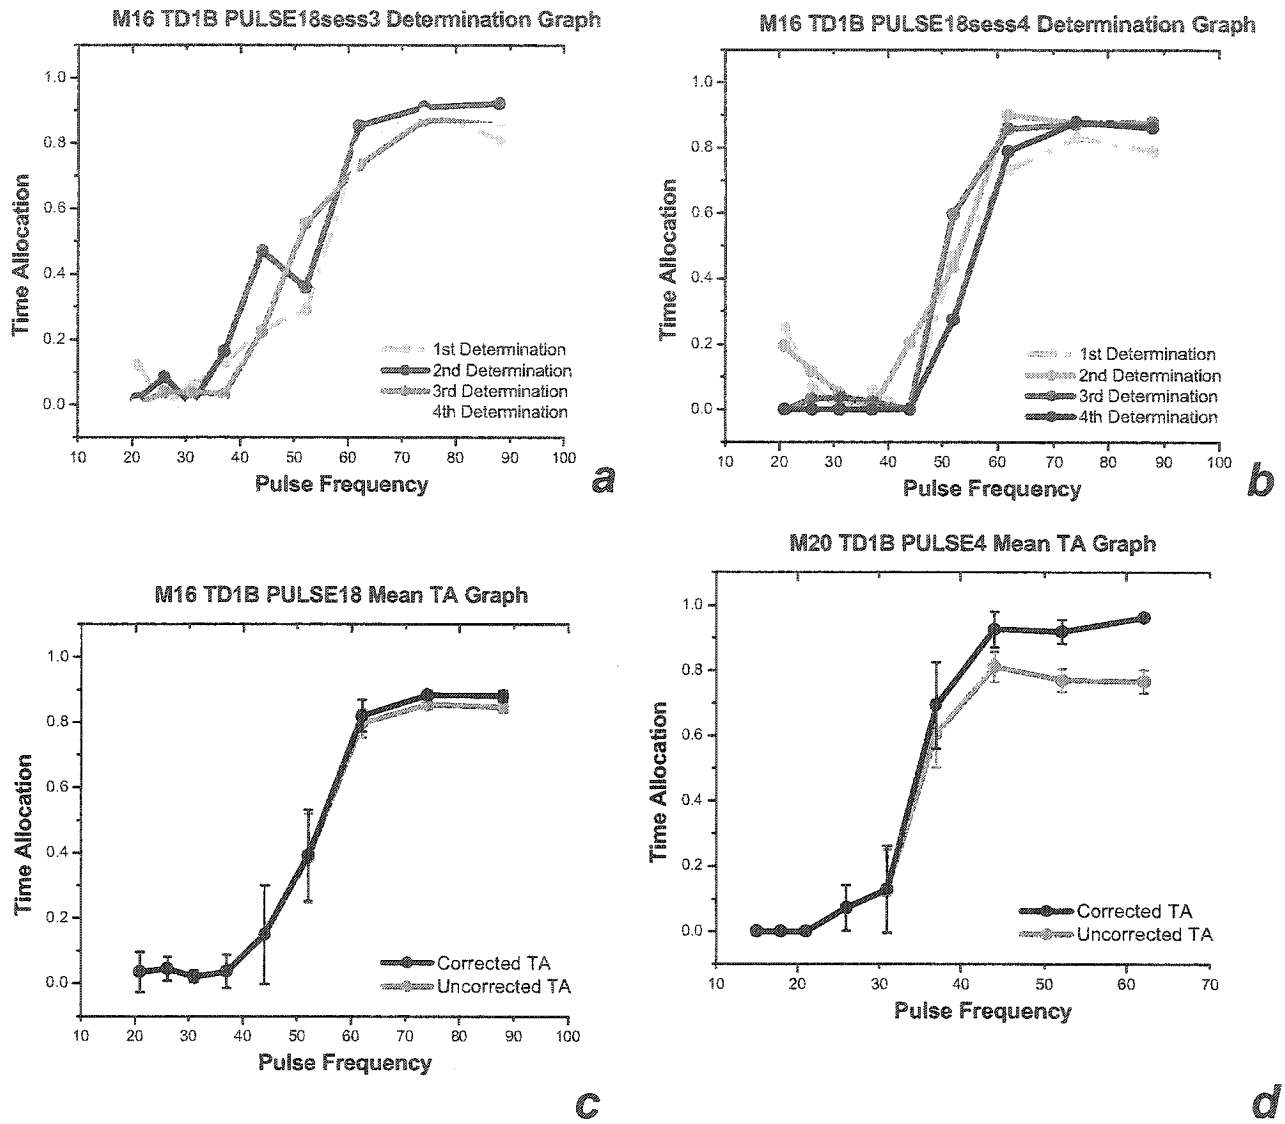

Figure 15a & b. At the higher price two sessions were combined (multiple sessions required to obtain 6 final curves). The grey lines in each sweep graph were dropped as the first curve of every set of frequency *sweeps* was considered to be a warm-up.

Figure 15c. This is the mean time-allocation graph. It is the mean of the raw data from the two sessions above with the warm-up determination removed.

The data was corrected for a bias, releases of the lever that were less than 1 sec in duration were re-classified as work time. The black line on the graph represents the time allocated to working at the lever for BSR after the correction whereas the grey line is uncorrected.

Figure 15d. M20 is an example of a subject who's interaction with the lever involved many more short breaks (more so than subject M16) that are corrected for. The effect of the correction in such a case is greater than the effect of the correction in Figure R2c.

Data acquisition and analysis

Every change in the state of the lever was recorded. Each depression of the lever was considered "work," and conversely each release was categorized as "leisure." The tendency of the rats to press the lever stochastically in bouts introduced a bias in the measurement of work and leisure time, causing work time to be underestimated and leisure time to be overestimated. It is unreasonable to consider the shorter time intervals (< 1 sec) separating work bouts as leisure time. There is not enough time available to engage in other activities when the lever is released so briefly, and the rat remains at the lever. To correct for this bias, releases that were less than 1 second in duration were re-classified as "work" (Conover & Shizgal, 2002).

Time allocation was computed by dividing the corrected work time by the trial time (the sum of work and leisure time).

Exploratory analyses were conducted daily in order to assess the stability of performance. Two-dimensional plots of the data (from a set of either frequency sweeps or price sweeps) were prepared with RS/1 software (Brooks Automation, Chelmsford, MA) and examined visually. If the curves were well clustered, then the data were considered to be stable, even if there were one or two outlying data points. Once the data appeared to be stable the rat was moved on to the next phase of the experiment.

Each phase of the experiment consisted, successively, of a low-price frequency sweep, a price sweep at a high frequency, and a high-price frequency sweep. The frequency chosen for the price sweep was the highest one the rat could tolerate in the low-price frequency-sweep tests. Once the price sweeps were completed, the high-price

frequency sweeps were conducted on the basis of a price that corresponded to the falling portion of the price-sweep function (Figure 16).

Fitting the mountain surface.

The goal of this experiment was to determine if it was possible to shift the mountain along one of its axes but not the other. The position of the mountain in each condition was determined by fitting the 3D psychometric function to the data through the application of a weighted, non-linear, least-square minimization routine (Statistica, StatSoft, Inc.). The following sections describe the stages of this procedure.

Weighting procedure

The conventional least-squares method for fitting a model to data is based on the assumptions that the data are normally distributed and that their variance is homogeneous. These assumptions are not tenable in the case of the current data set. Two factors, skew and the presence of outliers, cause the distribution of time-allocation values to deviate from normality. The skew arises from the fact that time-allocation values cannot exceed one and cannot be less than zero; the source of the outliers is not known, but it is likely that conditions during the long test sessions are not completely stationary. The S-shape of the 2D psychometric curves and their tendency to shift slightly along the price or frequency axes cause the variance of points that fall along the steep portion to be greater than the variance of points that fall near the upper or lower asymptotes. To address the issues of skew, the outliers, and heteroskedasticity, a weighting procedure based on Tukey's bisquare estimator was implemented. This procedure is described in the appendix.

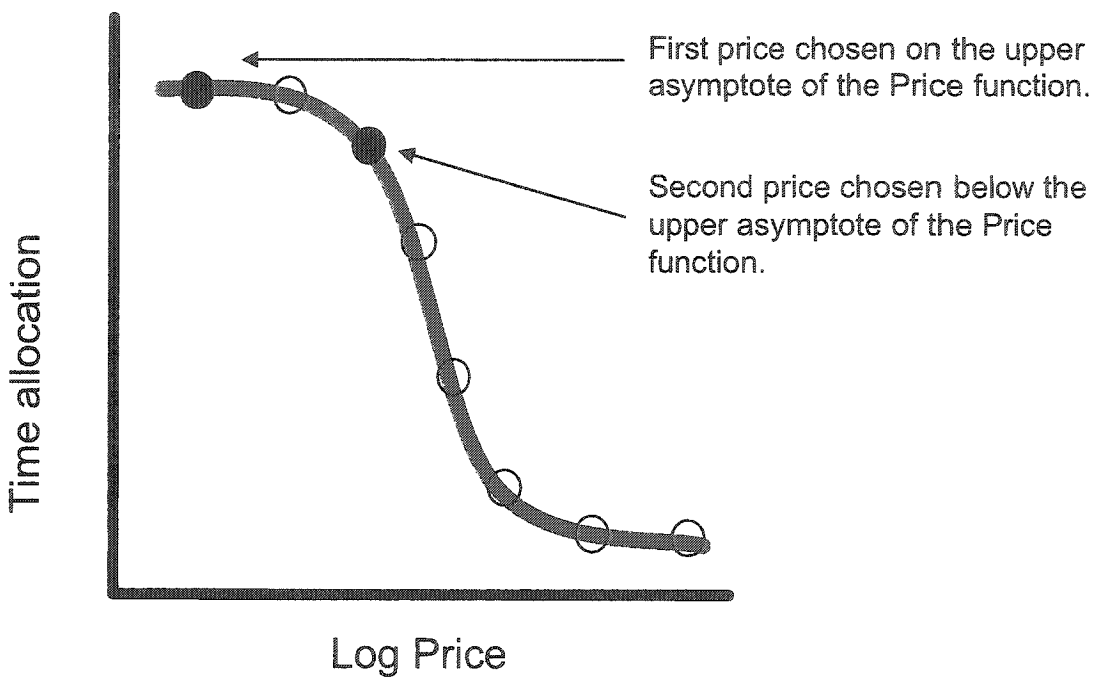

Figure 16. A graphical representation of how the price was chosen for the high-price frequency sweeps.

The bi-square weights, raw means, and standard errors of the mean were computed using the RS/1 software package (Brooks Automation, Chelmsford, MA). These values were then supplied to the Statistica package (StatSoft, Inc.), which carried out the non-linear surface fitting.

Results

An A-B-A experimental design was implemented to assess the effect of decreasing the duration of BSR trains. The baseline conditions were executed with a train duration of 1 second. The first baseline is the TD1A condition, and the second is TD1B, where TD is train duration, 1 is 1 second, A is the first determination of the baseline, and B the second determination. The experimental condition is TDP25, where P25 stands for a train duration of 0.25 seconds. Each condition comprises three successive sets of sweeps: a set of frequency sweeps conducted at a fixed low price, a set of price sweeps at a fixed high frequency, and a set of frequency sweeps conducted at a fixed high price. The presentation of the baseline data proceeds, step by step, from the raw data for a single subject and condition to a 3D representation that combines the data from the set of price sweeps with the data from both sets of frequency sweeps. The analysis assesses test-retest reliability first through qualitative comparison of 2D graphs and then by quantitative comparison of fitted 3D surfaces. Finally, the analysis examines the effects of decreasing the train duration to 0.25 seconds -- initially for one subject and then for the remaining four subjects.

A two-dimensional representation of a single baseline data set

Figure 17 presents all of the baseline data for one rat, subject M16. Examination of these results reveals the forms of the psychometric functions, the consistency of the time-allocation values obtained when the same price and frequency values were tested in the context of price or frequency sweeps, and the effect on the position of the frequency-sweep curve of changing the price.

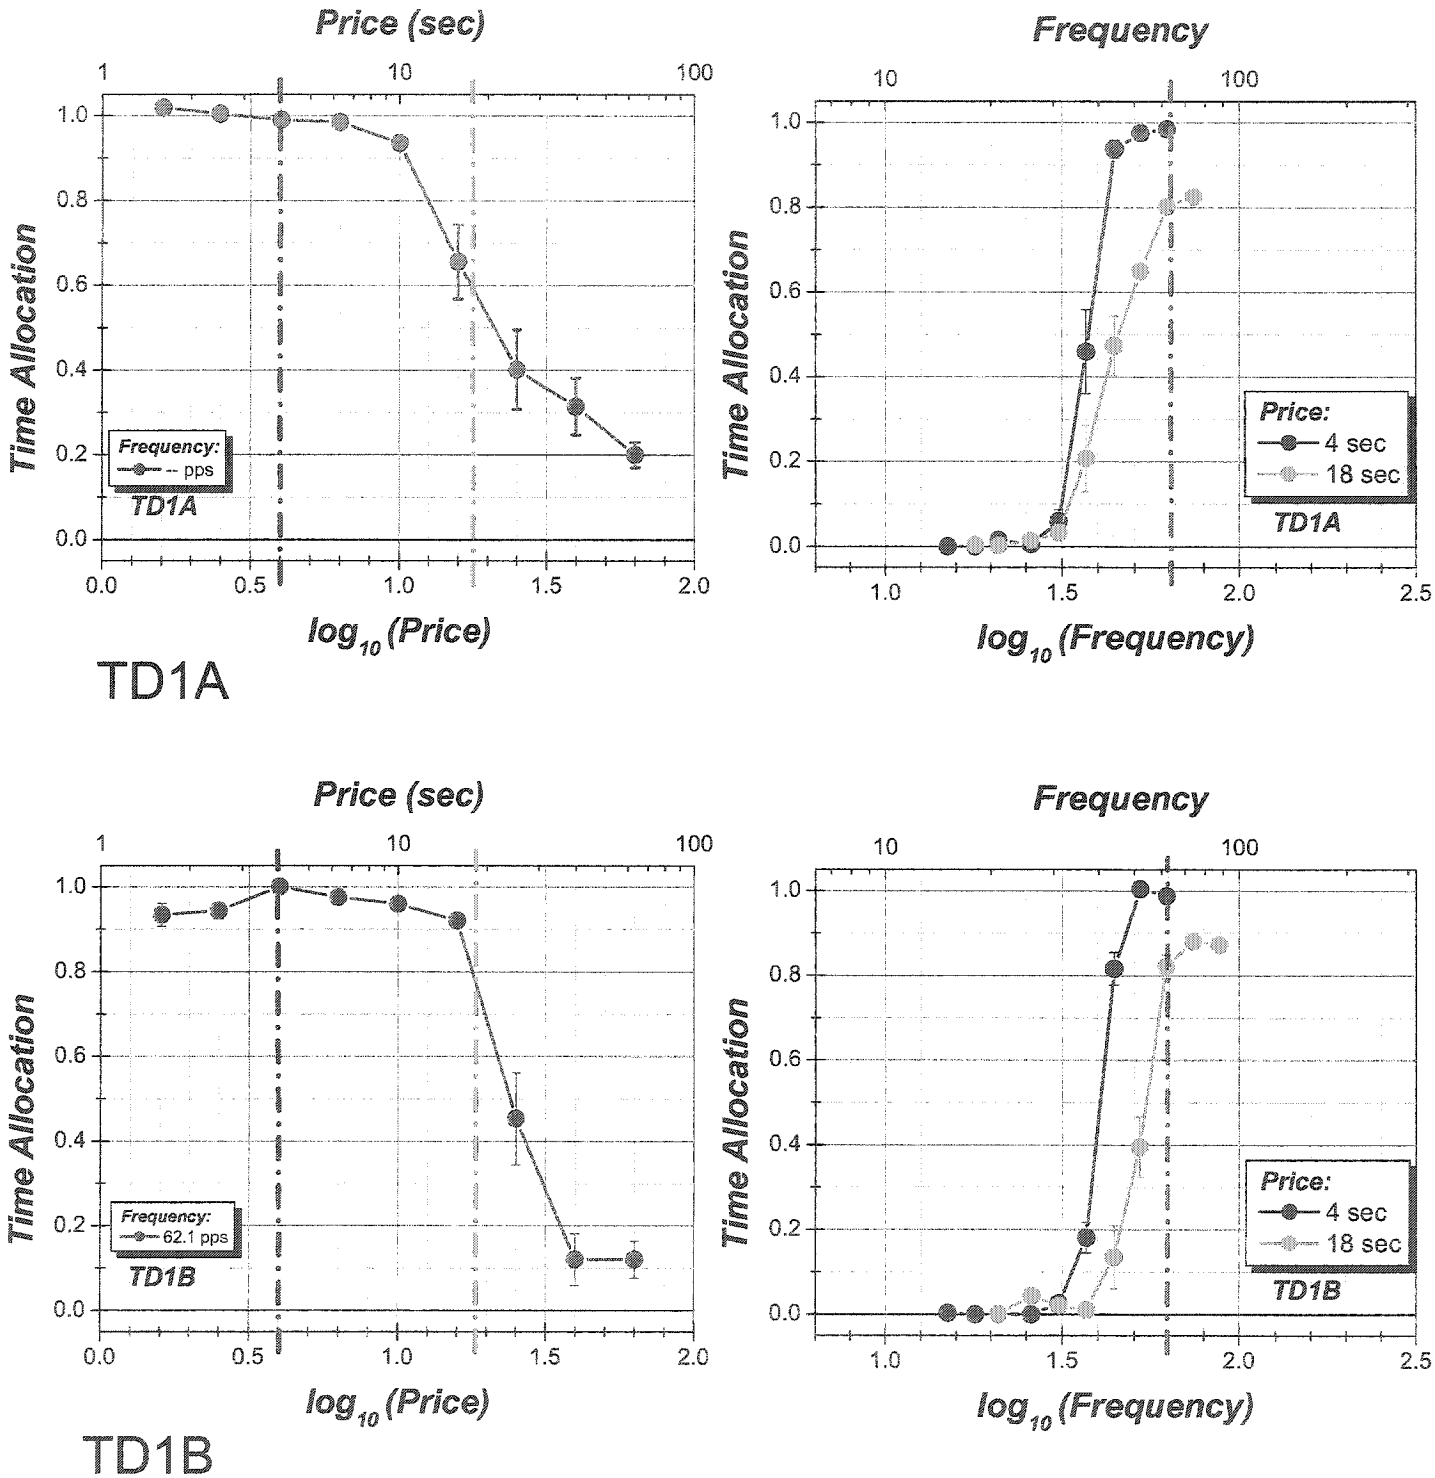

Figure 17. Price effect for subject M16. A higher price results in a rightward shift of the data along the frequency axis.

The top two panels in Figure 17 show the complete data set for the TD1A condition, and the bottom two show the data for the TD1B condition. Data from the TD1B condition were collected after the shorter train duration had been tested. All of the figures from this point on show corrected time-allocation values (see "Methods", above). Both of the left-hand panels in Figure 17 display price-sweep data. The dashed red and green lines on these graphs represent the prices at which the two frequency sweeps were run. The low-price frequency sweep was conducted with a 4-second price (red line), and the price of the stimulation during the high-price frequency sweep (green line) was 18 seconds. The two right-hand panels present the two frequency sweeps for the baseline conditions, TD1A and TD1B. The dashed blue line on both right-hand panels represents the frequency employed to obtain the price-sweep data shown in the left-hand panel (62.1 pulses per second).

As the two lower panels of Figure 17 show, the price-sweep and frequency-sweep data from the TD1B condition are consistent: similar time-allocation values were obtained when the same pair of prices and frequencies were tested during both the price sweep and the frequency sweep. Note that the intersection of the green dashed line and the blue solid line in the lower-left panel occurs at a time-allocation value similar to that of the intersection of the blue dashed line and the green solid line in the lower-right panel; at both intersections, the price is 18 seconds, and the frequency is 62.1 pulses per second. A small discrepancy emerges, however, in the high-price frequency-sweep data for the TD1A condition (the upper panels in Figure 17). The intersection of the blue dashed line with the green curve in the upper-right panel occurs at a higher value of time allocation than does the intersection of the green dashed line with the blue curve in the

upper-left panel. However, the standard errors of the points on the price-sweep curve (blue) are large in the vicinity of the intersection with the dashed green line.

The two right-hand panels in Figure 17 show the effect of manipulating the price of BSR. An increase in the price of a stimulation train from 4 seconds (red data points) to 18 seconds (green data points) shifted the psychometric functions to the right and decreased their slopes. The maximum time allocation was lower in the high-price condition (green) than it was in the low-price condition (red).

Building the mountain

In order to integrate the data from all three sets of sweeps into a single structure, the data were replotted in three-dimensional space. The upper panels of Figure 18 show the data from the first (TD1A) and second (TD1B) determinations of the baseline for rat M16. Each of these graphs positions the three sets of data points along the three axes of the mountain: time allocation, price, and frequency. The bottom two panels of Figure 18 present the same data along with a wire mesh surface that represents the best-fitting parameters of the mountain equation (Equation 8). A later section will describe the surface-fitting results in detail.

Reliability

The analysis compared the results from the two baseline conditions for each rat in order to assess the reliability of the measurements and to determine whether it was possible in terms of statistical significance to pool these results before comparing them with the results from the short train-duration condition. For four of the five subjects, the results of the return to the baseline condition were not large; they met the criterion for

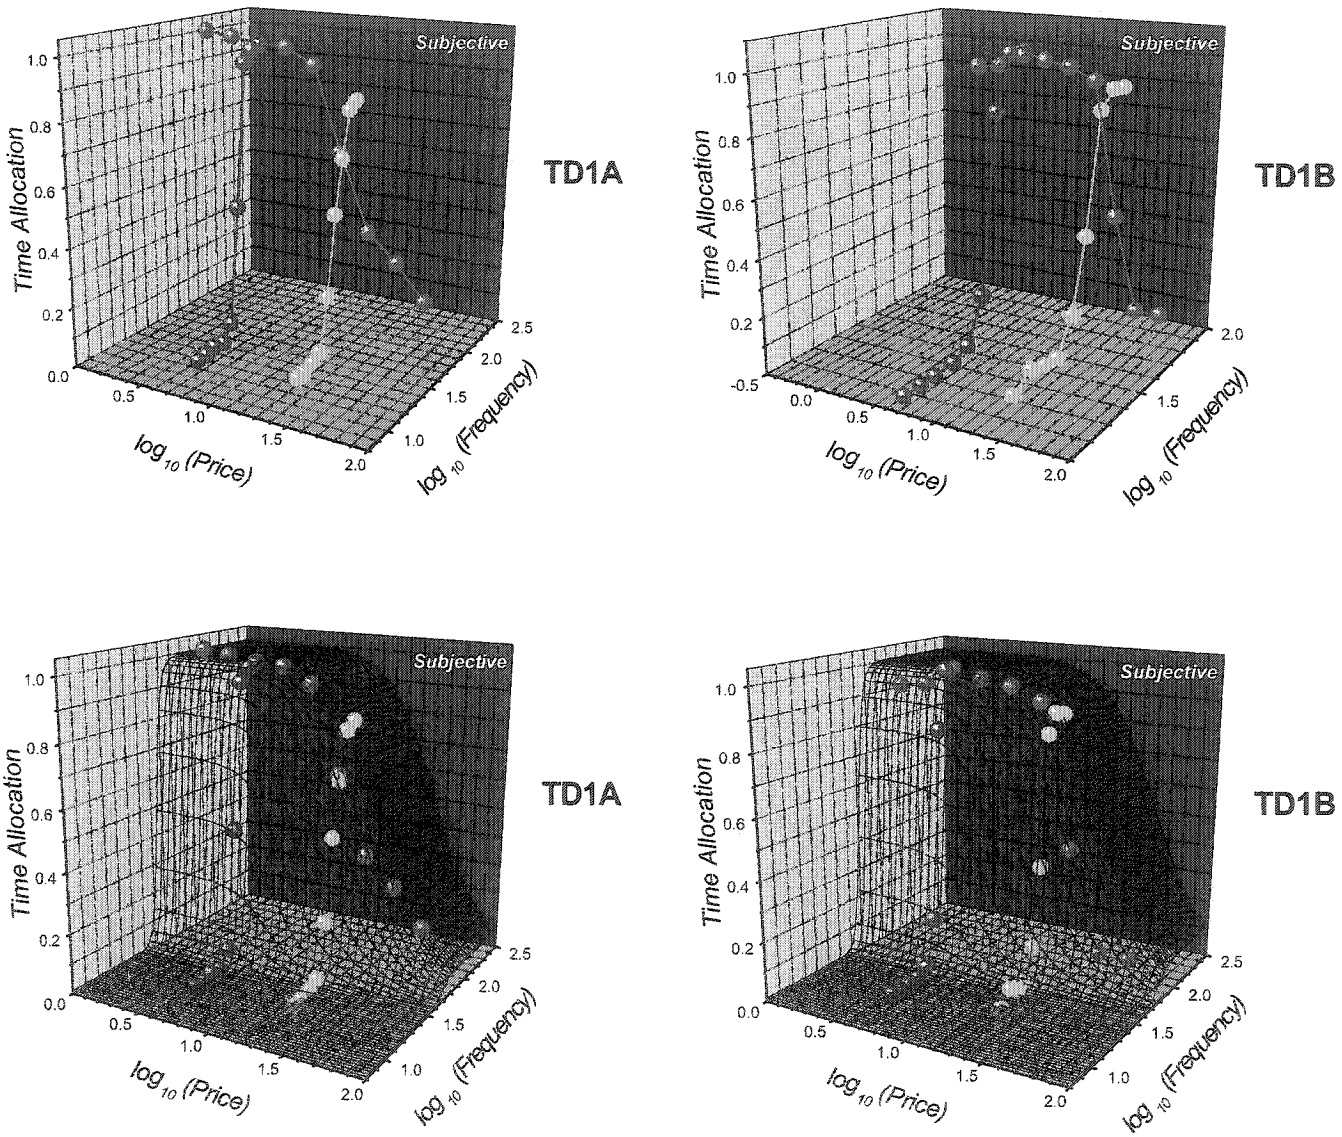

Figure 18. The integration of the data from all 3 sets of sweeps into a single structure, in 3D space, are seen in the upper panels for subject M16. The left most graphs represent the TD1A condition whereas the rightmost the TD1B condition. The bottom panels present the same data with a fitted surface.

statistical significance, however, and the data from the baseline conditions were not combined for these subjects.

Figure 19 presents the data for the two baseline conditions for subject M14. The top two panels in the figure represent the frequency comparisons, and the bottom panel shows the price-sweep data. In the case of the low-price frequency-sweep comparison (upper-left panel), apart from a small turn-down at the highest frequency, the two baseline data sets line up very well. There are very few differences between the data sets. (All three panels present the TD1A condition as the darker color and the TD1B condition as the lighter color.)

Figures 20 to 23 represent the baseline comparisons for the remaining four subjects. The test-retest reliability for subject M11 is very good for two of the three panels in Figure 20. The only substantial discrepancy between baselines occurs in the upper-left panel. A visual comparison suggests that the shift is approximately 0.05 log units.

The red points on the price function graphs for subjects M11 and M16 (the lower panels in both Figures 20 and 21, respectively) appeared to reflect a response bias and were, accordingly, omitted from the surface-fitting analyses. This "turn-down" is suspected to be an artifact of the very high level of behavioural activation that is seen when the price is very low and the frequency is quite high. When the price was 1 second the subject was only given 20 seconds in which to reap all 20 of the available rewards. It is possible that the powerful inexpensive reward produced a strong tendency to abandon the lever in order to explore, thereby reducing even further the relatively short time available in which to harvest rewards. At higher prices, brief exploratory excursions

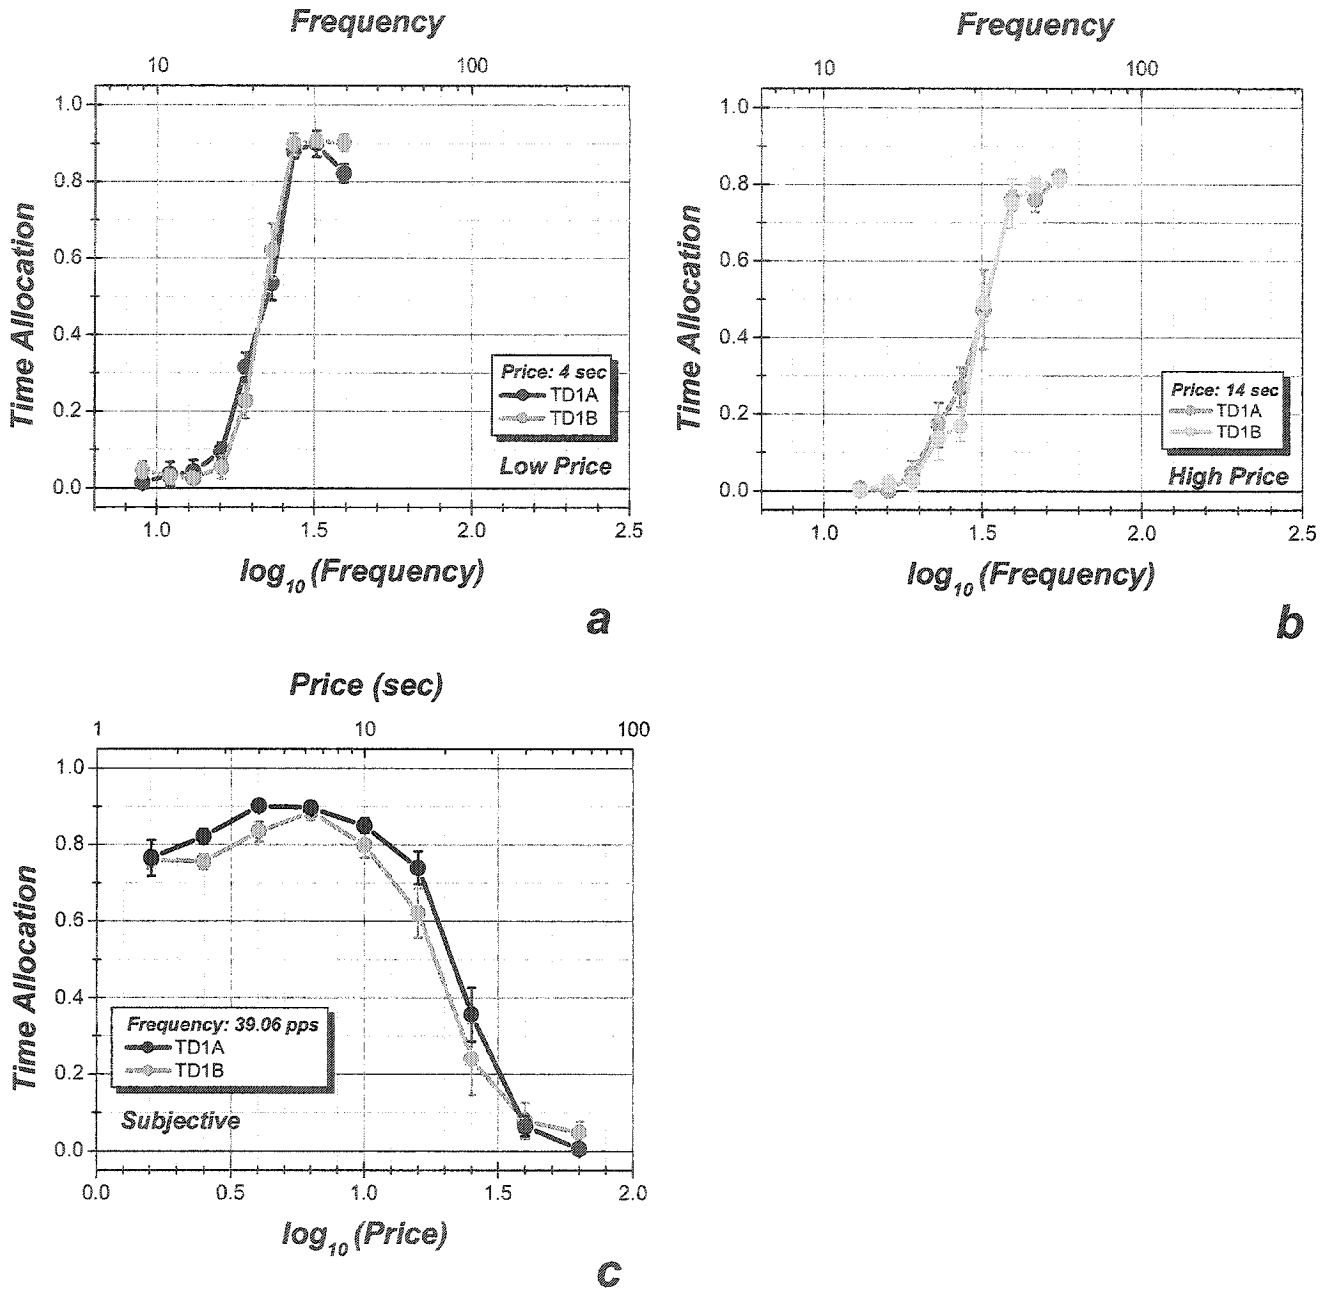

Figure 19. Here the two baseline conditions are compared for subject M14.
The two baseline conditions for M14 were completely unchanged. The graph on the left shows the low price frequency sweep results for the TD1B condition (light red) graphed together with the TD1A condition (dark red). The graph on the right is the high price frequency sweep condition where the data from TD1A is presented as dark green and TD1B is light green.

The lower left panel presents the price sweep data for both baselines. Again, the first baseline, TD1A, is the darker color.

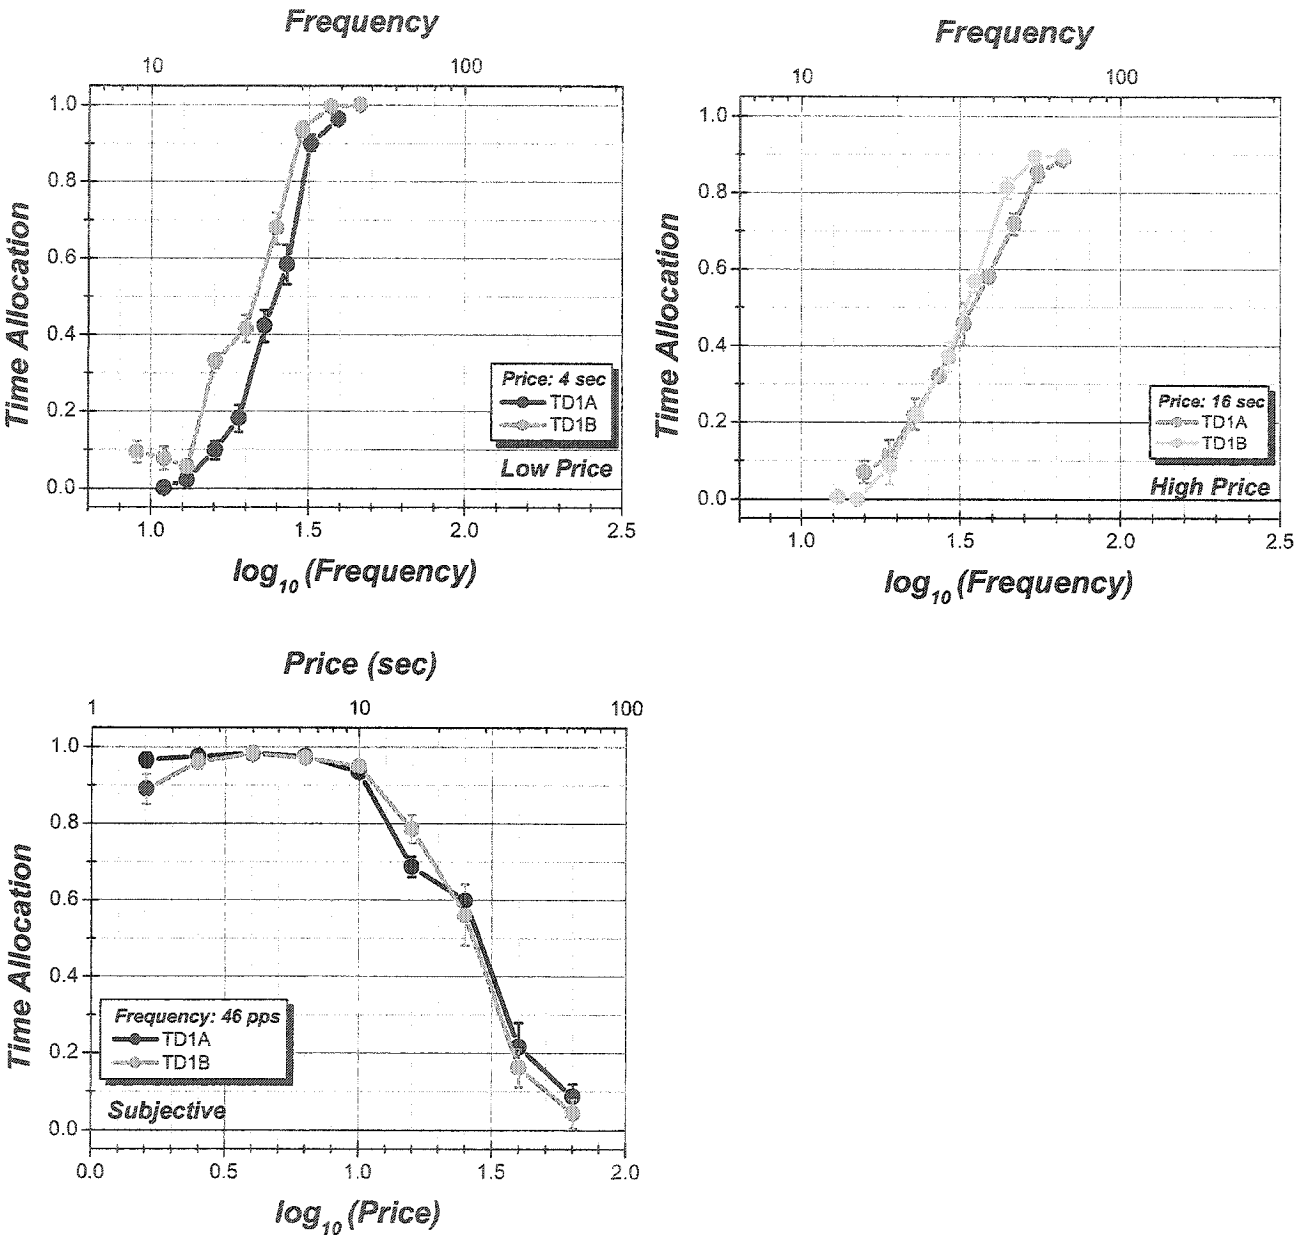

Figure 20. Baseline comparisons for subject M11.
The graph on the upper left shows the low price frequency sweep results for the TD1B condition (light red) graphed together with the TD1A condition (dark red). The graph on the upper right is the high price frequency sweep condition where the data from TD1A is presented as dark green and TD1B is light green. The graph to the left is the Price sweep for both baselines. Again, the first baseline, TD1A, is the darker color.

would have less impact on the time-allocation ratio. Further analyses of the distribution of leisure times might shed additional light on the "turn-down" phenomenon. Including points subject to the putative response bias would have inflated the error variance by adding a source of systematic variance.

The baseline data for subject M16 is presented in Figure 21. The low-price frequency-sweep results are statistically, very reliable (upper-left panel). There is some variation, however, in the curves yielded by the high-price frequency sweeps (upper-right panel) and in the slopes of the price-sweep curves (lower-left panel). Figures 22 and 23, respectively, present the most variable baseline data, those from subjects M20 and M22. In the case of subject M20, the high-price frequency-sweep shifted to the right for the second baseline condition (upper-right panel of Figure 22), whereas the low-price frequency-sweep (upper-left panel) was stable. The largest discrepancy between baselines for rat M22 occurs in the price-sweep data (lower panel of Figure 23). Overall, the discrepancies tend to be rather small.

The analysis to this point has described qualitative differences between the data sets. The following section describes the surface-fitting approach to a quantitative comparison of the two sets of baseline data.

Quantitative results of the baseline comparisons.

The analysis used a non-linear surface-fitting procedure to determine whether there was a statistically reliable difference between the two baseline datasets for each subject. This procedure would determine whether it was possible to pool the baseline data in order to establish a combined baseline that could then be compared with the short-train duration data. It would be possible to pool the data if there were no statistically

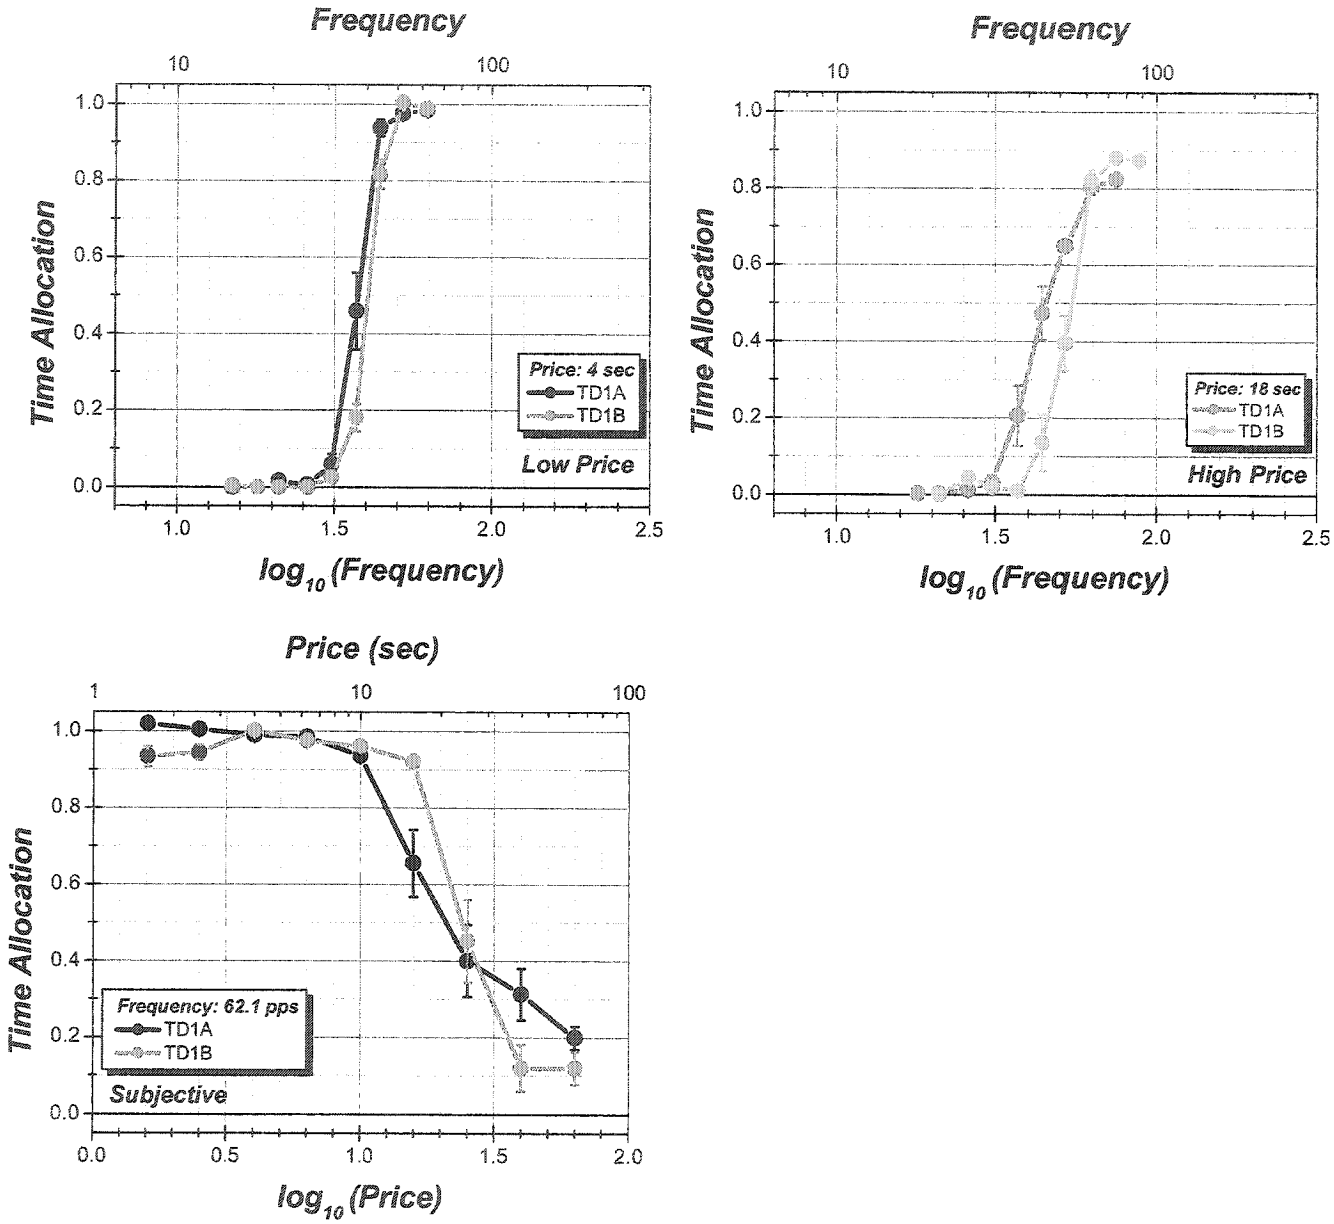

Figure 21. Baseline comparisons for subject M16.
The graph on the upper left shows the low price frequency sweep results for the TD1B condition (light red) graphed together with the TD1A condition (dark red). The upper-right panel is the high-price frequency-sweep condition, where the data from TD1A is presented as dark green and TD1B is light green. The graph to the left is the Price sweep for both baselines. Again, the first baseline, TD1A, is the darker color.

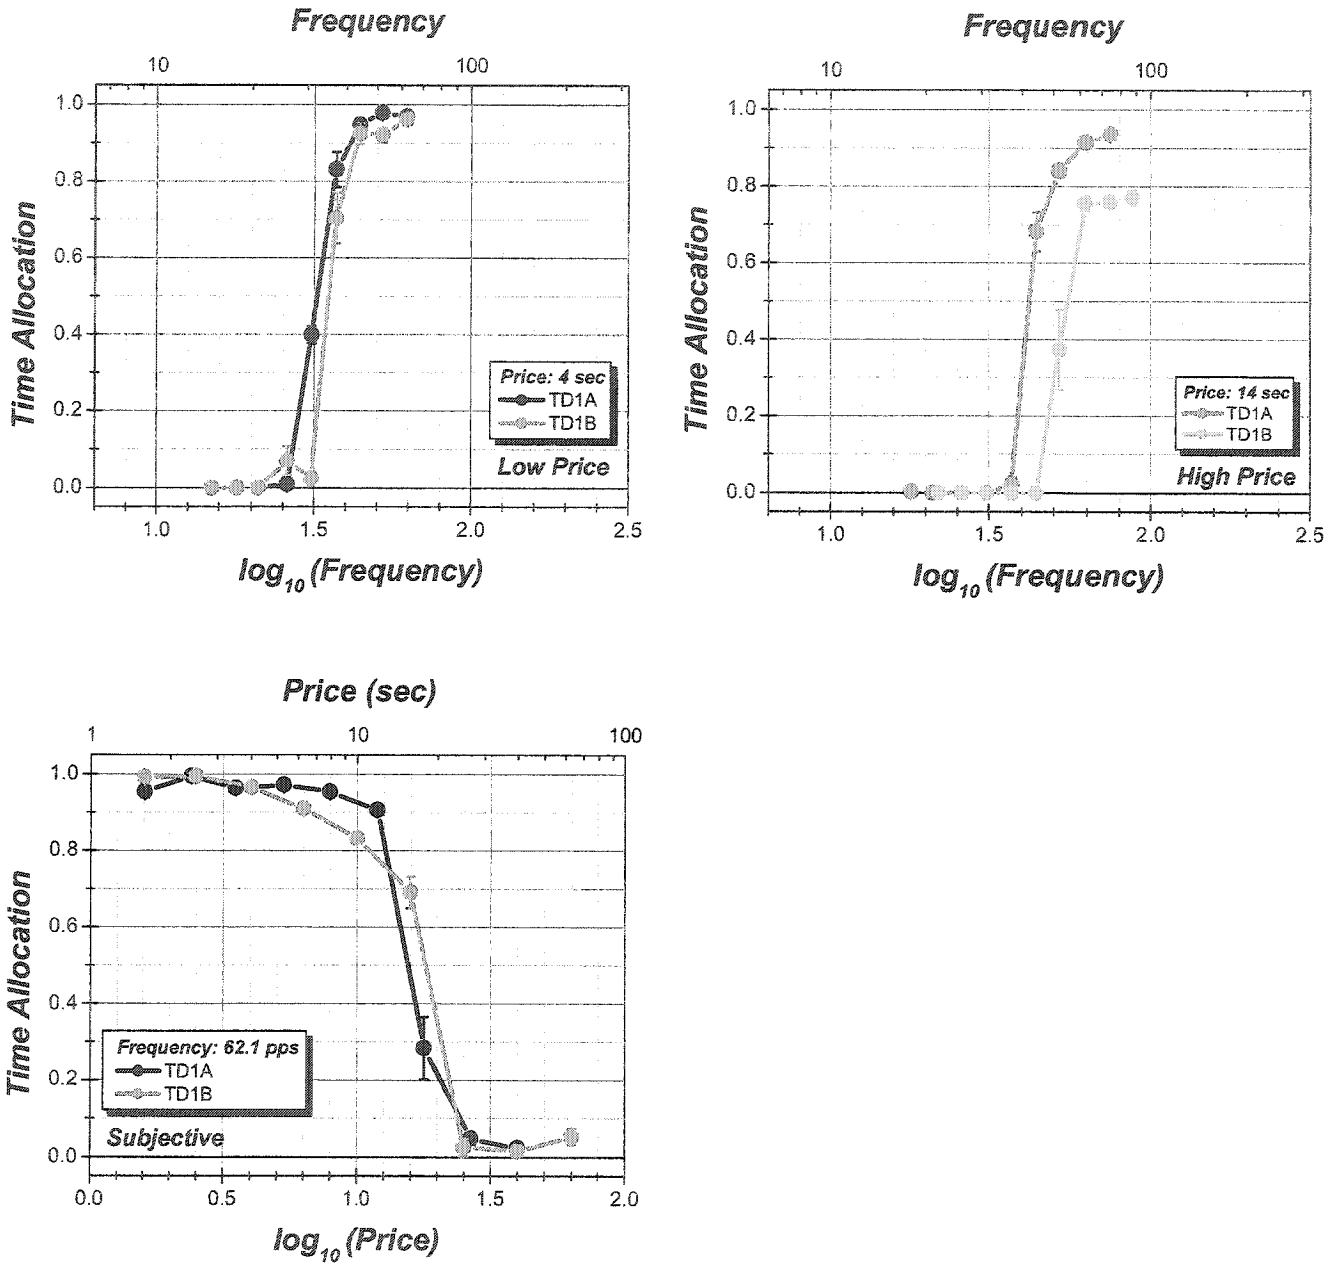

Figure 22. Baseline comparisons for subject M20.
The graph on the upper left shows the low price frequency sweep results for the TD1B condition (light red) graphed together with the TD1A condition (dark red). The graph on the upper right is the high price frequency sweep condition where the data from TD1A is presented as dark green and TD1B is light green. The graph to the left is the Price sweep for both baselines. Again, the first baseline, TD1A, is the darker color.

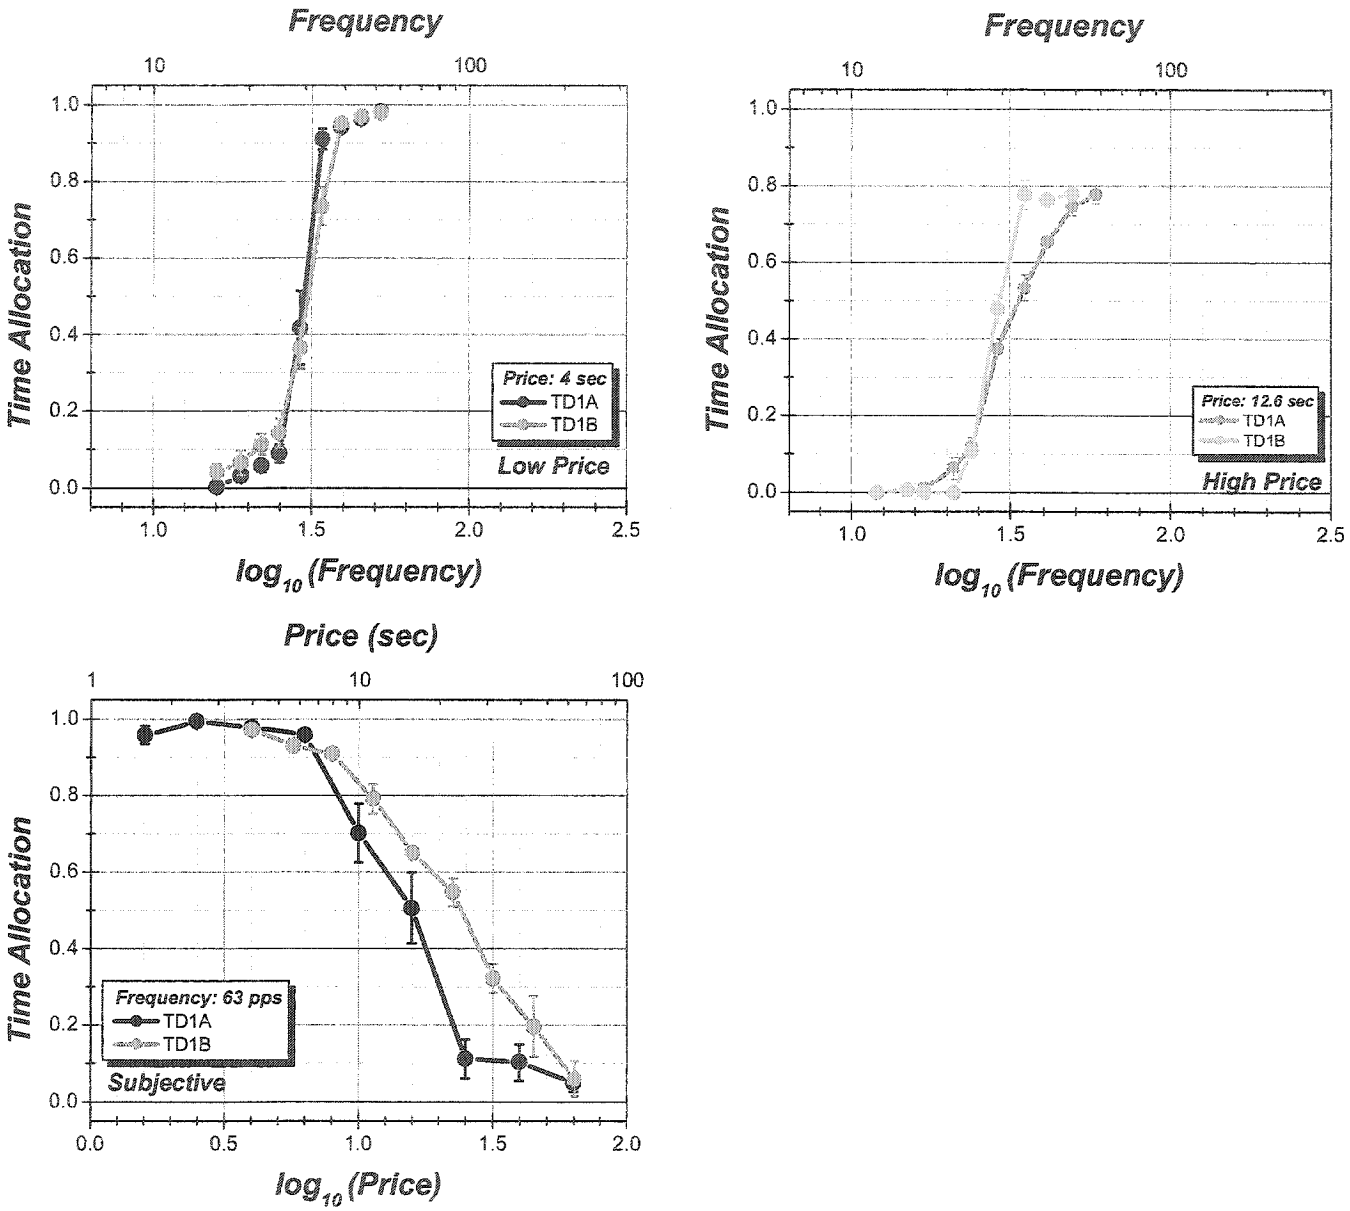

Figure 23. Baseline comparisons for subject M22
The graph on the upper left shows the low price frequency sweep results for the TD1B condition (light red) graphed together with the TD1A condition (dark red). The graph on the upper right is the high price frequency sweep condition where the data from TD1A is presented as dark green and TD1B is light green. The graph to the left is the Price sweep for both baselines. Again, the first baseline, TD1A, is the darker color.

significant difference between the values for each of the location parameters, f_{hm} and P_E, for the baseline mountain TD1A and the values of the parameters for the baseline mountain TD1B. To fit separate surfaces to each baseline data set would have entailed the simultaneous estimation of a large number of parameters (as many as 16 or 8 for each baseline mountain). Moreover, only the values of the location parameters were germane to testing the hypothesized effect of varying the train duration. Thus, common values were estimated for all parameters except the two that determine the position of mountain, f_{hm} and P_E. The shift in the position of the mountain from the TD1A condition to the TD1B condition was represented by Δf_{hm}, which gave the displacement along the frequency axis, and ΔP_e, which gave the displacement along the price axis. When the 95 percent confidence interval around either of these shift parameters did not include zero, the mountain was deemed to have shifted.

The baseline comparisons will be presented initially for one case, that of subject M14. Figure 24 presents a bar graph that represents the size of the shift along each axis and the corresponding 95 percent confidence intervals for the shift parameters. Note that the confidence intervals for each parameter includes zero; consequently, neither shift is considered to be statistically reliable.

The amount of variance accounted for by the fitted surface was calculated for all subjects by dividing the sum of squared deviations that were attributable to regression by the total sum of squares as presented in the corresponding analysis of variance (ANOVA) in Table 1. The r^2 for subject M14 is 0.989. Table 2, which presents the parameter estimates for subject M14, shows that the confidence interval in each of the parameters ΔP_e and Δf_{hm} includes zero. This outcome is consistent with the position of the

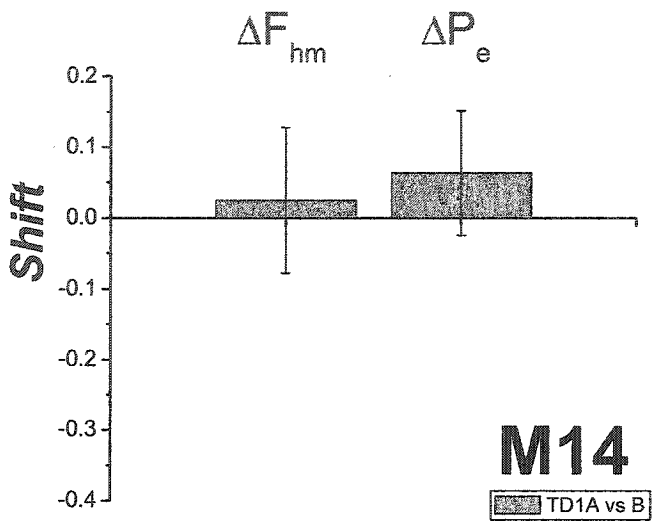

Figure 24. Baseline comparisons for subject M14. A bar graph that represents the size of the shift along each axis and the corresponding 95% confidence intervals for the shift parameters is presented. Note that the confidence intervals include zero and as such neither shift is considered to be statistically reliable. ΔF_{hm}, presents the displacement along the frequency axis, and ΔP_e, presents the displacement along the price axis.

	Sum of Squares	DF	Mean Squares	F-value	p-value
Regression	97.531	10	9.753	2378.886	0.000
Residual	1.066	260	0.004		
Total	98.597	270			
r^2	0.989				

Table 1. ANOVA table for subject M14

M14 DOUBLE MOUNTAIN (Baselines):	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPe: Delta Price [shift along the price axis for TD1A from TD1B]	0.063	0.045	1.418	0.157	-0.025	0.151
ΔFhm: Delta Frequency [shift along the frequency axis for TD1A from TD1B]	0.025	0.027	0.906	0.366	-0.029	0.078
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.021	0.007	2.954	0.003	0.007	0.035
T_{MAX}: Maximum time Allocation	0.864	0.008	106.219	0.000	0.848	0.880
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	5.514	0.426	12.940	0.000	4.675	6.353
P_{MIN}: minimum subjective price	0.870	0.025	35.448	0.000	0.821	0.918
P_{E2}: parameter determining location along the price axis [for TD1B]	0.009	163.927	0.000	1.000	-322.786	322.803
g: growth exponent [determines the slope of mountain on Price façade]	1.547	0.066	23.409	0.000	1.416	1.677
F_{HM2}: parameter determining location along the frequency axis [for TD1B]	1.544	0.052	29.799	0.000	1.442	1.646

Table 2. Parameter estimates for subject M14. the confidence intervals for both of the shift parameters include zero (highlighted yellow). There is no statistically significant difference between TD1A and TD1B.

confidence intervals in the bar graph in Figure 24. M14 was the only subject whose baseline conditions did not shift along either the price axis or the frequency axis.

The scatterplot in Figure 25 is a graphical comparison of the observed time allocation with the time allocation predicted by the fitted surface. The diagonal red line represents a perfect fit between the observed values and the predicted values. The actual data points line up nicely around the line of best fit. If the function did not fit the data well, the data points would follow a curved path rather than a straight one.

One subject whose baseline mountain did shift is M16. Thus the graph for subject M16 in Figure 26 differs from the graph for subject M14 in Figure 24 because in subject M16's case the confidence intervals surrounding the shifts along each axis do not include zero. Table 3 presents the parameter values for this subject, and Figure 27 presents the corresponding scatterplot. The r^2 for subject M16 is 0.997.

Table 4 presents the results of the baseline mountain comparison for subject M11. As in the case of subject M16, neither the set of confidence intervals for Δf_{hm} nor the set for ΔP_e includes zero. Consequently, the shifts along both axes are statistically significant and the baseline mountains for this subject cannot be combined. Figure 28 plots the observed and predicted time-allocation values for subject M11, illustrating the excellent fit of the surfaces. The corresponding r^2 value for this subject is 0.990.

Table 5 shows the results of the comparison for subject M20. The shift along the frequency axis (Δf_{hm}) is significant, but the shift along the price axis (ΔP_e) is not. Figure 29 shows the scatterplot for M20's observed and predicted time-allocation values. The scatterplot's apparent sparseness is a result of the fact that many of the 250 data points fell on top of others. The r^2 for M20 is 0.994.

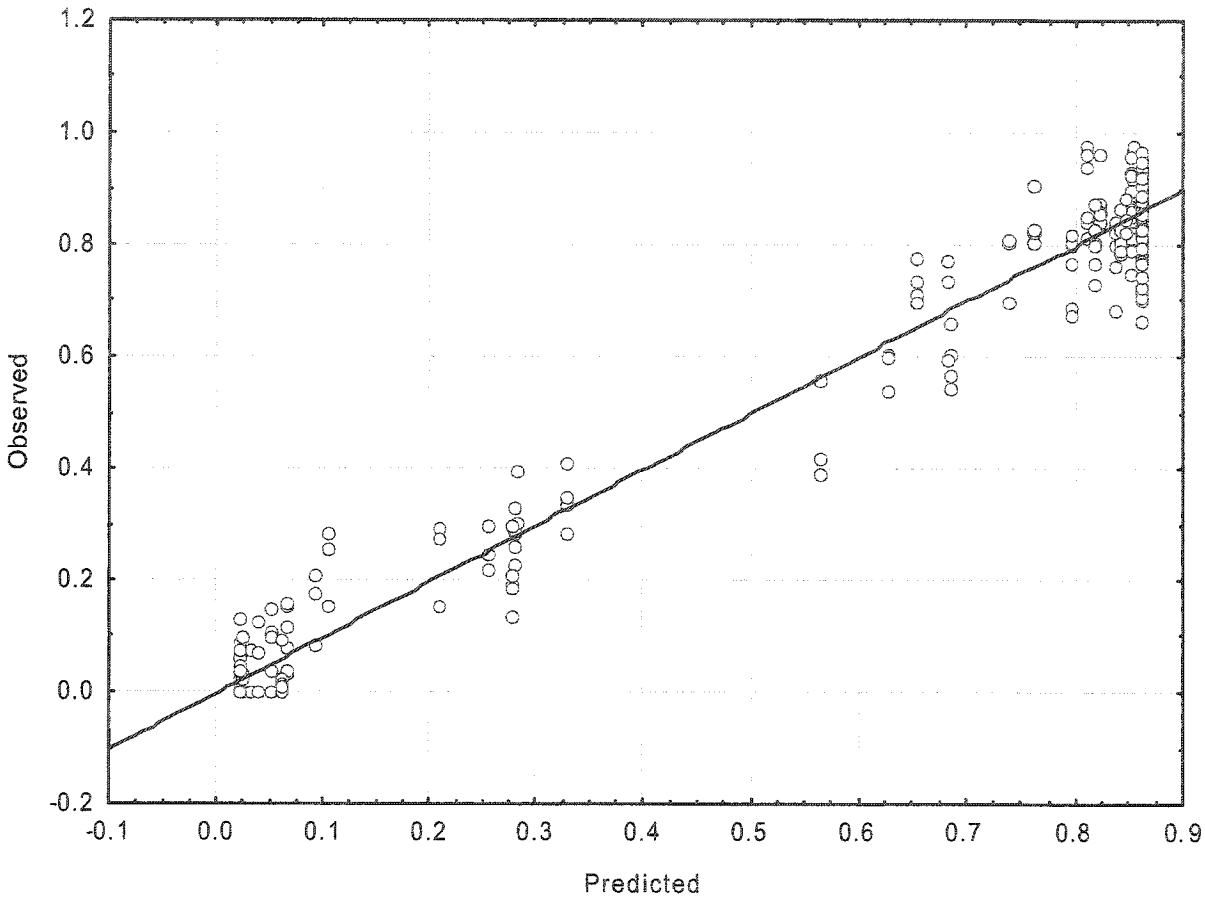

Figure 25. A graphical comparison of the observed time allocation and the time allocation predicted by the fitted surface for subject M14. The line for a perfect fit between the observed and predicted values is drawn in red diagonally through the graph. The actual data points line up nicely around the line of best fit. If the function did not fit the data well the data points would follow a curvilinear rather than a linear path in this space.

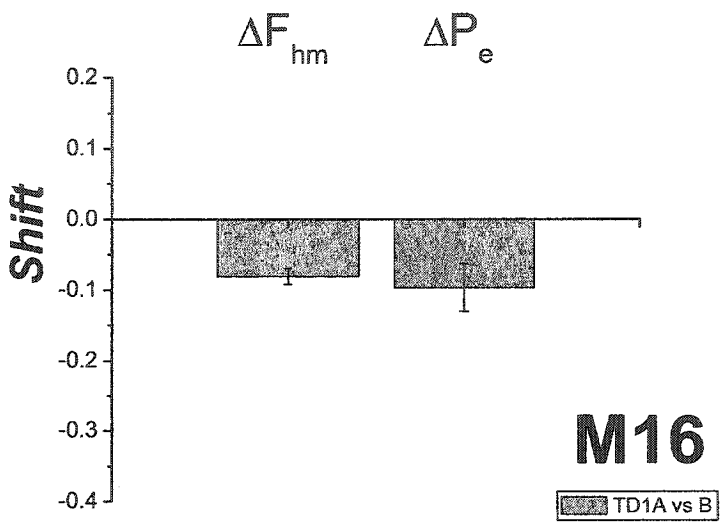

Figure 26. A bar graph that represents the size of the shift along each axis and the corresponding 95% confidence intervals for the shift parameters for subject M16. Note that the confidence intervals do not include zero for subject M16 and as such both shifts are considered to be statistically reliable. Δf_{hm}, presents the displacement along the frequency axis, and ΔP_e, presents the displacement along the price axis.

M16 DOUBLE MOUNTAIN (Baselines):	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPe: Delta Price [shift along the price axis for TD1A from TD1B]	-0.097	0.017	-5.607	0.000	-0.131	-0.063
ΔFhm: Delta Frequency [shift along the frequency axis for TD1A from TD1B]	-0.081	0.006	-13.994	0.000	-0.092	-0.069
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.010	0.004	2.637	0.009	0.003	0.017
T_{MAX}: Maximum time Allocation	0.997	0.005	186.223	0.000	0.986	1.007
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	2.679	0.102	26.284	0.000	2.478	2.880
P_{MIN}: minimum subjective price	0.486	0.138	3.508	0.001	0.213	0.758
P_{E2}: parameter determining location along the price axis [for TD1B]	1.610	0.019	86.011	0.000	1.573	1.646
g: growth exponent [determines the slope of mountain on Price façade]	7.135	0.335	21.267	0.000	6.474	7.796
F_{HM2}: parameter determining location along the frequency axis [for TD1B]	1.735	0.006	273.104	0.000	1.722	1.747

Table 3. Parameter estimates for subject M16. the confidence intervals for both of the shift parameters do not include zero (highlighted yellow).

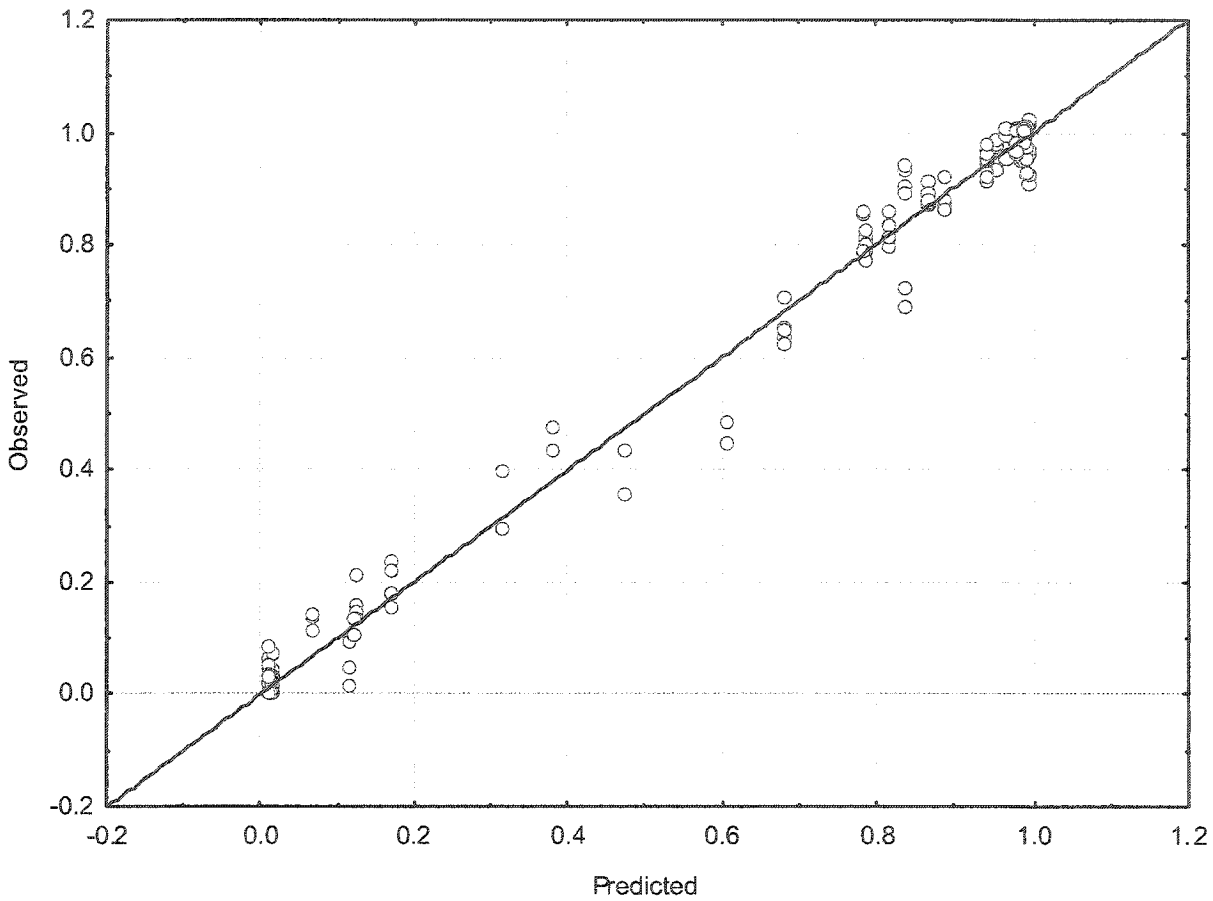

Figure 27. A graphical comparison of the observed time allocation and the time allocation predicted by the fitted surface for subject M16.

M11 DOUBLE MOUNTAIN (Baselines):	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPe: Delta Price [shift along the price axis for TD1A from TD1B]	0.088	0.043	2.028	0.044	0.003	0.173
ΔFhm: Delta Frequency [shift along the frequency axis for TD1A from TD1B]	0.104	0.027	3.932	0.000	0.052	0.157
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.019	0.013	1.453	0.148	-0.007	0.045
T_{MAX}: Maximum time Allocation	1.052	0.020	53.836	0.000	1.013	1.090
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	2.477	0.162	15.310	0.000	2.159	2.796
P_{MIN}: minimum subjective price	0.909	0.021	43.755	0.000	0.868	0.950
P_{E2}: parameter determining location along the price axis [for TD1B]	1.609	0.047	33.920	0.000	1.515	1.702
g: growth exponent [determines the slope of mountain on Price façade]	2.460	0.165	14.901	0.000	2.135	2.785
F_{HM2}: parameter determining location along the frequency axis [for TD1B]	1.583	0.037	43.216	0.000	1.511	1.656

Table 4. Shift values for subject M11. ΔP_e is the delta Price, the shift of the mountain along the price axis. Whereas Δf_{hm} is the shift parameter along the frequency axis.

In the case of M11 zero is not included between the lower confidence interval or the upper confidence interval for either of the shift parameters. So, we are 95% confident that there is a difference between baseline conditions. So the baseline conditions for this subject may not be combined.

M20 DOUBLE MOUNTAIN (Baselines):	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPe: Delta Price [shift along the price axis for TD1A from TD1B]	-0.006	0.016	-0.385	0.700	-0.040	0.026
ΔFhm: Delta Frequency [shift along the frequency axis for TD1A from TD1B]	-0.066	0.010	-6.846	0.000	-0.090	-0.047
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.017	0.005	3.667	0.000	0.010	0.027
T_{MAX}: Maximum time Allocation	0.941	0.006	166.049	0.000	0.930	0.952
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	14.320	3.237	4.424	0.000	7.940	20.697
P_{MIN}: minimum subjective price	0.870	0.070	12.460	0.000	0.730	1.007
TP: Transition Parameter	0.009	590.213	0.000	1.000	-1162.650	1162.669
P_{E2}: parameter determining location along the price axis [for TD1B]	1.288	0.046	27.844	0.000	1.200	1.379
g: growth exponent [determines the slope of mountain on Price façade]	4.373	1.366	3.201	0.002	1.680	7.065
F_{HM2}: parameter determining location along the frequency axis [for TD1B]	1.607	0.029	56.453	0.000	1.550	1.663

Table 5. Shift values for subject M20.

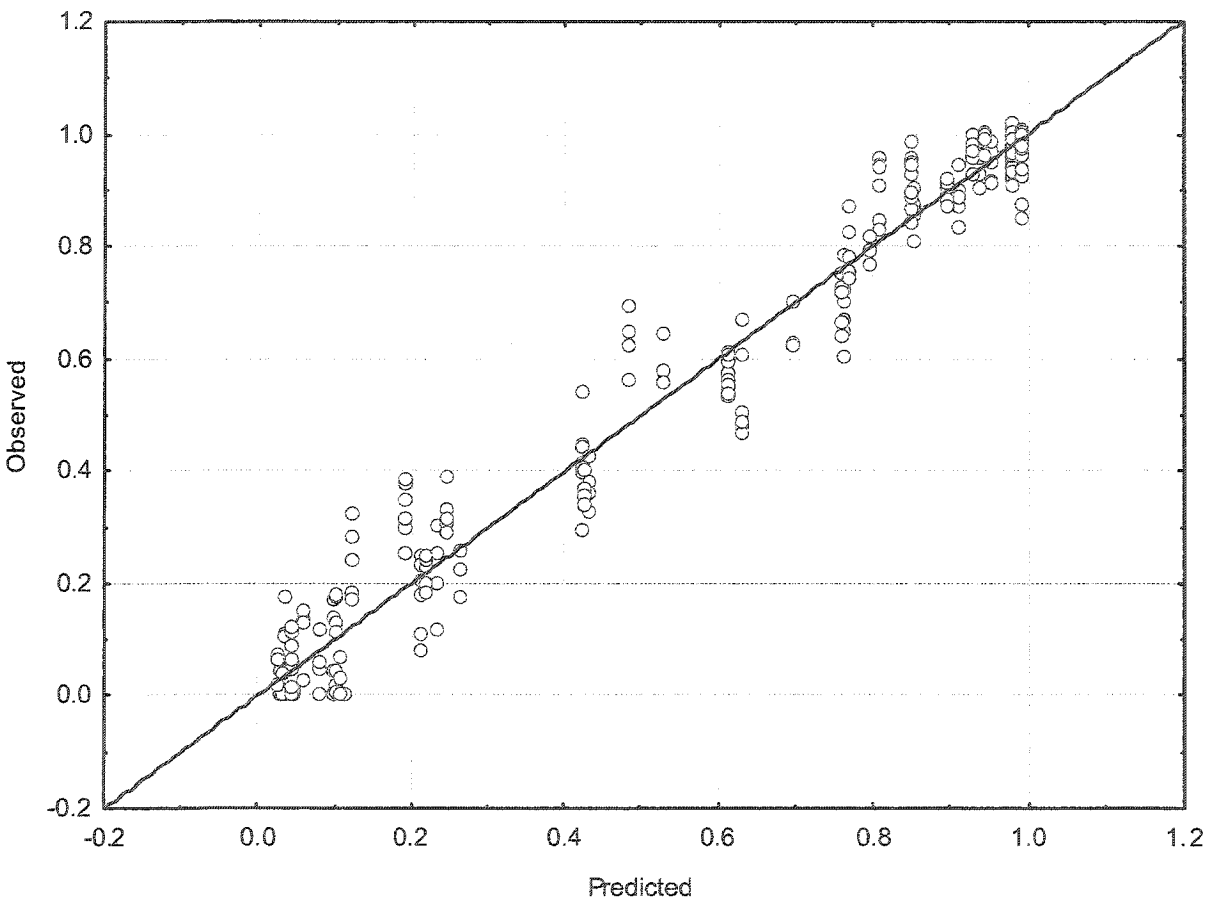

Figure 28. Scatterplot comparing the predicted time-allocation values for subject M11 to the observed time-allocation values.

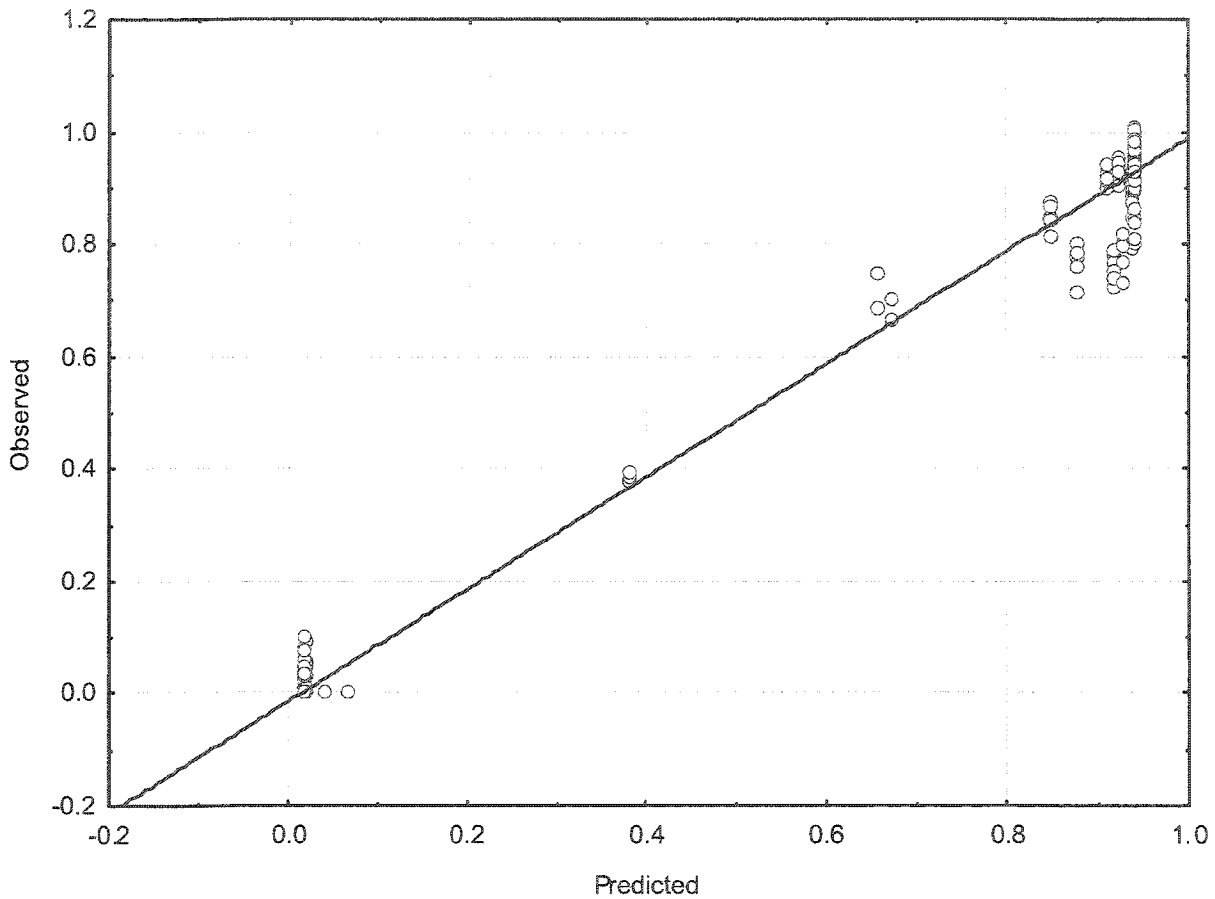

Figure 29. Scatterplot comparing the predicted time-allocation values for subject M20 to the observed time-allocation values.

The difference between the two baseline conditions for subject M22 is small but statistically significant. As Table 6 shows, neither the confidence interval for Δf_{hm} nor the interval for ΔP_e includes zero. The amount of variance accounted for, r^2, is 0.985. The Figure 30 presents the corresponding scatterplot.

As these results show, the only baseline datasets that could be combined for purposes of examining the effect of changing the train duration were those for subject M14. For subjects M11, M16, M20, and M22, separate baseline mountains were fitted to each baseline condition and compared individually with the experimental condition: the mountain obtained at the shorter train duration.

Quantifying the effect of manipulating train duration

Two-dimensional depiction.

The next step was to use the mountain model to address the fundamental question posed in this thesis: What is the effect of manipulating the train duration on the position of the mountain?

In the case of the four subjects whose data sets involved significant shifts in the location of the second baseline mountain, the "dual-mountain" model developed for assessing the baseline shifts was extended to apply to each of the two baseline conditions and the short train duration condition. To be more specific, the surface that fitted to the data from the experimental condition, TDP25, was compared separately with each of the surfaces that fitted to the data from the TD1A and TD1B baseline conditions. Two additional sets of shift parameters were added to the dual-mountain equation. The parameters $\Delta f_{hm}1$ and ΔP_e1 describe the shift from the experimental condition to the first baseline condition, and the parameters $\Delta f_{hm}3$ and ΔP_e3 describe the shift from the

M22 DOUBLE MOUNTAIN (Baselines):	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPe: Delta Price [shift along the price axis for TD1A from TD1B]	-0.082	0.018	-4.463	0.000	-0.118	-0.046
ΔFhm: Delta Frequency [shift along the frequency axis for TD1A from TD1B]	-0.019	0.008	-2.369	0.019	-0.034	-0.003
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.033	0.009	3.805	0.000	0.016	0.050
T_{MAX}: Maximum time Allocation	0.981					
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	3.784	0.552	6.858	0.000	2.697	4.870
P_{MIN}: minimum subjective price	0.955	0.055	17.378	0.000	0.847	1.063
TP: Transition Parameter	0.081	0.074	1.094	0.275	-0.065	0.226
P_{E2}: parameter determining location along the price axis [for TD1B]	1.323	0.017	77.253	0.000	1.289	1.356
g: growth exponent [determines the slope of mountain on Price façade]	6.994	0.747	9.364	0.000	5.523	8.465
F_{HM2}: parameter determining location along the frequency axis [for TD1B]	1.468	0.013	112.510	0.000	1.442	1.493

Table 6. Shift values for subject M22.

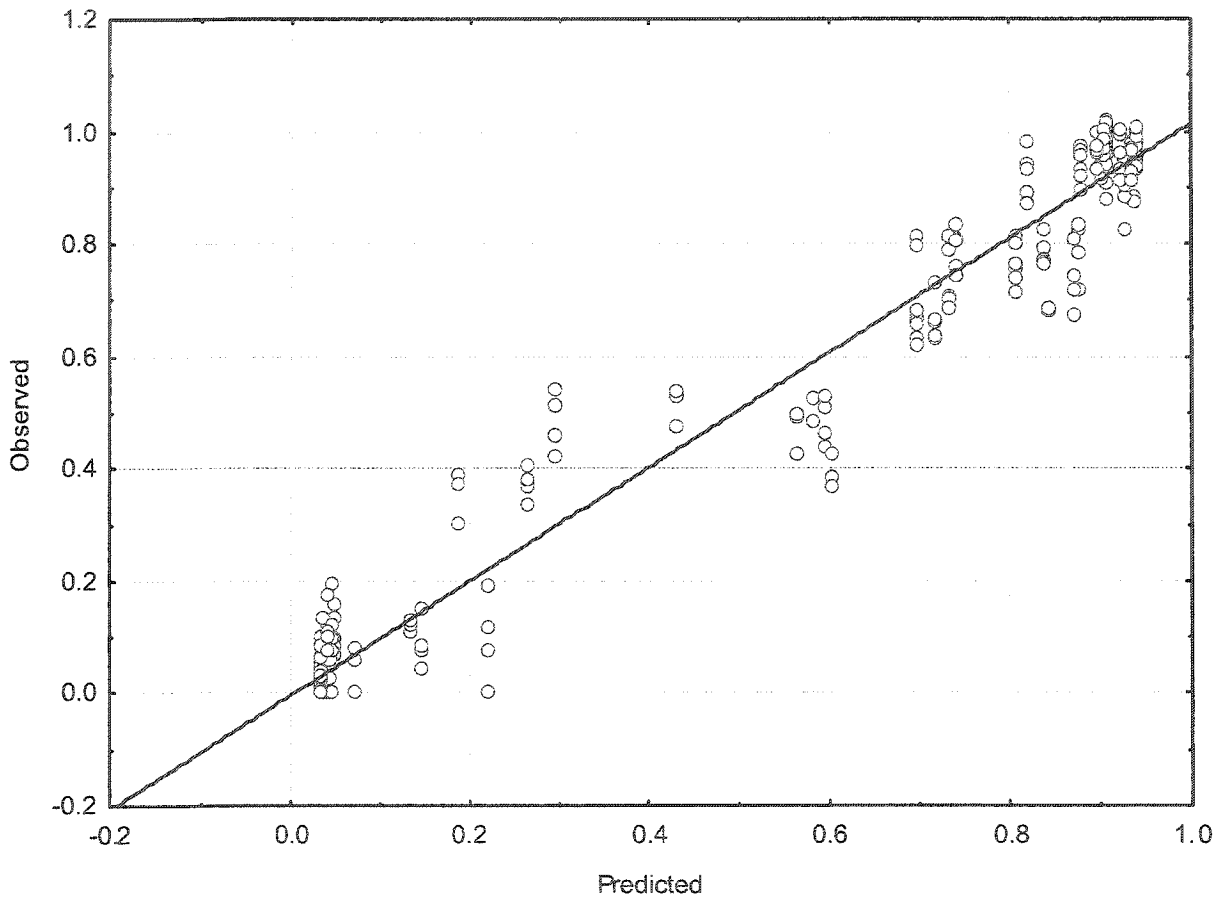

Figure 30. Scatterplot comparing the predicted time-allocation values for subject M22 to the observed time-allocation values.

experimental condition to the second baseline condition. In the case of the one rat whose baseline data were combined, subject M14, a dual mountain was fitted in order to allow a comparison of the combined baseline data with the data from the experimental condition. In the case of the dual-mountain fit, the assessment considered only one set of shift parameters, Δf_{hm} and ΔP_e.

In all subjects, reducing the train duration from 1 second to 0.25 seconds resulted in a marked rightward shift along the frequency axis but little shift along the price axis. This outcome will be illustrated below by 2D graphs for all subjects and by 3D graphs for one subject. Figure 31 presents the 2D graphs for subject M11. As each of the two upper panels of the figure shows, reducing the train duration from 1 second to 0.25 seconds produced a large shift in the data points and the fitted function along the frequency axis. On the other hand, as the lower panel shows, the reduction in the train duration failed to produce a reliable displacement of the data and the fitted functions along the price axis. Note that although there is some systematic deviation of data points from the fitted functions (upper-right panel) in general the model describes the data well.

Figures 32, 33 and 34 show that the shifts in the data sets for subjects M16, M20, and M22 are very similar to the shifts in the set for subject M11. As was mentioned previously, M14 was the only subject whose baseline conditions could be combined. Figure 35 shows the results for subject M14 in 2D space. As in the case of the other four subjects, there is a large shift in the position of the mountain along the frequency axis and very little shift if any along the price axis. The only difference between subject M14's graphs and the graphs for the other subjects is that the former require only one line to

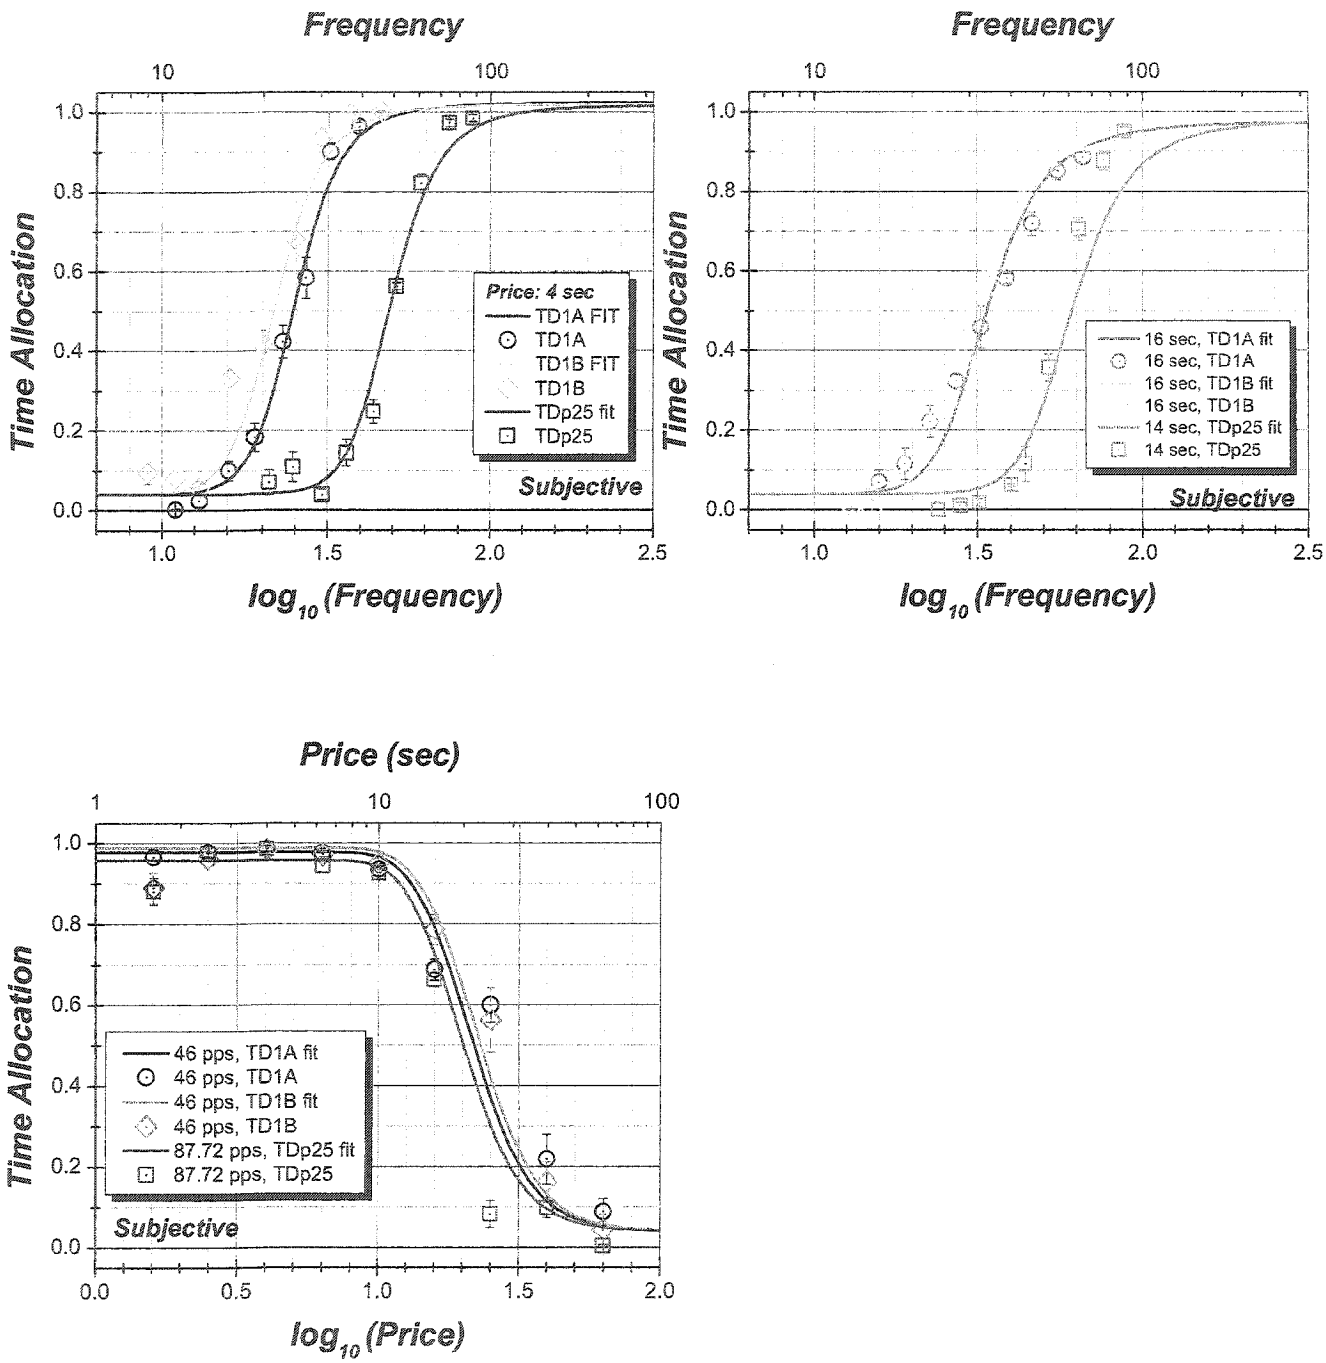

Figure 31. Reducing the train duration from 1 sec to 0.25 seconds resulted in a rightward shift along the frequency axis, seen in the frequency graphs above for M11. The price function graph (lower panel) shows no such shift in response to the decrease in train duration.

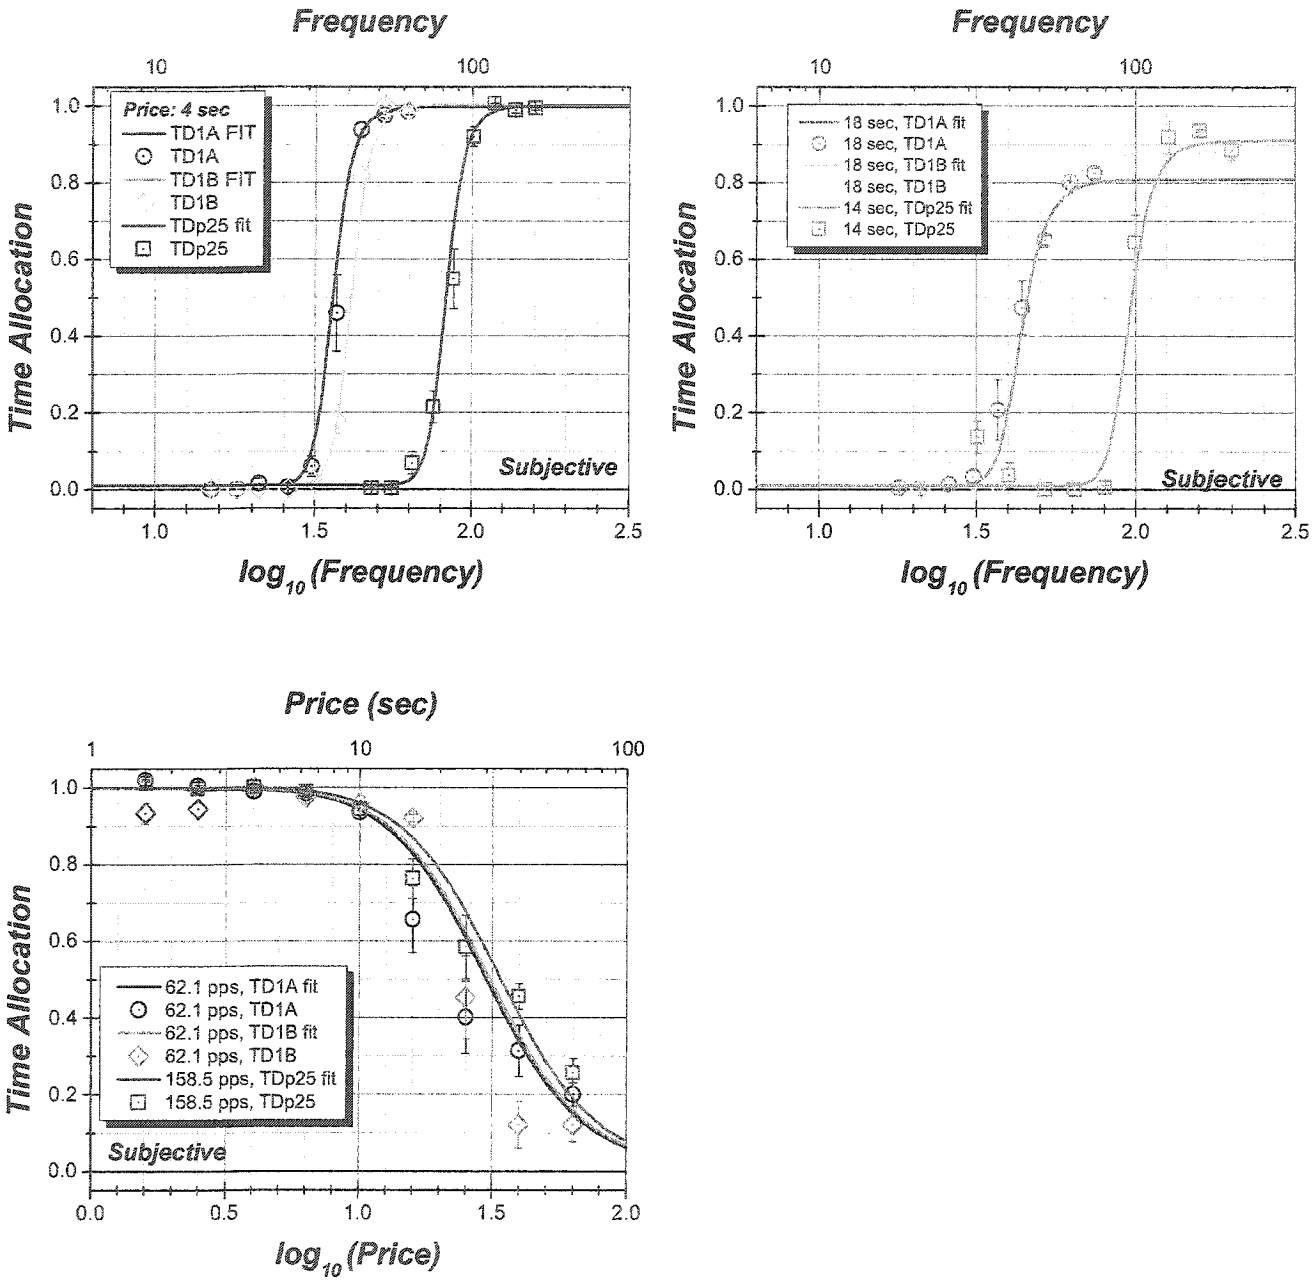

Figure 32. Decreasing the train duration from 1 sec to 0.25 seconds resulted in a rightward shift along the frequency axis, seen in the frequency graphs above for M16. The price function graph (bottom panel) shows no such shift in response to the decrease in train duration.

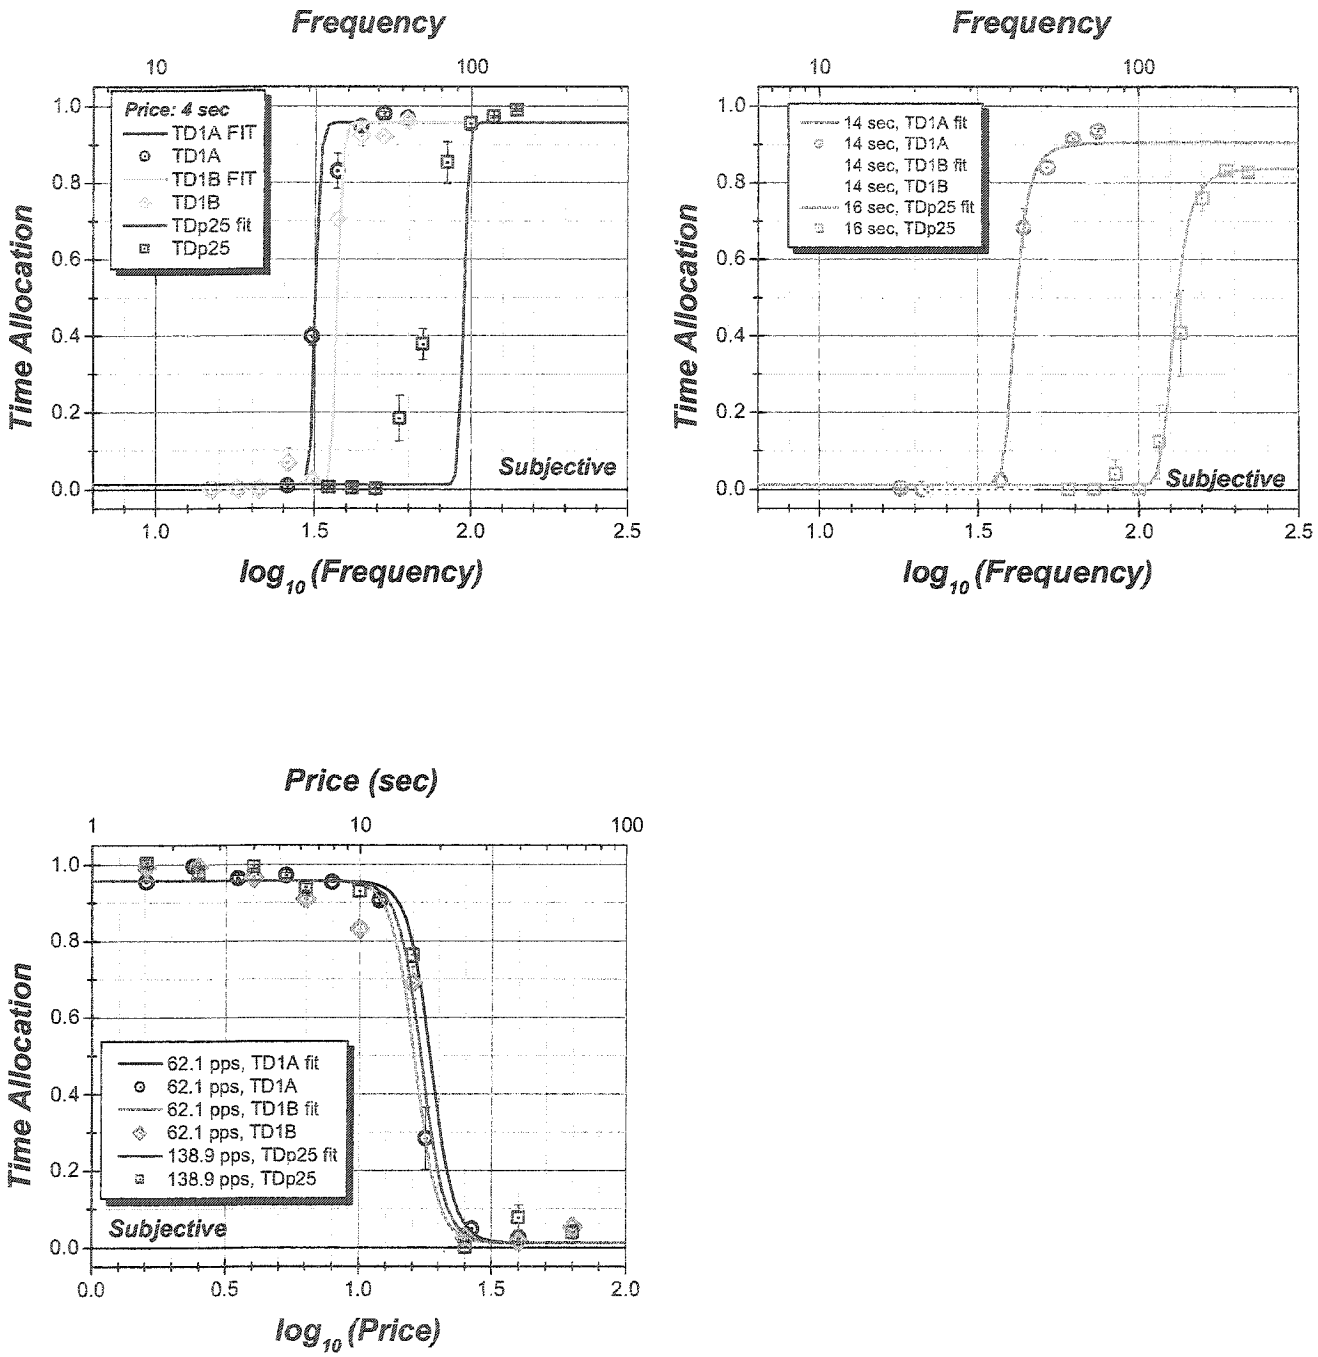

Figure 33. Decreasing the train duration from 1 sec to 0.25 seconds resulted in a rightward shift along the frequency axis, seen in the frequency graphs above for M20. The price function graph (lower panel) shows no such shift in response to the decrease in train duration.

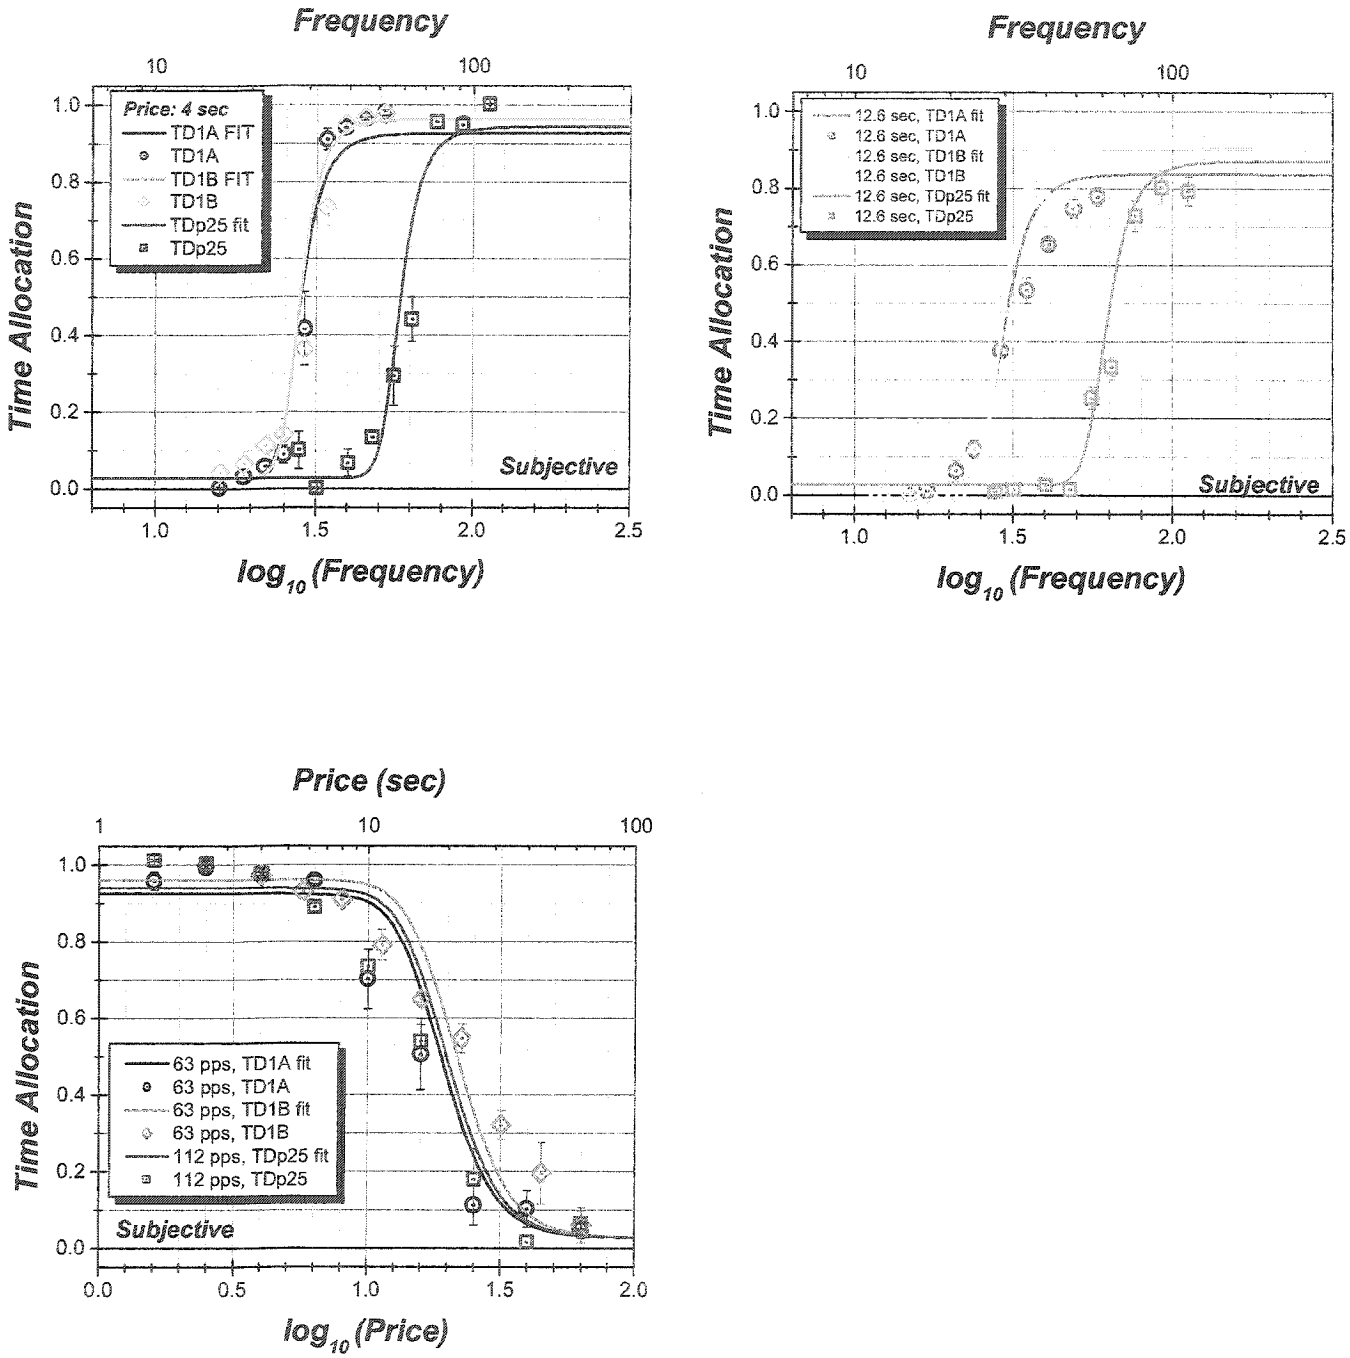

Figure 34. Decreasing the train duration from 1 sec to 0.25 seconds resulted in a rightward shift along the frequency axis, seen in the frequency graphs above for M22. The price function graph (lower panel) shows no such shift in response to the decrease in train duration.

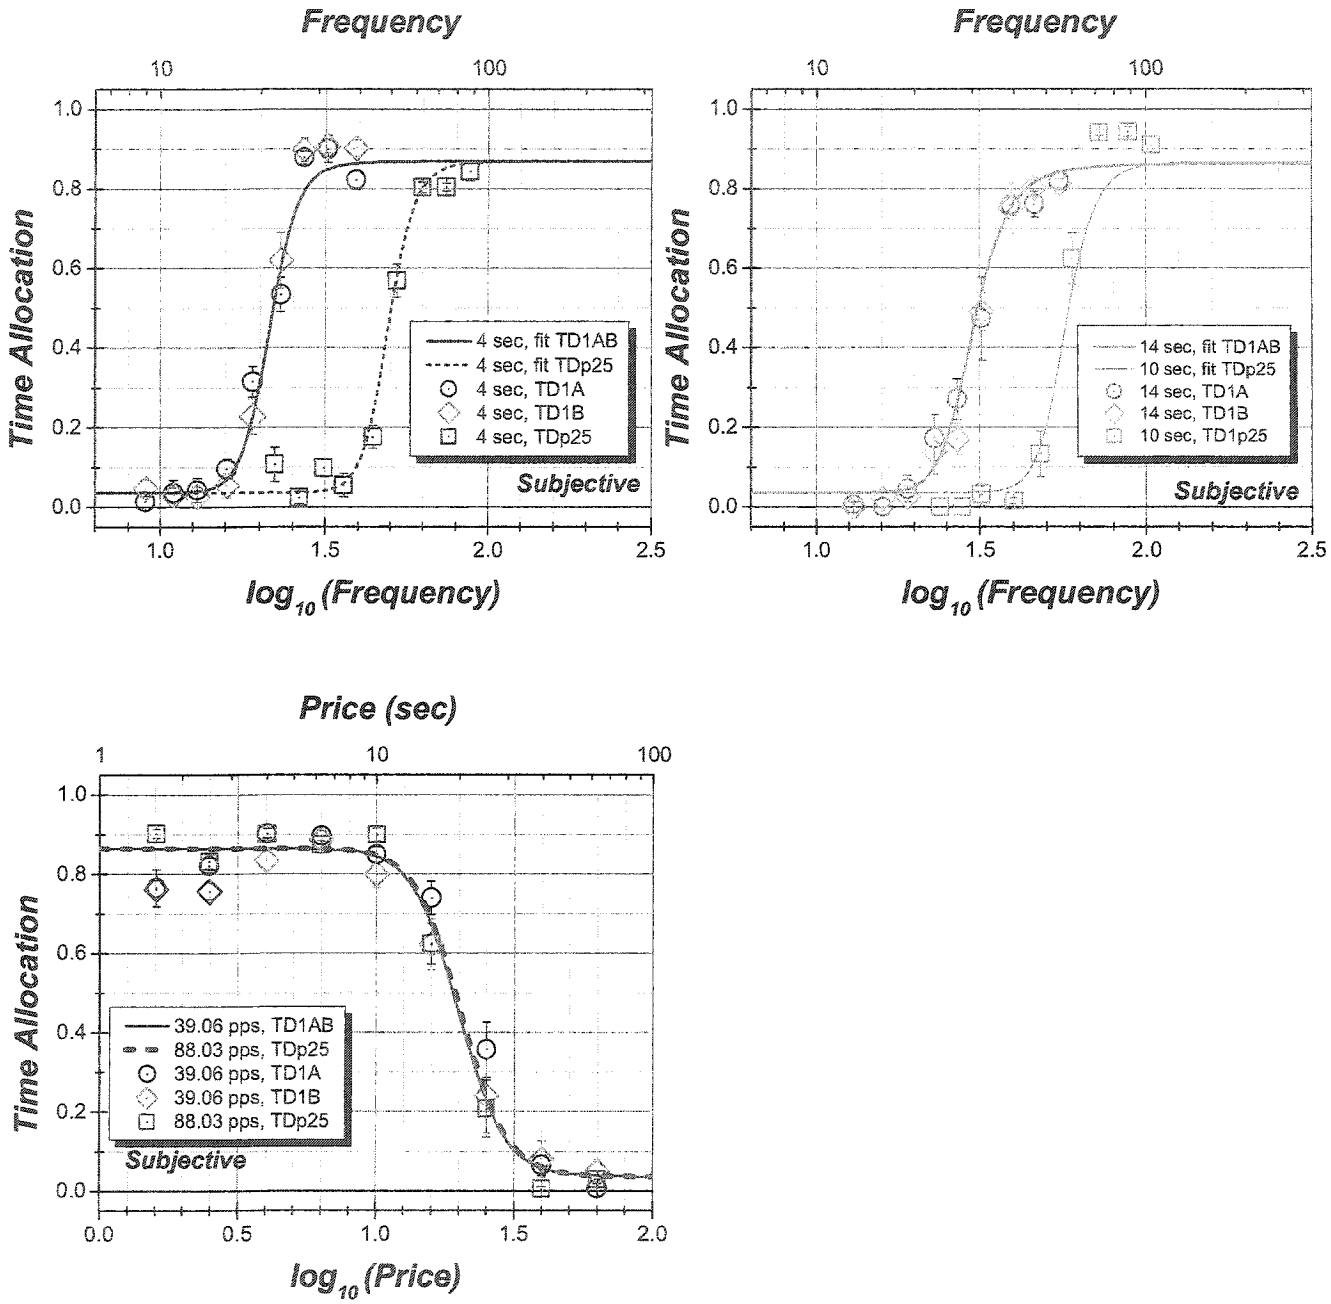

Figure 35. Decreasing the train duration from 1 sec to 0.25 seconds resulted in a rightward shift along the frequency axis, seen in the frequency graphs above for M14. The price function graph (lower panel) shows no such shift in response to the decrease in train duration.

represent the combined baseline condition. The solid lines in Figure 35 represent the combined baseline, and the dashed lines represent data for the shorter train duration.

Three-dimensional depiction.

The four mountains portrayed in Figure 36 represent M16's data in 3D space -- that is, they integrate the data from all three sweeps within each condition. The position of each mountain is best assessed by examining the two back walls behind the mountain. The amount of grey wall visible in the TDP25 condition (the upper right panel in Figure 36) is much greater than the amount visible in either of the baseline conditions; in the TDP25 condition, in other words, the mountain has receded along the frequency axis. On the other hand, a comparison of the purple walls in the baseline condition with the purple wall in the TDP25 condition reveals very little movement along the price axis.

Figure 37 presents contour graphs for subject M16 that correspond to the mountains in the previous figure. The position of the mountain is indicated here by the shifts in the position of the white dashed lines superimposed on the contour graphs. In order to facilitate comparison, the figure introduces black dashed lines that connect the white dashed lines in each graph. Thus the shift in the line for the half-maximal frequency value (f_{hm}) between the baseline condition TD1A and the TDP25 condition is substantial, whereas the shift in the P_E line is very small. In addition, the gap between the price lines for TD1A and TD1B and the gap between the frequency lines for the two baseline conditions are both very small relative to the shift that corresponds to the TDP25 condition.

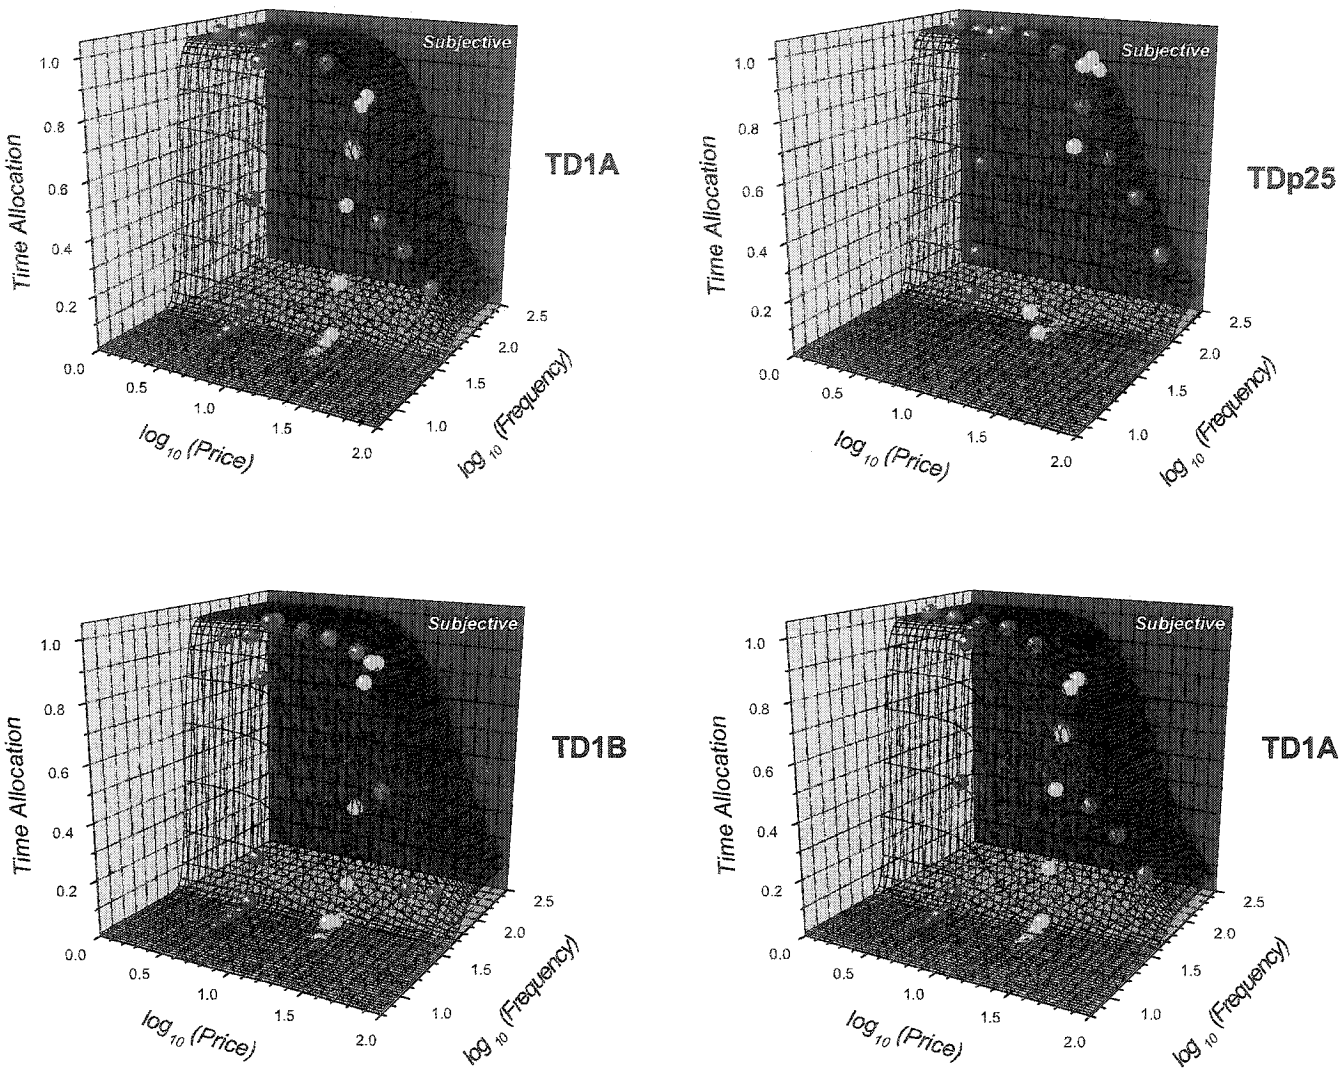

Figure 36. The four mountains above (for M16) enable a comparison between the baseline conditions (TD1A vs. TD1B) and the experimental condition, TDP25. The position of the mountain is best assessed by examining the size of the back wall on either side of the structure. The amount of grey wall visible in the TDP25 condition (top right graph) is much greater than either of the baseline conditions, the mountain has shifted back (along the frequency axis).

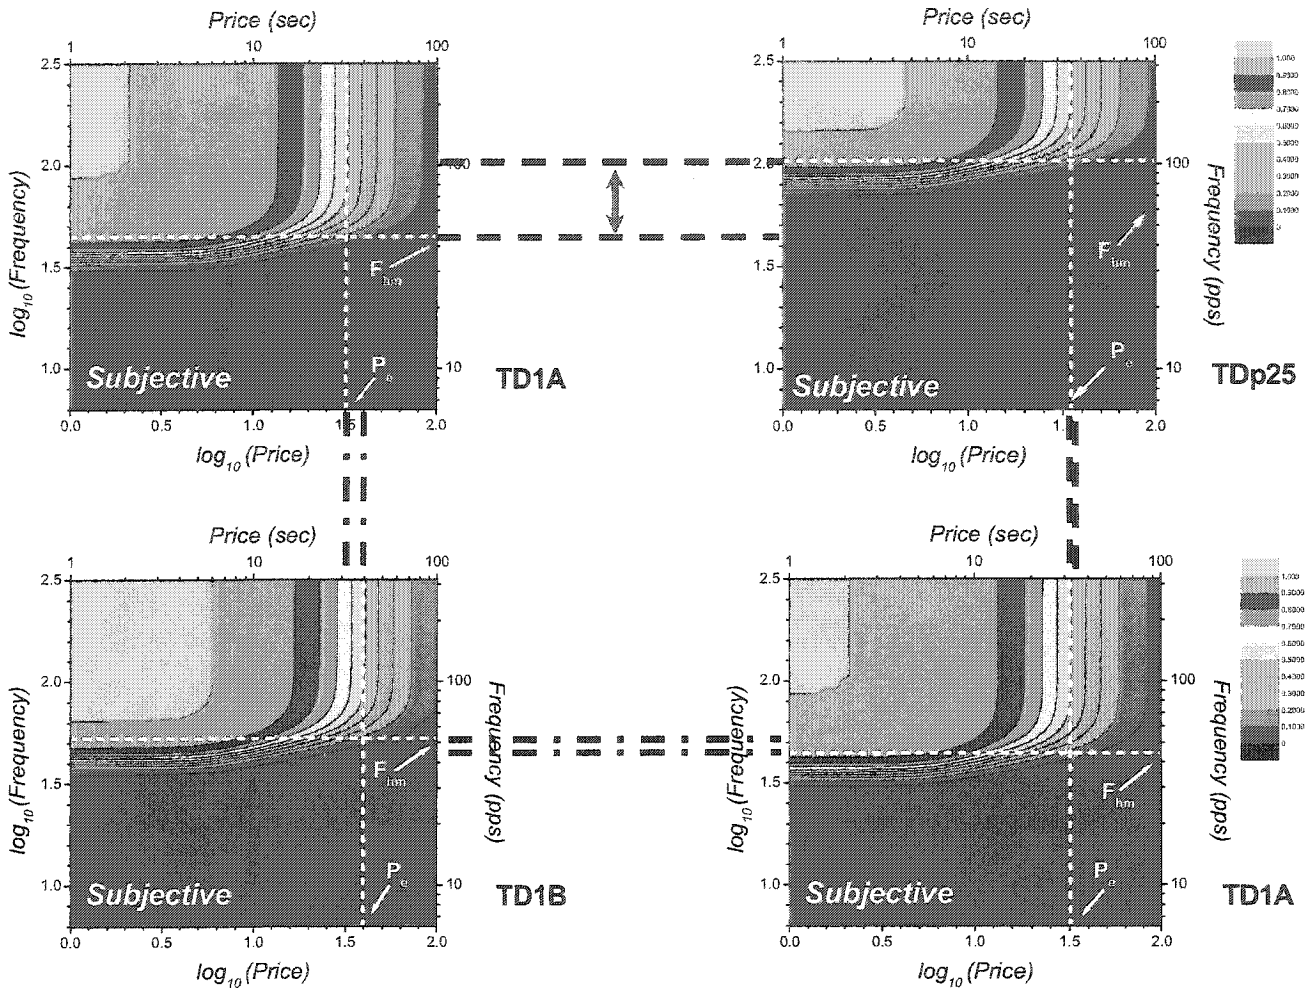

Figure 37. The contour graphs for subject M16 that correspond to the mountains in the previous figure. The position of the mountain is assessed here by the shifts of the white dashed lines. The position of the Frequency half-max line is greatly shifted in the top right figure compared to the baseline figure (top left) for TD1A. The white dashed lines have been extended with black dashed lines to facilitate the comparison between contour graphs.

Surface-fitting statistics.

Table 7 presents the statistical assessment for all five subjects of the effect on the position of the mountain of reducing the train duration. Table 7 is an analysis of variance (ANOVA) table; that is, it shows the partition of variance between the portion accounted for by the fitted surface and the portion accounted for by residual error. The corresponding *F*-ratio values are all very large, and the *p* values in the table for all five subjects exceed the statistical threshold of $p = 0.05$. These results are consistent with the scatterplot generated for each subject for the purpose of comparing the observed and predicted time-allocation values from the triple-mountain surface fit. Figure 38 shows the scatterplot for subject M16, which is representative of the scatterplots for the group of subjects as a whole. Table 8 displays the complete parameter set for the comparison between the two baseline conditions on the one hand and the experimental condition on the other for subject M16. The significant shift parameters and the corresponding confidence intervals are highlighted in yellow. Table 9 shows the fitting results for the four remaining subjects. The bar graphs in Figure 39 summarize for all five subjects the shifts in the position of the mountains that the reduction in the train duration produced. The bar graph for subject M16 shows clearly that reducing the train duration produced a much larger shift along the frequency axis than it did along the price axis. The final bar graph is the shift graph for subject M14, whose baseline data sets were combined to yield a dual-mountain rather than the triple-mountain fits obtained for the other subjects. In the case of this subject, the analysis compared only two shift parameters. The results for the dual-mountain fit are consistent with the results for the triple mountain fits: The shift along the frequency axis is very large, and the shift along the price axis is negligible.

m11		Sum of Squares	DF	Mean Squares	F-value	p-value
	Regression	183.03	12.00	15.25	2172.31	0.00
	Residual	2.97	423.00	0.01		
	Total	186.00	435.00			
	Regression vs.Corrected Total	183.03	12.00	15.25	97.14	0.00
	R^2	0.9840				
m14		Sum of Squares	DF	Mean Squares	F-value	p-value
	Regression	146.06	10.00	14.61	3332.94	0.00
	Residual	1.71	390.00	0.00		
	Total	147.77	400.00			
	Regression vs.Corrected Total	146.06	10.00	14.61	94.40	0.00
	R^2	0.9884				
m16		Sum of Squares	DF	Mean Squares	F-value	p-value
	Regression	160.67	12.00	13.39	6896.63	0.00
	Residual	0.66	341.00	0.00		
	Total	161.33	353.00			
	Regression vs.Corrected Total	160.67	12.00	13.39	66.93	0.00
	R^2	0.9959				
m20		Sum of Squares	DF	Mean Squares	F-value	p-value
	Regression	147.82	12.00	12.32	4194.71	0.00
	Residual	1.04	353.00	0.00		
	Total	148.86	365.00			
	Regression vs.Corrected Total	147.82	12.00	12.32	59.68	0.00
	R^2	0.9930				
m22		Sum of Squares	DF	Mean Squares	F-value	p-value
	Regression	157.34	10.00	15.73	2641.18	0.00
	Residual	2.27	381.00	0.01		
	Total	159.61	391.00			
	Regression vs.Corrected Total	157.34	10.00	15.73	94.19	0.00
	R^2	0.9858				

Table 7. The ANOVA results for all subjects comparing the non-linear curve-fitter results to the observed data.

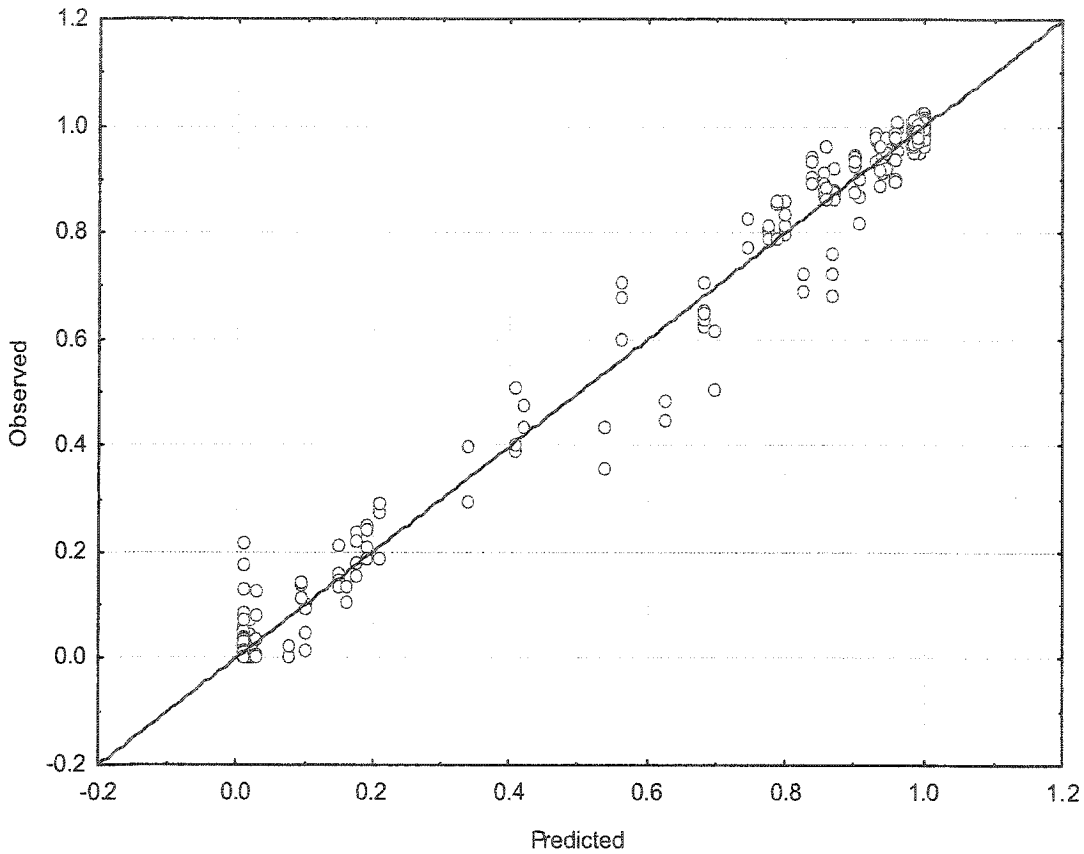

*Figure 38.*Observed time allocation versus Predicted time allocation from the curve-fitter for subject M16. Like all of the other subjects this is a good fit.

M16 TRIPLE MOUNTAIN:	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPe1: Delta Price [shift along the price axis for TD1A from TDP25]	-0.037	0.017	-2.153	0.032	-0.071	-0.003
ΔFhm1 : Delta Frequency [shift along the frequency axis for TD1A from TDP25]	-0.366	0.006	-56.616	0.000	-0.379	-0.354
ΔPe3: Delta Price [shift along the price axis for TD1B from TDP25]	0.055	0.020	2.777	0.006	0.016	0.093
ΔFhm 3: Delta Frequency [shift along the frequency axis for TD1B from TDP25]	-0.293	0.006	-52.999	0.000	-0.304	-0.282
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.010	0.004	2.571	0.011	0.002	0.018
T_{MAX}: Maximum time Allocation	1.014	0.006	156.928	0.000	1.001	1.027
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	2.376	0.086	27.730	0.000	2.208	2.545
P_{MIN}: minimum subjective price	0.723	0.076	9.535	0.000	0.574	0.872
TP: Transition Parameter	0.085	0.133	0.642	0.522	-0.176	0.346
P_{E2}: parameter determining location along the price axis [for TDp25]	1.542	0.014	106.869	0.000	1.513	1.570
g: growth exponent [determines the slope of mountain on Price façade]	7.317	0.302	24.194	0.000	6.722	7.912
F_{HM2}: parameter determining location along the frequency axis [for TDP25]	2.015	0.005	396.514	0.000	2.005	2.025

Table 8. M16's surface-fitting results. All parameters were held constant for the dual mountain fit and the shift parameters were assessed. A statistically significant shift of the mountain along the frequency axis would exist if the confidence intervals did not include zero, see the highlighted rows below.

M11 TRIPLE MOUNTAIN:	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPE1: Delta Price [shift along the price axis for TD1A from TDP25]	0.048	0.028	1.676	0.094	-0.008	0.103
Δfhm1 : Delta Frequency [shift along the frequency axis for TD1A from TDP25]	-0.266	0.024	-11.206	0.000	-0.313	-0.219
ΔPE3: Delta Price [shift along the price axis for TD1B from TDP25]	0.019	0.028	0.679	0.498	-0.037	0.075
Δfhm3: Delta Frequency [shift along the frequency axis for TD1B from TDP25]	-0.343	0.025	-13.929	0.000	-0.391	-0.295
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.040	0.009	4.222	0.000	0.021	0.058
T_{MAX}: Maximum time Allocation	1.035	0.018	56.052	0.000	0.999	1.071
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	3.923	0.257	15.278	0.000	3.418	4.428
P_{MIN}: minimum subjective price	1.011	0.015	68.427	0.000	0.982	1.040
TP: Transition Parameter	0.031	0.052	0.598	0.551	-0.072	0.134
P_{E2}: parameter determining location along the price axis [for TDp25]	1.459	0.043	34.299	0.000	1.375	1.543
g: growth exponent [determines the slope of mountain on Price façade]	2.317	0.197	11.775	0.000	1.931	2.704
F_{HM2}: parameter determining location along the frequency axis [for TDP25]	1.809	0.038	47.785	0.000	1.735	1.884

M14 DOUBLE MOUNTAIN:	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPE1: Delta Price [shift along the price axis for TD1A from TDP25]	-0.052	0.050	-1.059	0.290	-0.150	0.045
Δfhm1 : Delta Frequency [shift along the frequency axis for TD1A from TDP25]	-0.390	0.028	-13.797	0.000	-0.445	-0.334
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.036	0.006	6.094	0.000	0.025	0.048
T_{MAX}: Maximum time Allocation	0.869	0.007	116.830	0.000	0.854	0.884
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	5.176	0.327	15.819	0.000	4.533	5.819
P_{MIN}: minimum subjective price	0.870	0.017	50.067	0.000	0.836	0.904
TP: Transition Parameter	0.009	88.435	0.000	1.000	-173.860	173.878
P_{E2}: parameter determining location along the price axis [for TDp25]	1.615	0.087	18.545	0.000	1.443	1.786
g: growth exponent [determines the slope of mountain on Price façade]	2.528	0.248	10.202	0.000	2.041	3.016
F_{HM2}: parameter determining location along the frequency axis [for TDP25]	1.952	0.063	31.224	0.000	1.829	2.075

M20 TRIPLE MOUNTAIN:	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPE1: Delta Price [shift along the price axis for TD1A from TDP25]	-0.017	0.013	-1.386	0.167	-0.040	0.007
Δfhm1 : Delta Frequency [shift along the frequency axis for TD1A from TDP25]	-0.477	0.010	-50.029	0.000	-0.500	-0.458
ΔPE3: Delta Price [shift along the price axis for TD1B from TDP25]	-0.066	0.007	-9.098	0.000	-0.080	-0.052
Δfhm3: Delta Frequency [shift along the frequency axis for TD1B from TDP25]	-0.417	0.008	-52.027	0.000	-0.430	-0.401
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.012	0.004	2.747	0.006	0.000	0.020
T_{MAX}: Maximum time Allocation	0.957	0.006	171.767	0.000	0.950	0.968
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	9.723	1.067	9.114	0.000	7.630	11.822
P_{MIN}: minimum subjective price	0.521	3510.728	0.000	1.000	-6904.050	6905.094
TP: Transition Parameter	0.007	634.991	0.000	1.000	-1248.830	1248.849
P_{E2}: parameter determining location along the price axis [for TDp25]	1.291	0.012	108.158	0.000	1.270	1.314
g: growth exponent [determines the slope of mountain on Price façade]	8.212	1.113	7.382	0.000	6.020	10.401
F_{HM2}: parameter determining location along the frequency axis [for TDP25]	2.040	0.010	194.553	0.000	2.020	2.060

M22 TRIPLE MOUNTAIN:	Estimate	Standard	t-value	p-level	Lo. Conf	Up. Conf
SHIFT PARAMETERS:						
ΔPE1: Delta Price [shift along the price axis for TD1A from TDP25]	-0.026	0.013	-2.056	0.040	-0.051	-0.001
Δfhm1 : Delta Frequency [shift along the frequency axis for TD1A from TDP25]	-0.327	0.008	-40.565	0.000	-0.343	-0.312
ΔPE3: Delta Price [shift along the price axis for TD1B from TDP25]	0.036	0.015	2.386	0.018	0.006	0.066
Δfhm3: Delta Frequency [shift along the frequency axis for TD1B from TDP25]	-0.310	0.008	-37.733	0.000	-0.326	-0.294
LOCATION, SHAPE, AND SCALE PARAMETERS:						
T_{MIN}: Min Time Allocation	0.027					
T_{MAX}: Maximum time Allocation	1.012					
a: behavioral allocation exponent [determines the slope of mountain on Freq façade]	4.215	0.872	4.835	0.000	2.501	5.929
P_{MIN}: minimum subjective price	1.025	0.087	11.738	0.000	0.853	1.197
TP: Transition Parameter	0.134	0.055	2.429	0.016	0.026	0.243
P_{E2}: parameter determining location along the price axis [for TDp25]	1.295	0.043	30.376	0.000	1.211	1.379
g: growth exponent [determines the slope of mountain on Price façade]	6.153	0.566	10.875	0.000	5.040	7.265
F_{HM2}: parameter determining location along the frequency axis [for TDP25]	1.761	0.021	84.354	0.000	1.720	1.802

Table 9. The curve-fitting results for the remaining subjects: comparing the non-linear curve-fitter results to the observed data.

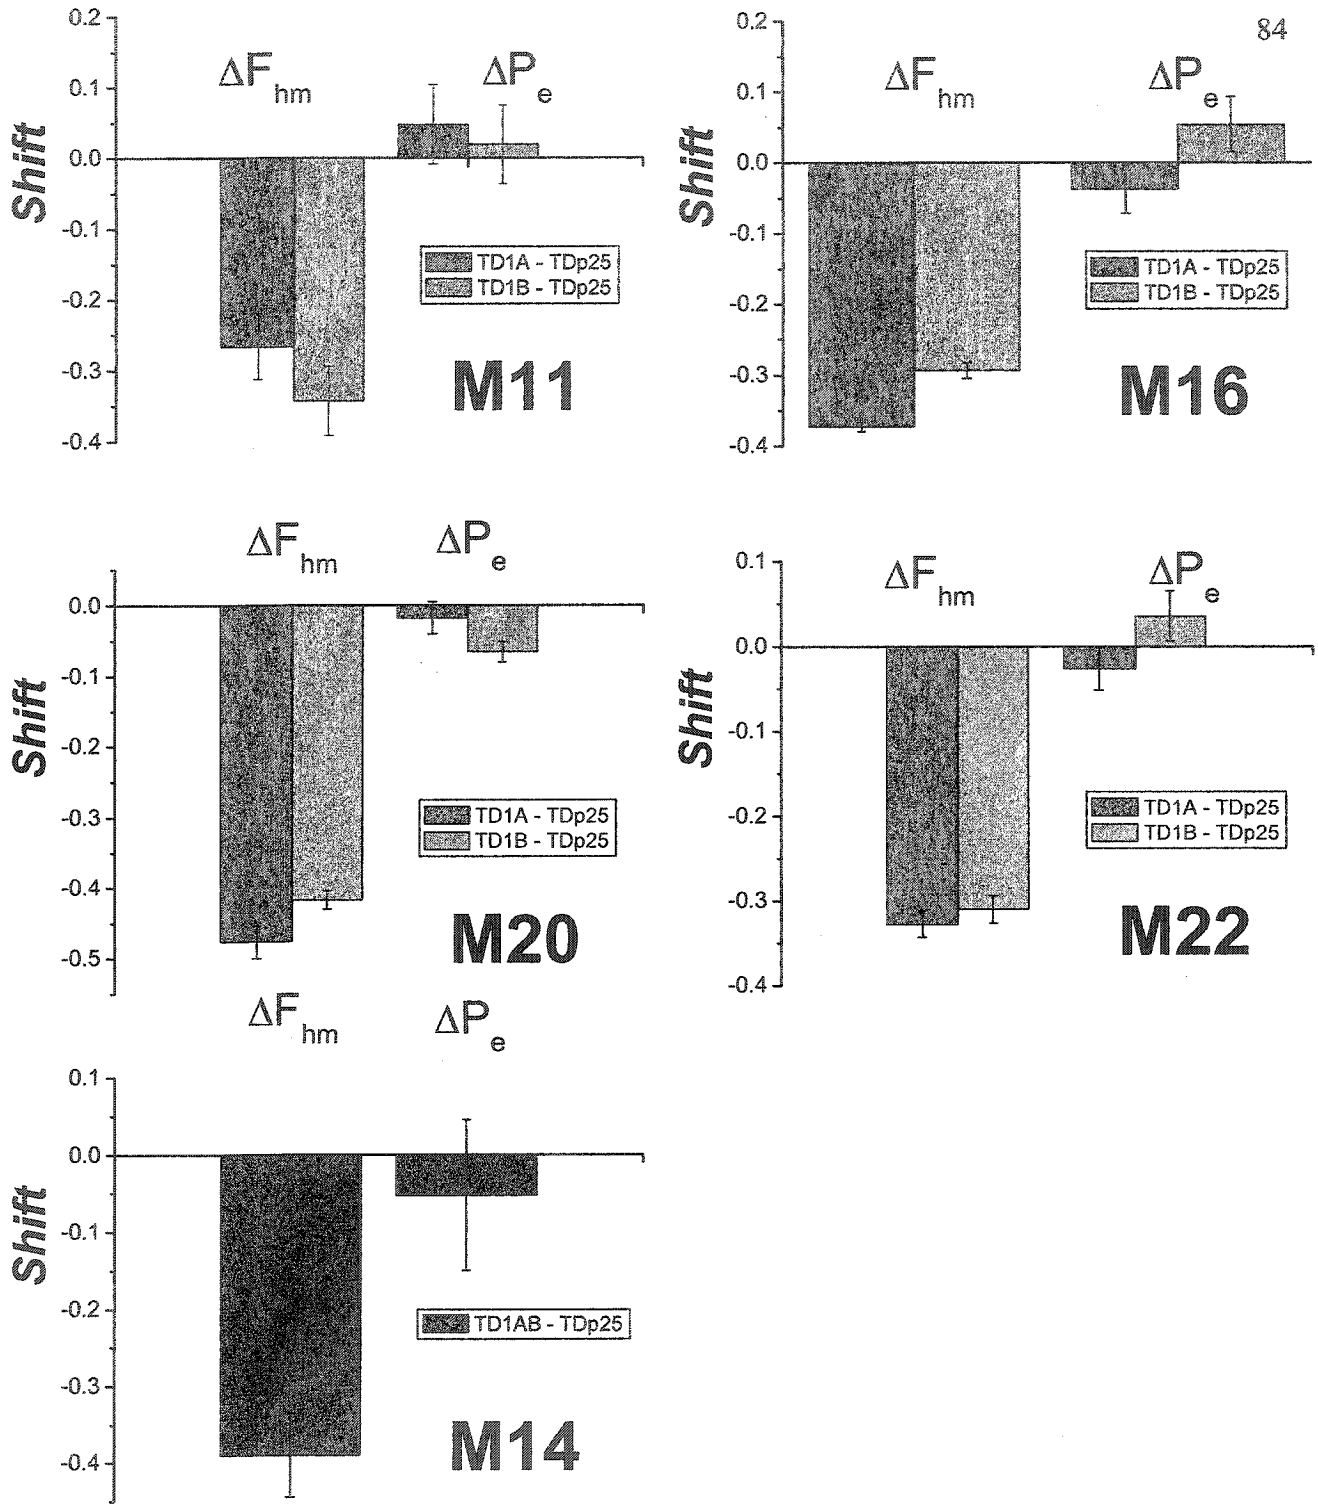

Figure 39. The four top-most bar graphs above portray the shifts of the triple mountain whereas the bottom left graph is the shift graph for M14 whose baseline conditions were combined, resulting in a dual mountain fit.

Discussion

The model tested in this experiment specifies how the relationship between stimulation strength and price determines reward strength (where time allocation is the corresponding behaviour). This thesis asks: Does this model work? To be more specific, does the three-dimensional psychometric function, the mountain, adequately describe the relationship between performance for brain stimulation reward (BSR), subjective reward intensity, and price? The experiment that this thesis describes tests the model by reducing the train duration. Does reducing the train duration result in a shift along the mountain's frequency axis alone, as the model predicts, or does it result in shifts along both the frequency axis and the price axis?

The results of the experiment support the original hypothesis – namely, that decreasing the duration of the pulse train will result in a shift along the mountain's frequency axis alone. The demonstration that the mountain can shift along one axis without shifting along the other supports the notion that the integrative model on which the mountain is based can distinguish manipulations that affect the reward system at different stages. As predicted, changing the train duration appears to affect the circuitry before the output of the reward-growth function and does not influence the behaviour-allocation function.

The first section below examines two technical issues, the leisure-time correction and test-retest reliability. The second examines two areas in which the model proved problematic. First, the model failed to reflect accurately the effect on the time-allocation versus frequency functions of varying the price of BSR. Second, it failed to provide a satisfactory mechanism for transforming objective prices into subjective prices. These

problems did not seriously degrade the fit of the model to the data, but they are important to understand so that they can be overcome in future work. The third section below discusses the principal findings and their implications. The fourth section discusses opportunities for future research. The final section offers some general conclusions.

Technical issues

Leisure-time correction.

The raw data – namely, the length of time that the lever was depressed during a trial -- was corrected for a bias toward under sampling work time. All releases of the lever that were less than 1 second long, were re-classified as work time. The rationale for this correction is as follows: informal observation revealed that during releases shorter than a second the rat remained near the lever, often with its paws extended, and was effectively still engaging in work, not in leisure. The time-allocation correction takes into account these bouts of bar-tapping and thus brings the performance measure into closer accord with the animal's observed behaviour.

The left-hand panel of Figure 40 shows the effect of the tapping correction on the low-price frequency sweep condition for subject M14. The difference between the corrected time allocation and the uncorrected time allocation is considerable. The right-hand panel of Figure 40 shows the effect of the correction on the high-price frequency-sweep condition for the same subject. The difference between the uncorrected time-allocation value and the corrected value is much less here than it is in the left-hand panel. The effect of this correction on the parameters of the fitted surface is not large for any of the subjects, but it does increase the dynamic range of the data by raising the position of the upper asymptote. It follows that had the correction not been made the fit of the data

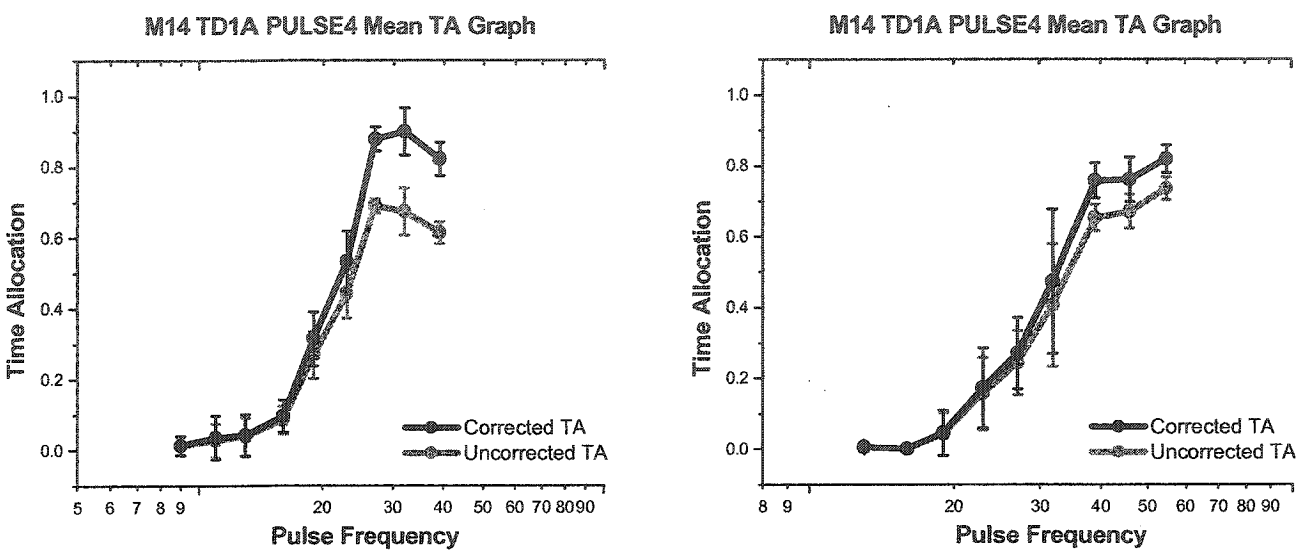

Figure 40. Presents the effect of the correction for the low-price frequency sweep condition for subject M14. There is a considerable difference between the corrected time allocation and the uncorrected time allocation. The right panel is the data for the same subject for the high price frequency sweep condition. The difference between the corrected time allocation value is much less here than it was in the left panel.

to the mountain would have been somewhat degraded. A shorter mountain would have resulted in a poorer fit, since it would have decreased the systematic variance attributable to the independent variable. Had the correction not been made, the ratio of variance accounted for by the fit to the total variance would likely have decreased.

Test-retest reliability.

Collecting the data for one complete mountain is time-consuming. As described earlier, it involves conducting two sets of frequency sweeps and one set of price sweeps, a process that can take up to 14 days. In order to rule out the possibility that the results could be explained by the passage of time rather than the decrease in train duration, the results were tested for internal consistency. If the mountain returned to its original position during the post-treatment baseline condition, then any shifts that accompanied the reduction in train duration would not be attributable to the passage of time. A comparison was made between the first baseline condition, TD1A, and the second baseline condition, TD1B. As was noted earlier, the data were stable overall. A comparison of the curves collected during the TD1A phase with those collected during the TD1B phase provides a rough qualitative impression of the remarkable test-retest reliability of the results (see Figures 19 to 23). The shift from one baseline condition to the other produced small changes for most subjects, but these changes were very small relative to the effects of manipulating train duration. Thus, it is reasonable to ascribe the effect of varying the train duration to the temporal integration of the reward signal and not to a drift of the baseline over time.

Less successful predictions of the mountain model

Price effect.

According to the model, time allocation depends upon both subjective reward intensity and reward price. In turn, reward intensity is determined by the frequency of reward. To hold time allocation constant, one can vary price and reward such that one variable compensates for the other. When the price is high, motivating the animal to work for stimulation requires higher frequencies. Consequently, the model predicts that an increase in price should shift the position of frequency-sweep curves to the right: the higher frequencies yield higher reward intensities, thus offsetting the effect of the higher price. Reward intensity cannot grow to infinity, however, and it must saturate as frequency increases to higher and higher values. Thus, the logistic function used to model the growth of reward intensity levels off at very high frequencies. This outcome ultimately produces a failure in the ability of the stimulation frequency to compensate for price increases. At sufficiently high prices, the reward intensity required to sustain time allocation at a high level may exceed the maximum that can be produced. In other words, reward growth will have saturated. As a result, the model predicts that the upper asymptote of responding for these high-priced frequency sweeps will fall below the animal's maximum time allocation and that the slope of the function will become increasingly gradual.

The price effect observed was generally, but not always, as predicted. This outcome can be seen in Figure 41, which replots the data in Figures 31 through 35. At the longer train duration, the high-price frequency-sweep function was shifted to the right for all of the subjects and in all three phases of the experiment, with the exception of the

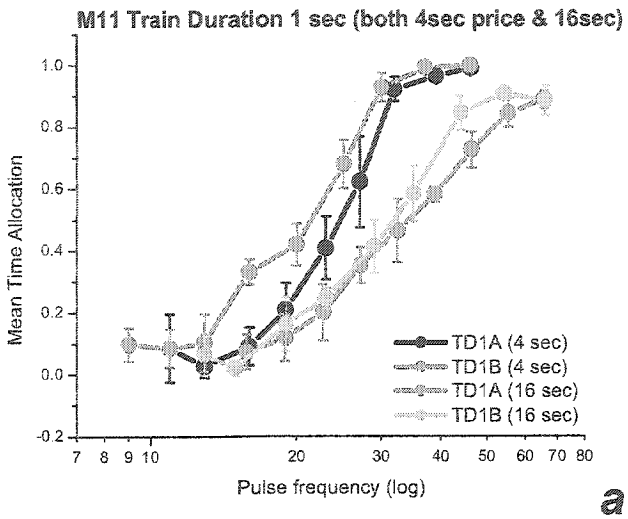

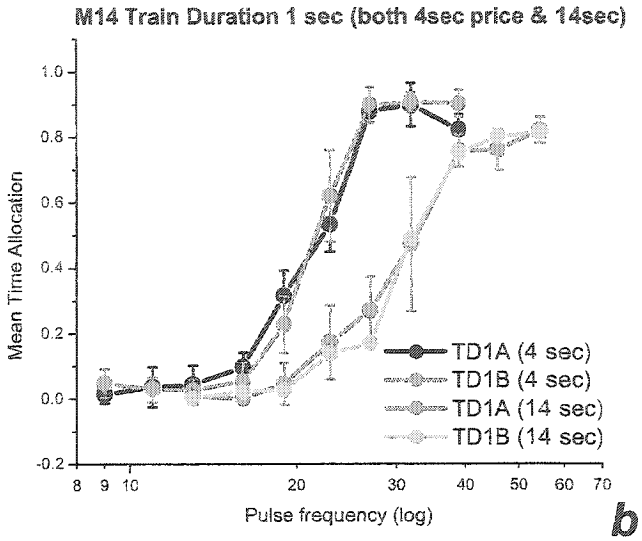

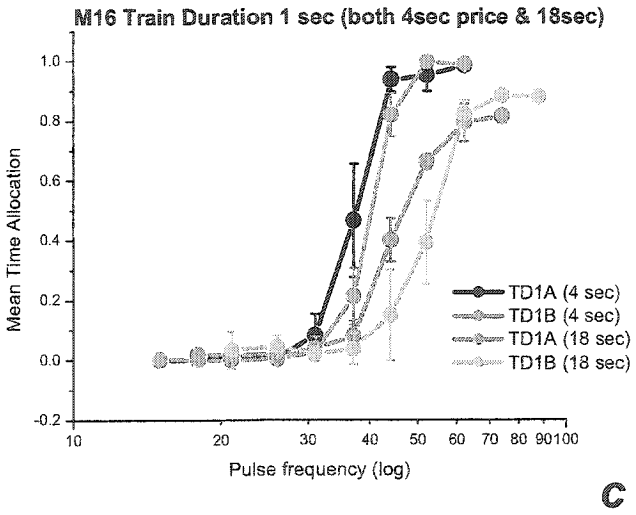

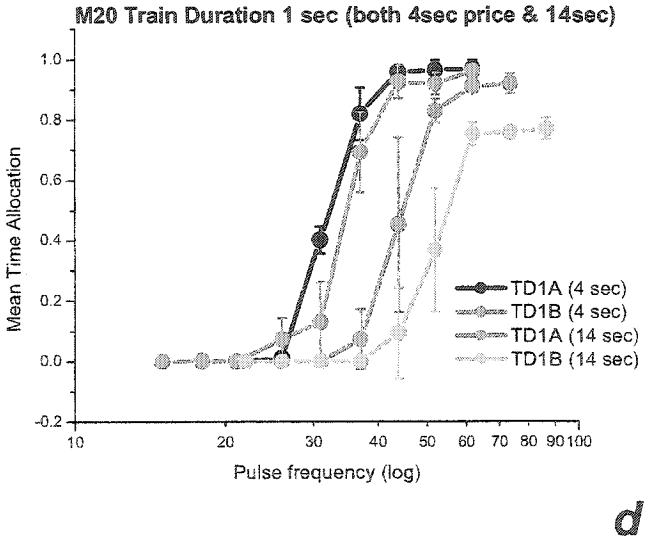

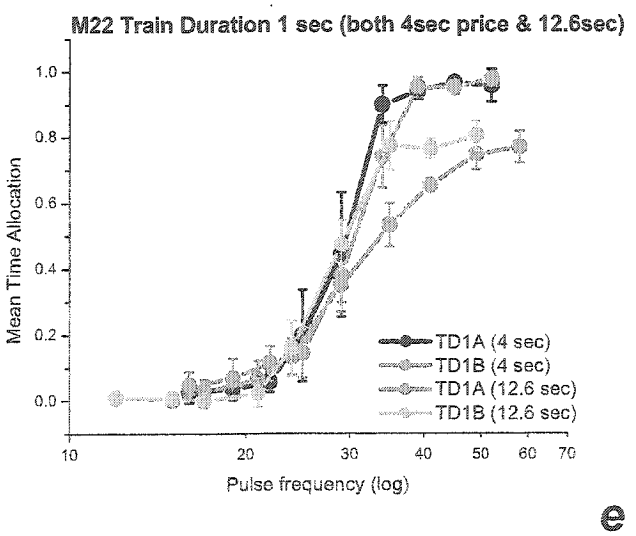

Figure 41. Replotted data from Figures 31-35. A summary of the effect of increasing the price for all 5 subjects at the longer train duration. The dark red and light red data represent the low-price frequency sweeps whereas the dark green and the light green data represent the high-price frequency sweeps.

TD1B condition for rat M22 (Figure 41e). However, at the point on the psychometric functions where time allocation begins to rise (the toe), the low- and high-price curves overlapped for three of the five subjects (M11, M16, and M22). At the shorter train duration, there were clear shifts in only three of the five sets of curves (Figure 42). Moreover, the fact that the toe or the position of the function appears to be unaltered by the increase in price, for some subjects at some conditions, is a deviation from the model. The model predicts that to obtain the same time allocation value with an increase in price requires a corresponding increase in frequency. Subsequent work (Breton, 2004) has suggested that the values of the parameters of the behaviour-allocation function may drift during prolonged testing at high prices. In other words, the act of measuring the mountain may sometimes deform it! Remedies currently being tested include the randomized selection of prices and frequencies and the bracketing of experimental trials by fixed trials in which the payoff from BSR is either maximal or minimal. In the case of the present study, the deviation of the price effect from predictions did degrade the fit of the model. Nonetheless, the fit was generally very good which suggests that the impact of the problem was minor.

Subjective price effect.

As the introduction noted, the version of the mountain model used here includes a function that translates objective prices into their subjective equivalents (Equation 8). The reason for including this function is the expectation that once the price is sufficiently low, beyond a certain point further decreases in price no longer matter to the rat. The subjective-price function has two parameters, a minimum subjective price (P_{min}) and a transition parameter (TP). The function looks like a hockey stick: P_{min} sets both the

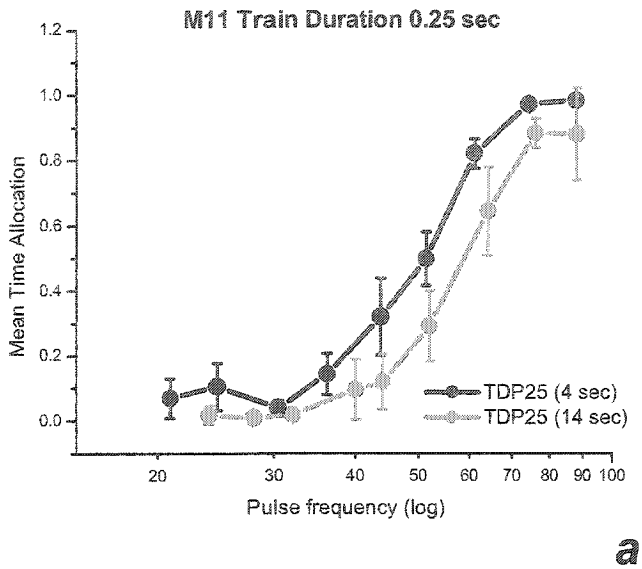

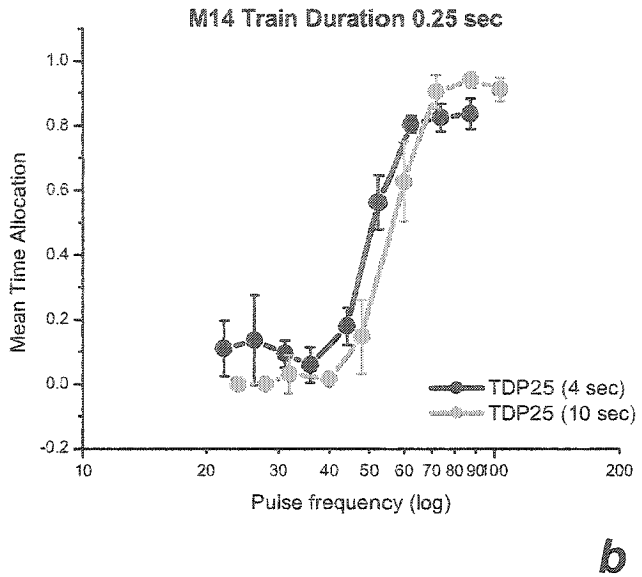

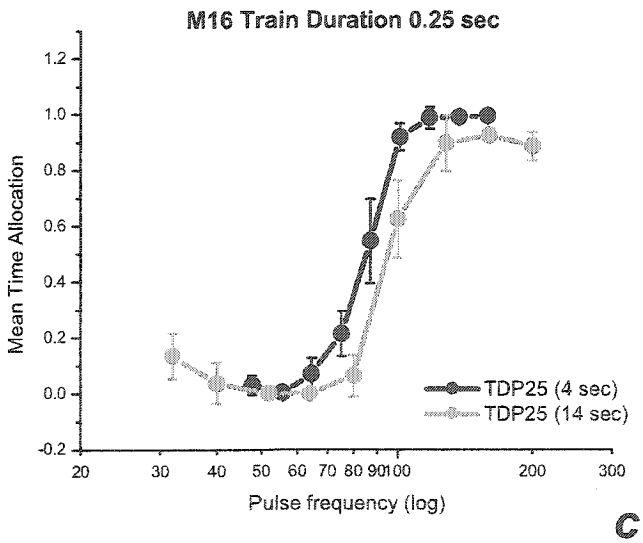

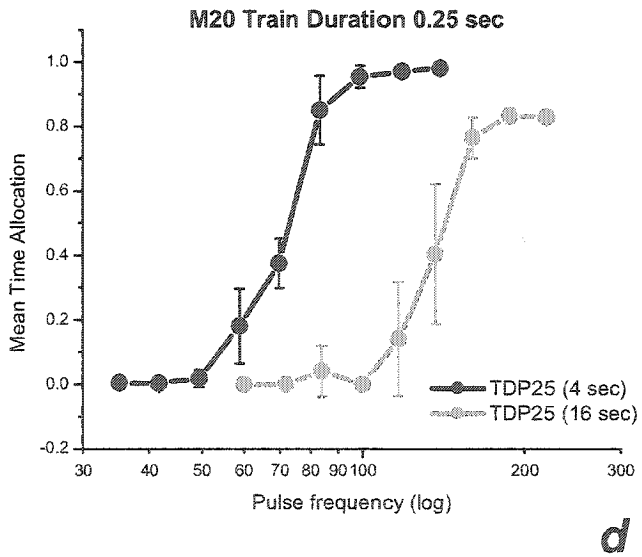

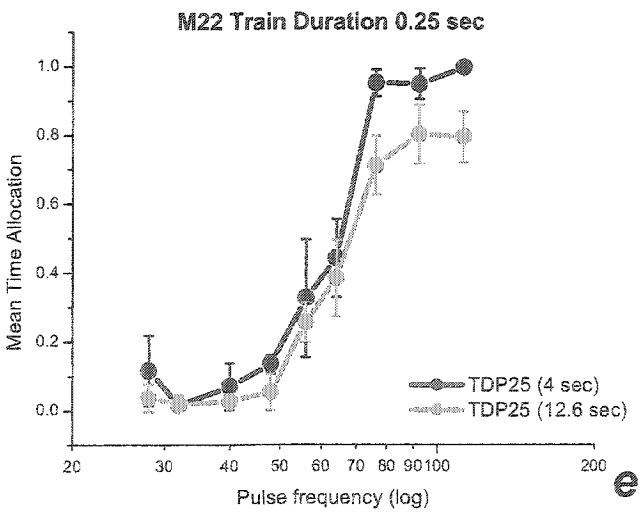

Figure 42. Replotted data from Figures 31-35. A summary of the effect of increasing the price for all 5 subjects at the shorter train duration. The red data represent the low-price frequency sweeps whereas the green data represent the high-price frequency sweeps.

height of the blade and the location of the transition between the blade and the handle; *TP* sets the abruptness of the transition.

The fitted values of P_{min} and *TP* were listed in Tables 8 and 9, above. For ready comparison, they have been listed together in Table 10. It is hardly surprising that the *TP* parameter often fails to deviate significantly from zero or that in two cases the model could not be fitted if this parameter was left free to vary. In order to accurately describe the hockey stick and the abruptness of the transition between the blade and the handle, many prices would have to be tested; this study has used only two different prices.

The fitted values of P_{min} range from 3.32 to 10.59 seconds; all but the smallest values deviate significantly from zero. These values would seem to imply that for subjects M11 and M22, there is no difference between prices of 1 second and 10 seconds. This interpretation contradicts everyday experience in the laboratory. Moreover, research carried out after the collection of the data reported here shows clearly that rats are able to distinguish between prices as low as 0.5 seconds and 1 second (Solomon, 2004). Indeed, inspection of the price-sweep data in Figure 23 (lower panel) shows that the allocation of time by subject M22 fell over a range of prices lower than 10 seconds. (That the allocation shown by the other subjects was maximal over this range does not imply that the prices were all equivalent. Keep in mind that the price sweeps were carried out at very high pulse frequencies.) Thus, time allocation was driven to its asymptotic value at low prices. Had price sweeps been carried out at lower frequencies, time allocation would likely have fallen over the 1 to 10 second range (Breton, 2004).

How can one explain the discrepancy between the fitted values of P_{min} and the empirical results (both those reported here and those collected subsequently)? Note that

Subject	log(P_{min})	P_{min}	p(P_{min})	Low price	High price	log(TP)	TP	TP fit?	p(TP)
M11	1.011	10.25	0.000	4	16.0	0.031	1.07	TRUE	0.551
M14	0.870	7.41	0.000	4	14.0	1.000	10.00	FALSE	1.000
M16	0.723	5.29	0.000	4	18.0	0.085	1.22	TRUE	0.522
M20	0.521	3.32	1.000	4	14.0	0.007	1.02	FALSE	1.000
M22	1.025	10.59	0.000	4	12.6	0.134	1.36	TRUE	0.016

Table 10. Fitted values obtained for all subjects for minimum price, P_{min}, and the transition parameter, *TP*.

in the case of subject M20 the P_{min} parameter did not deviate significantly from zero. Moreover, the surface fitting failed when the *TP* parameter was included and it became necessary to substitute a value selected by the experimenter. Thus, there was no need to have included the subjective-price function in the model for this subject; the objective price would have sufficed. It is interesting that in the case of this subject the price effect is as predicted by the model: the data from the frequency sweeps obtained at the high price are shifted substantially to the right of the data from the low-price sweeps (see panel d in both Figure 41 and Figure 42). Thus, in the case of the one subject for which the price effect was as predicted, there was no need for a subjective-price function. This outcome suggests that the significance of one parameter or both parameters of the subjective function in the case of the other subjects may be related to the deviation of the high-price frequency sweep from its predicted position.

Consider the case of subject M22, for which the low price was 4 seconds and the high price was12.6 seconds. Given the P_{min} value of 10.59 seconds, the model predicts very little rightward displacement of the curve from the high-price frequency sweep relative to the curve from the low-price sweep. Indeed, there is essentially no displacement of the empirical curves (see the lower-left panel of Figure 41), only a decrease in slope for the TD1A condition. The inclusion of the subjective-price function has allowed the model to compensate largely for its failure to predict correctly the effect of raising the price on the position of the frequency-sweep curve. The subjective-price function allowed the model to "pretend" that the subjective value of the high price was not meaningfully greater than the subjective value of the low price and thus to position the fitted curves for the two prices very close together.

Data collected in an ongoing experiment argue against the need for the subjective-price function in future work (Solomon, 2004). Remedies for the failure of the model to predict accurately the position of the high-price frequency sweep have been discussed above (bracketing and random selection of stimulation parameters). Data from another ongoing experiment (Breton, 2004) suggest that these remedies would be highly effective.

Now that the problematic aspects of the model have been discussed, it is possible to consider its successful aspects and the main objective of the thesis: to describe the effect of varying the train duration on the position of the mountain. As the goodness-of-fit statistics (Table 7) attest, the problems associated with the price and subjective-price effects (or the lack thereof) were not sufficiently serious to degrade the fit of the model to the data significantly.

More successful predictions of the model

The problems associated with the subjective-price component of the model notwithstanding, the proportion of variance accounted for by the fitted surface was above 0.975 for all subjects. Thus, the model does provide a good account of how time allocation depends on the strength (frequency) and price of BSR. Moreover, the fits were sufficiently accurate to describe the position of the mountain along the price and frequency axes, so that it was feasible to test the principal hypotheses: that it was possible for the mountain to move along only one of the two axes, and that in the case of manipulations of train duration this movement would be confined to the frequency axis.

Dissociating pre- and post-summation effects.

The three-dimensional analysis permits a more thorough assessment of the effect of varying the train duration than is possible with traditional rate-frequency measures. Unlike a two-dimensional representation, a mountain analysis can distinguish the effect of manipulations that alter the underlying circuitry before the output of the integrator (which shift the mountain along the frequency axis) from those that act after the output of the integrator (which shift the mountain along the price axis). If this portrayal is correct, it should be possible to shift the location of the mountain along only one of the two axes (frequency or price). The study tested this prediction by varying the train duration, a manipulation that should shift the mountain along the frequency axis alone.

Train duration effect.

Reducing the duration of a stimulation train reduces the time during which stimulation-induced post-synaptic excitation can build up in the reward substrate. The experimenter can compensate for the resulting decrease in reward intensity by increasing the strength of the stimulation. A higher-frequency train with a short duration can achieve the same level of reward intensity as a weaker long-duration train. Thus, decreasing the train duration should shift the reward-growth function to the right, and this movement, in turn, will shift the mountain along the frequency axis. Figures 36, 37, and 39 illustrate that shifts of this nature were indeed obtained.

The position of the mountain along the price axis was expected to remain stationary, since the impact of varying the train duration manifests itself at a stage of processing before the behaviour-allocation function.

A three-dimensional psychometric function can identify and discriminate the manipulations that correspond to either stage of the reward pathway. Figure 37 presented the contour graphs for subject M16 that correspond to the position of the 3D structure for the shorter train-duration condition (TDP25) and the baseline conditions (TD1A and TD1B). Decreasing the train duration caused a very large shift in f_{hm}, the frequency corresponding to a half-maximal reward; as the red arrow indicates, this shift along the frequency axis is 0.366 $\log_{10}$ units with respect to the TD1A condition. In contrast, there is a shift of only 0.037 $\log_{10}$ units between the position of the mountain in the short train-duration condition and the first baseline condition (upper-right and lower-right panels) and a shift of only ~0.1 $\log_{10}$ units along both axes. The effect of decreasing the train duration results in a very large shift (red double-headed arrow in Figure 37) in the position of the 3D structure, whereas returning to the baseline does not (see the closely spaced dot-dashed lines denoting the values of f_{hm} in the two baseline conditions).

Figure 39 showed that the shifts of the mountain in response to the decrease in train duration are consistently large for all subjects and are observed only along the frequency axis of the mountain. Much smaller and often inconsistent shifts were observed along the price axis. The average shift along the frequency axis is 8.8 times greater than the average shift along the price axis (Table 11). Thus, the results are highly supportive of the predictions of the mountain model and support previous findings by Arvanitogiannis (1997).

Subject	Average Frequency Shift	Average Price Shift	Difference
M11	0.305	0.034	9x
M14	0.390	0.050	8x
M16	0.330	0.046	7x
M20	0.447	0.042	10x
M22	0.319	0.031	10x

Table 11. Average shift sizes along both the frequency and price axes that correspond to the decrease in train duration.

Future research

The ability to determine at what stage different manipulations of the reward substrate exert their effects paves the way for future experiments that link different populations of neurons to particular stages in the processing of the reward signal. Such studies will facilitate the identification of the first-stage neurons responsible for BSR. In order to identify the primary neurons responsible for reward, several approaches need to be combined. A valid tool with which to measure alterations in reward-processing is required. The mountain analysis is a very good candidate for this role. Lesion methods are also required. The third approach required is to record from the candidate neurons. Accomplishing all three of these goals (stimulate/measure, lesion, and record) should lead to the identification of the first-stage neurons associated with reward.

Parallel experiments will be conducted in Dr. Shizgal's laboratory to assess whether the mountain can be shifted orthogonally, along the price axis alone. A shift of this nature would correspond to a manipulation that took place after the output of the integrator. For example, decreasing the value of leisure time would result in a shift along the price axis of the mountain. A proposed study involves setting up a bright light in the operant chamber that is on only during the subject's leisure time. Rats do not like bright lights, and the presence of a light is expected to alter the way in which the subjects choose to spend their time in the operant chamber. It is predicted that a manipulation of this kind would affect the output of the integrator and correspond to a shift of the mountain solely along the price axis.

Why is the mountain important?

The mountain may be a crucial tool for determining the structure of brain-reward circuitry, which is believed to play a key role in foraging behaviour and decision making. Decision processes are governed by reward assessment. Thus an animal will base its decision to remain in a patch for the purposes of foraging or to leave for another on relative reward values and costs (Gallistel, 1994). One can also apply mountain analysis to determine how the reward system is affected by different physiological manipulations, such as food restriction, and at what stage of processing these manipulations act to alter reward processing.

The role dopamine plays in BSR remains to be understood fully. Changes in dopaminergic neurotransmission profoundly alter performance for BSR. It is not known, however, at which stage(s) of processing dopamine acts to bring about these changes in performance. Application of the mountain model can identify these stages, thus helping to resolve a problem that has been prominent in the literature on brain-reward circuitry since the 1970s.

Over the past decade, a new subdicipline, neuroeconomics, has begun to emerge. The quantitative model at the core of this thesis fuses some functional ideas from behavioural economics (e.g., price sensitivity, payoff) with mechanistic concepts from behavioural neuroscience (e.g., temporal summation, rate codes). It is encouraging that the mountain model was largely successful in describing the relationship between behaviour allocation and the cost and strength of reward and in predicting the effect of varying train duration. For more than 50 years, students of brain-reward circuitry have wrestled with problems such as identifying the primary neurons that subserve BSR,

determining the role of dopaminergic neurons, and understanding how changes in physiological state modulate BSR. If application of the mountain model were to help solve problems of this kind, it would promote quantitative modelling in general and neuroeconomic modelling in particular as key tools for the study of the psychological and neural processes responsible for reward and decision making.

References

Arvanitogiannis, A. (1997). *Mapping the substrate for brain stimulation reward: New approaches to an old problem.* Unpublished PhD, Concordia University, Montreal.

Breton, Y. A. (2004). *Using the reference-sweep method to infer the subjective value of rewarding brain stimulation in the rat.* Unpublished Undergraduate Thesis, Concordia University, Montreal.

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

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

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

Conover, K., & Shizgal, P. (2002, May 24-25, 2002). *Telling work from leisure time: quantitative analysis of the temporal distributions of behavior for the validation of a labor supply model of single operant choice.* Paper presented at the Society for Quantitative Analyses of Behavior, Toronto, Ontario.

Conover, K., & Shizgal, P. (2004). Employing labor supply theory to measure the reward value of electrical brain stimulation. *Games and Economic Behavior, In press.*

Edmonds, D. E., & Gallistel, C. R. (1974). Parametric analysis of brain stimulation reward in the rat: III. Effect of performance variables on the reward summation function. *Journal of Comparative and Physiological Psychology, 87*, 876-883.

Frank, R. A., Preshaw, R. L., & Stutz, R. M. (1982). The effect of train duration of rewarding stimulation on food self-deprivation. *Physiological Psychology, 10*(1), 158-162.

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

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

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

Gallistel, C. R. (1994). Foraging for brain stimulation: Toward a neurobiology of computation. *Cognition, 50*(1-3), 151-170.

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

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

Hoaglin, D. C. (1983). Understanding Robust and Exploratory Data Analysis. In D. C. Hoaglin, F. Mosteller & J. W. Tukey (Eds.), *Understanding Robust and Exploratory Data Analysis* (pp. 415-425). New York: Wiley.

Hoebel, B. G. (1969). feeding and self-stimulation. *Annals of the New York Academy of Sciences, 157*, 758-778.

Miliaressis, E., Rompre, P. P., Laviolette, P., Philippe, L., & Coulombe, D. (1986). The curve-shift paradigm in self-stimulation. *Physiology and Behavior, 37*, 85-91.

Olds, J. (1958). Self-stimulation of the brain. *Science, 127*, 315-324.

Rachlin, H., Kagel, J. H., & Battalio, R. C. (1980). Substitutability in time allocation. *Psychological Review, 87*(4), 355-374.

Shizgal, P., & Murray, B. (1989). Neuronal basis of intracranial self-stimulation. In J. M. Liebman & S. J. Cooper (Eds.), *The Neuropharmacological Basis of Reward* (pp. 106-163). Oxford: Claredon Press.

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

Shizgal, P. (1999). On the neural computation of utility: implications from studies of brain stimulation reward. In D. Kahneman, E. Diener & N. Schwarz (Eds.), *Well-Being: The Foundations of Hedonic Psychology*. New York: Russell Sage Foundation.

Shizgal, P. (September 15, 2004). Proposal submitted to the Canadian Institute of Health Research (CIHR).

Simmons, J. M., & Gallistel, C. R. (1994). Saturation of subjective reward magnitude as a function of current and pulse frequency. *Behavioral Neuroscience, 108*(1), 151-160.

Solomon, R. (2004). *Interpreting Subjective and Objective Price of Brain Stimulation Reward.* Unpublished Undergraduate Thesis, Concordia University, Montreal.

Appendix

The following procedures were implemented in order to calculate the bi-square weights. The distance between the median time allocation value in each condition (m_j) and each data point, d_i (at condition $_j$) was determined with Equation 1.

$$d_{i,j} = \text{abs}\left(m_j - x_{i,j}\right), \forall x_{i,j} \in condition\left(j\right) \tag{1}$$

The median absolute deviation (MAD) was then determined for all conditions within each set of sweeps, as described in Equation 2.

$$MAD = \text{median}\left(d_{i,j}\right), \forall d_{i,j} \tag{2}$$

The scaling factor for the deviations, u, is then calculated by multiplying MAD by the tuning constant, k (in our experiment a tuning constant of 6 was used, a commonly employed value for this constant (Hoaglin, 1983).

$$u = k\,MAD \tag{3}$$

The bi-square equation (Equation 4) determines the weight of each data point. In the case where the data value is less than or equal to the scaling factor a weight is determined. In the case where d is greater than the scaling factor, a weight of zero is given.

$$w_{i,j} = \begin{pmatrix} \forall d_{i,j} \le u : \left(1 - \left(\frac{d_{i,j}}{u}\right)^2\right)^2 \\ \forall d_{i,j} > u : 0 \end{pmatrix} \tag{4}$$

The weighting function assigns weights close to 1 for points lying near the median; as the distance from the median grows, the weights fall off smoothly and eventually reach zero. Thus, the influence of outliers is eliminated or reduced without having to identify the outliers *a priori*. The tuning constant determines how far an outlier must deviate from the median in order to be assigned a weight of zero. Given that MAD is based on all the data for a given set of sweeps, points from conditions yielding low-variance data (e.g., points near the upper or lower asymptotes of the 2D psychometric functions) receive higher average weight than points from conditions yielding high-

variance data (e.g., points on the steep portion of the 2D psychometric functions). This differential weighting tends to compensate for the heteroskedasticity in the data. The effect of skew is reduced by the fact that the weighting procedure draws the fitted surface toward the median of the dataset for each condition.

Figure A1 illustrates the need for bi-square weighting. Results from a single price sweep (top panel) and a single frequency sweep (lower panel) are shown. Outliers (red) were observed at two of the prices tested in the session depicted in the upper panel (data points circled in yellow). If conventional means had been calculated, they would have been strongly biased by the outliers. However, the bi-square weighting procedure assigns these points weights of zero, thus removing the contaminating influence.

Heteroskedasticity is illustrated in the lower panel of Figure A1. The largest errors bars are associated with the rising or falling portions of the sigmoidal curve, and the smallest error bars are consistently seen at the upper asymptote of the data set. The bi-square weighting procedure compensates for the heteroskedasticity of the data by down-weighting the extreme values. This is illustrated in the lower panel of Figure A1 by colour-coding the data points on the basis of their associated bi-square weights. As shown in the weighting scale provided to the right of the figure, blue points were given a full weight of 1 whereas the red points were assigned weights of 0.1 or less.

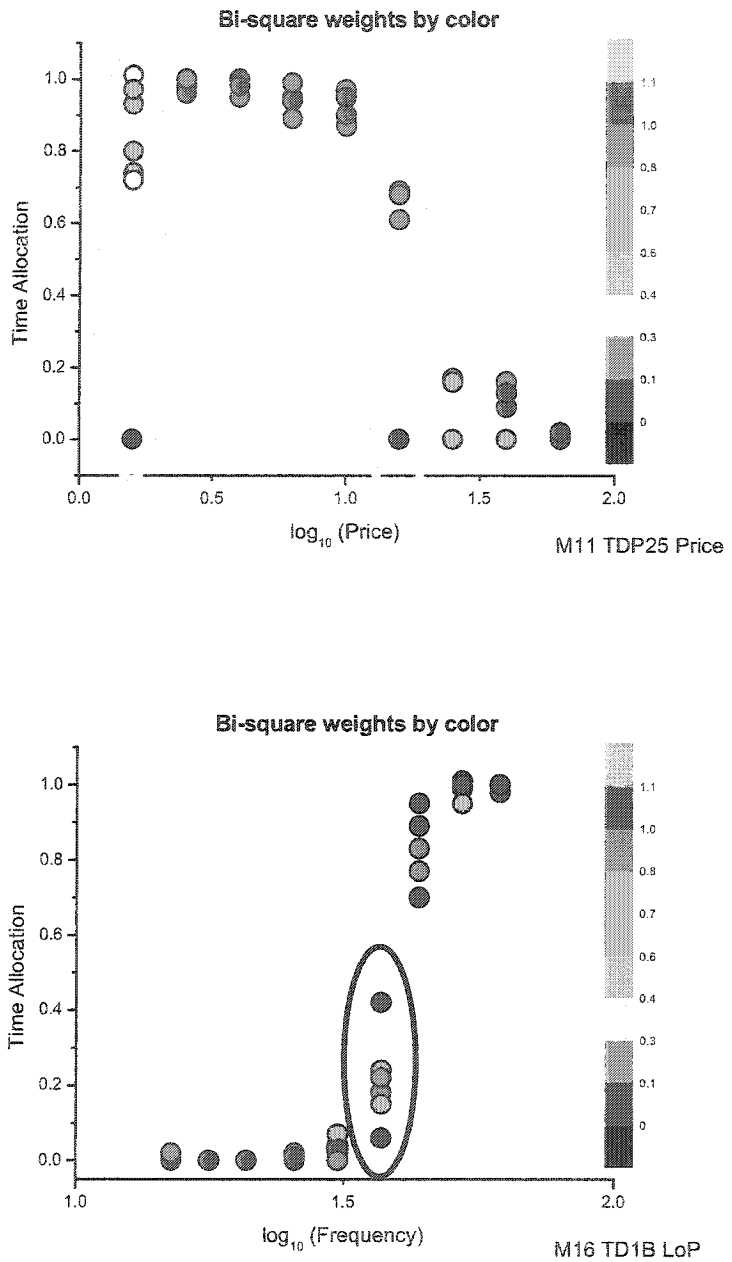

Figure A1. Illustration of the need for bi-square weighting. The cloud of points in the upper panel illustrates the reason for weighting the data set. If each of these points were given equal weight then we would know very little about the position of the mean of this group on the slope of the frequency function. The lower panel illustrates the extent of the heteroskedasticity of the data.