

Modulation of reward-seeking by changes in energy balance: a 3D perspective

Sarah Nolan-Poupart

A Thesis
in
The Department
Of
Psychology

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

August, 2016

© Sarah Nolan-Poupart, 2016

CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the thesis prepared

By: Sarah Nolan-Poupart

Entitled: Modulation of reward-seeking by changes in energy balance: a 3D
perspective

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

Master of Arts (Psychology)

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

Signed by the final Examining Committee:

Dr. Andrew Chapman	Chair
Dr. Andreas Arvanitogiannis	Examiner
Dr. Uri Shalev	Examiner
Dr. Peter Shizgal	Supervisor

Approved by _____

Chair of Department or Graduate Program Director

Dean of Faculty

Date August 29, 2016

ABSTRACT

Modulation of reward seeking by changes in energy balance: a 3D perspective

Sarah Nolan-Poupart

Concordia University, 2016

Peripheral signals of energy balance modulate central reward processing and reward-seeking. Prior work has shown that negative energy balance potentiates the reward effectiveness of electrical stimulation evoked from the lateral hypothalamus (LH) when the electrode is located in the perifornical region of the LH but not neighbouring LH sites. Recent work has revealed however that the measurement methods by which these findings were obtained are relatively ambiguous and insensitive. Here we employed a new three-dimensional (3D) method, called the reward-mountain paradigm, which offers additional information on reward processing by measuring operant behaviour as a function of both the strength and cost of the reward. Ten male Long-Evans rats were trained to hold down a lever for electrical stimulation of the LH. Subjects performed reward-mountain sessions during four phases designed to manipulate long-term energy balance: a baseline phase during which subjects were fed ad libitum; a chronic food restriction phase that lasted until subjects reached 75% of their baseline body weight; a stable restriction phase during which subjects were maintained at the target weight; a recovery phase during which subjects were returned to an ad lib diet. During stable restriction, short-term energy balance was varied by feeding rats either before or after test sessions. Both chronic food restriction and meal time yielded mixed effects, supporting the conclusion drawn by prior studies that there exist functionally heterogeneous reward substrates in the LH which are differentially modulated by signals of energy availability.

Acknowledgements

I would first like to express my sincerest gratitude to Dr. Shizgal for his guidance, encouragement, and support. His deep and careful thinking, his understanding, and his do-it-yourself attitude have been a true inspiration to me. I feel very fortunate to have been given the opportunity to learn from him. I also thank my committee members, Dr. Andreas Arvanitogiannis and Dr. Uri Shalev, for their helpful comments and for taking the time to review my thesis. I am also very thankful to Kent Conover, Dave Munro, and Steve Cabilio for all the help they have given me over the years, without which this project would not have been possible, and for their patience. I also wish to give my sincerest thanks my fellow lab members for the tremendous amount of help and inspiration they have provided over the years: Marie-Pierre, Cossette, Rebecca Solomon, Brian Dunn, Ivan Trujillo-Pisanty, and Czarina Evangelista. I also thank my friends and family for their love and support. Finally, I would like to acknowledge the financial support provided by the Natural Sciences and Engineering Research Council of Canada.

Contribution of Authors

Sarah Nolan-Poupart: Conducted the experiment, analyzed the training data, suggested adding low price frequency sweeps, conducted the 2D analysis, and wrote the thesis.

Kent Conover: Programmed all MATLAB procedures, created the model permutations that were used for data analysis, ran the 3D analyses, and provided invaluable advice throughout the course of the study.

Peter Shizgal: Designed the experiment and created the reward-mountain models upon which the experiment is based. He supervised the execution of the entire study, from experiment design to behavioural training, data collection and analysis, providing invaluable guidance and suggestions throughout. He revised the present thesis.

Table of Contents

List of Figures.....	viii
List of Tables.....	ix
List of Abbreviations.....	x
Introduction.....	1
Materials and Method.....	6
Subjects.....	6
Surgical Procedure.....	6
Apparatus and Stimulation.....	7
Reward-Mountain Training.....	7
Screening.....	7
Task Details.....	8
Reinforcement schedule.....	8
Trial.....	8
Intertrial Interval.....	9
Sweep training.....	9
Reward-Mountain Testing.....	11
Statistical Analysis.....	12
Raw Data.....	12
Effect of Chronic FR: 3D Dynamic Analysis	13
Reward-Mountain Models	14
Model Fitting	17
Akaike Information Criterion	20

Effect of Chronic FR: 2D Analysis	20
Effect of Pre-Feeding: Pre-vs. Post-Fed Analysis.....	21
Results.....	22
Weight Loss.....	22
Effect of Chronic Food Restriction: 3D Dynamic Analysis	24
Effect of Chronic Food Restriction: 2D Analysis	31
Effect of Pre-Feeding: Pre-vs. Post-Fed Analysis.....	34
Electrode Placemens.....	39
Discussion.....	41
Effect of Chronic FR.....	41
Comparison of the 2D and 3D Analyses.....	47
Effect of Pre-Feeding versus Post-Feeding.....	48
Limitations.....	51
Conclusion.....	51
References.....	53

List of Figures

Figure 1.	The reward-mountain model.	4
Figure 2.	Change in body weight across study days.	23
Figure 3.	3D dynamic analysis of data from an example rat.	25
Figure 4.	3D dynamic analysis: maximal shifts in location parameters.	27
Figure 5.	Winning 3D dynamic models.	29
Figure 6.	Regression of M-50 values on body weight.	32
Figure 7.	Effect of pre-feeding vs. post-feeding in an example rat.	35
Figure 8.	Shifts in location parameters as a function of pre-feeding.	36
Figure 9.	Electrode placements.	40

List of Tables

Table 1.	3D dynamic analysis: location parameter shifts.	28
Table 2.	3D dynamic analysis: AIC values from an example rat.	30
Table 3.	Regression statistics.	33
Table 4.	Pre-fed vs. post-fed analysis: location parameter shifts.	37
Table 5.	Pre-fed vs. post-fed analysis: AIC values from an example rat.	38

List of Abbreviations

AIC	Akaike Information Criterion
BSR	Brain Stimulation Reward
F_{hm}	Pulse frequency that produces half-maximal reward intensity
FR-1	Fixed-ratio 1
ITI	Inter-Trial Interval
LH	Lateral Hypothalamus
P_e	Price that produces half-maximal time allocation for maximally rewarding stimulation
TA	Time Allocation

1. Introduction

Physiological need states affect motivated behaviour. The relationship between peripheral energy balance and reward-seeking in particular has long been studied in rats performing operant tasks for delivery of electrical stimulation to the lateral hypothalamus (LH). The neural activity provoked by the electrical stimulation causes a rewarding effect, called brain stimulation reward (BSR), which animals will vigorously work to obtain and which is thought to mimic components of the neural processing of natural rewards (Conover & Shizgal, 1994; Conover, Woodside & Shizgal, 1994).

It has repeatedly been observed that depleting long-term energy stores by means of chronic food restriction increases the reward effectiveness of BSR in a subset of cases (Abrahamsen, Berman & Carr, 1995; Abrahamsen & Carr, 1996; Fulton, Richard, Woodside, & Shizgal, 2002; Fulton, Richard, Woodside, & Shizgal, 2004). This effect of chronic food restriction typically occurs when the electrode that delivers electrical stimulation is located in the dorsal or dorsolateral perifornical region of the LH, but not when it is placed outside of that region (Cabeza de Vaca, Holiman & Carr, 1998; Carr & Wolinsky, 1993; Carr, Kim, & Cabeza de Vaca, 2000; Fulton, Woodside & Shizgal, 2000; Fulton, Woodside & Shizgal, 2002; Fulton, Woodside & Shizgal, 2006). In contrast, short-term energy challenges, such as 48-hour food deprivation, have typically not been found to affect the reward effectiveness of BSR at any stimulation site (Cabeza de Vaca et al., 1998; Fulton et al., 2000; Fulton et al., 2004; Fulton et al., 2002; but see Abrahamsen et al., 1995).

Such findings have led to the proposal that the anatomical substrate that supports BSR comprises at least two discrete neural subpopulations –a chronic food restriction-sensitive and a food restriction-insensitive population – which process functionally distinct, reward-related

information (Fulton et al., 2006; Shizgal, Fulton & Woodside, 2001). According to this interpretation, the information processed by food restriction-sensitive sites relates to long-term peripheral energy stores, but not to short-term energy balance or to hunger (Fulton, 2010).

That this interpretation may not be the last word on the topic is suggested by recent studies showing that the methods employed in the original work are relatively insensitive and ambiguous (Arvanitogiannis & Shizgal, 2008; Hernandez, Breton, Conover, & Shizgal, 2010). The studies that reported the chronic food restriction effect primarily made use of the "curve-shift" paradigm, a two-dimensional (2D) measurement strategy which measures the vigour with which subjects perform an operant task for BSR as a function of a single independent variable (the strength of the electrical stimulation). This method quantifies the effect of a manipulation such as chronic food restriction based on the change in the value of the independent variable required to support half-maximal performance.

Operant behaviour is determined by multiple dimensions of reward computed at separate neural stages (Breton, Mullett, Conover, & Shizgal, 2013; Hernandez et al., 2010). Reward dimensions include the subjective evaluation of the strength of the reward, the opportunity and effort costs associated with working for the reward, the probability of obtaining the reward, and the value of competing activities (Shizgal, 1997). By measuring operant behaviour as a function of a single reward component, 2D strategies provide only an incomplete account of reward-seeking. Due to this inherent methodological limitation, previous studies have not been capable of distinguishing which reward component(s) is (are) affected by long-term food restriction in food restriction-sensitive cases. Changes in reward effectiveness could result from changes to any or all of the various dimensions of reward-seeking. Moreover, it has been shown that in cases where a manipulation simultaneously causes opposite effects on multiple reward

components, the half-maximal performance criterion used by 2D measurements to track effects on BSR may fail to change. This raises the possibility that BSR is also affected by chronic food restriction in food restriction-insensitive cases and by short-term energy challenges.

Recently, a 3D measurement strategy called the “reward-mountain method” has been introduced which provides information about the stages of neural processing of reward (Arvanitogiannis & Shizgal, 2008; Breton, Conover, & Shizgal, 2014; Breton et al., 2013; Hernandez et al., 2010). Previous studies have successfully implemented this method to disambiguate the effects of various pharmacological challenges on the different reward components of BSR (Hernandez, Trujillo-Pisanty, Cossette, Conover, & Shizgal, 2012; Trujillo-Pisanty et al., 2011; Trujillo-Pisanty, Conover, & Shizgal, 2014).

The reward-mountain method is based on a theoretical model that describes how the neural activity induced by electrical stimulation is translated into reward-seeking behaviour (Gallistel, Shizgal, & Yeomans, 1981; see Figure 1). First, electrical pulses delivered through an electrode tip cause nearby cell bodies and fibers of passage to fire. This volley of action potentials undergoes spatiotemporal integration, yielding the subjective intensity of the electrical reward. The neural signal for the subjective reward intensity is combined with the probability of the reward and scaled by subjective measures of opportunity and effort costs. This combination results in a representation of the total payoff of BSR (Shizgal, 1997). Finally, the model predicts that rats allocate their time between pursuit of BSR and competing activities, such as grooming and resting, based on their relative payoffs.

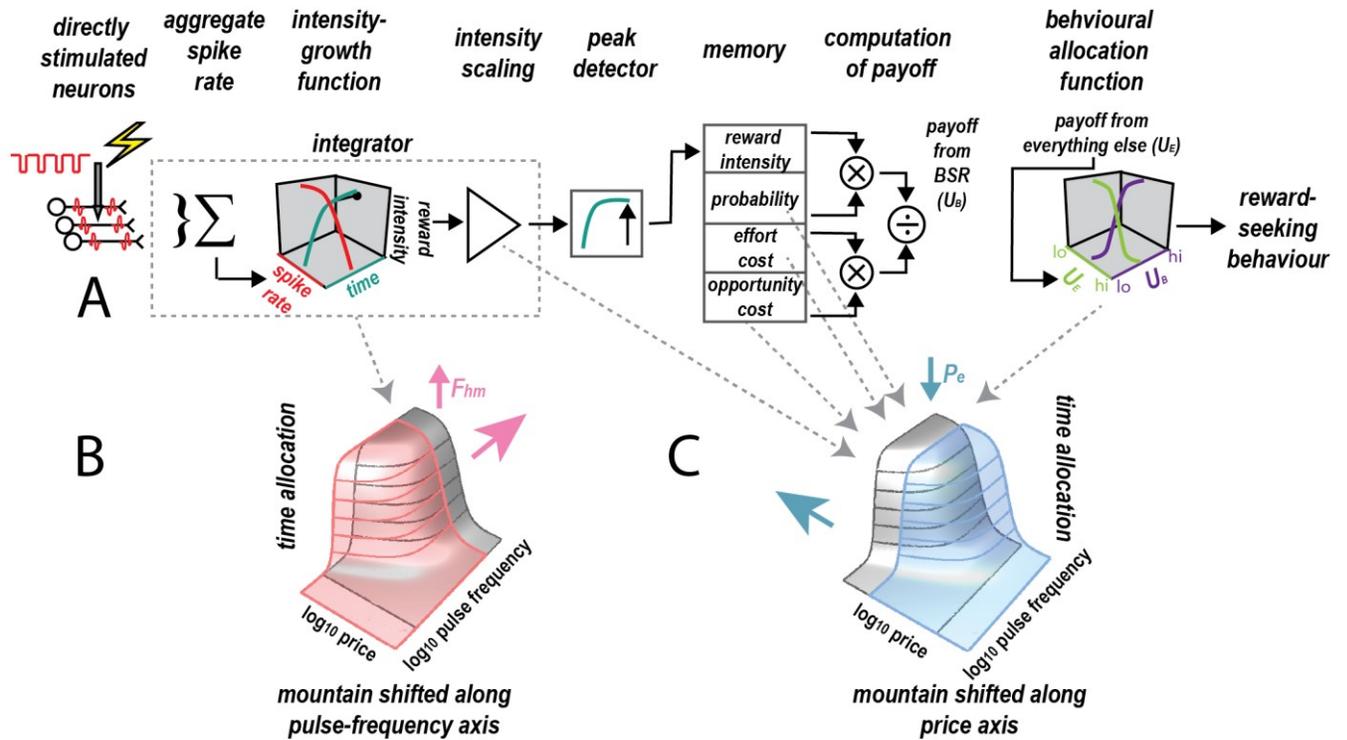


Figure 1. The reward-mountain model. (A) Electrical pulses cause the directly-stimulated neurons to fire. The stimulation-induced firing rate is spatially and temporally integrated and transformed into the subjective reward intensity of the stimulation. The peak reward intensity is recorded in memory. Subjective opportunity cost is estimated based on the objective price of the stimulation. The peak reward intensity of the stimulation is combined with the probability of obtaining the reward, its subjective opportunity and effort costs; this scalar combination results in the total payoff of the electrical stimulation. Time allocation (TA) depends on the relative payoffs of the electrical stimulation and of competing activities such as grooming and resting. (B) Manipulations affecting the circuitry prior to the output of the integrator result in changes in the pulse-frequency that supports half-maximal reward intensity and in the location of the reward-mountain along the pulse-frequency axis. (C) Changes at or beyond the output of the integrator result in shifts of the reward-mountain along the opportunity cost axis. (Adapted from Breton et al., 2013; Hernandez et al., 2010; Trujillo-Pisanty et al., 2014).

The dependent measure in the reward-mountain testing procedure is “time allocation” (TA). The procedure measures TA based on the proportion of time rats spend holding a lever for electrical stimulation as a function of two independent variables, strength and opportunity cost. Reward strength is controlled by varying pulse frequencies, whereas opportunity cost is controlled by varying the cumulative number of seconds during which the rat must hold down the lever in order to obtain the reward (called the “price”). Fitting the reward-mountain model to the behavioural data generates a 3D structure (referred to as the “reward mountain”) that is positioned in a space defined by one axis representing reward strength and one axis representing opportunity cost (or price). Two location parameters, F_{hm} and P_e , describe the mountain’s position along the pulse-frequency and price axes, respectively. Crucially, the model predicts that if a manipulation affects reward processing prior to the output of the circuitry subserving the spatiotemporal integration, the 3D mountain structure will shift along the pulse-frequency axis (measured as a change in F_{hm}). If a manipulation affects processing at a later stage, such as the rat’s evaluation of the subjective cost of working for the reward, the mountain will shift along the price axis (measured as a change in P_e).

In the present study, we had three objectives. First, we aimed to determine which stage(s) in the neural processing of BSR is affected by chronic food restriction in food restriction-sensitive cases. Our second aim was to verify whether use of the 3D measurement method would reveal effects of food restriction at sites previously labeled as food restriction-insensitive and effects of change in short-term energy balance. Our third aim was to examine whether the 3D method would also reveal effects of short-term energy balance. To address these questions, we implanted 10 rats with electrodes aimed at the perifornical region of the LH and tested them using the reward-mountain procedure under different long-term energy balance states: while fed ad

libitum, under chronic food restriction until they lost 25% of their body weight at baseline, and while regaining body weight. We also tested subjects while maintained at 75% of their original body weight under two short-term energy balance states: a hungry and a sated state.

2. Materials and method

2.1. Subjects

Ten male Long-Evans rats (Charles-River, St. Constant, QC, Canada) weighing 300-350g at arrival were maintained on a 12 hour light/dark reverse cycle (lights “off” at 8AM; lights “on” at 8PM). Behavioural procedures were performed during the dark cycle. Prior to surgery, subjects were paired-housed in Plexiglas shoebox cages with unrestricted access to food. Following surgery, they were housed individually and access to food was contingent upon the experimental phase. All procedures were performed in accordance with the principles outlined by the Canadian Council on Animal Care.

2.2. Surgical Procedure

Prior to surgery, monopolar electrodes were assembled by soldering a stainless steel insect pin (0.25 mm diameter), insulated with Formvar to within 0.5 mm of the tip, to one end of a copper wire that was attached to a gold-plated connector. A return electrode was fashioned from a copper wire attached to a gold-plated connector. Surgeries were performed once subjects weighed at least 450g to ensure that they had accumulated a substantial adipose mass before the start of the experiment. Anesthesia was first induced with a mixture of ketamine-xylazine (10/100 mg/kg, i.p.). Atropine sulfate (0.05 mg/kg, s.c.) and penicillin (0.3 ml/kg, s.c.) were administered to inhibit bronchial secretions and prevent infection, respectively. The rat’s head was then immobilized in a stereotaxic frame and general anesthesia was maintained by administering isoflurane (3%) through a nose cone. An incision was made in the scalp to expose

the skull, and six stainless steel jeweller's screws were driven into the skull through pilot holes. Two monopolar electrodes were aimed bilaterally at the perifornical region of the lateral hypothalamus (AP: - 3.0 or 3.12, ML: \pm 1.3 or 1.4, DV: -8.7 or 8.8). The copper wire portion of the return electrode was coiled around two jeweller's screws. All gold-plated connectors were then inserted into an externally-threaded, nine-pin connector (Scientific Technology Center, Carleton University, Ottawa, ON, Canada), and the head cap was secured with dental acrylic. Buprenorphine (0.05 mg/kg, s.c.) was administered to reduce postoperative pain.

2.3. Apparatus and Stimulation

Testing took place in an operant conditioning chamber (30 cm X 21 cm X 51 cm) with four Plexiglas walls and a hinged Plexiglas front door. The chamber was equipped with a retractable lever (1.5 cm X 5 cm) (ENV-112B, MED Associates, St. Albans, Vermont) located at the centre of one wall, a cue light (1.5 cm in diameter) located 2 cm above the lever, and an amber flashing light (5 cm X 10 cm) located 10 cm above the wire-mesh floor on the back wall.

Electrical stimulation was delivered through a lead cable attached to an electrical rotary joint located on the ceiling. Electrical stimulation consisted of 0.5 s trains of 0.1 ms-long rectangular, cathodal, constant-current pulses. A constant-current amplifier and digital pulse generator, controlled by a computer program written by Stephen Cabilio (Montreal, QC, Canada), were used to adjust the current amplitude and pulse frequency of the stimulation.

2.4. Reward-Mountain Training

2.4.1. Screening

Following a 7-day post-surgery recovery period, rats were shaped to lever press for electrical stimulation on a fixed ratio-1 reinforcement schedule. The stimulating electrode, current amplitude, and pulse frequency that supported the most vigorous lever-pressing in the

absence of motoric or aversive side effects were determined for each rat individually. The most effective electrode and current amplitude were used during the remainder of the experiment.

2.4.2. Task details

2.4.2.1. Reinforcement schedule

Rats were trained on a fixed, “cumulative handling-time” schedule of reinforcement (Breton et al., 2009). In accordance with this schedule, subjects learned to hold down the lever for a fixed, cumulative number of seconds (the “price”) in order to obtain a single stimulation train (the “reward”). The price established the opportunity cost of the reward: whilst working for BSR, rats had to simultaneously forgo the opportunity to engage in competing activities, such as exploring, grooming and resting.

2.4.2.2. Trial

Each trial consisted of an experimenter-controlled duration during which both the strength of the reward and the price were kept constant. Reward strength was set by the frequency with which current pulses were delivered during the 0.5 s stimulation train, with higher pulse frequencies resulting in greater reward strength.

The cue light was illuminated while subjects held down the lever. Once the cumulative hold-down time criterion was reached, the lever was retracted from the chamber, and a reward was delivered. The lever remained retracted during a 2 s period, called the “black-out delay,” before being extended back into the chamber.

Subjects could earn a maximum of 20 rewards during every trial. The duration of each trial was determined based on the total number of available rewards (20) multiplied by the price of the reward. The black-out delay was excluded from the total trial duration.

2.4.2.3. Intertrial interval

The intertrial interval (ITI) consisted of a 10 s-long period of time during which the lever was retracted from the chamber and a non-contingent (“priming”) stimulation train was delivered. For each individual rat, the pulse frequency of the priming stimulation was set to the maximal value that did not result in motoric or aversive side-effects. ITIs were signalled by the flashing amber light.

2.4.3. Sweep training

During every trial, rats allocated their time between the pursuit of BSR and alternative activities on the basis of their relative payoffs. The dependent variable, “time allocation” (TA), the proportion of time spent working for BSR during a trial, measured rats’ choices between these relative payoffs. The payoff obtainable from BSR depended on the two experimenter-controlled variables: opportunity cost (price) and reward strength (pulse frequency). High pulse frequencies and low prices are known to result in high payoff from BSR and high TA measures, whereas low pulse frequencies and high prices result in low payoff from BSR and low TA measures.

After being trained to hold down the lever for a 4 s price, subjects were gradually trained to perform four types of trial sequences called “sweeps”. Each sweep consisted of a series of 11 trials separated by ITIs. The sweeps were designed such that TA measures progressed from maximal to minimal values over the course of a single sweep. To accomplish this, one or both experimenter-controlled variables were varied in equal proportional steps across each subsequent trial in the sweep.

Rats were first trained to perform “frequency sweeps”. During the first three trials of each frequency sweep, the pulse frequency of the reward was set to the value that supported the most

vigorous lever-pressing without motoric or aversive side-effects, while the price was set to a relatively low value (3 or 4 s, depending on the subject). Over the course of the remaining 8 trials, pulse frequency was decreased in equal proportional steps while price was kept constant at 4 s. Rats performed a minimum of 4 frequency sweeps per daily training session. Following each session, TA measures were plotted against pulse frequency for each rat individually. The starting pulse frequency and logarithmic step size separating each trial's pulse frequency were modified until TA measures plotted across pulse frequency formed a sigmoidal curve with well-defined upper and lower asymptotes.

Once a rat's performance was stabilized, "price sweeps" were added to the daily training sessions. During the first three trials of a price sweep, the pulse frequency was set to its maximal value, while the price was set to 3 or 4 s (depending on the subject). Over the course of the remaining trials in the sweep, pulse frequency was kept constant while price was increased in equal proportional steps. After each training session, the TA measures obtained during price sweeps were plotted against price. The logarithmic step size separating the price of each trial was adjusted until a sigmoidal TA curve was achieved.

Next, "radial sweeps" were added to the training sessions. The first three trials of the radial sweep were identical to the first trials of the frequency and price sweeps. Across the remaining trials, pulse frequency was systematically decreased while price was systematically increased. The starting pulse-frequency value and logarithmic step sizes of both experimenter-controlled values were modified as necessary until sigmoidal TA curves were achieved.

Finally, "low-price frequency sweeps" were added to the training sessions of all subjects with the exception of rats B12, B15 and B19. During the first three trials of this sweep type, the pulse frequency was set to its maximal value while the price was set to 1 s. Pulse frequency was

systematically decreased across the last 8 trials while price was kept constant at 1 s. TA measures were examined daily and the pulse-frequency values of each trial were modified as described above in order to achieve sigmoidal TA curves.

At this stage of training, full reward-mountain sessions were in effect. These sessions consisted of two “surveys”, with each survey comprising one frequency sweep, one price sweep, one radial sweep and (with the exception of rats B12, B15 and B19) one low-price frequency sweep. The order of presentation of the sweeps was randomized within, but not between, surveys.

Once a rat had performed three or more reward-mountain sessions, the data were further examined by fitting the standard version of the reward-mountain model (see Statistical Analysis section for a detailed explanation of the fitting procedure). The aim was to obtain a good fit between the model and the data, in which case the confidence-interval bands around the location parameters were narrow, the radial sweep traversed the intersection of the two location parameters, and the price sweep crossed a vertical portion of the contour lines that defined the reward mountain. The experimenter-controlled parameters (the logarithmic step sizes separating pulse-frequency and price values across trials; the pulse-frequency and price values of the first trial of each sweep) were modified using a simulator developed by Yannick Breton (MATLAB, MathWorks, Natick, MA) as often as needed to optimize the fit. Once good fits were obtained, the reward-mountain testing phase of the experiment began.

2.4.4. Reward-mountain testing

During testing, subjects performed reward-mountain sessions over four consecutive phases designed to manipulate energy balance. During each phase, rats were tested every day or

every two days. Each session consisted of a warm-up frequency sweep followed by two surveys. Rats were weighed immediately after every testing session.

During the “baseline” phase, rats were given free access to food in their home cage and performed a total of 8 daily testing sessions. Subjects then entered the “chronic food restriction” phase during which access to food was restricted to 15 g a day so as to cause a steady decrease in body weight. Rats were given their food portion in their home cage at the same time every day; on testing days, this occurred after the testing session. Once rats reached 75% of their mean weight at baseline, they entered the “stable restriction” phase during which their weight was maintained at the 75% target value by titrating the daily food portion according to weight. During this phase, rats were fed either 45 min before (“pre-fed”) or immediately after (“post-fed”) their testing session. If rats had not consumed the full meal within the 45 min period preceding the testing session, they were kept in their home cage until having eaten approximately 95% of the full portion. This procedure was adapted from Abrahamsen, Berman and Carr (1995). Subjects performed 10 interdigitated pre-fed and post-fed sessions. During the final “recovery” phase, rats were once again given free access to food and tested as they regained weight.

After completion of the last behavioural testing phase, subjects were perfused and their brains were sliced to determine electrode placement.

2.5. Statistical Analysis

2.5.1. Raw data

The dependent variable TA was calculated based on the lever holds and releases recorded during each trial. The first two trials of every sweep functioned as warm-ups and were therefore excluded from all analyses. Time spent holding down the lever and releases that were under 1 s-long were classified as “work” time (i.e. time spent working for BSR). Releases under 1 s-long

were also classified as work time as they were considered too brief for subjects to have engaged in alternative activities. Conversely, “leisure” time consisted of releases that were over 1 s long (Breton, Marcus, & Shizgal, 2009). Leisure time was assumed to reflect the time spent engaging in competing activities, such as grooming, exploring, or resting. The mixture of holds and releases (work and leisure) that preceded the delivery of a reward defined each “reward encounter”.

2.5.2. Effect of chronic food restriction: 3D dynamic analysis

We wished to assess whether changes in long-term energy balance and body weight shift the 3D reward mountain along the pulse frequency and/or price axes. To do so, we developed the “dynamic” analysis. This analysis was based on an approach to model fitting that entailed independently fitting multiple permutations of the reward-mountain model to every individual subject’s dataset (Anderson, 2008). Each model permutation represented a reasonable hypothesis regarding the possible relationship between long-term energy balance and the 3D reward mountain. The Akaike information criterion (AIC) was then computed for each fit and used to compare the likelihood of each model’s hypothesis (Akaike, 1974).

We fit a total of 16 model permutations to the data. Four model permutations tested whether the location parameters varied with long-term energy balance. Each of these model permutations corresponded to one of four possible effects of the long-term energy balance manipulation on the 3D reward-mountain: according to one model permutation, the energy manipulation results in a change in both location parameters; according to a second model permutation, F_{hm} alone is affected by the energy manipulation; in a third model, P_e alone is affected; finally, according to a fourth model, the energy manipulation does not affect either location parameter. We fit each of these model permutations twice, once using the “standard”

reward-mountain model and once using the “conditioned-reward” model, resulting in 8 model permutations (see the following section for a description of the standard and conditioned-reward models).

Finally, we reasoned that shifts in the location of the 3D reward mountain could occur in one of two ways: a location parameter shift could occur gradually over the course of an experimental phase as a result of the gradual change in body weight; conversely, a location parameter shift could occur suddenly after the beginning of a new experimental phase, as a result of the sudden change in the physiological need state of the subject. The “slope” and “step” model permutations were designed to test each of these possibilities, respectively. Both of these models were fit once for each of the 8 aforementioned model permutations.

2.5.2.1. Reward-mountain models. Two versions of the reward-mountain model, the standard and the conditioned-reward models (Breton et al., 2013; Hernandez et al., 2010; Hernandez et al., 2012; Trujillo-Pisanty et al., 2011), were employed.

The standard reward-mountain model is defined by the following equation (see Figure 1 for a graphical representation):

$$TA = TA_{min} + \left[(TA_{max} - TA_{min}) \times \frac{\left(\frac{FFg}{FFg + FF_{hm}^g} \right)^a}{\left(\frac{FFg}{FFg + FF_{hm}^g} \right)^a + \left(\frac{SP}{SP_e} \right)^a} \right] \quad (1)$$

Where:

a = the price-sensitivity constant;

g = the reward-growth constant;

F = the pulse frequency;

FF = the stimulation-induced firing frequency of the directly-stimulated neurons;

FF_{hm} = the stimulation-induced firing frequency that produces a rewarding effect

of half-maximal intensity;

SP = the subjective price of the stimulation train;

SP_e = the subjective price at which TA for a maximally intense reward falls halfway between TA_{min} and TA_{max} ;

TA_{min} = minimum TA;

TA_{max} = maximum TA;

The firing frequency (FF) is obtained from pulse frequency with the following equation:

$$FF = F_{bend} \times \left[\text{Ln} \left(1 + e^{\frac{F_{NearMax}}{F_{bend}}} \right) - \left(1 + e^{\frac{F_{NearMax}-F}{F_{bend}}} \right) \right] \quad (2)$$

Where:

F_{bend} = the parameter governing the abruptness of the transition between the rising and flat segments of the function;

$F_{NearMax}$ = the midpoint of the transitional region;

Subjective price (SP) is converted from objective price using the following equation:

$$SP = SP_{bend} + \text{Ln} \left(1 + e^{\frac{P-SP_{min}}{SP_{bend}}} \right) \quad (3)$$

Where:

SP_{min} = minimum subjective price;

SP_{bend} = parameter determining the abruptness of the transition between SP_{min} and the rising portion of the subjective price function

The conditioned-reward reward-mountain model incorporates one additional parameter to adjust for systematically higher than expected TA measures during the lower pulse-frequency trials. This model version accounts for the possibility that such high TA measures are a reflection of conditioned reward rather than BSR (Hernandez et al., 2010):

$$TA = TA_{min} + \left[(TA_{max} - TA_{min}) \times \frac{\left(\frac{FF^g + FF_{CR}^g}{FF^g + FF_{CR}^g + FF_{hm}^g} \right)^a}{\left(\frac{FF^g + FF_{CR}^g}{FF^g + FF_{CR}^g + FF_{hm}^g} \right)^a + \left(\frac{SP}{SP_e} \right)^a} \right] \quad (4)$$

Where:

FF_{CR} = the contribution of conditioned reward, expressed as the pulse frequency required to produce an unconditioned reward of equal intensity to the conditioned reward.

When fit to data, the reward-mountain model yields the 3D surface known as the “reward mountain”. The reward-mountain surface is fitted to the TA measures obtained from the 4 sweep types; the latter are plotted separately in three-dimensions (3D) as a function of both independent variables, pulse frequency and price (See Figure 3 for an illustration of the 3D reward-mountain structure.)

Two location parameters, F_{hm} and P_e , describe the position of the 3D reward-mountain along the pulse-frequency and price axes. The location parameter that positions the reward-mountain along the pulse frequency axis, F_{hm} , corresponds to the pulse frequency that induces a half-maximal reward intensity. It is obtained by back-solving Equation 2 (Breton et al., 2013):

$$F_{hm} = F_{NearMax} - F_{bend} \times \ln \left(e^{\frac{F_{NearMax} - F_{hm}(D)}{F_{bend}}} + e^{-\frac{F_{hm}(D)}{F_{bend}}} - 1 \right) \quad (5)$$

Where:

$F_{NearMax}$ = the parameter that positions the function that relates the firing frequency of the directly-stimulated neurons to the pulse frequency of the stimulation; this value is near the maximal firing frequency of the substrate;

F_{bend} = the parameter that determines the abruptness of the bend of the function relating the firing frequency of the directly-stimulated neurons to the pulse frequency of the stimulation;

FF_{hm} = firing frequency of the directly-stimulated neurons that produces a half-maximal reward intensity;

The location parameter that positions the reward-mountain along the price axis, P_e , is equivalent to the price at which time allocation is half-maximal. It is obtained by back-solving the Equation 3 (Breton et al., 2013):

$$P_e = SP_{min} + SP_{bend} \times \text{Ln} \left(e^{\frac{SP_e - SP_{min}}{SP_{bend}}} - 1 \right) \quad (6)$$

Where:

SP_{min} = the minimum subjective price;

SP_{bend} = the parameter determining the abruptness of the transition between SP_{min} and the rising portion of the subjective-price function;

SP_e = the subjective price at which time allocation for a maximal BSR falls halfway between the maximum TA and the minimum TA;

2.5.2.2. *Model fitting.* The model fitting procedure entailed fitting all 16 model permutations to each individual subject's data independently. Data obtained from pre-fed sessions were excluded from the 3D dynamic analysis.

The model fits were performed using a procedure developed by Kent Conover (Montreal, QC, Canada) that combines nonparametric bootstrapping (resampling with replacement) with the

non-linear least-squares routine in MATLAB (MATLAB Optimization Toolbox, The Mathworks, Natick, MA). First, the program generated 100 resampled mountain datasets for every reward-mountain session by randomly resampling with replacement as many “reward encounters” as there were in each original trial.

Common shape (a and g), scale (TA_{max} , TA_{min}) and, in the case of the 7-parameter model, conditioned-reward parameter estimates, were then fit to the pooled resampled datasets from all sessions 100 times, giving a total distribution of 100 estimates of each of these parameters. The value of each parameter was calculated as the mean of its respective estimate distribution; 95% confidence intervals were calculated by excluding the lowest 2.5% and highest 2.5% estimates.

To generate location parameter estimates, the least-squares routine fit a model to each session’s 100 resampled mountains in a series of iterations intended to minimize the squared difference between the model predictions and the data. This resulted in a distribution of 100 location parameter estimates for every session. The mean estimates and confidence interval bands were calculated as described above. This procedure was repeated for each of the 16 model permutations independently.

The four slope-model permutations tested whether location parameters F_{hm} and P_e changed as a linear effect of change in body weight. To test this, the proportional change in body weight from the mean weight at baseline was first calculated in logarithmic units for every session ($\Delta BodyWeight$). The body weight values recorded during all baseline sessions were normalized to zero.

Next, the slope model was fit using the following equations to obtain location parameter estimates as described above:

$$F_{hm} = Base_{F_{hm}} + (F_{hmdv} \times \Delta BodyWeight)$$

Where:

$Base_{F_{hm}}$ = the F_{hm} estimate obtained from the baseline sessions;

F_{hmdv} = the scale parameter that determines the influence of the change in body weight on F_{hm} ;

$\Delta BodyWeight$ = change in body weight (logarithmic units) from mean baseline weight;

$$P_e = Base_{P_e} + (P_{edv} \times \Delta BodyWeight)$$

Where:

$Base_{P_e}$ = the P_e estimate obtained from the baseline sessions;

P_{edv} = the scale parameter that determines the influence of the change in body weight on P_e ;

In all four model permutations, F_{hm} and P_e were kept fixed across baseline sessions (i.e. they equalled $Base_{F_{hm}}$ and $Base_{P_e}$, respectively). Whether or not the location parameters were free to vary with change in body weight over the ensuing sessions depended on the model permutation.

When a location parameter was free to vary, the least-squares routine estimated the optimal value of the associated scale parameter (e.g. F_{hmdv} in the case of the model permutation that allows only F_{hm} to vary). When location parameters were kept fixed, the parameter determining the magnitude of its shift was assigned a value of zero. The F_{hm} shifts were calculated by subtracting the F_{hm} value of the baseline sessions from the F_{hm} value of each subsequent session. The P_e shifts were calculated in the same manner. The shifts described by the best fitting model permutation were plotted across session day.

The same procedure was performed in the case of the step model, with the exception that change in body weight ($\Delta BodyWeight$) was replaced by a logical value.

2.5.2.3. Akaike Information Criterion. The AIC was calculated for every model fit in the dynamic analysis. The AIC penalizes more complex models, requiring that any additional complexity account for substantive information.

The AIC is defined by the following equation:

$$AIC = -2 \ln \text{likelihood} + 2 K$$

Where:

- 2 K (ln (likelihood)) = the probability of the data given a model;

+ 2 K = the number of free parameters in the model;

A more negative AIC implies a better performing model. The model permutation that yielded the smallest AIC was retained.

2.5.3. Effect of chronic food restriction: 2D analysis

Prior studies on the effect of chronic food restriction on BSR relied on 2D methods, primarily the “curve-shift” paradigm. This method measures FR-1 response rates for BSR as a function of a single independent variable, pulse frequency, and assesses the impact of the manipulation by measuring change in “reward effectiveness”: the amount of change in pulse frequency needed to maintain half-maximal performance (the “M-50”). A reduction in the M-50 (which can be visualized as a leftward displacement of the M-50 along the pulse-frequency axis) has been interpreted to imply greater reward effectiveness (Abrahamsen et al., 1995; but see Arvanitogiannis et al. 2008; Hernandez et al., 2010). Using this method, it has been repeatedly observed that in a subset of “food restriction-sensitive” cases, the M-50, and therefore the pulse frequency needed to maintain it, decreases during chronic food restriction, implying an increase

in reward effectiveness. In “food restriction-insensitive” cases, however, no detectable change in M-50 is observed.

We wished to determine whether the subjects of our study would be classified as food restriction-sensitive or -insensitive by a 2D paradigm. To do so, change in reward effectiveness was assessed as a function of change in body weight. This 2D analysis was performed using only data from the low-price frequency sweeps on account of their similarity to the FR-1 data obtained from the 2D curve-shift procedure. The analysis therefore excluded rats B12, B15 and B19. The 2D analysis was performed using the data from the baseline, restriction, stable restriction and recovery phases. Data from pre-fed sessions were excluded from the analysis.

Using the TA data from the raw 2D low-price frequency sweep curves, the pulse frequency that supported half-maximal TA (the “M-50”) was computed by interpolation. Shifts in M-50 values were calculated by subtracting the mean of the M-50 values obtained during the baseline sessions from the M-50 values of each subsequent session. All the M-50 values of the baseline sessions were normalized to zero. The M-50 and the corresponding body weight (both in log units) were plotted across session day and a simple linear regression of the M-50 on body weight was performed.

2.5.4. Effect of pre-feeding: pre-fed versus post-fed analysis

We also wished to assess the effects of pre-feeding on the location of the reward mountain. To test this, we conducted a 3D model analysis using data from only pre-fed and post-fed stable restriction sessions.

The pre-fed vs. post-fed model was fit to the data using the same model fitting procedure as that described for the step model of the dynamic analysis. Data collected on pre-fed session

days were assigned a logical value of 0 and post-fed session data a logical value of 1. This procedure was conducted using both the standard and the conditioned-reward model versions.

The AIC was calculated for every model fit in the pre-fed vs. post-fed analysis. The model permutation that yielded the smallest AIC was retained.

3. Results

3.1. Weight loss

All 10 rats lost weight while undergoing chronic food restriction and reached 25% of their mean baseline body weight within 3-4 weeks on average (Fig. 2). All subjects regained weight after being returned to an ad libitum diet, with the exception of rats B26 and B32, who stopped working for BSR at the beginning of the recovery phase. The change in the behaviour of rats B26 and B32 may have been caused by a change in the location of their electrode.

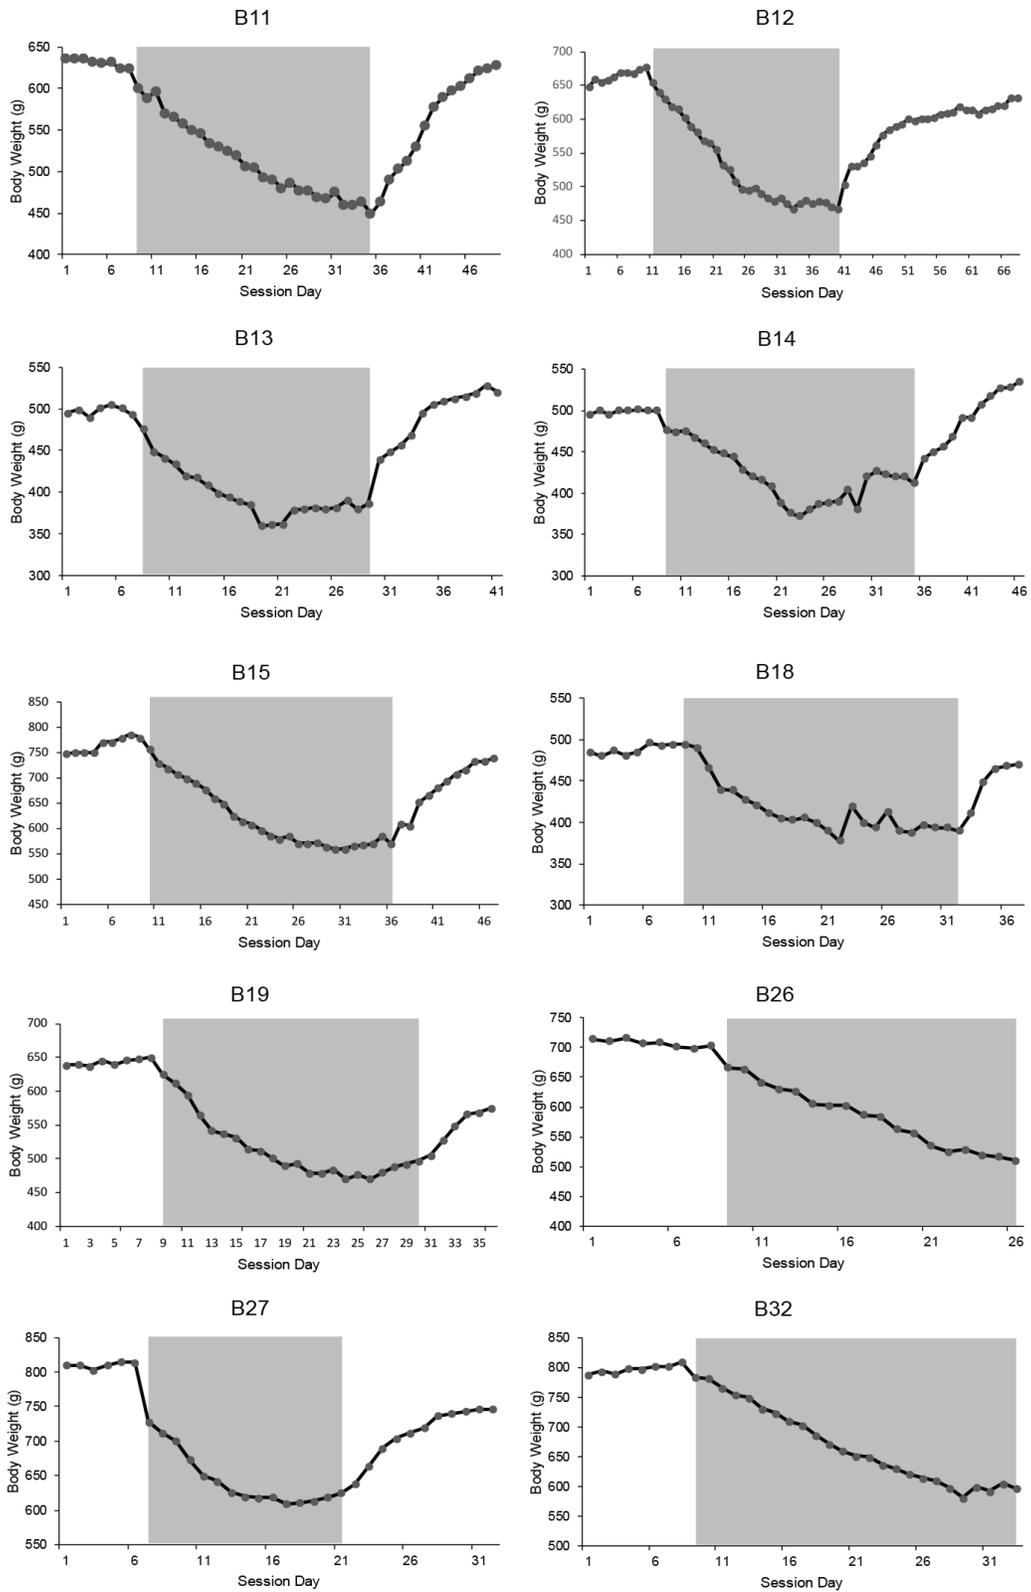


Figure 2. Change in body weight (in grams) across study days. Shaded area represents data points collected on days when the subject was fed a restricted diet.

3.2. Effect of chronic food restriction: 3D dynamic analysis

All subjects learned to perform reward-mountain sessions. Sixteen model permutations were independently fit to the dataset (which excluded pre-fed data) of every individual rat. For each analysis, the winning model permutation was selected based on the AIC.

Figure 3 illustrates the winning model permutation (in this case, the step model) of an example rat (subject B18). The four panels display the cumulative TA data plotted as a function of both pulse frequency and price, and the fitted surface of the winning model in 3D format (Figure 3A) and in a contour plot (Figure 3B). The location parameters F_{hm} and P_e are represented by the horizontal red line and the vertical blue line, respectively. The left panels display the location parameter estimates obtained during the baseline phase, while the right panels display the location parameter estimates obtained when the rat was at 75% of his mean baseline body weight at the end of the restriction phase. Comparison of the left and right panels therefore illustrates the shift in location parameters that occurred between the baseline and restriction phases (specifically, a decrease of 0.138 log units of parameter F_{hm} from baseline to restriction).

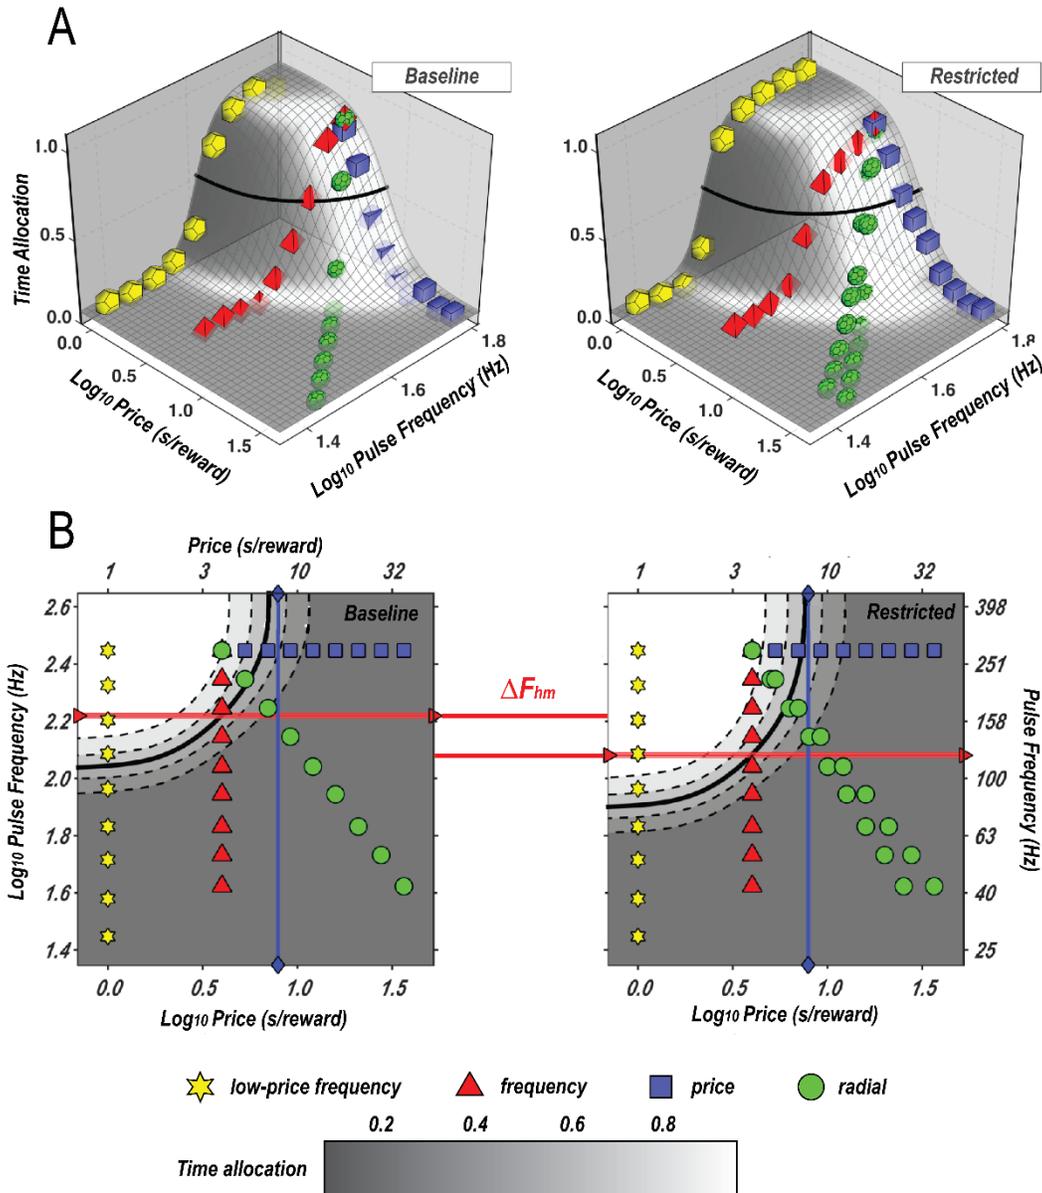


Figure 3. Effect of chronic food restriction: 3D dynamic analysis of data from rat B18. Each panel represents the surface of the winning reward mountain model permutation to the data from the baseline phase (left panels, A and B) and the end of the restriction phase (right panels, A and B). Data from the low-price frequency, frequency, price and radial sweeps are shown as yellow, red, blue and green points, respectively. A: The mountain surface fit to the data is shown in grey; paler grey corresponds to higher TA values and darker grey to lower TA values. The solid, horizontal red lines represent the F_{hm} estimates. The solid, vertical blue lines represent the P_e estimates. The paler red and blue bands framing the F_{hm} and P_e estimates represent the corresponding 95% confidence interval bands. The horizontal red lines between the right and left B panels illustrate the F_{hm} shift.

In 9 out of 10 rats, the reward mountain shifted as a result of chronic food restriction. Figure 4 and Figure 5 summarize the results of the Dynamic analysis. Figure 4 illustrates the maximal shifts in both location parameter estimates between the baseline and the restriction phases for all subjects: with decreasing body weight, F_{hm} decreased (between -0.128 and -0.217 log units) in the case of 3 rats (B11, B18, B26), increased (by 0.037 and 0.083 log units) in the case of 2 rats (B14, B32), and did not shift in the case of 5 rats (B12, B13, B15, B19, B27); P_e decreased (between -0.042 and -0.189 log units) in 4 rats (B11, B12, B13, B15), increased (0.096 log unit shift) in the case of 1 rat (B27), and did not shift in 5 cases (B14, B18, B19, B26, B32). See Table 1 for a more detailed list of the chronic FR effects.

Figure 5 illustrates the winning models for each subject: the 6-parameter model generated the lowest AIC value in the case of 4 subjects (B14, B18, B27, B32), whereas the 7-parameter model yielded the lowest AIC value for the 6 remaining subjects (B11, B12, B13, B14, B19, B26); the slope model yielded the lowest AIC value the case of 3 rats (B11, B27, B32), whereas the step model yielded the lowest AIC value in the case of 6 rats (B12, B13, B14, B15, B18, B26). Table 2 details the AIC values which determined the winning model permutation of one example rat (rat B12).

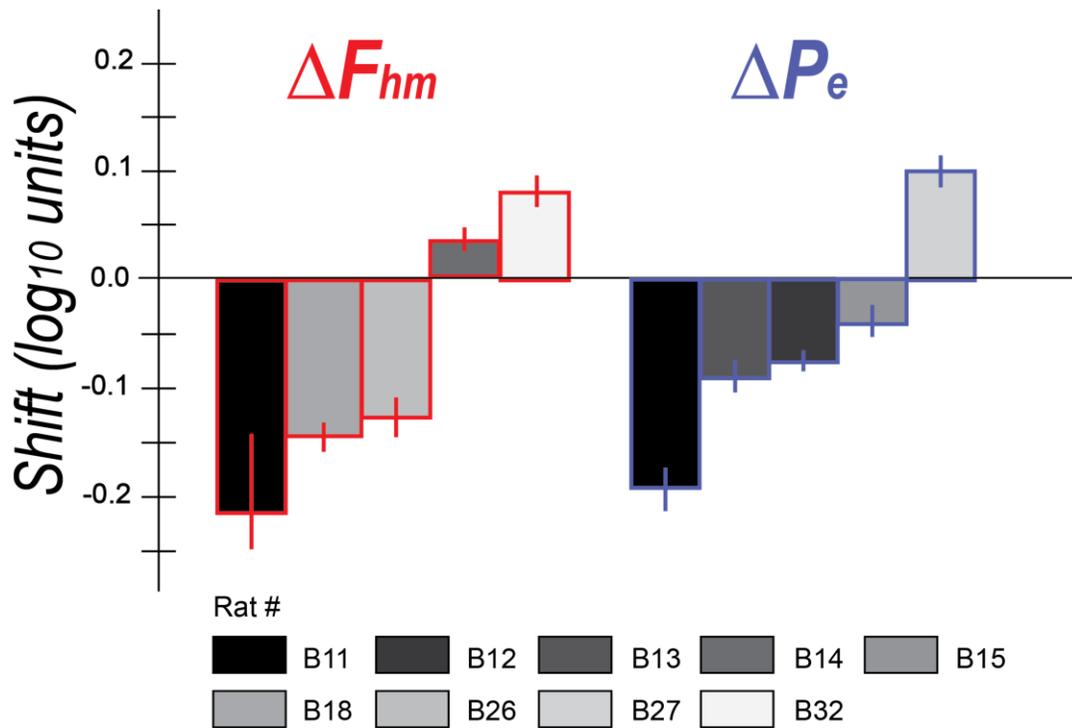


Figure 4. Maximal shifts in location parameters as a function of chronic food restriction. Shift of the reward-mountain along the pulse frequency (F_{hm}) and price (P_e) axis between baseline sessions and last restriction session. Error bars represent 95% confidence intervals.

Table 1. 3D dynamic analysis: maximal location parameter shifts (in \log_{10} units) between baseline and restriction phases.

Rat	F_{hm} shift	P_e shift
B11	-0.217 ± 0.026	-0.189 ± 0.025
B12	-	-0.075 ± 0.005
B13	-	-0.094 ± 0.007
B14	0.037 ± 0.006	-
B15	-	-0.042 ± 0.006
B18	-0.138 ± 0.004	-
B19	-	-
B26	-0.128 ± 0.009	-
B27	-	0.096 ± 0.015
B32	0.083 ± 0.005	-

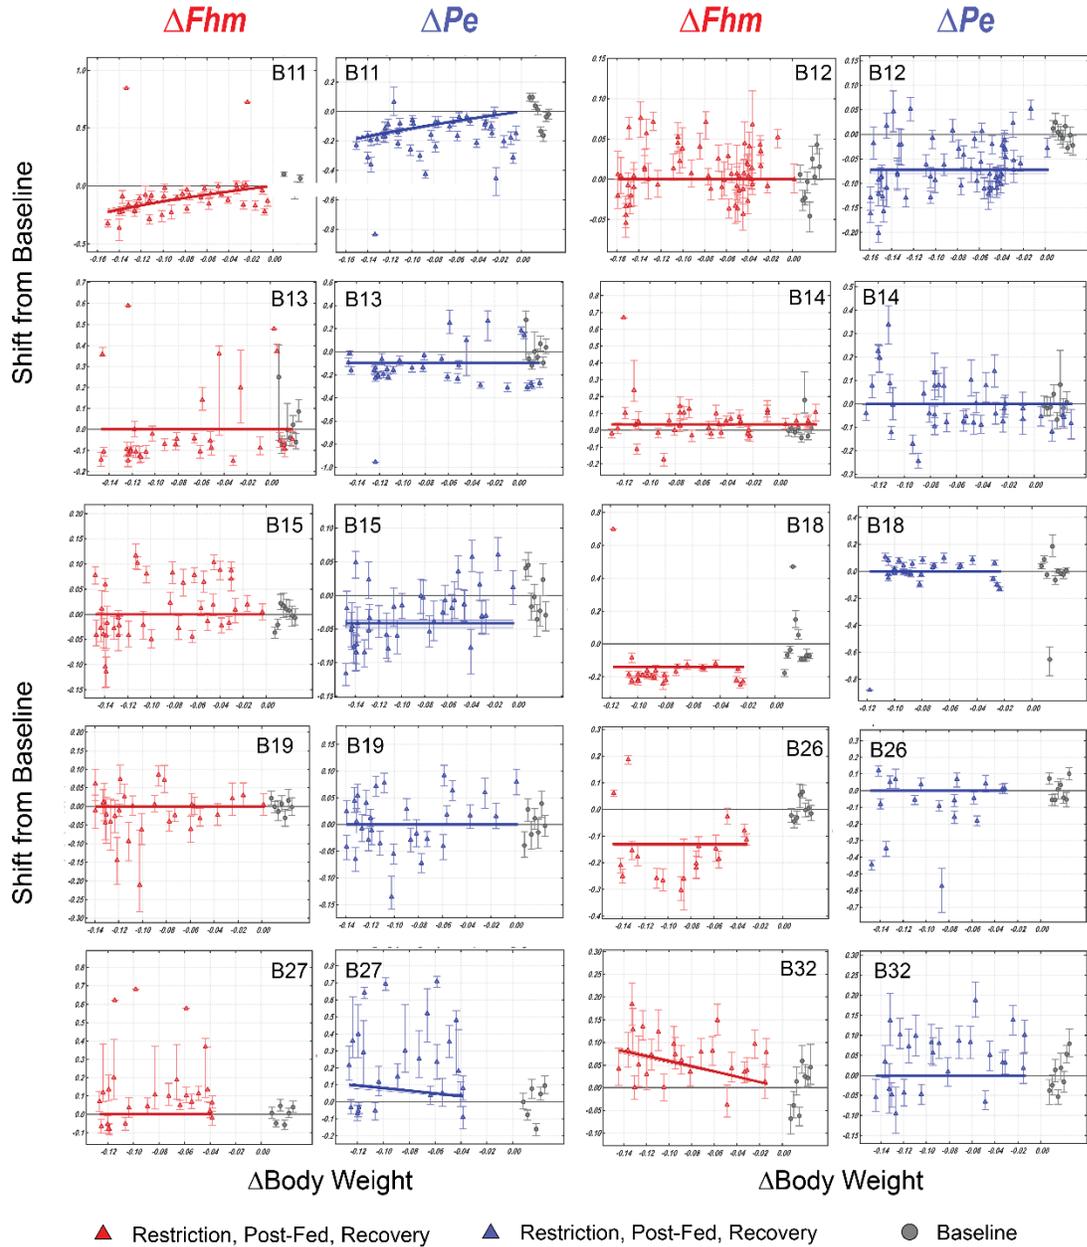


Figure 5. Winning 3D dynamic models. Each panel illustrates the location parameter shifts that were yielded by the dynamic model with the lowest AIC value for each rat. Red triangles represent the daily F_{hm} parameter estimates for restriction, post-fed stable restriction, and recovery days; blue triangles represent the daily P_e estimates for restriction, post-fed stable restriction, and recovery days; grey circles represent the location parameter estimates obtained during baseline testing. The size and slope of F_{hm} and P_e shifts are illustrated by the solid red or blue line, respectively.

Table 2.

AIC values for rat B12. Includes the four models with the lowest AIC values, all of which include a shift along the price axis (P_e).

Model permutation	AIC value	Rank
Conditioned reward; Step; P_e free to vary	-14634.79	1
Standard reward-mountain; Step; P_e free to vary	-14606.65	2
Conditioned reward; Slope; P_e free to vary	-14595.25	3
Standard reward-mountain; Step; P_e free to vary	-14562.95	4

3.3. Effect of chronic food restriction: 2D analysis

Figure 6 plots change in M-50 against change in body weight and the fitted regression line in all rats (except for rats B12, B15, and B19 who were excluded from the analysis). In the case of 4 rats, the M-50 value decreased with decrease in body weight; in the case of 2 rats, the M-50 value did not change systematically with change in body weight; in the case of 1 rat, the M-50 increased with decreasing body weight. Table 3 summarizes the results of the 2D regression analysis.

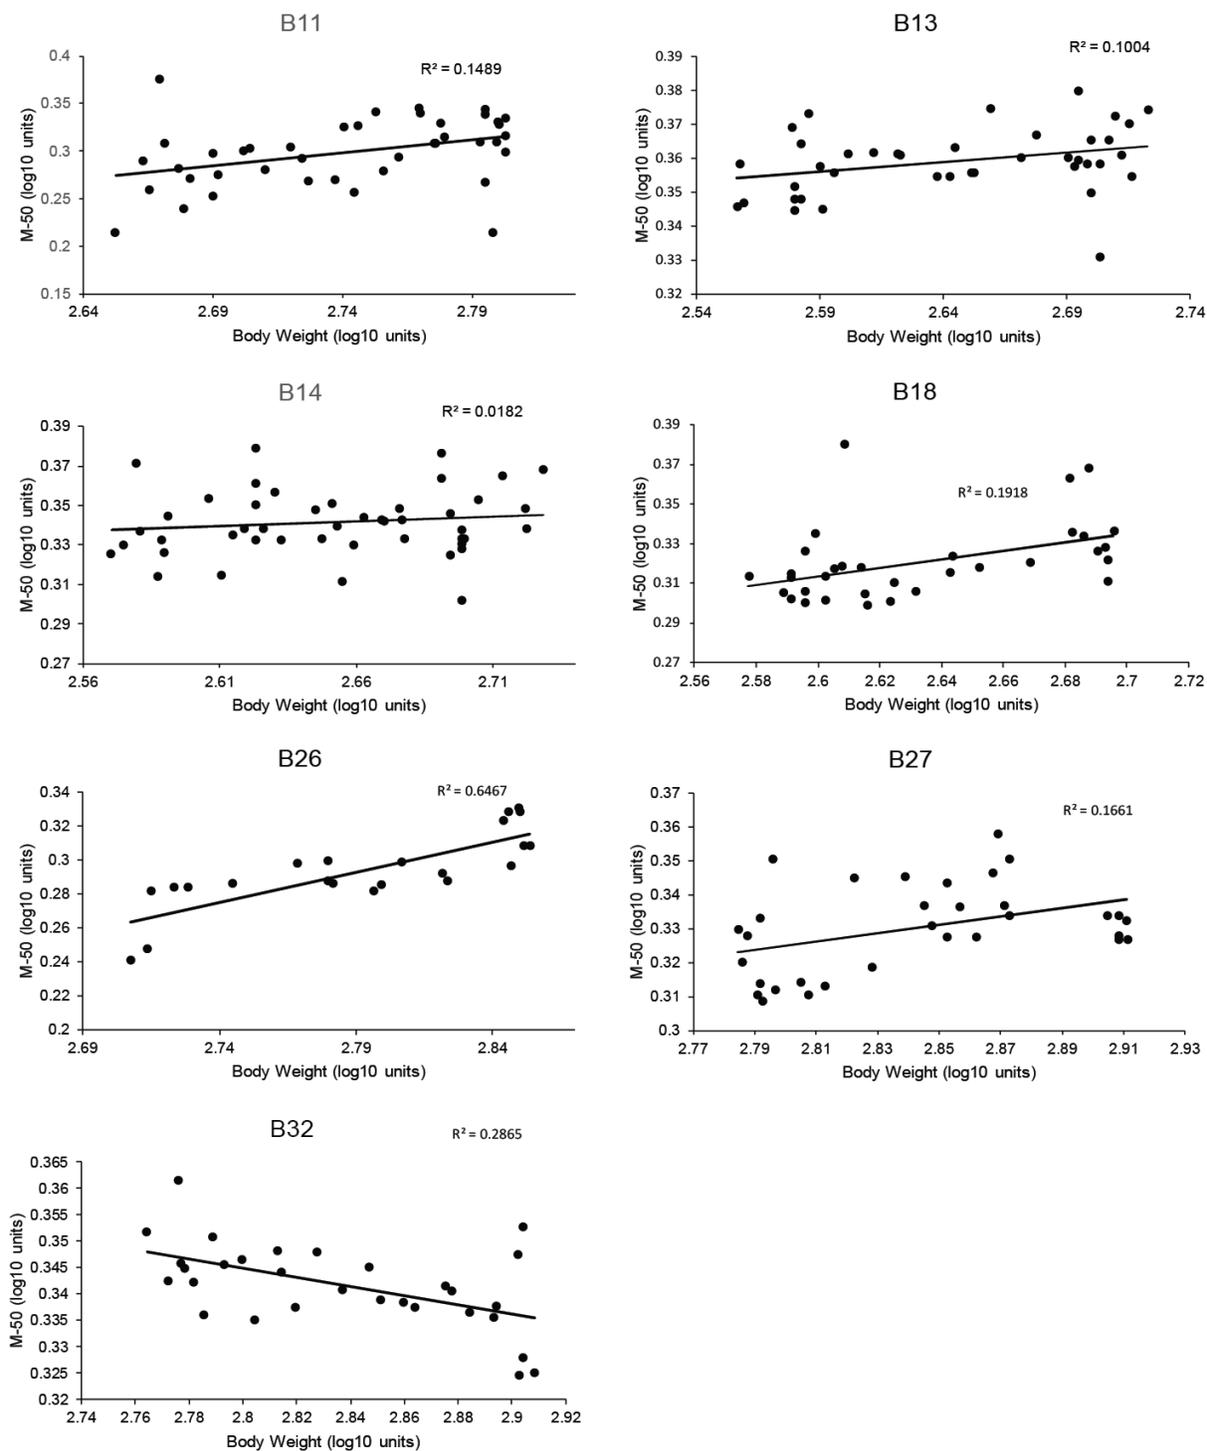


Figure 6. Regression of M-50 values on body weight for each rat. The M-50 values were derived from the results of the 2D, low-price, frequency sweeps.

Table 3.
Regression statistics for each subject.

Rat	Slope of regression	<i>p</i>	R squared
B11	0.539	0.013	0.149
B13	1.794	0.055	0.076
B14	0.372	0.371	0.018
B18	0.906	0.011	0.192
B26	1.829	6.49×10^{-6}	0.647
B27	1.361	0.021	0.166
B32	-3.41	4.96×10^{-4}	0.328

3.4. Effect of pre-feeding: Pre-fed vs.-post-fed analysis

Figure 7 illustrates the winning model permutation of an example rat (subject B26) in the pre-vs.-post analysis. F_{hm} (represented by the horizontal red line) increased by 0.116 log units from the post-fed to the pre-fed condition.

Figure 8 illustrates the maximal location parameter shifts between the pre-fed and post-fed conditions for all subjects: F_{hm} increased (between 0.05 and 0.116 log units) from post-fed to pre-fed days in the case of 3 rats (B13, B15, B26); P_e decreased (between -0.031 and -0.079 log units) from post-fed to pre-fed days in 2 cases (B12, B27); F_{hm} and P_e did not shift between the two conditions in 5 rats (B11, B14, B18, B19, B32). Table 4 details the pre-fed vs. post-fed location parameter shifts.

For 4 subjects (rats B13, B14, B18 and B32), the 6-parameter model generated a lower AIC value, whereas the 7-parameter model yielded the lowest AIC in 6 cases (rats B11, B12, B15, B19, B26, and B27). Table 5 illustrates the AIC values that determined the winning model permutation for an example rat (rat B12).

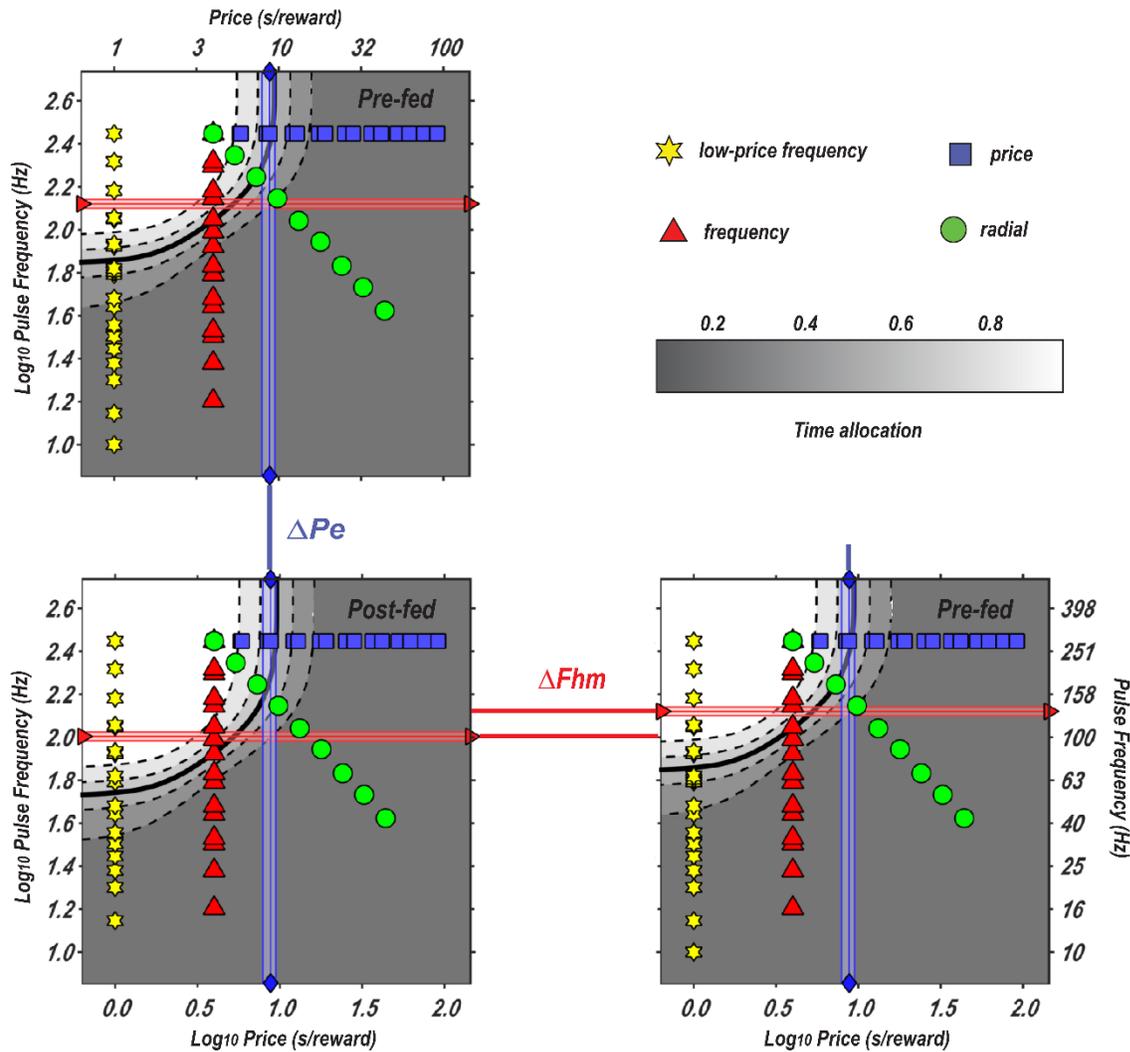


Figure 7. Effect of pre-feeding vs. post-feeding in rat B12. Each panel shows a contour graph representing the surface of the winning reward mountain model permutation fitted to the data from post-fed (lower left panel) and pre-fed (upper left and lower right panels) sessions. Data from the low-price frequency, frequency, price and radial sweeps are shown as yellow stars, red triangles, blue squares and green circles, respectively. The mountain surface fit to the data is shown in grey; paler grey corresponds to higher TA values and darker grey to lower TA values. The solid, horizontal red lines represent the F_{hm} estimates. The solid, vertical blue lines represent the P_e estimates. The paler red and blue bands framing the F_{hm} and P_e estimates represent the corresponding 95% confidence interval bands. The horizontal red lines between the lower right and left panels illustrate the F_{hm} shift: F_{hm} was lower on post-fed than on pre-fed days.

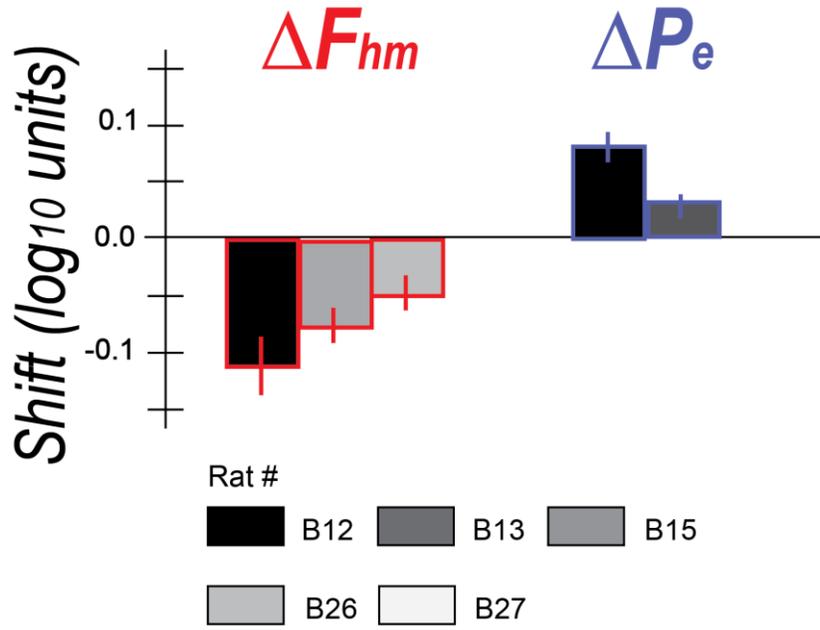


Figure 8. Shifts in location parameters from pre-fed to post-fed sessions. Error bars represent 95% confidence intervals.

Table 4

Pre-fed versus post-fed analysis: maximum location parameter shifts (in \log_{10} units) between pre-fed minus post-fed sessions.

Rat	F_{hm} shift	P_e shift
B11	-	-
B12	-	0.031 ± 0.005
B13	-0.05 ± 0.008	-
B14	-	-
B15	-0.077 ± 0.011	-
B18	-	-
B19	-	-
B26	-0.116 ± 0.017	-
B27	-	0.079 ± 0.017
B32	-	-

Table 5.

Pre-fed versus post-fed analysis: AIC values for rat B12. Includes the two models with the lowest AIC values, both of which include a shift along the price axis (P_e)

Model permutation	AIC value	Rank
Conditioned reward; Step; P_e free to vary	-8584.72	1
Standard reward-mountain; Step; P_e free to vary	-8550.56	2

3.5. Electrode Placements

Figure 9 illustrates the location of the electrode tips. In 7 subjects, the electrode was located in the perifornical region of the LH; in 2 rats, the electrode tip was located in the medial forebrain bundle at the level of the LH but fell outside of the perifornical region; in the remaining rat, the electrode tip was located outside the LH. As previously mentioned however (see section 3.1), there is reason to believe that the location of the electrode of rats B26 and B32 changed at the beginning of the recovery phase.

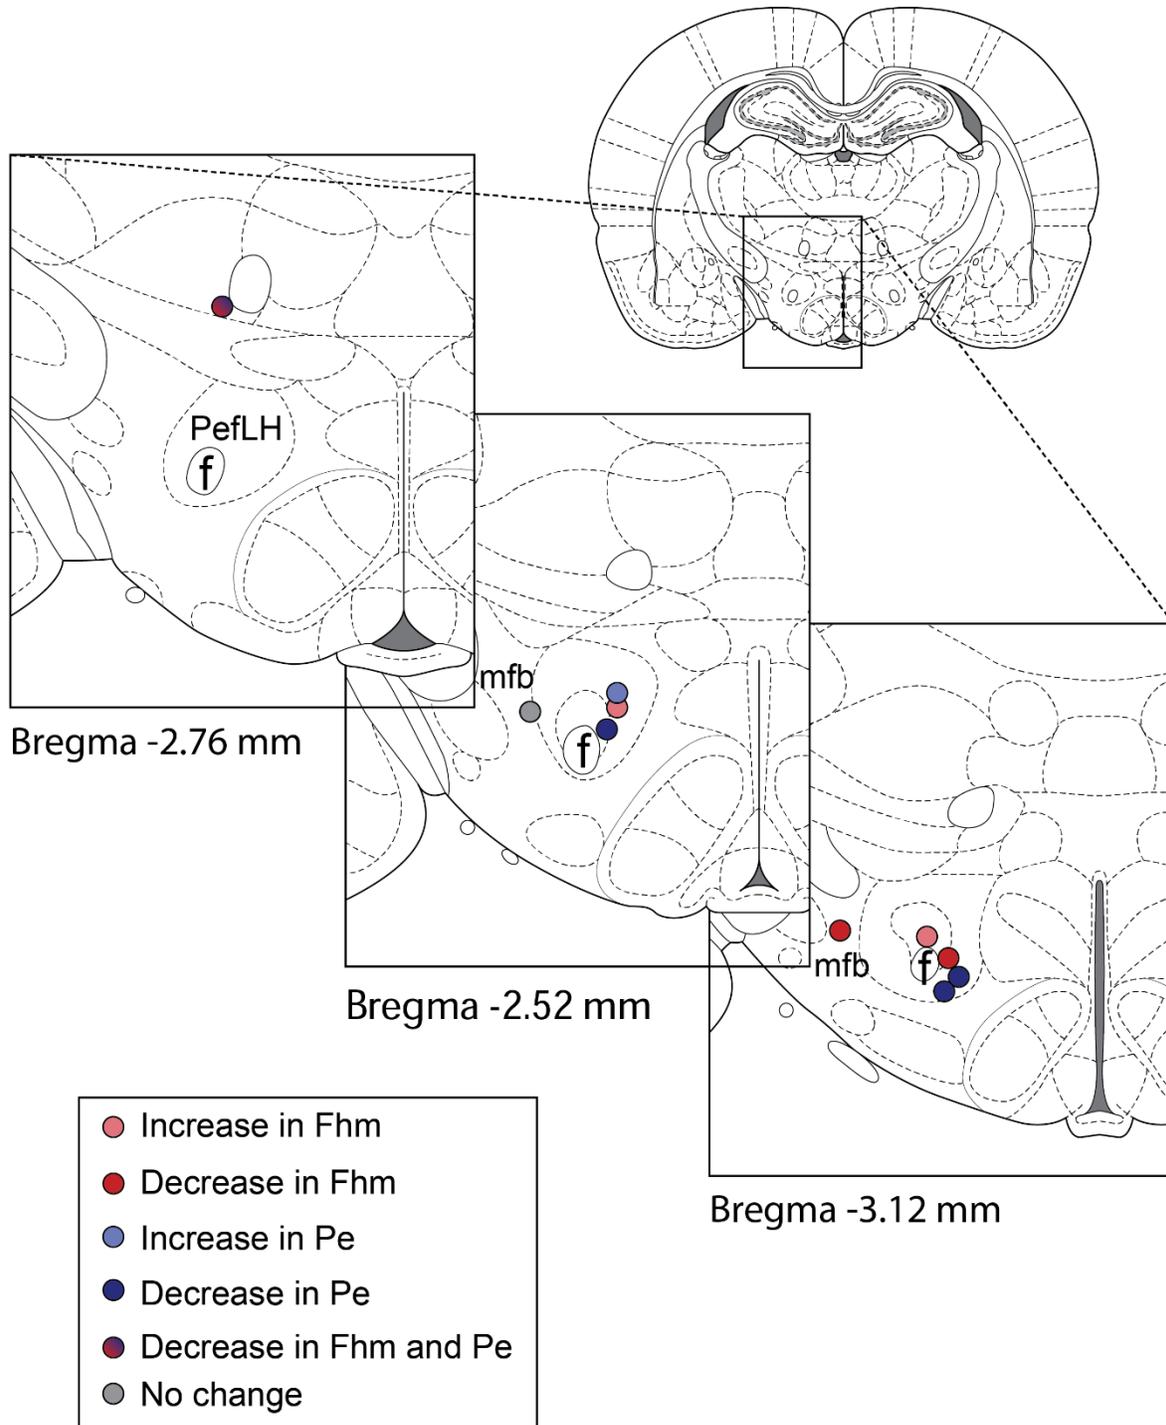


Figure 9. Electrode placements. Location of each electrode tip, as determined by low magnification microscopy and the Paxinos and Watson atlas (2007). Placements are colour-coded on the basis of the location parameter shifts obtained from the 3D dynamic analysis of the effect of chronic food restriction.

4. Discussion

Using the 3D reward-mountain strategy, the present study examined the effects of chronic food restriction on intracranial self-stimulation (ICSS). In 9 out of 10 subjects, chronic food restriction affected operant performance for BSR. The nature of these effects varied considerably between subjects: effects included both leftward and rightward displacements of the reward mountain along both the pulse-frequency and price axes.

The reward-mountain paradigm was also employed to examine whether pre-feeding, a short-term energy manipulation, could affect operant performance for BSR. Five out of 10 subjects showed small effects of meal time on BSR.

4.1. Effect of Chronic FR

The primary objective of the current study was to determine which reward components in the neural processing of BSR are affected by chronic food restriction. Previous investigations of the effect of long-term energy depletion on BSR employed 2D measurement strategies, primarily the curve-shift paradigm, which measures response rate as a function of a single independent variable. In the case of the curve-shift method, response rate is measured as a function of the pulse frequency of the electrical stimulation (graphically represented using 2D response rate-frequency curves), and the impact of chronic food restriction is assessed by measuring the change in pulse frequency needed to maintain half-maximal performance. Curve-shift experiments have consistently found that, in a subset of subjects, chronic food restriction decreased the pulse frequency needed to maintain the response-rate criterion. This effect was interpreted as an increase in reward effectiveness (Abrahamsen et al., 1995; Fulton et al., 2000; Fulton et al., 2002; Fulton et al., 2006).

More recent work however has demonstrated that 2D measurements of reward effectiveness are inherently ambiguous (Arvanitogiannis & Shizgal, 2008; Hernandez et al., 2010). Key to understanding this limitation is the fact that operant performance for BSR is determined by several variables. One of these multiple reward components is the sensitivity of the reward substrate. Reward sensitivity refers to a property of the circuitry which determines the strength of the electrical stimulation needed to generate a given reward intensity (such as a half-maximal reward intensity). Changes in reward sensitivity are comparable to changes in the K_m of an enzyme. ICSS can also be affected by changes in the gain of the reward circuitry. Gain determines the maximal rewarding impact of the stimulation and is comparable to the V_{max} of an enzyme-catalyzed reaction. Finally, operant behaviour for BSR is also determined by the subjective evaluation of the opportunity and effort costs associated with working for a given reward, the probability of obtaining the reward, and the subjective evaluation of competing activities such as resting and grooming.

A computational model has been developed to describe the manner in which the neural signal that arises from electrical stimulation of the LH relates to these various reward components (Figure 1). According to the model, the aggregate spike rate of the directly-stimulated neurons, which fire proportionally to the pulse frequency of the stimulation, is transformed into a subjective reward-intensity signal by a spatiotemporal integrator. This early processing stage determines the reward sensitivity of the substrate and is indexed by the pulse-frequency value required to produce half-maximal reward intensity (F_{hm}). A change in the pulse-frequency value needed to sustain half-maximal reward intensity (i.e. a change in F_{hm}) therefore reflects a change in the neural signalling which occurs prior to the output of the integrator and determines reward sensitivity. Next, the output of the integrator is rescaled (reflecting the gain of the reward

circuitry) and then discounted by the probability and costs of obtaining the reward, resulting in the total payoff from BSR. Comparison of the payoffs obtainable from BSR versus alternate activities results in the proportion of time spent working for BSR. The model predicts that any change in the reward processing that occurs at or beyond the output of the integrator (including a change in gain, subjective evaluation of costs and probability of BSR, and subjective evaluation of competing activities) would affect the price that produces half-maximal time allocation for a maximally-rewarding stimulation train (P_e).

In the reward-mountain method, the 3D reward-mountain structure is defined by plotting the dependent variable (time allocation) along two axes representing the independent variable. F_{hm} locates the mountain along the pulse-frequency axis, whereas P_e locates the mountain along the price axis. Displacements of the 3D reward mountain can therefore distinguish between changes in the neural signalling that occurs prior to and after the output of the integrator. The response rate-frequency curves obtained from the curve-shift method correspond to the silhouette of the 3D mountain structure projected onto a 2D plane (Hernandez et al., 2010; Trujillo-Pisanty et al., 2014). The curve-shift method consequently does not isolate the F_{hm} or P_e but instead collapses the information contained in the 3D space onto a 2D plane. As a result, identical 2D curve shifts may arise from orthogonal displacements of the 3D mountain.

To distinguish between the effects of chronic food restriction on the various reward components of BSR, the 3D strategy was employed. We found that operant performance for BSR was altered by chronic food restriction in the majority of subjects (9 out of 10 rats). The relationship between BSR and chronic food restriction that was detected is not a simple one however. The same effect of chronic food restriction was never observed in more than 3 out of 10 subjects. Instead, we found extensive between-subjects variation: in different cases, chronic

food restriction caused only a leftward displacement of the reward mountain along the pulse-frequency axis (3/10 rats), a rightward shift of the mountain along the pulse-frequency axis (2/10 rats), a leftward displacement of the mountain along the price axis (3/10 rats), a rightward displacement of the mountain along the price axis (1/10 rats), both a leftward shift along the pulse-frequency axis and a rightward shift along the price axis (1/10 rats), and no discernable change in the location of the reward mountain (1/10 rats).

Prior studies have consistently reported heterogeneous effects of chronic food restriction on operant performance for BSR (Cabeza de Vaca et al., 1998; Carr & Wolinsky, 1993; Carr et al., 2000; Fulton et al., 2000; Fulton et al., 2002; Fulton et al., 2006). Typically, roughly half of tested subjects were found to be “food restriction-sensitive” and the other half “food restriction-insensitive”, with chronic food restriction causing a leftward lateral displacement of response rate-frequency curves in the former but no change in the latter. Studies have also occasionally reported small but statistically significant rightward rate-frequency curve shifts (Carr & Wolinsky, 1993; Fulton et al., 2002). The current study reveals that chronic food restriction affects operant performance for BSR even more heterogeneously than previously believed, suggesting that the 2D perspective obscured the degree to which the influence of food restriction on ICSS varies across stimulation site. These findings illustrate the advantage of employing a measurement strategy that measures reward-seeking as a function of multiple reward components rather than a single independent variable.

In 5 out of 10 subjects, estimates of the parameter that locates the reward mountain along the pulse-frequency axis (F_{hm}) was altered by chronic food restriction. This indicates that in at least some cases long-term energy availability affects the circuitry that lies prior to the spatiotemporal integrator. However, different subjects displayed both decreased and increased

F_{hm} estimates, suggesting that chronic food restriction both increased and decreased reward sensitivity in different subjects. Similarly, both decreased and increased estimates of P_e were found in different subjects. Changes in P_e are indicative of changes in a reward component computed at or beyond the output of the spatiotemporal integrator. These include gain and the subjective evaluation of opportunity cost, effort cost, the probability of obtaining the reward, and of the value of competing activities. The varying effects of P_e could therefore be due to chronic food restriction affecting the same reward component in opposite directions in different subjects, or different reward components in different subjects. In the latter case, leftward shifts of the reward mountain along the price axis could for example result from an increased subjective evaluation of the opportunity cost of the reward, whereas rightward shifts could result from an decrease in the subjective evaluation of the value of alternate activities such as resting.

To further disentangle these later reward components, future experiments could employ a method that also measures TA as a function of effort cost and plots reward seeking in a four-dimensional (4D) space. The consequence of adding a third independent variable could be similar to the known advantage of transitioning from 2D to 3D measurements: chronic food restriction could potentially reveal that shifts along the opportunity cost axis in the 3D space in fact consist of separate effects along opportunity and effort cost axes in a 4D model. This suggests that a measurement method incorporating an additional number of independent reward components might uncover even greater heterogeneity in the behavioural effects of food restriction.

Prior work has strongly suggested that the variability in chronic food restriction effects is due to the anatomical location of the stimulating electrode and not to individual subject factors (Fulton et al., 2006). Food restriction-sensitive effects have typically been obtained from

electrodes located in the dorsal and dorsolateral perifornical region of the LH, whereas food restriction-insensitive effects usually occurred in cases where the electrode was placed outside of this region. Although the present findings are certainly suggestive of distinct, functionally heterogeneous reward subpopulations, there are insufficient data points for each of the observed food restriction effects to establish a correlation between the anatomical placement of the electrodes and the effects of chronic food restriction on the 3D mountain. However, what is noteworthy is that several adjacent anatomical electrode placements belong to rats with orthogonal and opposite 3D reward mountain shifts (Figure 9).

Finally, the 3D dynamic analysis produced a wholly unexpected finding: it revealed that the “step” model provided the best fit to the data in 6 out of the 9 rats for whom the long-term energy manipulation displaced the reward mountain. The step model is distinguished from the slope model in that the former assumes that location parameter estimates shift suddenly after the beginning of a new experimental phase while the latter predicts that location parameter estimates shift gradually over the course of an experimental phase as a result of the gradual change in body weight. Results suggest that the cause of the better fit provided by step model is unrelated to the specific type of food restriction effect: step effects were observed along both axes and, in the case of F_{hm} shifts, in both directions. It should be noted however that there are too few data points to assess this question statistically.

Significantly, the success of the step model is suggestive of a previously undiscovered effect in which operant performance for BSR changes abruptly at the beginning of a chronic FR phase. This is especially surprising given that previous studies have insisted on the chronic nature of the FR effect and its dependence on long-term change in body weight (Fulton et al., 2006). Moreover, the inconsistency between within-subject dynamic and the meal time results

(see section 4.2) suggests that the step effect does not relate to changes in short-term energy balance. Instead, chronic FR effects may result from subjects suddenly shifting into an entirely distinct internal need state, which may be signalled by significant changes in circulating long-term energy hormones.

4.2 Comparison of the 2D and 3D analyses

The second aim of the current study was to determine whether the 3D measurement method would reveal effects of chronic food restriction at sites previously labeled as food restriction-insensitive. To address this question, we first examined whether subjects' raw, 2D, low-price frequency-sweep data changed systematically as a function of the change in body weight. This analysis was comparable to those employed by curve-shift studies and therefore provided an estimate of whether the rats in this study would have been classified as food restriction-sensitive or insensitive in curve-shift testing. By comparing the 2D and 3D results of each individual rat, it was therefore possible to determine whether food-restriction "insensitive" rats displayed 3D shifts.

The 2D analysis revealed that 4 out of 7 subjects would have been classified as food restriction-sensitive by a curve-shift procedure: in the case of those rats, decreased body weight was accompanied by leftward lateral shifts of the 2D low-price frequency curves (Figure 6; rats B11, B18, B26 and B27). Amongst the 3 remaining rats, two showed no reliable change in 2D low-price frequency curves with change in body weight and are therefore classified as food restriction-insensitive (Figure 6; rats B13 and B14), whereas the third showed a rightward lateral shift of the 2D curves with decreased body weight (Figure 6; rat B32). These results are in line with prior studies which have typically reported that roughly half of subjects are food restriction-

sensitive, another half food restriction-insensitive, and that an occasional subject displays rightward 2D curve shifts.

Given the small number of subjects (7) in which a measure analogous to a 2D curve-shift could be computed, the question of whether the reward-mountain method is more sensitive in detecting the presence of food restriction effects cannot be reliably answered. However, it should be noted that both of the food restriction-insensitive subjects showed changes in the location of their reward mountain as a result of long-term energy deprivation. This supports the claim that the 3D strategy is more sensitive.

4.3. Effect of pre-feeding versus post-feeding

We loosely replicated an experiment by Abrahamsen et al. (1995) by assessing whether a manipulation of short-term energy balance in which restricted subjects were fed either before or after their testing sessions would alter the effect of chronic FR on BSR. Using the reward-mountain model, we have found evidence for an effect of pre-feeding in 5 out of 10 subjects.

In the case of 3 rats, the parameter that locates the reward mountain along the frequency axis, F_{hm} , was lower when the rat was fed after the testing session (post-fed). In the case of 2 rats, the parameter that locates the mountain along the price axis, P_e , was higher on post-fed testing days.

Supporting our results, Abrahamsen et al. (1995) reported that M-50 values obtained from testing sessions performed after a meal were slightly elevated in the case of both food-restricted and free-feeding control rats. As described above (see section 4.1), rightward shifts of 2D response rate-frequency curves can be caused by either a rightward shift of the mountain along the pulse-frequency axis or by a leftward shift of the mountain along the price axis. It should be noted that Abrahamsen et al. (1995) also compared the effect of the short-term manipulation on M-50 values in chronically food restricted versus free-feeding rats and concluded that meal time

did not significantly interact with the effect of chronic food restriction on BSR. Since we alternated meal time only in food restriction subjects that had lost 25% of their baseline body weight, we cannot presently determine whether the observed effects were due to an interaction between long- and short-term energy balance, or due to short-term signals alone. It should be noted however that there does not appear to be a strong relationship between individual subjects' 3D dynamic and pre-vs.-post results. Subjects whose reward mountain was displaced on the price axis by the long-term energy manipulation showed meal time effects on the pulse-frequency axis, for example. Though this observation should be analyzed statistically in a follow-up study with a larger sample size, it nonetheless suggests that the effects of meal time are separate from the chronic food restriction effects.

In contrast to the findings of the present experiment and Abrahamsen and colleagues (1995), other studies have reported that short-term energy manipulations including acute (48-hour) food deprivation, glucoprivation, and lipoprivation fail to alter the reward effectiveness of BSR (Cabeza de Vaca et al., 1998; Fulton et al., 2000; Fulton et al., 2004). The inconsistency between those reports and our findings may be due the insensitivity of the 2D measurement method previously used.

In the case of 3 subjects, the parameter that locates the reward mountain along the pulse-frequency axis was lower on post-fed (hungry) testing days. This shift implies that meal time affected an early stage in the neural processing of reward. It further suggests increased reward sensitivity on days when the animal was fed after the session: a lower pulse frequency was needed to maintain half-maximal reward intensity on post-fed (hungry) than on pre-fed days. In the case of 2 rats, meal time affected the reward mountain along the price axis, indicating that meal time affected reward-seeking for BSR at or beyond the output of the integrator.

Specifically, the parameter that locates the reward mountain along the price axis, P_e , was lower on post-fed than on pre-fed days, suggesting that rats were willing to pay a higher price for a maximally-rewarding stimulation when hungry than when sated.

These results are both consistent with a presumably higher state of hunger being associated with greater overall payoff from BSR, albeit due to effects on different psychological components and at different stages of neural processing. Consistent with these findings, studies have suggested that motivation and the reinforcing effects of food, drugs of abuse and the cues associated with them are also enhanced by acute manipulations of short-term energy balance (Jewett, Cleary, Levine, Schaal, & Thompson, 1995; Reilly, 1999; Shalev, Yap, & Shaham, 2001).

It should be noted however that the interpretation of these results is highly dependent upon the perspective from which they are viewed. For example, in the case of subjects who displayed smaller P_e parameter estimates on post-fed testing days, it is possible that circulating levels of hormones associated with acute hunger (the GI tract-released hormone ghrelin, for example) affect certain stimulation sites in a manner that results in increased willingness to work for BSR on post-fed testing days. On the other hand, it may be that post-ingestive effects (for example, metabolic satiety signals such as the hormone cholecystokinin) cause a decreased willingness to work for BSR on pre-fed testing days. Caution should therefore be employed before attributing the cause of meal-time effects to hunger as opposed to post-ingestive satiety or other unknown variables. The separate short-term energy signals that underlie each of these possible causes could offer an avenue for further research.

4.3 Limitations

One limitation of the 3D dynamic analysis is that it assumes that data obtained during the restriction and recovery phases can be grouped together. It should be noted however that subjects are in fundamentally different physiological need states during these two phases (a negative energy state during the restriction phase versus a positive energy state during recovery). The unexpectedly abrupt effect of internal need state on BSR has been highlighted by the step effect: the finding that most chronic food restriction effects occurred abruptly at the beginning of the restriction phase rather than gradually as a result of gradual loss in body weight. Grouping data collected during different need states together may therefore be masking a more complex relationship between chronic food restriction and reward-seeking for BSR.

Moreover, the current fitting procedure also imposes common scale (maximum and minimum TA) and shape (a and g) parameters onto the entire dataset, regardless of experimental phase or body weight. Since the current analyses' use of common scale and shape parameters is mainly for statistical rather than theoretical reasons, this practice may also be subject to revisions.

4.4 Conclusion

Application of the reward-mountain method in the present study has confirmed the importance of the 3D measurement strategy: the current findings have revealed a host of chronic FR effects that were previously undetectable using 2D methods. Across subjects, chronic FR affected reward processing both upstream and downstream to the spatiotemporal integrator in a manner suggestive of both increased and decreased reward sensitivity and willingness to work for BSR. These findings lend further support to the conclusion of prior studies that there exist

functionally heterogeneous reward substrates either within or coursing through the LH which are differentially modulated by signals of peripheral energy stores.

It should be noted that the electrical stimulation employed in the present study is known to activate cell body populations and fibers of passage non-specifically. An optogenetic approach involving subjects working for optical activation of specific LH subpopulations may both yield more homogeneous results and help identify the neurons responsible for each of the food restriction effects.

Surprisingly, the current study found that most changes in operant performance for BSR occurred abruptly at the beginning of the food restriction regimen rather than gradually as a result of gradual loss in body weight. Finally, results also support the existence of at least two separate effects of short-term energy balance on BSR.

References

- Abrahamsen, G. C., Berman, Y., & Carr, K. D. (1995). Curve-shift analysis of self-stimulation in food-restricted rats: relationship between daily meal, plasma corticosterone and reward sensitization. *Brain research*, 695(2), 186-194.
- Abrahamsen, G. C., & Carr, K. D. (1996). Effects of corticosteroid synthesis inhibitors on the sensitization of reward by food restriction. *Brain research*, 726(1), 39-48.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, 19, 716–723.
- Anderson, D.R. (2008). *Model based inference in the life sciences: A primer on evidence*. New York, NY: Springer.
- Arvanitogiannis, A., & Shizgal, P. (2008). The reinforcement mountain: allocation of behavior as a function of the rate and intensity of rewarding brain stimulation. *Behavioral neuroscience*, 122(5), 1126.
- Blundell, J. E., & Herberg, L. J. (1968). Relative effects of nutritional deficit and deprivation period on rate of electrical self-stimulation of lateral hypothalamus. *Nature*.
- Breton, Y., Conover, K., & Shizgal, P. (2014). The effect of probability discounting on reward seeking: a three-dimensional perspective. *Frontiers in Behavioral Neuroscience*, 8, 284.
- Breton, Y. A., Marcus, J. C., & Shizgal, P. (2009). Rattus Psychologicus: construction of preferences by self-stimulating rats. *Behavioural brain research*, 202(1), 77-91.
- Breton, Y. A., Mullett, A., Conover, K., & Shizgal, P. (2013). Validation and extension of the reward-mountain model. *Frontiers in behavioral neuroscience*, 7.
- Cabeza de Vaca, S., Holiman, S., & Carr, K. D. (1998). A search for the metabolic signal that

- sensitizes lateral hypothalamic self-stimulation in food-restricted rats. *Physiology & behavior*, 64(3), 251-260.
- Cabeza de Vaca, S., Krahne, L. L., & Carr, K. D. (2004). A progressive ratio schedule of self-stimulation testing in rats reveals profound augmentation of d-amphetamine reward by food restriction but no effect of a “sensitizing” regimen of d-amphetamine. *Psychopharmacology*, 175(1), 106-113.
- Carr, K. D. (2007). Chronic food restriction: enhancing effects on drug reward and striatal cell signaling. *Physiology & behavior*, 91(5), 459-472.
- Carr, K. D., Kim, G. Y., & Cabeza de Vaca, S. (2000). Hypoinsulinemia may mediate the lowering of self-stimulation thresholds by food restriction and streptozotocin-induced diabetes. *Brain research*, 863(1), 160-168.
- Carr, K. D., & Wolinsky, T. D. (1993). Chronic food restriction and weight loss produce opioid facilitation of perifornical hypothalamic self-stimulation. *Brain research*, 607(1), 141-148.
- 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.
- 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.
- Cornish, E. R., & Mrosovsky, N. (1965). Activity during food deprivation and satiation of six species of rodent. *Animal behaviour*, 13(2), 242-248.
- D’Cunha, T. M., Sedki, F., Macri, J., Casola, C., & Shalev, U. (2013). The effects of chronic

- food restriction on cue-induced heroin seeking in abstinent male rats. *Psychopharmacology*, 225(1), 241-250.
- 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(5), 876.
- Figlewicz, D. P., Evans, S. B., Murphy, J., Hoen, M., & Baskin, D. G. (2003). Expression of receptors for insulin and leptin in the ventral tegmental area/substantia nigra (VTA/SN) of the rat. *Brain research*, 964(1), 107-115.
- Fulton, S. (2010). Appetite and reward. *Frontiers in neuroendocrinology*, 31(1), 85-103.
- Fulton, S., Pissios, P., Manchon, R. P., Stiles, L., Frank, L., Pothos, E. N., ... & Flier, J. S. (2006). Leptin regulation of the mesoaccumbens dopamine pathway. *Neuron*, 51(6), 811-822.
- Fulton, S., Richard, D., Woodside, B., & Shizgal, P. (2002). Interaction of CRH and energy balance in the modulation of brain stimulation reward. *Behavioral neuroscience*, 116(4), 651.
- Fulton, S., Richard, D., Woodside, B., & Shizgal, P. (2004). Food restriction and leptin impact brain reward circuitry in lean and obese Zucker rats. *Behavioural brain research*, 155(2), 319-329.
- Fulton, S., Woodside, B., & Shizgal, P. (2000). Modulation of brain reward circuitry by leptin. *Science*, 287(5450), 125-128.
- Fulton, S., Woodside, B., & Shizgal, P. (2002). Does neuropeptide Y contribute to the modulation of brain stimulation reward by chronic food restriction? *Behavioural brain research*, 134(1), 157-164.

- Fulton, S., Richard, D., Woodside, B., & Shizgal, P. (2004). Food restriction and leptin impact brain reward circuitry in lean and obese Zucker rats. *Behavioural brain research*, *155*(2), 319-329.
- Fulton, S., Woodside, B., & Shizgal, P. (2006). Potentiation of brain stimulation reward by weight loss: evidence for functional heterogeneity in brain reward circuitry. *Behavioural brain research*, *174*(1), 56-63.
- Frutiger, S. A. (1986). Changes in self-stimulation at stimulation-bound eating and drinking sites in the lateral hypothalamus during food or water deprivation, glucoprivation, and intracellular or extracellular dehydration. *Behavioral neuroscience*, *100*(2), 221.
- Frutiger, S. A., & Drinkwine, P. (1992). Effect of glucoprivation on self-stimulation rate-frequency functions. *Physiology & behavior*, *52*(2), 313-319.
- Gallistel, C. R., Shizgal, P., & Yeomans, J. S. (1981). A portrait of the substrate for self-stimulation. *Psychological review*, *88*(3), 228.
- Goto, M., Canteras, N. S., Burns, G., & Swanson, L. W. (2005). Projections from the subfornical region of the lateral hypothalamic area. *Journal of Comparative Neurology*, *493*(3), 412-438.
- Hahn, J. D. (2010). Comparison of melanin-concentrating hormone and hypocretin/orexin peptide expression patterns in a current parceling scheme of the lateral hypothalamic zone. *Neuroscience letters*, *468*(1), 12-17.
- Hahn, J. D., & Swanson, L. W. (2012). Connections of the lateral hypothalamic area juxtadorsomedial region in the male rat. *Journal of Comparative Neurology*, *520*(9), 1831-1890.

- Hernandez, G., Breton, Y. A., Conover, K., & Shizgal, P. (2010). At what stage of neural processing does cocaine act to boost pursuit of rewards?. *PLoS One*, 5 (11), e15081.
- Hernandez, G., Trujillo-Pisanty, I., Cossette, M. P., Conover, K., & Shizgal, P. (2012). Role of dopamine tone in the pursuit of brain stimulation reward. *The Journal of Neuroscience*, 32(32), 11032-11041.
- Hodos, W. (1961). Progressive ratio as a measure of reward strength. *Science*, 134(3483), 943-944.
- Jewett, D. C., Cleary, J., Levine, A. S., Schaal, D. W., & Thompson, T. (1995). Effects of neuropeptide Y, insulin, 2-deoxyglucose, and food deprivation on food-motivated behavior. *Psychopharmacology*, 120(3), 267-271.
- Koubi, H. E., Robin, J. P., Dewasmes, G., Le Maho, Y., Frutoso, J., & Minaire, Y. (1991). Fasting-induced rise in locomotor activity in rats coincides with increased protein utilization. *Physiology & behavior*, 50(2), 337-343.
- Leininger, G. M., Jo, Y. H., Leshan, R. L., Louis, G. W., Yang, H., Barrera, J. G., ... & Myers Jr, M. G. (2009). Leptin acts via leptin receptor-expressing lateral hypothalamic neurons to modulate the mesolimbic dopamine system and suppress feeding. *Cell metabolism*, 10(2), 89-98.
- Margules, D. L., & Olds, J. (1962). Identical "feeding" and "rewarding" systems in the lateral hypothalamus of rats. *Science*, 135(3501), 374-375.
- Miliaressis, E., Rompre, P. P., Laviolette, P., Philippe, L., & Coulombe, D. (1986). The curve-shift paradigm in self-stimulation. *Physiology & behavior*, 37(1), 85-91.
- Niv, Y., Daw, N. D., & Dayan, P. (2005). How fast to work: Response vigor, motivation and

- tonic dopamine. In *NIPS* (Vol. 18, pp. 1019-1026).
- Niv, Y., Joel, D., & Dayan, P. (2006). A normative perspective on motivation. *Trends in cognitive sciences*, *10*(8), 375-381.
- Paxinos, G. & Watson, C. (2007). *The rat brain in stereotaxic coordinates* 6 ed, Elsevier/Academic Press.
- Pirke, K. M., Broocks, A., Wilckens, T., Marquard, R., & Schweiger, U. (1993). Starvation-induced hyperactivity in the rat: the role of endocrine and neurotransmitter changes. *Neuroscience & Biobehavioral Reviews*, *17*(3), 287-294.
- Reilly, S. (1999). Reinforcement value of gustatory stimuli determined by progressive ratio performance. *Pharmacology Biochemistry and Behavior*, *63*(2), 301-311.
- Rossi III, J., & Panksepp, J. (1992). Analysis of the relationships between self-stimulation sniffing and brain-stimulation sniffing. *Physiology & behavior*, *51*(4), 805-813.
- Shalev, U., Yap, J., & Shaham, Y. (2001). Leptin attenuates acute food deprivation-induced relapse to heroin seeking. *J Neurosci*, *21*(4), 1-5.
- Shalev, U. (2012). Chronic food restriction augments the reinstatement of extinguished heroin-seeking behavior in rats. *Addiction biology*, *17*(4), 691-693.
- Shizgal, P. (1997). Neural basis of utility estimation. *Current opinion in neurobiology*, *7*(2), 198-208.
- Shizgal, P., Fulton, S., & Woodside, B. (2001). Brain reward circuitry and the regulation of energy balance. *International journal of obesity and related metabolic disorders: journal of the International Association for the Study of Obesity*, *25*, S17-21.
- Solomon, R.B., Trujillo-Pisanty, I., & Shizgal, P. (2010). "The maximum firing frequency of the

neurons subserving brain stimulation reward,” in *Society for Neuroscience Abstract 716.3* (San Diego, CA). [online].

Trujillo-Pisanty, I., Conover, K., & Shizgal, P. (2014). A new view of the effect of dopamine receptor antagonism on operant performance for rewarding brain stimulation in the rat. *Psychopharmacology*, *231*(7), 1351-1364.

Trujillo-Pisanty, I., Hernandez, G., Moreau-Debord, I., Cossette, M. P., Conover, K., Cheer, J. F., & Shizgal, P. (2011). Cannabinoid receptor blockade reduces the opportunity cost at which rats maintain operant performance for rewarding brain stimulation. *The Journal of Neuroscience*, *31*(14), 5426-5435.

Trujillo-Pisanty, I., Conover, K., & Shizgal, P. (2014). A new view of the effect of dopamine receptor antagonism on operant performance for rewarding brain stimulation in the rat. *Psychopharmacology*, *231*(7), 1351-1364.

Vaccarino, F. J., & Koob, G. F. (1984). Microinjections of nanogram amounts of sulfated cholecystokinin octapeptide into the rat nucleus accumbens attenuates brain stimulation reward. *Neuroscience letters*, *52*(1), 61-66.

Waraczynski, M. A., & Kaplan, J. M. (1990). Frequency-response characteristics provide a functional separation between stimulation-bound feeding and self-stimulation. *Physiology & behavior*, *47*(5), 843-851.