

Cognition and Behavior

Regional Activity in the Rat Anterior Cingulate Cortex and Insula during Persistence and Quitting in a Physical-Effort Task

ID Blake S. Porter, Kunling Li, and Kristin L. Hillman

<https://doi.org/10.1523/ENEURO.0243-20.2020>

Department of Psychology, Brain Health Research Centre, University of Otago, 9016, Dunedin, New Zealand

Abstract

As animals carry out behaviors, particularly costly ones, they must constantly assess whether or not to persist in the behavior or quit. The anterior cingulate cortex (ACC) has been shown to assess the value of behaviors and to be especially sensitive to physical effort costs. Complimentary to these functions, the insula is thought to represent the internal state of the animal including factors such as hunger, thirst, and fatigue. Using a novel weight-lifting task for rats, we characterized the local field potential (LFP) activity of the ACC and anterior insula (AI) during effort expenditure. In the task, male rats are challenged to work for sucrose reward, which costs progressively more effort over time to obtain. Rats are able to quit the task at any point. We found modest shifts in LFP theta (7–9 Hz) activity as the task got progressively more difficult in terms of absolute effort expenditure. However, when the LFP data were analyzed based on the relative progress of the rat toward quitting the task, substantial shifts in LFP power in the theta and gamma (55–100 Hz) frequency bands were observed in ACC and AI. Both ACC and AI theta power decreased as the rats got closer to quitting, while ACC and AI gamma power increased. Furthermore, coherency between ACC and AI in the delta (2–4 Hz) range shifted alongside the performance state of the rat. Overall, we show that ACC and AI LFP activity changes correlate to the relative performance state of rats in an effort-based task.

Key words: anterior cingulate cortex; anterior insula; cost–benefit decision-making; effort; quitting

Significance Statement

Animals need to assess whether or not a behavior is worth pursuing based on their internal states (e.g., hunger, fatigue) and the costs and benefits of the behavior. However, internal states often change as behaviors are conducted, such as becoming fatigued, necessitating constant reassessment as to whether to continue the behavior or quit. We characterized brain activity in the anterior cingulate cortex and insula, brain regions involved in cost–benefit decision-making and internal state representations, respectively, as rats conducted a challenging physical-effort task. Both brain regions showed significant shifts in activity as the rats approached their quitting point. Our study provides one of the first characterizations of neural activity as an animal decides to quit an effortful task.

Introduction

When tasks become effortful, animals must evaluate whether additional physical and/or mental effort expenditures are worthwhile. This evaluation requires both a cost–benefit assessment of the task and an interoceptive

assessment—an animal must gauge physiological and/or psychological preparedness to handle any continued effort expenditure. While multiple brain regions are involved in these processes, the following two regions appear

Received June 8, 2020; accepted August 13, 2020; First published August 28, 2020.

The authors declare no competing financial interests.

Author contributions: B.S.P. and K.L.H. designed research; B.S.P. and K.L. performed research; B.S.P. contributed unpublished reagents/analytic tools; B.S.P. analyzed data; B.S.P. and K.L.H. wrote the paper.

central: the anterior cingulate cortex (ACC) and the anterior insula (AI).

Based on lesion and electrophysiological studies, the ACC is implicated in effort-based cost–benefit analysis. For example, when rats are presented behavioral options that include varying levels of effort, ACC single-unit activity is reflective of the behavior with highest utility (Hillman and Bilkey, 2010; Cowen et al., 2012; Hart et al., 2019; Porter et al., 2019). Human EEG recordings indicate that frontal midline theta, thought to originate from the ACC, correlates to task effort and difficulty (Smit et al., 2005; Cavanagh and Frank, 2014). Together, heightened population activity in the ACC may help drive action toward high-utility goals when effort is required (Amiez et al., 2006; Kennerley and Wallis, 2009; Cowen et al., 2012; Hillman and Bilkey, 2012).

The AI is implicated in representing the interoceptive state of the body. The insula receives hypothalamic, autonomic, and visceral inputs making the AI well positioned to process interoceptive information on states such as hunger and pain, which can influence cost–benefit decision-making (Cechetto and Saper, 1987; Craig, 2003; Critchley, 2005). In rats, lesions to the AI result in maladaptive persistence and perseveration behaviors in reward devaluation scenarios (Balleine and Dickinson, 2000; Parkes et al., 2015; Moschak et al., 2018). One interpretation of this is that AI functionality is needed during ongoing cost–benefit decision-making to spur effort reallocation—in the form of task switching and/or quitting—when a behavior has decreased utility relative to the needs of the body. In turn, the interoceptive state representation of the AI could be used by the ACC to evaluate the costs and benefits of persisting at or quitting a behavior.

White matter tracts between the cingulate and insula have been traced in humans (Moisset et al., 2010), and resting-state functional connectivity between the regions has been reported (Sridharan et al., 2008; Taylor et al., 2009). As outlined by Medford and Critchley (2010), the regions could function in tandem, as follows: the AI integrating external and internal sensory inputs to represent the current state of the organism, with the ACC then using this state representation to drive relevant behaviors. If external and internal sensory inputs to the AI increasingly signal effort costs (fatigue, energy depletion), the ACC could bias behaviors away from continued energy expenditure. Both brain regions have been shown to encode the expected energetic cost of physically demanding behaviors (Prévost et al., 2010), and as tasks get increasingly effortful the ACC and AI show conjoint BOLD activation (Engström et al., 2014). It is not well

understood, however, how dynamic changes in conjoint ACC–AI activity correlate to changes in effortful task performance. Furthermore, few studies have examined ACC–AI activity to the point of quitting.

To address these questions we recorded local field potentials (LFPs) from the ACC and AI of rats as they performed a novel weight-lifting task (WLT). The WLT allows systematic manipulation of physical-effort costs while maintaining the same motor pattern, spatial context, and reward. LFPs were analyzed for shifts in power across the following major frequency bands: delta (2–4 Hz), theta (7–9 Hz), beta (15–25 Hz), and gamma (55–100 Hz). Oscillation frequency is thought to reflect the number of synchronous neurons engaged in a given process—higher frequencies reflect local activity while slower rhythms reflect the synchronous activity of neurons across brain regions (von Stein and Sarnthein, 2000; Buzsáki and Draguhn, 2004). We predicted that the ACC and AI may be communicating via low-frequency oscillation during an effortful task, and that low-frequency LFP power from the ACC would inversely correlate to effort costs to reflect the decreasing utility signal generated by the ACC. In contrast, low-frequency LFP power from the AI would positively correlate to effort costs as fatigue increases and internal states drive behavior away from task engagement. In addition, we investigated the potential involvement of the beta rhythm, often associated with sensorimotor systems (Salenius and Hari, 2003) and working memory (Engel and Fries, 2010), and the gamma rhythm, thought to reflect local circuit activity (Buzsáki and Draguhn, 2004; Bartos et al., 2007; Merker, 2016). We also characterized ACC–AI LFP activity at the quitting point, a time period that to our knowledge has been minimally investigated at the electrophysiological level.

Materials and Methods

Subjects

Male Sprague Dawley rats ($n = 16$, weighing 400–575 g at the time of experiment; Hercus Taieri Resource Unit) were single housed in individually ventilated cages (38 × 30 × 35 cm; Tecniplast). Animals were maintained on a 12 h reverse light cycle, with experimentation occurring during the dark, active phase of the animals. Water was available *ad libitum*. Daily chow (Teklad diet, Envigo) was rationed to promote interest in food reward during experimentation. Rats were weighed twice per week, and daily food rations were adjusted to maintain the body weight of each animal at $\geq 85\%$ of their free-feeding body weight. All procedures were approved by the University of Otago Animal Ethics Committee (protocol 91/17).

Training

Before surgery, animals were trained in the WLT as previously described (Porter and Hillman, 2019). Briefly, the WLT occurs in a 120 × 90 × 60 cm arena (Fig. 1a,b); an animal is progressively shaped to pull a rope 30 cm to trigger the dispensing of 0.25 ml of a 20% sucrose solution. The rope can be weighted from 0 to 225 g to increase the effort demands of the task; however, only 0 and 45 g are

This research was supported by the Royal Society of New Zealand Te Apārangi and Marsden Fund Grant U001617 (to K.L.H.) from the Royal Society of New Zealand.

Correspondence should be addressed to Blake S. Porter at blakeporterneuro@gmail.com.

<https://doi.org/10.1523/ENEURO.0243-20.2020>

Copyright © 2020 Porter et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

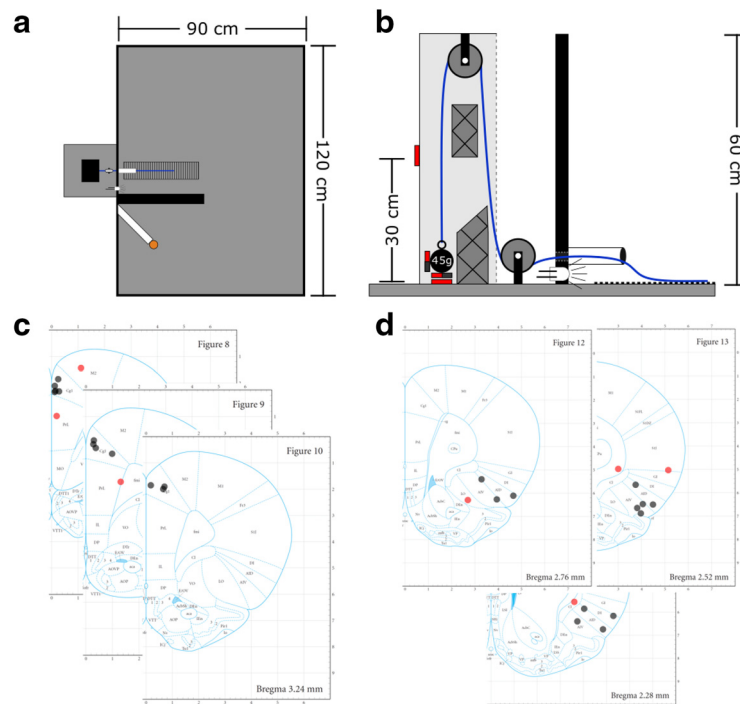


Figure 1. Experimental overview. **a**, Weight-lifting task apparatus. The pulley system is depicted by the black box on the left, outside the arena. The rope is shown in blue extending into the arena over the rubber grip pad (vertical lines). The 20% sucrose reward is delivered to the orange cap separated from the rope by the black barrier. **b**, The weight-lifting pulley schematic. The rope is shown in blue with one end attached to a 45 g weight with two magnets and the other end extending into the arena through the conduit. The hash mark shapes indicate Styrofoam inserts to prevent the weight from swinging while pulled. The magnetic sensors are shown in red. **c**, **d**, Histologic reconstructions illustrating electrode tip locations for the ACC (**c**) and AI (**d**). Black dots were on-target placements while off-target placements are red. Off-target electrodes were not used for LFP analyses. **a** and **b** were adapted from Porter and Hillman (2019); **c** and **d** were adapted from Paxinos and Watson (2007).

used in the training phase. Once an animal can perform 20 successful pulls (10 pulls on 0 g followed immediately by 10 pulls on 45 g) in under 5 min, they are considered trained in the WLT and are ready for surgery.

Surgery

Animals were implanted with two 200 μm insulated nichrome electrodes (Johnson Matthey) under isoflurane anesthesia as previously described (Porter et al., 2019). Stereotaxic coordinates, based on the study by Paxinos and Watson (2007), were set as follows: for the ACC, we targeted cingulate cortex area 1 (Cg1; equivalent to region 32d) at +3.7 mm anteroposterior (AP) from bregma, +0.4 mm mediolateral (ML) from bregma, -1.0 mm dorsoventral (DV) from bregma, and AI +2.7 mm AP from bregma, +2.0 mm from ML, and -5.8 mm from DV at an angle of 20° ML. All implants were in the right hemisphere. A 4 mm stainless steel screw placed over the cerebellum served as the ground connection. Five to seven additional 3 mm stainless steel screws were placed in the skull for dental cement to anchor to. Electrodes were secured with UV cure self-etching dental cement (Maxcem Elite, Kerr Corporation) followed by additional self-cure dental acrylic (Lang Dental Manufacturing Company) for additional strength. Ten days after surgery and postoperative care, rats were retested on the WLT with 0 and 45 g to

ensure no loss of function from surgery and to check the quality of the electrophysiological signals. Within 3 d, all rats returned to the baseline performance criterion (i.e., 10 successful trials on 0 g and 10 successful trials on 45 g, all within 5 min). Once the criterion performance was confirmed, experimentation began.

Weight-lifting task

To test persistent, effortful behavior, a progressive weight-lifting paradigm was used. This paradigm has been detailed previously (Porter and Hillman, 2019). Briefly, the rats need to pull a rope with a weight attached for 30 cm to trigger the release of 0.25 ml of a 20% sucrose solution. In the progressive WLT, an animal is challenged to pull progressively heavier weights: the weight starts at 0 g and after every 10 successful, rewarded pulls, 45 g of weight is added to the rope. The animal continues until they quit or a rope weight of 225 g is reached. If 10 successful pulls are achieved on 225 g, then the weight is made “impossible” by affixing it to the setup such that the weight can only be pulled up ~15 cm, thus making it impossible to trigger a reward. Quitting was defined as 2 min of no attempted rope pulls.

Each task session was bookended with 2 min of open field exploration in the arena. During these open field periods, the rope was not available and no sucrose was

dispensed. The initial 2 min exploration provided a baseline measure of LFP and locomotor activity before beginning the WLT. Likewise, the 2 min exploration at the end of the task provided post-task LFP and locomotor measurements. At the end of the 2 min post-task exploration, a sucrose reward check was performed to test for satiation. The sucrose satiation check entailed two manually released sucrose rewards, and the rats had 1 min to approach and consume the sucrose. Whether or not they approached and consumed the sucrose was noted by the experimenter.

Recordings

Local field potentials were recorded using a Neuralynx acquisition system (either a Digital Lynx SX or SX-M) at 6400 or 5000 Hz, respectively, with a bandpass filter set between 0.1 and 500 Hz. Rats were tethered via a Neuralynx Saturn-1 commutator to either a Neuralynx HS-36-LED or HS-36-mux-LED preamp headstage. Rat position was tracked via two headstage-mounted LEDs with Neuralynx Cheetah software using an overhead camera. Rats were ran in a darkened room to allow tracking of the LEDs. The Neuralynx acquisition system also recorded TTL (transistor–transistor logic) signals from the WLT Arduino Uno Microcontroller that was time-stamped when the weight was lifted off the base, when the weight returned down to the base, and when the weight was successfully pulled up 30 cm and a reward was released.

Histology

At the end of the study, animals were transcardially perfused with 4% paraformaldehyde in PBS and histology performed as previously described (Porter et al., 2019). Thionin staining confirmed 12 of 16 correct placements in the ACC (Fig. 1c), and 12 of 16 in the AI (Fig. 1d). Every rat included for analyses (behavioral and electrophysiological) had at least one electrode in one of the target regions (ACC or AI). Of those excluded, one ACC electrode could not be confirmed because of the tissue being damaged during histologic preparation; three ACC-targeted electrodes were confirmed to be misplaced; and four AI targeted electrodes were confirmed to have missed the AI. LFP data from electrodes with incorrect placements were excluded from analysis. Behavioral data were analyzed for all 16 rats. Note that there is continued debate on the rodent equivalent of the primate ACC (for review, see Laubach et al., 2018), and our ACC electrodes are primarily in Cingulate area 1/32d.

Experimental design and statistical analysis

Statistical analysis was conducted in MATLAB with custom and native scripts. All statistical analyses were double checked in Prism (version 8.4.1; GraphPad Software). Behavioral and LFP data were parsed based on the timestamps relayed by the Arduino TTL signals and experimenter notes using MATLAB. Methods for behavioral analyses have been reported in detail previously (Porter and Hillman, 2019). Briefly, data from a given behavioral metric was first tested for normality using a one-

sample Kolmogorov–Smirnov test. Depending on the distribution of the data, the appropriate statistical test was selected, either a one-way ANOVA or Kruskal–Wallis test, with a p value of 0.05. The Dunn–Sidak method was used for making *post hoc* multiple comparisons.

LFP analyses

Each rat ($n=16$) ran the progressive WLT for five to eight recording sessions, one session per day. In total, there were 83 recording sessions for ACC data and 82 for AI data. Each recording session was broken up into epochs based on the different rope weights (e.g., 0, 45, 90 g). Each epoch began when the respective weight was attached to the rope and ended when the next weight was attached after 10 successful trials or until the rat quit. The 0 g epoch began immediately following the 2 min prebaseline period when the rope was extended into the apparatus. Our LFP analyses focused on rope-pull attempts. LFP attempt epochs were defined as the 2 s preceding the rat lifting the weight (i.e., the initiation of task execution). If attempts overlapped in series (e.g., the rat pulled the weight, dropped it, and then pulled again within 2 s), only the first attempt was used. Attempts were analyzed regardless of their outcomes (successful or failed pull). If the rat did not have at least five attempts for a given weight, that weight was not included in the LFP analysis.

LFP signals were downsampled to 200 Hz before analysis using the MATLAB `downsample` function. LFP data were then analyzed using the Chronux toolbox (Bokil et al., 2010). After downsampling, LFP signals were locally detrended for low-frequency movement artifacts using the Chronux `locdetrend` function with a 500 ms moving window and 100 ms overlap. Following local detrending, the noise of the 50 Hz electrical mains was removed with the Chronux `rmlinesc` function. LFP power spectra [power spectrum densities (PSDs)] and spectrograms were then calculated via the multitaper spectral estimation method using the Chronux toolbox with a time–bandwidth product of three with five tapers, no FFT padding, and a bandpass filter of 0–100 Hz. The resulting power spectra covered 0–100 Hz in ~ 0.5 Hz increments. For a given recording session, all power spectra and spectrograms were z-scored based on 100 randomly sampled 2 s epochs from the prebaseline period. Any trial where $>15\%$ of the frequency spectrum (0–48 Hz) measured over ± 3 SDs was removed. Likewise, trials were excluded if mean power measured over ± 3 SDs within a frequency band of interest. While rare, these noisy trials tended to be caused by the rats hitting the headstage on the wall or arena floor, or from grooming behaviors before a pull attempt. Coherency, a measure of oscillatory synchronization, between the ACC and AI was analyzed using the Chronux `coherency` function. Coherency was averaged across trials, and data were pooled across sessions and rats. Spectrograms were created with the Chronux `mtspecgram` function using a time–bandwidth product of 3, five tapers, a padding factor of -1 (no padding), and a 2 s moving window with 0.05 s step. One additional second of data was added before and after the 2 s epoch of interest to compute the spectrograms. The normalized

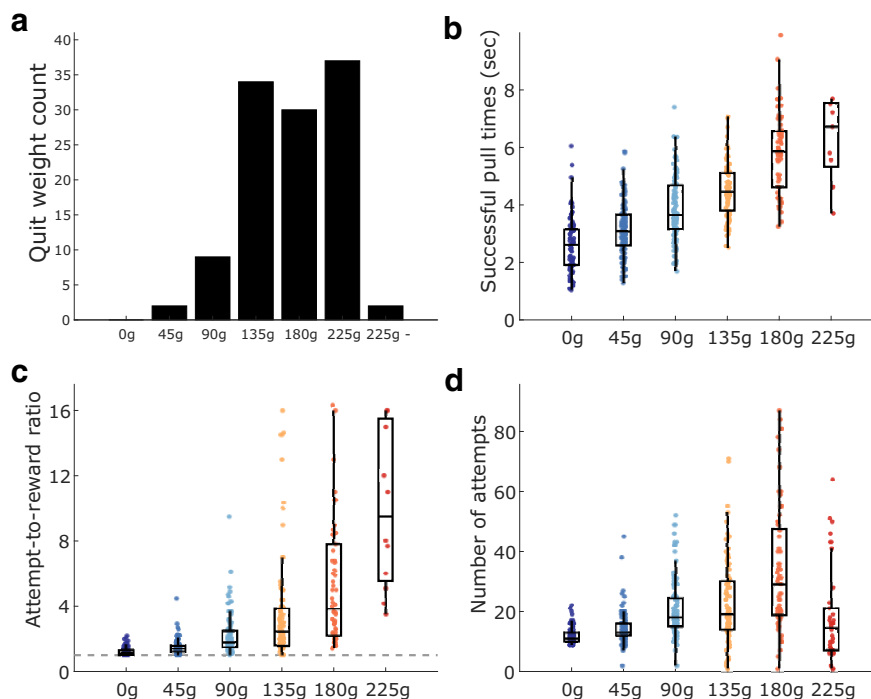


Figure 2. Progressive weight task behavior. **a**, Histogram of quit weight occurrence. “225 g -” indicates the impossible phase where no rewards could be earned; see Materials and Methods. **b**, The time it took rats to pull the weights up 30 cm from the base, sufficient to trigger sucrose reward. Heavier weights took longer to lift. **c**, The attempt-to-reward ratio for each weight showing the average number of attempts rats made before completing a successful attempt. Dotted line represents a ratio of 1.0 where every attempt would be successful. Heavier weights took more pull attempts before a reward was successfully triggered. **d**, Average number of attempts that rats made on each weight. Rats tended to make more attempts as weights got heavier; however, on heavier weights some rats made very few attempts before quitting.

LFP data from every attempt except for the first attempt within a weight block was averaged, and then data were pooled across rats and sessions. We did not include the first attempt of a weight (e.g., trial 11, trial 21) as the rats were likely still expecting the previous weight before making an attempt. However, we also cannot rule out with certainty that they were expecting the previous weight (e.g., they potentially registered the noises made by the experimenter while changing out the weight). Thus, we have taken out the first attempt of every weight block (including 0 g) to control for the uncertainty in the expectations of the rats.

One-way ANOVA was used to test whether normalized power differed within a frequency band across rope weights with a significance value of 0.05. In addition to absolute weights, we also analyzed the data in terms of the relative performance state of the rat. For each individual session, we assessed LFPs on the following: the start weight of 0 g, the highest weight the rat achieved (i.e., completed 10 trials on), and the weight the rat quit on. One-way repeated-measures ANOVA was used to test for changes in power and coherency across these three performance states within the session. Tukey’s HSD procedure was used to test for *post hoc* multiple comparisons using the MATLAB `multcompare` function.

Our graphs attempt to show all the data points, where feasible. However, we chose to programmatically limit the y-axis based on the minimum and maximum quartiles of

the data. As a result, some individual data points outside of these ranges are not shown graphically. However, for coherency analyses, we hard coded the y-axis for all frequency bands for easier visual comparisons as coherency values can only range from 0 to 1. Color templates were made using ColorBrewer (<https://colorbrewer2.org/>).

Data availability

This study was not preregistered. Data are not openly accessible but can be made available on request. Analysis was conducted with custom MATLAB scripts. All code is available on request to the corresponding author.

Results

Behavioral impact of progressive WLT

We first set out to characterize how rats behaved on the progressive WLT. Sixteen rats were tested on the progressive weight paradigm for a total of 114 sessions; each rat contributed five to eight sessions. Within a session, rats generally completed the first three weights (0, 45, and 90 g) with ease before quitting on one of the latter three weights (135, 180, or 225 g; Fig. 2a). The most common quit weight was 225 g, occurring in 33% of sessions. The average body weight of the rats was not correlated to the average weight they quit on ($r = 0.02$, $p = 0.93$). In 55% of sessions (51 of 114 sessions), rats completed at least one successful 30 cm pull on their quit weight before quitting,

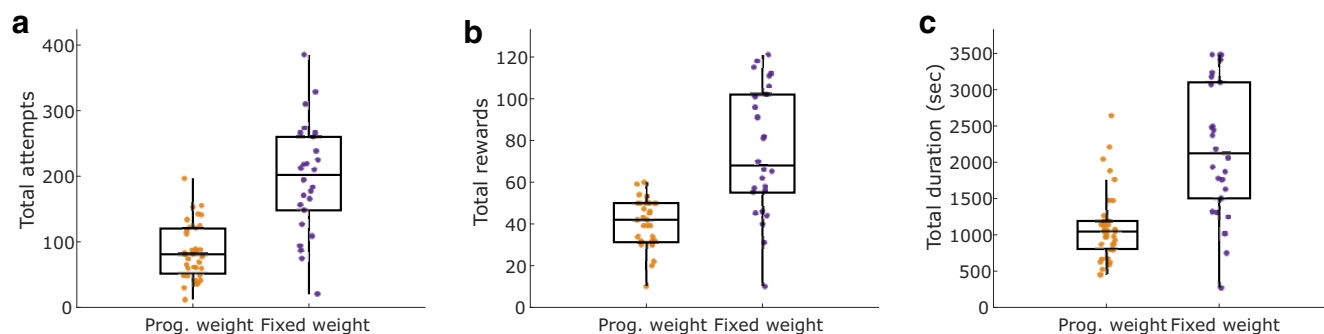


Figure 3. Behavioral comparisons between paradigms. In the progressive-weight task, 45 g of weight was added to the rope after every 10 successful pulls. In the fixed-weight task, a 180 g rope weight was accessible for up to 1 h. **a–c**, In the fixed-weight paradigm, rats made significantly more attempts (**a**), earned more rewards (**b**), and performed for longer (**c**), compared with their performance in the progressive-weight paradigm.

indicating that they were capable of successfully performing the action of lifting that weight up 30 cm but chose not to continue on. In 98% of sessions (112 of 114 sessions), rats were capable of at least lifting their quit weight off of the base. Together, the pulling behavior of the rats indicated that they likely quit the task because of the increasing effort cost caused by the increasing weight. Furthermore, in 100% of sessions rats consumed the sucrose rewards during the end-of-session satiation check, suggesting that they did not quit the rope-pulling task because of sucrose satiation.

The increasing effort load of the weights had a significant impact on the behavior of the rats. As the weight got progressively heavier, it took rats significantly longer to successfully pull the weight up 30 cm from the base ($H_{(5,469)} = 232, p < 0.0001$; Fig. 2*b*). Pairwise comparisons between each weight and its preceding weight revealed significant increases in successful pull times across most weight pairs (p values < 0.02), however, pull times between 0 and 45 g ($p = 0.06$) and 180 and 225 g ($p = 1.0$) were similar.

The increase in successful pull time with heavier weights indicates that the heavier weights are indeed more effortful for the rats. Successful pull times, however, can only be calculated for rewarded attempts and therefore do not fully capture the persistence of the rats at the task as it becomes increasingly difficult. To characterize the persistence of the rats, we calculated an attempts-to-reward ratio of the average number of attempts made for each rewarded attempt for each weight. A ratio of 1.0 indicates that every attempted pull was successful in triggering reward. Ratios > 1.0 indicate that the rats failed to pull the weight 30 cm on some attempts. Rats made significantly more attempts per reward as the weights increased ($H_{(5,482)} = 268, p < 0.0001$; Fig. 2*c*). Pairwise comparisons revealed that rats generally made more attempts per reward as the weights got heavier (all p values < 0.02), however, their performance was similar between 90 and 135 g ($p = 0.36$) and 180 and 225 g ($p = 0.93$). Thus, as the weights increase, rats fail more often to successfully lift the weight.

On 69% of sessions where rats reached the 225 g weight, they were not able to complete a single successful, rewarded pull (and thus we could not compute an attempt-to-

reward ratio). To get a better idea of the persistence of the rats in the face of increasing effort demands, we measured the total number of attempts rats made on each weight. Weight had a significant impact on the number of attempts that rats made ($H_{(5,543)} = 169, p < 0.0001$; Fig. 2*d*). The average number of attempts increased weight-to-weight up to 180 g ($p < 0.05$), with the exception of 90–135 g ($p = 1.0$), until 225 g, where attempts significantly decreased on 225 g compared with 180 g ($p < 0.0001$). Overall, these behavioral metrics indicate that the progressively increasing weights are more and more difficult for the rats to lift, requiring more effort (pull time, number of attempts) and persistence (attempt-to-reward ratio).

To further ensure that rats were not quitting due to finishing a specific number of pulls, sucrose satiation, or time-on-task, we ran five rats on a modified, fixed-weight version of the WLT in the week after they ran the progressive weight paradigm. Rather than the weight progressively increasing, rats had 10 trials of 0 g followed by unlimited trials on 180 g. The 180 g weight was chosen as it was the maximal weight that all five rats had completed 10 trials on in the progressive weight paradigm. Rats ran the fixed-weight paradigm until they quit (no attempt for 2 min) or for 1 h, whichever came first. Each rat completed five or six sessions of the fixed-weight paradigm; 26 sessions were recorded in total. Compared with their progressive weight performance, during the fixed-weight task the rats made significantly more attempts (median, 202 attempts; vs 81 attempts in the progressive weight; $z = -5.11, p < 0.0001$; Fig. 3*a*), earned significantly more rewards (68 vs 42 in the progressive weight; $z = -4.73, p < 0.0001$; Fig. 3*b*), and performed the task for significantly longer before quitting (2123 vs 1045 in the progressive weight; $z = -4.62, p < 0.0001$; Fig. 3*c*). Overall, these data further demonstrate that rats are quitting the progressive weight task because of the increase in effort (weights) and not because of repetitive rope-pulling attempts, becoming satiated on sucrose, or quitting after an elapsed time.

LFP correlates of progressive effort loading

We investigated the neural correlates of increasing effort costs by recording LFPs in the ACC and AI as rats

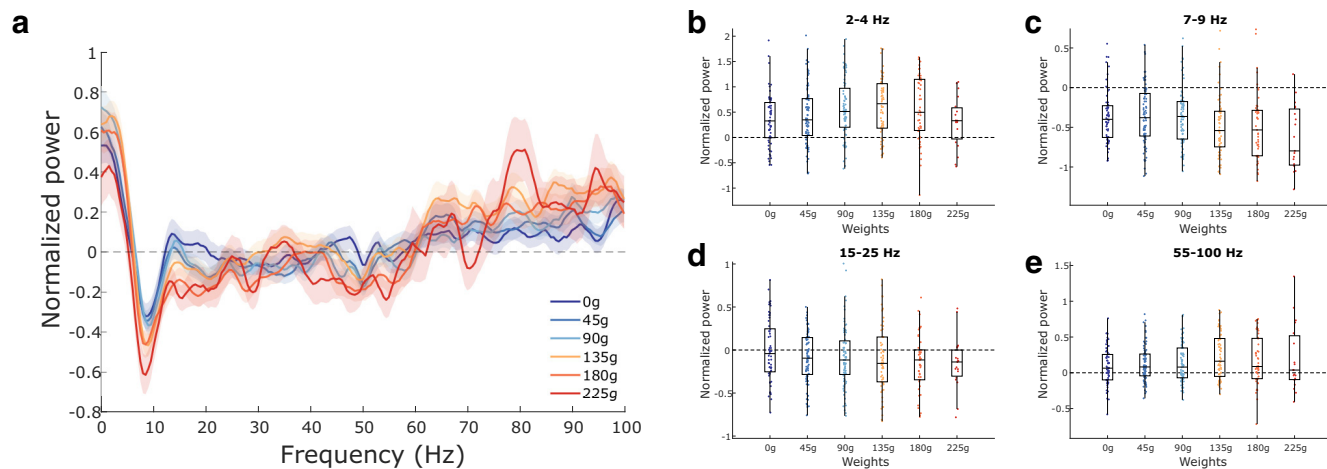


Figure 4. Changes in ACC LFP rhythms with increasing effort. **a**, Multitaper power spectrum of ACC LFP across the six different weights. Shaded error bars represent ± 1 SEM. **b–e**, Normalized power across the four rhythms of interest (delta, theta, beta, and gamma, respectively). Box plots show first, second (median; middle line), and third quartiles. Whiskers are 1.5 times the interquartile range. All power calculations were normalized using the 2 min prebaseline period (see Materials and Methods).

conducted the progressive WLT. We specifically focused our analyses on the 2 s before the rat lifting the weight up. This period reflects the rat choosing to actively engage in the task and make an attempt at expending energy to earn a reward. Changes in power in the delta (2–4 Hz), theta (7–9 Hz), beta (15–25 Hz), and gamma (55–100 Hz) frequency ranges were assessed. We predicted that low-frequency LFP power from the ACC would inversely correlate to effort costs, while low-frequency LFP power from the AI would positively correlate to effort costs.

The overall normalized ACC PSD for each weight is shown in Figure 4a. Weight had a significant influence on delta-band ($F_{(5,350)} = 2.75$, $p = 0.019$; Fig. 4b) and theta-band power ($F_{(5,350)} = 2.97$, $p = 0.012$; Fig. 4c). *Post hoc* multiple comparisons on delta and theta power using Tukey's test revealed no significant pairwise differences

(p values > 0.05). Beta ($F_{(5,350)} = 1.32$, $p = 0.257$; Fig. 4d) and gamma ($F_{(5,350)} = 1.89$, $p = 0.095$; Fig. 4e) power were not significantly impacted by weight.

The overall normalized AI PSD for each weight is shown in Figure 5a. Weight had no effect on power within the delta band ($F_{(5,352)} = 0.97$, $p = 0.44$; Fig. 5b) or beta band ($F_{(5,352)} = 1.23$, $p = 0.29$; Fig. 5d). Weight did have a significant effect on power within the theta band ($F_{(5,352)} = 2.38$, $p = 0.04$; Fig. 5c) and gamma band ($F_{(5,352)} = 4.80$, $p = 0.0003$; Fig. 5e). No *post hoc* comparisons were significant for theta power (all p values > 0.05). *Post hoc* pairwise comparisons revealed significant differences between gamma power on 0 g compared with the other weights (p values < 0.05), with the exception of 225 g ($p = 0.054$). No other pairwise comparisons were significant in the gamma band (p values > 0.05).

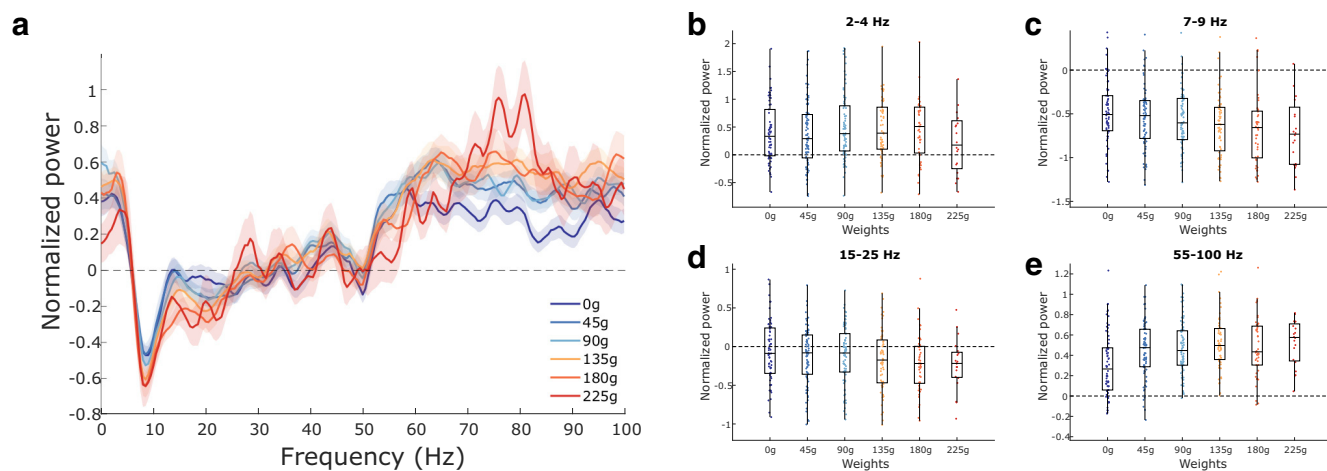


Figure 5. Changes in AI LFP rhythms with increasing effort. **a**, Multitaper power spectrum of AI LFP across the six different weights. Shaded error bars represent ± 1 SEM. **b–e**, Normalized power across the four rhythms of interest (delta, theta, beta, and gamma, respectively). Box plots show first, second (median; middle line), and third quartiles. Whiskers are 1.5 times the interquartile range. All power calculations were normalized using the 2 min prebaseline period.

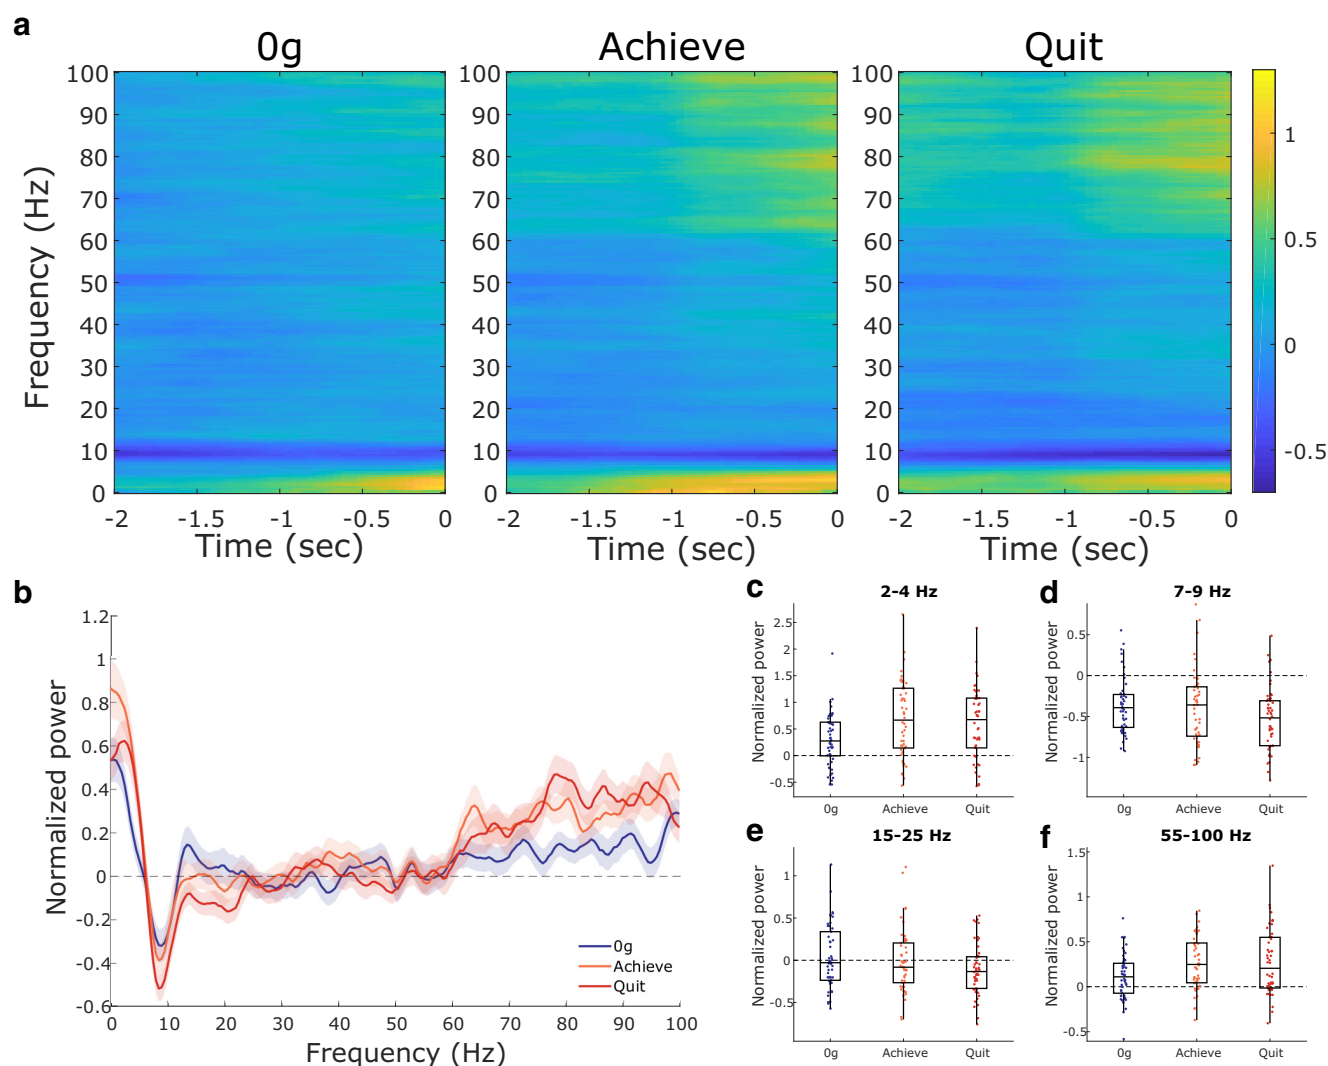


Figure 6. Changes in ACC LFP rhythms across performance states. **a**, Average ACC spectrograms for the three performance states showing mean normalized power in the 2 s before the attempt. **b–f**, The 0 g weight (blue). The achievement weight (orange; **b–f**) is the last complete (10 successful pulls) weight, while the quit weight (red; **b–f**) is the weight the rats quit on. **b–f**, Multitaper power spectrum of ACC LFP across the three performance states. Shaded error bars are ± 1 SEM. **c–f**, Normalized power across the four rhythms of interest (delta, theta, beta, and gamma, respectively). Box plots show first, second (median; middle line), and third quartiles. Whiskers are 1.5 times the interquartile range. All power calculations were normalized using the prebaseline period.

LFP correlates of quitting behavior

In addition to examining how increases in effort load (weight) affected regional activity in the ACC and AI, we were interested in examining how the performance state of an animal correlated with regional neural activity. As the WLT is a self-driven task, the animal can quit at any point, on any weight. This generates interindividual and intraindividual variability in quit points, and quit point metrics, likely reflecting variability in motivational states. Because of the quitting behaviors of the rats, the motivational state of a rat on a given weight could be different from that of another rat on the same weight. For example, one rat may quickly complete 10 successful pulls on the 135 g weight, while another rat may quit on 135 g after only a few attempts. To address session-to-session variability, we divided up the performance of the rats based

on their relative achievement for a given session. For each individual session, we assessed LFPs from the following three blocks of interest: (1) we analyzed the specific frequency bands during the low-effort start weight of 0 g; (2) the achievement weight of the rat (i.e., the maximum weight that the rat completed 10 trials on); and (3) the weight the rat quit on.

The decision to quit is likely a dynamic process, occurring as animals assess their current internal states with the costs and benefits of the currently available task. We were interested in visualizing the LFP temporal dynamics as rats chose to carry out their attempts across the three performance states (0 g, achievement weight, and quit weight). Figure 6a shows the averaged normalized spectrograms across all recordings of the three performance states. Prominent increases in theta and gamma power

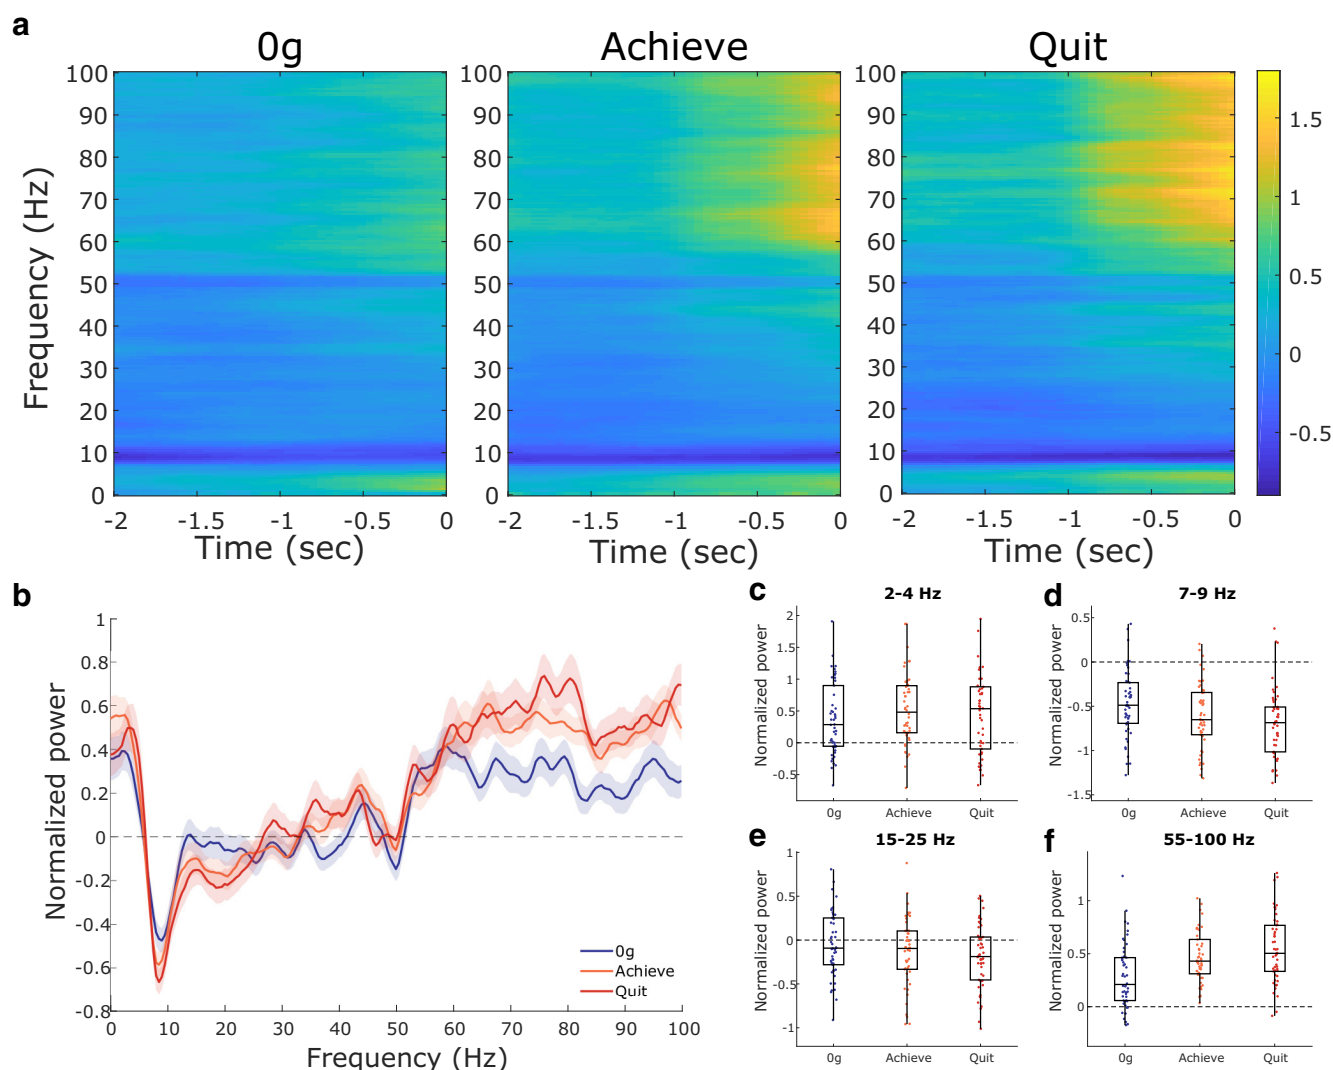


Figure 7. Changes in AI LFP rhythms across performance states. All power in normalized based on the prebaseline period. **a**, Average AI spectrograms for the three performance states showing the mean normalized power in the 2 s before the attempt. **b–f**, The 0 g weight (blue). The achievement weight (orange; **b–f**) is the last complete (10 successful pulls) weight, while the quit weight (red; **b–f**) is the weight the rats quit on. **b**, Multitaper power spectrum of AI LFPs across the three motivational states. Shaded error bars are ± 1 SEM. **c–f**, Normalized power across the four rhythms of interest (delta, theta, beta, and gamma, respectively). Box plots show first, second (median; middle line), and third quartiles. Whiskers are 1.5 times the interquartile range.

can be seen in the achievement and quit weights compared with the 0 g weight. In contrast, theta appears to decrease in power across performance states. The overall normalized ACC PSD across the three states can be seen in Figure 6b. When each frequency band of interest was independently assessed, significant differences in power were observed between performance states in the delta rhythm ($F_{(1.96,95.9)} = 16.8$, $p < 0.0001$; Fig. 6c), theta rhythm ($F_{(1.64,80.1)} = 5.45$, $p = 0.01$; Fig. 6d), beta rhythm ($F_{(1.58,77.3)} = 4.22$, $p = 0.026$; Fig. 6e), and gamma rhythm ($F_{(1.53,74.8)} = 11.49$, $p = 0.0002$; Fig. 6f). *Post hoc* comparisons revealed significant increases in ACC delta and gamma power between the 0 g and the achievement weights ($p < 0.0001$ and $p = 0.0002$, respectively). Both theta and beta power significantly decreased from the achievement weight to the quit weight ($p = 0.03$ and

$p = 0.02$). Every brain rhythm showed a significant difference from 0 g to the quit weight where theta ($p = 0.02$) and beta ($p = 0.03$) power decreased and delta ($p = 0.0003$) and gamma ($p = 0.004$) power increased.

Figure 7a shows the average normalized AI spectrograms from 0 g, achievement weight, and quit weight. Clear changes from the LFP baseline are apparent at the delta, theta, and gamma rhythms. These shifts in power seem to be most prominent just before the rat pulling the rope. The overall normalized AI PSD across performance states is shown in Figure 7b. Power in the delta frequency range ($F_{(1.73,84.8)} = 2.1$, $p = 0.14$; Fig. 7c) did not change significantly across the three performance states. In contrast, AI theta ($F_{(1.73,84.8)} = 17.5$, $p < 0.0001$; Fig. 7d), beta ($F_{(1.73,84.8)} = 5.55$, $p = 0.0077$; Fig. 7e), and gamma rhythms ($F_{(1.74,85.4)} = 19.8$, $p < 0.0001$; Fig. 7f) did show

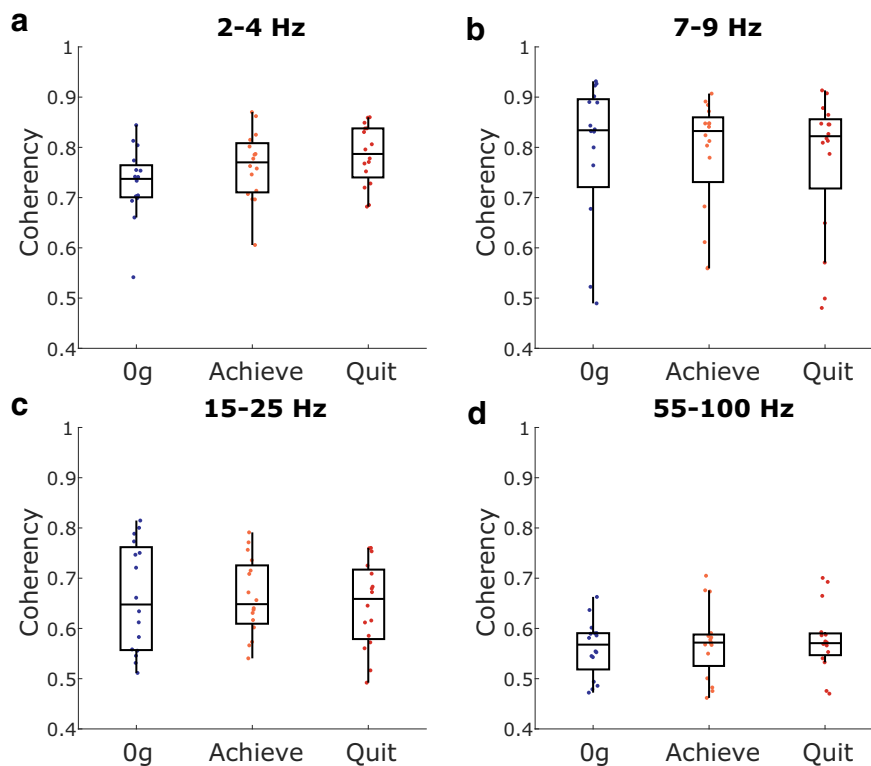


Figure 8. ACC–AI coherency measures. **a–d**, Coherency changes across performance states in the delta (**a**), theta (**b**), beta (**c**), and gamma (**d**) rhythms. Box plots show first, second (median; middle line), and third quartiles. Whiskers are 1.5 times the interquartile range.

significant differences across performance states. *Post hoc* pairwise tests revealed a significant decrease in theta power from 0 g to the achievement weight ($p = 0.0007$) and quit weight ($p < 0.0001$) as well as from the achievement weight to the quit weight ($p = 0.042$). Beta power decreased significantly only from 0 g to the quit weight ($p = 0.009$). In contrast, gamma power increased from both 0 g to the achievement weight ($p < 0.0001$) and from 0 g to the quit weight ($p < 0.0001$).

ACC–AI coherency shifts across performance states

The activity of distant brain regions can become synchronous during specific behaviors, possibly reflecting task-specific information processing in both regions (Harris and Gordon, 2015). We were interested in whether or not ACC–AI coherency increased as the animal got closer to quitting, with the ACC evaluating effort costs in the context of the interoceptive information of AI. To determine whether ACC–AI coherency changes occurred across WLT performance states, we assessed coherency during the 2 s attempt epoch across our three states of interest, as follows: 0 g, achievement weight, and quit weight. Coherency in the delta range showed a significant effect across states ($F_{(1.82,27.4)} = 11.38$, $p = 0.0004$; Fig. 8a). *Post hoc* testing showed coherence significantly increased from 0 g to the achievement weight ($p = 0.03$) and quit weight ($p = 0.002$). There were no significant changes in ACC–AI coherency in the theta ($p = 0.67$; Fig. 8b), beta ($p = 0.38$; Fig. 8c), or

gamma ($p = 0.10$; Fig. 8d) rhythms when performance states were compared.

Discussion

Behaviorally, this study demonstrates the efficacy of our novel WLT for probing effort exertion in laboratory rats. Progressively increasing the weight that the rats must lift to earn a fixed reward resulted in distinct behavioral shifts, reflective of increasing effort cost (Fig. 2). Heavier weights (≥ 135 g) proved to be particularly effortful, and resulted in high rates of quitting. As the reward amount was held constant, the observed quitting behaviors were most likely because of a decrease in perceived task utility driven by the increased physical effort costs. Time expenditure in the task was not likely to account for the quitting behaviors we observed, as experiments using a fixed-weight paradigm demonstrated that some rats will rope pull 180 g for up to 1 h (Fig. 3c). Motivated sucrose intake at the end of the session (sucrose satiation check; see Materials and Methods) and rats consuming significant more sucrose on the fixed-weight paradigm (Fig. 3b) further support the interpretation that quitting the progressive WLT was associated with mounting task costs, rather than reward devaluation because of satiation.

Only male rats were used in this study. However, previous research has demonstrated sex differences in decision-making behaviors (Orsini and Setlow, 2017). For example, Orsini et al. (2016) used a risky decision-making task and found that male rats were more risk seeking than

female rats. While no direct sex difference comparisons have been conducted on effort discounting in rats, [Uban et al. \(2012\)](#) demonstrated that estrogen receptor agonists shift behavioral preferences away from high-cost, high-reward choices toward low-cost, low-reward choices. As our present study only used male rats, our results may not generalize to female rats on the WLT.

Neurophysiologically, we predicted that low-frequency LFP power from the ACC would inversely correlate to effort costs, while low-frequency LFP power from the AI would positively correlate to effort costs. When the progressive WLT was examined based on absolute weight categories, our hypotheses were only partially supported. In the ACC, both delta and theta power had a main effect of weight but no *post hoc* tests were significant. Thus, we cannot conclude whether increasing effort costs increases low-frequency rhythms in the ACC. In contrast, AI LFP power did not show a progressive increase in any of the low-frequency bands examined, as we had predicted. There was a suspected decrement in theta power in response to rope weight on first visual inspection ([Fig. 5a](#)), but this was not significant. Regional gamma power in the AI did significantly increase from 0 g to the heavier rope weights, notably in the 75–85 Hz range.

While we found changes in ACC delta and theta bands along with the AI gamma band across the six weight conditions of the WLT session, one shortcoming of analyzing the six weight conditions as a single factor is that there is considerable behavioral variability occurring, notably on the higher weights. For example, 135 g may be a “midway” weight for rat A, who proceeds all the way to the 225 g in a given session, but 135 g may be the “quit point” for rat B in a given session. Accordingly, the neural states would likely be very different between rat A and rat B in those sessions, despite the 135 g condition being nominally the same. For this reason, we performed a secondary analysis on the dataset on the basis of performance state rather than specific weight.

Performance state analysis demonstrated a wider range of impacts across multiple neural rhythms. In the ACC, delta power was lower during the initial 0 g weight compared with the achievement and quit weights, and there was an increase in delta rhythm coherency between the ACC and AI across performance states. In general, slower rhythms are thought to reflect network-level coordination across many different brain regions ([von Stein and Sarnthein, 2000](#); [Buzsáki and Draguhn, 2004](#)). The increase in coherency on the challenging weights (achieve and quit) compared with the 0 g weight may reflect the increased network activity within the salience network as the WLT becomes increasingly demanding ([Uddin, 2015](#)). Furthermore, delta rhythms have been shown to coordinate prefrontal regions with VTA dopamine circuitry ([Fujisawa and Buzsáki, 2011](#); [Elston and Bilkey, 2017](#)). Dopamine may help drive continued task performance in the WLT in the face of increasing effort conditions by biasing ACC networks to discount effort costs and/or amplify reward benefits ([Schweimer and Hauber, 2006](#)).

In both the ACC and AI, theta and gamma power showed significant but opposing changes in relation to

performance state, as follows: regional theta power decreased and gamma power increased as the animals progressed toward their quitting point for that session ([Figs. 6, 7](#)). In both the ACC and AI, theta power dropped significantly on the quit weight compared with 0 g and the achievement weight. The decrease in theta power on the quit weight may reflect the shift in task utility and the task no longer being “worth-it” compared with the 0 g and achievement weights where utility is high. In contrast, the gamma power of both brain regions showed a significant increase from task onset (0 g) to achievement weight and quit weight. Thus, gamma power appears to become elevated earlier on in the task (achievement weight) and stay elevated relative to 0 g. In general, the elevated gamma power may reflect the increased local neuronal activity as the task becomes more challenging. Overall, these shifts in theta and gamma power may reflect changing motivational states as the rat approaches the decision to quit the task.

Previous studies investigating the neural correlates of effortful or persistence-like behaviors have tended to run subjects for a fixed time or for a fixed number of trials ([Croxson et al., 2009](#); [Hillman and Bilkey, 2012](#); [Engström et al., 2014](#); [McGuire and Kable, 2015](#); [Hart et al., 2017](#)). However, these studies do not capture what occurs in the brain when subjects voluntarily quit a task. During physically demanding tasks, ACC single-unit population activity in rodents and nonhuman primates has been shown to reflect the current value of a given behavior ([Kennerley et al., 2009](#); [Cowen et al., 2012](#); [Porter et al., 2019](#)). Our LFP data are in line with these findings, showing that in the rat ACC and AI, theta power decreases from the start of the task (0 g) to the self-selected end of the task (achievement weight and quit weight). Theta rhythm may reflect an overall net utility of the action at hand and, relatedly, the motivation to continue performing the task or to quit. [Elston and Bilkey \(2017\)](#) found a similar result in rats using a fixed-reward, jumpable-barrier effort task, as follows: ACC theta power increased on low-effort (no barrier) trials and decreased on high-effort trials (barrier present). Overall, ACC theta power may reflect the current value of the behavior being conducted and the cognitive control exerted by the ACC to bias behavior toward particular actions ([Womelsdorf et al., 2010](#)).

These rat-based observations and our interpretation as to why theta power decreases across the progressive WLT may initially appear at odds with human data. In EEG studies, frontal midline theta reliably increases as task difficulty and/or duration increases ([Paus et al., 1997](#); [Smit et al., 2005](#); [Mitchell et al., 2008](#)). One possible concordant interpretation links back to the notion of running subjects for a fixed duration versus allowing subjects to self-select the end point. If theta power in the ACC is reflective of the utility of a current action, theta power progressively decreases in our rat task as the task gets difficult, yet the reward amount remains the same; utility progressively declines. Human subjects, however, have additional factors feeding into cost–benefit utility calculations and these include completion goals. Social contract pressure to not “give up” or quit during a research task of fixed duration—

particularly if under researcher supervision—would increase the utility of performing and persisting in difficult tasks. Likewise, nearing the end of a session of known duration may confer lowered cost, and/or added benefit to task trials, lending heightened utility to trials proximal to the end point. Such factors may in part account for why midline frontal theta increases in human EEG studies and would be an interesting area for future investigation.

In contrast to the decrement in theta power across performance states, power in the gamma band in the ACC and AI increases across performance states. Cortical gamma has been proposed to be generated by interneurons and the need to balance excitation and inhibition during periods of activation (Bartos et al., 2007; Merker, 2016). Increased gamma power in our WLT may reflect the increasing cognitive processing demands of the task; mounting costs must be dynamically integrated with interoceptive signals, such as hunger and fatigue, to compute ongoing utility. For example, gamma rhythms can selectively amplify or dampen incoming signals dependent on when in the gamma cycle the information arrives (Fries et al., 2001; Sohal et al., 2009). Such a mechanism could reflect the increased information processing needed as the task becomes more challenging. Indeed, gamma power in the ACC and AI shows an increase in power during the achievement and quit weights compared with the easiest, and highest utility, 0 g weight. This may be reflective of the role of AI in integrating the physical cost of behaviors, and in particular the negative affective aspects of those physical costs, with the dynamic changes in internal states as the rats approach their quitting point (Williamson et al., 1999; Prévost et al., 2010).

Similarly, the heightened gamma power of the ACC alongside the increasing effort costs may be reflective of its role in effort-based cost–benefit decision-making (Kennerley et al., 2006; Croxson et al., 2009; Klein-Flügge et al., 2016) and cognitive control (Lorist et al., 2005; Woodward et al., 2008; Holroyd and McClure, 2015). For example, foraging tasks in humans (Kolling et al., 2012) and nonhuman primates (Hayden et al., 2011) demonstrate increasing ACC activity as subjects approach the decision to search for another foraging patch. While subjects in these foraging-based tasks do not quit the testing session, these data provide evidence that the ACC may determine the value of “quitting” the current behavioral option. Since the gamma rhythm is thought to reflect local circuit activity and gate incoming information (Fries et al., 2001; Sohal et al., 2009), the shifts in gamma activity across the task may reflect the ongoing integration of many sources of information (e.g., fatigue, satiation, reward history, current effort cost) to determine whether or not the task is worth persisting at or quitting.

Our WLT demonstrates that the behavior of the rat is altered as effort costs increase while reward stays constant. As shown in Figure 2, rats struggle to lift heavier weights, fail more often on their attempts, and, as result of failing more often, make more attempts as the weight increases. These factors likely contribute to the choice of the rat to quit and may also contribute to the changes in LFP power we report here. For example, because of the rats failing

more often, they earn rewards at a slower rate. These failed attempts may also mean that the rat perceives the task as becoming more volatile as their attempts become less likely to result in a reward. Both reward history (Kolling et al., 2016) and task volatility (Behrens et al., 2007) have been shown to be represented by the ACC.

Both ACC and AI activity changes may likewise be altered by the increasingly infrequent sucrose consumption (Horst and Laubach, 2013; Becker et al., 2017). In particular, a proposed role of the beta rhythm is in the maintenance of the status quo in that the current sensorimotor plan is expected to be maintained (Engel and Fries, 2010). However, as the task becomes more difficult and rewards more infrequent, and fatigue sets in, the status quo of carrying out the WLT may be challenged. Indeed, we observed a decrease in beta power on the quit weight compared with 0 g in both brain regions. Thus, factors associated with the changes in weight, changes in behavior, and changes in task outcome may be contributing to the observed shifts in ACC and AI LFPs.

Here we have characterized LFP activity in the ACC and AI during a physical effort-based task, where rats have the choice to engage in the task or quit the session. Motivational performance state, rather than absolute effort cost (weight), appears to better capture the dynamic changes in ACC and AI LFP power observed across the session. As rats approach the decision to quit, ACC and AI theta power decreases while gamma power increases. Our results may guide future effort-based cost–benefit decision studies in targeting the source and content of these rhythms.

References

- Amiez C, Joseph JP, Procyk E (2006) Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex* 16:1040–1055.
- Balleine BW, Dickinson A (2000) The effect of lesions of the insular cortex on instrumental conditioning: evidence for a role in incentive memory. *J Neurosci* 20:8954–8964.
- Bartos M, Vida I, Jonas P (2007) Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nat Rev Neurosci* 8:45–56.
- Becker CA, Flaisch T, Renner B, Schupp HT (2017) From thirst to satiety: the anterior mid-cingulate cortex and right posterior insula indicate dynamic changes in incentive value. *Front Hum Neurosci* 11:234.
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF (2007) Learning the value of information in an uncertain world. *Nat Neurosci* 10:1214–1221.
- Bokil H, Andrews P, Kulkarni JE, Mehta S, Mitra PP (2010) Chronux: a platform for analyzing neural signals. *J Neurosci Methods* 192:146–151.
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929.
- Cavanagh JF, Frank MJ (2014) Frontal theta as a mechanism for cognitive control. *Trends Cogn Sci* 18:414–421.
- Cechetti DF, Saper CB (1987) Evidence for a viscerotopic sensory representation in the cortex and thalamus in the rat. *J Comp Neurol* 262:27–45.
- Cowen SL, Davis GA, Nitz DA (2012) Anterior cingulate neurons in the rat map anticipated effort and reward to their associated action sequences. *J Neurophysiol* 107:2393–2407.
- Craig AD (2003) Interoception: the sense of the physiological condition of the body. *Curr Opin Neurobiol* 13:500–505.
- Critchley HD (2005) Neural mechanisms of autonomic, affective, and cognitive integration. *J Comp Neurol* 493:154–166.

- Crosson PL, Walton ME, O'Reilly JX, Behrens TE, Rushworth MF (2009) Effort-based cost-benefit valuation and the human brain. *J Neurosci* 29:4531–4541.
- Elston TW, Bilkey DK (2017) Anterior cingulate cortex modulation of the ventral tegmental area in an effort task. *Cell Rep* 19:2220–2230.
- Engel AK, Fries P (2010) Beta-band oscillations—signalling the status quo? *Curr Opin Neurobiol* 20:156–165.
- Engström M, Karlsson T, Landtblom AM, Craig AD (2014) Evidence of conjoint activation of the anterior insular and cingulate cortices during effortful tasks. *Front Hum Neurosci* 8:1071.
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563.
- Fujisawa S, Buzsáki G (2011) A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. *Neuron* 72:153–165.
- Harris AZ, Gordon JA (2015) Long-range neural synchrony in behavior. *Annu Rev Neurosci* 38:171–194.
- Hart EE, Gerson JO, Zoken Y, Garcia M, Izquierdo A (2017) Anterior cingulate cortex supports effort allocation towards a qualitatively preferred option. *Eur J Neurosci* 46:1682–1688.
- Hart EE, Blair GJ, O'Dell TJ, Blair HT, Izquierdo A (2019) Anterior cingulate cortex activity regulates effort-based decision making. *bioRxiv*. Advance online publication. Retrieved October 4, 2019. doi:10.1101/792069.
- Hayden BY, Pearson JM, Platt ML (2011) Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14:933–939.
- Hillman KL, Bilkey DK (2010) Neurons in the rat anterior cingulate cortex dynamically encode cost-benefit in a spatial decision-making task. *J Neurosci* 30:7705–7713.
- Hillman KL, Bilkey DK (2012) Neural encoding of competitive effort in the anterior cingulate cortex. *Nat Neurosci* 15:1290–1297.
- Holroyd CB, McClure SM (2015) Hierarchical control over effortful behavior by rodent medial frontal cortex: a computational model. *Psychol Rev* 122:54–83.
- Horst N, Laubach M (2013) Reward-related activity in the medial prefrontal cortex is driven by consumption. *Front Neurosci* 7:56.
- Kennerley SW, Wallis JD (2009) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci* 29:2061–2073.
- Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF (2006) Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 9:940–947.
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21:1162–1178.
- Klein-Flügge MC, Kennerley SW, Friston K, Bestmann S (2016) Neural signatures of value comparison in human cingulate cortex during decisions requiring an effort-reward trade-off. *J Neurosci* 36:10002–10015.
- Kolling N, Behrens TE, Mars RB, Rushworth MF (2012) Neural mechanisms of foraging. *Science* 336:95–98.
- Kolling N, Wittmann MK, Behrens TE, Boorman ED, Mars RB, Rushworth MF (2016) Value, search, persistence and model updating in anterior cingulate cortex. *Nat Neurosci* 19:1280–1285.
- Laubach M, Amarante LM, Swanson K, White SR (2018) What, if anything, is rodent prefrontal cortex? *eNeuro* 5:ENEURO.0315-18.2018.
- Lorist MM, Boksem MA, Ridderinkhof KR (2005) Impaired cognitive control and reduced cingulate activity during mental fatigue. *Brain Res Cogn Brain Res* 24:199–205.
- McGuire JT, Kable JW (2015) Medial prefrontal cortical activity reflects dynamic re-evaluation during voluntary persistence. *Nat Neurosci* 18:760–766.
- Medford N, Critchley HD (2010) Conjoint activity of anterior insular and anterior cingulate cortex: awareness and response. *Brain Struct Funct* 214:535–549.
- Merker BH (2016) Cortical gamma oscillations: details of their genesis preclude a role in cognition. *Front Comput Neurosci* 10:78.
- Mitchell DJ, McNaughton N, Flanagan D, Kirk IJ (2008) Frontal-midline theta from the perspective of hippocampal “theta”. *Prog Neurobiol* 86:156–185.
- Moisset X, Bouhassira D, Denis D, Dominique G, Benoit C, Sabaté JM (2010) Anatomical connections between brain areas activated during rectal distension in healthy volunteers: a visceral pain network. *Eur J Pain* 14:142–148.
- Moschak TM, Wang X, Carelli RM (2018) A neuronal ensemble in the rostral agranular insula tracks cocaine-induced devaluation of natural reward and predicts cocaine seeking. *J Neurosci* 38:8463–8472.
- Orsini CA, Setlow B (2017) Sex differences in animal models of decision making. *J Neurosci Res* 95:260–269.
- Orsini CA, Willis ML, Gilbert RJ, Bizon JL, Setlow B (2016) Sex differences in a rat model of risky decision making. *Behav Neurosci* 130:50–61.
- Parkes SL, Bradfield LA, Balleine BW (2015) Interaction of insular cortex and ventral striatum mediates the effect of incentive memory on choice between goal-directed actions. *J Neurosci* 35:6464–6471.
- Paus T, Zatorre RJ, Hofle N, Caramanos D, Gotman J, Petrides M, Evans AC (1997) Time-related changes in neural systems underlying attention and arousal during the performance of an auditory vigilance task. *J Cogn Neurosci* 9:392–408.
- Paxinos G, Watson C (2007) The rat brain in stereotaxic coordinates. London: Academic.
- Porter B, Hillman KL (2019) A novel weight lifting task for investigating effort and persistence in rats. *Front Behav Neurosci* 13:275.
- Porter BS, Hillman KL, Bilkey DK (2019) Anterior cingulate cortex encoding of effortful behavior. *J Neurophysiol* 121:701–714.
- Prévost C, Pessiglione M, Météreau E, Cléry-Melin ML, Dreher JC (2010) Separate valuation subsystems for delay and effort decision costs. *J Neurosci* 30:14080–14090.
- Salenius S, Hari R (2003) Synchronous cortical oscillatory activity during motor action. *Curr Opin Neurobiol* 13:678–684.
- Schweimer J, Hauber W (2006) Dopamine D1 receptors in the anterior cingulate cortex regulate effort-based decision making. *Learn Mem* 13:777–782.
- Smit AS, Eling PA, Hopman MT, Coenen AM (2005) Mental and physical effort affect vigilance differently. *Int J Psychophysiol* 57:211–217.
- Sohal VS, Zhang F, Yizhar O, Deisseroth K (2009) Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature* 459:698–702.
- Sridharan D, Levitin DJ, Menon V (2008) A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proc Natl Acad Sci U S A* 105:12569–12574.
- Taylor KS, Seminowicz DA, Davis KD (2009) Two systems of resting state connectivity between the insula and cingulate cortex. *Hum Brain Mapp* 30:2731–2745.
- Uban KA, Rummel J, Floresco SB, Galea LA (2012) Estradiol modulates effort-based decision making in female rats. *Neuropsychopharmacology* 37:390–401.
- Uddin LQ (2015) Salience processing and insular cortical function and dysfunction. *Nat Rev Neurosci* 16:55–61.
- von Stein A, Sarnthein J (2000) Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *Int J Psychophysiol* 38:301–313.
- Williamson JW, McColl R, Mathews D, Ginsburg M, Mitchell JH (1999) Activation of the insular cortex is affected by the intensity of exercise. *J Appl Physiol* 87:1213–1219.
- Womelsdorf T, Johnston K, Vinck M, Everling S (2010) Theta-activity in anterior cingulate cortex predicts task rules and their adjustments following errors. *Proc Natl Acad Sci U S A* 107:5248–5253.
- Woodward TS, Metzack PD, Meier B, Holroyd CB (2008) Anterior cingulate cortex signals the requirement to break inertia when switching tasks: a study of the bivalency effect. *Neuroimage* 40:1311–1318.