

Sensitivity of Locus Ceruleus Neurons to Reward Value for Goal-Directed Actions

Sebastien Bouret^{1,2} and Barry J. Richmond¹

¹Laboratory of Neuropsychology, National Institute of Mental Health, National Institutes of Health, Department of Health and Human Services, Bethesda, Maryland 20892, and ²Motivation, Brain Behavior (MBB) team, Institut du Cerveau et de la Moelle épinière (ICM), Groupe Hospitalier Pitié-Salpêtrière, 75013 Paris, France

The noradrenergic nucleus locus ceruleus (LC) is associated classically with arousal and attention. Recent data suggest that it might also play a role in motivation. To study how LC neuronal responses are related to motivational intensity, we recorded 121 single neurons from two monkeys while reward size (one, two, or four drops) and the manner of obtaining reward (passive vs active) were both manipulated. The monkeys received reward under three conditions: (1) releasing a bar when a visual target changed color; (2) passively holding a bar; or (3) touching and releasing a bar. In the first two conditions, a visual cue indicated the size of the upcoming reward, and, in the third, the reward was constant through each block of 25 trials. Performance levels and lipping intensity (an appetitive behavior) both showed that the monkeys' motivation in the task was related to the predicted reward size. In conditions 1 and 2, LC neurons were activated phasically in relation to cue onset, and this activation strengthened with increasing expected reward size. In conditions 1 and 3, LC neurons were activated before the bar-release action, and the activation weakened with increasing expected reward size but only in task 1. These effects evolved as monkeys progressed through behavioral sessions, because increasing fatigue and satiety presumably progressively decreased the value of the upcoming reward. These data indicate that LC neurons integrate motivationally relevant information: both external cues and internal drives. The LC might provide the impetus to act when the predicted outcome value is low.

Key words: decision-making; locus ceruleus; motivation; neuromodulation; noradrenaline; reward

Introduction

The two biochemically similar neuromodulators, noradrenaline and dopamine, are speculated to have two different roles related to behavior. The dopaminergic system is related strongly to the incentive effects of reward (Liu et al., 2004; Pessiglione et al., 2006; Berridge 2007; Flagel et al., 2011), whereas the central noradrenergic system is associated with arousal, attention, or cognitive flexibility (Berridge and Waterhouse, 2003; Arnsten and Li, 2005; Aston-Jones and Cohen, 2005; Bouret and Sara, 2005; Yu and Dayan, 2005; McGaughy et al., 2008; Sara and Bouret 2012). Several recent studies show that noradrenergic neuronal responses might also be related to motivation (Ventura et al., 2007, 2008; Bouret and Richmond, 2009; Bouret et al., 2012).

Here, we have measured the reactivity of noradrenergic neurons in the locus ceruleus (LC) from monkeys in relation to internal and external predictors of value (knowledge of the task

structure and visual stimuli predicting reward). We measured two behavioral responses lipping, an appetitive pavlovian reflex, and bar release, a goal-directed (operant) action, as we manipulated the size of the expected reward and whether or not the monkey had to make an action to obtain the reward. We also monitored the progression through the session to estimate decreasing drive as satiation and/or fatigue increased. As seen before, LC responses varied directly with the value of a cue presented at the beginning of a trial, in which the cue predicted reward size (Bouret and Richmond 2009). Here we find that the magnitude of LC activation was related inversely to the size of the expected reward when monkeys initiated the operant action leading to reward delivery. The relation between cue and action-related activity were related inversely on a trial-by-trial basis. Altogether, it appears that LC activity reflects the integration of external and internal aspects of the expected outcome value. We propose that one role of the LC activity is to provide the energy needed to perform the reward-directed action when the value of the outcome is expected to be low.

Materials and Methods

Subjects

Two male rhesus monkeys, L (9.0 kg) and T (9.5 kg), were used. The experimental procedures followed the Institute of Laboratory Animal Research *Guide for the Care and Use of Laboratory Animals* and were approved by the National Institute of Mental Health Animal Care and Use Committee.

Received Nov. 4, 2014; revised Jan. 15, 2015; accepted Jan. 24, 2015.

Author contributions: S.B. and B.J.R. designed research; S.B. and B.J.R. performed research; S.B. and B.J.R. analyzed data; S.B. and B.J.R. wrote the paper.

This work was supported by the Intramural Research Program of the National Institute of Mental Health. S.B. is supported by European Research Council Starting Grant BIOMOTIV.

The authors declare no competing financial interests.

Correspondence should be addressed to Sebastien Bouret, Motivation Brain Behavior Team, Institut du Cerveau et de la Moelle épinière, Groupe Hospitalier Pitié-Salpêtrière, 47, bd de l'Hôpital, F-75013 Paris, France. E-mail: sebastien.bouret@icm-institute.org.

DOI:10.1523/JNEUROSCI.4553-14.2015

Copyright © 2015 the authors 0270-6474/15/354005-10\$15.00/0




Figure 1. Experimental design. Monkeys perform three types of trials: cued-active (left), cued-passive (middle), and self-initiated (right) trials. Every trial starts when the monkey touches a bar. In cued trials, a visual cue (black and white pattern) indicates a combination of two factors: reward size (1, 2, or 4 drops of fluid) and action (active or passive trial). Left, In the cued-active trials, monkeys must release the bar when a red spot (wait signal, 500 ms after cue onset) turns green (go signal) after a variable time (500–1500 ms, jagged arrows), in which case a feedback (blue spot) appears, followed by the reward after 400–600 ms. Middle, In cued-passive trials, the feedback appears 2 s after cue onset independently of the monkey's behavior. Right, Self-initiated trials simply required touching and releasing a bar; no cue is present. Self-initiated trials are run in randomly alternating blocks of ~60 trials with constant reward size (1, 2 or 4 drops).

Behavior

Each monkey squatted in a primate chair positioned in front of a monitor on which visual stimuli were displayed. A touch-sensitive bar was mounted on the chair at the level of the monkey's hands. Liquid rewards were delivered from a tube positioned with care between the monkey's lips but away from the teeth. With this placement of the reward tube, the monkey did not need to protrude its tongue to receive a reward. The tube was equipped with a force transducer to monitor the movement of the lips (referred to as "lipping," as opposed to licking, which we reserve for the situation in which tongue protrusion is needed; Bouret and Richmond, 2009, 2010).

The monkeys performed the tasks depicted in Figure 1. They were first trained to perform a simple color discrimination task, in which each trial began when the monkey touched the bar. A red target point (wait signal) appeared in the center of the cue 500 ms after the appearance of the cue. After a random interval of 500–1500 ms, the target turned green (go signal). If the monkey released the touch bar 200–800 ms after the green target appeared, the target turned blue (feedback signal) 10–20 ms after bar release, and a liquid reward was delivered 400–600 ms later. Once the monkeys mastered this task (>80% correct trials), we introduced cued-active trials. In cued-active trials, the reward sizes of one, two, or four drops of liquid were related to a visual cue that appeared at the beginning of each trial, i.e., after the monkey touched the bar, 500 ms before the wait signal appeared. If the monkey released the bar before the go signal appeared or after the go signal disappeared, an error was registered. No explicit punishment was given for an error in either condition, but the monkey had to perform a correct trial to move on in the task. That is, the monkey had to repeat the same trial with a given reward size until the trial was completed correctly. Performance of the operant bar-release response was quantified by measuring error rates and reaction times.

Once the monkeys adjusted their operant performance as a function of reward-predicting cues (1–2 d), they were exposed to cued-passive trials.

In these passive trials, a different set of cues was used. The monkeys still had to touch the bar to initiate a trial, but, once the cue had appeared on the screen, releasing or touching the bar had no effect. Two seconds after cue onset, the blue point (identical to the feedback signal in cued-active trials) was presented, and water was delivered 400–600 ms later as a reward. The 2 s delay between cue onset and feedback signal was chosen to match the average interval between these two events in active trials. After 2–3 d of training with passive trials alone, monkeys had virtually stopped releasing the bar in cued-passive trials. In the final version, the six trial types with a combination of reward size (one, two, or four drops) and action contingency (active or passive) were presented in randomly interleaved order. Monkeys were exposed to this final version of the task for ~1 week before we started electrophysiological recordings.

For self-initiated trials, animals were placed in the same environment as before, except that the background of the screen had a large green rectangle present. Monkeys rapidly (1 d) learned to hold and release the bar to get the reward without any conditioned cue signaling reward size or timing of actions. To facilitate comparison with cued trials, a blue point was also used as a feedback signal on bar release, and reward was delivered within 400–600 ms. Before the neurophysiological recordings were taken, monkeys were trained for 1 week with alternating blocks of different reward sizes (one, two, or four drops). Each block comprised 50–70 trials with a given reward size and blocks alternated randomly and abruptly, without explicit signaling. Block-wise conditions were also used during recording.

Electrophysiology

After initial behavioral training, a 1.5 T MR image was obtained to determine the location of the LC by its relation to known MRI distinguished landmarks (including the inferior colliculus and the rostral part of the fourth ventricles) to guide recording well placement. Then, a sterile surgical procedure was performed under general isoflurane anesthesia in a fully equipped and staffed surgical suite to place the recording well and head fixation post. The well was positioned at the level of the interaural line, with an angle of ~15° in the coronal plane (the exact angle was determined for each monkey from its MR image).

Electrophysiological recordings were made with tungsten microelectrodes (impedance, 1.5 M Ω ; FHC or Microprobe). A guide tube was positioned using a stereotaxic plastic insert with holes 1 mm apart in a rectangular grid (Crist Instruments). The electrode was inserted through a guide tube. After a few recording sessions, MR scans were obtained with the electrode at one of the recording sites; the position of the recording sites was reconstructed based on relative position in the stereotaxic plastic insert and on the alternation of white and gray matter based on electrophysiological criteria during recording sessions. LC units were identified based on a combination of anatomical and physiological criteria similar to those used previously (Bouret and Richmond, 2009). These criteria included the following: (1) broad spikes (>600 μ s); (2) slow firing rate (1–3 spikes/s); (3) biphasic (activation/inhibition in response to salient stimuli); (4) complete absence of activity during sleep; and (5) reversible inhibition by clonidine (20 μ g/kg, i.m.), an α 2 adrenergic agonist.

Data analysis

All data analyses were performed in the R statistical computing environment (R Development Core Team 2004).

Lipping behavior. The lipping signal was monitored continuously and digitized at 1 kHz. For each trial, a response was defined as the first of three successive windows in which the signal displayed a consistent increase in voltage of at least 100 mV from a reference epoch of 250 ms taken right before the event of interest (cue or feedback).

Single-unit activity. In cued trials, the encoding of the factors reward size (one, two, or four drops), action (active or passive trial), or their interaction was studied using a sliding window procedure. For each neuron, we counted spikes in a 200 ms test window that was moved in 25 ms increments around the onset of the cue (from –400 to 1300 ms), around the feedback signal (from –800 to 1000 ms). At each point, a two-way ANOVA was performed with spike count as the dependent variable. The two factors were reward size (three levels) and action (two levels). In

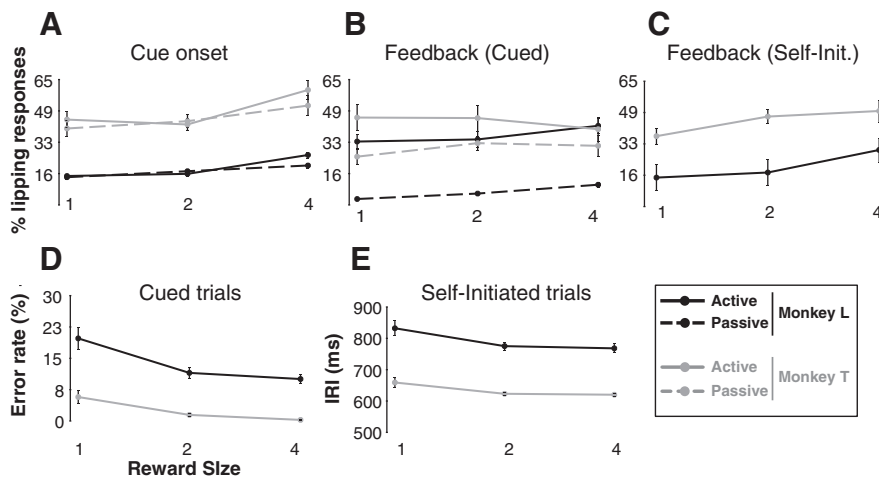


Figure 2. Behavior across task conditions. **A–C**, Lipping behavior (mean \pm SEM). **A**, In cued trials, the proportion of lipping responses to cues increased with reward size, with no difference between active and passive trials. **B**, At the feedback in cued trials, lipping responses were significantly more frequent in active than in passive trials, with little effect of reward size. **C**, In self-initiated trials, lipping increased for larger rewards. **D, E**, Bar-release behavior (mean \pm SEM). **D**, In cued trials, error rates were inversely related to reward size. **E**, In self-initiated trials, the latency to release the bar from the end of the preceding trial (release interval) decreased for larger rewards.

self-initiated trials, the same procedure was used to study the encoding of reward size around the feedback signal and reward.

Results of the sliding window analysis were used to define several epochs of interest. In cued trials, we considered two epochs: (1) cue onset (from 100 to 400 ms after cue onset) and (2) pre-feedback (from 300 to 0 ms before the feedback). In self-initiated trials, we only used a pre-feedback epoch (from 300 to 0 ms before the feedback). In each epoch, response latency was defined as the beginning of the first of three successive windows showing a significant effect ($p < 0.05$) of a given factor.

To quantify the dynamics of neuronal responses in the task, we measured the latency of event-related changes in activity across all conditions. For event-related changes in firing rate, the onset latency was defined as the middle of the first of three successive windows showing a significantly higher spike count compared with the baseline firing. The peak latency was defined as the middle of the window in which the firing rate was maximal, around the event of interest. To quantify the dynamics of firing modulations across conditions, we measured the onset latency of a given type of modulation (reward, action, or their interaction), defined as the middle of the first of three windows showing a significant effect of the factor of interest. The peak latency captures the time at which the effect was the largest for each neuron. It was defined as the middle of the window in which the effect was maximal for that event (largest variance explained), whether or not it was significant.

Results

Behavior

We measured two behaviors: (1) an appetitive pavlovian response (lipping) and (2) an operant response (bar release). In cued trials, lipping occurred around two task events: (1) just after the appearance of the cue and (2) at the time of the feedback for success (blue light). At the cue, lipping was related to the size of the expected reward, with more lipping when the cue indicated that a larger reward would be forthcoming regardless of whether it was an active or passive trial (Fig. 2A; monkey T, $F_{(2)} = 5.9$, $p = 0.004$; monkey L, $F_{(2)} = 18.8$, $p < 10^{-4}$). There was no significant difference between active and passive trials (no effect of action: monkey T, $F_{(1)} = 1.0$, $p = 0.3$; monkey L, $F_{(1)} = 1.8$, $p = 0.2$) and no significant interaction (monkey T, $F_{(2,78)} = 0.6$, $p = 0.5$; monkey L, $F_{(2,246)} = 2.9$, $p = 0.06$). At the feedback in cued trials, there was a significant effect of action on lipping, with greater lipping in active than in passive trials (Fig. 2B; monkey T, $F_{(1)} = 9.3$, $p =$

0.003; monkey L, $F_{(1)} = 144$, $p < 10^{-4}$). For monkey L, there was a significant small effect of reward ($F_{(2)} = 3.6$, $p = 0.03$) but not for monkey T ($F_{(2)} = 0.3$, $p = 0.7$), and there was no significant interaction between action and reward (monkey T, $F_{(2,78)} = 0.6$, $p = 0.5$; monkey L, $F_{(2,246)} = 2.9$, $p = 0.06$). At the feedback in self-initiated trials, lipping increased with reward size for both monkeys (Fig. 2C; monkey T, $F_{(2)} = 3.8$, $p = 0.03$; monkey L, $F_{(2)} = 6.4$, $p = 0.003$).

In both cued-active and self-initiated trials, the operant performance (releasing a bar) was related to the size of the expected reward. Error rates decreased as reward increased in cued-active trials (Fig. 2D; monkey T, $F_{(2)} = 8.3$, $p = 4 \times 10^{-4}$; monkey L, $F_{(2)} = 7.6$, $p = 8.0 \times 10^{-4}$), and the inter-release interval became shorter as reward increased (Fig. 2E; monkey T, $F_{(2)} = 3.9$, $p = 0.02$; monkey L, $F_{(2)} = 3.1$, $p = 0.05$).

Electrophysiology: identification of LC neurons

We recorded the activity of 121 single LC neurons in two monkeys ($n = 63$ in monkey T, $n = 58$ in monkey L). Single LC units were identified outside the task periods based on their anatomical localization and their electrophysiological signature (slow rate $< 5/s$, wide spikes ~ 2 – 3 ms, strong modulation with arousal), as described previously (Bouret and Richmond 2009). We report the activity of all neurons for which the activity was tested in at least 20 trials in the cued ($n = 100$ neurons) and/or self-initiated conditions ($n = 86$ neurons), with 65 recorded in both. The average firing rate of all 121 neurons was 2.5 ± 0.1 spikes/s over the entire recording session. We never observed any change in “tonic” activity during the course of these experiments, aside from the well known relation to the sleep–wake cycle (Bouret and Richmond, 2009).

LC neurons respond to major task events

All of the putative LC neurons showed phasic activation during the task. The activation occurred at the cue and/or the feedback in cued trials and at the feedback in self-initiated trials, as shown in the example (Fig. 3A–C). Generally, the phasic responses were modulated according to task conditions. Globally, the firing of LC neurons was modulated by reward size after cue onset (Fig. 3A) and, to a lesser extent, before the bar release (see below). In cued trials, the activation around the feedback was modulated strongly by the factor action (Fig. 3B): in cued-active trials, in which the reward was contingent on an action (bar release), there was a strong LC activation just before the action that preceded the feedback by a few milliseconds. In cued-passive trials, in which no action was required, there was a small activation after the feedback announcing the imminent reward delivery. Thus, the activation of LC neurons occurred mainly in two circumstances: (1) significant visual signals (the cue in cued trials and, to a much lesser extent, the feedback in cued-passive trials); and (2) the initiation of the bar release (in both cued-active and self-initiated trials). In cued-active trials, the feedback occurred so soon after the action (10–20 ms) that the two events are in practice superimposed and thereby can be regarded as equivalent for the purpose of our analysis.

Dynamics of neuronal responses

We used the sliding window analysis to characterize the dynamics of the firing modulation for each neuron across the six task conditions, defined by the two factors reward size and action (Fig. 4). These z-scored responses were subjected to two-way ANOVA across our sliding windows with factors reward size (three levels for the three reward sizes) and action (two levels representing the active and passive conditions). The strength of the reward size modulation is represented by the variance explained for reward size (Fig. 4A–C), and the strength of the action modulation is represented by the variance explained by action (Fig. 4D,E). At the time of the cue, there is a substantial amount of power in the reward signal for about half of the neurons (Fig. 4A), whereas the power related to the action is concentrated around the feedback signal (Fig. 4E). The response power related to the action is biphasic because of the activation occurring before the feedback in cued-active conditions and after the feedback in cued-passive conditions (Fig. 3B for a single-unit example; see Fig. 6B for population activity). In cued-active trials, the activation occurs before the bar release on which the feedback signal and the reward are contingent. In cued-passive trials, LC neurons are activated after the feedback, which predicts the reward and occurs regardless of the behavior. In the self-initiated trials, a one-way ANOVA reveals that approximately one-fifth of neurons show a weak influence of the reward size factor (Fig. 4C). Thus, the neuronal activation at the cue is almost exclusively modulated as a function of the upcoming reward, with no difference between trials in which monkeys make an action or wait to obtain it. The activity around the feedback is more strongly modulated by the way in which the trial is completed (actively or passively) than by the reward size.

We quantified the dynamics of these firing modulations (see Materials and Methods). In cued trials, we focused on the dominant effects (reward at the cue, action at the feedback). When we measured the onset and peaks of activity at the cue in cued trials (Table 1), the firing rate increased for all cues, ~ 100 ms before differential firing across reward sizes appeared (Wilcoxon's test, $p < 10^{-4}$). That is, the initial part of the phasic activation had the same intensity across reward conditions, and the divergence began after 100 ms. The average firing rate and the firing modulation across rewards peaked at approximately the same time (Wilcoxon's test, $p > 0.05$). Thus, at cue onset, the difference in firing rate across

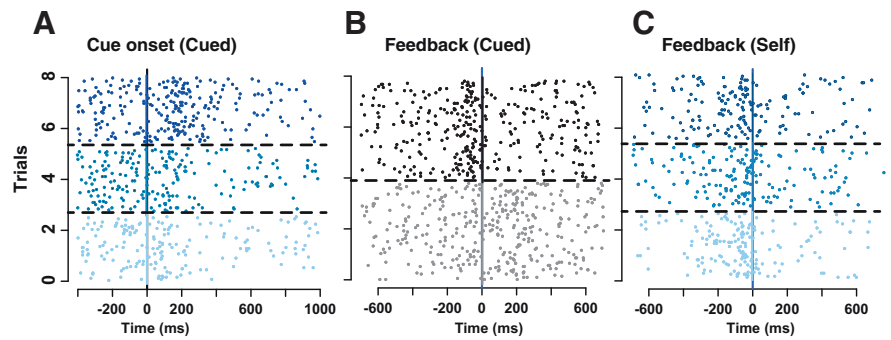


Figure 3. Activity of a representative LC unit. Raster display for the activity of a representative LC neuron aligned on the cue onset in cued trials (**A**) and the feedback in cued (**B**) and self-initiated (**C**) trials. In **A** and **C**, trials are sorted by decreasing reward sizes, with the three levels of blue corresponding to the three levels of reward. Trials corresponding to different conditions are separated by broken lines. For this neuron, expected reward modulates (positively) the firing at the cue (**A**) but not at the action in self-initiated trials (**C**). At the feedback in cued trials (**B**), trials are sorted by trial type (black, active trials; gray, passive trials). The activation before the feedback only occurs in active trials, when monkeys need to perform an operant action (bar release) to get the reward. In passive trials, there is a small activation after the visual feedback.

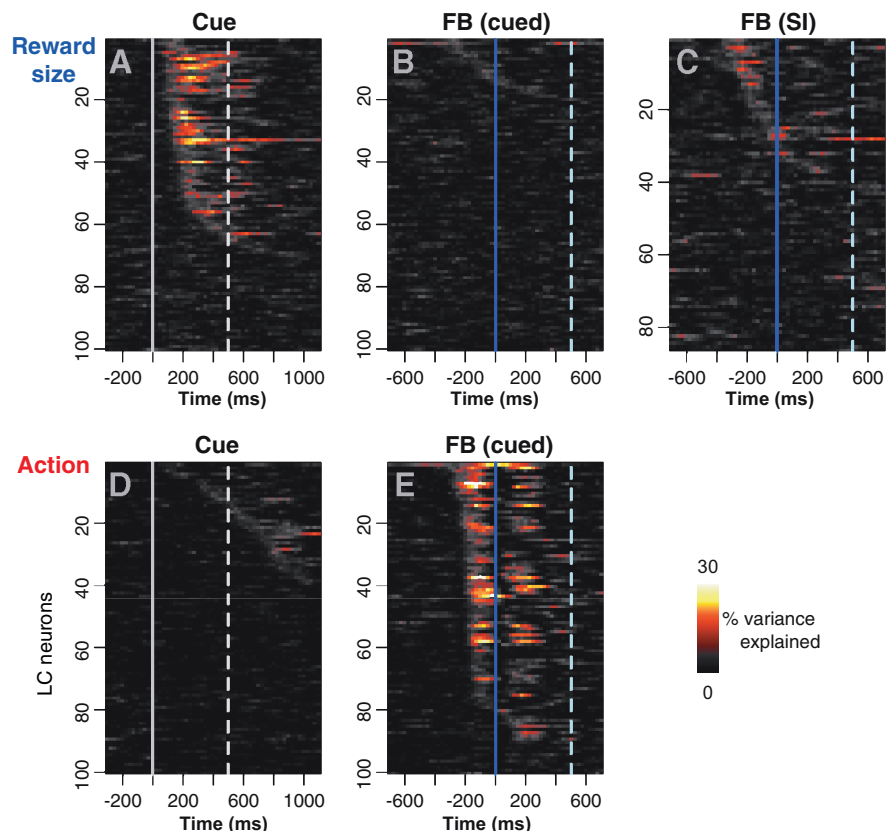


Figure 4. Sliding window analysis of the firing modulation across task condition. Distribution of the firing modulation across neurons (ordinates) and time around an event (abscissas). The firing modulation, quantified using percentage variance explained, was calculated using a sliding window ANOVA for the reward size factor (**A–C**) and the action factor (**D–E**). FB, Feedback. At the cue, LC neurons are clearly sensitive to the amount of expected reward (**A**) but not to the way in which it can be obtained (**D**). The broken line indicates the time at which the red wait signal appears in cued-active trials. At the feedback in cued trials, there the reward has very little influence on firing (**B**), but neurons clearly distinguish between active and passive trials (**E**). At the feedback in self-initiated trials (**C**; SI), a few neurons display a limited modulation as a function of the expected reward. In **B**, **C**, and **E**, the broken line indicates the average time of the reward delivery.

conditions was biggest at the time of the peak in firing rate. At the feedback, in both cued and self-initiated trials, the dynamics of the global changes in firing rate were the same as the firing modulation across conditions (Wilcoxon's test, $p > 0.05$). In other words, the difference between cued-active and

Table 1. Latencies of LC responses (median and interquartile range across all neurons)

Latencies (ms)	Cue (reward effect)	Feedback (action effect)	Self-Initiated (reward effect)
Rate onset	50 (0/125)	−225 (−275/−175)	−250 (−350/−175)
Rate peak	150 (75/150)	−100 (−125/13)	−75 (−100/−50)
Effect onset	150 (100/225)	−200 (−225/−150)	−275 (−425/−100)
Effect peak	200 (150/431)	−138 (−200/125)	−100 (−325/263)

cued-passive trials was constant in time during a trial, with an activation occurring only before the feedback (and the corresponding action) in active trials and only after the feedback in passive trials.

Modulation of LC activity: single neurons and population effects

We calculated the firing modulations related to reward and action for each neuron using two-way ANOVAs in three epochs: (1) after the cue; (2) before the feedback in cued trials; and (3) before the feedback in self-initiated trials. These epochs (300 ms windows around the peak responses) were chosen based on the observed dynamics of these modulations (Fig. 4; for additional details, see Materials and Methods). For each epoch, we calculated the proportion of neurons showing a significant effect ($p < 0.05$) of the two factors (reward and action) and their interaction (Fig. 5A). We also measured the variance explained by the two factors and their interaction for each neuron, whether or not the corresponding effect was significant (Fig. 5B). Both measures of firing modulation were related to expected rewards and action, and they showed a similar pattern: at the cue, the firing was most strongly modulated as a function of expected reward size. At the feedback in cued trials, firing was mostly modulated by the action factor, i.e., the response was different in active and passive trials, but there was also a significant proportion of neurons (10%) displaying an interaction between the factors reward and action. In self-initiated trials, firing at the feedback was modulated according to the expected reward.

To examine the population as a whole, we standardized individual neuronal responses by converting them to z-scores and ran the previous ANOVA using the activity of all recorded LC units (all trials for all conditions). At the cue, the standardized firing of the population was affected significantly by the factor reward size ($F_{(2)} = 254, p < 10^{-4}$), but there was no effect of action ($F_{(1)} = 0.2, p = 0.6$) nor was the interaction between action and reward significant ($F_{(2,594)} = 2.1, p = 0.1$). Before the feedback in cued trials, there was a significant effect of action ($F_{(1)} = 748, p < 10^{-4}$) and a significant interaction between action and reward ($F_{(2,594)} = 13, p < 10^{-4}$). After the feedback in cued trials, there was a significant effect of action only ($F_{(1)} = 418, p < 10^{-4}$) and no effect of reward and no interaction. At the bar release in self-initiated trials, there was no significant effect of reward ($F_{(2)} = 0.2, p = 0.8$). Our analyses were generally performed using ANOVAs, a quite robust statistical approach. However, LC firing rates are low, and, because spike trains appear as if they arose from a point process, it might be appropriate to examine the spike count data using a generalized linear model (GLM) having a Poisson link function (Truccolo et al., 2005). All the effects reported above using ANOVAs with spike counts as the dependent variable were unchanged using the GLM with Poisson link for the same data. The positively signed modulation of cue-evoked activity by the expected reward size ($p < 10^{-4}$), the greater activation before the feedback (and the action) in cued-active vs cued-passive trials (no action, $p < 10^{-4}$) and the negatively signed modulation by

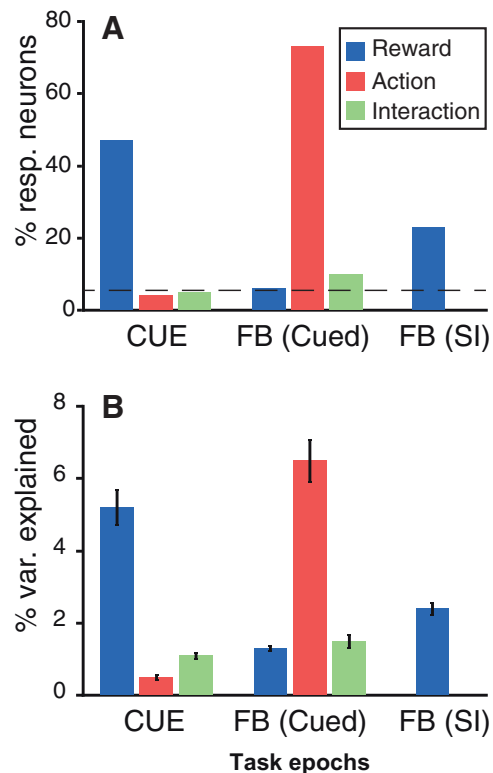


Figure 5. Modulation of neuronal activity. Influence of the task factors (reward, 1, 2, or 4 drops; action, active vs passive trials) and their interaction across epochs: cue onset, feedback in cued trials, and feedback in self-initiated trials. We considered three key epochs for this analysis: after the cue onset, before the feedback in cued trials, and before the feedback in self-initiated trials. **A**, Percentage of neurons (with all recorded neurons included) showing an effect; the broken line indicates the percentage of neurons expected by chance. **B**, Percentage of variance explained (mean \pm SEM across all recorded neurons). At cue onset, LC neurons only show an effect of reward size. At the feedback (FB) in self-initiated trials (SI), LC neurons show a clear effect of action (difference between active and passive trials) and a small interaction between action and reward. At the feedback in self-initiated trials, there is a significant effect of the expected reward.

the expected reward size before the action in cued-active trials ($p < 10^{-4}$).

To visualize the dynamics of the firing of the entire population, we plotted the average z-scored firing of all recorded LC units across task conditions (Fig. 6). At cue onset, we only compared firing across the three reward sizes because there was no significant modulation according to the factor action (see preceding paragraph). The population firing increased with the size of the expected reward, with little variability across neurons (Fig. 6A). In line with the analysis of individual units (Fig. 4A; Table 1), the largest modulation by the factor reward occurred in the late phase of the response to the cue, when the firing rate started to decrease after the initial activation. The initial phase of the activation was identical across expected rewards. In cued-active trials, neurons were activated before the feedback, i.e., shortly before the monkey initiated the action, and Figure 6B (solid lines) shows that the amplitude of the action-related activation is inversely correlated with the expected reward. In cued-passive trials, in which the feedback appears without any behavioral requirement, the population of LC neurons showed a small activation that did not change across reward sizes (Fig. 6B, dotted lines). In self-initiated trials in which monkeys released the bar without an external trigger and when the information about the forthcoming reward was only available through the remembered

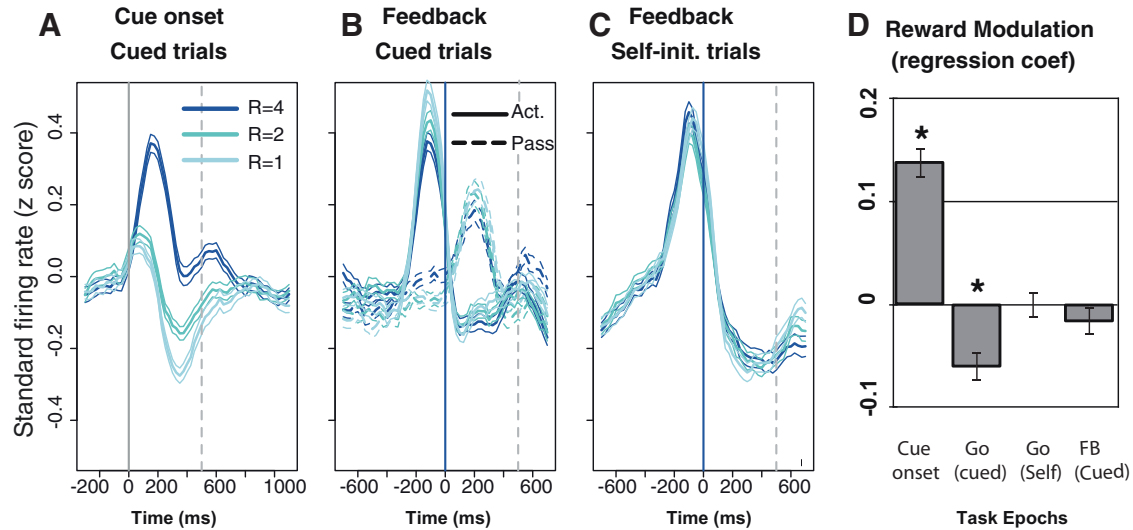


Figure 6. Average population activity. **A–C**, Population activity (z-scored by neuron, mean \pm SEM) for 100 neurons in cued trials and 86 neurons in self-initiated trials. At cue onset (**A**), the population is phasically activated in all three reward conditions. There is a strong positive modulation by the expected reward. Before the action in cued-active trials (**B**), there is a negative influence of the expected reward on the action initiation-related activation. In passive trials, the small activation after the feedback is not modulated by the expected reward. Act., Active; Pass, passive. In self-initiated trials (**C**), the population shows a strong activation before action initiation but no modulation by the expected reward. **D**, Reward modulation (mean \pm SEM across all neurons) of the regression coefficient in a linear model of the spike count as a function of reward size, across four epochs of the task: 300 ms windows after cue, before action in cued active, before action in self initiation, and after feedback (FB) in cued passive. * $p < 0.01$, t test. At the population level, the expected reward has a positive influence on cue-evoked activity and a negative influence on action initiation-related activity.

context (Fig. 6C), LC neurons displayed a strong phasic activation before the action that did not change across different expected rewards. All these effects of the expected reward size on firing rate were confirmed with a regression analysis across the epochs of interest (Fig. 6D). Again, we observed the same effects using a GLM with Poisson link.

To evaluate the average magnitude of the neuronal activation around task events, we measured the changes in firing rate in 300 ms windows around the events of interest relative to a reference period just before trial onset. At the cue (100–400 ms after cue onset), the sign of the firing modulation varied with the expected reward: when the expected reward was the smallest (one drop), the response was significantly smaller than the baseline activity (baseline, $87 \pm 3\%$, average across all 100 cells; $t_{(99)} = -4$, $p = 0.0001$). For the intermediate reward (two drops), the firing rate was indistinguishable from the baseline ($97 \pm 3\%$ from baseline; $t_{(99)} = -0.9$, $p = 0.33$), and for the large reward (four drops), the firing rate was elevated significantly above the baseline ($155 \pm 11\%$ from baseline; $t_{(99)} = 5.1$, $p = 1.3 \cdot 10^{-6}$). In cued-active trials, there was a significant increase in firing in the 300 ms preceding the action ($159 \pm 8\%$ from baseline; $t_{(99)} = 7.6$, $p = 1.8 \cdot 10^{-11}$) but not in the 300 ms after the feedback ($94 \pm 8\%$ from baseline; $t_{(99)} = -0.7$, $p = 0.5$). In cued-passive trials, there was a significant activation after the feedback ($131 \pm 4\%$ from baseline; $t_{(99)} = -7.5$, $p = 3.4 \cdot 10^{-11}$), but the firing before the feedback was indistinguishable from the baseline firing rate ($102 \pm 4\%$ from baseline; $t_{(99)} = 0.5$, $p = 0.6$). In self-initiated trials, the firing increased to $162 \pm 6\%$ right before the action compared with the preceding 300 ms. The magnitude of the action-related activation for the cued-active trial was indistinguishable from that in the self-initiated trials (t test, $p > 0.05$).

Several features of the population activity were consistent with the analysis of single units: at the cue, the analysis of individual neurons and that of the average firing consistently showed a strong modulation by reward size, but there was no

detectable difference between active and passive trials (Figs. 3A, 4, 5, 6A). In cued trials, neurons were only activated before the feedback in active trials (i.e., before the action) and after the feedback in passive trials (when there was no action).

Several features of the population activity could not be predicted from the analysis of individual units. Although only few individual neurons display a significant modulation by reward or an interaction between action and reward (Figs. 4, 5), before the feedback in cued trials, a vast majority of neurons shift their firing in the same direction (negative influence of reward), thereby providing a reliable signal in the population (Fig. 6B,D). Reciprocally, although a significant proportion of individual neurons display a firing modulation related to reward before the feedback in self-initiated trials (Figs. 4C, 5), there is no modulation related to the reward size at the population level (Fig. 6C,D). This implies that the effects observed on individual neurons are not coherent enough to emerge at the level of the population.

When we examined the relationship between LC activity and several behavioral variables, including lipping and reaction times, none of the trial-by-trial correlations were significant.

Overall, three features are seen in the analysis of event-related activity. (1) LC neurons were activated at the cue, and the late part of the cue response was related directly to the size of the expected reward. (2) LC neurons were activated shortly before the action (bar release) in both cued and self-initiated trials. (3) The size of population response in cued trials was inversely related to the reward size but not in self-initiated trials.

Relation between cue- and action-related activity

The finding above indicates that the reward-related response modulations at the time of the cue are on average inversely related to those at the time of the feedback (illustrated in Fig. 6). For a majority of LC neurons, the relationship between firing and reward was positive at the cue and negative before the action (Fig. 6D; t test at the second level, $p < 0.05$). The influence of reward

was significantly positive for 29 individual units at the cue and significantly negative for 13 units at the action. For five units, the effect was significant in both directions. To test more directly the relationship between activity at the cue and the action across neurons, we examined the correlation between the reward regression coefficients at the times of these two events. This correlation was negative ($r = -0.2, p = 0.004$), indicating that, across neurons, the larger rewards led to increases in LC firing at the cue and to decreased LC firing at the time of the action.

This brings up the question of whether the responses at those two times are inversely related on a trial-by-trial basis. We calculated the difference between the z-scored firing rate in 300 ms epochs after the cue and before the bar release on each trial for each neuron. The linear regression coefficients for the relation between this response difference (action cue) and the expected reward size were negative for 86 of 100 neurons (Fig. 7A), and this regression was significant for 42 of the 100 neurons. None of the 14 positive correlations reached significance. We measured the influence of reward size on response difference of all neurons using a repeated-measures ANOVA. In other words, we specifically measured the “within-neuron” effect of reward size on response difference, regardless of the random fluctuations across the 100 neurons. The effect of reward size was significant ($F_{(2)} = 87.9, p < 10^{-10}$), with the response difference (action cue) being positive for small rewards and negative for large rewards (Fig. 7B). Thus, the difference between cue and action-related activity in each trial depended strongly on the expected reward size for both individual neurons and the population as a whole. For large incentives, the activity was greater at the cue than at the action, and for smaller incentives, the firing was relatively smaller at the cue and greater before the action.

Influence of internal motivational factors

We also examined the influence of internal aspects of motivation on the behavior and the activity of LC neurons in this task. As reported previously, there was a gradual decrease in motivation as monkeys progressed through each daily session and accumulated water (Minamimoto et al., 2009; Bouret and Richmond, 2010). We compared operant performance between recordings that took place first thing during a daily session, when animals were presumably thirstier, and recordings that were obtained at the end of the session, when monkeys were presumably less thirsty and about to stop working on that day (Fig. 8). For both monkeys, in cued-active trials and self-initiated trials, there was a significant decrease in operant performance between the beginning and the end of the daily sessions (t test, $p < 0.05$). In cued trials, the average error rate went from 0.6 to 9% for monkey T and from 10 to 18% from monkey L (Fig. 8A,B). In self-initiated trials, the average inter-release interval went from 613 to 669 ms for monkey T and from 753 to 812 for monkey L.

We compared the activity of LC neurons recorded at the beginning to the activity recorded at the end of a daily session around three events of interest: (1) after cue onset; (2) before the feedback in cued trials; and (3) before the feedback in self-initiated trials (Fig. 8). We used the standardized (z-scored) spike counts in 300 ms epochs of interest for each neuron recorded at either the beginning or the end of a session for the neuronal response measures. At cue onset in a three-way ANOVA with factors reward (three levels), action (two levels), and advancement (two levels, beginning and end), there was a significant effect of reward size, no effect of action, and a significant interaction between the factors reward and advancement ($F_{(2,324)} = 5.5, p = 0.004$). The firing increased with expected reward size at both

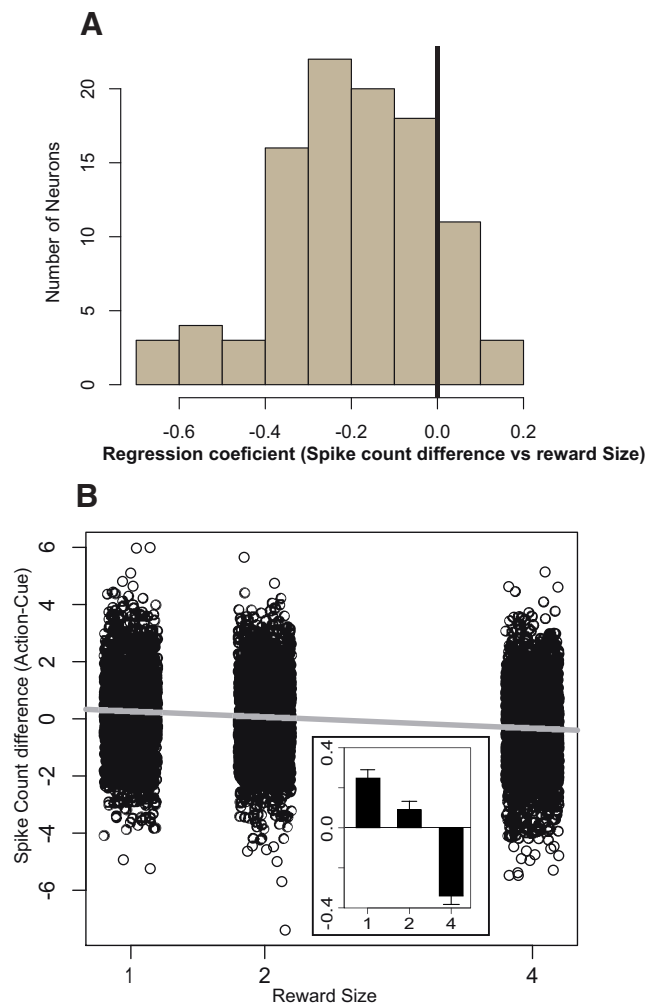


Figure 7. Within-trial dynamics of reward sensitivity. **A**, Distribution of the reward modulation of spike count difference between cue and action across all neurons. The regression coefficient of spike count difference (action-cue) as a function of reward size was estimated for each LC unit. The distribution is clearly shifted to the left (mean \pm SEM, -0.2 ± 0.02 ; significantly different from 0, $t_{(99)} = -11.5, p < 10^{-10}$). In addition, this negative relation between expected reward and spike count difference was significant for 42 units, and none showed a significant positive correlation. Thus, LC neurons show a homogenous tendency to have a reward-dependent modulation of firing between the onset of the cue and the triggering of the behavioral response. **B**, Influence of the expected reward size on the spike count difference between cue and action initiation-related activity for each neuron, on a trial-by-trial basis. The expected reward has a significant negative influence on that difference ($\beta = -0.2, r^2 = 0.03, p < 10^{-10}$). The inset shows the mean values of these differences across the three reward sizes. The spike count difference showed a significant effect of reward size in a one-way ANOVA, after removing the between-subject effects ($F_{(1)} = 0.2, p = 0.8; F_{(2)} = 87.9, p < 10^{-10}$). For small rewards, the activity increases between cue onset and action initiation. For large rewards, the activity decreases between cue and action.

the beginning (Fig. 8C) and the end of a session (Fig. 8D), but the firing difference across reward conditions was significantly smaller (less discriminative) at the end than at the beginning of a session. Before the feedback, there was a significant interaction among the three factors ($F_{(2,324)} = 4.3, p = 0.01$). The firing was greater in cued-active than in cued-passive trials (Fig. 8E,F). This action-related activity was affected by the expected reward at the end of a session, when the monkeys were least motivated (compare the four drop points in Fig. 8C,D for the active condition). In self-initiated trials, the action-related activity was unaffected by either by the reward or the advancement (all p values > 0.05).

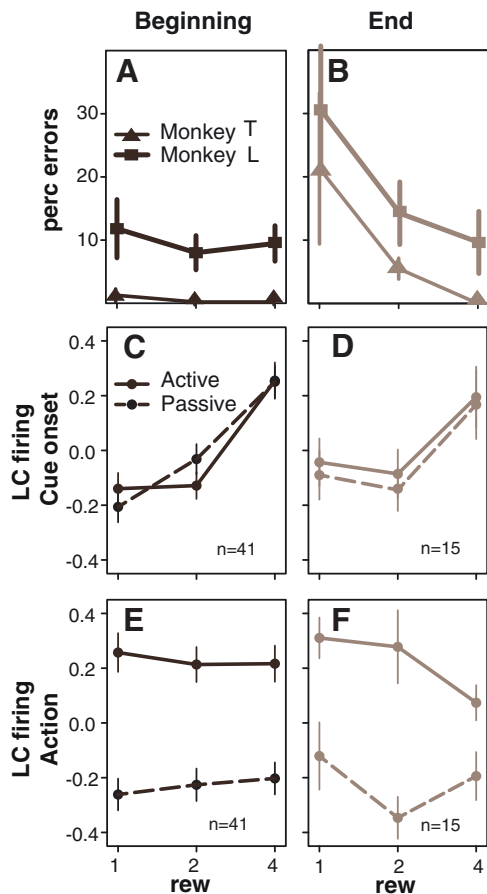


Figure 8. Modulation of LC responses by fatigue and satiety. At the behavioral level, the influence of reward on error rates (mean \pm SEM) clearly increases between the beginning of each daily session (**A**; when monkeys presumably are fresh and thirsty) and the last ones (**B**), at the end of the session, when the monkeys are presumably tired, satiated, and about to stop working. At cue onset, the firing of LC neurons (mean \pm SEM) is more strongly affected by the expected reward at the beginning (**C**) compared with the end (**D**) of the recording sessions (n indicates the number of neurons used in this specific analysis). There is no difference between active and passive trials at cue onset. This is in line with the idea that the positive effect of reward value decreases over time, as monkeys get satiated. At the time of action initiation (**E**, **F**), the firing of LC neurons is significantly higher in active trials (when monkeys actually perform an action) compared with the same window in passive trials (right before the feedback when there is no action). Here, the effect of reward (rew) is less pronounced at the beginning (**E**) than at the end (**F**) of the session: the negative influence of reward on action-related activity is more pronounced when monkeys are about to stop working. This effect only occurs when there is an action; it is not there in passive trials. The number in the bottom right corner of each box (**C**, **D**) indicates the number of neurons used for the corresponding analysis.

Overall, these data indicate that firing of LC neurons integrate both external (stimulus dependent) and internal (satiety) information about outcome value. At the end of the session, when monkeys are satiated and/or tired, the size of the expected reward has a weaker positive effect on cue-evoked activity but a stronger negative effect on action-related firing.

Discussion

LC neurons were activated around two behaviorally significant events: (1) cue onset and (2) bar release. In cued trials, the neuronal responses were reciprocally related around these two events. Activity increased as expected reward increased at the cue, whereas activity decreased as expected reward increased at bar release. As might be suspected from this pattern, the internal motivational factors (modulated with the progression through

each daily session) reflected the same reciprocal relation between LC activity at the cue and at the bar release.

What process does LC activation subserve?

Activation of LC neurons has long been associated with salient events and the associated orienting responses (Foot et al., 1980; Aston-Jones and Bloom, 1981; Abercrombie and Jacobs, 1987; Berridge and Waterhouse, 2003; Bouret and Richmond, 2009; Sara and Bouret, 2012). This activation probably arises via connections from subcortical areas, such as the paragigantocellularis nucleus and the central nucleus of the amygdala (Aston-Jones et al., 1991; Amorpant et al., 2000; Bouret et al., 2003; Valentino and Van Bockstaele, 2008). The novelty here is the difference in timing between the stimulus-induced activation from baseline activity and the modulation of the firing across conditions. The strongest reward modulation occurred \sim 100 ms after the peak of the activation, when the activity diverged across conditions, with firing falling faster in low reward conditions. One possible source for the signal giving rise to this late modulation of the cue-evoked activity is ventral prefrontal cortex areas, in which neurons also encode reward size in this task and project directly to the LC (Chiba et al., 2001; Vertes, 2004; Bouret and Richmond, 2010). In this task, the signal related to the reward at the cue appeared much earlier in the orbitofrontal cortex (\sim 60 ms) than in the LC (\sim 150 ms).

Although it is possible that the responses related to cue onset and bar release reflect two different functions, it seems more parsimonious to consider that they reflect a common underlying function. It has been shown several times that the activation of LC neurons is not simply primary sensory or motor (Bouret and Sara, 2004; Clayton et al., 2004; Aston-Jones and Cohen, 2005; Bouret and Sara, 2005; Kalwani et al., 2014). Because of the reciprocal relation between the LC activation and the behavioral measures of reward value at the cue and the bar release, it is difficult to construct a general interpretation in terms of costs or benefits. It also seems difficult to construct a general interpretation in terms of attention (Bouret and Richmond, 2009).

LC activation encodes behavioral energy?

Our interpretation is based on the intuition that triggering an action requires a given amount of resources to be mobilized. These resources include all the muscular, neuronal, and metabolic (such as oxygen, glucose, and ATP) components involved in the movement. In the task we used, the bar-release action is the same in all three active conditions, so we can reasonably assume that the amount of physical, and perhaps mental, energy required to trigger the action is the same. Behaviorally, energy can be mobilized by the incentive drive, which includes potential reinforcers as a function of the internal state. At the cue onset, both behavioral (lipping) and electrophysiological data indicate that the size of the expected reward has a positive incentive effect. In addition, the probability of generating the response at all in a given trial increases with the expected reward. In our framework, this would correspond to the cue mobilizing a given amount of energy, which would be correlated positively with the size of the expected reward. However, triggering the action requires an extra amount of energy to reach the necessary threshold. If the threshold is the same across conditions (an assumption we make given that the action is the same, and it is imposed), the complementary amount of energy required to reach the threshold would be inversely related to the amount of energy invested at the cue. In other words, the greater the incentive, the less energy would be needed at the time of the action to reach the threshold and trigger

the response. This is precisely what we observe in the activity of LC neurons. This is analogous to the subjective effort required to act when the expected reward has little value: something must complement the objective reward value to perform the action. More generally, this model of behavioral energy predicts that the relation between the cue-related and action-related activity would vary systematically with the expected reward, and again, this is precisely what we observed for LC neurons.

The modulation of LC responses as monkeys progress in the session is compatible with this hypothesis. LC responses to cues become less sensitive to reward, because the incentive power of the cues decreases when monkeys approach the breaking point because of satiety and/or fatigue (Bouret and Richmond, 2010). This effect mirrors the decrease in reward coding in the ventromedial prefrontal cortex, a key structure for the encoding of subjective value (Boorman et al., 2009; Lebreton et al., 2009; Bouret and Richmond, 2010; Lim et al., 2011). At the bar release, we observed a complementary process: as monkeys progressed through the session and were approaching their breaking points, the sensitivity to reward increased, with neurons displaying larger responses for smaller rewards. Again, this is in line with the interpretation in terms of energy, because triggering the action requires more compensation when monkeys approach the breaking point and the incentive effect of rewards decreases, especially for the least preferred option.

In self-initiated trials, monkeys initiate the action spontaneously, and because there is no cue and reward conditions are blocked, there is probably no extra cost for triggering the action when the expected reward remains constant from one trial to the next. In these trials, LC neurons are activated in relation to the amount of energy necessary to perform the action, but because there is no cue-dependent mobilization of energy, no reward-dependent compensation occurs at the action. Thus, interpreting the modulation of LC activity in terms of energetic mobilization seems to provide a reasonable account of the activity pattern in this task.

This interpretation in terms of mobilization of energy necessary for action could account for activity in other tasks in which LC activation occurs around goal-directed actions (Bouret and Sara, 2004; Clayton et al., 2004; Kalwani et al., 2014). In another related task, a reward schedule task, LC activation at the cue correlated positively with the incentive value measured using lip-ping and the response at the action was weaker in unrewarded compared with rewarded trials (Bouret et al., 2012). Thus, even if the negative influence of expected reward on action-related activity is small, it is reliable enough to emerge at the population level in two different tasks. Finally, other areas, such as the posterior cingulate cortex and the centromedial thalamus, show a similar response pattern, and they could form a functional network with the LC (Minamimoto et al., 2005; Heilbronner and Platt, 2013).

The close relationship between LC activity and physiological arousal is additional evidence that the firing of LC neurons is associated with the mobilization of physiological resources to face challenges (Abercrombie and Jacobs, 1987; Valentino and Van Bockstaele, 2008; Sara and Bouret, 2012). However, the function of LC neurons is not simply attributable to arousal, because the noradrenergic system plays a critical role in attention, inhibitory control, and cognitive flexibility (Arnsten and Li, 2005; Aston-Jones and Cohen, 2005; Yu and Dayan, 2005; Chamberlain et al., 2006; Tait et al., 2007; McGaughy et al., 2008; Robbins and Arnsten, 2009; Sara and Bouret, 2012). All these functions are cognitively demanding and can be cast in terms of cognitive effort in which LC

activation reflects the mobilization of energy necessary to perform not only physical actions (current work) but also cognitive operations, with the amount of noradrenergic mobilization increasing with the difficulty of the task at hand.

There are other ideas about the role of LC neurons for behavior. Aston-Jones and Cohen (2005) base their theory on a tonic/phasic mode switch, which has been difficult to reproduce, even when the animals performed reversal learning tasks (Bouret and Sara, 2004; Bouret et al., 2009; Kalwani et al., 2014). Yu and Dayan (2005) present a model that describes learning in the face of unexpected uncertainty. In our experiment, there is very little uncertainty so it is difficult to see a strong relation between our data and their model. Nonetheless, uncertainty is a challenge for the organism, and in that sense it would also require a significant mobilization of energy to act in the face of uncertainty (Sara and Bouret, 2012).

Here we have shown that the activity of LC neurons reflects both expected reward and action. Their firing pattern suggests that they are activated in relation to the energy needed to respond to behaviorally significant events as a function of the current behavioral state, anticipated actions, and corresponding outcomes. In this hypothesis, the noradrenergic system would be critical in effortful situations, in which physically and/or mentally difficult operations require a high level of energy to be completed.

References

- Abercrombie ED, Jacobs BL (1987) Single-unit response of noradrenergic neurons in the locus coeruleus of freely moving cats. I. Acutely presented stressful and nonstressful stimuli. *J Neurosci* 7:2837–2843. [Medline](#)
- Amarapant P, LeDoux JE, Nader K (2000) Different lateral amygdala outputs mediate reactions and actions elicited by a fear-arousing stimulus. *Nat Neurosci* 3:74–79. [CrossRef Medline](#)
- Arnsten AFT, Li BM (2005) Neurobiology of executive functions: catecholamine influences on prefrontal cortical functions. *Biol Psychiatry* 57:1377–1384. [CrossRef Medline](#)
- Aston-Jones G, Bloom FE (1981) Norepinephrine-containing locus coeruleus neurons in behaving rats exhibit pronounced responses to non-noxious environmental stimuli. *J Neurosci* 1:887–900. [Medline](#)
- Aston-Jones G, Cohen JD (2005) An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu Rev Neurosci* 28:403–450. [CrossRef Medline](#)
- Aston-Jones G, Shipley MT, Chouvet G, Ennis M, van Bockstaele E, Pieribone V, Shiekhhattar R, Akaoka H, Drolet G, Astier B, Charl  ry R, Valentino RJ, Williams JT (1991) Afferent regulation of locus coeruleus neurons: anatomy, physiology and pharmacology. *Prog Brain Res* 88:47–75. [CrossRef Medline](#)
- Berridge CW, Waterhouse BD (2003) The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Res Brain Res Rev* 42:33–84. [CrossRef Medline](#)
- Berridge KC (2007) The debate over dopamine's role in reward: the case for incentive salience. *Psychopharmacology* 191:391–431. [CrossRef Medline](#)
- Boorman ED, Behrens TEJ, Woolrich MW, Rushworth MFS (2009) How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* 62:733–743. [CrossRef Medline](#)
- Bouret S, Sara SJ (2004) Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning. *Eur J Neurosci* 20:791–802. [CrossRef](#)
- Bouret S, Sara SJ (2005) Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends Neurosci* 28:574–582. [CrossRef Medline](#)
- Bouret S, Richmond BJ (2009) Relation of locus coeruleus neurons in monkeys to Pavlovian and operant behaviors. *J Neurophysiol* 101:898–911. [Medline](#)
- Bouret S, Richmond BJ (2010) Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci* 30:8591–8601. [CrossRef Medline](#)
- Bouret S, Duvel A, Onat S, Sara SJ (2003) Phasic activation of locus coeruleus

- neurons by the central nucleus of the amygdala. *J Neurosci* 23:3491–3597. [Medline](#)
- Bouret S, Ravel S, Richmond BJ (2012) Complementary neural correlates of motivation in dopaminergic and noradrenergic neurons of monkeys. *Front Behav Neurosci* 6:40. [CrossRef Medline](#)
- Chamberlain SR, Müller U, Blackwell AD, Clark L, Robbins TW, Sahakian BJ (2006) Neurochemical modulation of response inhibition and probabilistic learning in humans. *Science* 311:861–863. [CrossRef](#)
- Chiba T, Kayahara T, Nakano K (2001) Efferent projections of infralimbic and prelimbic areas of the medial prefrontal cortex in the Japanese monkey, *Macaca fuscata*. *Brain Res* 888:83–101. [CrossRef Medline](#)
- Clayton EC, Rajkowski J, Cohen JD, Aston-Jones G (2004) Phasic activation of monkey locus ceruleus neurons by simple decisions in a forced-choice task. *J Neurosci* 24:9914–9920. [CrossRef Medline](#)
- Flagel SB, Clark JJ, Robinson TE, Mayo L, Czuj A, Willuhn I, Akers CA, Clinton SM, Phillips PE, Akil H (2011) A selective role for dopamine in stimulus-reward learning. *Nature* 469:53–57. [CrossRef Medline](#)
- Footo SL, Aston-Jones G, Bloom FE (1980) Impulse activity of locus coeruleus neurons in awake rats and monkeys is a function of sensory stimulation and arousal. *Proc Natl Acad Sci U S A* 77:3033–3037. [CrossRef Medline](#)
- Heilbronner SR, Platt ML (2013) Causal evidence of performance monitoring by neurons in posterior cingulate cortex during learning. *Neuron* 80:1384–1391. [CrossRef Medline](#)
- Kalwani RM, Joshi S, Gold JJ (2014) Phasic activation of individual neurons in the locus ceruleus/subceruleus complex of monkeys reflects rewarded decisions to go but not stop. *J Neurosci* 34:13656–13669. [CrossRef Medline](#)
- Lebreton M, Jorge S, Michel V, Thirion B, Pessiglione M (2009) An automatic valuation system in the human brain: evidence from functional neuroimaging. *Neuron* 64:431–439. [CrossRef Medline](#)
- Lim SL, O'Doherty JP, Rangel A (2011) The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *J Neurosci* 31:13214–13223. [CrossRef Medline](#)
- Liu Z, Richmond BJ, Murray EA, Saunders RC, Steenrod S, Stubblefield BK, Montague DM, Ginns EI (2004) DNA targeting of rhinal cortex D2 receptor protein reversibly blocks learning of cues that predict reward. *Proc Natl Acad Sci U S A* 101:12336–12341. [CrossRef Medline](#)
- McGaughy J, Ross RS, Eichenbaum H (2008) Noradrenergic, but not cholinergic, deafferentation of prefrontal cortex impairs attentional set-shifting. *Neuroscience* 153:63–71. [CrossRef Medline](#)
- Minamimoto T, Hori Y, Kimura M (2005) Complementary process to response bias in the centromedian nucleus of the thalamus. *Science* 308:1798–1801. [CrossRef Medline](#)
- Minamimoto T, La Camera G, Richmond BJ (2009) Measuring and modeling the interaction among reward size, delay to reward, and satiation level on motivation in monkeys. *J Neurophysiol* 101:437–447. [Medline](#)
- Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD (2006) Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 442:1042–1045. [CrossRef Medline](#)
- Robbins TW, Arnsten AFT (2009) The neuropsychopharmacology of fronto-executive function: monoaminergic modulation. *Annu Rev Neurosci* 32:267–287. [CrossRef Medline](#)
- Sara SJ, Bouret S (2012) Orienting and reorienting: the locus coeruleus mediates cognition through arousal. *Neuron* 76:130–141. [CrossRef Medline](#)
- Tait DS, Brown VJ, Farovik A, Theobald DE, Dalley JW, Robbins TW (2007) Lesions of the dorsal noradrenergic bundle impair attentional set-shifting in the rat. *Eur J Neurosci* 25:3719–3724. [CrossRef Medline](#)
- Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN (2005) A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 93:1074–1089. [CrossRef Medline](#)
- Valentino RJ, Van Bockstaele E (2008) Convergent regulation of locus coeruleus activity as an adaptive response to stress. *Eur J Pharmacol* 583:194–203. [CrossRef Medline](#)
- Ventura R, Morrone C, Puglisi-Allegra S (2007) Prefrontal/accumbal catecholamine system determines motivational salience attribution to both reward- and aversion-related stimuli. *Proc Natl Acad Sci U S A* 104:5181–5186. [CrossRef Medline](#)
- Ventura R, Latagliata EC, Morrone C, La Mela I, Puglisi-Allegra S (2008) Prefrontal norepinephrine determines attribution of “high” motivational salience. *PLoS One* 3:e3044. [CrossRef Medline](#)
- Vertes RP (2004) Differential projections of the infralimbic and prelimbic cortex in the rat. *Synapse* 51:32–58. [CrossRef Medline](#)
- Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. *Neuron* 46:681–692. [CrossRef Medline](#)