

Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function

MARK S. GILZENRAT

Princeton University, Princeton, New Jersey

SANDER NIEUWENHUIS AND MARIEKE JEPMA

Leiden University, Leiden, The Netherlands

AND

JONATHAN D. COHEN

Princeton University, Princeton, New Jersey

An important dimension of cognitive control is the adaptive regulation of the balance between exploitation (pursuing known sources of reward) and exploration (seeking new ones) in response to changes in task utility. Recent studies have suggested that the locus coeruleus–norepinephrine system may play an important role in this function and that pupil diameter can be used to index locus coeruleus activity. On the basis of this, we reasoned that pupil diameter may correlate closely with control state and associated changes in behavior. Specifically, we predicted that increases in baseline pupil diameter would be associated with decreases in task utility and disengagement from the task (exploration), whereas reduced baseline diameter (but increases in task-evoked dilations) would be associated with task engagement (exploitation). Findings in three experiments were consistent with these predictions, suggesting that pupillometry may be useful as an index of both control state and, indirectly, locus coeruleus function.

Is it better to spend your time revising a manuscript that is likely to be published or attending a seminar on a topic that is unfamiliar but relevant to your work? The former holds well-defined prospects for reward, whereas the latter may open you up to new opportunities. This dilemma illustrates what is often referred to as the *exploit/explore trade-off*: whether to continue pursuing a known source of reward (exploit) or search for new ones (explore). We confront it in all aspects of our lives and at all levels of our decision making: whether to revisit a favorite restaurant or try a new one, interact with a child in a familiar manner or shift to a new one, persist in a line of research or pursue a new one. Balancing this trade-off involves regulating the extent to which the control state of the system favors task engagement (exploitation) versus disengagement from the current task and the sampling of other inputs and/or behaviors (exploration). Although this regulation is a critical form of higher level control, little research has addressed how it is achieved (Cohen, McClure, & Yu, 2007). Here, we draw upon recent work suggesting that the locus coeruleus–norepinephrine (LC–NE) system plays an important role in this regulation. In particular, recent studies have revealed a close relationship between LC activity and behavioral performance indicative of shifts in control state

between exploitation and exploration. In addition, neurophysiological studies in nonhuman primates have revealed a tight correlation between LC activity and pupil diameter. On the basis of these observations, we reasoned that pupil diameter may provide a useful, noninvasive index of changes in control state. Below, we review research concerning the role of the LC–NE system in regulating the exploit/explore balance and its relationship to pupil diameter. We then report three experiments designed to test theory-based predictions regarding the relationship of pupillary dynamics to both naturally occurring and experimentally induced changes in control state and behavioral performance.

The Role of the LC–NE System in Regulating Control State

The LC is a neuromodulatory nucleus, situated in the dorsal pons, with widely distributed, ascending projections throughout the forebrain that are responsible for all of the NE release in the neocortex. Traditionally, investigators have associated the LC–NE system with relatively general functions, such as arousal and the sleep–wake cycle (Aston-Jones, Foote, & Bloom, 1984; Jouvet, 1969). However, recent neurophysiological studies (e.g., Aston-

Jones, Rajkowski, & Kubiak, 1997; Clayton, Rajkowski, Cohen, & Aston-Jones, 2004) and computational modeling studies (Gilzenrat, Holmes, Rajkowski, Aston-Jones, & Cohen, 2002; McClure, Gilzenrat, & Cohen, 2005; Shea-Brown, Gilzenrat, & Cohen, 2008; Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999; Yu & Dayan, 2005) have suggested a more specific role for the LC–NE system in regulating the exploit/explore balance. The adaptive gain theory (AGT; Aston-Jones & Cohen, 2005) formalizes this idea, proposing that the LC–NE system responds to changes in task utility by altering gain modulation of cortical processing mechanisms responsible for task performance, which, in turn, influences the exploit/explore balance. This theory builds on two key observations: one concerning the modes of LC function, and the other concerning the influence of NE release on cortical processing.

LC activity appears to exhibit two modes of function: phasic and tonic (Usher et al., 1999). Phasic firing typically occurs in response to task-relevant events during epochs of high performance and lower baseline LC activity (Aston-Jones, Rajkowski, Kubiak, & Alexinsky, 1994). In contrast, the tonic mode is associated with elevated baseline firing rate, absence of phasic responses, and degraded task performance (Aston-Jones et al., 1994).¹ On the basis of the neuromodulatory effects of NE release on cortical processing, AGT proposes that the phasic mode promotes exploitation, whereas the tonic mode promotes disengagement from the current task in favor of exploration (Aston-Jones & Cohen, 2005).

Although the precise biophysical mechanisms remain unclear, NE effects can be described computationally as increased responsivity (gain modulation) of units to their inputs (Moxon, Devilbiss, Chapin, & Waterhouse, 2007; Servan-Schreiber, Printz, & Cohen, 1990). As proposed by AGT and supported by modeling work (e.g., Gilzenrat et al., 2002; Usher et al., 1999), when this effect is timed to task-relevant events (LC phasic mode), it facilitates execution of task-relevant processes (exploitation). However, indiscriminate uniform increases in gain (LC tonic mode) facilitate processing of stimuli and actions not relevant to the current task. This effect is akin to introducing processing noise, thereby favoring exploration.

AGT has received support from direct LC neuronal recordings in both monkeys (Aston-Jones et al., 1997; Aston-Jones et al., 1994; Rajkowski, Kubiak, & Aston-Jones, 1993) and rats (Bouret & Sara, 2005). For example, Aston-Jones et al. (1994) showed that in monkeys performing a simple target detection task, decreased baseline LC activity and increased phasic responses to task-related events (LC phasic mode) were closely associated with improved task performance (e.g., as indicated by signal detection measures), whereas elevated baseline activity and the absence of phasic responses (LC tonic mode) were associated with degraded task performance and greater distractibility. AGT further proposes that LC-mediated shifts in control state are driven by changes in task utility, with high utility favoring exploitation (phasic mode) and low utility favoring

exploration (tonic mode). Consistent with this hypothesis, anatomical studies have shown that the primary neocortical projections to LC come from the orbitofrontal and anterior cingulate cortex (Aston-Jones et al., 2002; Rajkowski, Lu, Zhu, Cohen, & Aston-Jones, 2000; Zhu, Iba, Rajkowski, & Aston-Jones, 2004)—areas known to be responsive to task-related rewards and costs of performance (Botvinick, 2007; Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004)

The LC–NE System and Pupil Diameter

In follow-up analyses of the monkey target detection experiment discussed above, Rajkowski et al. (1993) found that baseline pupil diameter (recorded at the onset of each trial of the experiment) closely tracked LC tonic discharge frequency. Figure 1 provides an illustration of this strong correlation between LC discharge frequency and pupil diameter. Rajkowski and colleagues concluded that pupil diameter varies with LC mode, such that the LC tonic mode is marked by a relatively large baseline pupil diameter and the LC phasic mode is marked by a relatively smaller baseline pupil diameter. There is also more recent evidence that baseline pupil diameter increases with tonic LC activity: Spontaneous and drug-induced drowsiness and other low-arousal states, which are characterized by low tonic LC activity, are accompanied by a reduced baseline pupil diameter (e.g., Hou, Freeman, Langley, Szabadi, & Bradshaw, 2005; Morad, Lemberg, Yofe, & Dagan, 2000). Conversely, noradrenergic drugs that increase arousal and tonic LC activity also increase baseline pupil diameter (Phillips, Szabadi, & Bradshaw, 2000). Finally, a large number of studies have shown that task processing is accompanied by rapid and dramatic pupil dilation (Beatty, 1982a, 1982b; Einhäuser, Stout, Koch, & Carter, 2008; Richer & Beatty, 1987), which is consistent with the occurrence of an LC phasic response to task-relevant events. Together, these findings suggest that pupil diameter may track instantaneous LC activity, such that baseline pupil diameter corresponds to LC tonic firing rate, and task-evoked dilations correspond to LC phasic activity. Accordingly, we reasoned that baseline pupil diameter and task-evoked dilations should be inversely correlated, corresponding to the reciprocal relationship observed between LC tonic and phasic activity. Furthermore, pupillary responses should be correlated with the changes in control state and behavioral profiles associated with each mode of LC function, as predicted by AGT.

It is important to note that, despite the close relationship between pupil diameter and LC activity, the mechanisms underlying this relationship are not yet understood. There are presently no known anatomic pathways that might mediate this relationship directly, which suggests that these effects may reflect parallel downstream influences of a common source. An important candidate for such a source is the paragigantocellularis (PGi) nucleus of the ventral medulla. The PGi is a critical relay for the sympathoexcitatory efferents of the hypothalamus on their way to the autonomic centers that subserve re-




Figure 1. Relationship between tonic pupil diameter and baseline firing rate of a locus coeruleus (LC) neuron in the monkey (adapted with permission from Rajkowski, Kubiak, & Aston-Jones, 1993). Pupil diameter measurements were taken by a remote eyetracking camera at each instant in time at which the monkey achieved fixation of a visual spot during the target detection task (see the text for details). Note the close positive relationship between pupil diameter and the rate of LC activity.

flexive pupillary dilation (Hilton & Smith, 1984; Loewy, Wallach, & McKellar, 1981). Importantly, the PGI is also one of the major afferents to the LC (Aston-Jones, Ennis, Pieribone, Nickell, & Shipley, 1986). Thus, the PGI is thought to be a key afferent branch point for sympathetic activation that strongly innervates both autonomic nuclei and the LC (e.g., Aston-Jones et al., 1991). This notion is further supported by observations that LC discharge parallels sympathetic nerve activity (e.g., Elam, Svensson, & Thoren, 1986). In summary, although it does not appear that pupil diameter is under direct control of the LC, converging empirical evidence suggests that it is closely correlated with LC activity and, thus, may be useful as a *reporter variable*.

The Present Research

In this article, we present three experiments that investigated the relationship of pupil diameter to control state, corresponding to the relationship between LC activity and control state predicted by AGT.

In Experiment 1A, we used a target detection task comparable to that used in the monkey studies described above (Aston-Jones et al., 1994; Rajkowski et al., 1993) to test the prediction that, as in monkeys, trial-to-trial changes in baseline pupil diameter and task-evoked pupil dilations are negatively correlated with each other. We also tested the prediction that these pupillary responses are associated with behaviors indicative of differences in control state (i.e., task engagement vs. disengagement) corresponding to the two modes of LC function (phasic and tonic, respectively). In Experiment 2, we attempted to manipulate control state directly by regulating conflict and reward across blocks of trials in a pitch discrimination task. AGT proposes that the level of task

engagement is determined by online assessment of task utility, which is jointly determined by the costs (here, amount of conflict) and reward value of the task. Thus, we examined whether pupil size and task performance varied with changes in task utility, as predicted by AGT. Finally, in Experiment 3, we used a novel diminishing-utility task that set increasing task difficulty (producing conflict) against increasing potential reward, resulting in dynamic changes in task utility. Furthermore, participants were given an opportunity to disengage from the task and reset the level of task difficulty/reward when the perceived utility of remaining engaged began to decline (i.e., conflict began to exceed reward). The goal of this experiment was to provide a more ecologically valid context for inducing changes in control state and to allow participants to emit a behavior explicitly indicative of task disengagement. We tested predictions about the effect of changes in task utility on behavior and pupillary response.

EXPERIMENT 1A

In Experiment 1A, we attempted to reproduce and expand on the pupillometric results from monkey studies reported by the Aston-Jones group (Rajkowski et al., 1993). Following methods that have been used in a number of their neurophysiological studies (Aston-Jones et al., 1997; Aston-Jones et al., 1994; Rajkowski et al., 1993), we administered a target detection (*oddball*) task to humans and recorded pupil diameter throughout the task. In contrast, however, we presented our stimuli in the auditory modality to avoid potential confounds in pupil measurement caused by the presentation of visual stimuli (Janisse, 1977).

In contrast to monkeys, our human participants showed little variation or periodicity in performance during the task. Therefore, it was not feasible to determine periods of "good" and "poor" performance, as has been done in other studies (cf. Aston-Jones et al., 1994). Instead, we examined whether extremes in baseline pupil size showed a systematic relationship with performance and pupil dilation. To this end, the trials of each participant were segregated into those preceded by small baseline pupils (defined as observations in the 25th percentile of that participant's baseline pupil diameters or smaller) and large baseline pupils (defined as observations in the 75th percentile of that participant's baseline pupil diameters or larger).

We predicted that, as in Rajkowski et al. (1993), larger baseline pupil diameters would be associated with poorer performance (higher false alarm rate and longer, more variable response times [RTs]), indicative of lapses of engagement (reflecting periodic drifts into LC tonic mode). Conversely, we predicted that smaller baseline pupil diameters would be associated with better performance, indicative of task engagement (mediated by the LC phasic mode). In addition, we predicted that larger baseline diameters would be associated with smaller posttarget dilations, and vice versa.

Method

Participants. The participants were 23 Princeton University undergraduate and graduate students, 9 of them female, with ages ranging from 19 to 26 years.

Apparatus. Stimuli were delivered and controlled via a PC running E-Prime (Psychology Software Tools, Pittsburgh, PA) stimulus presentation software. Pupil diameter and eye position were recorded using the ETL-500 eyetracking laboratory from ISCAN (Cambridge, MA). The participants wore a lightweight cap outfitted with an infrared illuminator and a miniaturized infrared-sensitive camera that recorded an image of the eye via a dichroic "hot" mirror that sat in front of the participant's left eye. This mirror was transparent to visible light, permitting visibility, but reflective to infrared. A separate scene camera, also mounted on the cap, allowed the measurement of point of regard to confirm that gaze was maintained within a prescribed bounding box. Eye images were sampled by the eye camera at 60 Hz. Event markers transmitted to the eyetracking computer by the stimulus presentation computer allowed coordination between behavioral events and eye data for analysis. Pupil data were postprocessed and analyzed using MATLAB (The MathWorks, Natick, MA). Artifacts and blinks were removed using a linear interpolation algorithm based on Siegle (1999).

Stimuli and Procedure. The participants performed an auditory oddball discrimination task while their left pupil was continuously measured. The participants were seated at a distance of 60 cm from a computer monitor displaying a blank medium gray field (12 cd/m² luminance). They were instructed to hold gaze within a central fixation square delineated by a thin black border subtending 10° of visual angle. The participants were instructed that trials would not proceed unless their gaze was maintained within the fixation square, as confirmed by eyetracking data. Subsequently, they received several practice trials to familiarize them with the task. Practice was repeated until the participants performed all the practice trials correctly.

The stimuli consisted of a series of sinusoidal tones (72 dB, 250 msec), played through headphones. The participants were instructed to make speeded keypress responses with the dominant hand to target tones (880 Hz) and to ignore distractor tones (800 Hz). Each trial began with the presentation of one of these two tones, followed

by a 1-sec time window during which the response had to be made. The time intervals between subsequent tones were roughly normally distributed, with a mean of 4.2 sec and a range of 2.5–5.2 sec. RT and accuracy were recorded on each trial. No feedback was given during task performance.² The participants performed 750 trials: 150 (20%) target trials and 600 (80%) distractor trials. The time to complete the task was about 50 min.

For each trial, the baseline pupil diameter prior to trial onset was assessed, as was the peak magnitude of the pupil dilation following the tone. To determine baseline pupil diameter, the 1 sec of pupil data collected immediately preceding tone presentation was averaged. Dilation of the pupil was measured as the highest deviation from this baseline in the 2.5 sec following tone onset. All reported *t* tests are one-sided, except where indicated.

Because the interstimulus intervals were relatively short (to maintain homology with the animal experiments), measures of baseline pupil diameter on trial *n* were potentially confounded by pupil dilations on trial *n*–1. To directly address this concern, we computed the average stimulus-evoked pupil dilation on trials preceding the large baseline pupil trials (0.13 mm) and trials preceding the small baseline pupil trials (0.06 mm). Although this difference was highly significant [$t(22) = 4.5, p < .001$], it constituted only 5% of the corresponding difference in baseline pupil diameter on trial *n* (5.1 vs. 3.7 mm). Furthermore, individual differences in the magnitude of the pupil dilation effect on trial *n*–1 did not reliably correlate with any of the reported pupil or behavioral effects on trial *n* (all *ps* > .1). These results indicate that our selections of large and small baseline pupil trials, and their relation with target-evoked pupil dilation and behavioral measures, were largely uninfluenced by pupil dilations on the previous trial. Nonetheless, to correct for any carryover effects, we included the average magnitude of pupil dilations on trial *n*–1 as a covariate in the analyses that are reported below.

Results

The participants performed near ceiling overall, with a mean false alarm rate (i.e., percentage of responses to distractor tones) of 0.6% and a hit rate of 99.4%. Mean RT across all target trials was 484 msec.

The results of comparisons between the large and small baseline pupil trials were generally in line with our predictions. Large baseline pupils were associated with a reliably higher false alarm rate than were small pupils [1.3% vs. 0.3%; $t(22) = 1.9, p < .05$]. The number of misses (i.e., no response on target trials) was also higher for large baseline pupils than for small baseline pupils (0.8% vs. 0.6%), but this difference did not reach significance [$t(22) = 1.4, p = .08$]. Five participants emitted false alarms and misses during both large and small baseline pupil trials, allowing for the computation of the signal detection parameters *d'* (sensitivity) and β (criterion). For these participants, the average *d'* was higher for small than for large baseline pupils [4.86 vs. 4.50; $t(4) = 2.3, p < .05$]. The average β s for small and large pupils showed the predicted ordering (1.32 vs. 0.82), but this was not reliable [$t(4) = 1.6, p = .09$].

The RT data were also consistent with our predictions. Figure 2 depicts the RT distributions for small and large baseline pupils, averaged across participants. The RT distributions were constructed using Vincentizing (Van Zandt, 2000; Vincent, 1912), a quantile averaging method that produces representative group distributions while absorbing interparticipant variation. The two distributions were of comparable form, but the distribution for large baseline pupil trials had a somewhat longer tail, resulting




Figure 2. Variable-width histograms constructed from Vincentized response time distributions (Ratcliff, 1979) in Experiment 1A for trials associated with small (top panel) and large (bottom panel) baseline pupil diameters.

in a slightly larger mean and greater dispersion. Average RTs were reliably shorter for small than for large baseline pupils [483 vs. 499 msec; $t(22) = 1.7, p = .05$]. Furthermore, the dispersion of RTs was tighter for small than for large baseline pupils [SDs 93 vs. 103 msec; $t(22) = 1.7, p < .05$].

Finally, we consider the pupil dilations evoked by the target and distractor tones. Consistent with basic findings in pupillometry (Steinhauer & Zubin, 1982; Van Olst, Heemstra, & Ten Kortenaar, 1979; reviewed in Beatty, 1982b), average maximal pupil dilations were far larger for detected targets (hits) than for correctly rejected distractors [0.21 vs. 0.06 mm; $t(22) = 12.3, p < .0001$; see Figure 3]. We further predicted that small baseline pupils would be accompanied by larger task-evoked dilations, as compared with large baseline pupils. As is shown in Figure 3, this prediction was confirmed: Peak dilations to targets were larger for small than for large baseline pupil trials [0.41 vs. 0.09 mm; $t(22) = 7.0, p < .0001$].

Discussion

The results of Experiment 1A are in line with our predictions. Specifically, we observed that small baseline pupils were associated with better performance (i.e., fewer false alarms and shorter, less variable RTs), as well as with larger task-evoked dilations. Conversely, large baseline pupils were associated with poorer performance and attenuated task-evoked dilations. The observed relationship between behavior and pupil dynamics parallels the results obtained in studies with monkeys, using a similar target detection task (Aston-Jones et al., 1994;

Rajkowski et al., 1993). Furthermore, the relationship of pupillary responses (baseline diameter and task-evoked dilations) to behavior observed here closely paralleled the relationship of LC pattern of activity (tonic vs. phasic) to behavior observed in the monkey. This provides evidence—albeit indirect—that pupillary responses may track LC activity in the human, as has been observed directly in the monkey.

A potential weakness of this experiment concerned our choice of relatively short intertrial intervals. As a result, our measures of baseline pupil diameter may have been confounded by pupil dilations on the preceding trial. Although we verified that the magnitude of the previous trial dilation explained only a small part of the variance in baseline pupil size and our dependent measures and we statistically controlled for this confounding variable (see the Method section), the fact remains that baseline pupil diameter could not be properly estimated. In Experiment 2, we therefore substantially increased the intertrial intervals (4.0–8.0 sec). The intertrial intervals in Experiment 3 were similar to those in Experiment 1A, but the results in Experiment 3 explicitly demonstrated a negative correlation between pupil dilation on the previous trial and baseline pupil size on the current trial—a pattern that is inconsistent with the notion of confounding carryover effects but consistent with our prediction that these variables show opposite trends as control state gradually changes.

Another limitation of this experiment is that the level of task engagement was not experimentally controlled. Rather, engagement was allowed to vary spontaneously throughout the task in a fashion homologous to that in

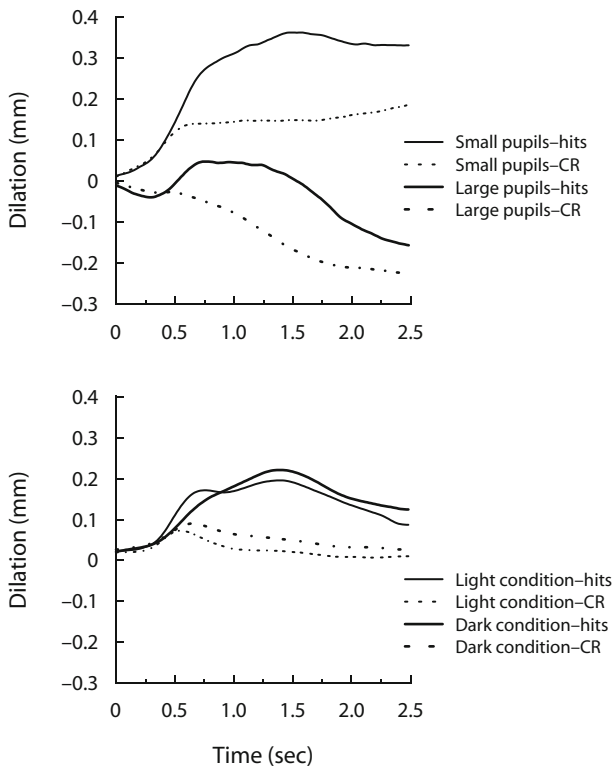


Figure 3. Time courses of grand-averaged poststimulus pupil dilations for correctly detected targets (hits) and correctly rejected (CR) distractors. For Experiment 1A (top panel), dilations are plotted separately for trials that were preceded by large and small baseline pupil diameters. For Experiment 1B (bottom panel), dilations are plotted separately for the light and dark conditions.

animal experiments (Aston-Jones et al., 1997; Aston-Jones et al., 1994; Rajkowski et al., 1993). However, our human participants showed far less variability in performance and maintained performance near ceiling, exhibiting fairly few instances of behavior that we would consider indicative of a change in control state. Although we observed the predicted direction of behavior when comparing trials with the largest and smallest baseline pupil sizes, the effects were at best modest. To address this shortcoming, in Experiments 2 and 3, we attempted to manipulate control state through independent variables that we hypothesized would directly bear on task engagement. But first, we report a control experiment that we conducted to rule out an alternative explanation of the results in Experiment 1A.

EXPERIMENT 1B

Although the inverse relationship between baseline and task-evoked pupil measures observed in Experiment 1A conforms to our predictions regarding changes in control state (and underlying changes in the LC mode of activity), the law of initial values (Lacey, 1956) offers an alternative explanation for the differential dilations observed in small versus large baseline pupil trials. According to this law, the size of a physiological response to a stimulus is often af-

ected by the prestimulus baseline level of that response system. If the baseline level is high, the expected increase in the physiological response to a stimulus will be smaller, because of a ceiling effect; in contrast, if the baseline level is low, a larger response may be produced, because of the greater potential “upward” range available. Thus, in Experiment 1A, the pupil may not have dilated to a greater extent in large baseline trials simply because the pupil was already relatively dilated and its physiology limited further dilation. Indeed, since the area of the pupil increases as a square of its radius, a much smaller change in diameter is required at larger pupil sizes to obtain an equivalent change in area, meaning that an equivalent dilatory impulse may have far less impact in terms of diameter in larger baseline pupils (thus appearing as a smaller dilation).

To address this possibility, we conducted a control experiment examining whether we would find a similar inverse relationship between baseline pupil size and the magnitude of pupil dilations for groups of trials that differed in baseline pupil size due to a systematic manipulation of luminance (as opposed to trial-to-trial fluctuations in control state). If we failed to replicate the inverse relationship in this control experiment, this would be consistent with our explanation that the pattern observed in Experiment 1A specifically reflected an effect of control state. In contrast, if the results of the control experiment showed a clear inverse relationship, this would support an alternative explanation in terms of intrinsic properties of pupillary responses, rather than a relationship to control state. In Experiment 1B, participants performed the target detection task used in Experiment 1A. During the experiment, we adaptively changed the screen luminance using a staircase-tracking procedure. The aim of this procedure was to induce average baseline pupil sizes in the first and second halves of the experiment that differed by the same amount as the critical effect size in Experiment 1A. We then compared the task-evoked pupil dilations for the *dark* and *light* conditions.

Method

All details were the same as those in Experiment 1A, with the following exceptions. The participants were 8 volunteers, 6 of them female, with ages ranging from 18 to 26 years (mean age = 20.9). Pupil diameter and eye position were recorded using a Tobii T120 eyetracker, which was integrated into a 17-in. TFT monitor (Tobii, 2007). Pupil data were postprocessed and analyzed using BrainVision Analyzer. The participants performed two blocks of 190 trials each: one with the light condition and one with the dark condition. The order of the two conditions was counterbalanced across participants. Each block contained 38 (20%) target trials and 152 (80%) distractor trials. The time to complete the task was about 30 min.

Before each block, we adapted the RGB values of the screen color to determine the desired baseline pupil diameter: 3.7 mm in the light condition and 5.1 mm in the dark condition. The R, G, and B values were kept equal such that only grayish colors were used. If needed, we also adapted the room illumination. For some participants, it was not possible to reach a pupil size of 5.1 mm. For these participants, we used smaller pupil criteria, but the difference between the light and dark conditions was always 1.4 mm.

To maintain the desired baseline pupil diameters throughout each block, we used a staircase tracking procedure. On each trial, the pupil diameter during the 1 sec preceding each tone was measured, and a running average of these values was kept. If the running aver-

age value was within ± 0.05 mm of the desired pupil diameter, the screen color was not changed. If the running average went above this range, the RGB values were increased by 10 points (lighter). If the running average value went below this range, the RGB values were decreased by 10 points (darker).

Results

As in Experiment 1A, the participants performed near ceiling, with a mean overall false alarm rate (i.e., percentage of responses to distractor tones) of 0.4% and a hit rate of 96.9%. Mean RT across all target trials was 588 msec. There were no reliable differences in performance between the light and the dark conditions, with one exception: The dispersion of RTs was tighter in the light than in the dark condition [*SDs* 86 vs. 109 msec; $t(7) = 2.6$, $p < .05$, two-sided].

As guaranteed by our procedure, the average baseline pupil diameters in the light and dark conditions differed by approximately 1.4 mm: 3.26 mm ($SD = 0.37$ mm) in the light condition and 4.64 mm ($SD = 0.35$ mm) in the dark condition. Thus, the difference in baseline pupil size between the two conditions was the same as the difference between the small and large baseline pupil quartiles in Experiment 1A. Importantly, as is shown in Figure 3, peak dilations to targets did not differ between the light and dark conditions [0.26 vs. 0.25 mm; $t(7) = 0.56$, $p = .59$, two-sided].

Discussion

In this experiment, we tested whether the law of initial values (Lacey, 1956) would provide an alternative explanation for the results of Experiment 1A. The findings are inconsistent with the idea that the inverse relationship between baseline pupil size and the magnitude of task-evoked dilations in Experiment 1A was an artifact of the dynamic range of the pupil. If this had been the case, we would have found a similar inverse relationship when comparing pupil dynamics in dark and light luminance conditions. However, we found no sign of an inverted relationship under such conditions (for similar results, see Bradshaw, 1969). Instead, we propose that the results in Experiment 1A provide evidence of a systematic relationship between the pupillary response pattern and control state. Once again, this relationship conforms to predictions made by AGT about the relationship between the LC mode of activity and control state, under the assumption that pupillary responses reflect LC activity.

In the two experiments reported below, we also attempted to rule out the law of initial values as an alternative interpretation of the results. To ensure that the transient pupil dilations observed over the course of the experiments were well within the pupil's dynamic range, we took reference measurements of the pupil diameter in darkness for a subset of our participants. The experiments themselves were conducted at a moderate level of illumination. To foreshadow the results, the reference measurements confirmed that the maximal dilations observed during the experiment were well within the physiological range of the pupil. Although it is still possible that the

pupil's range is bounded for a given level of illumination, these reference measures provided at least a first-order check that the effects of our manipulations were well within the pupil's overall dynamic range.

EXPERIMENT 2

The goal of Experiments 2 and 3 was twofold. First, we sought to demonstrate that, as predicted by AGT, control state can be directly influenced by manipulating task utility (i.e., the costs and rewards of performance). Second, we sought to determine whether manipulations of task utility would also produce the pupillary dynamics predicted to accompany corresponding changes in control state. In each case, we used AGT to make predictions about the expected relationship between changes in task utility and the LC mode of activity associated with the corresponding control state and then translated these into predictions about patterns of pupil response under the assumption that these tracked predicted changes in LC activity.

In Experiment 2, we attempted to induce particular control states by manipulating the conflict (i.e., costs) and reward experienced by participants in each of four blocks of a pitch discrimination task—a modified version of the task used by Kahneman and Beatty (1967). We manipulated the degree of conflict (high or low) experienced in a block by manipulating the frequency of occurrence of impossible-discrimination trials—trials in which the reference tone and comparison tone were of equal pitch and, thus, discrimination was at chance. Reward was manipulated by varying the feedback participants received on these impossible-discrimination trials. Specifically, in positive feedback blocks, participants always received positive feedback to their responses on these impossible-discrimination trials, whereas in negative feedback blocks, participants always received negative feedback on these trials.

We predicted that both the high-conflict/positive-feedback block and the low-conflict/negative-feedback block would promote task engagement (LC phasic mode), but for complementary reasons. In the former condition, an additional recruitment of control is encouraged, because the high amount of experienced conflict is presumably offset by the reward associated with the positive feedback. In the latter condition, the negative feedback received on impossible-discrimination trials may be interpreted as a signal that current levels of control are not sufficient, but these conflict trials are not encountered frequently enough to indicate that control is futile. Thus, we predicted that in both of these blocks, an increase in control would be an adaptive response, given the perceived task utility (see Aston-Jones & Cohen, 2005, Figure 10).

In contrast, we predicted that the high-conflict/negative-feedback block would promote task disengagement (LC tonic mode). This response would be adaptive, given that potential attempts to recruit additional control would consistently and frequently fail. Under such circumstances of low task utility, the participant might ordinarily benefit from task disengagement and exploration (although, as we will discuss below, the task used here does not allow such

behaviors). Finally, for completeness of the fully factorial design, we included a low-conflict/positive-feedback block. This block does not strongly motivate a change in control, since impossible-discrimination trials are encountered infrequently and are consistently rewarded. Thus, we had no strong prediction about the control state engendered by this block, except that it would likely be intermediate between the two engagement-promoting blocks and the disengagement-promoting block.

We expected to observe the pattern of pupillary responses associated with the control states predicted for each block. Thus, in the engagement-promoting blocks, we predicted relatively small baseline pupil diameters and relatively large transient dilations to the comparison tones (corresponding to the LC phasic mode). Conversely, in the disengagement-promoting blocks, we predicted relatively large baseline pupil diameters and relatively small task-evoked dilations (corresponding to the LC tonic mode). In the low-conflict/positive-feedback block, we expected intermediate pupil responses. Finally, we expected that RTs would be shorter and accuracy would be higher in the engagement-promoting blocks than in the disengagement-promoting block.

Method

Participants. The participants were 24 Princeton University undergraduate students, 14 of them female, with ages ranging from 19 to 22 years.

Apparatus. The same setup was used as that in Experiment 1A.

Stimuli and Procedure. The experimental procedure was the same as that in Experiment 1A, except as noted below. The participants were presented sequences of two tones and were instructed to indicate whether the second of the two tones was higher or lower in pitch than the first. Each trial began with an 850-Hz reference tone presented for 250 msec. This tone was followed 3 sec later by the comparison tone, which ranged from 820 to 880 Hz in steps of 10 Hz. The participants were instructed to respond as quickly and accurately as possible upon hearing the comparison tone. Four seconds after the comparison tone, the participants received a 250-msec feedback sound that informed them of their accuracy. The feedback was followed by a variable intertrial interval, chosen randomly between 4 and 8 sec. Prior to the start of the experimental session, the participants viewed visually presented instructions and were given a short block of practice trials at easiest discriminability to familiarize them with the stimuli. Practice was repeated until the participants performed correctly on all the practice trials.

The participants performed four blocks of 36 trials, with each block lasting 9–10 min. Conflict (low vs. high) and feedback (positive vs. negative) were varied independently across blocks, employing a fully factorial 2×2 design. Block order was randomized across participants, using a balanced Latin square design. The first 12 trials of each block were designed to extinguish any carryover effects from the previous block and were discarded from the analyses. During these trials, the participants never received an impossible-discrimination (i.e., 850-Hz) comparison tone. The remaining 24 trials of each block contained the experimental manipulation. In two of the four blocks, the participants received 6 trials (25%) with impossible-discrimination comparison tones (low-conflict condition). The other two blocks contained 12 (50%) impossible-discrimination trials (high-conflict condition). On all other trials, the comparison tone was selected randomly without replacement from the set [820, 830, 840, 860, 870, and 880 Hz], so that the participants encountered all of these comparison tones with equal likelihood. Furthermore, in two of the four blocks, the participants always re-

ceived positive feedback on impossible-discrimination trials. In the other two blocks, the participants always received negative feedback on impossible-discrimination trials.

The pretrial baseline pupil diameter and the maximal pupil dilation evoked by the comparison tone were determined as in Experiment 1A. Pupil dilation and RT are influenced by the pitch difference between the comparison tone and the reference tone (Kahneman & Beatty, 1967). Because the four blocks differed with regard to the distribution of pitch differences, this presented a confound in our design. To control for this confound, we computed weighted averages of pupil dilation size and RT by averaging impossible-discrimination trials and all other trials separately and then averaging the resulting two averages. For baseline pupil diameter, we computed unweighted averages, because the baseline occurred before the tones were presented. Average accuracy was determined after excluding impossible-discrimination trials.

For 12 of the participants, an additional reference procedure was conducted prior to the beginning of the experiment. Pupil diameter was recorded continuously for 30 sec in darkness (room illumination, <1 lux; monitor luminance, 0.02 cd/m²). We then averaged the 5 sec of maximal pupil size recorded during this period, allowing for full acclimation to the low level of illumination, and discarded artifacts. The experiment itself was conducted at a slightly dimmed illumination level (room illumination, 77 lux; monitor luminance, 12 cd/m²).

Results

We conducted a series of 2×2 (conflict \times feedback) repeated measures ANOVAs and a number of planned contrasts, which we describe for each dependent measure of interest.

Grand-averaged baseline pupil diameters for each condition are plotted in Figure 4. Neither the main effect of conflict [$F(1,23) = 3.1, p = .09$] nor the main effect of feedback [$F(1,23) = 0.7, p = .41$] was significant. The conflict \times feedback interaction also was not significant [$F(1,23) = 2.3, p = .14$], although it did show the expected trend. As a more sensitive test of whether the observed pattern of results was in line with our predictions, we carried out a planned contrast with the following weights for the four conditions: 1.0 for high-conflict/negative-feedback, -0.5 for low-conflict/negative-feedback, -0.5 for high-conflict/positive-feedback, and 0 for low-conflict/positive-feedback. This contrast yielded a significant result [$t(23) = 2.8, p = .005$], indicating that the ordering of baseline pupil values was consistent with our predictions. Subsequent pairwise comparisons indicated that high-conflict/negative-feedback differed from high-conflict/positive-feedback [$t(23) = 1.9, p < .05$] and low-conflict/negative-feedback [$t(23) = 2.6, p = .008$].

Grand-averaged pupil dilations evoked by the comparison tone are shown in Figure 4. Neither the two main effects nor the interaction effect was significant (all $F_s < 1$). As before, we then conducted a planned contrast among the four conditions, except now with the inversely signed weights, to reflect the hypothesized inverse relationship between baseline and task-evoked pupil diameter. Thus, we used the weights 1.0 for high-conflict/negative-feedback, -0.5 for low-conflict/negative-feedback, -0.5 for high-conflict/positive-feedback, and 0 for low-conflict/positive-feedback. The trend described by this contrast was significant [$t(23) = 2.0, p < .05$]. Subsequent pair-




Figure 4. Grand-averaged baseline pupil diameter (top panel) and pupil dilation magnitude (bottom panel) as a function of conflict (low vs. high) and feedback (negative vs. positive), the two independent variables in Experiment 2.

wise comparisons indicated that high-conflict/negative-feedback differed from low-conflict/negative-feedback [$t(23) = 2.0, p < .05$], but not from high-conflict/positive-feedback, although the observed trend was in the expected direction [$t(23) = 1.3, p = .098$].

A repeated measures ANOVA of the behavioral results revealed no reliable differences between the conditions in RT and accuracy. However, the main effect of feedback exhibited a trend toward longer RTs in the negative-feedback condition than in the positive-feedback condition [1,201 vs. 1,133 msec; $F(1,23) = 3.1, p = .09$]. Contrasts using the same weights as those for the pupil analysis also yielded null results.

Finally, we examined the reference pupil measurements recorded in dark conditions. For each participant, we compared this reference measurement against the maximal pupil magnitude observed following a comparison tone (under normal lighting conditions). The mean pupil diameter re-

corded in darkness was 7.20 mm, and the mean maximal pupil size attained after a dilation to a comparison tone was 6.34 mm. Two of the 12 participants achieved a pupil diameter larger than their dark reference measurement after a comparison tone at some point during the experiment. However, on average, the pupil in darkness was reliably larger than the size attained due to the experimental manipulation [$t(11) = 3.5, p < .005$]. These results, along with the results of Experiment 1B, mitigate the concern that the smaller dilations observed in the disengagement-promoting block were due to a physiological ceiling in pupil diameter.

Discussion

The pattern of pupillometric results obtained in Experiment 2 provides additional support for a relationship between the pattern of pupillary response and control state. Insofar as the pattern of pupillary responses conforms to the pattern of LC activity predicted by AGT, these findings provide further indirect evidence that pupil diameter may be a useful index of LC activity. Accordingly, the results are also consistent with the influence of task utility on control state predicted by AGT. Although we did not obtain the overarching statistical main effects of conflict or reward or the predicted conflict \times reward interaction, the ordering of both baseline and peak-dilation pupillary responses across the four blocks showed a systematic pattern in line with our predictions. Namely, the pupils exhibited smaller baseline diameters and larger dilations when the amount of both conflict and reward value were low and, likewise, when the amount of both conflict and reward value were high (nonsignificant trend for dilations). According to AGT, both conditions should promote task engagement (LC phasic mode), since both signal a need for recruitment of control (due to either high conflict via increased discrimination difficulty or low reward via negative feedback) in circumstances in which the required additional effort pays off (either in the form of positive feedback or through a reduction in conflict). Conversely, we observed larger baseline pupil diameters and smaller task-evoked dilations in the block with protracted conflict and low reward. As predicted by AGT, this block should have promoted task disengagement (LC tonic mode) and, hence, an adaptive reduction in the recruitment of control, since conflict remained high and feedback remained negative despite effortful performance.

We had expected to find more accurate performance and faster, tighter RT distributions in the engagement-promoting blocks than in the disengagement-promoting block. However, contrary to this prediction, we did not observe any reliable difference in behavioral performance across the various blocks. It is possible that our manipulations—although leading to the expected pupillometric results—were not powerful enough to engender large enough differences in control states to produce a reliable difference in behavior. Indeed, in general, the relative weakness of the obtained effects underscores some potential limitations of this experiment. Our operationalization of costs/conflict (i.e., the percentage of impossible-discrimination trials) and of reward (the va-

lence of feedback on these trials) was based largely on face-valid assumptions of what experimental conditions might induce the control state of interest. Although this was an improvement over Experiment 1A, in which control state was expected to vary spontaneously, individual subjective experience may have introduced a great deal of variability in the efficacy of these manipulations. More specifically, the demands on control were not adaptively titrated (i.e., akin to adjustment methods used in psychophysics). Furthermore, although the high-conflict/negative-feedback block produced pupil dynamics more consistent with disengagement than did the other blocks, the task did not actually provide for the overt expression of exploratory behaviors or true task disengagement. We remedied the latter limitation in Experiment 3.

EXPERIMENT 3

According to AGT, the utility of the LC tonic mode lies in its ability to allow disengagement from the current task context. In Experiment 3, we used a novel task, the diminishing-utility task, in which the task protocol gave the participants such an opportunity to disengage. The goal was to provide a more ecologically valid context in which to elicit and be able to measure behavioral responses explicitly indicative of task disengagement.

The participants performed a series of pitch discriminations of progressively increasing difficulty, with rewards for correct performance that increased in value with increasing task difficulty. Initially, the increases in reward value outpaced increases in difficulty (and associated increases in errors), so that the expected value of performance progressively increased and the participants remained engaged in the task. However, after several trials, the increases in difficulty led to sufficient numbers of errors so as to reduce the expected value of performance even in the face of the increasing reward value of correct responses. At the beginning of every trial, the participants were allowed to press an escape button, which would start a new series of discriminations, beginning again with low difficulty and low reward value. Thus, the participants were motivated to persevere with the current series of discriminations while the expected value increased. When the utility (expected value) of persevering with the current series began to diminish and, eventually, bottomed out, the participants could reset the task, trading off larger but less attainable rewards for smaller rewards attained more easily (i.e., yielding greater expected value).

We assumed that as task difficulty increased, the participants would experience increasing conflict, as evidenced by impaired performance. Initially, this higher conflict would promote greater task engagement (the LC phasic mode), resulting in the allocation of more control. However, as task difficulty became insurmountable, the protracted high conflict in the context of decelerating reward would favor a transition to task disengagement (the LC tonic mode) and, eventually, result in exercising the escape option. We predicted that baseline and transient pupil diameter would track the change from engagement to disengagement (relying, once again, on predictions

made by AGT about corresponding changes in LC activity) and would anticipate the escape event. Specifically, we predicted that baseline pupil diameter would increase over the trials leading up to an escape, peak on the escape trial, and rapidly decrease on the trials following the escape as achievable rewards increased. For pupil dilations, we predicted the opposite pattern. More precisely, we predicted that pupillary responses should track changes in the expected value of the task. As a result, escape behaviors should coincide both with pupillary indices of task disengagement (increased baseline diameter and diminished stimulus-evoked dilations) and with a minimum in expected value for the trial. We used a participant's history of performance at each of the levels of difficulty to calculate the probability of responding correctly at that level and used this probability in conjunction with the associated reward value for that difficulty level to generate an approximate measure of expected value.

Method

Participants. The participants were 31 Princeton University undergraduate students, 13 of them female, with ages ranging from 19 to 22 years.

Apparatus. The experimental setup was identical to that in Experiments 1 and 2.

Task and Procedure. The participants performed a pitch discrimination task similar to the one used in Experiment 2. However, in this experiment, the participants earned points for each trial judged correctly. If a participant responded correctly on a particular trial, the value of that trial was added to the participant's total score. In addition, in the next trial, the reward that could be earned increased by 5 points, and the pitch discrimination was made more difficult by halving the difference in pitch between the reference and comparison tones. Following an incorrect response, the reward value of the subsequent trial decreased by 10 points (but with a floor value of 0 points), and the level of task difficulty remained the same. Importantly, prior to each trial, the participants had the opportunity to "escape" from the current series of discriminations without penalty and receive a new discrimination task (i.e., comparison against a new reference tone), with the point value reset to 5 points and the easiest pitch discriminability. The participants were instructed to maximize their total score over the limited time of the experimental session, which was 30 min.

The task procedure and timing of stimulus events are illustrated in Figure 5. At the start of each trial, the participants were shown a score/value screen that displayed the total score accumulated thus far and the point value of the next trial (displayed within the bounding box for maintaining gaze; see Experiment 1A). The participants then indicated with a keypress whether they wanted to *accept* this trial or *escape*. If the participant elected to accept the trial, a reference/comparison tone pair followed after a delay of 5 sec: A 250-msec reference tone was played, followed by an interstimulus interval of 4 sec, followed by a 250-msec comparison tone. The participants were instructed to indicate as quickly and accurately as possible whether the comparison tone was higher or lower in pitch than the reference tone. After an additional delay of 3 sec, the accuracy of the participant's response was indicated by a 250-msec feedback sound: a bell sound for correct responses and a buzzer sound for incorrect responses. After an interval of 4 sec, the next trial started. If the participants pressed the *escape* button at the score/value screen, a 250-msec *escape sound* was played, immediately followed by a new score/value screen.

We refer to a series of trials accepted by a participant up to and including an escape trial as an *epoch* of play. Electing to escape began a new epoch. On the first trial of each epoch, the difference in pitch between the two tones was 64 Hz. As was noted above,

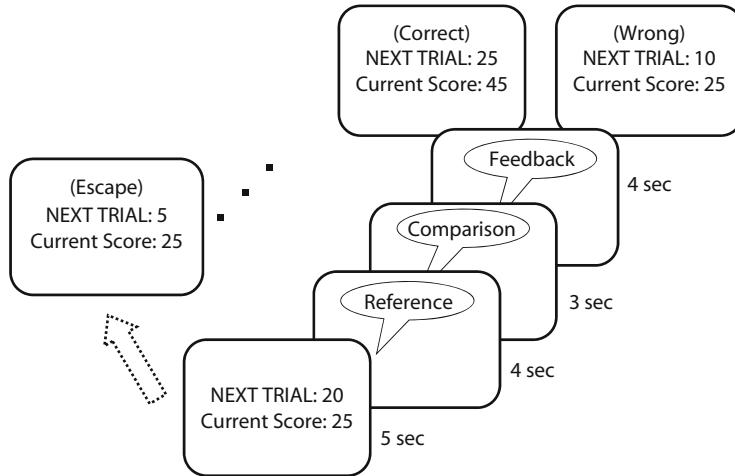


Figure 5. Illustration of a sample trial in the diminishing-utility task used in Experiment 3. All the stimuli were separated by delays of sufficient length to allow resolution of any pupillary response to the preceding event. Baseline pupil diameters were measured in the 1 sec immediately preceding the onset of the score/value screen. Dilations of the pupil were measured as peak deviations from this baseline in the 2.5 sec following onset of the comparison tone. See the text for further details.

this difference was halved following each correct response. If the participants correctly discriminated a 0.25-Hz difference, the tones presented on the next trial were impossible to discriminate (i.e., 0-Hz difference), and impossible-discrimination trials continued to be presented until the participant elected to escape. Accordingly, the participants would exhaust any real discriminable differences between the reference and comparison tones after 9 correct trials; the 10th and subsequent trials within an epoch were impossible to discriminate. The feedback signal on impossible-discrimination trials (correct or incorrect) was randomly picked. The same reference tone was presented on each trial within a given epoch. After an escape, a new reference tone was selected randomly without replacement from the set (400, 550, 700, and 850 Hz). The set was replenished if all the reference tones were exhausted. On 50% of the trials, the comparison tone was higher in pitch, and on the remaining trials, it was lower in pitch than the reference tone. All the presented sounds had an intensity of 72 dB.

Data analysis. Baseline pupil diameter was defined as the average diameter during the 1-sec interval prior to the onset of the score/value screen. Peak dilations of the pupil were measured as the maximal deviations from this baseline in the 2.5 sec following onset of the comparison tone. Because we were interested in dependent measures leading up to and following an escape, trials were averaged as a function of their position relative to these escape events. For each participant, we considered only escape events that were preceded and followed by a minimum of four regular (i.e., nonescape) trials. This constraint helped to ensure that epochs of sufficient task engagement were considered. In addition to behavioral and pupil measures, we calculated an estimate of expected value for each of the eight trials flanking an escape and the escape trial itself. For a given trial, expected value was computed individually for each participant by multiplying the point value of the trial (representing the potential reward value if the trial was accepted) by the expected accuracy on that trial for that participant. Expected accuracy was defined as the probability that the participant would give a correct response, given the level of difficulty of the required pitch discrimination. To determine this, we averaged the accuracy of all other trials for that participant with the same frequency difference between reference and comparison tones. Prior to the start of the experimental session, reference measures of the pupils were

taken from a subset of 18 participants, using the same procedure as that in Experiment 2.

Results

During the 30 min allotted to playing the game, the participants completed an average of 92 trials (range, 78–105) and elected to escape an average of 8.7 times (range, 2–22). The average epoch contained 11.4 trials. The average game score attained by the participants was 1,017 (range, 345–1,645). The average point value on single trials was 18.5, and the average frequency difference between the reference and comparison tones was 14.4 Hz. The highest point value attained on any given trial was 70 points.

We now consider our dependent measures as observed in the vicinity of an escape event, with trial-averaging locked to the escape event, as described in the Method section. Figure 6 (upper panel) shows escape-locked average RTs and accuracy for the trials flanking escapes. As was expected, performance was best on the trial immediately following the escape and then gradually decreased as task difficulty increased. Responses were slower [1,583 vs. 1,054 msec; $F(1,30) = 39.2, p < .001$] and less accurate [46.1% vs. 93.5%; $F(1,30) = 301.5, p < .001$] prior to an escape than afterward. Following the first postescape trial, RTs rapidly increased over subsequent trials [$F(1,90) = 47.3, p < .001$]. RTs continued to increase over the trials leading up to an escape [$F(1,90) = 9.1, p < .005$] and peaked on the trial immediately prior to the escape (1,875 msec). Accuracy values showed the opposite pattern, with a modest but reliable decrease during the first four trials of a new epoch [$F(1,90) = 47.8, p < .001$] and a sharp decrease during the trials leading up to the escape. As was expected, minimal accuracy was reached on the trial immediately prior to the escape (5.9%).

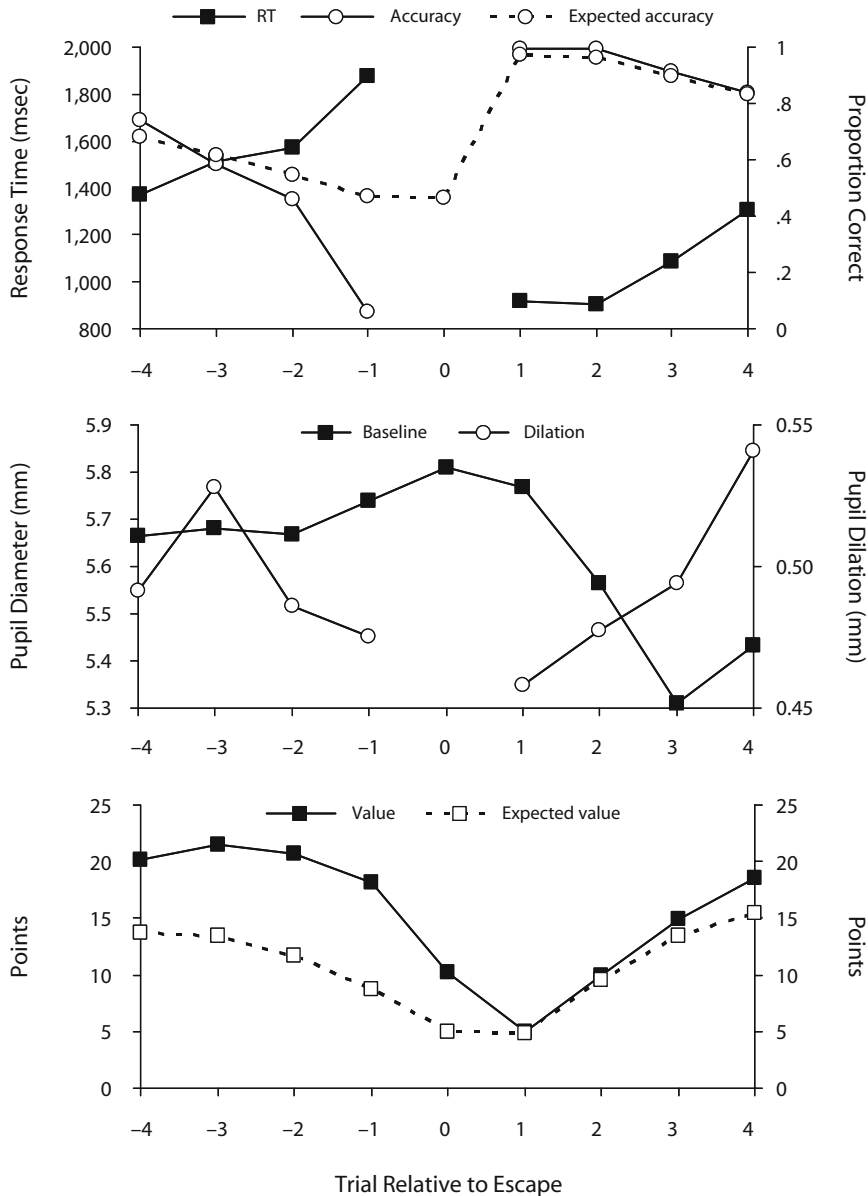


Figure 6. Grand-averaged dependent measures for peri-escape trials in Experiment 3. Trial number “0” indicates the escape trial. Top panel: Response time (RT), accuracy, and expected accuracy. Middle panel: Baseline pupil diameter and pupil dilation magnitude. Note that no measures of RT, accuracy, and reflexive pupil dilation are available for escape trials, because, on these trials, no comparison tone was presented. Bottom panel: Trial value and its computed expected value. Expected value most closely tracks the pattern of the pupil data.

Figure 6 (middle panel) shows pupil measures, averaged across participants, for escape trials and the four flanking trials preceding and following an escape. Average baseline pupil diameter showed a modest but significant increase over the four consecutive trials leading to an escape, as revealed by a repeated measures linear trend analysis [$F(1,120) = 4.7, p < .05$]. Baseline pupil diameter peaked on the escape trial and then showed a sharp decrease over the next few trials, as participants started a new epoch [$F(1,120) = 50.0, p < .001$]. The minimum baseline pupil diameter was observed on the third postescape

trial. The overall inverted-U-shaped trend, centered on the escape trial, was significant, as indicated by a quadratic trend analysis [$F(1,30) = 32.7, p < .001$]. Pupil dilations evoked by the comparison tone did not follow the predicted decreasing trend prior to an escape ($F < 1$). However, the predicted increasing trend for trials after an escape was highly significant [$F(1,90) = 12.3, p < .001$].

We next examined how our metric of expected value varied across the trials surrounding an escape. Figure 6 (lower panel) shows that, on average, the participants elected to escape when expected value was at a minimum.

Expected value showed a significant quadratic trend centered on the escape that mirrored the pattern observed in baseline pupil diameter [$F(1,240) = 522.3, p < .0001$]. Indeed, a series of regression analyses confirmed that our expected value metric was a much better predictor of pupil diameter (both baseline and dilations) than were the other performance metrics, including the subcomponents of expected value (actual point value and expected accuracy; see Figure 6) and actual accuracy achieved on a given trial. Regressing expected value onto baseline pupil diameter produced a linear fit of $r^2 = .48$ ($p < .05$), as compared with unreliable fits for actual point value ($r^2 = .04, p = .60$), expected accuracy ($r^2 = .30, p = .13$), and actual accuracy ($r^2 = .32, p = .12$). Regressing expected value onto pupil dilation yielded a fit of $r^2 = .73$ ($p < .01$), as compared with unreliable fits for actual point value ($r^2 = .04, p = .60$), expected accuracy ($r^2 = .02, p = .75$), and actual accuracy ($r^2 = .0001, p = .98$).

Finally, we examined the reference pupil measurements recorded in the dark conditions. As in Experiment 2, we compared this reference measurement against the maximal pupil magnitude observed following a comparison tone. The mean pupil diameter recorded in darkness was 7.79 mm, and the mean maximal pupil size attained after a dilation to a comparison tone was 6.57 mm. Two of the 18 participants achieved a pupil diameter larger than their dark reference measurement after a comparison tone at some point during the experiment. However, on average, the pupil in darkness was reliably larger than the size attained due to the experimental manipulation [$t(17) = 4.4, p < .001$].

Discussion

In Experiment 3, we administered a task that was specifically designed to modulate task engagement over the course of an experimental session. Most important, the task provided an overt behavioral means (the escape) by which the participants could disengage from the current series of tone discriminations and opt for a new series with potentially higher utility. The participants behaved optimally on average, choosing to escape when the expected value of the discriminations began to decline. The data observed prior to the escape (as task engagement presumably waned) and after the escape (as engagement built again) provide additional evidence that the pattern of pupillary responses shows a consistent relationship to control state. And, as in Experiments 1 and 2, this relationship parallels the relationship between patterns of LC activity and control state predicted by AGT. In particular, we observed reliably increasing baseline pupil diameters (with a peak on the escape trial) and a trend toward diminished transient dilations (paralleling the LC tonic mode) as expected value began to decline and the participants approached a decision to escape. Similar pupil dynamics have been reported under different conditions of processing overload. Granholm, Asarnow, Sarkin, and Dykes (1996) found that pupil diameter increased systematically with increasing memory load (in a digit span recall task) until the limit of available memory capacity (7 ± 2 dig-

its), when it reached asymptote, and then declined with memory overload (>9 digits). Pooch (1973) found that pupil diameter systematically increased with increasing stimulus presentation rate in a forced choice RT task until the presentation rate exceeded an individual's maximum processing rate, when pupil diameter started to decline again. Unfortunately, the short interstimulus intervals in these two studies did not allow a distinction to be made between baseline pupil diameter and task-evoked dilations.

After an escape, we observed reliably decreasing baseline pupil diameters and increasing transient dilations (paralleling the LC phasic mode). These results also replicated the hallmark finding in the classic work of Kahneman, Beatty, and others: Increases in task difficulty and corresponding mental effort result in increased task-evoked pupil dilations (Beatty, 1982b; Kahneman, 1973).

AGT proposes that control state is driven by online assessments of utility, which are used to determine whether task engagement or disengagement is most adaptive. To test for this, and whether it was reflected in pupillary responses, we used a measure of expected value to gauge whether pupil dynamics tracked the trial-to-trial changes in utility experienced by the participants. However, the metric we used had some flaws. To determine expected value, we had to determine a measure of probability of correct performance on the current trial. For this purpose, we used participants' performance on all trials of equal difficulty, which assumed unrealistically that they had information about trials in the future as well as the past. Furthermore, we assumed that judgments involving the same absolute difference in frequency between reference and comparison tones were all of equal difficulty, irrespective of the absolute frequencies of the tones. This clearly did not take into account the proportional nature of perceived differences (i.e., Weber's law). However, because all our reference tones were of similar magnitude, the proportion of comparison tone to reference tone was not dramatically different for the same absolute difference in frequency for another comparison/reference tone pair (except for extreme values). Correcting these two flaws would have severely underpowered our ability to compute expected value for a given trial. Nevertheless, even with these shortcomings, our measure of expected value proved to be reliable in predicting pupil baseline and pupil dilation. Indeed, as a predictor of pupillary dynamics, it proved superior to any other measure of performance, suggesting some validity of the measure. This result provides suggestive evidence that pupillary responses (and, by inference, LC activity) track utility, although the specific algorithm by which this is done remains to be determined.

GENERAL DISCUSSION

We have reported three experiments that investigated the relationship between pupillary responses and behaviors reflective of two different control states: task engagement (exploitation) and task disengagement (exploration). The overarching goal of this work was to understand the neural mechanisms that govern shifts in control state

from exploitation to exploration. Our experiments were inspired by several recent developments: the identification of a relationship between LC mode of activity and behaviors reflective of different control states; the formulation of AGT, which relates changes in LC activity and control state to changes in task utility; and observations in the monkey that there is a close correlation between LC activity and pupil diameter. In the present experiments, we exploited these developments to make specific predictions about the relationship among changes in task utility, changes in control state, and corresponding patterns of pupillary responses.

Although the effect sizes obtained were generally modest, the results were consistently in line with our predictions. In Experiment 1A, we used a target detection task that preserved close homology with the original work with monkeys (Aston-Jones et al., 1994; Rajkowski et al., 1993). This experiment demonstrated the predicted inverse relationship between baseline and task-evoked pupil diameter and the relationship between pupil dynamics and behavior predicted by AGT if pupil diameter is indicative of LC activity. In Experiments 2 and 3, we manipulated conflict and reward in order to induce theoretically predicted changes in control state and showed that these manipulations produced the corresponding pattern of predicted pupil dynamics. Furthermore, in Experiment 3, we found that trial-to-trial changes in task utility (measured as expected value) were reliably correlated with baseline and task-evoked changes in pupil diameter and accurately predicted overt indications of task engagement and disengagement.

At the least, these findings provide the first evidence for a close relationship between pattern of pupillary response and control state, suggesting that pupillometry may provide a useful index of control state in future research. We suggest that two additional inferences can be made from our findings: (1) Pupillary responses track LC activity in humans in a manner similar to that observed in monkeys, and (2) on the basis of this, changes in task utility are associated with changes in LC activity and control state predicted by AGT. We recognize that our data do not provide direct support for these conclusions. However, given the specificity, detail, and nonintuitive character of the predictions, it seems highly unlikely that they reflect the effects of some less specific set of processes or mechanisms. Rather, we suggest that the confirmation of these predictions provided by our data lends support to the validity both of the method (i.e., pupillometry can be used to index LC-mediated control state) and of the theory that gave rise to the predictions (i.e., AGT; Figure 7).

An important theoretical issue that is not directly addressed by the present work is the question of which brain structure(s) mediate the relationship between LC activity and pupil diameter, given that there is no known direct connection between the LC and the mechanisms that regulate pupil size.³ Earlier, we suggested that the relationship between the two variables may reflect the downstream influences of a common afferent source—perhaps the PGI nucleus of the ventral medulla, which is both a critical

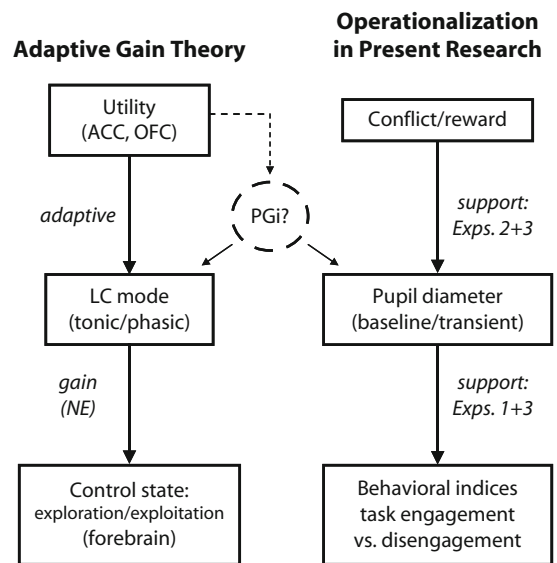


Figure 7. Schematic outline of adaptive gain theory (AGT) and its relationship to the present research. ACC, anterior cingulate cortex; OFC, orbitofrontal cortex; LC, locus coeruleus; NE, nor-epinephrine; PGI, paraventricular nucleus of the ventral medulla. One central assumption of AGT is that the LC firing mode is adaptively driven by online assessments of task utility. This part of the theory received empirical support in Experiments 2 and 3, which showed the predicted effect of experimental manipulations of conflict (costs) and feedback valence (reward) on pupillary dynamics. The second core assumption of the theory is that the LC uses this information about utility to make compensatory adjustments in control state. This part of the theory received support in Experiments 1 and 3, which showed the predicted relationship between pupillary dynamics and behavioral indices associated with the two LC modes.

relay for sympathoexcitatory signals to the autonomic nervous system and one of the major afferents to the LC (Aston-Jones et al., 1986; Aston-Jones et al., 1991). An interesting corollary of this hypothesis is that utility signals originating in frontal structures should influence the LC and the pupil by way of this common source. This possibility is consistent with anatomical studies that have revealed widespread afferents to the PGI from numerous brain areas, including the medial prefrontal cortex, insula, hypothalamus, and periaqueductal gray (Van Bockstaele, Pieribone, & Aston-Jones, 1989). Furthermore, fMRI studies in humans and single-cell stimulation/recording studies in animals have shown that activity in this afferent network (including the prefrontal cortex) is related to changes in pupil diameter (Critchley, Tang, Glaser, Butterworth, & Dolan, 2005; Loewenfeld, 1993).

Although the exact neuroanatomical underpinnings remain to be determined, the hypothesis that pupillary responses track instantaneous LC firing rate, reflecting both its tonic and phasic activity profiles, is consistent with a wide array of studies in cognitive psychology. This includes ones that have reported task-evoked phasic dilations in pupil diameter following virtually any effortful cognitive act, including memory recall, problem solving, language processing, perceptual tasks, and attentional tasks

(for reviews, see Beatty, 1982b; Steinhauer & Hakerem, 1992). Typically, the size and duration of these dilations are positively correlated with task difficulty and, thus, may be regarded as a reliable index of the relative effort exerted on a given trial (Kahneman, 1973). For example, Kahneman and Beatty (1967) showed that the magnitude of task-evoked pupil dilations varied parametrically with the difficulty of a pitch discrimination task, with the largest dilations occurring for impossible discriminations. It is important to note that these studies in humans have focused mostly on the transient, phasic changes in pupil size associated with cognitive acts. In contrast, the findings of Rajkowski et al. (1993) with monkeys described above pertain to changes in baseline pupil size, recorded just prior to stimulus presentation; stimulus-related changes were not recorded. Taken together, these previous human and animal studies have shown that task-related, effortful processing is associated with tonic constrictions (in the monkey) and phasic dilations (in the human) of the pupil. An important contribution of the work presented here is that it replicates and extends these results within the same subjects. Furthermore, our work suggests that if task difficulty is further increased until task utility falls below an acceptable value, the pattern of pupil dynamics reverses: tonic dilation and phasic constriction, mirroring the dynamics of the tonic LC mode.

Arousal Theory

Both the pupil (Kahneman, 1973) and the LC (Foote, Aston-Jones, & Bloom, 1980) have traditionally been cast as mechanisms related to arousal. Because this article proposes a close relationship between the two and draws on AGT to relate them to behavior, the question arises as to how AGT relates to arousal theory. This is a complicated question to address, since the literature on arousal is so vast and heterogeneous and is still contentious (see Staal, 2004). It is generally accepted that there is not a single, unitary construct of arousal but, rather, a constellation of brain and somatic systems that subservise distinct but often overlapping functions (Neiss, 1988; Pribram & McGuinness, 1975; Robbins, 1997). The LC–NE system is one of these systems and exhibits a relationship with performance that reflects one of the most important components of arousal theory: the Yerkes–Dodson law, which asserts that the function relating performance to arousal has an inverted-U shape (Duffy, 1957; Yerkes & Dodson, 1908). However, given the multifaceted nature of arousal, many have questioned the validity of such a simple unidimensional construct, and there is substantial disagreement as to whether the literature justifies the Yerkes–Dodson law's lawful status (Landers, 1980; Neiss, 1988). Nonetheless, the Yerkes–Dodson law captures the key features of the relationship between LC activity and performance: Optimal performance is observed in the middle of the LC tonic firing rate range (i.e., the LC phasic mode), with increasing drowsiness and inattention below this (low or absent tonic firing rate; Hobson, McCarley, & Wyzinski, 1975) and poor, distractible performance above it (the LC tonic mode; Aston-Jones et al., 1997; Aston-Jones et al., 1994).

To the extent that the LC–NE system mediates a Yerkes–Dodson relationship between arousal and performance, its properties must be reconciled with Easterbrook's (1959) cue utilization hypothesis, the most enduring psychological account of the Yerkes–Dodson law. Easterbrook argued that arousal impacts performance by restricting attentional selection. Thus, as arousal increases, the scope of selection narrows, and task-relevant information is more efficiently selected and processed. However, if arousal increases too much, selection becomes so restricted that some task-relevant cues fail to be processed, resulting in impaired performance. This account proved to have robust support in the subsequent literature (e.g., Broadbent, 1979; Hockey, 1978).

The cue utilization hypothesis as Easterbrook (1959) cast it dovetails well with some aspects of AGT. Inasmuch as the LC phasic mode filters the processing of irrelevant information and selectively facilitates task-relevant information, AGT accords with Easterbrook's hypothesis quite well. However, on this account, the LC tonic mode increases the likelihood of selecting and processing task-irrelevant sources of information, in the service of exploration. On the surface, this does not seem compatible with Easterbrook's restriction premise. However, the cue utilization hypothesis has some inherent inconsistencies. First, it suggests that poor performance at low arousal is due to excessive processing of information when drowsy or inattentive, which seems highly unlikely. Furthermore, it suggests that at very high arousal, performance is poor because of a rigidity in focus, another claim that is not valid on its face, given observable distractibility in highly aroused subjects (and associated with the LC tonic mode observed in monkeys). Kahneman (1973) noted this paradox of simultaneous attentional rigidity and lability at high arousal and suggested that high arousal may impart a system with local rigidity but high global distractibility, such that narrow selective attention is allocated in a labile fashion. AGT may provide a mechanistic grounding for this insight.

AGT associates high arousal (the LC tonic mode) with indiscriminate (vs. task-selective) increases in processing gain throughout the system. This indiscriminate increase in gain increases the chances that task-irrelevant information will be processed and compete with task-relevant information for control of behavior and, thus, might have effects similar to the labile allocation of attention proposed by Kahneman (1973). Moreover, a number of authors (e.g., Hanoch & Vitouch, 2004; Zajonc, 1980) have argued that, from an ecological perspective, the adaptiveness of the high-arousal state has been neglected. These authors note that under high-arousal conditions, thorough information processing is sacrificed for quick action, self-preservation behavior, and efficient selection of alternate behaviors. This argument is consistent with an account of the adaptiveness of the LC tonic mode in terms of disengagement and exploration, as proposed by AGT. In sum, to the extent that the LC is an important determinant of arousal, AGT offers a mechanistic account of the utility of the right-hand side of the Yerkes–Dodson curve, as well as the far more heavily studied center.

Future Directions

An important claim of AGT is that the LC tonic mode encourages task disengagement and exploration through tonically (and indiscriminately) increased levels of gain, resulting in destabilization of task-related representations and lability of behavior. In Experiment 3, we found empirical support for this claim by showing that disengagement in the diminishing-utility task was associated with pupil dynamics that we predicted should be indicative of LC tonic activity. However, we note that this experiment provided only an opportunity for participants to disengage, but no opportunity to explore per se. We predict that the LC tonic mode should be associated with exploration in a more natural environment in which many behavioral options are available. Thus, an important extension of this work would be to investigate a more ecologically valid paradigm, in which participants are encouraged to explore multiple task options or strategies before choosing to engage in any one. We expect that the LC tonic mode will be evident throughout such exploratory phases, and not just at punctate moments of disengagement. A promising candidate would be the *n*-armed bandit paradigm, in which agents must decide which of a number of probabilistically rewarding options (each with its own reward probability structure) to exploit and for how long (see Sutton & Barto, 1998). The neural and computational basis for behavior in such paradigms has already begun to be explored by others (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006), and the findings seem largely consistent with AGT (Cohen et al., 2007).

As was noted above, the LC is thought to function largely in parallel with the sympathoexcitatory pathways involved in somatic arousal. Indeed, the LC-NE system has been likened to a "central sympathetic nervous system" (Aston-Jones, 1985; Nieuwenhuis, De Geus, & Aston-Jones, in press; Van Bockstaele & Aston-Jones, 1995): It is driven by the same hypothalamic pathways that drive the sympathetic nervous system, but its own efferents do not have peripheral sympathetic targets. As a consequence, it is likely that our pupillometric measures are indices of an efference copy of an afferent common to both LC and sympathetic targets. Future studies that cross-validate our pupil findings with other sensitive measures of psychophysiological arousal, such as the galvanic skin response, may prove instructive with regard to this issue.

Another direction for future research whose aim is to clarify the pupil-LC relationship in humans entails simultaneous recordings of pupil diameter and scalp electrophysiological potentials—in particular, the P3(00). We have recently proposed that the P3 reflects the electrophysiological consequences of phasic activity of the LC-NE system, a hypothesis that is supported by a wealth of data from electrophysiology, lesion studies, psychopharmacology, and other methods (Nieuwenhuis, Aston-Jones, & Cohen, 2005). Previous studies have shown that pupil diameter and P3 amplitude are sensitive to similar experimental variables (Steinhauer & Hakerem, 1992). One objective for future research would be to investigate whether stimulus-evoked P3 amplitude is dependent on

prestimulus pupil diameter in a way that reflects the inverse relationship between phasic and tonic LC activity.

Of course, the strongest prediction following from the present work is that simultaneous recordings of LC activity and pupil size should bear out the hypothesized relationship. One approach to acquiring such data is direct neuronal recordings in nonhuman species. As has been discussed, preliminary data from a visual target detection experiment in monkeys have already confirmed this prediction for tonic LC activity and baseline pupil diameter (Rajkowski et al., 1993). In follow-up work, an auditory or haptic variant of the target detection task should be used, so that the prediction can also be tested for phasic LC activity and stimulus-evoked pupil dilations. Another approach is to exploit recently developed fMRI methods for measuring BOLD responses from brainstem neuro-modulatory nuclei in humans. Recent success conducting such measurements from the ventral tegmental area (D'Ardenne, McClure, Nystrom, & Cohen, 2008) offers the promise that they may be usefully applied to the measurement of LC activity. Finally, establishing the pupil as a convenient, noninvasive marker of LC activity may allow it to be used as a diagnostic tool in the assessment of pathologies thought to involve the LC-NE system, such as attention deficit disorder (Solanto, 1998).

AUTHOR NOTE

This research was supported by the Netherlands Organization for Scientific Research (grant to S.N.), a National Institute of Mental Health Sylvio O. Conte Center Grant for Neuroscience Research (Grant MH062196 to M.S.G. and J.D.C.), and the Center for the Study of Brain, Mind, and Behavior at Princeton University. We thank Gary Aston-Jones, Leigh Nystrom, and Daniel Kahneman for valuable discussions related to this work, and Henk van Steenberg and Rinus Verdonshot for technical advice. Correspondence concerning this article should be addressed to S. Nieuwenhuis, Cognitive Psychology Unit, Institute for Psychological Research, Leiden University, Wassenaarseweg 52, 2333 AK Leiden, The Netherlands (e-mail: snieuwenhuis@fsw.leidenuniv.nl).

REFERENCES

- ASTON-JONES, G. (1985). Behavioral functions of locus coeruleus derived from cellular attributes. *Physiological Psychology*, *13*, 118-126.
- ASTON-JONES, G., & COHEN, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403-450.
- ASTON-JONES, G., ENNIS, M., PIERIBONE, V. A., NICKELL, W. T., & SHIPLEY, M. T. (1986). The brain nucleus locus coeruleus: Restricted afferent control of a broad efferent network. *Science*, *234*, 734-737.
- ASTON-JONES, G., FOOTE, S. L., & BLOOM, F. E. (1984). Anatomy and physiology of locus coeruleus neurons: Functional implications. In M. Ziegler & C. R. Lake (Eds.), *Norepinephrine: Frontiers of clinical neuroscience* (Vol. 2, pp. 92-116). Baltimore, MD: Williams & Wilkins.
- ASTON-JONES, G., RAJKOWSKI, J., & KUBIAK, P. (1997). Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience*, *80*, 697-715.
- ASTON-JONES, G., RAJKOWSKI, J., KUBIAK, P., & ALEXINSKY, T. (1994). Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *Journal of Neuroscience*, *14*, 4467-4480.
- ASTON-JONES, G., RAJKOWSKI, J., LU, W., ZHU, Y., COHEN, J. D., & MORECRAFT, R. J. (2002). Prominent projections from the orbital prefrontal cortex to the locus coeruleus in monkey. *Society for Neuroscience Abstracts*, *28*, 86-89.

- ASTON-JONES, G., SHIPLEY, M. T., CHOUVET, G., ENNIS, M., VAN BOCKSTAELE, E., PIERIBONE, V., ET AL. (1991). Afferent regulation of locus coeruleus neurons: Anatomy, physiology and pharmacology. In C. D. Barnes & O. Pompeiano (Eds.), *Progress in brain research* (Vol. 88, pp. 47-75). New York: Elsevier.
- BEATTY, J. (1982a). Phasic not tonic pupillary responses vary with auditory vigilance performance. *Psychophysiology*, **19**, 167-172.
- BEATTY, J. (1982b). Task-evoked pupillary responses, processing load, and the structure of processing resources. *Psychological Bulletin*, **91**, 276-292.
- BOTVINICK, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cognitive, Affective, & Behavioral Neuroscience*, **7**, 356-366.
- BOURET, S., & SARA, S. J. (2005). Network reset: A simplified overarching theory of locus coeruleus noradrenergic function. *Trends in Neurosciences*, **28**, 574-582.
- BRADSHAW, J. L. (1969). Background light intensity and the pupillary response in a reaction time task. *Psychonomic Science*, **14**, 271-272.
- BROADBENT, D. E. (1979). Human performance and noise. In C. M. Hartis (Ed.), *Handbook of noise control* (pp. 17.1-17.20). New York: McGraw-Hill.
- CLAYTON, E. C., RAJKOWSKI, J., COHEN, J. D., & ASTON-JONES, G. (2004). Phasic activation of monkey locus coeruleus neurons by simple decisions in a forced choice task. *Journal of Neuroscience*, **24**, 9914-9920.
- COHEN, J. D., MCCLURE, S. M., & YU, A. J. (2007). Should I stay or should I go? Exploration versus exploitation. *Philosophical Transactions of the Royal Society B*, **362**, 933-942.
- CRITCHLEY, H. D., TANG, J., GLASER, D., BUTTERWORTH, B., & DOLAN, R. J. (2005). Anterior cingulate activity during error and autonomic response. *NeuroImage*, **27**, 885-895.
- D'ARDENNE, K., MCCLURE, S. M., NYSTROM, L. E., & COHEN, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, **319**, 1264-1267.
- DAW, N. D., O'DOHERTY, J. P., DAYAN, P., SEYMOUR, B., & DOLAN, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, **441**, 876-879.
- DUFFY, E. (1957). The psychological significance of the concept of "arousal" or "activation." *Psychological Review*, **64**, 265-275.
- EASTERBROOK, J. A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological Review*, **66**, 187-201.
- EINHÄUSER, W., STOUT, J., KOCH, C., & CARTER, O. (2008). Pupil dilation reflects perceptual selection and predicts subsequent stability in perceptual rivalry? *Proceedings of the National Academy of Sciences*, **105**, 1704-1709.
- ELAM, M., SVENSSON, T. H., & THOREN, P. (1986). Locus coeruleus neurons and sympathetic nerves: Activation by cutaneous sensory afferents. *Brain Research*, **366**, 254-261.
- FOOTE, S. L., ASTON-JONES, G., & BLOOM, F. E. (1980). Impulse activity of locus coeruleus neurons in awake rats and monkeys is a function of sensory stimulation and arousal. *Proceedings of the National Academy of Sciences*, **77**, 3033-3037.
- GILZENRAT, M. S., HOLMES, B. D., RAJKOWSKI, J., ASTON-JONES, G., & COHEN, J. D. (2002). Simplified dynamics in a model of noradrenergic modulation of cognitive performance. *Neural Networks*, **15**, 647-663.
- GRANHOLM, E., ASARNOW, R. F., SARKIN, A. J., & DYKES, K. L. (1996). Pupillary responses index cognitive resource limitations. *Psychophysiology*, **33**, 457-461.
- HANOCH, Y., & VITOUCH, O. (2004). When less is more: Information, emotional arousal and the ecological reframing of the Yerkes-Dodson law. *Theory & Psychology*, **14**, 427-452.
- HILTON, S. M., & SMITH, P. R. (1984). Ventral medullary neurones excited from the hypothalamic and mid-brain defence areas. *Journal of the Autonomic Nervous System*, **11**, 35-42.
- HOBSON, J. A., MCCARLEY, R. W., & WYZINSKI, P. W. (1975). Sleep cycle oscillation: Reciprocal discharge by two brainstem neuronal groups. *Science*, **189**, 55-58.
- HOCKEY, G. R. J. (1978). Effects of noise on human work efficiency. In D. N. May (Ed.), *Handbook of noise assessment* (pp. 335-372). New York: Van Nostrand-Reinhold.
- HOU, R. H., FREEMAN, C., LANGLEY, R. W., SZABADI, E., & BRADSHAW, C. M. (2005). Does modafinil activate the locus coeruleus in man? Comparison of modafinil and clonidine on arousal and autonomic functions in human volunteers. *Psychopharmacology*, **181**, 537-549.
- JANISSE, M. P. (1977). *Pupillometry: The psychology of the pupillary response*. Washington, DC: Hemisphere.
- JOUVET, M. (1969). Biogenic amines and the states of sleep. *Science*, **163**, 32-41.
- KAHNEMAN, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.
- KAHNEMAN, D., & BEATTY, J. (1967). Pupillary responses in a pitch-discrimination task. *Perception & Psychophysics*, **2**, 101-105.
- LACEY, J. I. (1956). The evaluation of autonomic responses: Toward a general solution. *Annals of the New York Academy of Sciences*, **67**, 125-163.
- LANDERS, D. M. (1980). The arousal-performance relationship revisited. *Research Quarterly for Exercise & Sport*, **51**, 77-90.
- LEVITT, P., & MOORE, R. Y. (1979). Origin and organization of brainstem catecholamine innervation in the rat. *Journal of Comparative Neurology*, **186**, 505-528.
- LOEWENFELD, I. (1993). *The pupil: Anatomy, physiology, and clinical applications*. Detroit: Wayne State University Press.
- LOEWY, A. D., WALLACH, J. H., & MCKELLAR, S. (1981). Efferent connections of the ventral medulla oblongata in the rat. *Brain Research*, **228**, 63-80.
- MCCLURE, S. M., GILZENRAT, M. S., & COHEN, J. D. (2005). An exploration-exploitation model based on norepinephrine and dopamine activity. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems 18* (pp. 867-874). Cambridge, MA: MIT Press.
- MORAD, Y., LEMBERG, H., YOFE, N., & DAGAN, Y. (2000). Pupilligraphy as an objective indicator of fatigue. *Current Eye Research*, **21**, 535-542.
- MOXON, K. A., DEVILBISS, D. M., CHAPIN, J. K., & WATERHOUSE, B. D. (2007). Influence of norepinephrine on somatosensory neuronal responses in the rat thalamus: A combined modeling and in vivo multichannel, multi-neuron recording study. *Brain Research*, **1147**, 105-123.
- NEISS, R. (1988). Reconceptualizing arousal: Psychological states in motor performance. *Psychological Bulletin*, **103**, 345-366.
- NIEUWENHUIS, S., ASTON-JONES, G., & COHEN, J. D. (2005). Decision making, the P3, and the locus coeruleus-norepinephrine system. *Psychological Bulletin*, **131**, 510-532.
- NIEUWENHUIS, S., DE GEUS, E. J., & ASTON-JONES, G. (in press). The anatomical and functional relationship between the P3 and autonomic components of the orienting response. *Psychophysiology*.
- PHILLIPS, M. A., SZABADI, E., & BRADSHAW, C. M. (2000). Comparison of the effects of clonidine and yohimbine on spontaneous pupillary fluctuations in healthy human volunteers. *Psychopharmacology*, **150**, 85-89.
- POOCK, G. K. (1973). Information processing vs pupil diameter. *Perceptual & Motor Skills*, **37**, 1000-1002.
- PRIBRAM, K. H., & MCGUINNESS, D. (1975). Arousal, activation, and effort in the control of attention. *Psychological Review*, **82**, 116-149.
- RAJKOWSKI, J., KUBIAK, P., & ASTON-JONES, G. (1993). Correlations between locus coeruleus (LC) neural activity, pupil diameter and behavior in monkey support a role of LC in attention. *Society for Neuroscience Abstracts*, **19**, 974.
- RAJKOWSKI, J., LU, W., ZHU, Y., COHEN, J. D., & ASTON-JONES, G. (2000). Prominent projections from the anterior cingulate cortex to the locus coeruleus (LC) in rhesus monkey. *Society for Neuroscience Abstracts*, **26**, 2230.
- RATCLIFF, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, **86**, 446-461.
- RICHER, F., & BEATTY, J. (1987). Contrasting effects of response uncertainty on the task-evoked pupillary response and reaction time. *Psychophysiology*, **24**, 258-262.
- RIDDERINKHOF, K. R., ULLSPERGER, M., CRONE, E. A., & NIEUWENHUIS, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, **306**, 443-447.
- ROBBINS, T. W. (1997). Arousal systems and attentional processes. *Biological Psychology*, **45**, 57-71.
- SERVAN-SCHREIBER, D., PRINTZ, H., & COHEN, J. D. (1990). A network

- model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, **249**, 892-895.
- SHEA-BROWN, E., GILZENRAT, M., & COHEN, J. D. (2008). Optimization of decision making in multilayer networks: The role of locus coeruleus. *Neural Computation*, **20**, 2863-2894.
- SIEGLE, G. J. (1999). *Cognitive and physiological aspects of attention to personally relevant negative information in depression*. Unpublished doctoral dissertation, University of California, San Diego.
- SOLANTO, M. V. (1998). Neuropsychopharmacological mechanisms of stimulant drug action in attention-deficit hyperactivity disorder: A review and integration. *Behavioural Brain Research*, **94**, 127-152.
- STAAL, M. A. (2004). *Stress, cognition, and human performance: A literature review and conceptual framework* (NASA Tech. Memorandum 212824). Moffett Field, CA: NASA Ames Research Center.
- STEINHAEUER, S. R., & HAKEREM, G. (1992). The pupillary response in cognitive psychophysiology and schizophrenia. In D. Friedman & G. Bruder (Eds.), *Psychophysiology and experimental psychopathology: A tribute to Samuel Sutton* (Annals of the New York Academy of Sciences, Vol. 658, pp. 182-204). New York: New York Academy of Sciences.
- STEINHAEUER, S. R., & ZUBIN, J. (1982). Vulnerability to schizophrenia: Information processing in the pupil and event-related potential. In E. Usdin & I. Hanin (Eds.), *Biological markers in psychiatry and neurology* (pp. 371-385). Oxford: Pergamon.
- SUTTON, R. S., & BARTO, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- USHER, M., COHEN, J. D., SERVAN-SCHREIBER, D., RAJKOWSKI, J., & ASTON-JONES, G. (1999). The role of locus coeruleus in the regulation of cognitive performance. *Science*, **283**, 549-554.
- VAN BOCKSTAELE, E. J., & ASTON-JONES, G. (1995). Integration in the ventral medulla and coordination of sympathetic, pain and arousal functions. *Clinical & Experimental Hypertension*, **17**, 153-165.
- VAN BOCKSTAELE, E. J., PIERIBONE, V. A., & ASTON-JONES, G. (1989). Diverse afferents converge on the nucleus paragigantocellularis in the rat ventrolateral medulla: Retrograde and anterograde tracing studies. *Journal of Comparative Neurology*, **290**, 561-584.
- VAN OLST, E. H., HEEMSTRA, M. L., & TEN KORTENAAR, T. (1979). Stimulus significance and the orienting reaction. In H. D. Kimmel, E. H. Olst, & J. F. van Orlebeke (Eds.), *The orienting reflex in humans* (pp. 521-547). Hillsdale, NJ: Erlbaum.
- VAN ZANDT, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review*, **7**, 424-465.
- VINCENT, S. B. (1912). The function of the vibrissae in the behavior of the white rat. *Behavioral Monographs*, **1**, 7-81.
- YERKES, R. M., & DODSON, J. D. (1908). The relation of strength of stimulus to rapidity of habit formation. *Journal of Comparative Neurology & Psychology*, **18**, 459-482.
- YU, A. J., & DAYAN, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, **46**, 681-692.
- ZAJONC, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, **2**, 151-175.
- ZHU, Y., IBA, M., RAJKOWSKI, J., & ASTON-JONES, G. (2004). Projection from the orbitofrontal cortex to the locus coeruleus in monkeys revealed by anterograde tracing. *Society for Neuroscience Abstracts*, **30**, 211-213.

NOTES

1. It is important to note that, whereas experimental studies have focused on these as two distinct modes of LC function, it is possible that they represent the ends of a continuum. Modeling work suggests that, even while lying along a continuum, gradual changes in the relevant parameters are punctuated by relatively discrete transitions in function. For this reason and for clarity of presentation, we will continue to refer to these as distinct modes of function. However, it remains an empirical question whether transitions between these modes are continuous or more discrete.

2. Pilot testing of this task revealed that feedback itself evoked a large pupillary dilation. To maximize the number of trials obtained in a session, we eliminated feedback in Experiment 1A to avoid the extra time delay necessary to allow the pupil to return to baseline. With reliable results in this task, we reintroduced feedback in Experiments 2 and 3, trading off trial number for more sophisticated manipulations.

3. Authors occasionally assume a direct influence of the LC on the pupil, on the basis of the following premises: (1) There is direct NE innervation of the Edinger-Westphal nucleus, a major center for pupillary control; and (2) the LC is the main source of NE. However, this argument is incorrect, because the NE innervation of the Edinger-Westphal nucleus originates from lower medullary NE cell groups, and not from the LC (Levitt & Moore, 1979). Furthermore, the α -adrenoceptors on the iris dilator muscle receive stimulation from peripheral sympathetic fibers, not from LC fibers (Loewenfeld, 1993).

(Manuscript received May 26, 2008;
revision accepted for publication November 2, 2009.)