Article

Current Biology

Strategic stabilization of arousal boosts sustained attention

Graphical abstract

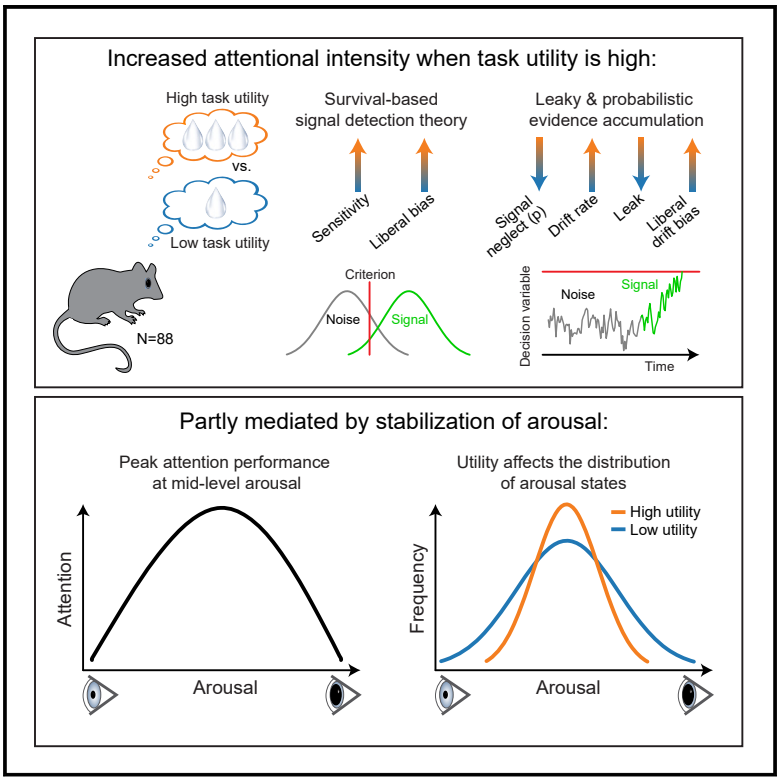

Highlights

- Mice increase their attentional intensity when task utility is high
- Moderate arousal is optimal for sustained feature-based attention
- Stabilizing arousal near optimality partly mediates the utility effect on attention
- New modeling approaches strongly implicate stochastic signal neglect

Authors

Jan Willem de Gee, Zakir Mridha, Marisa Hudson, ..., Hong Jiang, Wenhao Zhang, Matthew J. McGinley

Correspondence

j.w.degee@uva.nl (J.W.d.G.), matthew.mcginley@bcm.edu (M.J.M.)

In brief

de Gee et al. monitor pupil-linked arousal and behavioral performance of mice in a feature-based, sustained-attention task with intermittently shifting task utility. A suite of new survival analysis and diffusion-based decision models show that arousal and utility have major, partially intersecting impacts on multiple aspects of choice behavior.

de Gee et al., 2024, Current Biology *34*, 4114–4128
September 23, 2024
https://doi.org/10.1016/j.cub.2024.07.070

CellPress

CellPress
OPEN ACCESS

Current Biology

Article

Strategic stabilization of arousal boosts sustained attention

Jan Willem de Gee,[1,2,3,4,*] Zakir Mridha,[1,2] Marisa Hudson,[1,2] Yanchen Shi,[1,2] Hannah Ramsaywak,[1,2] Spencer Smith,[1,2] Nishad Karediya,[1,2] Matthew Thompson,[1,2] Kit Jaspe,[1,2] Hong Jiang,[1,2] Wenhao Zhang,[1,2] and Matthew J. McGinley[1,2,5,6,*]

[1]Department of Neuroscience, Baylor College of Medicine, 1 Baylor Plaza, Houston, TX 77030, USA
[2]Jan and Dan Duncan Neurological Research Institute, Texas Children's Hospital, 1250 Moursund Street, Houston, TX 77030, USA
[3]Cognitive and Systems Neuroscience, Swammerdam Institute for Life Sciences, University of Amsterdam, Science Park 904, Amsterdam 1098 XH, the Netherlands
[4]Research Priority Area Brain and Cognition, University of Amsterdam, Science Park 904, Amsterdam 1098 XH, the Netherlands
[5]Department of Electrical and Computer Engineering, Rice University, 6100 Main Street, Houston, TX 77005, USA
[6]Lead contact
*Correspondence: j.w.degee@uva.nl (J.W.d.G.), matthew.mcginley@bcm.edu (M.J.M.)
https://doi.org/10.1016/j.cub.2024.07.070

SUMMARY

Arousal and motivation interact to profoundly influence behavior. For example, experience tells us that we have some capacity to control our arousal when appropriately motivated, such as staying awake while driving a motor vehicle. However, little is known about how arousal and motivation jointly influence decision computations, including if and how animals, such as rodents, adapt their arousal state to their needs. Here, we developed and show results from an auditory, feature-based, sustained-attention task with intermittently shifting task utility. We use pupil size to estimate arousal across a wide range of states and apply tailored signal-detection theoretic, hazard function, and accumulation-to-bound modeling approaches in a large cohort of mice. We find that pupil-linked arousal and task utility both have major impacts on multiple aspects of task performance. Although substantial arousal fluctuations persist across utility conditions, mice partially stabilize their arousal near an intermediate and optimal level when task utility is high. Behavioral analyses show that multiple elements of behavior improve during high task utility and that arousal influences some, but not all, of them. Specifically, arousal influences the likelihood and timescale of sensory evidence accumulation but not the quantity of evidence accumulated per time step while attending. In sum, the results establish specific decision-computational signatures of arousal, motivation, and their interaction in attention. So doing, we provide an experimental and analysis framework for studying arousal self-regulation in neurotypical brains and in diseases such as attention-deficit/hyperactivity disorder.

INTRODUCTION

Arousal state and motivational state both profoundly impact behavior, but how they interact at the level of decision computations is largely unknown. Autonomic arousal, controlled by neuromodulators released from the reticular activating system, fluctuates on multiple time scales, shaping sensory processing and decision-making behavior.[1–4] Fluctuations in pupil size at constant luminance track the global arousal state[4,5] and the activity of the underlying neuromodulatory systems, including noradrenaline[6–11] and acetylcholine.[7,10,12] Optimal detection of simple sounds occurs at intermediate levels of spontaneously fluctuating ("tonic") pupil-linked arousal.[13–15] Large baseline pupil size is associated with task disengagement, including exploratory behaviors like locomoting,[14,16,17] and small baseline pupil is associated with drowsiness or sleep.[14,18] This "inverted-U" shape of the effects of pupil-linked arousal on task exploitation versus task disengagement as exploration or rest is referred to as the Yerkes-Dodson relationship and relates closely to the "adaptive gain" theory of noradrenergic tone.[1,19–24]

Sustained, feature-based attention is an important cognitive function that is potently influenced by motivational state and can be effectively modeled in mice.[25,26] Motivation level, for example, driven by the magnitude of reward for attending effectively (called "task utility"), impacts the intensity aspect of attention (also called "attentional effort").[27,28] In lay terms, an individual's motivation level can turn their attention level up and down, often referred to as "paying more attention." In conventional tasks, signal-detection theory (SDT) analyses can capture motivated shifts in attentional intensity as changes in the discrimination sensitivity, in contrast to orthogonal changes in behavioral strategy, which are captured as changes in bias (also called "criterion").[29] In perceptual decisions, sensitivity and bias are thought to arise from the underlying process of accumulating decision-relevant sensory evidence, implemented in a distributed network of brain areas.[30–32] Thus, accumulation-to-bound models of

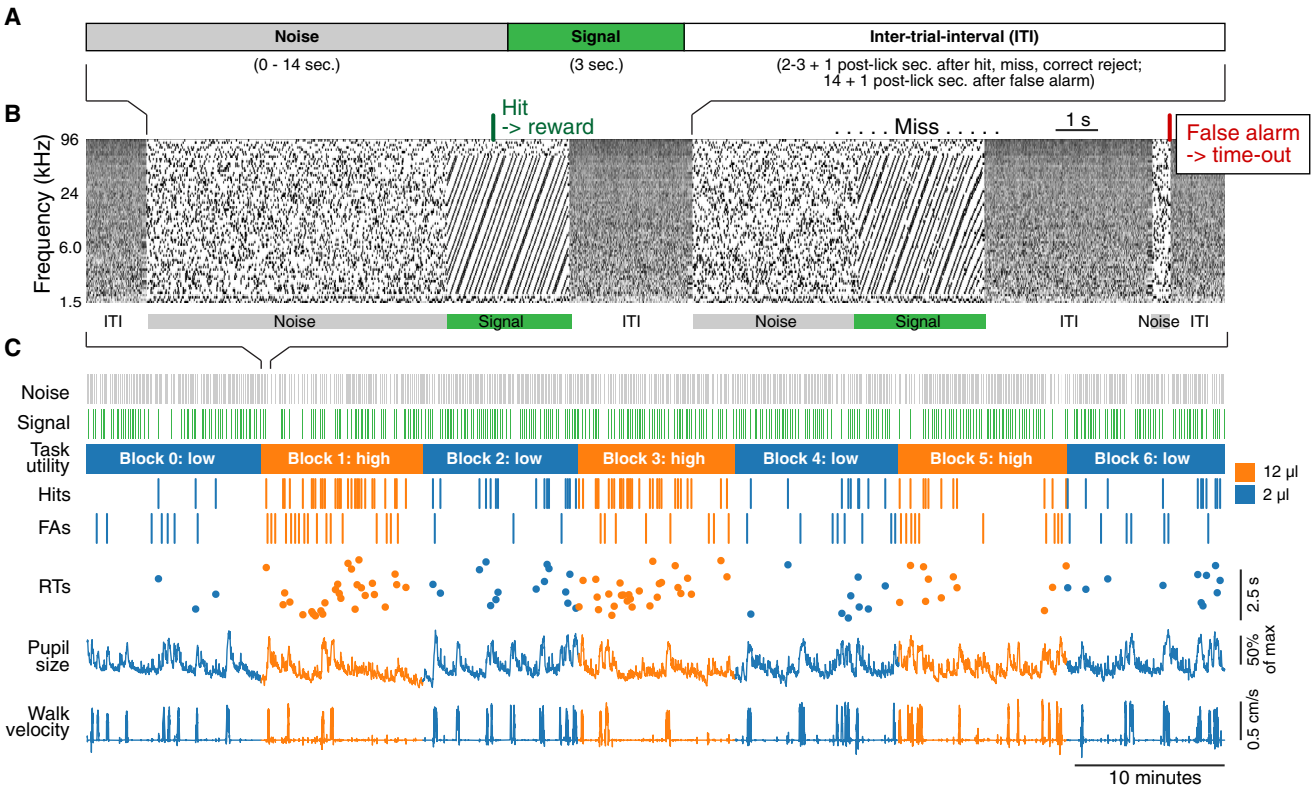

Figure 1. Monitoring performance and pupil-linked arousal during a sustained-attention value task

(A) Each trial consisted of three consecutive intervals: (i) the "noise" or tone cloud interval (exponentially distributed between 0 and 14 s), (ii) the "signal" (temporal coherence) interval (3 s), and (iii) the inter-trial interval (pink noise, uniformly distributed between 2 and 3 s).

(B) Spectrogram of the sound during an example of three consecutive trials. Correct-go responses (hits, green text, and line at top) were followed by a sugar water reward. Incorrect-go responses (false alarms, red text, and line at top) terminated the trial and were followed by a timeout (14 s ITI).

(C) Example session. From top to bottom: noise stimulus timing, signal stimulus timing, task utility (alternating between "low" and "high" in blocks of 60 trials), correct-go responses (hits), incorrect-go responses (false alarms), reaction times (RTs) on hit trials, pupil size, and walking velocity. Color (blue or orange) indicates utility context.

See also Figure S1.

judgments about weak time-varying sensory signals in noise can capture distinct aspects of decision processes and relate them to brain activity.[30,33–35]

How motivated shifts in attentional intensity interact with the inverted-U model of pupil-linked arousal is not known. This is an important gap in knowledge because dysregulated self-control of arousal is a prominent aspect of attention-deficit/hyperactivity disorder (ADHD) and other developmental disabilities and psychiatric disorders.[36–38] According to the adaptive gain theory, when task utility is high, one should regulate toward a moderate arousal state to facilitate optimal task engagement; by contrast, when utility is low, one should upregulate to a high-arousal state to facilitate exploration of alternatives or downregulate to a low arousal state to rest and consolidate. In line with this proposed adaptability of arousal self-control, the locus coeruleus receives strong projections from frontal regions including the orbital frontal cortex and anterior cingulate cortex,[1,7,39–41] reflecting a more complex organization than previously assumed.[6,42–44]

We here sought to understand how fluctuations in arousal and shifts in task utility interact to influence decision computations supporting sustained feature-based attention. We developed what we call a "sustained-attention value" task for head-fixed mice and report behavioral and pupillary signatures from a large cohort. We develop and apply a survival analysis-based SDT approach, related hazard function modeling, as well as tailored accumulation-to-bound models, and assess the interacting effects of task utility and pupil-linked arousal on task performance. We find that both arousal state and motivational state substantially impact decision computations. During periods of high task utility, mice exhibit multiple signatures of heightened attentional intensity, which are partially mediated by stabilization of arousal closer to optimal, moderate levels.

RESULTS

A sustained-attention value task to study attentional intensity allocation

To study how arousal and motivation interact to shape decision computations underlying the allocation of sustained attention, we developed a quasi-continuous auditory detection task for head-fixed mice (Figure 1). Mice were trained to detect coherent

time-frequency motion (the "signal"; called temporal coherence[45]) embedded at unpredictable times in an ongoing and otherwise random cloud of tones (the "noise") (Figures 1A and 1B). The task required sustained, attentive listening to achieve high detection performance due to the perceptual difficulty of noticing the temporally unpredictable emergence of a high-order acoustic feature. Mice were motivated to elicit lick responses by being food scheduled and by administration of sugar water reward when they licked during the signal. To suppress excessive licking, mice received a 14 s timeout if they licked during the noise (Figures 1A and 1B). We manipulated task utility by shifting the reward size back and forth between 2 and 12 μL in seven consecutive blocks of 60 trials. We recorded the timing of correct ("Hit") and incorrect ("False Alarm") responses (Figure 1C, top rows), the diameter of the pupil as a readout of arousal, and walking speed as an additional behavioral state measure (Figure 1C, bottom rows). Hereafter, we refer to the task as the sustained-attention value task (see STAR Methods for additional details).

Upregulation of task engagement and attentional intensity during high task utility

Because the sustained-attention value task is, by design, demanding of attentional resources, we hypothesized that mice would increase their attentional intensity during high task utility and reduce attentional intensity during low task utility to conserve cognitive resources and/or engage in other activities. We computed signal-detection theoretic bias as a measure of task engagement and sensitivity as (feature-based) attentional intensity (STAR Methods). Mice rapidly learned the task structure but slowly learned the feature-based attention and adaptive attentional intensity allocation (Figures S1H–S1K; STAR Methods), supporting that the task was perceptually difficult. We observed that mice spent most of the 1st block (termed block "0"; always low utility) becoming gradually engaged in the task (Figures 2A, 2C, 2E, 2G, S5E, and S5N). Therefore, we focused all analyses on the subsequent six blocks (termed blocks 1 to 6).

During high versus low task utility, we found that mice exhibited a more liberal choice bias (licked more during both noise and signal sounds; Figures 2A and 2B) and, critically, were also more sensitive (Figures 2C and 2D). The increased sensitivity indicates that mice could adapt feature-based attentional intensity to the current task utility. Across the session duration, mice became more conservative and less sensitive (Figures 2A–2D), probably resulting from effects of fatigue and/or satiety, both of which would progressively decrease task utility. Mirroring the patterns in bias and sensitivity, reaction times (RTs) were substantially shorter during high versus low task utility and gradually increased during the session (Figures 2E and 2F). These block-based analyses were performed starting 14 trials after the first hit in the block (STAR Methods) when performance in the block had stabilized (Figures S2A–S2C). We observed qualitatively similar results when considering all trials within a block (Figure S2D). In sum, increases in task utility boosted both task engagement, as indicated by the shift toward a liberal bias, as well as feature-based attentional intensity, as indicated by increased sensitivity and decreased reaction times. The liberal bias may also partially reflect increased attentional intensity (see decision computations improve during high task utility and moderate arousal, below).

An animal or human optimally performing a reward-harvesting task should allocate cognitive resources in a way that maximizes the rate of returns. We thus wondered if mice collected more rewards during high task utility. This is not trivially the case because in our quasi-continuous task, a noise stimulus always precedes each signal stimulus. Thus, unlike conventional go/no-go tasks, a liberal bias in our task, leading to more false alarms, lowers the probability that signals are presented, and thus lowers the potential reward rate (Figure S1L). Mice did ignore the noise and detected the signal on a larger fraction of trials during high task utility and thus collected more rewards (Figures 2G and 2H). Mice on average (±SEM) sustained a stable reward probability of 0.28 ± 0.01 across the entire session during the high task utility blocks (odd numbered), while reward probability declined monotonically from 0.26 ± 0.01 in the first low task utility block (block 2) to 0.21 ± 0.01 in the last (block 6; Figure 2H).

We next used a multiple regression hazard-modeling framework[46] (STAR Methods) to understand how these within-block and across-block effects interact to shape behavior. This allowed us to pinpoint the unique contribution of time-on-task (trial number) and task utility (high versus low) on choice behavior (Figure 2I), as well as the unique contribution of the previous trial's reward delivery (0, 2, or 12 μL) for up to seven trials back in time[47] (Figure 2I), as a more proximal cue for task utility. For overall responsiveness (closely related to bias), we observed the same behavioral patterns as in our stratification-based analyses (Figures 2A–2H, S2E, and S2F). Namely, there was a negative effect of time-on-task and a positive effect of task utility (Figure 3J). Additionally, the previous trial's reward delivery increased the overall responsiveness, but this effect decayed quickly with trial number (Figure 3J). For signal-selective responsiveness (closely related to sensitivity), we similarly observed a negative effect of time-on-task (Figure 3K). Additionally, previous trial's reward delivery strongly increased signal-selective responsiveness, with a much longer timescale that supplanted the effects of explicit task utility, which was now no longer significant (Figure 3K). Note that the lack of effect of reward delivery on the immediately previous trial on signal-selective responsiveness likely results from mice continuing to lick after rewards through the inter-trial interval into the subsequent trial (STAR Methods).

We used the same modeling framework to precisely compare the magnitude of the utility-driven changes in criterion and sensitivity. Changes in task utility had a bigger effect on responsiveness than signal-selective responsiveness ($t(87) = 13.98$, $p < 0.001$), also in a model with only trial number and task utility as covariates ($t(87) = 17.14$, $p < 0.001$).

Using several control analyses, we (1) ruled out effects of spurious hits from random licking with simulations (Figure S1L), (2) ruled out a timing strategy for matching licking to the noise duration statistics using a psychometric behavioral variant (Figures S3A–S3C), and (3) ruled out interpretative issues with false alarm rate due to the quasi-continuous nature of the task, using survival function analysis (Figure S3D). Finally, we asked if the block-wise changes in reward probability were predicted

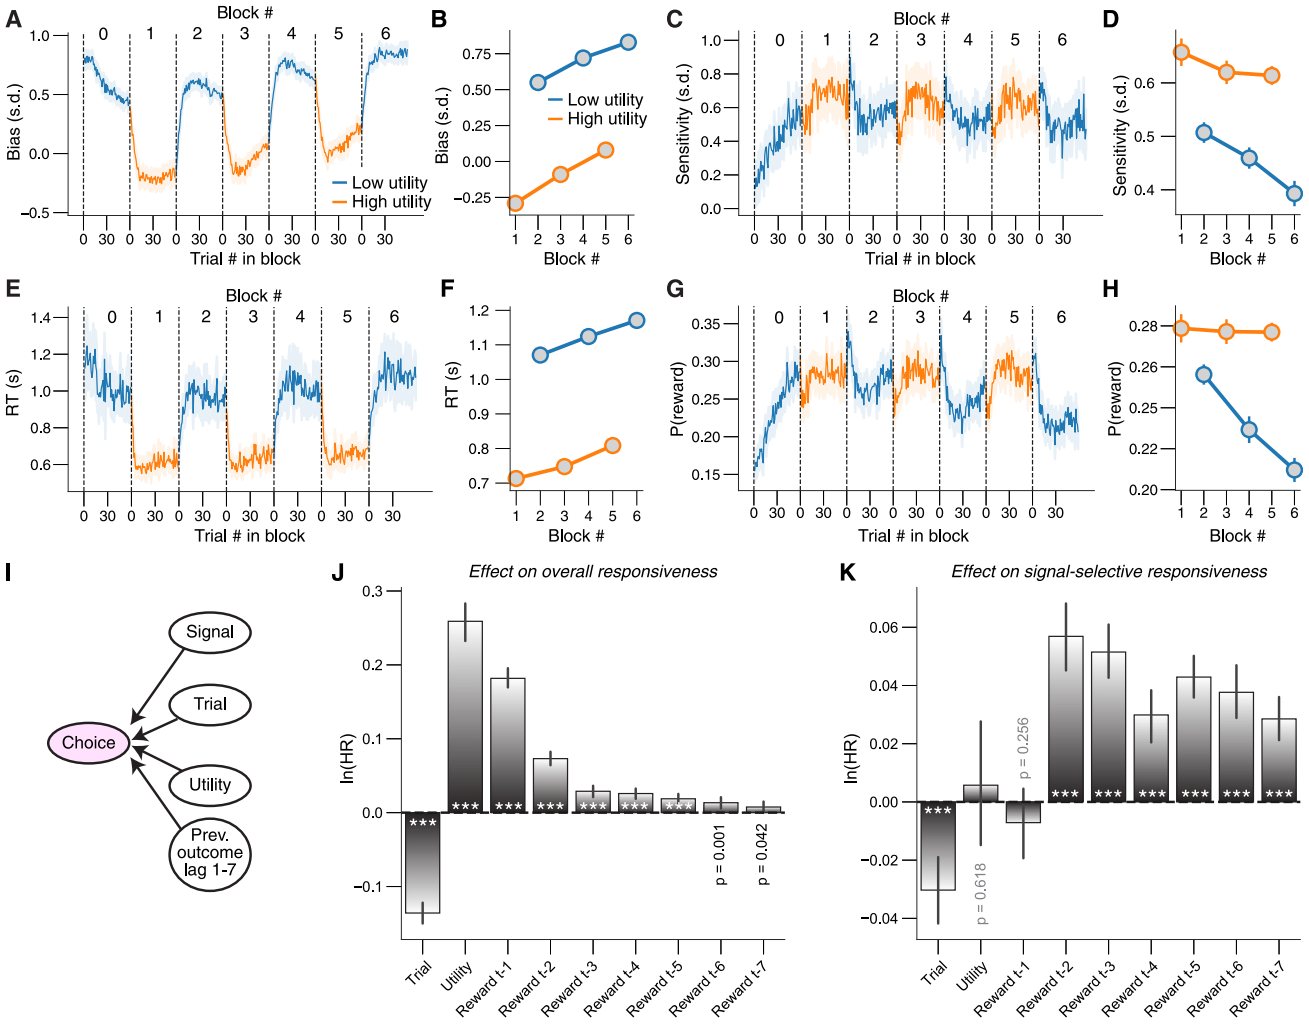

Figure 2. Rapid and adaptive changes in performance after shifts in task utility

(A) Bias across low and high task utility blocks within experimental sessions.

(B) As (A) but collapsed across trials within each block. Statistics, two-way repeated-measures ANOVA (factors, task utility [high versus low] and time-on-task [early, middle, and late]); main effect task utility, $F_{1,87} = 657.2$, $p < 0.001$; main effect time-on-task, $F_{2,174} = 124.5$, $p < 0.001$; interaction effect, $F_{2,174} = 4.0$, $p = 0.020$.

(C and D) As (A) and (B) but for sensitivity. Main effect task utility, $F_{1,87} = 33.2$, $p < 0.001$; main effect time-on-task, $F_{2,174} = 7.8$, $p = 0.001$; interaction effect, $F_{2,174} = 2.6$, $p = 0.076$.

(E and F) As (A) and (B) but for RT. Main effect task utility, $F_{1,87} = 499.3$, $p < 0.001$; main effect time-on-task, $F_{2,174} = 27.4$, $p < 0.001$; interaction effect, $F_{2,174} = 0.3$, $p = 0.741$.

(G and H) As (A) and (B) but for reward probability. Main effect task utility, $F_{1,87} = 19.7$, $p < 0.001$; main effect time-on-task, $F_{2,174} = 26.0$, $p < 0.001$; interaction effect, $F_{2,174} = 22.1$, $p < 0.001$.

(I) Schematic of multiple regression approach (STAR Methods).

(J) Hazard ratios (natural logarithm) from multiple regression model (I), capturing the effects of time-on-task, utility, and reward magnitude obtained in the previous trials (up to seven trials back) on overall responsiveness (closely related to bias). Statistics, Wilcoxon signed-rank test; ***$p < 0.001$.

(K) As (J) but for interaction effects between each factor and signal, capturing each factor's effect on signal-selective responsiveness (closely related to sensitivity).

All panels: shading or error bars, 68% confidence interval across animals ($n = 88$).

See also Figures S1–S3.

by behavioral adjustments in bias or sensitivity by analyzing individual differences in behavior. The results favored the latter. First, utility-driven shifts in bias and sensitivity were not positively correlated to each other across mice (Figure 3A), and neither were their overall, per-animal estimates (Figure S3E), suggesting separate underlying processes. Second, the magnitude of utility-driven shifts in reward probability was strongly correlated across mice with shifts in sensitivity (Figure 3B, left) but not bias (Figure 3B, right).

In sum, we conclude that mice increase their feature-based attentional intensity during periods of high task utility. When utility is high, mice are more sensitive in discriminating signals from

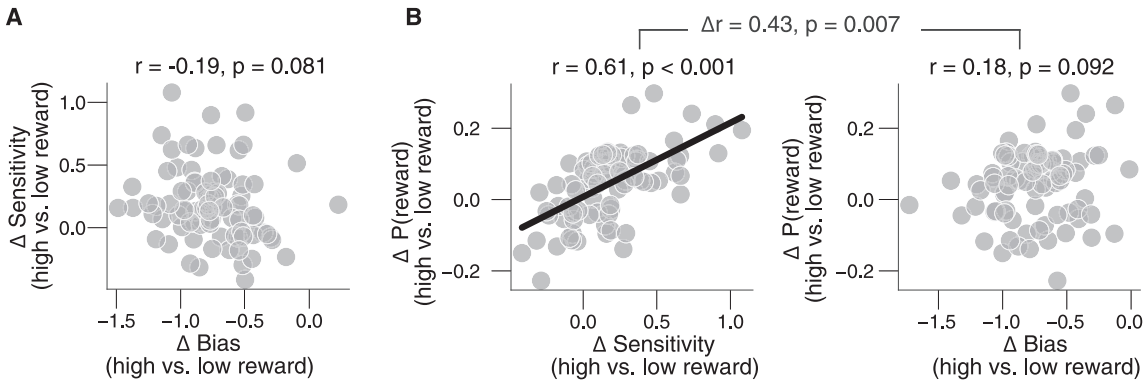

Figure 3. Adaptive allocation of attentional intensity is apparent in patterns of individual difference

(A) Change in sensitivity (high versus low task utility) plotted against change in bias. Every data point is a unique animal.

(B) Left: change in reward probability (high versus low task utility) plotted against change in sensitivity (high versus low task utility). Every data point is a unique animal. Right: as left, but for change in bias on the x axis. Statistics for difference in correlation, permutation test.

See also Figure S3.

noise and are faster at doing so, resulting in the collection of more rewards.

Mid-level pupil-linked arousal is optimal for fast and accurate performance

We previously showed that optimal signal-detection behavior in a simple tone-in-noise detection task occurred at intermediate levels of spontaneously fluctuating pupil-linked arousal.[14,15] Based on these results, we sought to test the hypothesis that the elevated task engagement (liberal bias shift) and attentional intensity (increased sensitivity) we observed during high task utility were partly caused by stabilization of arousal near its optimal state. To do so, we first characterized task performance as a function of arousal, independent from utility.

We quantified arousal as the diameter of the pupil (Figure 1C) measured immediately before each trial. On trials characterized by a mid-size pre-trial pupil diameter, we observed the smallest bias, highest sensitivity, shortest RTs, and highest reward probability (Figures 4A–4D; STAR Methods). Based on these results, we defined the optimal level of arousal as the pre-trial pupil size bin with maximal reward probability (green vertical line in Figure 4D), which also was the arousal-defined bin for which bias was minimal, sensitivity was maximal, and RTs were short (dashed green lines in Figures 4A–4C). Across animals, this optimal pre-trial pupil size was on average ~27.5% of its maximum in the session. Hit rates and false alarm rates also peaked at this pupil size (Figures S4D and S4E), and similar patterns were observed early and late in blocks (Figure S4F), across learning (Figure S4G), and for pre-trial pupil size measures derived with or without having regressed out effects of time-on-task and previous hit (Figure S4H; STAR Methods).

Locomotor status is another widely used marker of behavioral state.[14,48] Mice were considerably more liberal and less sensitive on trials associated with pre-trial walking (asterisks in Figures 4A and 4B), and reward probability was at its lowest (Figure 4D). These results suggest that feature-based attentional intensity was low during locomotion. We extended our regression model by including three additional predictors: walking (walk versus still), pupil-linked arousal (pre-trial pupil size), and pre-trial pupil size raised to the power of two (Figure 4E; STAR Methods). The resulting regression coefficients (Figures 4F and 4G) recapitulate the stratification-based effects of pupil (Figures 4A and 4B). We also verified that none of our results can be explained by whether mice walked on the previous trial (Figures S4I and S4J).

In sum, on trials characterized by a mid-size pre-trial pupil size, we observed the smallest bias, highest sensitivity, shortest RTs, and highest reward probability. We conclude that mid-level pupil-linked arousal is the optimal state for fast and accurate feature-based sustained attention.

Utility-related performance improvements are partially mediated by stabilization of pupil-linked arousal near optimality

Having identified the optimal state for performance of the task, we next determined whether mice spent more time in this optimal state during periods of high task utility. During high versus low task utility, the pre-trial pupil size was overall smaller (Figures S5A–S5C), and mice walked less frequently (Figures S5E–S5G). These results suggest an overall reduction in arousal during high task utility. Since the optimal pupil-linked arousal level of 27.5% was below the average arousal level for all blocks (range, 37%–39%), this reduction in arousal was a shift toward optimality. To capture how close each animal's arousal state on each trial was to the optimal level, we computed the absolute difference between pre-trial pupil size and the optimal size (STAR Methods). The distance from optimality was lower during high versus low task utility (Figures 5A and 5B). Furthermore, as the session progressed, their arousal state became less optimal, but importantly, mice maintained their arousal closer to the optimal state during high versus low task utility (Figures 5A and 5B).

We sought to determine if the changes in state in high-utility context were characterized by uniform downregulation of arousal or by stabilization near the optimal level. To test this, we first compared the distributions of state occupancy during high versus low task utility. Strikingly, mice spent less time in both the low- and high-arousal states, both of which are suboptimal (Figures 4A–4D), and upregulated a narrow range of states around optimality (Figures 5C and S5I). We next tested whether

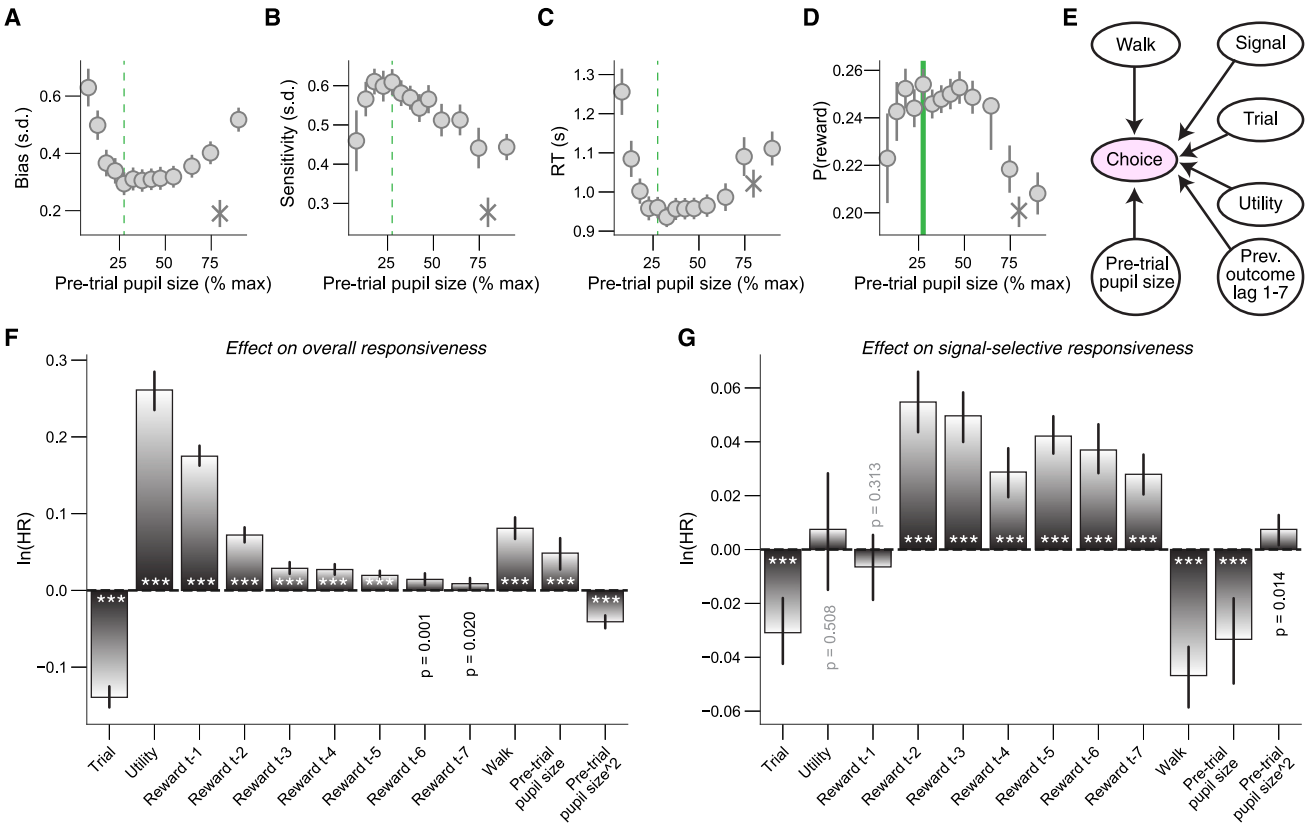

Figure 4. Optimal performance occurs at intermediate levels of pupil-linked arousal

(A) Relationship between pre-trial pupil size (time-on-task and previous hit regressed out; STAR Methods) and bias, irrespective of task utility. A first-order (linear) fit was superior to a constant fit ($F_{1,12}$ = 12.7, p = 0.004), and a second-order (quadratic) fit was superior to the first-order fit ($F_{1,12}$ = 7.8, p = 0.016; sequential polynomial regression; STAR Methods). Asterisk, walking trials (STAR Methods). Dashed green line, optimal arousal state (maximum probability peaks; D). See Figure S4K for results separately for high and low task utility.

(B) As (A) but for sensitivity. First-order fit, $F_{1,12}$ = 16.6, p = 0.002; second-order fit, $F_{1,12}$ = 0.7, p = 0.430.

(C) As (A) but for RT. First-order fit, $F_{1,12}$ = 3.2, p = 0.101; second-order fit, $F_{1,12}$ = 5.7, p = 0.034.

(D) As (A) but for reward probability. First-order fit, $F_{1,12} \sim 0.0$, p = 0.926; second-order fit, $F_{1,12}$ = 7.2, p = 0.020. Green line, optimal arousal state (maximum probability peaks).

(E) Schematic of multiple regression approach (STAR Methods).

(F) Hazard ratios (natural logarithm) from multiple regression model (E; STAR Methods), additionally capturing the effects of pre-trial walking and pre-trial pupil size on overall responsiveness (closely related to bias). Statistics, Wilcoxon signed-rank test; ***$p < 0.001$.

(G) As (F) but for interaction effects between each factor and signal, capturing effects on signal-selective responsiveness (closely related to sensitivity).

All panels: error bars, 68% confidence interval across animals (n = 88).

See also Figure S4.

changes in task utility were more associated with changes in the average pre-trial pupil size or with deviations of the pre-trial pupil size from the optimal size. To do so, we first compared the effect sizes of the main effects of task utility on both pupil-linked arousal measures. The partial η^2 was 0.33 for pre-trial pupil size and 0.55 for its distance from optimal, indicating a larger effect size for distance from optimal compared with raw pre-trial size. Second, we performed a logistic regression of block-wise utility (indicated as 0 or 1) on either Z scored block-wise pre-trial pupil size or its distance from optimal. The fitted coefficients were negative in both cases, but significantly more so for the measure of distance from optimality (Figure S5J). Thus, during heightened task utility, mice do not stereotypically downregulate their arousal state but instead up- and downregulate their arousal closer to its optimal level.

Having observed and quantified that epochs of high task utility are associated with both a more optimal pupil-linked arousal state and increased behavioral performance, we wondered to what extent the arousal stabilization contributed the utility-related performance effects. We addressed this with two complementary approaches. In a first approach, we first computed the pupil-linked and walk-related arousal-predicted behavioral performance during high and low task utility by plugging each trial's pre-trial pupil size into the previously observed relationship between arousal bins and behavior (Figures 4A–4D; STAR Methods). We then computed the difference in behavioral performance between high and low task utility that would have occurred solely because of the differing pupil-linked and walk-related arousal state sampling between utility contexts. This result showed that a small, but highly significant, portion of the

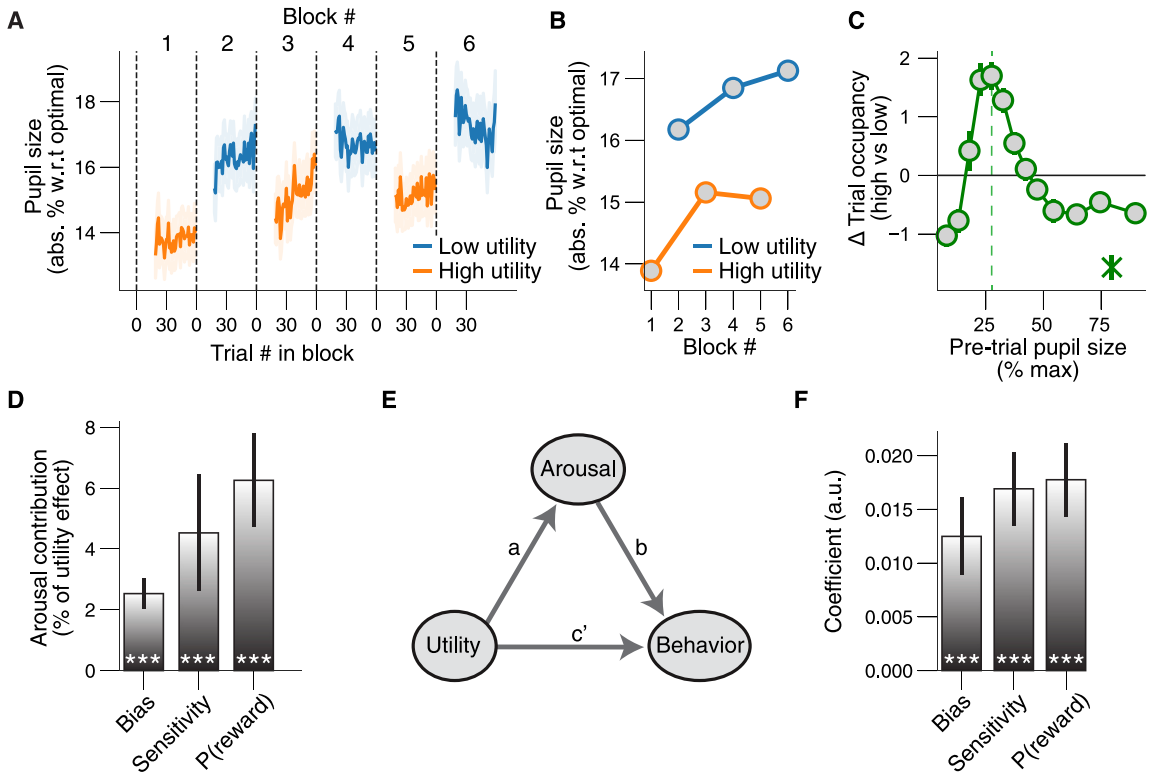

Figure 5. Mice regulate their pupil-linked arousal toward optimality when task utility is high

(A) Distance of pupil size from the optimal level (STAR Methods) across low and high task utility blocks within an experimental session.

(B) As (A) but collapsed across trials within a block. Statistics, two-way repeated-measures ANOVA (factors, task utility [high versus low] and time-on-task [early, middle, and late]); main effect task utility, $F_{1,87} = 125.6$, $p < 0.001$; main effect time-on-task, $F_{2,174} = 28.5$, $p < 0.001$; interaction effect, $F_{2,174} = 2.7$, $p = 0.070$.

(C) Change in trial density after increases in task utility, separately for pupil-defined arousal states; asterisk, walking trials.

(D) Percentage of the purely arousal-predicted shift in behavior compared with the total shift in behavior after changes in task utility (see main text). Statistics, Wilcoxon signed-rank test; ***$p < 0.001$.

(E) Schematic of mediation analysis of task utility to behavior via distance with regard to optimal measure (STAR Methods). Arrows, regressions; coefficient $a \times b$ quantifies the "indirect" (mediation) effect; coefficient c' quantifies the "direct effect."

(F) Fitted regression coefficients of the indirect path ($a \times b$; mediation). Statistics, Wilcoxon signed-rank test; ***$p < 0.001$.

All panels: error bars, 68% confidence interval across animals ($n = 88$).

See also Figure S5.

observed utility-based shift in performance could be accounted for solely by the shifts in pupil-linked and walk-related state occupancy (Figure 5D). We observed similar effects when considering pupil-linked arousal but not whether mice walked or remained still (Figure S5L).

In the second approach, we tested for statistical mediation[49] of arousal in the apparent effect of task utility on the different performance metrics (Figure 5E). We found that (the indirect path of) block-wise increases in task utility predicting block-wise decreases in distance from the optimal pupil-linked arousal state, in turn driving block-wise increases in sensitivity and reward probability, partially mediate the apparent effect of task utility on behavior (Figure 5F). We observed similar effects when only considering trials during which mice did not walk (Figure S5M) and when using pre-trial pupil size (Figure S5D) or walk probability (Figure S5H) instead of distance from the optimal arousal state as a mediator. Finally, we observed similar effects when repeating all analyses without having regressed out effects of time-on-task and previous outcome from the pre-trial pupil size measures (Figures S5N–S5S; STAR Methods). Taken together, we conclude that regulating pupil-linked arousal toward an optimal level partially implements the adaptive behavioral adjustments that match attentional intensity to its utility.

Accumulation-to-bound modeling of decision-making in the sustained-attention value task

Because the signal stimulus in our task was a high-order spectro-temporal statistic that emerged at unpredictable times in ongoing noise, correct detection required accumulation across time of partial evidence. Consistent with this perspective, reaction times were overall long and variable (Figures 6D and 6E). We therefore applied accumulation-to-bound sequential sampling modeling, as has extensively been used with similar tasks in the primate visual system[50] or for auditory click discrimination in rodents.[51] Although it has been shown that rodents can perform acoustic evidence accumulation,[51,52] it is not known if and how evidence accumulation is shaped by task utility or arousal.

The most widely used accumulation-to-bound models describe the complete accumulation (i.e., without forgetting) of

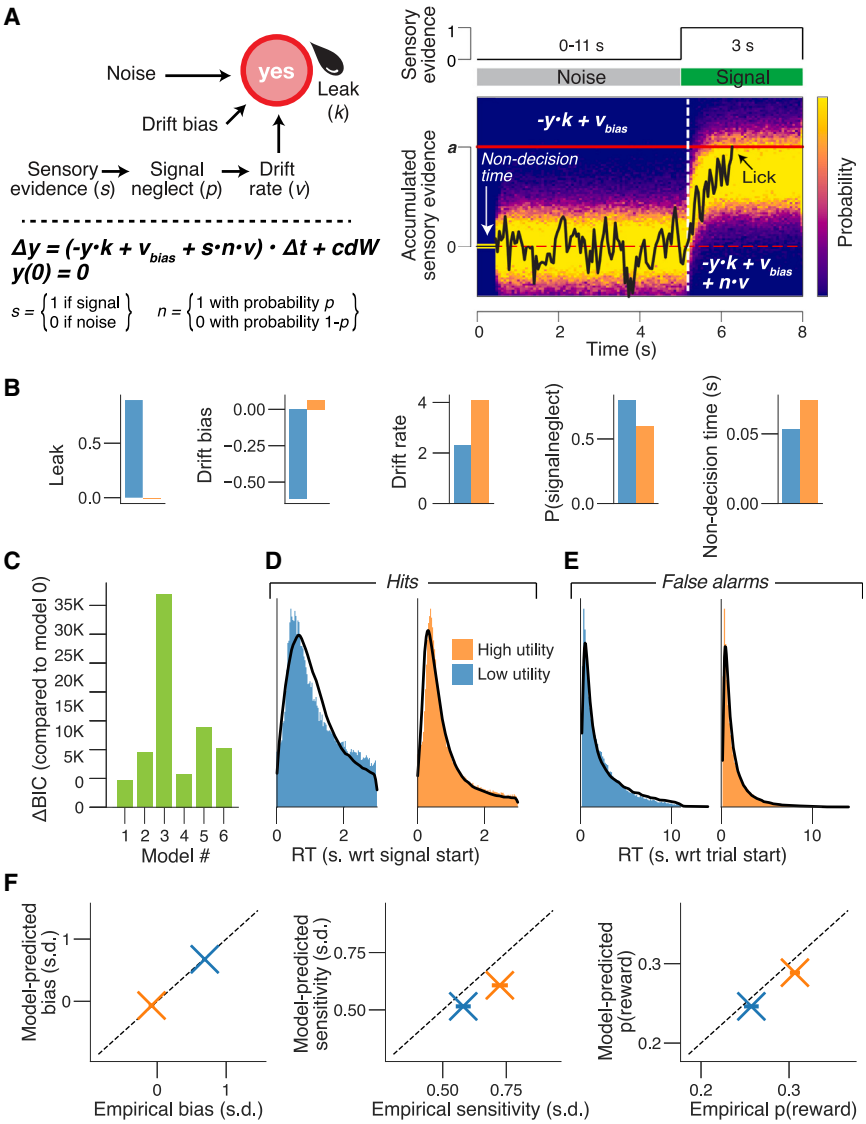

Figure 6. A leaky accumulation-to-bound model accounts for behavior in the sustained-attention value task

(A) Left: schematic of bounded accumulation model accounting for the fraction of go responses and their associated reaction times (RT). Right: the decision dynamics were governed by leak (k), drift bias (v_{bias}), and gaussian noise during the tone cloud, and additionally, signal neglect probability (n) and drift rate (v) during the signal stimuli. The decision terminated at bound height a.

(B) Parameter estimates.

(C) We compared the Bayesian information criterion (BIC) between seven models. The BIC for the winning model was used as a baseline. Lower BIC values indicate a model that is better able to explain the data, considering the model complexity; a ΔBIC of 10 is generally taken as a threshold for considering one model a sufficiently better fit.

(D) RT distribution for correct responses (hits) during low (left) and high (right) task utility. Black line, model fit.

(E) As (D), but for incorrect responses (false alarms).

(F) Model-predicted bias (left), sensitivity (middle), and reward probability (right) in during low and high task utility plotted against the empirical estimates. Dashed line, identity line.

(C)–(E) Pooled data across animals (n = 88).

See also Figure S6.

noisy sensory evidence as a decision variable that drifts to one of two decision bounds. Crossing a decision bound triggers a response, specifying a reaction time.[33,53–55] This model produces the fastest decisions for a fixed error rate when the evidence is stationary (i.e., its summary statistics are constant across time), as occurs in typical go/no-go or two-alternative forced choice tasks.[33] In our task, however, like perceptual decisions in most natural settings, the relevant evidence is not stationary. In this case, complete integration is suboptimal because it results in an excessive number of false alarms due to excessive integration of pre-signal noise.[56] Furthermore, a no-go bound is suboptimal in our task since the animal should learn that a signal occurs at some time point on all trials. We thus employed a diffusion modeling framework with leaky (i.e., forgetful) integration to a single decision bound[57] (STAR Methods).

We compared seven plausible models. The winning model contained six main parameters (Figures 6A and 6B; STAR Methods), the choices of which were motivated by task design and observed behavioral patterns. These parameters were (1) bound height, which determines how much evidence needs to be accumulated before committing to a go response; (2) non-decision time, which captures the combined duration of pre-decisional evidence encoding and post-decisional translation of choice into motor response; (3) leak, which controls the timescale of evidence accumulation; (4) drift bias, which is an evidence independent constant that is added to the drift; (5) mean drift rate, which controls the efficiency of accumulation of the relevant sensory feature (temporal coherence, in this case); and (6) signal neglect probability, which is the fraction of signal epochs for which the task-relevant sensory evidence is not accumulated. The fitted model (Figure 6B) accounted well for the behavior in the task, making accurate predictions for utility context-sorted RTs (Figures 6D and 6E), bias, sensitivity, and reward probability (Figure 6F).

The six alternative models (STAR Methods), including one with variable bound height, provided worse fits, both quantitatively (Figure 6C) and qualitatively (Figure S6). Two alternative models established that including the signal neglect probability parameter was necessary to account for the observation that correct responses (hits) were faster than expected for their frequency of occurrence (Figures S6E–S6H). A drift bias was necessary to account for frequent short latency errors (false alarms), especially during low task utility (Figures S6I–S6L). Unlike conventional two-bound models of go/no-go tasks, in the single-bound

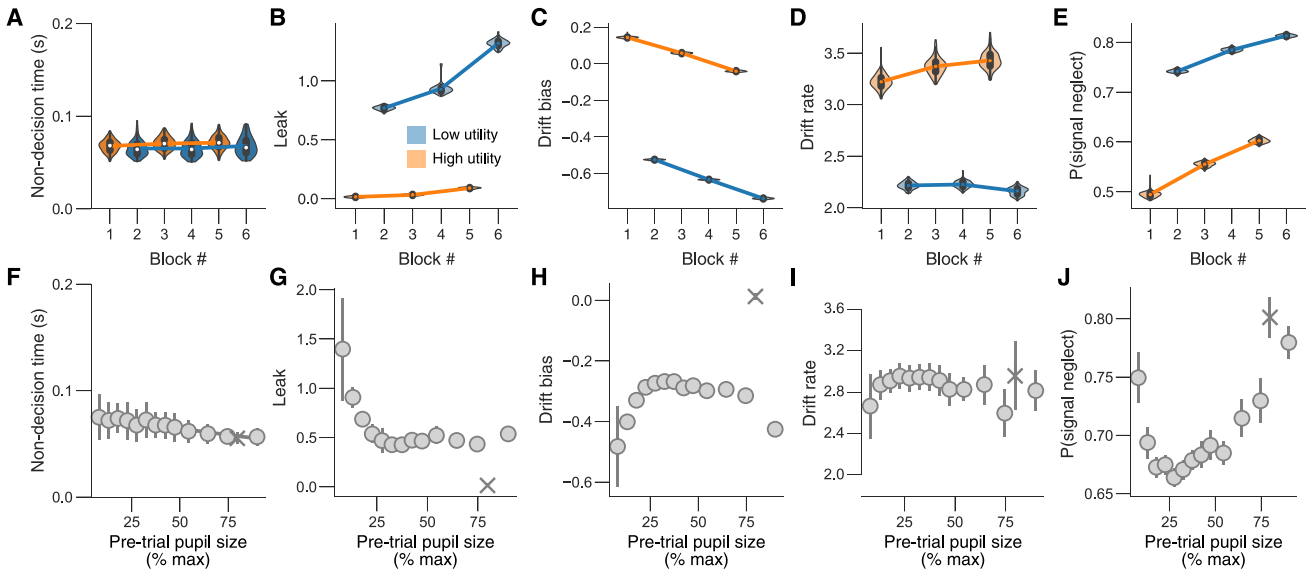

Figure 7. Task utility and pupil-linked arousal impact intersecting aspects of the decision computation

(A) Fitted non-decision time estimates (kernel density estimate of 100 bootstrapped replicates) separately per block number. Main effect task utility (fraction of bootstrapped parameter estimates in the low-utility blocks higher than in the high-utility blocks): $p = 0.32$. Main effect time-on-task (fraction of bootstrapped parameter estimates in the first two blocks higher than in the last two blocks): $p = 0.41$.

(B) As (A) but for leak. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p < 0.01$.

(C) As (A) but for drift bias. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p < 0.01$.

(D) As (A) but for drift rate. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p = 0.46$.

(E) As (A) but for signal neglect probability. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p < 0.01$.

(F) Fitted non-decision estimates (100 bootstrapped replicates) separately per arousal state (same pupil size defined bins as in Figures 4A–4D; irrespective of task utility; STAR Methods). A first-order (linear) fit was superior to a constant fit ($F_{1,12} = 128.8$, $p < 0.001$), and a second-order (quadratic) fit was not superior to the first-order fit ($F_{1,12} = 2.2$, $p = 0.168$; sequential polynomial regression; STAR Methods). Asterisk, walking trials (STAR Methods).

(G) As (F) but for leak. First-order fit, $F_{1,12} = 8.9$, $p = 0.012$; second-order fit, $F_{1,12} = 5.4$, $p = 0.039$.

(H) As (F) but for drift bias. First-order fit,$F_{1,12} = 8.8$, $p = 0.391$; second-order fit, $F_{1,12} = 7.7$, $p = 0.017$.

(I) As (F) but for drift rate. First-order fit, $F_{1,12} = 0.8$, $p = 0.378$; second-order fit, $F_{1,12} = 3.1$, $p = 0.101$.

(J) As (F) but for signal neglect probability. First-order fit, $F_{1,12} = 2.2$, $p = 0.165$; second-order fit, $F_{1,12} = 7.3$, $p = 0.019$.

See also Figure S7.

models we used here, increases in drift rate (or reductions in signal neglect probability) will affect SDT criterion as well as sensitivity because an increase in drift rate affects hit rate but not false alarm rate. Thus, effects of attentional intensity on these attention-related diffusion-model parameter will affect both SDT parameters (Figure 2).

Decision computations improve during high task utility and moderate arousal

We next sought to use our accumulation-to-bound model to dissociate distinct elements of the decision-making process that may underlie the effects of task utility on SDT behavioral metrics (Figure 2). Because time-on-task, previous trial outcome, and task utility had prominent effects on SDT measures (Figures 2 and 3), we first fitted the full model separately per block number, separately for each previous trial's outcome; however, bound height was fixed across conditions (STAR Methods). We found that the drift bias and drift rate were higher, and the leak and signal neglect probability were lower, during high versus low task utility (Figures 7B–7E). There was no significant effect of task utility on non-decision time (Figure 7A). The leak and signal neglect probability increased, and drift bias decreased, with time-on-task (Figures 7B, 7C, and 7E). These effects were similar when additionally accounting for previous trial outcome (Figures S7A–S7E). Furthermore, the main patterns did not depend on the specifics of the model: for each of the alternative models, leak and signal neglect probability were lower, and drift bias and drift rate were higher, during high versus low task utility (Figures S7U–S7Z).

Mirroring the above SDT analysis, we next sought to understand how the accumulation-to-bound model parameters changed with pupil-linked arousal state. We thus fitted the full model separately for each of the pupil-linked arousal bins. Non-decision time and leak decreased monotonically with arousal (Figures 7F, 7G, S7F, and S7G), suggesting faster sensory/motor processing and longer-lasting evidence accumulation in higher arousal states. Drift bias, drift rate, and signal neglect probability all exhibited U- or inverted-U shaped dependencies on arousal (Figures 7H–7J and S7H–S7J), indicating that the mice more reliably accumulated evidence from and acted upon the signal at moderate arousal levels. As with the utility-related effects, we observed similar effects of pupil-linked arousal whether or not we regressed out effects of time-on-task and previous outcome from the pre-trial pupil size measures (Figures S7K–S7T; STAR Methods).

In sum, increases in task utility resulted in more efficient evidence accumulation (higher drift rate), more reliable evidence accumulation (lower signal neglect probability), a longer accumulation time constant (lower leak), and a stronger urge to respond (higher drift bias). Furthermore, attention to the relevant feature is maximal at intermediate levels of pupil-linked arousal and worst during walking.

DISCUSSION

To efficiently meet their survival needs, organisms must regulate both of what Daniel Kahneman termed the "selective" and "intensive" aspects of attention.[58] Characterizing the selective aspect of attention has been a cornerstone of systems neuroscience,[59–61] but the intensive aspect has received comparatively little scrutiny. Recent work has emphasized the importance of motivational factors in driving attentional intensity (also called "attentional effort").[27,28,62–65] Indeed, heightened reward expectation increases perceptual sensitivity and reduces reaction times in humans and non-human primates.[27,66,67]

Parallel work has focused on the effects of arousal level on task performance, typified by the Yerkes-Dodson inverted-U dependence of performance on arousal level.[14,15,24] But how arousal and motivation interact to shape attentional intensity has not been explored. We have address this by here applying signal-detection theoretic, hazard function, and accumulation-to-bound modeling to a large dataset and showing that during high task utility, mice (1) collect more rewards; (2) accumulate perceptual evidence more efficiently, reliably, and across longer timescales; (3) suppress task-irrelevant locomotor behavior; and (4) stabilize their pupil-linked arousal closer to an optimal level.

Non-stationarity in task performance and its causes

Growing evidence supports that the neural computations underlying behavioral performance are not stationary within a session. Some of this non-stationarity seems to arise from fluctuations in internal state, as observed in neural activity patterns in primary sensory cortices[14,68–70] and sensory-guided behavior.[14] More recently, spontaneous shifts between engaged, biased, and/or disengaged states have been inferred from behavior.[71–73] However, the underlying causes or behavioral functions of these non-stationarities are not unclear.

Sequential-trial and time-on-task effects are two well-established sources of non-stationarity in behavior.[47,74–76] Indeed, our hazard function modeling showed that previous trial outcomes lasting up to seven trials impacted behavior, imparting some of the effects of utility context. Fatigue and/or satiety result in declining task utility and associated performance across a behavioral session. In our results, mice sustained their highest level of performance across the three high task utility blocks interspersed across long sessions in the difficult sustained-attention value task. Performance was comparably high in only the first low-utility block and then declined precipitously in subsequent low-utility blocks. These results suggest that later in a session, when satiated and fatigued, only large rewards were sufficient to motivate them to increase their attentional intensity.[77] Our accumulation-to-bound modeling showed that mainly changes in leak, drift bias, and signal neglect probability were responsible for these large time-on-task effects. Because we included no context cue, beyond the experienced rewards, mice needed to keep performing during low task utility at a sufficiently high level to detect the switch to high task utility. Thus, even larger effects of utility on attentional intensity might occur with a context cue.

Motivated shifts in attentional intensity

Task utility-based shifts in behavioral performance, but not attention-related, have been observed in rodents.[78,79] Other than in humans, the regulation of attentional intensity by utility has only been studied in monkeys.[27] Rodents can match their rate of learning to the statistics of a dynamic environment,[80] perform cost-benefit analysis,[78] and adapt their response vigor to task utility.[79] A recent study demonstrated that mice can consider their own information processing limitations to adaptively allocate the selectivity of their visual attention,[81] although effects of changes in overall task utility were not studied.

Our results show that in the presence of nonstationary overall task utility but stable selectivity requirements, both task engagement (as indicated by choice bias) and attentional intensity (as indicated by increased sensitivity) increased when task utility was high. In our accumulation-to-bound modeling, the dependencies of leak, drift rate, and signal neglect probability on task utility all support the conclusion that attentional intensity was increased during high versus low task utility. The higher drift rate and lower signal neglect probability in the high-utility blocks are likely responsible for the increased sensitivity and reward probability, as well as reduced RTs. The lower leak during high task utility indicates high sensory stimulus engagement, in the form of integrating the stimulus for a longer time, and thus is a logical component of sensory engagement. However, low leak may, or may not, improve reward harvesting because more noise is also accumulated.

The Yerkes-Dodson inverted-U model in motivated attention

Our results add to the growing evidence for an inverted-U model for the role of pupil-linked arousal in behavioral performance.[4,13–15,24,72] As in our previous work using a perceptually and procedurally simpler detection task and stationary task utility,[14,15] we here found that the optimal state for task performance was not at either extreme of arousal but rather was at a low-to-mid level. Small pupil size has been associated with sharp waves in the hippocampus and slow oscillations in auditory neocortex,[14] which are classic signatures of low arousal. In our task, compared with mid pupil size, small pre-trial pupil was associated with the following: conservative SDT bias, low SDT sensitivity, long RTs, lower reward probability, high sensory leak, conservative drift bias, and high probability of signal neglect. All of these patterns are consistent with drowsiness or some other resting form of disengagement. Large pupil size in our task was associated with conservative SDT bias (except during walking), low SDT sensitivity, longer RTs, decreased reward probability, increased walking, and a large increase in signal neglect probability; all of these are consistent with a hyper-aroused form of task disengagement. Importantly, mice spent less time in both the under- and over-aroused pupil states when in the high task utility context, which partially mediated the effects of utility context.

Our findings that both pupil-linked arousal and extent of walking were higher during low versus high task utility indicate that mice did not use the low-utility blocks simply to rest but

instead to also engage in alternative, aroused, and perhaps exploratory behaviors. This is in line with a recent observation that lapses in perceptual decisions reflect exploration,[82] which is consistent with the broad notion of an exploration-exploitation trade-off,[1,16,17] contributing to the right side of the inverted U. Our pupil results on the high-arousal right side of the inverted U are partially consistent with results of a recent study,[81] which found that large baseline pupil was associated with task disengagement. However, we observed that large pupil also was associated with a large reduction in accumulation of signal-related sensory evidence (a high signal neglect probability), whereas that study did not find effects of pupil on sensory encoding precision. This difference may result from their use of temporally predictable and simple (visual grating) signal-detection process versus our use of a more complex acoustic feature (temporal coherence), tracked across time.

Performance adaptations to utility versus difficulty

Our observed role for pupil-linked arousal in mediating adaptive adjustments in attentional intensity, based on task *utility*, is in contrast with the large literature on pupil dilation as a readout of attentional capacity (also called effort) driven by fluctuations in task *difficulty*,[83–87] including the extensive work on pupil-linked listening effort.[88,89] In this literature, the magnitude of pupil response to a stimulus is measured and typically compared between conditions that differ in perceptual or cognitive difficulty.[83,90] Motivational factors are customarily neglected, which has been widely acknowledged in the listening effort field,[89] but largely unaddressed. In our task, perceptual difficulty and selectivity were constant across the session, whereas motivation (driven by task utility) was changed in blocks. Furthermore, we focus on the pre-stimulus, so-called tonic, pupil-linked arousal measured before each trial rather than peri-stimulus, so-called phasic, stimulus-related pupil dilation.[91] Future work is needed to determine the interaction of these complementary arousal functions in behavior, perhaps by combining non-stationarities in both task utility and perceptual difficulty or selectivity.

Neural substrates for self-regulating arousal in attention

It is an open question which neuromodulatory systems contribute to the effects of pupil-linked arousal on attentional intensity. Our finding that pre-stimulus pupil size is lower during high versus low task utility is in line with the adaptive gain theory of locus coeruleus function[1,16,17] (but see Bari et al.[92]), implicating norepinephrine, which plays a major role in pupil control.[6–8,10,11] On the other hand, our results are not in line with the idea that increased acetylcholine improves attention[28,93] (but see Robert et al.[94]). The willingness to exert behavioral control is thought to be mediated by tonic mesolimbic dopamine[79,95,96] and/or serotonin,[97] but these effects are more related to behavioral strategy, whereas attentional intensity is more related perceptual processes. Serotonin and orexin/hypocretin have both been implicated in pupil control.[98,99]

Our results suggest a cost-benefit analysis being used to adapt pupil-linked arousal to an evolving motivational state.[65,100] Orbitofrontal cortex (OFC) and dorsal anterior cingulate cortex (dACC) are likely candidate effectors since both perform value computations related to optimizing behavior.[65,100–102] These structures are strongly connected to sensory cortices and to neuromodulatory nuclei, including LC.[1,7,39–41] Impaired frontal regulation of arousal is one of the hallmarks of ADHD[103] and also plays a role in autism[38] and a wide array of psychiatric disorders.[36,37] Future work could determine which neuromodulatory systems and frontal effectors support adaptive attentional intensity allocation, including in mouse models of neurological disorders.

Analysis foundations for quasi-continuous behavior

The interpretation of behavioral and physiological data regarding decision-making processes, including attention, has been greatly facilitated by using randomized, fixed-duration, trial-based task designs. These designs permit the straightforward application of a trio of powerful analysis approaches: (1) SDT to separate sensitivity and choice bias,[29] (2) logistic or probit regression models to assess the impact of behavioral or physiological covariates on sensitivity and bias,[104] and (3) accumulation-to-bound modeling to further parse decision processes.[33,53–55] Such a foundation of analysis approaches has been lacking for more natural behaviors. We here have developed adaptations of the analysis trio to allow comparable interpretability of quasi-continuous signal-detection behavior. Specifically, we (1) developed an SDT approach to estimate false alarm rate during variable duration noise epochs, derived from the survival function; (2) applied regression modeling to the hazard rate for response for the first time in decision-making; and (3) developed a family of diffusion models for quasi-continuous behavior. These approaches should greatly facilitate the interpretation of a broad range of more natural behaviors.

In sum, we found that mice can increase attentional intensity when it is useful, in part by regulating their pupil-linked arousal toward an optimal level. These results suggest that at least a part of the large behavioral and neural variability that can be explained by fluctuations in pupil-linked arousal state serves an adaptive function; states conducive to a particular biological need (attention to a rewarded stimulus) are upregulated at appropriate times (when the reward is large).

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead contact
  - Materials availability
  - Data and code availability
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
  - Animal Care & Sources
- METHOD DETAILS
  - Head post implantation
  - Behavioral task
  - Data acquisition
  - Pupil size
  - Walking speed
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Analysis and modeling of choice behavior
  - Trial exclusion criteria

- Model-free behavioral metrics
- Signal-detection theoretic modeling
- Time-varying survival regression
- Mediation analysis
- Accumulation-to-bound modeling
- Time course of learning
- Simulation of optimal signal-independent response rate
- Analysis of pupil data
- Preprocessing
- Quantification of pre-trial pupil size and distance w.r.t optimal
- Utility context-based pupil resampling simulations
- Analysis of walking data
- Statistical comparisons

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.cub.2024.07.070.

ACKNOWLEDGMENTS

We thank Daeyeol Lee for help designing the behavioral paradigm; Anton Banta for technical assistance; Sarim Aleem for help with pupil size analysis; many Rice University undergraduate students for help with animal behavioral training; Max Shinn, Konstantinos Tsetsos, and Peter Murphy for helpful discussions about the accumulation-to-bound modeling; and Nelson Totah for helpful comments on an early draft of the manuscript. Funding: NIDCD R01DC017797, NIDCD R03DC015618, NSF-NCS 2319492, and IDDRC P50 HD103555.

AUTHOR CONTRIBUTIONS

Conceptualization, J.W.d.G. and M.J.M.; funding acquisition, M.J.M.; data acquisition software, W.Z. and H.J.; investigation, Z.M., M.H., Y.S., H.R., S.S., N.K., M.T., and K.J.; data curation, J.W.d.G., M.H., Y.S., H.R., S.S., N.K., M.T., and K.J.; formal analysis, J.W.d.G. and M.J.M.; methodology, J.W.d.G. and M.J.M.; writing – original draft, J.W.d.G. and M.J.M.; writing – review and editing, J.W.d.G., Z.M., and M.J.M.; supervision, M.J.M.; project administration, M.J.M.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: January 18, 2023
Revised: July 17, 2024
Accepted: July 18, 2024
Published: August 15, 2024

REFERENCES

1. Aston-Jones, G., and Cohen, J.D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. Annu. Rev. Neurosci. *28*, 403–450. https://doi.org/10.1146/annurev.neuro.28.061604.135709.
2. Harris, K.D., and Thiele, A. (2011). Cortical state and attention. Nat. Rev. Neurosci. *12*, 509–523. https://doi.org/10.1038/nrn3084.
3. Lee, S.-H., and Dan, Y. (2012). Neuromodulation of brain states. Neuron *76*, 209–222. https://doi.org/10.1016/j.neuron.2012.09.012.
4. McGinley, M.J., Vinck, M., Reimer, J., Batista-Brito, R., Zagha, E., Cadwell, C.R., Tolias, A.S., Cardin, J.A., and McCormick, D.A. (2015). Waking state: rapid variations modulate neural and behavioral responses. Neuron *87*, 1143–1161. https://doi.org/10.1016/j.neuron.2015.09.012.
5. Joshi, S., and Gold, J.I. (2020). Pupil size as a window on neural substrates of cognition. Trends Cogn. Sci. *24*, 466–480. https://doi.org/10.1016/j.tics.2020.03.005.
6. Breton-Provencher, V., and Sur, M. (2019). Active control of arousal by a locus coeruleus GABAergic circuit. Nat. Neurosci. *22*, 218–228. https://doi.org/10.1038/s41593-018-0305-z.
7. de Gee, J.W., Colizoli, O., Kloosterman, N.A., Knapen, T., Nieuwenhuis, S., and Donner, T.H. (2017). Dynamic modulation of decision biases by brainstem arousal systems. eLife *6*, 309. https://doi.org/10.7554/eLife.23232.
8. Joshi, S., Li, Y., Kalwani, R.M., and Gold, J.I. (2016). Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. Neuron *89*, 221–234. https://doi.org/10.1016/j.neuron.2015.11.028.
9. Murphy, P.R., O'Connell, R.G., O'Sullivan, M., Robertson, I.H., and Balsters, J.H. (2014). Pupil diameter covaries with BOLD activity in human locus coeruleus. Hum. Brain Mapp. *35*, 4140–4154. https://doi.org/10.1002/hbm.22466.
10. Reimer, J., McGinley, M.J., Liu, Y., Rodenkirch, C., Wang, Q., McCormick, D.A., and Tolias, A.S. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. Nat. Commun. *7*, 13289. https://doi.org/10.1038/ncomms13289.
11. Varazzani, C., San-Galli, A., Gilardeau, S., and Bouret, S. (2015). Noradrenaline and dopamine neurons in the reward/effort trade-off: a direct electrophysiological comparison in behaving monkeys. J. Neurosci. *35*, 7866–7877. https://doi.org/10.1523/JNEUROSCI.0454-15.2015.
12. Mridha, Z., de Gee, J.W., Shi, Y., Alkashgari, R., Williams, J., Suminski, A., Ward, M.P., Zhang, W., and McGinley, M.J. (2021). Graded recruitment of pupil-linked neuromodulation by parametric stimulation of the vagus nerve. Nat. Commun. *12*, 1539. https://doi.org/10.1038/s41467-021-21730-2.
13. Schriver, B.J., Bagdasarov, S., and Wang, Q. (2018). Pupil-linked arousal modulates behavior in rats performing a whisker deflection direction discrimination task. J. Neurophysiol. *120*, 1655–1670. https://doi.org/10.1152/jn.00290.2018.
14. McGinley, M.J., David, S.V., and McCormick, D.A. (2015). Cortical membrane potential signature of optimal states for sensory signal detection. Neuron *87*, 179–192. https://doi.org/10.1016/j.neuron.2015.05.038.
15. Beerendonk, L., Mejías, J.F., Nuiten, S.A., de Gee, J.W., Fahrenfort, J.J., and van Gaal, S. (2024). A disinhibitory circuit mechanism explains a general principle of peak performance during mid-level arousal. Proc. Natl. Acad. Sci. USA *121*, e2312898121. https://doi.org/10.1073/pnas.2312898121.
16. Gilzenrat, M.S., Nieuwenhuis, S., Jepma, M., and Cohen, J.D. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. Cogn. Affect. Behav. Neurosci. *10*, 252–269. https://doi.org/10.3758/CABN.10.2.252.
17. Jepma, M., and Nieuwenhuis, S. (2011). Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory. J. Cogn. Neurosci. *23*, 1587–1596. https://doi.org/10.1162/jocn.2010.21548.
18. Yüzgeç, Ö., Prsa, M., Zimmermann, R., and Huber, D. (2018). Pupil size coupling to cortical states protects the stability of deep sleep via parasympathetic modulation. Curr. Biol. *28*, 392–400.e3. https://doi.org/10.1016/j.cub.2017.12.049.
19. Arnsten, A.F.T., and Li, B.-M. (2005). Neurobiology of executive functions: catecholamine influences on prefrontal cortical functions. Biol. Psychiatry *57*, 1377–1384. https://doi.org/10.1016/j.biopsych.2004.08.019.
20. Berridge, C.W., and Waterhouse, B.D. (2003). The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. Brain Res.Brain Res. Rev. *42*, 33–84. https://doi.org/10.1016/s0165-0173(03)00143-7.
21. Kane, G.A., Vazey, E.M., Wilson, R.C., Shenhav, A., Daw, N.D., Aston-Jones, G., and Cohen, J.D. (2017). Increased locus coeruleus tonic activity causes disengagement from a patch-foraging task. Cogn. Affect. Behav. Neurosci. *17*, 1073–1083. https://doi.org/10.3758/s13415-017-0531-y.

22. Pfeffer, T., Ponce-Alvarez, A., Tsetsos, K., Meindertsma, T., Gahnström, C.J., van den Brink, R.L., Nolte, G., Engel, A.K., Deco, G., and Donner, T.H. (2021). Circuit mechanisms for the chemical modulation of cortex-wide network interactions and behavioral variability. Sci. Adv. *7*, eabf5620. https://doi.org/10.1126/sciadv.abf5620.

23. Usher, M., Cohen, J.D., Servan-Schreiber, D., Rajkowski, J., and Aston-Jones, G. (1999). The role of locus coeruleus in the regulation of cognitive performance. Science *283*, 549–554. https://doi.org/10.1126/science.283.5401.549.

24. Yerkes, R.M., and Dodson, J.D. (1908). The relation of strength of stimulus to rapidity of habit-formation. J. Comp. Neurol. Psychol. *18*, 459–482. https://doi.org/10.1002/cne.920180503.

25. Robbins, T.W. (2002). The 5-choice serial reaction time task: behavioural pharmacology and functional neurochemistry. Psychopharmacology *163*, 362–380. https://doi.org/10.1007/s00213-002-1154-7.

26. Wang, L., and Krauzlis, R.J. (2018). Visual selective attention in mice. Curr. Biol. *28*, 676–685.e4. https://doi.org/10.1016/j.cub.2018.01.038.

27. Ghosh, S., and Maunsell, J.H.R. (2021). Single trial neuronal activity dynamics of attentional intensity in monkey visual area V4. Nat. Commun. *12*, 2003. https://doi.org/10.1038/s41467-021-22281-2.

28. Sarter, M., Gehring, W.J., and Kozak, R. (2006). More attention must be paid: The neurobiology of attentional effort. Brain Res. Rev. *51*, 145–160. https://doi.org/10.1016/j.brainresrev.2005.11.002.

29. Green, D.M., and Swets, J.A. (1966). Signal Detection Theory and Psychophysics (Peninsula Publishing).

30. Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791–806. https://doi.org/10.1016/j.neuron.2013.10.047.

31. Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266–273. https://doi.org/10.1038/s41586-019-1787-x.

32. van Vugt, B., Dagnino, B., Vartak, D., Safaai, H., Panzeri, S., Dehaene, S., and Roelfsema, P.R. (2018). The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. Science *360*, 537–542. https://doi.org/10.1126/science.aar7186.

33. Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. *113*, 700–765. https://doi.org/10.1037/0033-295X.113.4.700.

34. Siegel, M., Engel, A.K., and Donner, T.H. (2011). Cortical network dynamics of perceptual decision-making in the human brain. Front. Hum. Neurosci. *5*, 21. https://doi.org/10.3389/fnhum.2011.00021.

35. Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. Neuron *60*, 215–234. https://doi.org/10.1016/j.neuron.2008.09.034.

36. de Lecea, L., Carter, M.E., and Adamantidis, A. (2012). Shining light on wakefulness and arousal. Biol. Psychiatry *71*, 1046–1052. https://doi.org/10.1016/j.biopsych.2012.01.032.

37. Sander, C., Hensch, T., Wittekind, D.A., Böttger, D., and Hegerl, U. (2015). Assessment of wakefulness and brain arousal regulation in psychiatric research. Neuropsychobiology *72*, 195–205. https://doi.org/10.1159/000439384.

38. Zhao, S., Liu, Y., and Wei, K. (2022). Pupil-linked arousal response reveals aberrant attention regulation among children with autism spectrum disorder. J. Neurosci. *42*, 5427–5437. https://doi.org/10.1523/JNEUROSCI.0223-22.2022.

39. Arnsten, A.F.T., and Goldman-Rakic, P.S. (1984). Selective prefrontal cortical projections to the region of the locus coeruleus and raphe nuclei in the rhesus monkey. Brain Res. *306*, 9–18. https://doi.org/10.1016/0006-8993(84)90351-2.

40. Joshi, S., and Gold, J.I. (2022). Context-dependent relationships between locus coeruleus firing patterns and coordinated neural activity in the anterior cingulate cortex. eLife *11*, e63490. https://doi.org/10.7554/eLife.63490.

41. Porrino, L.J., and Goldman-Rakic, P.S. (1982). Brainstem innervation of prefrontal and anterior cingulate cortex in the rhesus monkey revealed by retrograde transport of HRP. J. Comp. Neurol. *205*, 63–76. https://doi.org/10.1002/cne.902050107.

42. Poe, G.R., Foote, S., Eschenko, O., Johansen, J.P., Bouret, S., Aston-Jones, G., Harley, C.W., Manahan-Vaughan, D., Weinshenker, D., Valentino, R., et al. (2020). Locus coeruleus: a new look at the blue spot. Nat. Rev. Neurosci. *21*, 644–659. https://doi.org/10.1038/s41583-020-0360-9.

43. Totah, N.K., Neves, R.M., Panzeri, S., Logothetis, N.K., and Eschenko, O. (2018). The locus coeruleus is a complex and differentiated neuromodulatory system. Neuron *99*, 1055–1068.e6. https://doi.org/10.1016/j.neuron.2018.07.037.

44. McKinney, A., Hu, M., Hoskins, A., Mohammadyar, A., Naeem, N., Jing, J., Patel, S.S., Sheth, B.R., and Jiang, X. (2023). Cellular composition and circuit organization of the locus coeruleus of adult mice. eLife *12*, e80100. https://doi.org/10.7554/eLife.80100.

45. Shamma, S.A., Elhilali, M., and Micheyl, C. (2011). Temporal coherence and attention in auditory scene analysis. Trends Neurosci. *34*, 114–123. https://doi.org/10.1016/j.tins.2010.11.002.

46. Therneau, T.M., and Grambsch, P.M. (2000). The Cox Model. In Modeling Survival Data: Extending the Cox Model, T.M. Therneau, and P.M. Grambsch, eds. (Springer), pp. 39–77. https://doi.org/10.1007/978-1-4757-3294-8_3.

47. Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. J. Vis. *14*, 9. https://doi.org/10.1167/14.7.9.

48. Polack, P.-O., Friedman, J., and Golshani, P. (2013). Cellular mechanisms of brain state-dependent gain modulation in visual cortex. Nat. Neurosci. *16*, 1331–1339. https://doi.org/10.1038/nn.3464.

49. Baron, R.M., and Kenny, D.A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. J. Pers. Soc. Psychol. *51*, 1173–1182. https://doi.org/10.1037//0022-3514.51.6.1173.

50. Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038.

51. Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98. https://doi.org/10.1126/science.1233912.

52. Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223. https://doi.org/10.1038/nature14066.

53. Brody, C.D., and Hanks, T.D. (2016). Neural underpinnings of the evidence accumulator. Curr. Opin. Neurobiol. *37*, 149–157. https://doi.org/10.1016/j.conb.2016.01.003.

54. Laming, D.R.J. (1968). Information Theory of Choice-Reaction Times (Academic Press).

55. Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. *20*, 873–922. https://doi.org/10.1162/neco.2008.12-06-420.

56. Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T.H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. Curr. Biol. *23*, 981–986. https://doi.org/10.1016/j.cub.2013.04.039.

57. Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. Psychol. Rev. *108*, 550–592. https://doi.org/10.1037/0033-295x.108.3.550.

58. Kahneman, D. (1973). Attention and Effort (Prentice-Hall), pp. 218–226.

59. Carrasco, M. (2011). Visual attention: The past 25 years. Vision Res. *51*, 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012.

60. Fritz, J.B., Elhilali, M., David, S.V., and Shamma, S.A. (2007). Auditory attention—focusing the searchlight on sound. Curr. Opin. Neurobiol. *17*, 437–455. https://doi.org/10.1016/j.conb.2007.07.011.

61. Maunsell, J.H.R., and Treue, S. (2006). Feature-based attention in visual cortex. Trends Neurosci. *29*, 317–322. https://doi.org/10.1016/j.tins.2006.04.001.

62. Brehm, J.W., and Self, E.A. (1989). The intensity of motivation. Annu. Rev. Psychol. *40*, 109–131. https://doi.org/10.1146/annurev.ps.40.020189.000545.

63. Kurzban, R., Duckworth, A., Kable, J.W., and Myers, J. (2013). An opportunity cost model of subjective effort and task performance. Behav. Brain Sci. *36*, 661–679. https://doi.org/10.1017/S0140525X12003196.

64. Richter, M., Gendolla, G.H.E., and Wright, R.A. (2016). Three decades of research on motivational intensity theory: what we have learned about effort and what we still don't know. Chapter Five. In Advances in Motivation Science, A.J. Elliot, ed. (Elsevier), pp. 149–186. https://doi.org/10.1016/bs.adms.2016.02.001.

65. Shenhav, A., Botvinick, M.M., and Cohen, J.D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron *79*, 217–240. https://doi.org/10.1016/j.neuron.2013.07.007.

66. Engelmann, J.B., and Pessoa, L. (2014). Motivation sharpens exogenous spatial attention. Motiv. Sci. *1*, 64–72. https://doi.org/10.1037/2333-8113.1.S.64.

67. Locke, H.S., and Braver, T.S. (2008). Motivational influences on cognitive control: behavior, brain activation, and individual differences. Cogn. Affect. Behav. Neurosci. *8*, 99–112. https://doi.org/10.3758/CABN.8.1.99.

68. Goris, R.L.T., Movshon, J.A., and Simoncelli, E.P. (2014). Partitioning neuronal variability. Nat. Neurosci. *17*, 858–865. https://doi.org/10.1038/nn.3711.

69. Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. Nat. Neurosci. *22*, 1677–1686. https://doi.org/10.1038/s41593-019-0502-4.

70. Nestvogel, D.B., and McCormick, D.A. (2022). Visual thalamocortical mechanisms of waking state-dependent activity and alpha oscillations. Neuron *110*, 120–138.e4. https://doi.org/10.1016/j.neuron.2021.10.005.

71. Ashwood, Z.C., Roy, N.A., Stone, I.R., Urai, A.E., Churchland, A.K., Pouget, A., and Pillow, J.W.; International Brain Laboratory (2022). Mice alternate between discrete strategies during perceptual decision-making. Nat. Neurosci. *25*, 201–212. https://doi.org/10.1038/s41593-021-01007-z.

72. Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D.A., and Jaramillo, S. (2023). Decision-making dynamics are predicted by arousal and uninstructed movements. Preprint at bioRxiv. https://doi.org/10.1101/2023.03.02.530651.

73. Weilnhammer, V., Stuke, H., Standvoss, K., and Sterzer, P. (2023). Sensory processing in humans and mice fluctuates between external and internal modes. PLoS Biol. *21*, e3002410. https://doi.org/10.1371/journal.pbio.3002410.

74. Warm, J.S., Parasuraman, R., and Matthews, G. (2008). Vigilance requires hard mental work and is stressful. Hum. Factors *50*, 433–441. https://doi.org/10.1518/001872008X312152.

75. Akrami, A., Kopec, C.D., Diamond, M.E., and Brody, C.D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature *554*, 368–372. https://doi.org/10.1038/nature25510.

76. Fischer, J., and Whitney, D. (2014). Serial dependence in visual perception. Nat. Neurosci. *17*, 738–743. https://doi.org/10.1038/nn.3689.

77. Hernández-Navarro, L., Hermoso-Mendizabal, A., Duque, D., de la Rocha, J., and Hyafil, A. (2021). Proactive and reactive accumulation-to-bound processes compete during perceptual decisions. Nat. Commun. *12*, 7148. https://doi.org/10.1038/s41467-021-27302-8.

78. Reinagel, P. (2021). Rational regulation of water-seeking effort in rodents. Proc. Natl. Acad. Sci. USA *118*. e2111742118. https://doi.org/10.1073/pnas.2111742118.

79. Wang, A.Y., Miura, K., and Uchida, N. (2013). The dorsomedial striatum encodes net expected return, critical for energizing performance vigor. Nat. Neurosci. *16*, 639–647. https://doi.org/10.1038/nn.3377.

80. Grossman, C.D., Bari, B.A., and Cohen, J.Y. (2022). Serotonin neurons modulate learning rate through uncertainty. Curr. Biol. *32*, 586–599.e7. https://doi.org/10.1016/j.cub.2021.12.006.

81. Grujic, N., Brus, J., Burdakov, D., and Polania, R. (2022). Rational inattention in mice. Sci. Adv. *8*, eabj8935. https://doi.org/10.1126/sciadv.abj8935.

82. Pisupati, S., Chartarifsky-Lynn, L., Khanal, A., and Churchland, A.K. (2021). Lapses in perceptual decisions reflect exploration. eLife *10*, e55490. https://doi.org/10.7554/eLife.55490.

83. Alnæs, D., Sneve, M.H., Espeseth, T., Endestad, T., van de Pavert, S.H.P., and Laeng, B. (2014). Pupil size signals mental effort deployed during multiple object tracking and predicts brain activity in the dorsal attention network and the locus coeruleus. J. Vis. *14*, 1. https://doi.org/10.1167/14.4.1.

84. Hess, E.H., and Polt, J.M. (1964). Pupil size in relation to mental activity during simple problem-solving. Science *143*, 1190–1192. https://doi.org/10.1126/science.143.3611.1190.

85. Kahneman, D., Beatty, J., and Pollack, I. (1967). Perceptual deficit during a mental task. Science *157*, 218–219. https://doi.org/10.1126/science.157.3785.218.

86. Kahneman, D., and Beatty, J. (1966). Pupil diameter and load on memory. Science *154*, 1583–1585. https://doi.org/10.1126/science.154.3756.1583.

87. Laeng, B., Sirois, S., and Gredebäck, G. (2012). Pupillometry: a window to the preconscious? Perspect. Psychol. Sci. *7*, 18–27. https://doi.org/10.1177/1745691611427305.

88. Peelle, J.E. (2018). Listening effort: how the cognitive consequences of acoustic challenge are reflected in brain and behavior. Ear Hear. *39*, 204–214. https://doi.org/10.1097/AUD.0000000000000494.

89. Pichora-Fuller, M.K., Kramer, S.E., Eckert, M.A., Edwards, B., Hornsby, B.W.Y., Humes, L.E., Lemke, U., Lunner, T., Matthen, M., Mackersie, C.L., et al. (2016). Hearing impairment and cognitive energy: the framework for understanding effortful listening (FUEL). Ear Hear. *37* (*Supplement 1*), 5S–27S. https://doi.org/10.1097/AUD.0000000000000312.

90. Zekveld, A.A., Heslenfeld, D.J., Johnsrude, I.S., Versfeld, N.J., and Kramer, S.E. (2014). The eye as a window to the listening brain: neural correlates of pupil size as a measure of cognitive listening load. NeuroImage *101*, 76–86. https://doi.org/10.1016/j.neuroimage.2014.06.069.

91. de Gee, J.W., Tsetsos, K., Schwabe, L., Urai, A.E., McCormick, D., McGinley, M.J., and Donner, T.H. (2020). Pupil-linked phasic arousal predicts a reduction of choice bias across species and decision domains. eLife *9*, e54014. https://doi.org/10.7554/eLife.54014.

92. Bari, A., Xu, S., Pignatelli, M., Takeuchi, D., Feng, J., Li, Y., and Tonegawa, S. (2020). Differential attentional control mechanisms by two distinct noradrenergic coeruleo-frontal cortical pathways. Proc. Natl. Acad. Sci. USA *117*, 29080–29089. https://doi.org/10.1073/pnas.2015635117.

93. Hasselmo, M.E., and McGaughy, J. (2004). High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. In Progress in Brain Research Acetylcholine in the Cerebral Cortex (Elsevier), pp. 207–231. https://doi.org/10.1016/S0079-6123(03)45015-2.

94. Robert, B., Kimchi, E.Y., Watanabe, Y., Chakoma, T., Jing, M., Li, Y., and Polley, D.B. (2021). A functional topography within the cholinergic basal forebrain for encoding sensory cues and behavioral reinforcement outcomes. eLife *10*, e69514. https://doi.org/10.7554/eLife.69514.

95. Hamid, A.A., Pettibone, J.R., Mabrouk, O.S., Hetrick, V.L., Schmidt, R., Vander Weele, C.M., Kennedy, R.T., Aragona, B.J., and Berke, J.D.

(2016). Mesolimbic dopamine signals the value of work. Nat. Neurosci. *19*, 117–126. https://doi.org/10.1038/nn.4173.

96. Niv, Y., Daw, N.D., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. Psychopharmacology *191*, 507–520. https://doi.org/10.1007/s00213-006-0502-4.
97. Gutierrez-Castellanos, N., Sarra, D., Godinho, B.S., and Mainen, Z.F. (2022). Maturation of prefrontal input to dorsal raphe nucleus increases behavioral persistence in mice. Elife *13*, e93485. https://doi.org/10.7554/eLife.93485.
98. Cazettes, F., Reato, D., Morais, J.P., Renart, A., and Mainen, Z.F. (2021). Phasic activation of dorsal raphe serotonergic neurons increases pupil size. Curr. Biol. *31*, 192–197.e4. https://doi.org/10.1016/j.cub.2020.09.090.
99. Grujic, N., Tesmer, A., Bracey, E., Peleg-Raibstein, D., and Burdakov, D. (2023). Control and coding of pupil size by hypothalamic orexin neurons. Nat. Neurosci. *26*, 1160–1164. https://doi.org/10.1038/s41593-023-01365-w.
100. Botvinick, M., and Braver, T. (2015). Motivation and cognitive control: from behavior to neural mechanism. Annu. Rev. Psychol. *66*, 83–113. https://doi.org/10.1146/annurev-psych-010814-015044.
101. Akam, T., Rodrigues-Vaz, I., Marcelo, I., Zhang, X., Pereira, M., Oliveira, R.F., Dayan, P., and Costa, R.M. (2021). The anterior cingulate cortex predicts future states to mediate model-based action selection. Neuron *109*, 149–163.e7. https://doi.org/10.1016/j.neuron.2020.10.013.
102. Tremblay, L., and Schultz, W. (2000). Reward-related neuronal activity during go-nogo task performance in primate orbitofrontal cortex. J. Neurophysiol. *83*, 1864–1876. https://doi.org/10.1152/jn.2000.83.4.1864.
103. Barkley, R.A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. Psychol. Bull. *121*, 65–94. https://doi.org/10.1037/0033-2909.121.1.65.
104. DeCarlo, L.T. (1998). Signal detection theory and generalized linear models. Psychol. Methods *3*, 186–205. https://doi.org/10.1037/1082-989X.3.2.186.
105. Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. Nature *341*, 52–54. https://doi.org/10.1038/341052a0.
106. Kalbfleisch, J.D., and Prentice, R.L. (2011). The Statistical Analysis of Failure Time Data (John Wiley & Sons).
107. Macmillan, N.A., and Kaplan, H.L. (1985). Detection theory analysis of group data: estimating sensitivity from average hit and false-alarm rates. Psychol. Bull. *98*, 185–199. https://doi.org/10.1037/0033-2909.98.1.185.
108. Hautus, M.J. (1995). Corrections for extreme proportions and their biasing effects on estimated values ofd′. Behav. Res. Methods Instrum. Comput. *27*, 46–51. https://doi.org/10.3758/BF03203619.
109. Davidson-Pilon, C. (2019). lifelines: survival analysis in Python. J. Open Source Software *4*, 1317. https://doi.org/10.21105/joss.01317.
110. Vallat, R. (2018). Pingouin: statistics in Python. JOSS *3*, 1026. https://doi.org/10.21105/joss.01026.
111. Shinn, M., Lam, N.H., and Murray, J.D. (2020). A flexible framework for simulating and fitting generalized drift-diffusion models. eLife *9*, e56938. https://doi.org/10.7554/eLife.56938.
112. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. *21*, 1281–1289. https://doi.org/10.1038/s41593-018-0209-y.
113. Draper, N.R., and Smith, H. (1998). Applied Regression Analysis*326* (Wiley-Interscience).
114. Spiegelhalter, D.J., Best, N.G., and Carlin, B.P. (2002). Bayesian measures of model complexity and fit. J. R. *64*, 583–639.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Compiled behavior and pupil data	This paper	https://doi.org/10.5281/zenodo.11517468
Analysis code	This paper	https://doi.org/10.5281/zenodo.11517468
Experimental models: Organisms/strains		
Mouse; C57BL/6 strain	The Jackson Laboratory	JAX stock #000664, RRID: IMSR_JAX:000664
Mouse; Ai148	The Jackson Laboratory	JAX stock #030328, RRID: IMSR_JAX: 030328
Mouse; Ai162	The Jackson Laboratory	JAX stock #031562, RRID: IMSR_JAX: 031562
Mouse; ChAT-Cre	The Jackson Laboratory	JAX stock #006410, RRID: IMSR_JAX: 006410
Mouse; ChAT-Cre crossed with Ai162	In-bred	N/A
Software and algorithms		
LabVIEW	National Instruments	https://www.ni.com/en/shop/labview.html, RRID: SCR_014325
Python	N/A	https://www.python.org/

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources may be directed to, and will be fulfilled by the lead contact, Matthew McGinley (matthew.mcginley@bcm.edu).

Materials availability

This study did not generate new materials.

Data and code availability

All resources, including data and code used for the analyses on this paper (in Python), are publicly available at https://doi.org/10.5281/zenodo.11517468

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Animal Care & Sources

All surgical and animal handling procedures were carried out in accordance with the ethical guidelines of the National Institutes of Health and were approved by the Institutional Animal Care and Use Committee (IACUC) of Baylor College of Medicine. A total of 114 animals were trained through to at least 5 sessions of the final phase of the task (see behavioral task). We excluded 26 animals from the analysis (Figures S1E–S1G) who had less than 5 sessions worth of data per animal after excluding sessions with an overall reward probability (see analysis and modeling of choice behavior) of less than 15%. Thus, all remaining analyses are based on 88 mice (74 male, 14 female) aged 7-8 weeks at training onset. Wild-type mice were of C57BL/6 strain (Jackson Labs) (N=51; 1 female). Various heterozygous transgenic mouse lines used in this study were of Ai148 (IMSR Cat# JAX:030328; N=6; 3 females), Ai162 (IMSR Cat# JAX:031562; N=10, 3 females), ChAT-Cre (IMSR Cat# JAX:006410; N=3; all male), or ChAT-Cre crossed with Ai162 (N=18; 7 females). No differences were observed between genotypes or sexes, so results were pooled. This variety in genetic profile was required to target specific neural circuitries with two-photon imaging; the results of the imaging experiments are not reported, here. Mice received ad libitum water. Mice received ad libitum food on weekends but were otherwise placed on food restriction

to maintain ~90% normal body weight. Animals were trained Monday-Friday. Mice were individually housed and kept on a regular light-dark cycle. All experiments were conducted during the light phase.

METHOD DETAILS

Head post implantation

The surgical station and instruments were sterilized prior to each surgical procedure. Isoflurane anesthetic gas (2–3% in oxygen) was used for the entire duration of all surgeries. The temperature of the mouse was maintained between 36.5°C and 37.5°C using a homoeothermic blanket system. After anesthetic induction, the mouse was placed in a stereotax (*Kopf Instruments*). The surgical site was shaved and cleaned with scrubs of betadine and alcohol. A 1.5-2 cm incision was made along the scalp mid-line, the scalp and overlying fascia were retracted from the skull. A sterile head post was then implanted using dental cement.

Behavioral task

Each 'trial' of the behavior consisted of three consecutive intervals (Figures 1A and 1B): (i) the 'noise' or tone cloud interval, (ii) the 'signal' (temporal coherence) interval, and (iii) the inter-trial-interval (ITI). The duration of the noise interval was randomly drawn before each trial from an exponential distribution with mean of 5 seconds; this was done to ensure a flat hazard function for signal start time. In most sessions (82.8%), randomly drawn noise durations greater than 11 s were set to 11 s. In the remainder of sessions (17.2%), these trials were converted to a form of catch trial, consisting of 14 seconds of noise. Results were not affected by whether sessions included catch trials or not (Figures S3L and S3M), and thus results were pooled for all further analyses. The duration of the signal interval was 3 s. The duration of the ITI was uniformly distributed between 2 and 3 s, plus an additional second after the last lick during the ITI.

The noise stimulus was a 'tone cloud', consisting of consecutive chords of 20 ms duration (gated at start and end with 0.5 ms raised cosine). Each chord consisted of 12 pure tones, selected randomly from a list of semitone steps from 1.5-96 kHz. For the signal stimulus, after the semi-random (randomly jittered by 1-2 semitones from tritone-spaced set of tones) first chord, all tones moved coherently upward by one semitone per chord. The ITI-stimulus was pink noise, which is highly perceptually distinct from the tone cloud. Stimuli were presented free field in the front left, upper hemifield at an overall intensity of 55 dB SPL (root-mean square [RMS]) using an intermittently recalibrated Tucker Davis ES-1 electrostatic speakers and custom software system in LabVIEW.

Mice were head-fixed on a wheel and learned to lick for sugar water reward to report detection of the signal stimulus. Correct-go responses (hits) were followed by either 2 or 12 μL of sugar water, depending on block number. Reward magnitude alternated between 2 and 12 μL in blocks of 60 trials, each. Incorrect-go responses (false alarms) terminated the trial and were followed by a 14 s timeout with the same pink noise ITI-stimulus. Correct no-go responses (correct rejecting the full 14 s of noise) in the sessions that contained catch trails were also followed by 2 or 12 μL of sugar water in some sessions (see above).

Training mice to perform the sustained-attention value task involved three separate phases. In phase 1, the signal was louder than the noise sounds (58 and 52 dB, respectively), 'classical conditioning' trials (5 automatic rewards during the signal sounds for the first 5 trials) were included, and there were no block-based changes in reward magnitude (reward size was 5 μL after every hit, across the session). Phase 1 was conducted for four experimental sessions (Figure S1A). In phase 2, we introduced the block-based changes in task utility. Once mice obtained a reward probability higher than 0.25, and the fraction of trials resulting in a false alarm was below 0.5 for two out of three sessions in a row, they were moved up to the phase 3. Phase 2 lasted for 2 - 85 (median, 9) experimental sessions (Figure S1B). Phase 3 was the final version of the task, with signal and noise stimuli of equal loudness and without any classical conditioning trials.

In a subset of experiments, the signal quality was systematically degraded by reducing the fraction of tones that moved coherently through time-frequency space (Figure S3A). This is similar to reducing motion coherence in the classic random-dot motion task.[105] In these experiments, signal coherence was randomly drawn beforehand from six different levels: easy (100% coherence; as in the main task), hard (55-85% coherence), and four levels linearly spaced in between. In the main behavior, coherence ranged from 90-100% on each trial. Performance in the main behavior did not change with coherence level, so results were pooled across coherence levels.

After exclusion criteria (Figures S1E–S1G), a total of 88 mice performed between 5 and 60 sessions (2100–24,960 trials per subject) of the final version of the sustained-attention value task (phase 3), yielding a total of 1983 sessions and 823,019 trials. A total of 10 mice performed the experiment with psychometrically degraded signals; they performed between 5 and 28 sessions (2083–11,607 trials per subject), yielding a total of 142 sessions and 58,826 trials.

Data acquisition

Custom LabVIEW software was written to execute the experiments, and synchronize all sounds, licks, pupil videography, and wheel motion. Licks were detected using a custom-made infrared beam-break sensor.

Pupil size

We continuously recorded images of the right eye with a Basler GigE camera (acA780-75gm), coupled with a fixed focal length lens (55 mm EFL, f/2.8, for 2/3''; Computar) and infrared filter (780 nm long pass; Midopt, BN810-43), positioned approximately 8 inches from the mouse. An off-axis infrared light source (two infrared LEDs; 850 nm, Digikey; adjustable in intensity and position) was used to yield high-quality images of the surface of the eye and a dark pupil. Images (504 × 500 pixels) were collected at 15 Hz, using a

National Instruments PCIe-8233 GigE vision frame grabber and custom LabVIEW code. To achieve a wide dynamic range of pupil fluctuations, an additional near-ultraviolet LED (405-410 nm) was positioned above the animal and provided low intensity illumination that was adjusted such that the animal's pupil was approximately mid-range in diameter following placement of the animal in the set-up and did not saturate the eye when the animal walked.

Walking speed

We continuously measured treadmill motion using a rotary optical encoder (Accu, SL# 2204490) with a resolution of 8,000 counts/revolution.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analysis and modeling of choice behavior

All analyses were performed using custom-made Python scripts, unless stated otherwise.

Trial exclusion criteria

We excluded the first (low utility) block of each session, as mice spent this block (termed block '0') becoming engaged in the task (see Figures 2A, 2C, 2G, 2E, S5E, and S5N). We found that a small fraction of trials began during a lick bout that had already started during the preceding ITI. These trials were immediately terminated. These rare 'false start trials' (2.5±0.2 % s.e.m. of trials across mice), were removed from further analyses. When pooling data across trials within a block, we always excluded the first 14 trials after the first hit in each block (in both high and low-utility blocks; see also time course of behavioral adjustments, below).

Model-free behavioral metrics

Reward probability was defined as the fraction of trials that ended in a hit (a lick during the signal). Reaction time on hit trials was defined as the time from signal onset until the response (first lick).

Signal-detection theoretic modeling

We applied signal detection theory (SDT[29]) to assess perceptual sensitivity (d', also called discriminability) and bias (c, also called criterion). Two features of the quasi-continuous sustained-attention value task meant that care was required in calculating hit and false alarm rates: (i) on each trial, each signal was always preceded by a noise stimulus, and (ii) the duration of the noise varied across trials. Hit rate (HR) was calculated in the traditional manner, as the fraction of signals for which the mouse licked. Because signals only occurred if a mouse did not lick during the noise (did not make a false alarm), this HR is implicitly a conditional probability. The appropriate false alarm rate (FAR) was calculated as the time-in-noise matched conditional probability, but without signal present, which could be estimated using the Kaplan-Meier Survival function (see below for details). Sensitivity and bias were defined as the difference or average, respectively, of the z-scored HR and FAR.[29]

On each trial, if a signal epoch occurred, it was of fixed duration. Therefore, a hit rate could be calculated using the conventional definition:

$$HR = h / (h + m) \quad \text{(Equation 1)}$$

where 'h' is the total number of 'hits' (signals with a lick response), and 'm' is the total number of misses (signals without a lick response).

Because the signal on each trial started at a random time (drawn from an exponential distribution), this HR may depend on the time of the signal start. Therefore, we defined a signal-start time-dependent HR:

$$HR(t_s) = h(t_s) / (h(t_s) + m(t_s)) \quad \text{(Equation 2)}$$

where '$h(t_s)$' is the number of 'hits' that occurred for signals with a start time of 't_s', and '$m(t_s)$' is the number of misses for signals with a start time of 't_s'. Because t_s was drawn from a continuous distribution, rather than a discrete list of values, $HR(t_s)$ could be empirically estimated in bins of signal start time.

Because licks during the 'noise' epoch aborted the trial, resulting in a subsequent signal not being played, the observed signal start-time dependent HR corresponds to the following conditional probability:

$$HR(t_s) = p(t_s < t_l < (t_s + d_s) \ / \ t_l > t_s) \quad \text{(Equation 3)}$$

where 't_l' is the time of the first response (lick) and 'd_s' is the duration of the signal. In words, $HR(t_s)$ is the probability of licking during the signal, for a signal that starts at time 't_s' given that the mouse did not lick during the preceding noise.

Mirroring the above conditional probability definition of $HR(t_s)$, the corresponding conditional probability for time-varying FAR rate is:

$$FAR(t_{s*}) = p(t_{s*} < t_l < (t_{s*} + d_s) \ / \ t_l > t_{s*}) \quad \text{(Equation 4)}$$

where 't_{s*}' refers to the corresponding signal start time in the '$HR(t_s)$' calculation. The '*' is used to indicate that signal had not started yet. In words, '$FAR(t_{s*})$' is the time-varying FAR in the appropriate time windows matched to the '$HR(t_s)$' function.

Applying Bayes' rule to $FAR(t_{s*})$, and noting that $p(t_l > t_{s*} \;/\; t_{s*} < t_l < (t_{s*} + d_s)) = 1$, yields:

$$FAR(t_{s*}) = p(t_{s*} < t_l < (t_{s*} + d_s) \;/\; t_l > t_{s*}) = p(t_{s*} < t_l < (t_{s*} + d_s)) \;/\; p(t_l > t_{s*}) \quad \text{(Equation 5)}$$

And therefore that:

$$FAR(t_{s*}) = (S(t_{s*}) - S(t_{s*} + d_s)) / S(t_{s*}) \quad \text{(Equation 6)}$$

where $S(t)$is the survival function for licking during noise. $S(t)$could be empirically measured using the Kaplan-Meier estimate, with signal starts treated as right censored false alarm events.[106] Using these time-varying HR and FAR estimates, time-varying SDT measures could be calculated as:

$$d'(t_s) = z(HR(t_s)) - z(FAR(t_s)) \quad \text{(Equation 7)}$$

$$c(t_s) = -0.5 * (z(HR(t_s)) + z(FAR(t_s))) \quad \text{(Equation 8)}$$

These time-varying SDT estimates can be sampled from the empirical signal start-time distribution and averaged to yield single summary estimates.[107] All SDT estimates were adjusted by adding 0.5 to each outcome type (hit, miss, correct reject, false alarm) to prevent infinite values resulting '0' or '1' for HR or FAR.[108] For FAR estimates, this corresponded to an adjustment based on size of the risk set at the start of the analysis window ('t_s').[106]

Simulations (Figure S1L) and a lack of positive correlation between bias and sensitivity differences between blocks or animals (see Figures 3 and 4) support that our approach to SDT successfully encapsulates and orthogonalizes d' and c.

Time-varying survival regression

We fitted a number of Cox's time varying proportional hazard models[46] using the Python package *lifelines*.[109] The general formulation is given by:

$$\lambda(t) = \lambda_0(t) e^{(\beta' \boldsymbol{x} + \gamma' s(t) + \delta' \boldsymbol{x} : s(t))} \quad \text{(Equation 9)}$$

Where $\lambda(t)$ is the hazard rate for licking at time t, λ_0 is the non-parametric baseline hazard function (non-negative), $\boldsymbol{x}$ is a set of time-fixed covariates $\boldsymbol{x} = [x_1, \ldots, x_p]$, s is the time-varying covariate signal presence ('0', noise; '1', signal), the : operator indicates an interaction, and β', γ', and $\delta'$are the (vectors for the) fitted regression coefficients. Effects on bias should load onto β' coefficients and effects on sensitivity should load onto the δ' coefficients.

Model 1 included the time-varying covariate signal presence and the time-fixed covariates (i) trial number, (ii) task-utility (0, low task utility; 1, high task utility), (iii) reward delivery in the previous trial (0, 2 or 12 μL), and (iv-x) reward delivery in 2-7 trials back,[47] respectively. Each of the time-fixed covariates could interact with signal presence. Model 2 additionally included time-fixed covariates coding for (xi) whether the animal walked (0, still; 1, walk), (xii) pre-trial pupil size, and (xiv) pre-trial pupil size squared. Model 3 additionally included the time-fixed covariate coding for (xv) whether the animal walked on the previous trial (0, still; 1, walk).

Mediation analysis

We used mediation analysis[49] to characterize the interaction between task utility, arousal, and behavioral performance (Figures 5E and 5F) using the Python package *Pingouin*.[110] We fitted the following linear regression models based on standard mediation path analysis:

$$\mathbf{Y} = i_0 1 + c\mathbf{X} \quad \text{(Equation 10)}$$

$$\mathbf{M} = i_1 1 + a\mathbf{X} \quad \text{(Equation 11)}$$

$$\mathbf{Y} = i_2 1 + c'\mathbf{X} + b\mathbf{M} \quad \text{(Equation 12)}$$

where $\boldsymbol{Y}$ was a vector of the block-wise behavioral metric (e.g., sensitivity), $\boldsymbol{X}$ was binary a vector describing the block-wise task utility (0, low; 1, high), $\mathbf{M}$ was a vector of block-wise distance with respect to the optimal measures (see analysis of pupil data), and c, c', a, b, i_0, i_1 and i_2 were the free parameters of the fit. The parameters were fit

Accumulation-to-bound modeling

We fitted the choice and reaction time data, pooled from all animals, with accumulation-to-bound models of the decision variable. Models (Figures 6A and S6) were fitted based on continuous maximum likelihood using the Python package *PyDDM*.[111] The combination of model parameters determines the fraction of correct responses and their associated RT distributions (Figures 6D–6F). We employed a single accumulator model that describes the accumulation of noisy sensory evidence toward a single choice boundary for a go-response.

In the model, the decision dynamics were governed by leak, drift bias and gaussian noise during 'noise' (tone cloud) stimuli, and additionally by the drift rate and signal neglect probability during the 'signal' (temporal coherence) stimuli, based on the following equation:

$$\Delta y = (-y \cdot k + v_{bias} + s \cdot n \cdot v) \cdot \Delta t + cdW \quad \text{(Equation 13)}$$

where y is the decision variable (black example trace in Figure 6A, right), k is the leak and controls the effective time constant ($1/k$) for which accumulated evidence decays, v_{bias} is an evidence-independent constant that is added as a drift in the diffusion process, s is the stimulus category (0 during 'noise'; 1 during 'signal'), n, equal to $(1 - p)$, is a Bernoulli variable ('0' or '1'), determined with probability p as the fraction of signal presentations on which the relevant sensory evidence s was neglected (not accumulated), v is the drift rate and controls the overall efficiency of accumulation of relevant evidence (coherence), and cdW is Gaussian distributed white noise with mean 0 and variance $c^2\Delta t$. The value of c was chosen, based on the time step size, so that the standard deviation of the noise was '1'. Evidence accumulation terminated at the bound height (go response) or at the end of the trial (no-go response), whichever came first. The starting point of evidence accumulationwas fixed to '0'.

Changes in bound height, leak and drift predict similar enough changes in the fraction of go-responses and fits that it proved to be unstable when letting all these parameters vary freely with task utility. Therefore, we initially compared three different models: in model 0 all parameters could vary with task utility except bound height, in model 1 all parameters could vary except leak, and in model 2 all parameters could vary except drift bias. Model 0 produced the best fits, both quantitatively (Figure 6C) as well as qualitatively (compare Figures 6D–6F and S6A–S6D).

We fitted two additional alternative models to verify that the signal neglect probability was an essential parameter. Model 3 did not include signal neglect probability at all, and in model 4 signal neglect probability was fitted but could not vary with task utility. Both alternative models produced worse fits, both quantitatively (Figure 6C) and qualitatively (Figures S6E–S6H).

We fitted two additional alternative models to verify that drift bias was an essential parameter. Model 5 did not include drift bias at all, and in model 6 drift bias was fitted but could not vary with task utility. Both alternative models produced worse fits quantitatively (Figure 6C) and qualitatively (Figures S6I–S6L).

For the winning model (model '0' we let all parameters except bound height vary with block number (Figures 7A–7E). In a separate analysis, all parameters except bound height could vary with the different arousal states and for the low-utility and high-utility blocks (Figures S7F–S7J). For simplicity of illustration and to pinpoint the effects of arousal on decision-making that are independent of task utility we then averaged the fits across task utility (Figures 7F–7J).

Time course of learning

To characterize animal's learning, we fitted the following function:

$$B(s) = a \cdot e^{-b \cdot s} + c \quad \text{(Equation 14)}$$

where B is a behavioral metric of interest, s is session number with respect to start of phase 3, and a, b and c the free parameters of the fit.

Simulation of optimal signal-independent response rate

To characterize the theoretical relationship between an overall (signal-independent) Poisson response rate and our behavioral metrics we generated a simulated data set (Figure S1L). Specifically, we generated synthetic trials that matched the statistics of the empirical trials (noise duration was drawn from an exponential distribution with mean = 5 s; truncated at 11 s, the signal duration was 3 s). We then systematically varied the overall response rate by drawing random response times from exponential distributions with various means (1 / rate) and randomly assigned those response times to the synthetic trials. We varied the overall response rate from 0.05 to 1 responses/s, in 20 evenly spaced steps. For each response rate, the decision agent performed 500 thousand simulated trials. For each iteration we then calculated the resulting bias, sensitivity, RT, and reward probability.

Analysis of pupil data

All analyses were performed using custom-made Python scripts, unless stated otherwise.

Preprocessing

We measured pupil size and exposed eye area from the videos of the animal's eye using the Python package *DeepLabCut*.[12,112] In approximately 1000 training frames randomly sampled across all sessions, we manually identified 8 points spaced approximately evenly along the edge of the pupil, and 8 points along the edge of the eyelids. The network (resnet 110) was trained with default parameters. To increase the network's speed and accuracy when labeling (unseen) frames of all videos, we specified video-wise cropping values in the *DeepLabCut* configuration file that corresponded to a square around the eye. The pupil size (or exposed eye area) was computed as the area of an ellipse fitted to the detected pupil (exposed eye) points. If two or more points were labeled with a likelihood smaller than 0.1 (e.g., during blinks), we did not fit an ellipse, but flagged the frame as missing data. We then applied the following signal processing to the pupil (exposed eye) time series of each measurement session: (i) resampling to 10 Hz; (ii) blinks were detected by a custom algorithm that marked outliers in the z-scored temporal derivative of the pupil size time series; (iii) linear interpolation of missing or poor data due to blinks (interpolation time window, from 150 ms before until 150 ms after missing data); (iv)

low-pass filtering (third-order Butterworth, cut-off: 3 Hz); and (v) conversion to percentage of the 99.9 percentile of the time series. See Mridha et al.[12] and McGinley et al.[14] for additional details.

Quantification of pre-trial pupil size and distance w.r.t optimal

We quantified pre-trial pupil size as the mean pupil size during the 0.25 s before trial onset. Pre-trial pupil size was highest after previous hits (Figure S4A), likely because the phasic lick-related pupil response did not have enough time to return to baseline. Pre-trial pupil size also generally increased with time-on-task (Figure S5N). We thus removed (via linear regression) components explained by previous outcome (reward vs. no reward) and trial number. We obtained qualitatively similar results without doing so (Figures S5N–S5S and S7K–S7T). To capture how close the animal's arousal state on each trial was to the optimal level, we computed the absolute difference between each pre-trial's pupil size and the optimal size. Here, optimal size (which was found to be 27.5% of max) was defined as the pre-trial pupil size for which reward probability was maximal (green vertical line in Figure 4D).

Utility context-based pupil resampling simulations

Per animal and block type (high and low task utility) we counted the number of trials in each arousal-defined bin (same as bins as in Figure 5C); we then used these counts to compute the arousal-predicted behavioral performance (e.g., sensitivity), using the previously observed relationship between arousal states and behavioral performance (irrespective of task utility; Figures 4A–4D); per animal, we then computed the difference between the arousal-predicted behavioral performance during high versus low task utility. Finally, we divided the average (across animals) arousal-predicted difference by the average (across animals) total actually observed difference (Figure 2). To estimate uncertainty, we bootstrapped trials within animals and blocks (5K bootstraps). We then computed the fraction between the purely arousal-predicted change in behavior and the total change in behavior after changes in task utility (see Figure 5D).

Analysis of walking data

The instantaneous walking speed data was resampled to 10 Hz. We quantified pre-trial walking speed as the mean walking velocity during the 2 s before trial onset. We defined walking probability as the fraction of trials for which the absolute walking speed exceeded 1.25 cm/s (Figure S5B).

Statistical comparisons

We used a 3 × 2 repeated measures ANOVA to test for the main effects of task utility and time-on-task (block number of a given task utility), and their interaction (Figures 2B, 2D, 2F, 2H, and 5B). We used the non-parametric Wilcoxon signed rank test to test coefficients against 0 (Figures 2J, 2K, 4F, 4G, 5D, and 5F).

We used sequential polynomial regression analysis,[113] to quantify whether the relationships between pre-trial pupil size and behavioral measures were better described by a 1st order (linear) or 2nd order model (Figures 4A–4D and 7F–7J):

$$\mathbf{Y} = \beta_0 1 + \beta_1 \mathbf{X} + \beta_2 \mathbf{X}^2 \quad \text{(Equation 15)}$$

where **Y** was a vector of the dependent variable (e.g., bin-wise sensitivity), **X** was a vector of the independent variable (e.g. bin-wise pre-trial pupil size), and β as polynomial coefficients. To assess the amount of variance that each predictor accounted for independently, we orthogonalized the regressors prior to model fitting using QR-decomposition. Starting with the zero-order (constant) model and based on F-statistics,[113] we tested whether incrementally adding higher-order predictors improves the model significantly (explains significantly more variance). We tested 1st up to 5th order models.

We used Bayesian information criterion (BIC) for model selection and verified whether the complexity of the different variants of the accumulation-to-bound model was justified to account for the data (Figure 6C). A difference in BIC of 10 is generally taken as a threshold for considering one model a sufficiently better fit than another.[114] We directly compared bootstrapped distributions of the model parameter estimates to test for the main effects of task utility and time-on-task (Figures 7A–7E) and for the effects pupil-linked and walk-related arousal (Figures 7F–7J).

All tests were performed two-tailed. All error bars are 68% bootstrapped confidence intervals of the mean, unless stated otherwise.

Current Biology, Volume *34*

Supplemental Information

Strategic stabilization of arousal boosts sustained attention

Jan Willem de Gee, Zakir Mridha, Marisa Hudson, Yanchen Shi, Hannah Ramsaywak, Spencer Smith, Nishad Karediya, Matthew Thompson, Kit Jaspe, Hong Jiang, Wenhao Zhang, and Matthew J. McGinley

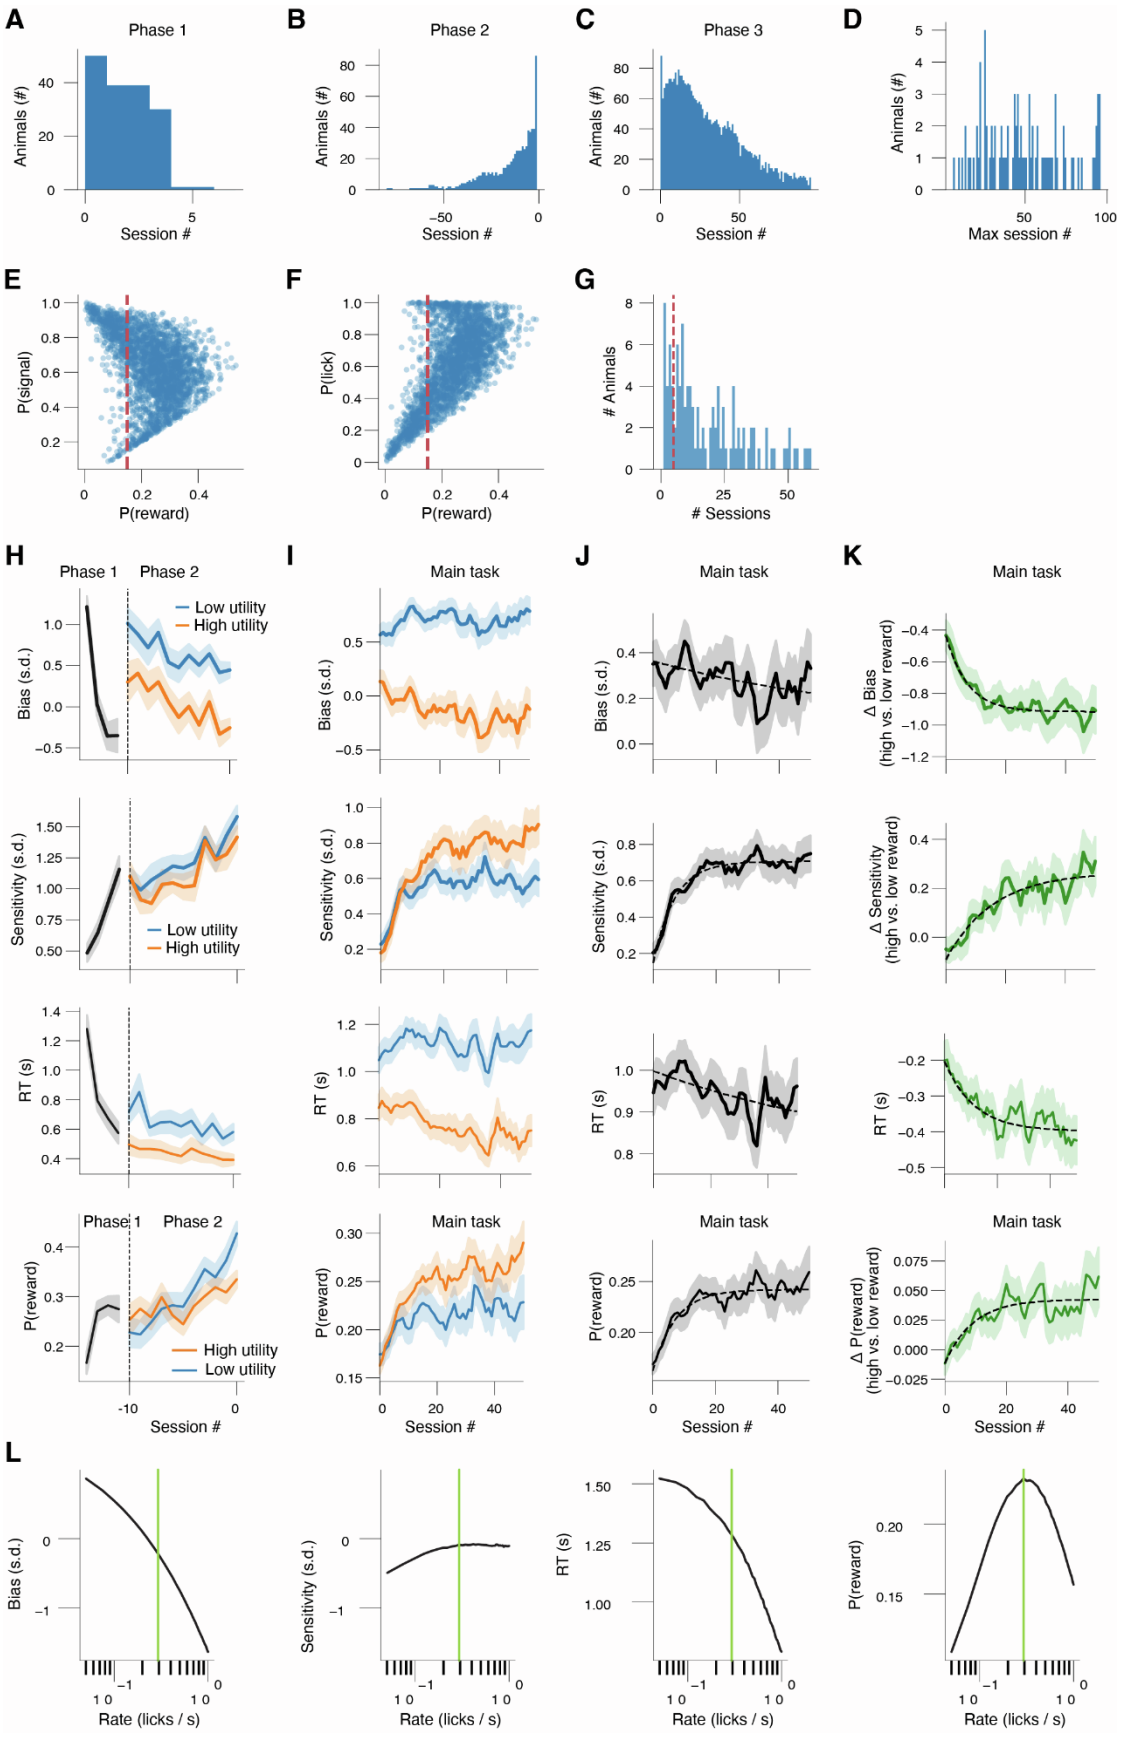
A
Phase 1
Animals (#)
Session #
B
Phase 2
Animals (#)
Session #
C
Phase 3
Animals (#)
Session #
D
Animals (#)
Max session #
E
P(signal)
P(reward)
F
P(lick)
P(reward)
G
Animals
Sessions
H
Phase 1
Phase 2
Low utility
High utility
Bias (s.d.)
Sensitivity (s.d.)
RT (s)
P(reward)
Session #
I
Main task
Bias (s.d.)
Sensitivity (s.d.)
RT (s)
P(reward)
Session #
J
Main task
Bias (s.d.)
Sensitivity (s.d.)
RT (s)
P(reward)
Session #
K
Main task
Δ Bias (high vs. low reward)
Δ Sensitivity (high vs. low reward)
RT (s)
Δ P(reward) (high vs. low reward)
Session #
L
Bias (s.d.)
Sensitivity (s.d.)
RT (s)
P(reward)
Rate (licks / s)

Figure S1. Data selection, sustained-attention value task learning, and behavioral simulations. Related to Figures 1 & 2.

(A) Histogram of experimental session number in learning phase 1.

(B) Histogram of experimental session number in learning phase 2 (with respect to last session number in phase 2).

(C) Histogram of experimental session number in phase 3.

(D) Histogram of the maximum session number per animal.

(E) Fraction of trials containing a signal plotted against reward probability. Every data point is a unique session. We excluded 455 sessions with a reward probability smaller than 0.15, thereby excluding 4 animals.

(F) As E, but for fraction of trials on which the animal licked (responded) on the y-axis.

(G) Histogram of number of sessions per animal (after excluding sessions in panel E,F). We excluded 22 additional animals with fewer than 5 remaining sessions.

(H) Top: Bias (Methods) across experimental sessions in learning phases 1 and 2 (Methods); session numbers are with respect to the last session in phase 2. Rows 2-4: same as top row, but for sensitivity, RT and reward probability (Methods). Shading, 68% confidence interval across animals (N=88, n=4473 sessions). In phase 1, the signal was 6 dB louder than the noise, several free-reward trials (not contingent on response) were administered, and reward magnitude was constant (5 µL) throughout the session. During this phase, mice learned to respond (lick) to harvest rewards: within just three sessions their bias changed from quite conservative (positive bias values) to slightly liberal (negative bias values). Sensitivity also increased across the first few sessions. In phase 2, we introduced the block-based shifts in task utility and trained the mice until they reached a performance threshold (Methods). Sensitivity gradually increased during phase 2.

(I) As H, but for phase 3 (main task). Dashed lines, exponential fit (Methods). Shading, 68% confidence interval across animals (N=88, n=4473 sessions). In phase 3, acoustic signal and noise stimuli were of equal loudness and the small number of classical conditioning trials were removed. By this time, mice were sufficiently experienced with the task to maintain engagement. However, since they could no longer rely on the small loudness difference between signal and noise sounds, their sensitivity dropped to near zero at the beginning of phase 3.

(J) As I, but for the mean across high and low utility blocks. Re-learning the, time-varying and now purely feature-based, signal detection was perceptually difficult, as intended, as indicated by the shallow learning curves for sensitivity and reward probability.

(K) As I, but for the difference between the high and low utility blocks.

(L) From left to right: simulation of bias, sensitivity, RT and reward probability as a function of Poisson lick rate (Methods). Vertical green line indicates the optimal lick rate (maximum reward probability) when lick rate during noise and signal are the same. Note that the optimal average lick rate will differ in conditions when the signal is discriminated from the noise. The max achievable random P(reward), when licking at the optimal random rate, is 0.231.

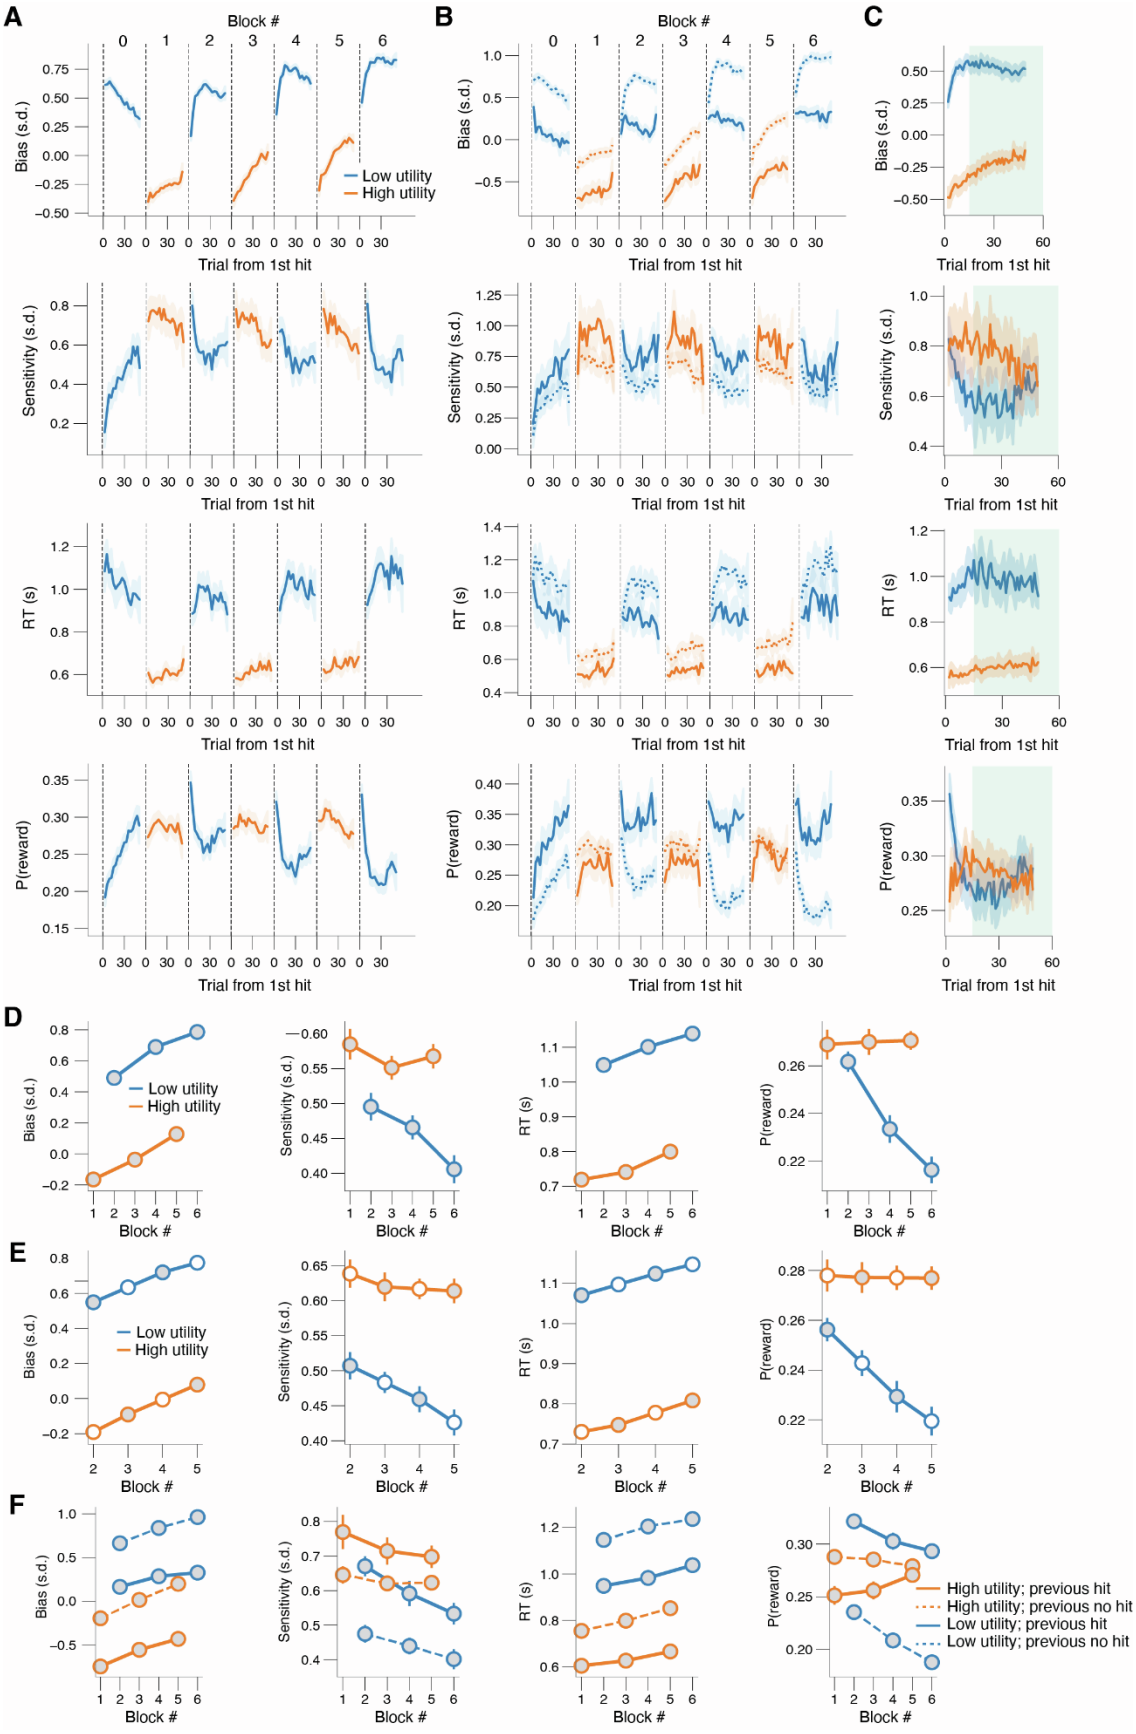
A
B
C
Block #
Bias (s.d.)
Sensitivity (s.d.)
RT (s)
P(reward)
Trial from 1st hit
Low utility
High utility
D
E
F
Block #
High utility; previous hit
High utility; previous no hit
Low utility; previous hit
Low utility; previous no hit

Figure S2. Behavioral control analyses, part 1. Related to Figure 2.

(A) Bias, sensitivity, RT and reward probability across high and low utility blocks within an experimental session, locked to first hit in block.

(B) As A, but separately after previous hits (solid lines) and no-hits (dashed lines).

(C) A, but after collapsing across previous hits and no-hits, and collapsed across blocks of the same task utility.

(D) Bias, sensitivity, RT and reward probability across high and low utility blocks within an experimental session, collapsed across *all* trials within a block. Stats, 2-way repeated measures ANOVA (factors task utility [high vs. low] and time-on-task [early, middle, late]). Main effects task utility are as follows: bias, $F_{1,87} = 747.0$, $p < 0.001$; sensitivity, $F_{1, 87} = 17.7$, $p < 0.001$; RT, $F_{1, 87} = 612.9$, $p < 0.001$; reward probability, $F_{1, 87} = 12.5$, $p = 0.001$.

(E) As D, but when controlling for time-on-task. The white circles are the average of the two adjacent blocks of the same task utility. Main effects task utility are as follows: bias, $F_{1,87} = 498.5$, $p < 0.001$; sensitivity, $F_{1, 87} = 26.8$, $p < 0.001$; RT, $F_{1, 87} = 365.1$, $p < 0.001$; reward probability, $F_{1, 87} = 15.3$, $p < 0.001$.

(F) As D, but when additionally stratifying on previous hit (Methods). Main effects previous hit are as follows: bias, $F_{1,87} = 733.7$, $p < 0.001$; sensitivity, $F_{1, 87} = 34.0$, $p < 0.001$; RT, $F_{1, 87} = 184.2$, $p < 0.001$; reward probability, $F_{1, 87} = 55.1$, $p < 0.001$. Main effects task utility are as follows: bias, $F_{1,87} = 770.2$, $p < 0.001$; sensitivity, $F_{1, 87} = 21.9$, $p < 0.001$; RT, $F_{1, 87} = 667.0$, $p < 0.001$; reward probability, $F_{1, 87} = 1.8$, $p = 0.179$.

All panels: shading or error bars, 68% confidence interval across animals (N=88; n=1983 sessions).

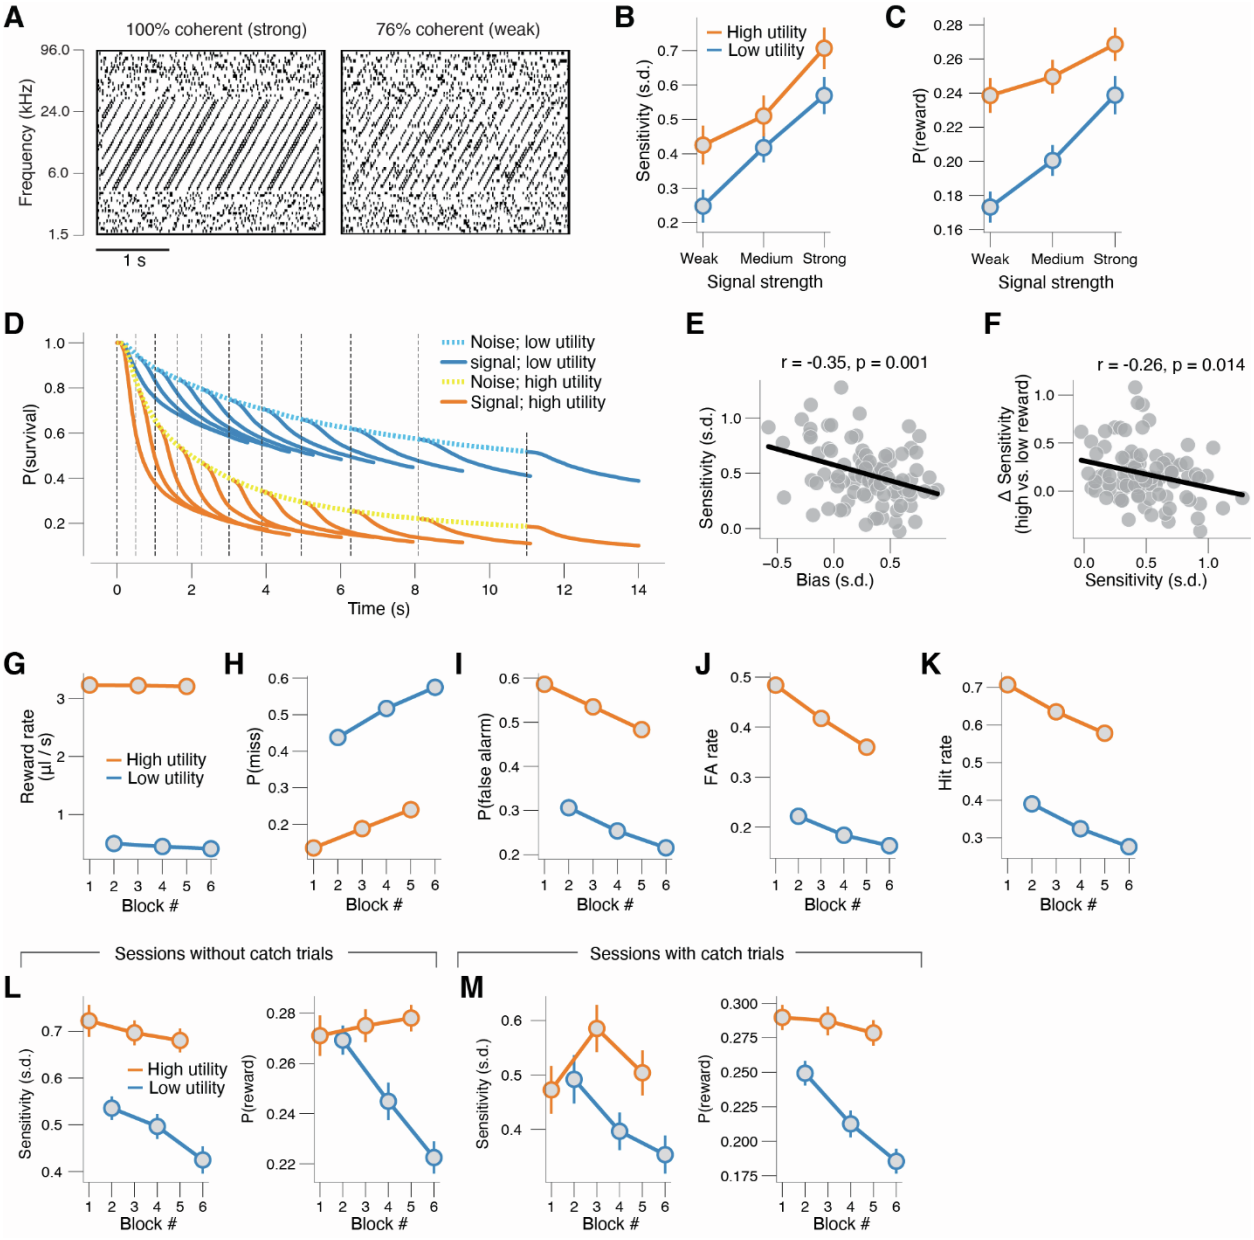
A
100% coherent (strong)
76% coherent (weak)
Frequency (kHz)
96.0
24.0
6.0
1.5
1 s
B
High utility
Low utility
Sensitivity (s.d.)
Weak
Medium
Strong
Signal strength
C
P(reward)
Signal strength
D
Noise; low utility
signal; low utility
Noise; high utility
Signal; high utility
P(survival)
Time (s)
E
r = -0.35, p = 0.001
Sensitivity (s.d.)
Bias (s.d.)
F
r = -0.26, p = 0.014
Δ Sensitivity (high vs. low reward)
Sensitivity (s.d.)
G
Reward rate (µl / s)
High utility
Low utility
Block #
H
P(miss)
Block #
I
P(false alarm)
Block #
J
FA rate
Block #
K
Hit rate
Block #
Sessions without catch trials
Sessions with catch trials
L
Sensitivity (s.d.)
High utility
Low utility
Block #
P(reward)
Block #
M
Sensitivity (s.d.)
Block #
P(reward)
Block #

Figure S3. Behavioral control analyses, part 2. Related to Figures 2 & 3.

Due to the temporal structure of the task, it is possible that mice could learn to match their lick times to the temporal statistics of signal occurrence to harvest at least some reward, rather than attend to the temporal coherence per se. We ruled out such a timing-based strategy, using a separate cohort of mice (N=10; n=142 sessions) in a psychometric variant of the task in which signal coherence was varied randomly from trial to trial, in addition to the block-based utility manipulation. Further supporting feature-based attention, rather than a temporal strategy, empirical sensitivity (panel B) and reward probability (panel C) were substantially higher than those predicted with random licking (Figure S1L) and signal start-time sorted survival functions for hits deviated strongly from the time-matched false-alarm survival functions (panel D).

(A) Example spectrograms of strong (left) and weak (right) signals (Methods). Signal coherence was degraded by manipulating the percentage of tones in each chord that moved coherently through time-frequency space. This manipulation is analogous to reducing spatial motion coherence in visual random-dot motion tasks [105].

(B) Sensitivity across three difficulty bins and separately during high and low task utility. Stats, 2-way repeated measures ANOVA (factors task utility [high vs. low] and signal strength (coherence) bin [weak, medium, strong]); main effect task utility: $F_{1,9} = 4.8$, $p = 0.056$; main effect signal strength: $F_{2,18} = 29.7$, $p < 0.001$; interaction effect: $F_{2,18} = 1.7$, $p = 0.201$. Error bars, 68% confidence interval across animals (N=10, n=142 sessions).

(C) As B, but for reward probability. Main effect task utility: $F_{1,9} = 17.8$, $p = 0.002$; main effect signal strength: $F_{2,18} = 8.0$, $p = 0.003$; interaction effect: $F_{2,18} = 2.0$, $p = 0.162$.

(D) Probability of survival (Kaplan-Meier estimator) stratified into groups with different ranges of noise duration, separately for low and high task utility and separately for noise and signal epochs. The Kaplan-Meier estimator is a non-parametric statistic used to estimate the 'survival' of not-licking. There is a steeper decline of p(survival) during high versus low task utility, corresponding to a more liberal bias. A steeper decline of p(survival) during signal versus noise epochs, corresponds to an overall positive sensitivity (d'). Critically, the signal-related decline in of p(survival) is more pronounced during high versus low task utility, corresponding to a higher sensitivity.

(E) Overall sensitivity (collapsed across task utility) plotted against overall bias (collapsed across task utility). Every data point is a unique session.

(F) Change in sensitivity between utility blocks plotted against overall sensitivity (collapsed across task utility). Every data point is a unique session.

(G) Reward rate (Methods) collapsed across trials within each block. Stats, 2-way repeated measures ANOVA (factors task utility [high vs. low] and time-on-task [1, 2, 3]); main effect task utility: $F_{1,87} = 1519.3$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 1.9$, $p = 0.155$; interaction effect: $F_{2,174} = 0.7$, $p = 0.519$.

(H) As G, but for miss probability (Methods). Main effect task utility: $F_{1,87} = 985.2$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 177.1$, $p < 0.001$; interaction effect: $F_{2,174} = 1.5$, $p = 0.228$.

(I) As G, but for false alarm probability (Methods). Main effect task utility: $F_{1,87} = 714.5$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 188.2$, $p < 0.001$; interaction effect: $F_{2,174} = 4.2$, $p = 0.016$.

(J) As G, but for false alarm rate (Methods). Main effect task utility: $F_{1,87} = 439.0$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 86.6$, $p < 0.001$; interaction effect: $F_{2,174} = 17.9$, $p < 0.001$.

(K) As G, but for hit rate (Methods). Main effect task utility: $F_{1,87} = 538.1$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 139.4$, $p < 0.001$; interaction effect: $F_{2,174} = 0.7$, $p = 0.512$.

(L) As main Figure 2D,H, but only for 82.8% of sessions without catch trials (Methods). Main effects task utility are as follows: sensitivity, $F_{1,69} = 31.9$, $p < 0.001$; reward probability, $F_{1,69} = 6.2$, $p = 0.015$.

(M) As main Figure 2D,H, but only for 17.2% of sessions with catch trials (Methods). Main effects task utility are as follows: sensitivity, $F_{1,43} = 6.6$, $p = 0.013$; reward probability, $F_{1,43} = 29.0$, $p < 0.001$.

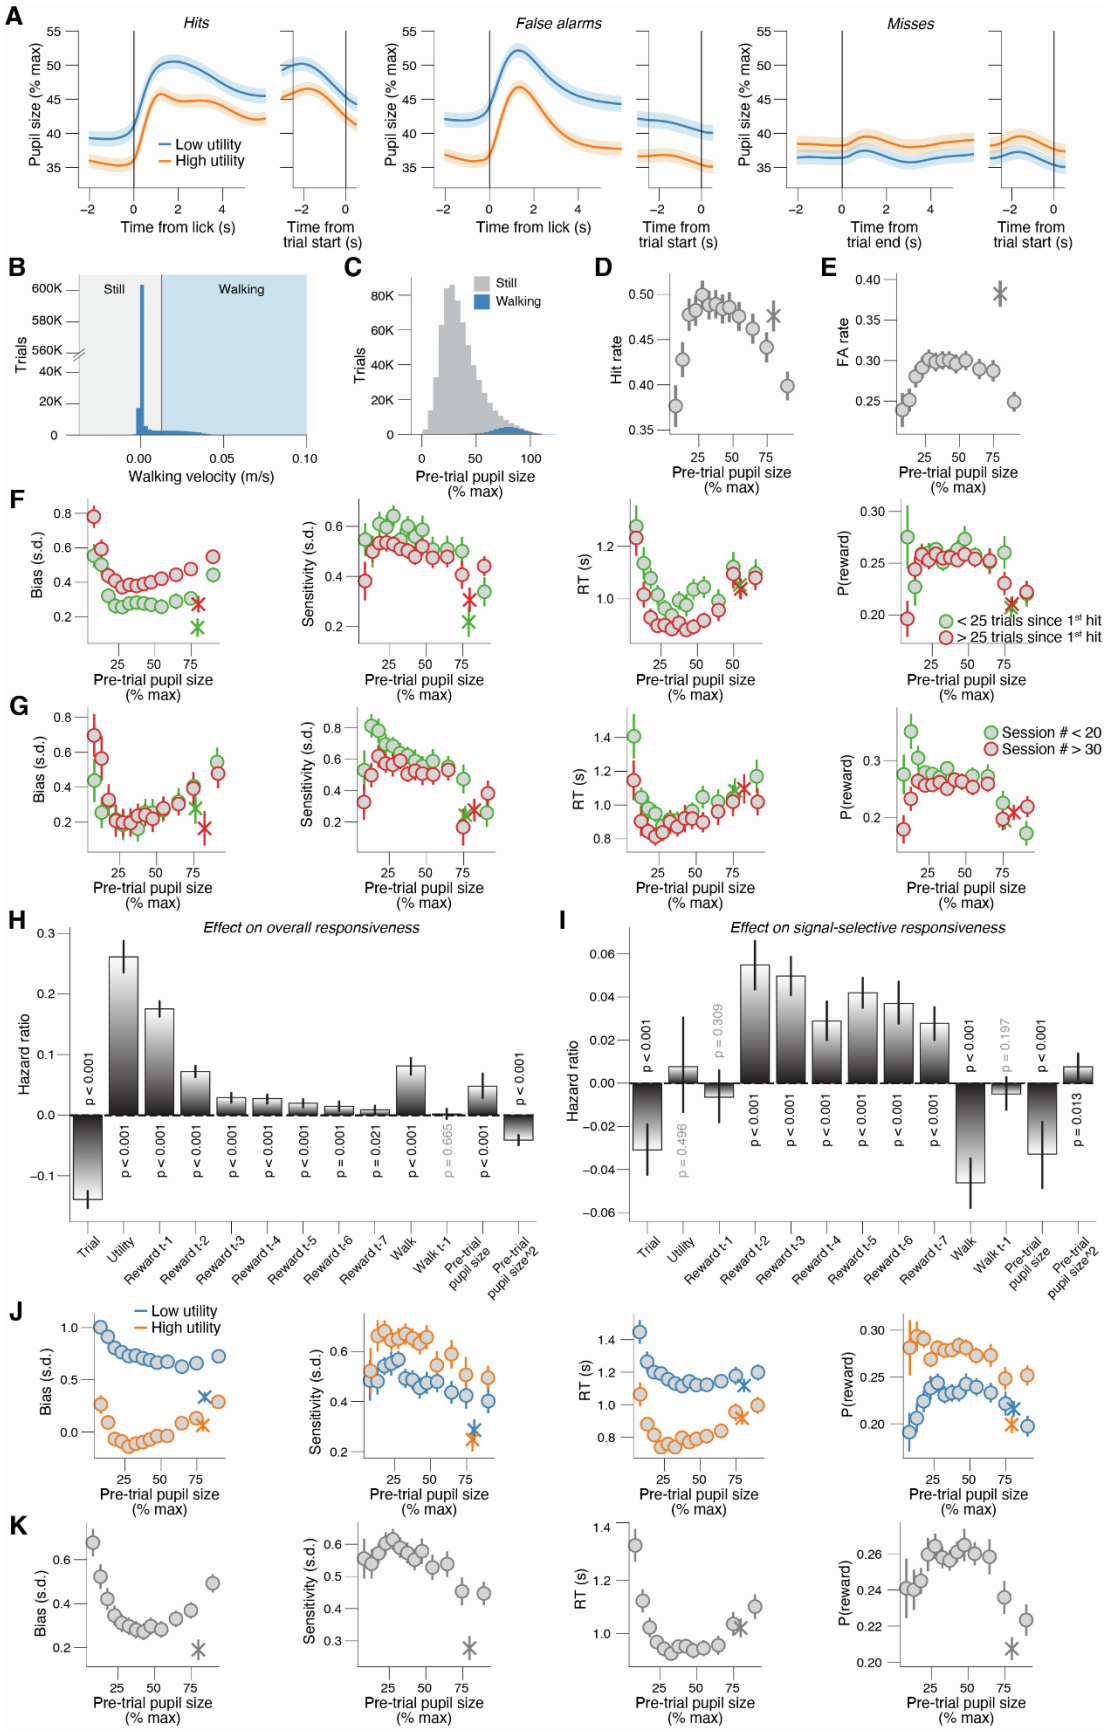
A
Hits
False alarms
Misses
Pupil size (% max)
Low utility
High utility
Time from lick (s)
Time from trial start (s)
Time from trial end (s)
B
Still
Walking
Trials
Walking velocity (m/s)
C
Pre-trial pupil size (% max)
D
Hit rate
E
FA rate
F
Bias (s.d.)
Sensitivity (s.d.)
RT (s)
P(reward)
< 25 trials since 1st hit
> 25 trials since 1st hit
G
Session # < 20
Session # > 30
H
Effect on overall responsiveness
Hazard ratio
Trial
Utility
Reward t-1
Reward t-2
Reward t-3
Reward t-4
Reward t-5
Reward t-6
Reward t-7
Walk
Walk t-1
Pre-trial pupil size
Pre-trial pupil size^2
I
Effect on signal-selective responsiveness
J
K

Figure S4. Optimal performance occurs at intermediate levels of pupil-linked arousal. Related to Figure 4.

(A) Left: Pupil size across time on hit trials, separately for high and low task utility. Left, locked to licks; right, locked to next trial's onset. Middle, as left, but on false alarm trials. Right: As left, but for miss trials. Left, locked to trial offset (end of 3-s signal sound).

(B) Histogram of pre-trial walking velocity (across all animals and experimental sessions; Methods). Red line, cutoff for defining walking.

(C) Histogram of pre-trial pupil size (across all animals and experimental sessions), separately for still and walking trials.

(D) Relationship between pre-trial pupil size and hit rate (irrespective of task utility; Methods). A 1st order (linear) fit was not superior to a constant fit ($F_{1,12} = 0.1$, $p = 0.717$) and a 2nd order (quadratic) fit was superior to the 1st order fit ($F_{1,12} = 6.5$, $p = 0.026$; sequential polynomial regression; Methods). Asterisk, walking trials (Methods).

(E) As D, but for false alarm rate. 1st order fit: $F_{1,12} = 0.2$, $p = 0.690$; 2nd order fit: $F_{1,12} = 7.9$, $p = 0.016$.

(F) As Figure main 4A-D, but separately for fewer or more than 25 trials after the first hit in each block.

(G) As main Figure 4A-D, but separately for the first 20 sessions of the main task, or for session numbers greater than 30. In this analysis, we only considered the first 50 sessions of 34 animals with at least 50 sessions worth of data in the main task.

(H,I) As main Figure 4F,G, but for a model with one additional predictor: previous walk.

(J) As main Figure 4A-D, but separately for high and low task utility.

(K) As main Figure 4A-D, but for pre-trial pupil size measures without having regressed out effects of time-on-task and previous hit (Methods). Stats (sequential polynomial regression; Methods) are as follows. Bias: 1st order fit: $F_{1,12} = 13.2$, $p = 0.003$; 2nd order fit: $F_{1,12} = 8.0$, $p = 0.016$. Sensitivity: 1st order fit: $F_{1,12} = 4.4$, $p = 0.058$; 2nd order fit: $F_{1,12} = 3.0$, $p = 0.111$. RT: 1st order fit: $F_{1,12} = 5.1$, $p = 0.043$; 2nd order fit: $F_{1,12} = 7.5$, $p = 0.018$. Reward probability: 1st order fit: $F_{1,12} = 1.3$, $p = 0.268$; 2nd order fit: $F_{1,12} = 7.3$, $p = 0.019$.

All panels: shading or error bars, 68% confidence interval across animals (N=88, n=1983 sessions).

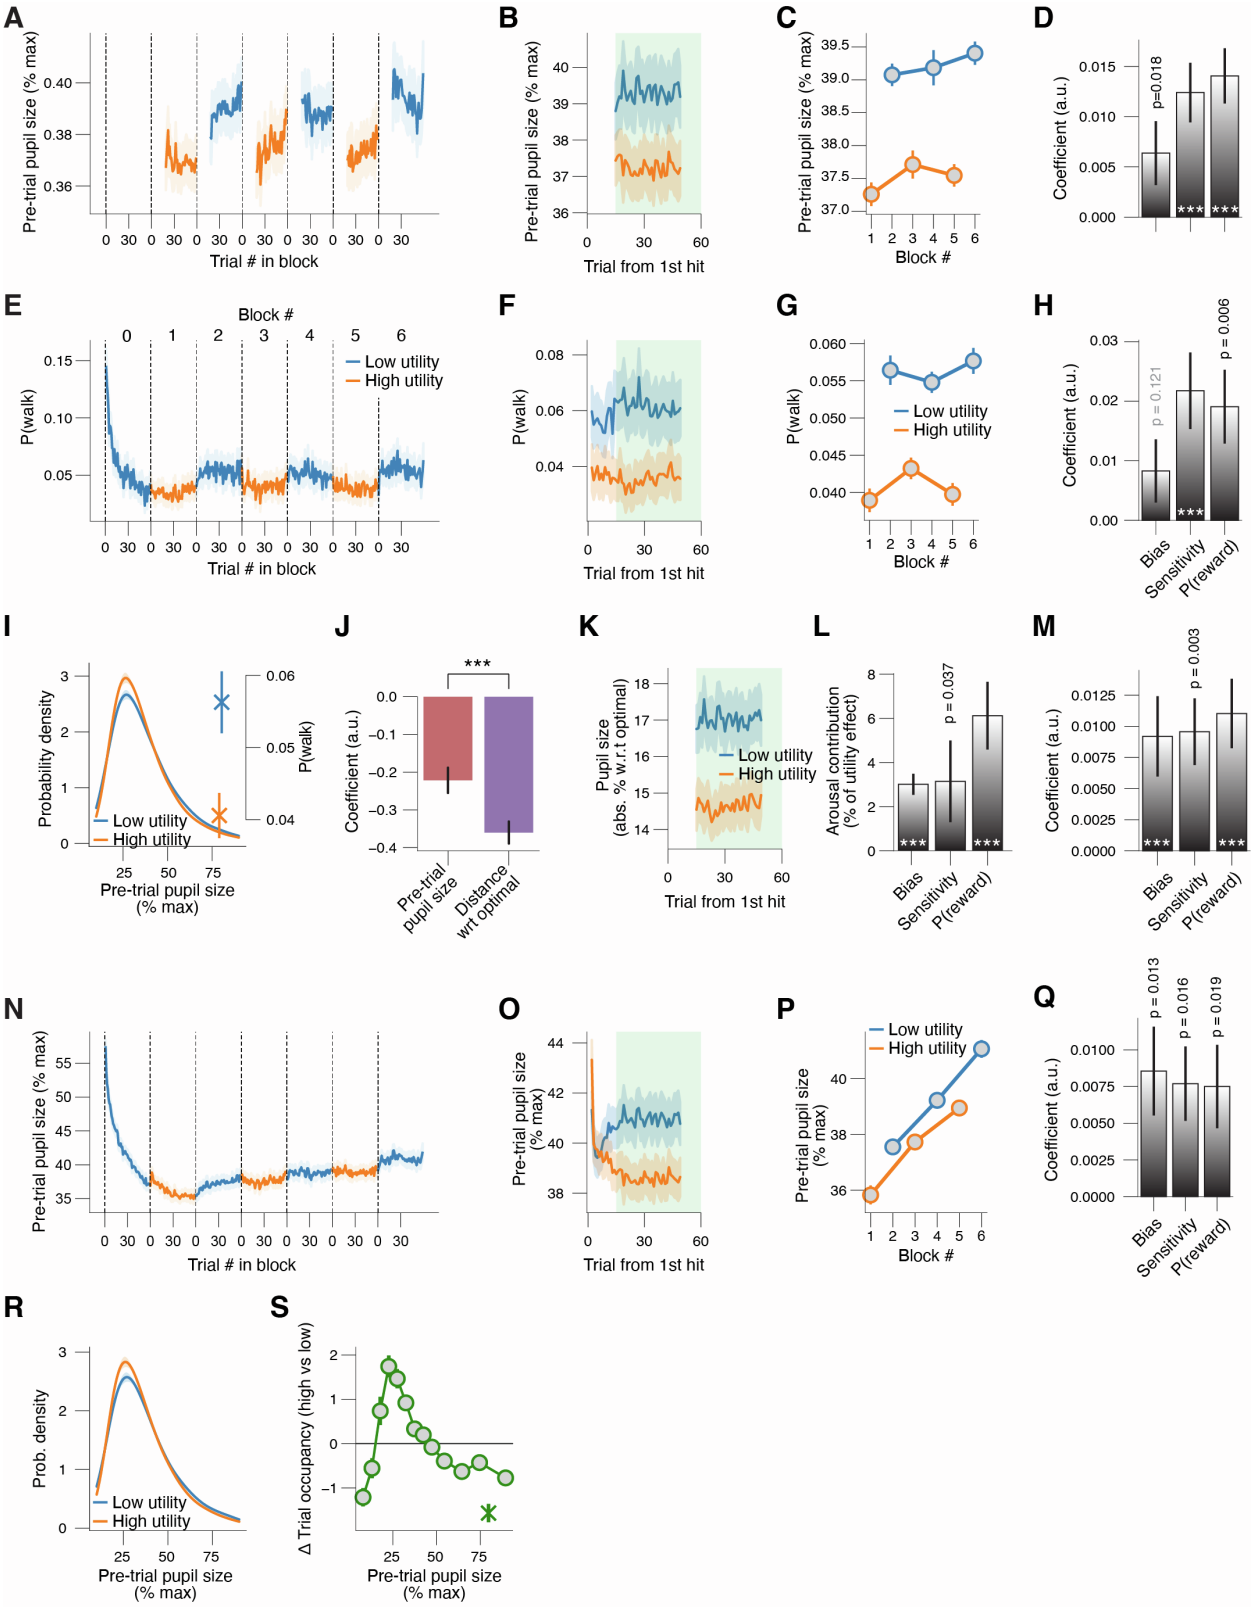
A
Pre-trial pupil size (% max)
Trial # in block
B
Pre-trial pupil size (% max)
Trial from 1st hit
C
Pre-trial pupil size (% max)
Block #
D
Coefficient (a.u.)
p=0.018
E
Block #
P(walk)
Low utility
High utility
Trial # in block
F
P(walk)
Trial from 1st hit
G
P(walk)
Low utility
High utility
Block #
H
Coefficient (a.u.)
p = 0.121
p = 0.006
Bias
Sensitivity
P(reward)
I
Probability density
P(walk)
Low utility
High utility
Pre-trial pupil size (% max)
J
Coefficient (a.u.)
Pre-trial pupil size
Distance wrt optimal
K
Pupil size (abs. % w.r.t optimal)
Low utility
High utility
Trial from 1st hit
L
Arousal contribution (% of utility effect)
p = 0.037
Bias
Sensitivity
P(reward)
M
Coefficient (a.u.)
p = 0.003
Bias
Sensitivity
P(reward)
N
Pre-trial pupil size (% max)
Trial # in block
O
Pre-trial pupil size (% max)
Trial from 1st hit
P
Pre-trial pupil size (% max)
Low utility
High utility
Block #
Q
Coefficient (a.u.)
p = 0.013
p = 0.016
p = 0.019
Bias
Sensitivity
P(reward)
R
Prob. density
Low utility
High utility
Pre-trial pupil size (% max)
S
Δ Trial occupancy (high vs low)
Pre-trial pupil size (% max)

Figure S5. Mice regulate their pupil-linked arousal towards optimality when task utility is high. Related to Figure 5.

(A) Pre-trial pupil size (time-on-task and previous hit regressed out; Methods) across low and high utility blocks within an experimental session.

(B) As A, but locked to first hit in block and after averaging over previous hits and no-hits (same logic as in Figure S2C).

(C) As A, but collapsed across trials within a block. Stats, 2-way repeated measures ANOVA (factors task utility [high vs. low] and time-on-task [early, middle, late]); main effect task utility: $F_{1,87} = 35.5$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 1.7$, $p = 0.187$; interaction effect: $F_{2,174} = 1.0$, $p = 0.376$.

(D) As main Figure 5F, but for pre-trial pupil size as mediator.

(E-H) As A-D, but for walk probability. Main effect task utility: $F_{1,87} = 58.5$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 0.3$, $p = 0.710$; interaction effect: $F_{2,174} = 3.3$, $p = 0.039$.

(I) Probability density function of pupil-linked arousal states, and walk probability, separately for high and low task utility. The difference between high and low task utility is reported in main Figure 5C.

(J) Coefficients from logistic regression of block-wise task utility [high vs. low] on pre-trial pupil size or on absolute distance from optimal pre-trial pupil size (Methods).

(K) As B, distance of pupil size from the optimal level (Methods).

(L) As main Figure 5D, but without modeling walking trials in a separate state bin.

(M) As main Figure 5F, but only for trials mice did not walk.

(N-Q) As A-D but for pre-trial pupil size without having regressed out time-on-task and previous hit (Methods). Main effect task utility: $F_{1,87} = 43.4$, $p < 0.001$; main effect time-on-task: $F_{2,174} = 37.1$, $p < 0.001$; interaction effect: $F_{2,174} = 3.3$, $p = 0.282$.

(R) As I, but for pre-trial pupil size without having regressed out time-on-task and previous hit (Methods).

(S) As main Figure 5C, but for pre-trial pupil size without having regressed out time-on-task and previous hit (Methods).

All panels: shading or error bars, 68% confidence interval across animals (N=88, n=1983 sessions). Panels D,H,L,M,Q: stats, Wilcoxon signed-rank test; ***, $p < 0.001$.

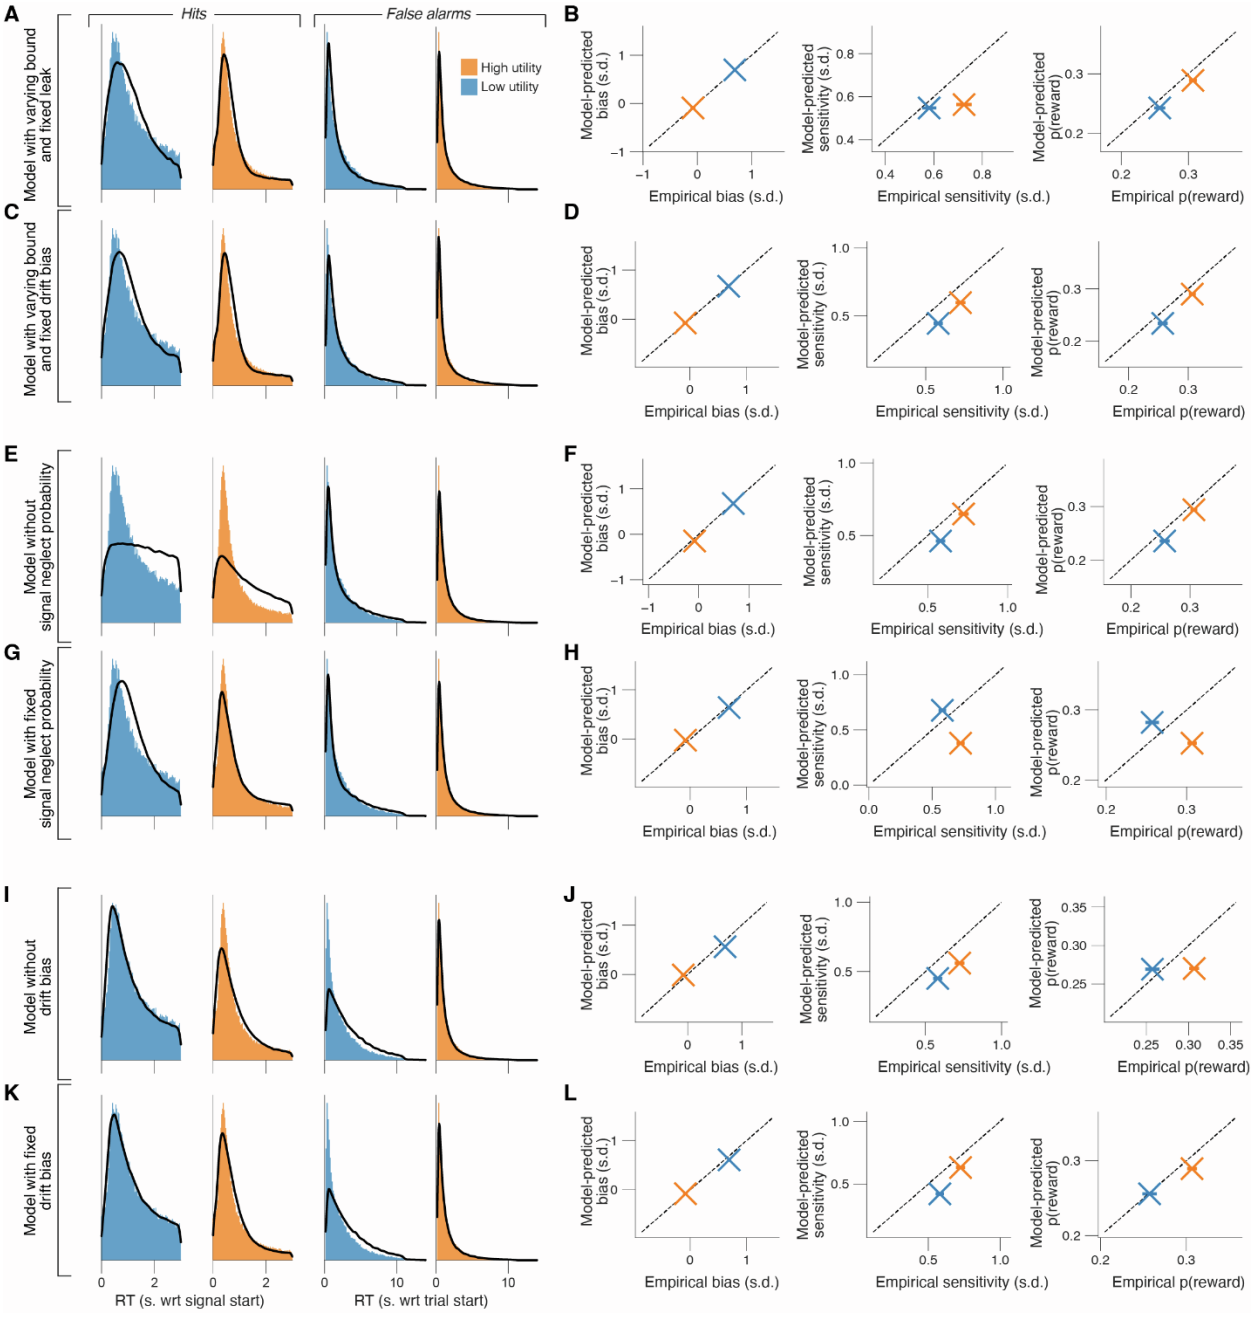
A
Hits
False alarms
High utility
Low utility
Model with varying bound and fixed leak
B
Model-predicted bias (s.d.)
Empirical bias (s.d.)
Model-predicted sensitivity (s.d.)
Empirical sensitivity (s.d.)
Model-predicted p(reward)
Empirical p(reward)
C
Model with varying bound and fixed drift bias
D
E
Model without signal neglect probability
F
G
Model with fixed signal neglect probability
H
I
Model without drift bias
J
K
Model with fixed drift bias
L
RT (s. wrt signal start)
RT (s. wrt trial start)

Figure S6. Accumulation-to-bound model fits of alternative models. Related to Figure 6.

(A) RT distribution for correct responses (hits; left) and incorrect responses (false alarms; right) in the low and high utility blocks. Black line, fit of Model 1 (with varying bound height and fixed leak; Methods).

(B) Model-predicted bias (left), sensitivity (middle) and reward probability (right) in the low and high utility blocks plotted against the empirical estimates. Dashed line, identity line.

(C,D) As A,B, but for Model 2 (with varying bound height and fixed drift bias; Methods).

(E,F) As A,B, but for Model 3 (without signal neglect probability; Methods).

(G,H) As A,B, but for Model 4 (with fixed signal neglect probability; Methods).

(I,J) As A,B, but for Model 5 (without drift bias; Methods).

(K,L) As A,B, but for Model 6 (with fixed drift bias; Methods).

All panels: pooled data across animals (N=88) and sessions (n=1983).

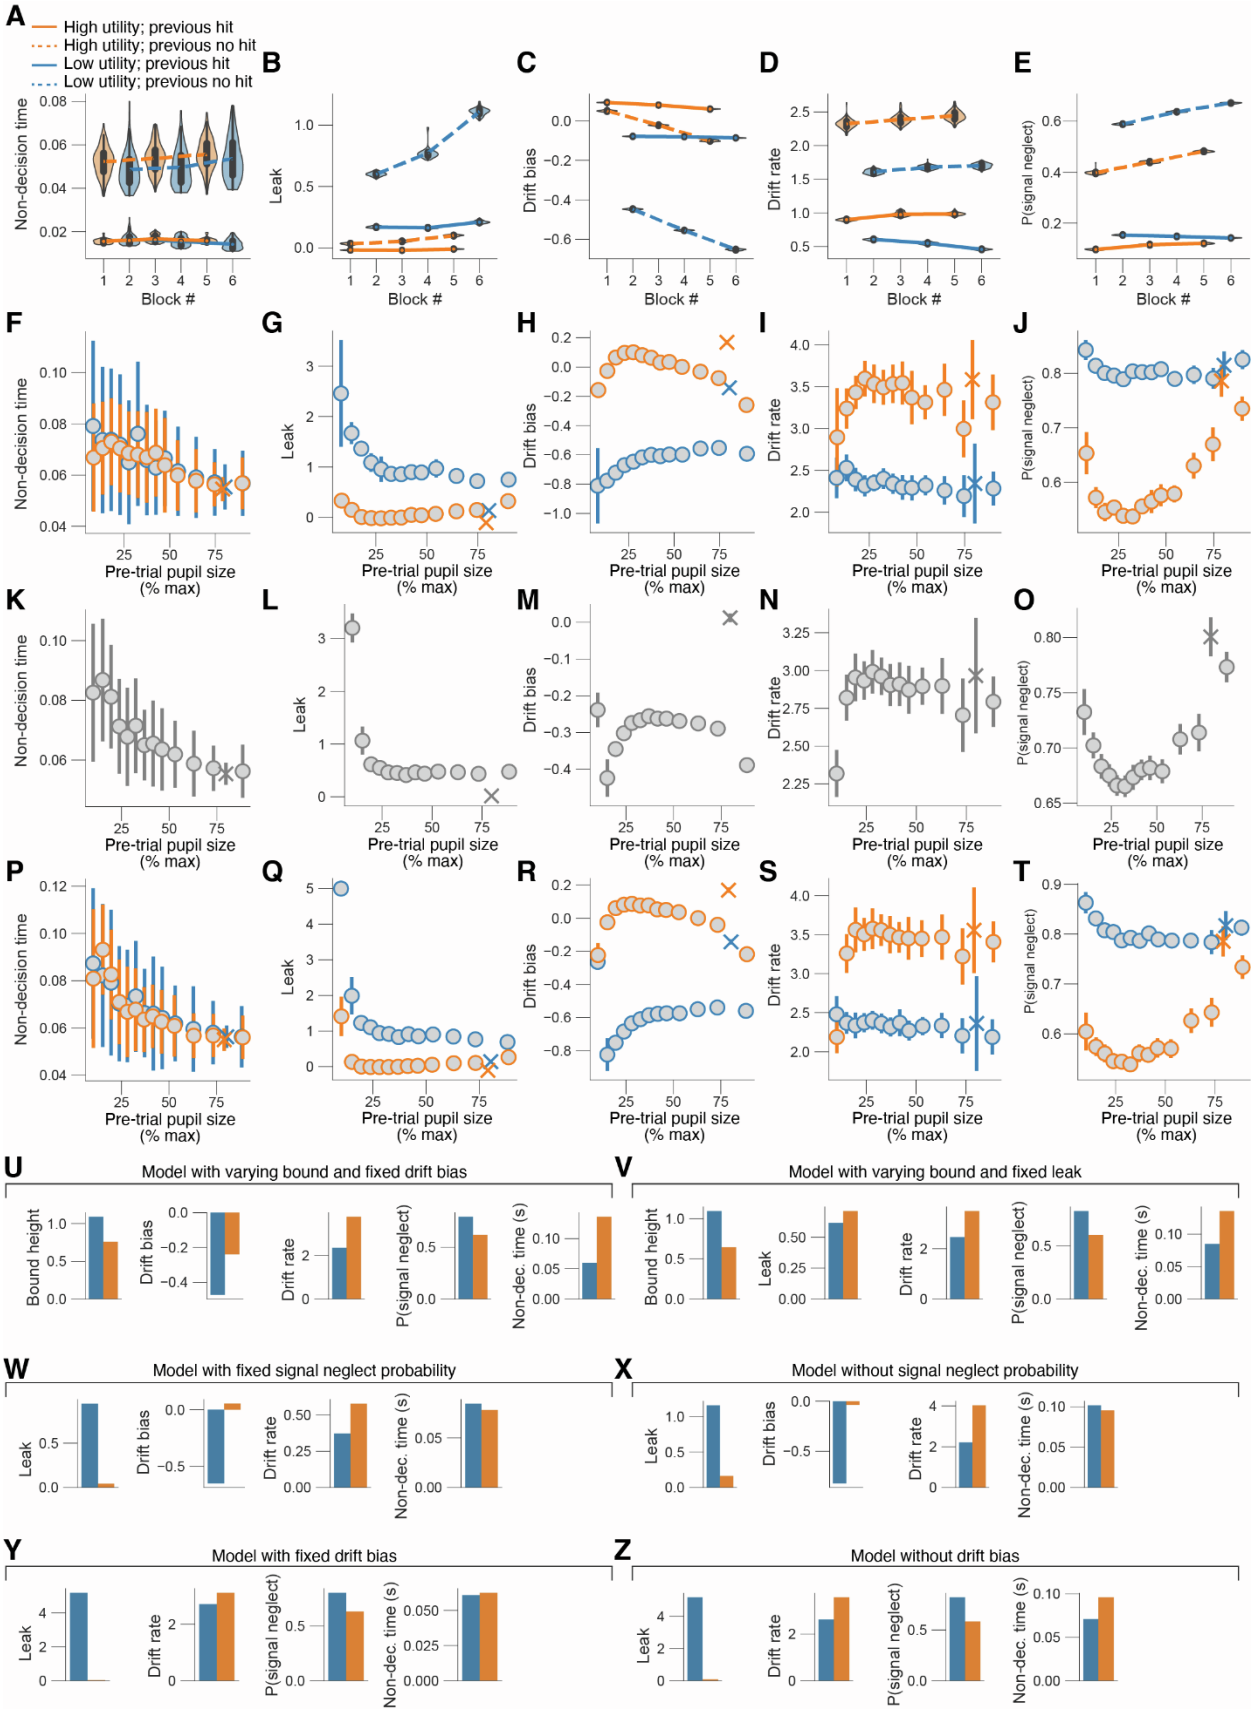
A
High utility; previous hit
High utility; previous no hit
Low utility; previous hit
Low utility; previous no hit
Non-decision time
Block #
B
Leak
C
Drift bias
D
Drift rate
E
P(signal neglect)
F
G
H
I
J
K
L
M
N
O
P
Q
R
S
T
Pre-trial pupil size (% max)
U
Model with varying bound and fixed drift bias
Bound height
Drift bias
Drift rate
P(signal neglect)
Non-dec. time (s)
V
Model with varying bound and fixed leak
W
Model with fixed signal neglect probability
X
Model without signal neglect probability
Y
Model with fixed drift bias
Z
Model without drift bias

Figure S7. Task utility and pupil-linked arousal impact intersecting aspects of the decision computation. Related to Figure 7.

(A) Fitted non-decision time estimates (kernel density estimate of 100 bootstrapped replicates) separately per block number. Main effect task utility (fraction of bootstrapped parameter estimates in the low utility blocks higher than in the high utility blocks): $p = 0.32$. Main effect time-on-task (fraction of bootstrapped parameter estimates in the first two blocks higher than in the last two blocks): $p = 0.41$.

(B) As A, but for leak. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p < 0.01$

(C) As A, but for drift bias. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p < 0.01$.

(D) As A, but for drift rate. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p = 0.46$.

(E) As A, but for signal neglect probability. Main effect task utility: $p < 0.01$. Main effect time-on-task: $p < 0.01$.

(F-J) As main Figure 7F-J, but separately for high and low task utility.

(K) Fitted non-decision estimates (100 bootstrapped replicates) separately per arousal state (same pupil size defined bins as in Figure 4A-D, but without having regressed out effects of time-on-task and previous hit; irrespective of task utility; Methods). A 1st order (linear) fit was superior to a constant fit ($F_{1,12} = 118.8$, $p < 0.001$) and a 2nd order (quadratic) fit was not superior to the 1st order fit ($F_{1,12} = 2.3$, $p = 0.154$; sequential polynomial regression; Methods). Asterisk, walking trials (Methods).

(L) As K, but for leak. 1st order fit: $F_{1,12} = 5.9$, $p = 0.032$; 2nd order fit: $F_{1,12} = 3.3$, $p = 0.093$.

(M) As K, but for drift bias. 1st order fit: $F_{1,12} \sim 0.0$, $p = 0.870$; 2nd order fit: $F_{1,12} = 1.3$, $p = 0.269$.

(N) As K, but for drift rate. 1st order fit: $F_{1,12} = 0.6$, $p = 0.469$; 2nd order fit: $F_{1,12} = 3.8$, $p = 0.076$.

(O) As K, but for signal neglect probability. 1st order fit: $F_{1,12} = 1.2$, $p = 0.290$; 2nd order fit: $F_{1,12} = 7.9$, $p = 0.016$.

(P-T) As 7F-J, but separately for high and low task utility.

(U-Z) Parameter estimates from alternative models (Methods), separately for high and low task utility.

All panels: pooled data across animals (N=88) and sessions (n=1983).