Regulation of evidence accumulation by pupil-linked arousal processes

Waitsang Keung*[1], Todd Hagen*[1], and Robert C. Wilson[1,2]

[1]Department of Psychology, University of Arizona;

[2]Cognitive Science Program, University of Arizona,

*equal contribution

Abstract

Integrating evidence over time is crucial for effective decision making. For simple perceptual decisions, a large body of work suggests that humans and animals are capable of integrating evidence over time fairly well, but that their performance is far from optimal. This suboptimality is thought to arise from a number of different sources including (1) noise in sensory and motor systems, (2) unequal weighting of evidence over time, (3) order effects from previous trials and (4) irrational side biases for one choice over another. In this work we investigated whether and how these different sources of suboptimality are related to pupil dilation, a putative correlate of norepinephrine tone. In particular, we measured pupil response in humans making a series of decisions based on rapidly-presented auditory information in an evidence accumulation task. We found that people exhibited all four types of suboptimality, but that only noise and the uneven weighting of evidence over time, the 'integration kernel', were related to the change in pupil response during the stimulus. Moreover, these two different suboptimalities were related to different aspects of the pupil signal, with the individual differences in pupil response associated with individual differences in integration kernel, while trial-by-trial fluctuations in pupil response were associated with trial-by-trial fluctuations in noise. These results suggest that different sources of suboptimality in human perceptual decision making are related to distinct pupil-linked processes possibly related to tonic and phasic norepinephrine activity.

Regulation of evidence accumulation by pupil-linked arousal processes

Introduction

The ability to integrate evidence over time is a crucial component of perceptual decision making. This is true whether we are integrating visual information from saccade to saccade as we scan a scene, or integrating auditory information from word to word as we listen to someone talk. In recent years much work has been devoted to understanding how humans and animals perform evidence integration over short time scales (one the order of one second) in simple perceptual tasks (Brunton, Botvinick, & Brody, 2013; Erlich, Brunton, Duan, Hanks, & Brody, 2015; Katz, Yates, Pillow, & Huk, 2016; Yates, Park, Katz, Pillow, & Huk, 2017). In a classic paradigm from this literature, known as the Random Dot Motion Task, participants are presented with a movie of randomly moving dots that have a weak tendency to drift in a particular direction (e.g. left or right) and they must decide which way the dots are drifting (Newsome & Pare, 1988). The optimal strategy in this task is to count, i.e. integrate, the number of dots moving to the left and right over the time course of the stimulus and keep track of the difference between them over time. Amazingly, this optimal strategy can account for many of the qualitative properties of human and animal behavior and neural correlates of integrated evidence can be found in several areas of the brain (Erlich et al., 2015; Gold & Shadlen, 2001; Hanks et al., 2015; Katz et al., 2016).

Despite the ability of the optimal model to qualitatively account for a number of experimental findings, the quantitative performance of even highly trained humans and animals is suboptimal. This suboptimality is thought to arise from at least four different sources: (1) neuronal noise, (2) unequal weighting of evidence over time, (3) order effects from previous trials and (4) side biases.

The first source of suboptimality is neuronal noise. While the exact cause of neuronal noise is subject to some debate (Beck, Ma, Pitkow, Latham, & Pouget, 2012; Drugowitsch, Wyart, Devauchelle, & Koechlin, 2016; Faisal, Selen, & Wolpert, 2008; Ma, Beck, Latham, & Pouget, 2006), it is thought that variability in neural firing noise impacts perceptual decision making in one of two ways. First, noisy sensory information

reduces the quality of the evidence going into the accumulator in the first place (Brunton et al., 2013; Kaufman & Churchland, 2013; Osborne, Lisberger, & Bialek, 2005; Smith & Ratcliff, 2004). Second, noisy action selection causes mistakes to be made even after the integration process is complete (Daw, O'doherty, Dayan, Seymour, & Dolan, 2006; Griffiths & Tenenbaum, 2006; Sutton & Barto, 1998).

The second source of suboptimality comes from the unequal weighting of evidence over time, which we call here the 'integration kernel'. In particular, while the optimal kernel in most perceptual decision making tasks is flat - i.e. all information is weighed equally over time - a number of studies have shown that humans and animals can have quite suboptimal kernels. For example, in the Random Dot Motion Task monkeys exhibit a "primacy" kernel, putting more weight on early part of the stimulus relative to the later part of the stimulus (Yates et al., 2017). Conversely, in a slightly different integration task, humans exhibit the opposite 'recency' kernel, weighing later information more than early information (Wyart, Myers, & Summerfield, 2015). Finally, in some experiments this second source of suboptimality is absent, with a 'flat' integration kernel being found in both rats and highly trained humans (Brunton et al., 2013).

The third source of suboptimality reflects the tendency to let previous decisions and outcomes interfere with the present choice. Thus, when making multiple perceptual decisions, the current decision is influenced by the choice we just made, for example by repeating an action when it was rewarded and choosing something else when it was not, a reinforcement learning effect (Abrahamyan, Silva, Dakin, Carandini, & Gardner, 2016; Daw et al., 2006; Drugowitsch et al., 2016; Sutton & Barto, 1998), or simply repeating a choice regardless of the outcome associated with it, a choice kernel effect(Abrahamyan et al., 2016; Drugowitsch et al., 2016). Such sequential dependence can be advantageous when there are temporal correlations between trials, as is the case in many reinforcement learning tasks (Daw et al., 2006; Sutton & Barto, 1998), but is suboptimal in most perceptual decision making tasks when each trial is independent of the past (Abrahamyan et al., 2016; Akrami, Kopec, Diamond, & Brody, 2018;

Barraclough, Conroy, & Lee, 2004; Drugowitsch et al., 2016).

Finally, the fourth suboptimality is an overall side bias where both humans and animals develop a preference for one option (e.g. left) even though that leads to more errors overall (Abrahamyan et al., 2016).

Evidence from a number of authors suggests that pupil-linked arousal processes, putatively driven by the locus coerulues norepinephrine system (Aston-Jones & Cohen, 2005; Eldar, Cohen, & Niv, 2013; Joshi, Li, Kalwani, & Gold, 2016; Rajkowski, 1993; Reimer et al., 2016), are well placed to modulate these four different sources of suboptimality. First, with regard to noise, changes in pupil response have been shown as indicators of various cognitive states and processes such as effort, arousal, mood, attention and memory (Aston-Jones & Cohen, 2005; Sara, 2009). Specifically, previous work on pupil and perceptual decision making suggests a role for pupil-linked arousal systems to modulate the overall neuronal noise, i.e. the signal-to-noise ratio of sensory cues (Cavanagh, Wiecki, Kochar, & Frank, 2014; Murphy, Vandekerckhove, & Nieuwenhuis, 2014).

Second, with regard to kernel and side bias, pupil response changes have also been used as indicators of norepinephrine processes that change the 'gain' of other neural systems, and thus modulate the strength of internal and external cognitive biases on decision making (Aston-Jones & Cohen, 2005; Eldar et al., 2013; Eldar, Felso, Cohen, & Niv, 2018; Mather, Clewett, Sakaki, & Harley, 2016). In addition, recent work has also suggested that norepinephrine modulates the urgency of making a decision in a sequential sampling task, that is, the higher the norepinephrine level, the more urgently a decision is made (Hauser, Moutoussis, Purg, Dayan, & Dolan, 2018). Taken together these studies point to the possibility of pupil-linked norepinephrine systems to modulate the integration kernel by changing the strength of bias during the integration process.

Third, with regard to sequential effects, pupil changes have been related to how humans integrate relevant information from previous trials to infer uncertainty and expectation, suggesting a role for pupil-linked arousal systems in modulating sequential effects (Krishnamurthy, Nassar, Sarode, & Gold, 2017; Nassar et al., 2012).

In this work we examined whether and how pupil-linked arousal indices influenced all of these sources of suboptimalities in a simple perceptual decision making task. We found that, within subject, trial-by-trial pupil response was associated with trial-by-trial changes to the signal-to-noise ratio. Conversely, across subjects, individual differences in pupil response were associated with individual differences in integration kernel. Neither sequential effects nor side bias was associated with any pupil measure. This dissociation suggests that different sources of suboptimality arise from distinct neural processes and that signal-to-noise ratio and integration kernel are regulated by different parts of the arousal system.

Results

Bernoulli Clicks Task

To study the effects of pupil response on evidence accumulation, we designed an auditory discrimination task based on the Poisson Clicks Task in Brunton et al. (2013). In this task, participants listened to two trains of clicks in the left and right ears, and were instructed to indicate which side they thought had more clicks (Figure 1A). Clicks in our task were generated according to a Bernoulli process, such that there was always a click every 50ms that was on the left with probability p_{left}, otherwise it was on the right. This process meant that the total number of clicks was always fixed at 20 clicks and the clicks occurred at a fixed frequency of 20 Hz. This generative process for the clicks represents a slight departure from (Brunton et al., 2013), in which clicks were generated by a Poisson process with a refractory period of 20ms. To indicate this difference we refer to our task as the Bernoulli Clicks Task.

Psychometric and chronometric curves

108 participants each performed between 666 and 938 trials (mean 760.7) of the Bernoulli Clicks Task. Basic behavior was consistent with behaviors in similar pulsed-accumulation tasks (Brunton et al., 2013; Yates et al., 2017). Choice exhibited a sigmoidal dependence on the net difference in evidence strength, i.e. the difference in

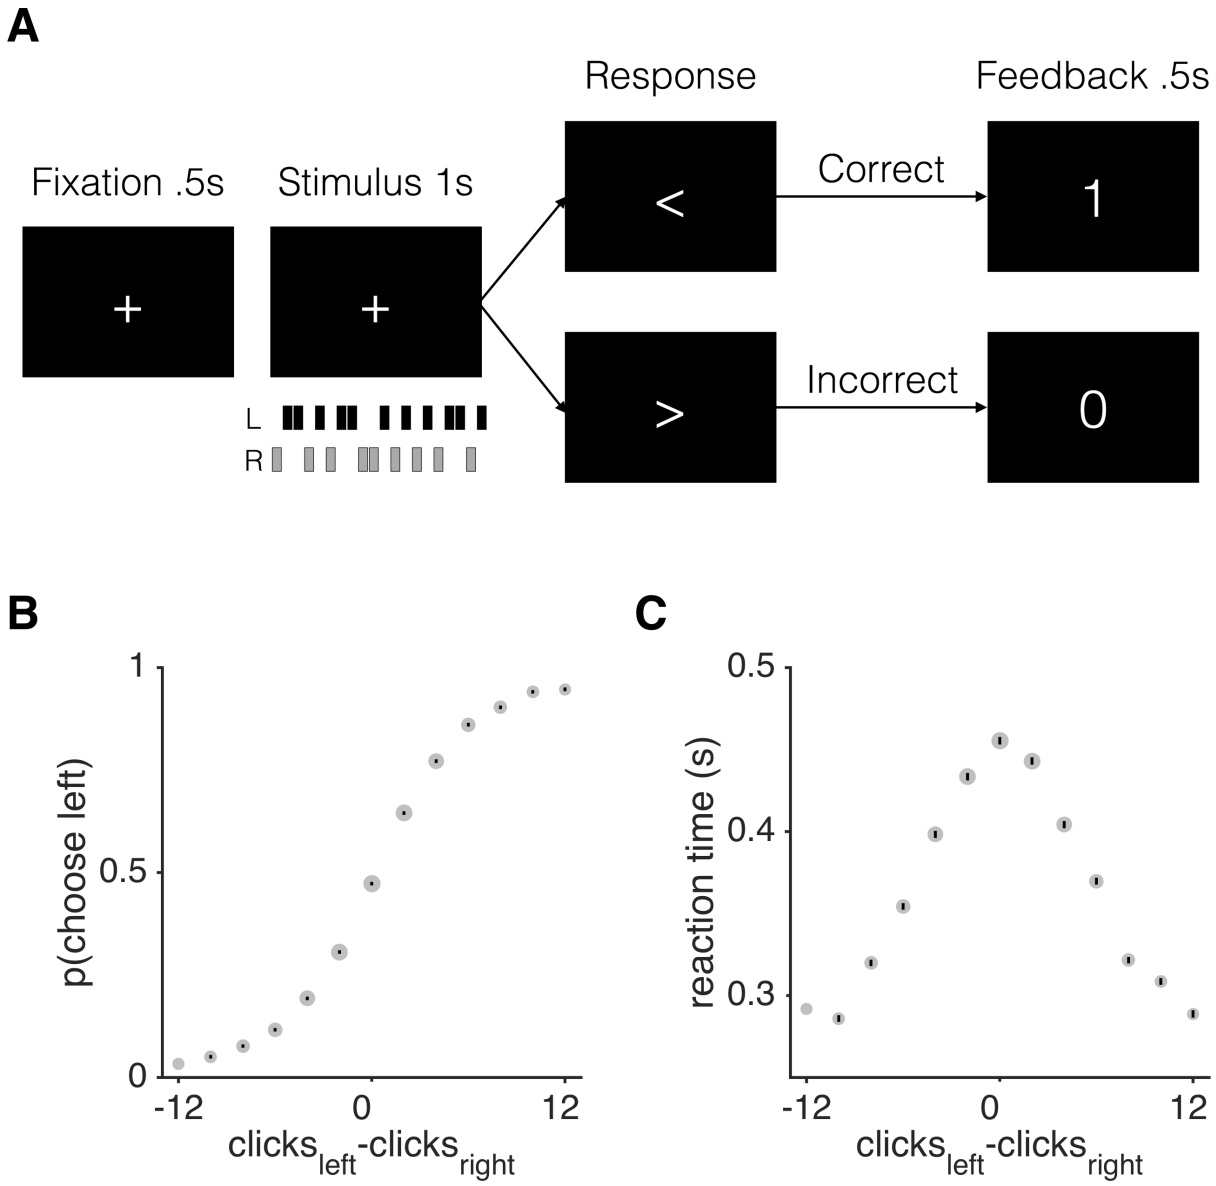

Figure 1. Basic behavior of 108 participants. (A) Participants listened to a train of twenty clicks coming in either the left (L, black bars) or right (R, grey bars) ear for one second, and decided which side had more clicks. (B) Choice probability (probability of choosing left) showed sigmoidal relationship with difficulty (the difference in number of clicks between left and right). Size of grey dots scaled by number of trials. Black bars indicate s.e.m. (C) Reaction times were higher on more difficult trials.

number of clicks between the right and left (Figure 1B). Reaction times were also modulated by net evidence strength (Figure 1C) such that participants were faster and more accurate when the difference of number of clicks was large (easy trials), and less accurate and slower when that difference was small (hard trials).

Humans exhibit all four suboptimalities in the Bernoulli Clicks Task

We used a logistic regression model to characterize the four different types of suboptimality in human decision making in our task. This model quantifies the impact of each click, the reinforcement learning and choice kernel effects from the five previous trials and the side bias on participants' choices. In particular, we assume that the probability of choosing left on trial t is given by

$$\text{logit}(p_{\text{left}} \text{ at trial } t) = \underbrace{\sum_{i=1}^{20} \beta_i^{\text{click}} \text{c}_i}_{\substack{\text{intergration}\\ \text{kernel}}} + \underbrace{\sum_{j=1}^{5} \beta_j^{\text{RL}} \text{a}_{t-j} \text{r}_{t-j}}_{\substack{\text{reinforcement}\\ \text{learning}}} + \underbrace{\sum_{j=1}^{5} \beta_j^{\text{CK}} \text{a}_{t-j}}_{\substack{\text{choice}\\ \text{kernel}}} + \underbrace{\beta^{\text{side}}}_{\substack{\text{side}\\ \text{bias}}} \tag{1}$$

where c_i is the ith click (+1 for a left click and -1 for right), a_{t-j} is the choice made on the $t-j$th trial (+1 for a left choice and -1 for right), and r_{t-j} is the "reward" on the $t-j$th trial (+1 for correct and -1 for incorrect). Therefore, $\text{a}_{t-j}\text{r}_{t-j}$ indicates the correct side on the $t-j$th trial (+1 when left was correct and -1 when right was correct). The relative effect of each of these terms on the decision is determine by the regression weights: β_i^{click}, β_j^{RL}, β_j^{CK} and β^{side}.

Each of the four suboptimalities can be quantified using different parameters from this model (Figure 2). First, the signal-to-noise ratio, corresponding to suboptimalities arising from neuronal noise, is quantified as the average weight given to all clicks ($\frac{1}{20}\sum_{i=1}^{20} \beta_i^{\text{click}}$). The higher the average click weight, the higher the signal-to-noise ratio or equivalently, the lower the relative level of the noise. This average was significantly different from zero ($T_{107} > 13.25$, FDR corrected for multiple comparisons $p < 0.00001$, Cohen's ds > 1.27) (Figure 2 leftmost panel), indicating that participants based their decision on (at least some of) the clicks and that each click increased the log odds of ultimately choosing that direction by about 0.4.

The second suboptimality, i.e. deviations from a flat integration kernel, is quantified as the deviation of the click weights from the average, i.e. $\beta_i^{\text{click}} - \frac{1}{20}\sum_{j=1}^{20}\beta_j^{\text{click}}$ (Figure 2 second from left panel). Here we found that participants did not weigh all the clicks equally (repeated measures ANOVA, $F_{19,2033} = 28.21$, $p < 0.00001$, partial $\eta^2 = 0.2086$) . This was not consistent with previous reports with a similar task where all clicks received equal weighting on average (Brunton et al., 2013).

Sequential effects, the third suboptimality, are captured by the effects from previous trials. Specifically, the terms β_j^{choice} and β_j^{correct} quantify the effect of the past choice and past outcome (which indicates which choice would have been correct on previous trials) for the past five trials on the current choice. In line with earlier work (Abrahamyan et al., 2016), we found significant effects of both correct sides and choices from previous trials on participants' choices (Figure 2 second from right panel). Notably, the positive regression weight of previous correct side demonstrates a reinforcement learning effect, in that participants tended to choose whichever side that was shown to be correct on the previous trial (FDR corrected for multiple comparisons $p < 0.00001$, Cohen's d $= 1.39$). The negative regression weight of previous choice indicated an alternating choice kernel — participants tended to choose the opposite of what they had chosen on the previous trial (FDR corrected for multiple comparisons $p < 0.00001$, Cohen's d $= -1.00$).

Finally, the side bias is quantified by the intercept term β^{side} in the model (Figure 2 rightmost panel). This term quantifies the extent to which a participant chose the left side on all trials regardless of which side is the correct side. Here we saw a significant right bias indicated by a significantly negative regression weight ($T_{107} = -4.11$, FDR corrected for multiple comparisons $p = 0.0001$, Cohen's d $= -0.40$) .

Individual differences in integration kernel correlate with individual differences in pupil response

To examine the interaction between individual differences in pupil response and integration behavior, we first computed the pupil diameter change during the

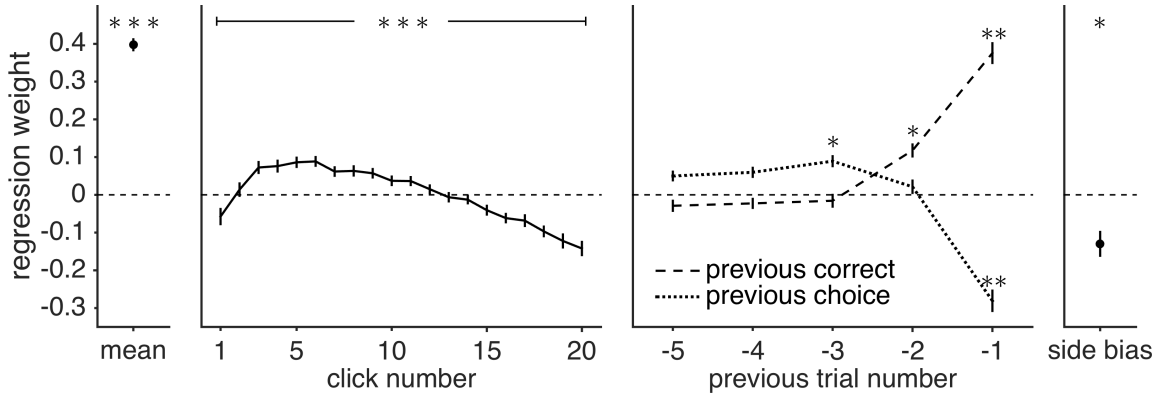

Figure 2. Regression model. (Leftmost) Mean of click weights. $***$: t-test against zero, $p < 0.00001$, Cohen's d $= 2.6606$. (Second from left) Deviation of click weights from the mean. $***$: repeated measures ANOVA $p < 0.00001$. (Second from right) Effect of previous trials. $**$: t-test against zero, FDR corrected for multiple comparisons $p < 0.00001$, Cohen's d > 1; $*$: t-test against zero, FDR corrected for multiple comparisons $p < 0.00001$, Cohen's d > 0.5. (Rightmost) Side bias. $*$: t-test against zero, FDR corrected for multiple comparisons $p = 0.0001$, Cohen's d $= -0.4$.

presentation of clicks stimulus. We time-locked the pupillary response to the onset of the clicks stimulus, and averaged the pupil diameter within each participant. We then took the difference between the peak and the trough of the pupil diameter within the clicks stimulus, which we called the change in pupil response for each participant (Figure 3A). As shown by a median split in Figure 3B, there were considerable individual differences in the pupil change with some participants showing almost no change while others changed a lot during the stimulus.

To examine the relationship of pupil change with overall signal-to-noise ratio and with integration kernel, we used a two way mixed ANOVA to compare the main effects of pupil change (coded as a continuous variable) and time, and the interaction effect between pupil change and time on participants' regression weights from equation (β_i^{click}s from equation (1)). If pupil change has an effect on the overall signal-to-noise ratio, we should see a main effect of pupil change on the regression weights, while if pupil change has an effect on the integration kernel, we should see a significant effect of interaction between pupil and time on regression weights. We found a significant positive effect of

interaction between pupil change and time on regression weights ($F(19, 2014) = 2.225$, $p = 0.0018$, partial $\eta^2 = 0.0206$) (Figure 3 second from left panel), but no significant main effect of pupil change on the regression weights (Figure 3 leftmost panel). This suggests that individual differences in pupil change affect the shape of the integration kernel but not the overall signal-to-noise ratio.

To further understand which click weights were affected by pupil change we performed a correlation analysis between individual differences in the regression weights for each click and individual differences in the pupil change. These post hoc tests suggest that the main change occurs in the second and third clicks, whose weights are increased in participants with high pupil change (correlation $r >= 0.29$, FDR corrected for multiple comparisons $p < 0.002$).

To examine the relationship between pupil change and sequential effects and side bias, we looked at the correlation between pupil change and regression weights of side bias and previous trials (previous correct sides and previous choices). We found no significant relationship between pupil change and either sequential effects (absolute correlation $r < 0.16$, FDR corrected for multiple comparisons $p > 0.48$) (Figure 3 second from right panel) or side bias (correlation $r = -0.01$, FDR corrected for multiple comparisons $p = 0.94$) (Figure 3 rightmost panel).

These results combined indicate that individual differences in pupil change are associated with individual differences in only one of the four suboptimalities, the kernel of integration such that participants with larger pupil change had more uneven integration kernels. For visualization purpose, we ranked the participants by pupil change and split them into high pupil change group and low pupil change group (Figure 3B), and compared their choice regression weights, which also showed a significant interaction between pupil change and time on regression weights ($F(19, 2014) = 2.465$, $p = 0.0004$, partial $\eta^2 = 0.0227$) (Figure 3C).

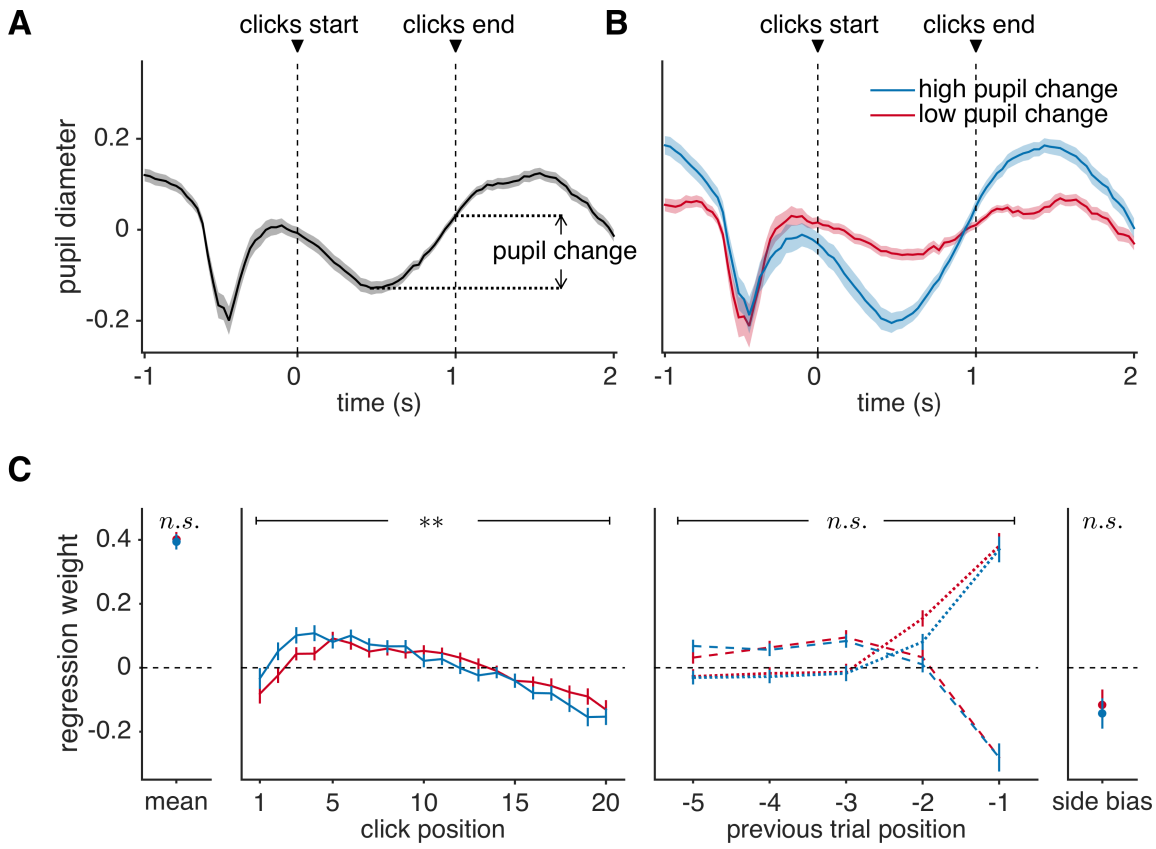

Figure 3. Interaction between pupil change and integration behavior across participants. (A) Pupil diameter time-locked to the onset of clicks, averaged within participants across trials and then across participants. Shaded area indicates s.e.m. across participants. (B) Averaged pupil response across participants split into two groups — high vs. low change in pupil response. Shaded area indicates s.e.m. across participants. (C) Regression weights averaged across participants split into high vs. low pupil change groups for visualization. (Leftmost) Mean regression weight showed no change across groups. (Second from left) Pupil had a significant interaction effect with time on regression weights. (Second from right) Effect of previous trials showed no differences across groups. $**$: two way mixed ANOVA, p=0.001. Error bars s.e.m. across participants. (Rightmost) Side bias showed no change across groups.

Trial-by-trial variability in pupil change correlates with trial-by-trial variability in signal-to-noise ratio, but not integration kernel, sequential effects or side bias

To quantify how trial-by-trial pupil change relates to the four suboptimalities in evidence accumulation, we modified the regression model (equation 1) to include interaction terms between clicks, previous trials, and trial-by-trial fluctuations in pupil:

$$\begin{aligned}\text{logit}(p_{\text{left}} \text{ at trial } t) = &\sum_{i=1}^{20} \beta_i^{\text{click}} \text{c}_i + \sum_{j=1}^{5} \beta_j^{\text{RL}} \text{a}_{t-j} \text{r}_{t-j} + \sum_{j=1}^{5} \beta_j^{\text{CK}} \text{a}_{t-j} + \beta^{\text{side}} \\ &+ \underbrace{\sum_{i=1}^{20} \beta_i^{\text{click}\times\Phi} \text{c}_i \Phi_t}_{\text{kernel}\times\text{pupil}} + \underbrace{\beta_j^{\text{RL}\times\Phi} \text{a}_{t-j} \text{r}_{t-j} \Phi_t}_{\substack{\text{reinforcement}\\ \text{learning}\times\text{pupil}}} + \underbrace{\beta_j^{\text{CK}\times\Phi} \text{a}_{t-j} \Phi_t}_{\substack{\text{choice kernel}\\ \times\text{pupil}}} + \underbrace{\beta_j^{\text{side}\times\Phi} \Phi_t}_{\substack{\text{side bias}\\ \times\text{pupil}}}\end{aligned} \tag{2}$$

where Φ_t is the pupil change measure at trial t. The first four terms in this model are the same as equation (1), and the last four terms are the respective interaction terms of clicks (integration kernel), previous correct side, previous choice, and side bias with pupil change. With repeated measures ANOVA, we did not find a significant effect of time on $\beta^{\text{click}_i \times \text{pupil}_t}$ ($F(19, 2033) = 0.72$, $p = 0.81$), suggesting that pupil change did not modulate the integration kernel on a trial-by-trial level. We did not find a signfiicant interaction effect between pupil and previous correct side ($T(107) = 1.78$, FDR corrected $p = 0.14$), previous choice ($T(107) = -0.78$, FDR corrected $p = 0.62$), or side bias ($T(107) = -1.63$, FDR corrected $p = 0.19$) either.

We then tested whether there is an interaction between pupil change and overall signal-to-noise ratio with a slightly modified version of equation (2):

$$\begin{aligned}\text{logit}(p_{\text{left}} \text{ at trial } t) = &\sum_{i=1}^{20} \beta_i^{\text{click}} \text{c}_i + \sum_{j=1}^{5} \beta_j^{\text{RL}} \text{a}_{t-j} \text{r}_{t-j} + \sum_{j=1}^{5} \beta_j^{\text{CK}} \text{a}_{t-j} + \beta^{\text{side}} \\ &+ \underbrace{\beta^{\Delta\text{click}\times\Phi} \Delta\text{c} \Phi_t}_{\text{SNR}\times\text{pupil}} + \underbrace{\beta_j^{\text{RL}\times\Phi} \text{a}_{t-j} \text{r}_{t-j} \Phi_t}_{\substack{\text{reinforcement}\\ \text{learning}\times\text{pupil}}} + \underbrace{\beta_j^{\text{CK}\times\Phi} \text{a}_{t-j} \Phi_t}_{\substack{\text{choice kernel}\\ \times\text{pupil}}} + \underbrace{\beta_j^{\text{side}\times\Phi} \Phi_t}_{\substack{\text{side bias}\\ \times\text{pupil}}}\end{aligned} \tag{3}$$

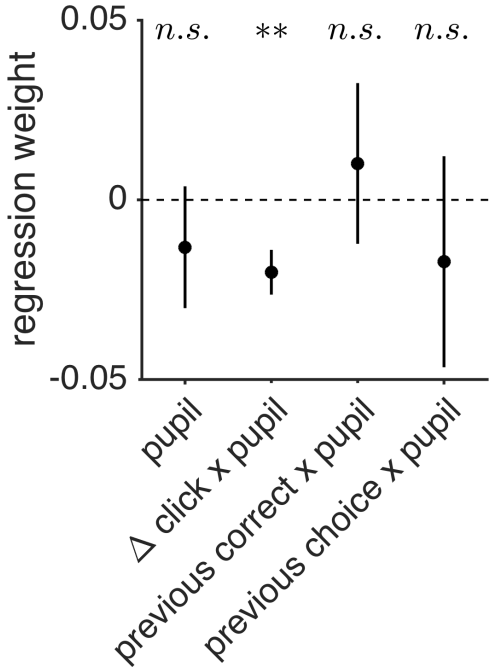

Figure 4. Trial-by-trial interaction between pupil change and integration behavior. $**$: FDR corrected for multiple comparisons $p = 0.0016$. Error bars s.e.m. across participants.

where Δc is number of clicks on the left minus number of clicks on the right, corresponding to the mean click regression weight in Figure 2, indicating the average signal-to-noise ratio. We found that trial-by-trial pupil change interacted significantly with Δc after correction for multiple comparisons ($T_{107} = -3.27$, FDR corrected for multiple comparisons $p = 0.0016$, Cohen's d $= -0.31$), but not with side bias, previous correct sides or choices (Figure 4). These results combined suggest that pupil change on a trial-by-trial level specifically modulates the overall signal-to-noise ratio, and not integration kernel or sequential effects.

Discussion

In this paper we investigated four sources of suboptimality in human evidence integration: neuronal noise (as reflected in the signal-to-noise ratio), uneven integration kernel, sequential effects and side bias, and their relationship with pupil diameter. Of these four suboptimalities only two were associated with pupil at the trial-by-trial and individual difference levels respectively. At the individual difference level, only the

integration kernel was associated with individual differences in pupil change, with a more uneven profile of integration being associated with larger pupil change during stimulus presentation (Figure 3). Conversely, at the trial-by-trial level only noise was associated with trial-by-trial changes in pupil diameter, with larger pupil change associated with smaller signal-to-noise ratio on that trial (Figure 4).

At first glance, the difference between the across- and within-participant results seems contradictory. However, there is reason to believe that these slightly different measures of pupil diameter - individual difference vs trial-by-trial fluctuation - may represent different neural measures, with trial-to-trial fluctuation representing transient, or phasic, locus coeruleus (LC) signals and the average relating to the baseline or tonic LC firing.

At the individual difference level, Eldar and colleagues (Eldar et al., 2013, 2018) have suggested that mean pupil response within a participant is an inverse measure of tonic LC activity. The reasoning here is that phasic and tonic LC have an inverse relationship — phasic activity is low when tonic activity is high, and when tonic activity is high phasic activity is low (Aston-Jones & Cohen, 2005). Thus people with high tonic LC activity should have a lower pupil change overall and, conversely, people with low tonic LC should have higher pupil change overall. In contrast, at the trial-by-trial level, work in monkeys and in mice has suggested that moment-to-moment pupil diameter changes track phasic LC firing (Joshi et al., 2016; Reimer et al., 2016).

Applying these interpretations to the present findings suggests that tonic LC activity changes the kernel of integration while phasic LC decreases the signal-to-noise ratio. The interpretation that tonic LC modulates the integration kernel across participants is consistent with previous work showing that individual differences in pupil change correlate with individual differences in susceptibility to a variety of cognitive biases including risk aversion (Yechiam & Telpaz, 2011), learning styles (Eldar et al., 2013), and decision biases such as framing effect (Eldar et al., 2018). On the other hand, the result that trial-by-trial pupil change modulates not integration kernel but signal-to-noise ratio is consistent with previous work on pupil and evidence

accumulation (Murphy et al., 2014).

The result that pupil diameter does *not* modulate sequential effects seems to be at odds with previous work that has shown pupil-linked systems modulating how information from previous trials affect choice (Krishnamurthy et al., 2017; Nassar et al., 2012). However, an important consideration is that in these tasks, trials were dependent over time such that participants *should* use information from past trials to maximize their reward. In contrast, in our task each trial is independent from the previous trial, so the sequential effects we found are suboptimal for the task and thus may not be regulated by pupil-linked arousal system, which was proposed to be involved in rational regulation of learning dynamics (Nassar et al., 2012). In addition, the sequential effects we observed are also rather small, having about the same effect on the decision as one click Figure 2, thus any correlation with pupil diameter may simply be too small to observe.

Finally, we note that one particularly intriguing feature of our behavioral results is the uneven integration kernel — specifically a 'bump' kernel where the stimulus in the middle is weighted more than the beginning or the end. This contrasts with previous work on perceptual decision making done by Brunton et al. (2013), who found that the integration kernel of rats and well-trained humans was flat, and Yates et al. (2017), who showed a purely primacy driven integration kernel in monkeys.

One obvious question is whether the bump kernel is a genuine feature of the integration process or some artifact of either the task or the analysis pipeline? One possibility is that the bump may be a mixture of primacy and recency kernels which average into a bump across participants. To test this we categorized the integration kernel for each participant into one of the following four shapes: bump, primacy, recency, and flat (for categorization method, see Supplementary Materials). All 108 of these are plotted in Supplementary Figure S1. From here it is easy to see that a large number of subjects (49%) exhibit the bump kernel. This suggests that at least on the level of individual participants, the bump kernel is a feature of the integration process, and not just a artifact of averaging analysis across participants. Of course, the

possibility remains that this bump kernel is a result of mixing a primacy and recency kernels within subject (e.g. some trials have primacy kernels, and some have recency). More detailed modeling work will be needed to tease these interpretations apart.

In summary, our findings show that some, but not all, of the suboptimalities in perceptual decision making are associated with trial-to-trial and individual differences in pupil dilation, putatively linked to tonic and phasic LC signalling. This leads to the tentative suggestion that LC acts to regulate these suboptimalities, helping to adapt them to the statistics of the task at hand. Future work to establish a direct and causal relationship between LC and these suboptimalities (e.g. with drug administration) will be needed to fully test this hypothesis.

Methods

Participants

188 healthy participants (University of Arizona undergraduate students) took part in the experiment for course credit. We excluded 55 participants due to poor performance (accuracy lower than 60%), and then another 25 participants due to poor eye tracking data (see Eye tracking section below). All participants provided informed written consent prior to the experiment, and both experiments were approved by the local ethics committee.

The Bernoulli Clicks Task

Participants made a series of auditory perceptual decisions. On each trial they listened to a series of 20 auditory "clicks" presented over the course of 1 second. Clicks could be either 'Left' or 'Right' clicks, presented in the left or right ear. Participants decided which ear received the most clicks. In contrast to the Poisson Clicks Task (Brunton et al., 2013), in which the click timing was random, clicks in our task were presented every 50 ms with a fixed probability ($p = 0.55$) of occurring in the 'correct' ear. The correct side was determined with a fixed 50% probability.

Participants performed the task on a desktop computer, while wearing

headphones, and were positioned in chin rests to facilitate eye-tracking and pupillometry. They were instructed to fixate on a symbol displayed in the center of the screen, where response and outcome feedback was also displayed during trials, and made responses using a standard keyboard. Participants played until they made 500 correct responses or 50 minutes of total experiment time was reached.

Behavioral analyses

We modeled the choice with logistic regression using equation (1). In particular we assumed that the probability of choosing left on trial t, is a sigmoidal function of the impact from each click, the impact from five previous trial correct sides, the impact from five previous trial choices, and an overall side bias. In this model, by giving the ith click its own weight, we could account for the overall integration kernel.

Eye tracking

A desk-mounted EyeTribe eye-tracker was used to measure participants' pupil diameter from both eyes at a rate of 30 samples per second while they were performing the behavioral task with their head fixed on a chin rest. Pupil diameter data were processed to detect and remove blinks and other artifacts. For each trial, pupil response was extracted time-locked to the start of the trial (Figure 3A). Change in pupil response was computed as the difference between the peak diameter and the minimum diameter during the 1s following trial onset. Pupil response measurements in which more than one-third of the samples contained artifacts were considered invalid and excluded from the analysis. Only participants with at least 200 valid trials were included in analysis ($n = 108$).

Across participants pupil analysis

For each participant, we took the mean pupil response across trials and computed the change in pupil diameter as described in previous section. We then compared this pupil change measurement with regression weights from equation (1). Specifically, we performed a two way mixed ANOVA in which pupil change is a between subject

variable, time is a within subject variable, and regression weight is the dependent variable. We inspected the main effect of pupil change, which informed whether the average regression weight changed with pupil change across participants. We also inspected the interaction effect between pupil change and time, which informed whether pupil change modulates the effect of time on regression weights (i.e. the integration kernel).

Trial-by-trial within participants pupil analysis

For each trial, we took the pupil response and computed the change in pupil diameter. We then modeled participants' choices with the logistic model in equations (3) and (2) to parse out trial-by-trial effects of pupil on integration. The first three terms in both equations were similar to equation (1). But in addition, we assumed that choice was also a function of the interaction between trial-by-trial pupil change and clicks, previous correct side, and previous choice.

Statistics

All data analyses and statistics were done in MATLAB and R. Repeated measures ANOVA, two way mixed ANOVA, and the corresponding post hoc tests done in R. All other analyses and statistical tests done in MATLAB.

Author contributions

W. Keung analyzed the data. T. Hagen collected and preprocessed the data. T. Hagen and R. Wilson designed the experiment. W. Keung and R. Wilson wrote the manuscript. All three authors contributed to intepretation of the results and critical discussion.

Acknowledgments

We thank Maxwell Alberhasky, Chrysta Andrade, Daniel Carrera, Kathryn Chung, Michael de Leon, Zamigul Dzhalilova, Asha Esprit, Abigail Foley, Emily Giron, Brittney Gonzalez, Anthony Haddad, Leah Hall, Maura Higgs, Marcus Jacobs, Min-Hee

Kang, Kathryn Kellohen, Neha Kwatra, Hannah Kyllo, Alex Lawwill, Stephanie Low, Colin Lynch, Alondra Ornelas, Genevieve Patterson, Filipa Santos, Shlishaa Savita, Catie Sikora, Vera Thornton, Guillermo Vargas, Christina West, and Callie Wong for help in running the experiments.

References

Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M., & Gardner, J. L. (2016). Adaptable history biases in human perceptual decisions. *Proceedings of the National Academy of Sciences*, *113*(25), E3548–E3557.

Akrami, A., Kopec, C. D., Diamond, M. E., & Brody, C. D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature*.

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.*, *28*, 403–450.

Barraclough, D. J., Conroy, M. L., & Lee, D. (2004). Prefrontal cortex and decision making in a mixed-strategy game. *Nature neuroscience*, *7*(4), 404.

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, *74*(1), 30–39.

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, *340*(6128), 95–98.

Cavanagh, J. F., Wiecki, T. V., Kochar, A., & Frank, M. J. (2014). Eye tracking and pupillometry are indicators of dissociable latent decision processes. *Journal of Experimental Psychology: General*, *143*(4), 1476.

Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876.

Drugowitsch, J., Wyart, V., Devauchelle, A.-D., & Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, *92*(6), 1398–1411.

Eldar, E., Cohen, J. D., & Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature neuroscience*, *16*(8), 1146.

Eldar, E., Felso, V., Cohen, J. D., & Niv, Y. (2018). A pupillary index of susceptibility to decision biases. *bioRxiv*, 247890.

Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D., & Brody, C. D. (2015).

Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife*, *4*.

Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature reviews neuroscience*, *9*(4), 292.

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in cognitive sciences*, *5*(1), 10–16.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological science*, *17*(9), 767–773.

Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, *520*(7546), 220.

Hauser, T. U., Moutoussis, M., Purg, N., Dayan, P., & Dolan, R. J. (2018). Noradrenaline modulates decision urgency during sequential information gathering. *bioRxiv*, 252932.

Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2016). Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron*, *89*(1), 221–234.

Katz, L. N., Yates, J. L., Pillow, J. W., & Huk, A. C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature*, *535*(7611), 285.

Kaufman, M. T., & Churchland, A. K. (2013). Cognitive neuroscience: sensory noise drives bad decisions. *Nature*, *496*(7444), 172.

Krishnamurthy, K., Nassar, M. R., Sarode, S., & Gold, J. I. (2017). Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nature human behaviour*, *1*(6), 0107.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature neuroscience*, *9*(11), 1432.

Mather, M., Clewett, D., Sakaki, M., & Harley, C. W. (2016). Norepinephrine ignites local hotspots of neuronal excitation: How arousal amplifies selectivity in

perception and memory. *Behavioral and Brain Sciences*, *39*.

Murphy, P. R., Vandekerckhove, J., & Nieuwenhuis, S. (2014). Pupil-linked arousal determines variability in perceptual decision making. *PLoS computational biology*, *10*(9), e1003854.

Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature neuroscience*, *15*(7), 1040.

Newsome, W. T., & Pare, E. B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (mt). *Journal of Neuroscience*, *8*(6), 2201–2211.

Osborne, L. C., Lisberger, S. G., & Bialek, W. (2005). A sensory source for motor variation. *Nature*, *437*(7057), 412.

Rajkowski, J. (1993). Correlations between locus coeruleus (lc) neural activity, pupil diameter and behavior in monkey support a role of lc in attention. *Soc. Neurosc., Abstract, Washington, DC, 1993*.

Reimer, J., McGinley, M. J., Liu, Y., Rodenkirch, C., Wang, Q., McCormick, D. A., & Tolias, A. S. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature communications*, *7*, 13289.

Sara, S. J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nature reviews neuroscience*, *10*(3), 211.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in neurosciences*, *27*(3), 161–168.

Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (Vol. 135). MIT press Cambridge.

Wyart, V., Myers, N. E., & Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. *Journal of neuroscience*, *35*(8), 3485–3498.

Yates, J. L., Park, I. M., Katz, L. N., Pillow, J. W., & Huk, A. C. (2017). Functional dissection of signal and noise in mt and lip during decision-making. *Nature*

neuroscience, *20*(9), 1285.

Yechiam, E., & Telpaz, A. (2011). To take risk is to face loss: a tonic pupillometry study. *Frontiers in psychology*, *2*, 344.

Supplementary Material

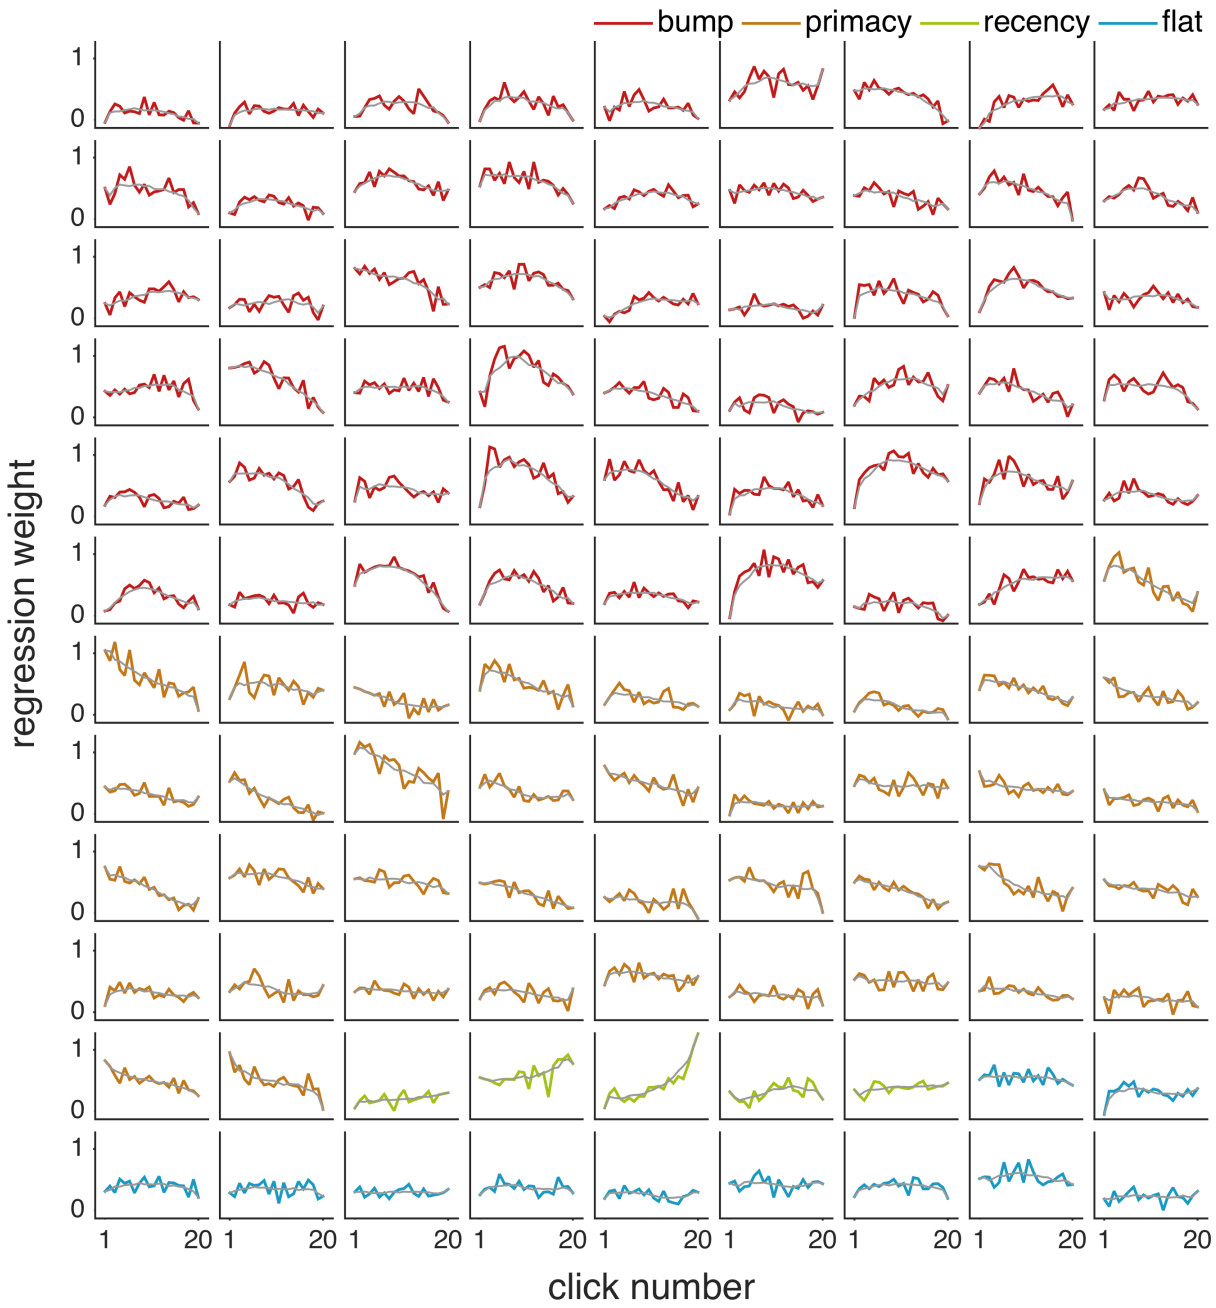

Figure S1. Individual integration kernel plots. Colored lines are regression weights of clicks from equation (1). Plots are sorted and color coded by kernel shape. Light grey line shows smoothed integration kernel.

To categorize the kernel for each participant into one of the four shapes, we fit polynomial functions with different degrees to participants' choices, and selected the best fitting model with model comparison using the Akaike Information Criterion

(AIC). Specifically, we assume that the probability of choosing left at trial t is (the logit of) the weighted sum of clicks, where the weights are from a polynomial function, as shown in the following equation:

$$\text{logit}(p_{\text{left}} \text{ at trial } t) = \sum_{i=1}^{20} \beta_i^{\text{poly}} \text{c}_i \text{ , where } \beta_i^{\text{poly}} = \sum_{n=0}^{N} \alpha_n i^n \tag{4}$$

We fitted three different polynomial functions by changing N from 0 to 2: constant, linear, and quadratic. We then selected the best fitting function for each participant by comparing the fits from different polynomials with AIC. We categorized each participant's integration kernel into one of the four shapes using the following criteria: (1) flat: kernel was best fit with the constant function; (2) primacy: kernel was best fit with linear function with a negative slope (α_1), or with quadratic function with a minimum ($\alpha_2 > 0$) and the minimum is located later than the 10th click, or with quadratic function with a maximum ($\alpha_2 < 0$) and the maximum is located earlier than the 2nd click; (3) recency: kernel was best fit with linear function with a positive slope, or with quadratic function with a minimum ($\alpha_2 > 0$) and the minimum is located earlier than the 10th click, or with quadratic function with a maximum ($\alpha_2 < 0$) and the maximum is located later than the 18th click; (4) bump: kernel that did not meet the previous three criteria (i.e. kernel was best fit with quadratic function and was neither primacy nor recency).