Trial-by-trial fluctuations in amygdala activity track motivational enhancement of desirable sensory evidence during perceptual decision-making

Abbreviated Title: Amygdala contributions to motivated seeing

Ren Paterson, Yizhou Lyu, Yuan Chang Leong*

Department of Psychology, University of Chicago, Chicago, IL 60637

* Correspondence should be addressed to Yuan Chang Leong at ycleong@uchicago.edu

Number of pages: 34

Number of figures: 4

Number of words for Abstract: 228; Introduction: 587; Discussion: 1421

The authors declare no competing financial interests.

Acknowledgments: We would like to thank members of the Motivation and Cognition Neuroscience Lab for their feedback on this work, and Howard Nusbaum for helpful comments on a previous version of this manuscript.

Abstract

People are biased towards seeing outcomes that they are motivated to see. For example, sports fans of opposing teams often perceive the same ambiguous foul in favor of the team they support. Here, we test the hypothesis that amygdala-dependent allocation of visual attention facilitates motivational biases in perceptual decision-making. Human participants were rewarded for correctly categorizing an ambiguous image into one of two categories while undergoing fMRI. On each trial, we used a financial bonus to motivate participants to see one category over another. The reward maximizing strategy was to perform the categorization task accurately, but participants were biased towards categorizing the images as the category we motivated them to see. Heightened amygdala activity preceded motivation consistent categorizations, and participants with higher amygdala activation exhibited stronger motivational biases in their perceptual reports. Trial-by-trial amygdala activity was associated with stronger enhancement of neural activity encoding desirable percepts in sensory cortices, suggesting that amygdala-dependent effects on perceptual decisions arose from biased sensory processing. Analyses using a drift diffusion model provide converging evidence that trial-by-trial amygdala activity was associated with stronger motivational biases in the accumulation of sensory evidence. Prior work examining biases in perceptual decision-making have focused on the role of frontoparietal regions. Our work highlights an important contribution of the amygdala. When people are motivated to see one outcome over another, the amygdala biases perceptual decisions towards those outcomes.

Significance Statement

Forming accurate perceptions of the environment is essential for adaptive behavior. People however are biased towards seeing what they want to see, giving rise to inaccurate perceptions and erroneous decisions. Here, we combined behavior, modeling, and fMRI to show that the bias towards seeing desirable percepts is related to trial-by-trial fluctuations in amygdala activity. In particular, during moments with higher amygdala activity, sensory processing is biased in favor of desirable percepts, such that participants are more likely to see what they want to see. These findings highlight the role of the amygdala in biasing visual perception, and shed light on the neural mechanisms underlying the influence of motivation and reward on how people decide what they see.

Introduction

People are biased towards seeing percepts that they are motivated to see, a phenomenon known as 'wishful seeing' (Dunning and Balcetis, 2013). For example, sports fans of opposing teams watching the same game often perceive an ambiguous foul differently, with each group of fans judging the foul in favor of the team they support. When motivated to see one interpretation of reality over another, people become less objective and are more prone to perceptual errors (Voss et al., 2008; Bromberg-Martin and Sharot, 2020; Leong et al., 2021). Individuals who score higher on trait paranoia are also more likely to exhibit wishful seeing (Rossi-Goldthorpe et al., 2021), suggesting that extreme cases of wishful seeing might contribute to biased beliefs about the world that are the hallmark of many psychiatric disorders (Kube and Rozenkrantz, 2021). Why are people more likely to see what they are motivated to see?

To formulate a mechanistic understanding of wishful seeing, we can draw on the extensive neurophysiological and behavioral modeling work characterizing how perceptual decisions are determined from noisy sensory information (e.g., Heekeren et al., 2008; Shadlen and Kiani, 2013). In particular, perceptual decisions are thought to involve the temporal accumulation of sensory evidence from the environment. When the level of sensory evidence reaches a threshold, the individual commits to a decision. Recent work suggests that wishful seeing is driven by the selective enhancement of neural activity encoding desirable percepts (Leong et al., 2019). This biases sensory evidence accumulation in favor of percepts one is motivated to see, giving rise to motivationally biased perceptual decisions. While prior work has shown that motivation enhances the neural representation of desirable percepts in sensory cortices, how this enhancement occurs is not known. Consequently, the neural mechanisms that give rise to wishful seeing remain underspecified.

The amygdala has extensive connections to the visual cortex (Freese and Amaral, 2005), and is well-positioned to mediate wishful seeing. Human neuroimaging and non-human primate electrophysiology studies have found that the amygdala plays a central role in biasing attention towards affectively salient aspects of the environment (Pessoa and Adolphs, 2010; Pourtois et al., 2013; Peck and Salzman, 2014a; Mather et al., 2016). For example, fluctuations in amygdala activity correlate with the enhanced neural representation of images that had acquired affective salience due to having been previously paired with an electrical shock (Lim et al., 2009). We hypothesize that wanting to see a percept ascribes affective significance to the associated sensory features, and the amygdala is similarly involved in enhancing the neural representation of these features. Past studies, however, have not examined the role of the amygdala in contexts where participants are motivationally biased to see one percept over another. It is thus unclear if and how the amygdala biases perception towards desirable percepts.

In the current work, we conducted a series of pre-registered analyses on a previously published dataset (Leong et al., 2019) to test the hypothesis that the amygdala facilitates wishful seeing by enhancing sensory evidence accumulation in favor of desirable percepts. We first assessed if amygdala activity was associated with motivationally biased perceptual decisions. Next, we tested if amygdala activity was associated with stronger neural representations of percepts participants were motivated to see. Finally, we incorporated participants' trial-by-trial amygdala activity, perceptual decisions and response times into a computational model to test the hypothesis that amygdala activity was specifically associated with motivational biases in evidence accumulation. Together, our work takes a convergent approach that combines behavioral, neural and modeling measures to characterize the neural mechanisms underlying motivational biases in perceptual decision-making.

Materials and Methods

Participants

Thirty-three participants provided informed consent and were compensated between $30-50 depending on their performance in the task. Data were discarded from 3 participants due to excessive head motion (i.e., >3 mm in any direction) during >1 scanning sessions, yielding a final sample of 30 participants (17 males, 13 females; age range, 18-43 years; mean age, 22.3 years). These data have been previously reported in Leong et al. (2019).

Stimuli

Stimuli consisted of grayscale images created by combining a scene image and face image in varying proportions. Face images were frontal photographs taken from the Chicago Face Database (Ma et al., 2015), and consisted of half female and half male faces with neutral expressions. Scene images consisted of half outdoor and half indoor scenes. For each participant, six stimulus sets were created. Each set contained 40 images: 1 × 100% scene, 3 × 65% scene, 5 × 60% scene, 7 × 55% scene, 8 × 50% scene, 7 × 45% scene, 5 × 40% scene, 3 × 35% scene and 1 × 0% scene. Images with more than 50% scene were considered "more scene" while images with less than 50% scene were considered "more face". Stimuli were presented to participants using MATLAB (Mathworks) and the Psychophysics Toolbox (Brainard, 1997).

Experimental Task

Participants were told that they would be performing a visual categorization task with a teammate ("Cooperation" condition) and an opponent ("Competition" condition) (Fig. 1). At the beginning of each trial, either the teammate or the opponent would make a bet on whether the upcoming image would consist of more face or more scene (i.e. whether the face image or scene image was of higher intensity; "Motivation Cue", presented for 4 seconds). Participants were then presented

with a face-scene composite image (see *Stimuli*). Participants would earn 40 cents if the teammate's bet was correct (e.g., if the teammate bet "more face" and the image actually consists of more face) and lose 40 cents if the teammate's bet was wrong. In contrast, participants would lose 40 cents if the opponent's bet was correct and earn 40 cents if the opponent's bet was wrong. As such, participants would be motivated to see the image as the category that the teammate had bet on and the category that the opponent had bet against.

Participants were then asked to indicate whether the image consists of more face or more scene and received a 10-cent reward if they correctly categorized the image. If participants did not respond in 4 seconds, the trial would time out and participants would lose the opportunity to earn the reward. The image remained on screen for the entire 4 seconds regardless of when participants made their response. Participants then indicated how confident they were in their categorization on a 1-5 rating scale. After a variable inter-trial-interval (ITI) of 2-4 seconds, participants proceeded to the next trial. Participants performed 2 fMRI runs of the Cooperation condition and 2 fMRI runs of the Competition condition in an interleaved order that was counterbalanced across participants. Each run was approximately 8 minutes long and consisted of 40 trials.

Importantly, the 40-cent bonus associated with the teammate or opponent's bet depended solely on the objective category of the image (i.e. whether the image actually contained more face or more scene) and was not affected by participants' responses. For example, if the teammate bet that the upcoming image contains more face but the image objectively contained more scene, participants would lose the 40 cents regardless of how they categorized the image. As such, participants would earn the most amount of money from the experiment if they ignored the bets and categorized the images accurately based on what they saw to earn the maximum reward from correctly categorizing each image. Participants were told that both the teammate and the opponent

did not have *a priori* information about the images and were making blind bets. Participants were also told that neither the teammate or opponent would be informed of their responses, and were thus not pressured to conform to the bets. Unbeknownst to participants, the teammate's and opponent's bets were pseudo-randomized to be correct on half of the trials.

The instructions were delivered verbally by the experimenter and also in written form on screen at the beginning of the experiment. It was intuitive to participants that the bet depended on the objective category of the images rather than their categorizations, and that they should categorize each image accurately to earn maximum reward. In the post-experiment survey, all participants responded that the task instructions were clear and that they had understood the task correctly.

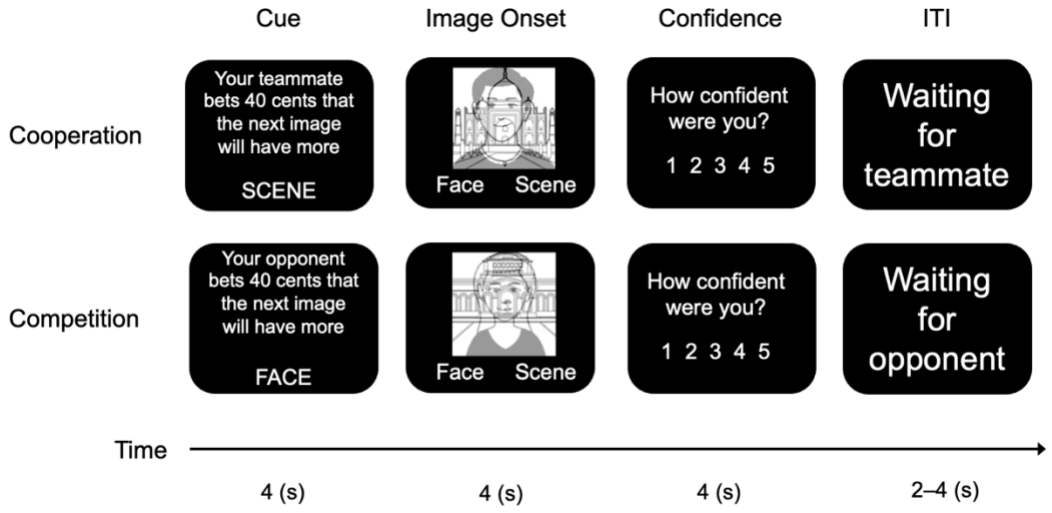

Figure 1. Experimental design. Each trial begins with a teammate (Cooperation condition) or an opponent (Competition condition) making a bet on whether the upcoming image would consist of more face or more scene. Participants earn 40 cents if the teammate's bet is correct, and lose 40 cents otherwise. Conversely, they lose 40 cents if the opponent's bet is correct and earn 40 cents otherwise. Participants are then presented with a face-scene composite image, and have to categorize whether the image has more face or more scene. Finally, participants rated how confident they were in their categorization. Images have been replaced with cartoon images to comply with *bioRxiv* policy.

Motivational Bias

In earlier work, we had found that motivational bias and response times did not differ between the Competition and Cooperation conditions - participants were just as likely and quickly to categorize the image as the category that the teammate had bet on as they were to categorize it as the category the opponent had bet against (Leong et al., 2019). Therefore, in the current study, we collapsed across the two conditions based on whether participants were motivated to see more face (i.e. when the teammate bet on more face or when the opponent bet on more scene) or more scene (i.e. when the teammate bet on more scene or when the opponent bet on more face).

We modeled participants' data using a generalized linear mixed effects model (GLMM) using the glmer function from the lme4 package in R:

response ~ % scene + motivation + (motivation | subj)

where *response* denotes the perceptual report (face = 0, scene = 1), *motivation* is the category that participants were motivated to see (face = 0, scene = 1), and *% scene* is the proportion of scene in an image. The model included *motivation* as a random effect to account for random variation between subjects. Importantly, the estimate of the random effect of a given participant provides a measure of the effect of *motivation* on that participants' perceptual categorizations. For each participant, we added the random effect to the corresponding fixed effect to obtain an individual difference measure of motivational bias.

We used a robust regression model to examine the relationship between motivational bias and amygdala activation (see fMRI analyses) using the rlm function from the MASS package in R with default settings (M-estimation and Huber weighting). Robust regression has been shown to be less sensitive to the influence of outliers than Pearson correlation and ordinary least squares

regression (Venables and Ripley, 2013). Statistical significance of regression coefficients were assessed using a robust F test as implemented in the sfsmisc package in R.

MRI Data Acquisition and Preprocessing

A 3T General Electric scanner was used to collect the MRI data. Functional images were obtained with T2*-weighted echo planar imaging (EPI) pulse sequences. Each volume comprised 46 transverse slices. Volumes were acquired in interleaved order, with the following imaging parameters: repetition time, 2 s; echo time, 25 ms; flip angle, 77º; voxel size, 2.9 mm^3. A T1-weighted pulse sequence was used to acquire anatomical images at the beginning of the session with the following parameters: repetition time, 7.2 ms; echo time, 2.8 ms; flip angle, 12º; voxel size, 1mm^3.

Preprocessing and analysis of EPI images were performed using FSL/FEAT v6.05 (FMRIB software library, FMRIB, Oxford, UK) and included motion correction (rigid-body realignment of functional volumes to the first volume), slice-timing correction, high-pass filtering of the data with a 100 ms cut-off, and spatial smoothing using a Gaussian kernel with a full-width at half-maximum of 4 mm. For multivoxel classification analyses, our classifier was trained and tested in each participant's native space (see fMRI analyses). For all other analyses, we first registered the functional volumes to participants' anatomical images (6 d.f.) and then registered the volumes to a template brain in the Montreal Neurological Institute (MNI) space (affine transformation with 12 d.f.).

Regions of interest (ROI) definition

We defined a bilateral amygdala ROI as the voxels that are estimated to have greater than 0.5 probability as being part of the amygdala according to the Harvard-Oxford Subcortical Structural Atlas. To measure sensory representations, we defined an occipito-temporal ROI by combining

the bilateral occipital and temporal lobe masks from the Harvard-Oxford Structural Atlas. This mask includes the visual cortex as well as the ventral temporal areas important for object recognition (Grill-Spector, 2003). We defined additional ROIs from an atlas derived from applying independent component analysis to resting-state data (Shirer et al., 2012). Specifically, the anterior insula and dorsal anterior cingulate cortex (dACC) ROIs were obtained from the set of regions labeled the "anterior salience network", while the dorsolateral prefrontal cortex (DLPFC) and inferior parietal lobule (IPL) ROIs were obtained from the set of regions labeled the "executive control network".

fMRI Analyses

We ran a general linear model (GLM) analysis contrasting neural activity on motivation consistent trials (i.e. trials where participants categorized the image as the desirable category) against that on motivation inconsistent trials (i.e. trials where participants categorized the image as the less desirable category). The two trial types were modeled as separate regressors, each created by convolving a stimulus function, consisting of an impulse response at the time of stimulus onsets, with a Gamma hemodynamic response function (HRF). For each participant, a separate model was estimated for each run. Response times (convolved with HRF) and extended motion parameters (i.e. standard motion parameters, their temporal derivatives, and the squares of both the standard motion parameters and the temporal derivatives; not convolved with HRF) were included as covariates of no interest. The results were then combined across the four runs in a fixed effects analysis to obtain a statistical parametric map for each participant. We then used FSL's *featquery* function to extract the average *z*-statistic of the contrast in the amygdala ROI for each participant, and ran a one-sample t-test against zero to assess if amygdala activity was higher on motivation consistent trials than on motivation inconsistent trials.

To examine if individual differences in amygdala activation tracked motivational biases in categorizations, we ran a second set of GLM analyses to estimate the average magnitude of amygdala activation for each participant. The modeling procedure was the same as before, except that motivation consistent and motivation inconsistent trials were not modeled separately. Instead, a single regressor was created by convolving a stimulus function, consisting of an impulse response at the time of stimulus onset of all trials, with a HRF. For each participant, we used FSL's *featquery* function to extract the average beta weights in the amygdala ROI as an estimate of the extent to which that participant's amygdala was activated by the task. We then used a robust regression analysis to assess if individual differences in amygdala activation were associated with the magnitude of motivational bias in categorizations (see *Motivational Bias*).

To examine the time course of amygdala activation and its relationship to motivation consistent responses, we extracted the time course in the amygdala ROI from the preprocessed BOLD data of each participant. The time courses were z-scored across each run, shifted by 4 seconds to correct for haemodynamic lag, then time locked to the stimulus onset of each trial. For each time point from -8 seconds prior to stimulus onset to 8 seconds post stimulus onset, we ran a generalized linear mixed effects model examining the effect of amygdala activity on whether participants made a motivation consistent categorization:

$$\text{MotCon} \sim \text{amygdala} + (\text{amygdala} \mid \text{subj})$$

where *MotCon* denotes whether participants made a motivation consistent or inconsistent response (motivation consistent = 1, motivation inconsistent = 0) and amygdala denotes amygdala activity at a given time point.

At the end of the experimental task, participants performed two runs of a localizer task where they were presented with unambiguous face and scene images one at a time (5 blocks of 15

images of each category; each image presented for 2 seconds with a 2-second interval). Participants had to indicate whether each face was male or female, and whether each scene was indoors or outdoors. We trained an L1-regularized logistic regression classifier (C = 1) to classify whether the participant was seeing a face or scene based on the patterns of activity in the occipito-temporal ROI. We then applied the trained classifier to the single-trial activation patterns in the experimental task. Specifically, all trials in a run were simultaneously estimated in a single GLM, with each trial modeled using a separate regressor consisting of an impulse response at the time of stimulus onset convolved with a HRF (i.e. a "Least-Squares All" model; Mumford et al., 2014). The beta weights of each trial were then used as inputs for the trained classifier. For each trial, the classifier outputs the probability that the participant was viewing a scene vs. a face based on the beta weights in the occipito-temporal ROI, thus providing us with a measure of the relative strength of scene vs. face selective neural activity for that trial.

We estimated anticipatory amygdala activity on each trial using a Least-Squares All GLM that modeled each trial as a separate regressor created by convolving a stimulus function, consisting of an impulse response at the time of cue onset, with a HRF. The resulting beta weights obtained for each regressor were then averaged within the amygdala ROI as the estimate of anticipatory amygdala activity for that trial. A linear mixed effects model was used to test the relationship between trial-by-trial anticipatory amygdala activity and sensory representations:

scene vs. face activity ~ % scene + motivation + amygdala + motivation*amygdala +
(motivation + amygdala + motivation*amygdala | subj)

where *scene vs. face activity* is the relative strength of scene vs. face-selective neural activity as indexed by the classifier output, *% scene* is the amount of scene in an image, *motivation* is the category participants were motivated to see (face = 0, scene = 1), and *amygdala* is amygdala

activation at the time of the cue. The coefficient for the interaction term, *motivation*amygdala*, would thus capture the extent to which motivational effects on sensory-related activity depended on anticipatory amygdala activation. To unpack the interaction effect, we divided each participant's data into trials with high and low amygdala activity based on a median split, and examined the effect of motivation on *scene vs. face activity* for the two trial types separately:

scene vs. face activity ~ % scene + motivation + (motivation | subj)

In a series of exploratory analyses, we also examined if motivational effects on scene and face-selective activity was moderated by activity in the anterior insula, dorsal anterior cingulate cortex (dACC), dorsolateral prefrontal cortex (DLPFC), and inferior parietal lobule (IPL). All models were estimated using the lmer function in the lme4 package in R, with P values computed from t-tests with Satterthwaite approximation for the degrees of freedom as implemented in the lmerTest package.

Drift Diffusion Model

We fit a drift diffusion model to participants' data using the HDDM toolbox with default priors (Wiecki et al., 2013). The model assumes that participants' decisions are determined by the stochastic accumulation of sensory evidence towards one of two decision thresholds. We estimated model parameters that determined the starting point of the decision process (parameter *z*), the rate of evidence accumulation (i.e. *drift rate*, parameter *v*), the distance between the two decision thresholds (parameter *a*) and time not related to the decision process *(e.g.,* time for stimulus encoding or motor response; *non-decision time*, parameter *t*). We modeled the "scene" decision threshold as the top boundary (i.e. scene threshold = *a*) and the "face" threshold as the bottom boundary (i.e. face threshold = *0*). HDDM implements hierarchical Bayesian estimation, which assumes that parameters for individual participants were randomly drawn from a group-level

distribution. Individual participant parameters and group-level parameters were jointly estimated using Markov Chain Monte Carlo sampling (20,000 samples; burn-in=2,000 samples; thinning=2). To account for outliers generated by a process other than that assumed by the model (e.g., accidental button press, lapses in attention), we estimated a mixture model where 5% of trials were assumed to be distributed according to a uniform distribution. We chose not to include inter-trial variability parameters as simulation studies have shown that they can compromise the reliability and accuracy of the estimates of individual parameters (Boehm et al., 2018). A model that included parameters for inter-trial variability in the starting point, drift rate and non-decision time yielded qualitatively similar results (see *Results*).

The drift rate was allowed to vary trial-by-trial as a function of the stimulus, the category participants were motivated to see, and amygdala activity:

$$v = \beta_{v1}\,\%scene\ +\ \beta_{v2}\,motivation + \beta_{v3}\,amygdala\ + \beta_{v4}\,amygdala * motivation + \beta_{v0}$$

where *% scene* refers to the percentage scene in an image, *motivation* refers to the category participants were motivated to see (face = -1; scene = 1), and *amygdala* refers to amygdala activation at cue onset estimated from a GLM (see fMRI analysis). The coefficient for the *amygdala*motivation* interaction reflects the extent to which the effect of motivation on the drift rate depended on amygdala activation. Specifically, a positive value of β_{v4} would indicate that the drift bias towards the motivation consistent category was stronger when amygdala activation was higher. To assess if β_{v4} was positive, we extracted the posterior distribution of β_{v4} estimated from participants' data and computed the proportion of the distribution that was greater than 0. This proportion denotes the probability the amygdala activation moderated motivational effects on the drift rate.

Within the same model, the starting point was also allowed to vary trial-by-trial as a function of the motivation consistent category and the amygdala activation at cue onset:

$$z = \frac{1}{1 + exp - (\beta_{z1} motivation + \beta_{z2} amygdala + \beta_{z3} amygdala * motivation + \beta_{z4})}$$

Here, the coefficient for the *amygdala*motivation* interaction reflects the extent to which the effect of motivation on the starting point depended on amygdala activation. Specifically, a positive value of β_{z3}would indicate that the starting point shifts more towards the motivation consistent category when amygdala activation was higher. To assess if β_{z3} was positive, we extracted the posterior distribution of β_{z3} estimated from participants' data and computed the proportion of the distribution that was greater than 0. This proportion denotes the probability the amygdala activation moderated motivational effects on the starting point.

Model convergence was assessed using the Gelman-Rubin $\hat{R}$ statistic. Model convergence metrics, posterior means and 95% credible intervals of model parameters are reported in Extended Table 4.1.

Preregistration

We pre-registered how we defined the amygdala and occipitotemporal ROIs. We had also pre-registered the following hypotheses: (1) participants with stronger amygdala responses will exhibit greater motivational biases in their perceptual categorizations, (2) participants would be more biased in their categorizations on trials with high amygdala activity, (3) there would be stronger motivational enhancement of desirable sensory representations on trials with high amygdala activity, and (4) there was be stronger motivational biases in evidence accumulation, as estimated by the DDM, on trials with high amygdala activity. The pre-registration can be viewed here: https://osf.io/2vbcn/?view_only=030ad17d15724af2bb5f6f02ca5b34e7

Results

Thirty participants were scanned using fMRI while they were presented with visually ambiguous images consisting of a face image and a scene image superimposed with one another. Participants had to categorize whether they saw more face or more scene in each image, and received a monetary reward for each correct categorization. At the beginning of each trial, we motivated participants to see one category over the other with a cue indicating that they would win extra money if the image was of one category (the "desirable" percept) and lose some of their earnings if the image was of the other category (the "less desirable" percept) (see *Materials and Methods*). The bonus or loss depended on the objective category of the image (i.e. whether the image actually contained more face or more scene) and not participants' categorizations. As such, participants would earn the most money from the task if they categorized the images accurately and were not influenced by what they were motivated to see. In earlier work, we had shown that participants were nevertheless more likely to categorize the images as the desirable percept, and that this bias was associated with enhanced neural activity selective to the desirable percept in the ventral visual stream (Leong et al., 2019).

Amygdala activity was associated with motivational bias in perceptual decisions

We divided trials into trials where participants categorized the image as the desirable percept (motivation consistent trials) and trials where they categorized the image as the less desirable percept (motivation inconsistent trials). We then ran a GLM that contrasted neural activity on motivation consistent trials against that on motivation inconsistent trials. In a bilateral amygdala ROI, the average z-statistic of the contrast was significantly above zero ($t(29) = 2.96$, $P = 0.006$; Fig. 2A, 2B), indicating that amygdala activity was higher when participants made motivation consistent categorizations than when they made motivation inconsistent categorizations.

To examine the time course of amygdala activation, we extracted the amygdala activity time course of each trial and used generalized linear mixed effects models to assess the relationship between each time point and participants' categorizations. The time courses were first shifted by 4 seconds to account for hemodynamic lag, and then time-locked to image onset. Amygdala activity was marginally associated with making a motivation consistent response at the time of the motivational cue (i.e. when the teammate or opponent made their bet; $b = 0.059$, $SE = 0.035$, $z = 1.71$, $p = 0.087$). This relationship was significant 2 seconds later ($b = 0.10$, $SE = 0.037$, $z = 2.80$, $p = 0.0051$), and remained significant until the image appeared on screen ($b = 0.112$, $SE = 0.036$, $z = 3.13$, $p = 0.0018$; Fig. 2C). As the relationship between amygdala activity and motivational consistent responses was present prior to image onset, the results suggest that the heightened amygdala activity did not reflect activation in response to participants having seen the desirable category. Instead, they suggest that the amygdala activation predisposes participants towards categorizing the image as the desirable category.

Next, we assessed if individual differences in motivational bias were associated with amygdala activation. We first fit a generalized linear mixed effects model to participants' data to jointly estimate the effect of motivation and % scene in the image on perceptual categorizations. For a given participant, the random effect of motivation reflects the extent to which that participant was more likely to make motivation consistent categorizations. We added each participant's random effect to the fixed effect of motivation on perceptual categorizations to obtain an individual difference measure of motivational bias. For each participant, we also ran a GLM to estimate the amygdala response to the task (see Methods). Using a robust regression analysis, we found that participants who exhibited stronger motivational biases also had stronger amygdala activation (Fig. 2C, $b = 4.3$, $SE = 0.86$, $F(1, 28) = 5.6$, $P = 0.025$; Fig. 2D).

In summary, motivation consistent categorizations were associated with higher amygdala activity, this amygdala activation preceded image onset, and participants who exhibited stronger motivational biases in their categorizations had a stronger amygdala response to the task. Together, these results suggest that heightened amygdala activity biased participants towards making motivationally consistent perceptual decisions.

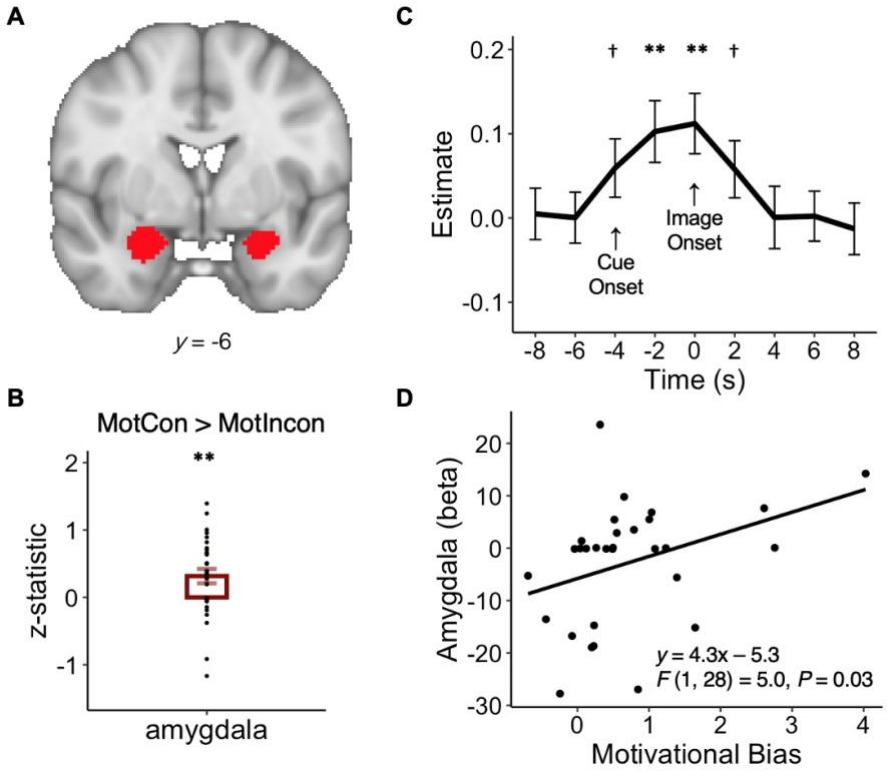

Figure 2. Amygdala activity was associated with motivational bias in perceptual decisions. A. Amygdala ROI defined using the Harvard-Oxford Subcortical Structural Atlas. **B.** Amygdala activity was higher when participants made motivation consistent categorizations (MotCon) than when they made motivation inconsistent categorizations (MotIncon). Data points denote average z-statistic of MotCon > MotIncon contrast in the amygdala ROI for individual participants. Error bars indicate s.e.m. **C.** Amygdala activation preceded image onset. For each timepoint, we ran a generalized linear mixed effects model estimating the effect of amygdala activity on whether participants would make a motivation consistent response. Plot shows estimates of regression coefficients with standard errors. Time points have been corrected for haemodynamic lag by shifting the BOLD data by 4s. †$P < 0.1$, **$P < 0.01$. **D.** Participants with stronger motivational biases in their categorizations had on average higher amygdala responses.

Anticipatory amygdala activity was associated with stronger motivational enhancement of category-selective neural activity

The preceding analyses indicated amygdala activity was associated with motivational biases in perceptual decisions. Did amygdala activity also moderate motivational enhancement of sensory representations? To answer this question, we used multivariate pattern analysis to quantify face- and scene-selective activity in the occipitotemporal cortex. We first trained a logistic regression classifier on BOLD data from a localizer task to categorize whether participants were seeing a face or a scene based on the pattern of activity in a occipitotemporal cortex ROI (see Methods). We then applied the trained classifier to data from the experimental task to measure the level of face and scene-selective activity on each trial.

In previous work, we had found that motivation enhances sensory representations in the occipitotemporal cortex — for a given image, face-selective activity was higher when participants were motivated to see more face, while scene-selective activity was higher when participants were motivated to see more scene (Leong et al., 2019). Here, we test if this motivational enhancement was dependent on amygdala activity. Given our earlier analyses indicating that amygdala activation preceded image onset, we focused on amygdala activity at the time of the cue (i.e. when participants were informed which category would win them more money). As the amygdala is known to respond to faces (Rutishauser et al., 2015; Visser et al., 2021), extracting the amygdala response to the cue also allowed us to isolate amygdala activation that is not contaminated by the response to the face-scene composite images.

We modeled each trial as a separate regressor in a GLM. Each regressor was created by convolving an impulse response at the time of cue onset with the HRF. The resulting beta weights were then averaged within the amygdala ROI to obtain amygdala activation at the time of the cue

for each trial. We then used a linear mixed effects model to examine the relationship between category-selective activity, motivation and amygdala activity. There was a significant interaction between amygdala activity and motivation on category-selective activity in the occipitotemporal cortex ($t(37) = 2.77$, $P < 0.009$, $b = 0.039$, $SE = 0.014$). To unpack this interaction, we divided each participant's data into trials with high and low amygdala activity. Motivation enhanced category-selective activity on trials with high amygdala activity ($t(30) = 2.22$, $P = 0.034$, $b = 0.034$, $SE = 0.015$), but not on trials with with low amygdala activity ($t(33) = 0.49$, $P = 0.63$, $b = 0.0077$, $SE = 0.016$; Fig. 3A), suggesting that motivational enhancement of sensory representations was dependent on amygdala activity.

Prior work implicates a network of frontoparietal regions in biasing perceptual decisions, including the dorsolateral prefrontal cortex (DLPFC), the dorsal anterior cingulate cortex (dACC), and the intraparietal lobule (IPL) (Domenech and Dreher, 2010; Summerfield and Koechlin, 2010; Mulder et al., 2012, 2014; White et al., 2012). In a series of exploratory analyses, we tested if these regions were associated with the motivational enhancement of sensory representations. We also considered the role of the anterior insula, which previous work had found to be involved in the accumulation of sensory evidence during perceptual decision-making (Ho et al., 2009). Activity in these regions did not moderate motivational effects on face and scene-selectively regions, suggesting that these regions were not involved in the motivational enhancement of desirable perceptual representations (DLPFC: $t(20) = -0.19$, $P = 0.852$, $b = -0.001$, $SE = 0.008$; dACC: $t(24) = 0.89$, $P = 0.385$, $b = 0.011$, $SE = 0.012$; IPL: $t(39) = 0.62$, $P = 0.551$, $b = 0.005$, $SE = 0.009$; anterior insula: $t(33) = 1.25$, $P = 0.221$, $b = 0.017$, $SE = 0.013$; Fig. 3B).

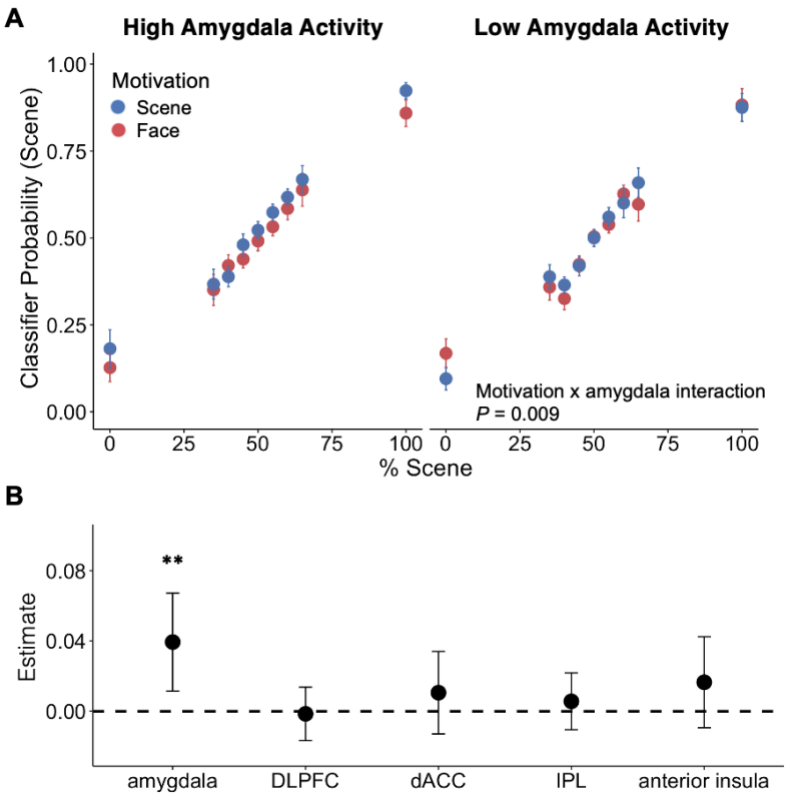

Figure 3. Amygdala activity was associated with stronger motivational enhancement of category-selective neural activity in occipitotemporal cortex. A. Classifier probability that the participant was seeing a scene rather than a face based on the pattern of BOLD activity in the occipito-temporal ROI, which we used as a proxy for category-selective activity. On trials with high amygdala activity, scene-selective activity was higher when participants were motivated to see more scene than when they were motivated to see more face. On trials with low amygdala activity, motivation had no effect on category-selective activity. Error bars indicate s.e.m. **B.** Regression coefficients of the motivation x activity interaction on classifier probability for the amygdala, DLPFC, dACC, IPL and anterior insula. The interaction was significant only in the amygdala. Error bars indicate 95% confidence intervals.

Amygdala activity was specifically associated with motivational biases in evidence accumulation

The drift diffusion model (DDM) provides a formal framework to examine the decision processes underlying perceptual decision-making (Forstmann et al., 2016). The model assumes that perceptual decisions arise from the noisy accumulation of sensory evidence from the external environment. When the level of evidence reaches a decision threshold, the corresponding decision is made. The starting point and rate of evidence accumulation, as well as the decision thresholds and time not related to decision processes (i.e. *non-decision time*, e.g., time needed to translate the decision into a motor response), are determined by model parameters that can be estimated by fitting the model to participants' data (Fig. 4A). In previous work, we had found that motivation biases perceptual decisions by (1) shifting the starting point towards the decision threshold of the desirable category, thereby reducing the amount of sensory evidence needed to make a motivation consistent response, and (2) by biasing sensory evidence in favor of the desirable category (Leong et al., 2019). To understand the amygdala's contribution to motivation biases, we assessed if either or both of these biasing mechanisms were dependent on trial-by-trial fluctuations in amygdala activity.

We took a linear regression approach to examine the effects of the amygdala on motivational biases in the starting point and rate of evidence accumulation. Specifically, the starting point (parameter z), was allowed to vary as a function of the motivation consistent category, amygdala activation at the time of the cue, and the interaction between the two (amygdala*motivation interaction; see Methods). The regression coefficient on the interaction term reflects the extent to which motivational effects on the starting point depended on amygdala activity. A positive coefficient would indicate that the starting point is more strongly biased

towards the decision threshold of the motivation consistent category on trials with higher amygdala activity. When we fit the model to participants' behavioral data, the posterior distribution of the regression coefficient of the interaction term was centered around zero ($P(z_{amygdala*motivation}>0)$ = 0.384, mean = -0.002, 95% credible interval = -0.020 to 0.018; Fig. 4B), indicating that motivational effects on the starting point did not depend on amygdala activity.

In the same model, we allowed the rate of evidence accumulation (i.e. *"drift"* rate, parameter *v*) to vary as a function of the percentage scene in the image, the motivation consistent category, amygdala activation at the time of the cue, and the interaction between the motivation consistent category and amygdala activation. Here, the regression coefficient on the interaction term reflects the extent to which motivational effects on the drift rate depended on amygdala activity. A positive coefficient would indicate a stronger bias in evidence accumulation in favor of the motivation consistent category on trials with higher amygdala activity. When we fit the model to participants' behavioral data, the posterior distribution of the regression coefficient of the interaction term indicated a high probability that the coefficient was positive ($P(v_{amygdala*motivation}>0)$ = 0.978, mean = 0.079, 95% credible interval = 0.002 to 0.160; Fig. 4B). These results suggest on trials with high amygdala activity, evidence accumulation is biased towards the decision threshold of the motivation consistent category, resulting in an increased bias towards making motivation consistent responses.

In exploratory analyses, we fit a model that included inter-trial variability parameters for the drift rate, starting point and non-decision-time, which allows for random trial-by-trial variability in model parameters. Consistent with our earlier results, we found that trial-by-trial amygdala activity moderated motivational biases in the drift rate ($P(v_{amygdala*motivation}> 0)$ = 0.997, mean = 0.138, 95% credible interval = 0.040 to 0.243) but not the starting point

($P(z_{amygdala*motivation} > 0)$ = 0.124, mean = -0.015, 95% credible interval = -0.039 to 0.011), suggesting that our results were not due to random trial-by-trial fluctuations in model parameters.

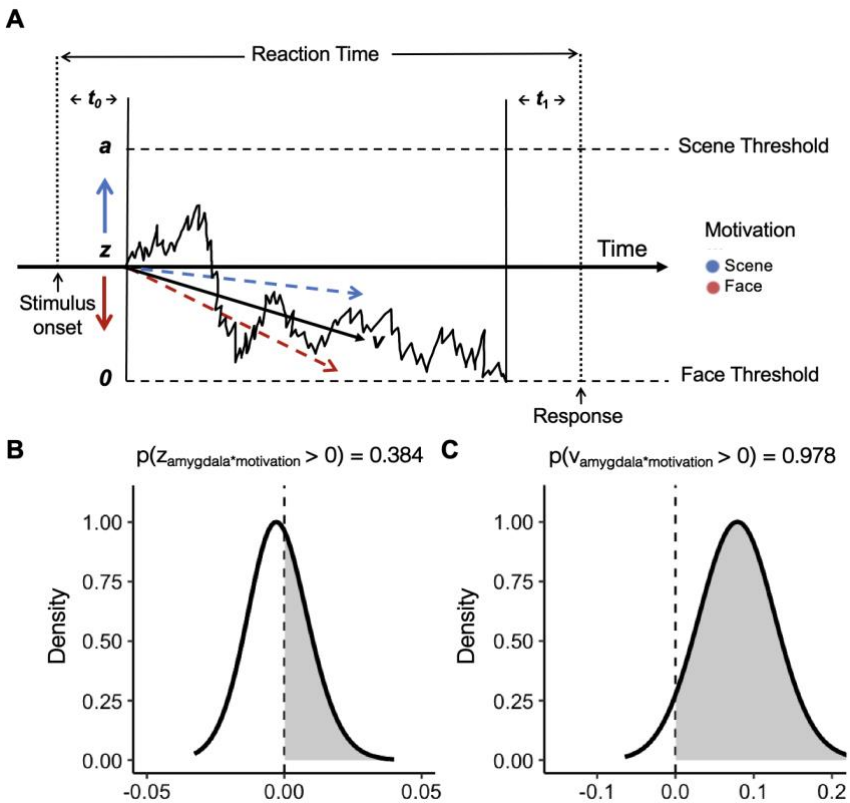

Figure 4. Modeling Results. A. Schematic diagram of the drift diffusion model (DDM). The DDM assumes that perceptual decisions arise from the stochastic accumulation of sensory evidence over time towards one of two decision thresholds. *z*: starting point, *v*: drift rate, *a*: decision threshold, t_0: sensory delay, t_1: time needed for response execution, the sum of t_0 and t_1 constitutes non-decision time. Motivation can bias perceptual decisions by biasing both the starting point and drift rate (red and blue lines). **B.** Posterior distribution of the regression coefficient of the amygdala*motivation interaction on the starting point. The distribution is centered around 0, indicating that motivational effects on the starting point did not depend on amygdala activity. **C.** Posterior distribution of the regression coefficient of the amygdala*motivation interaction on the drift rate. The distribution skews positive, indicating the motivational effects on the drift rate were dependent on amygdala activity.

Parameter	Prior	Coefficient	Estimate	Gelman-Rubin $\hat{R}$
a	$a \sim Gamma(2, 0.3)$	a	2.264[2.166, 2.366]	1.000
t	$t \sim N(0.5, 0.15)$	t	0.650[0.609, 0.694]	1.000
z	$z = invlogit(N(0.03, 0.04))$	$z_{intercept}$	0.486[0.468, 0.505]	1.000
		$z_{amygdala}$	0.007[-0.011, 0.026]	1.001
		$z_{motivation}$	0.022[0.007, 0.037]	1.000
		$z_{amygdala*motivation}$	-0.002[-0.020, 0.018]	1.002
v	$v = N(0.06, 0.2)$	$v_{intercept}$	0.191[0.065, 0.314]	1.000
		$v_{stimulus}$	0.920[0.833, 1.010]	1.000
		$v_{amygdala}$	0.087[-0.003, 0.179]	1.000
		$v_{motivation}$	0.085[0.002, 0.167]	1.000
		$v_{amygdala*motivation}$	0.081[0.002, 0.160]	1.003

Extended Table 4.1. Priors, posterior estimates, and model convergence metrics of model parameters. Estimates indicate the posterior means, with square brackets denoting 95% credible intervals. Model convergence was assessed using the Gelman-Rubin $\hat{R}$ statistic. $\hat{R}$ of all parameters were less than 1.1, suggesting that there were no issues with model convergence. *a*: decision threshold, *t*: non-decision time, z: starting point; *v*: drift rate, Gamma: a gamma distribution parameterized by mean and standard deviation, *N*: denotes a normal distribution parameterized by mean and standard deviation, invlogit: inverse logit transform.

Discussion

Psychophysical studies have previously shown that stimulus features that have been repeatedly paired with reward tend to "capture" attention and are preferentially processed by the visual system (Anderson et al., 2011; Le Pelley et al., 2015). Relatedly, fMRI studies have found that reward value modulates sensory-related neural activity, including activity in both early visual areas (Serences, 2008; Anderson, 2017) and activity in the object-selective ventral temporal cortex (Hickey and Peelen, 2015; Barbaro et al., 2017; Leong, Radulescu et al., 2017). These findings highlight the pervasive influence of reward on sensory representations in the brain. Building on this body of work, we recently demonstrated that using financial incentives to motivate participants to see a particular interpretation of an ambiguous image enhances neural activity encoding stimulus features associated with said interpretation (Leong et al., 2019). This selective enhancement biases participants towards seeing perceptual outcomes that they are motivated to see. The neural mechanisms driving the selective enhancement of desirable representations, however, remain unclear.

In the current work, we show that the motivational enhancement of desirable percepts is dependent on amygdala activity. In particular, motivation enhanced neural activity selective to desirable perceptual features, but only on trials with high amygdala activity. Heightened amygdala activity preceded participants categorizing an image as the category they were motivated to see, and participants with stronger amygdala activation also exhibited strong motivational biases in their perceptual categorizations. To understand the decision processes affected by the amygdala, we fit a DDM to participants' data and trial-by-trial amygdala activity. The model fits suggest that increased amygdala activity biases sensory evidence accumulation in favor of desirable percepts.

Together, the current findings provide converging behavioral, neural and modeling evidence of the amygdala's role in biasing perceptual decisions towards desirable perceptual outcomes.

Historically, the amygdala has been associated with the processing of fear and threat-related stimuli (LeDoux, 2000, c.f. Visser et al., 2021). More recent perspectives have emphasized a broader role of the amygdala in processing motivationally salient information, including information relevant for predicting and obtaining rewards (Pessoa and Adolphs, 2010; Pourtois et al., 2013). For example, activity in the amygdala has been shown to encode the appetitive value of stimuli (Gottfried et al., 2003; Jenison et al., 2011; Lichtenberg et al., 2017; Shanahan et al., 2021). In a series of studies with non-human primates, Peck and colleagues have also found that amygdala activity tracks spatial attention towards stimuli associated with reward (Peck et al., 2013; Peck and Salzman, 2014a, 2014b). Our findings are consistent with this broader view of amygdala function that considers its role in processing reward-relevant information. Notably, our results suggest that the amygdala is not only sensitive to the value of reward-relevant stimuli, but also flexibly enhances the neural representation of what is valuable at a given moment, which in turn gives rise to a perceptual bias towards percepts that are desirable.

The amygdala can enhance sensory representations both via direct connections to sensory areas (Freese and Amaral, 2005; Gschwind et al., 2012), as well as via indirect pathways involving frontoparietal regions and brainstem neuromodulatory systems (Lim et al., 2009; Pourtois et al., 2013). Recent theories have proposed that the amygdala interacts with the locus-coeruleus norepinephrine (LC-NE) system to selectively amplify the representation of physically and affectively salient stimuli (Markovic et al., 2014; Mather et al., 2016). According to these theories, the amygdala enhances cortical processing of salient stimuli through its direct projections to sensory cortices. At the same time, the amygdala recruits the LC which then releases NE into the

cortex. NE released in sensory areas increase the gain of sensory neurons selective to salient stimuli, such that they are more likely to fire to target inputs (Berridge and Waterhouse, 2003). In addition, NE released in the amygdala further enhances the saliency signal, resulting in even greater amplification of the neural response to salient stimuli (Gallagher and Holland, 1994).

Amygdala activity has indeed been found to be associated with reward-driven changes in pupil dilation (Watanabe et al., 2019), a known correlate of LC-NE activity (Murphy et al., 2014; Joshi et al., 2016). Using a variant of the current task, we recently demonstrated that motivational biases in sensory evidence accumulation were associated with increased pupil dilation (Leong et al., 2021). Thus, there is reason to hypothesize that the amygdala-dependent motivational enhancement of sensory representations observed in the current study were mediated by the LC-NE system. In particular, wanting to see a percept could have imbued affective salience to the associated perceptual features. The amygdala then amplifies the neural response to these features through coordinated action with the LC-NE system. The small size and location of the LC present challenges to imaging the LC using fMRI (Liu et al., 2017). Recent studies have been able to take advantage of novel approaches to functionally and anatomically localize the LC (de Gee et al., 2017; Grueschow et al., 2021). Future work can use these approaches to simultaneously measure activity in the amygdala, LC and sensory regions to further characterize the neural circuitry underlying motivational biases in perceptual decisions.

Prior work examining reward-related biases in perceptual decision-making have focused on the role of frontoparietal regions (Mulder et al., 2014). Here, we did not find evidence of the involvement of frontoparietal regions in enhancing desirable sensory evidence, suggesting that our results arose from distinct biasing mechanisms. We note that in earlier studies that assign asymmetric rewards to perceptual options, participants would earn the larger reward only if they

correctly categorized a stimulus as the category associated with larger reward (Summerfield and Koechlin, 2010; Mulder et al., 2012). In such settings, participants would earn more reward over the course of the experiment if they biased their responses towards the category associated with larger reward (Bogacz et al., 2006; Feng et al., 2009; Fan et al., 2018). Thus, reward-related biases in these earlier experiments have largely been interpreted as a shift in response strategy to maximize reward on the task and are not necessarily related to biases in sensory evidence accumulation. Consistent with this idea, these studies have found that reward biases perceptual responses without a corresponding effect on sensory representations.

In our task, however, the additional reward associated with the motivation consistent category was not contingent on participants' responses - participants won or lost the bonus based on the objective category of the image, regardless of their subjective categorizations. Hence, the bias observed in our task was unlikely to have resulted from a shift in response strategy to maximize reward. Instead, it was likely driven by participants wanting to see one percept over another. Using a drift diffusion model, we had previously found that this motivational bias reflects both an *a priori* response bias towards outcomes participants want to see (i.e. a bias in the starting point of evidence accumulation), and a perceptual bias in how sensory evidence accumulates over time (i.e. a bias in the drift rate) (Leong et al., 2019). The current results indicate that amygdala activity was specifically associated with the perceptual bias and not the response bias. Thus, the amygdala-dependent biases in perceptual decisions are potentially distinct from the response-related biases associated with frontoparietal regions documented in earlier studies.

In everyday life, people are rarely neutral observers indifferent to different perceptual outcomes. They are motivated to see some outcomes over others, and exhibit a bias to report seeing those outcomes over possible alternatives. The current work suggests that this bias is associated

with the amygdala-dependent enhancement of desirable neural representations in sensory cortices. Together, these results expand our understanding of the neural mechanisms underlying motivational influences on perceptual decision-making. Notably, similar motivational biases have been documented across a broad array of decision-making contexts (Hughes and Zaki, 2015). In one recent study, participants were asked to judge whether they were in a desirable state (i.e. an environment with greater rewards than loses) or an undesirable state (i.e. an environment with low probability of reward) by sampling evidence that was probabilistically associated with each state (Gesiarz et al., 2019). The authors found that participants accumulated evidence in a manner that was biased towards the belief they were in a desirable state, much like the bias towards desirable visual percepts observed in our study. Could the amygdala be involved in biases in evidence accumulation beyond sensory perception? Given that the amygdala has also been implicated in value-based learning and decision-making (Jenison et al., 2011; Lichtenberg et al., 2017; Prévost et al., 2013), this is certainly possible. We believe that exploring the role of the amygdala in mediating motivational biases in other reasoning and evaluative processes would be a fruitful direction for future research.

References

Anderson BA (2017) Reward processing in the value-driven attention network: reward signals tracking cue identity and location. Social Cognitive and Affective Neuroscience 12:461–467.

Anderson BA, Laurent PA, Yantis S (2011) Value-driven attentional capture. PNAS 108:10367–10371.

Barbaro L, Peelen MV, Hickey C (2017) Valence, Not Utility, Underlies Reward-Driven Prioritization in Human Vision. J Neurosci 37:10438–10450.

Berridge CW, Waterhouse BD (2003) The locus coeruleus–noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. Brain Research Reviews 42:33–84.

Boehm U, Annis J, Frank MJ, Hawkins GE, Heathcote A, Kellen D, Krypotos A-M, Lerche V, Logan GD, Palmeri TJ, van Ravenzwaaij D, Servant M, Singmann H, Starns JJ, Voss A, Wiecki TV, Matzke D, Wagenmakers E-J (2018) Estimating across-trial variability parameters of the Diffusion Decision Model: Expert advice and recommendations. Journal of Mathematical Psychology 87:46–75.

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks. Psychological Review 113:700–765.

Brainard DH (1997) The Psychophysics Toolbox. Spatial Vis 10:433–436.

Bromberg-Martin ES, Sharot T (2020) The Value of Beliefs. Neuron 106:561–565.

de Gee JW, Colizoli O, Kloosterman NA, Knapen T, Nieuwenhuis S, Donner TH (2017) Dynamic modulation of decision biases by brainstem arousal systems. eLife 6:e23232.

Domenech P, Dreher J-C (2010) Decision Threshold Modulation in the Human Brain. J Neurosci 30:14305–14317.

Dunning D, Balcetis E (2013) Wishful Seeing: How Preferences Shape Visual Perception. Curr Dir Psychol Sci 22:33–37.

Fan Y, Gold JI, Ding L (2018) Ongoing, rational calibration of reward-driven perceptual biases. eLife 7:e36018.

Feng S, Holmes P, Rorie A, Newsome WT (2009) Can Monkeys Choose Optimally When Faced with Noisy Stimuli and Unequal Rewards? PLOS Computational Biology 5:e1000284.

Forstmann BU, Ratcliff R, Wagenmakers E-J (2016) Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions. Annual Review of Psychology 67:641–666.

Freese JL, Amaral DG (2005) The organization of projections from the amygdala to visual cortical areas TE and V1 in the macaque monkey. Journal of Comparative Neurology 486:295–317.

Gallagher M, Holland PC (1994) The amygdala complex: multiple roles in associative learning and attention. PNAS 91:11771–11776.

Gesiarz F, Cahill D, Sharot T (2019) Evidence accumulation is biased by motivation: A computational account. PLOS Computational Biology 15:e1007089.

Gottfried JA, O'Doherty J, Dolan RJ (2003) Encoding Predictive Reward Value in Human Amygdala and Orbitofrontal Cortex. Science 301:1104–1107.

Grill-Spector K (2003) The neural basis of object perception. Current Opinion in Neurobiology 13:159–166.

Grueschow M, Stenz N, Thörn H, Ehlert U, Breckwoldt J, Brodmann Maeder M, Exadaktylos AK, Bingisser R, Ruff CC, Kleim B (2021) Real-world stress resilience is associated with the responsivity of the locus coeruleus. Nat Commun 12:2275.

Gschwind M, Pourtois G, Schwartz S, Van De Ville D, Vuilleumier P (2012) White-Matter Connectivity between Face-Responsive Regions in the Human Brain. Cerebral Cortex 22:1564–1576.

Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. Nature Reviews Neuroscience 9:467–479.

Hickey C, Peelen MV (2015) Neural Mechanisms of Incentive Salience in Naturalistic Human Vision. Neuron 85:512–518.

Ho TC, Brown S, Serences JT (2009) Domain General Mechanisms of Perceptual Decision Making in Human Cortex. J Neurosci 29:8675–8687.

Hughes BL, Zaki J (2015) The neuroscience of motivated cognition. Trends in Cognitive Sciences 19:62–64.

Jenison RL, Rangel A, Oya H, Kawasaki H, Howard MA (2011) Value Encoding in Single Neurons in the Human Amygdala during Decision Making. J Neurosci 31:331–338.

Joshi S, Li Y, Kalwani RM, Gold JI (2016) Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. Neuron 89:221–234.

Kube T, Rozenkrantz L (2021) When Beliefs Face Reality: An Integrative Review of Belief Updating in Mental Health and Illness. Perspect Psychol Sci 16:247–274.

Le Pelley ME, Pearson D, Griffiths O, Beesley T (2015) When goals conflict with values: Counterproductive attentional and oculomotor capture by reward-related stimuli. Journal of Experimental Psychology: General 144:158–171.

LeDoux JE (2000) Emotion Circuits in the Brain. Annu Rev Neurosci 23:155–184.

Leong YC, Dziembaj R, D'Esposito M (2021) Pupil-Linked Arousal Biases Evidence Accumulation Toward Desirable Percepts During Perceptual Decision-Making. Psychol Sci 32:1494-1509

Leong YC, Hughes BL, Wang Y, Zaki J (2019) Neurocomputational mechanisms underlying motivated seeing. Nature Human Behaviour 3:962–973.

Leong YC, Radulescu A, Daniel R, DeWoskin V, Niv Y (2017) Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments. Neuron 93:451–463.

Lichtenberg NT, Pennington ZT, Holley SM, Greenfield VY, Cepeda C, Levine MS & Wassum, KM (2017). Basolateral amygdala to orbitofrontal cortex projections enable cue-triggered reward expectations. J Neurosci, 37:8374-8384.

Lim S-L, Padmala S, Pessoa L (2009) Segregating the significant from the mundane on a moment-to-moment basis via direct and indirect amygdala contributions. PNAS 106:16841–16846.

Liu KY, Marijatta F, Hämmerer D, Acosta-Cabronero J, Düzel E, Howard RJ (2017) Magnetic resonance imaging of the human locus coeruleus: A systematic review. Neuroscience & Biobehavioral Reviews 83:325–355.

Ma DS, Correll J, Wittenbrink B (2015) The Chicago face database: A free stimulus set of faces and norming data. Behav Res 47:1122–1135.

Markovic J, Anderson AK, Todd RM (2014) Tuning to the significant: Neural and genetic processes underlying affective enhancement of visual perception and memory. Behavioural Brain Research 259:229–241.

Mather M, Clewett D, Sakaki M, Harley CW (2016) Norepinephrine ignites local hotspots of neuronal excitation: How arousal amplifies selectivity in perception and memory. Behavioral and Brain Sciences 39

Mulder MJ, van Maanen L, Forstmann BU (2014) Perceptual decision neurosciences – A model-based review. Neuroscience 277:872–884.

Mulder MJ, Wagenmakers E-J, Ratcliff R, Boekel W, Forstmann BU (2012) Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. J Neurosci 32:2335–2343.

Mumford JA, Davis T, Poldrack RA (2014) The impact of study design on pattern estimation for single-trial multivariate pattern analysis. NeuroImage 103:130–138.

Murphy PR, O'Connell RG, O'Sullivan M, Robertson IH, Balsters JH (2014) Pupil diameter covaries with BOLD activity in human locus coeruleus. Human Brain Mapping 35:4140–4154.

Peck CJ, Lau B, Salzman CD (2013) The primate amygdala combines information about space and value. Nat Neurosci 16:340–348.

Peck CJ, Salzman CD (2014a) The Amygdala and Basal Forebrain as a Pathway for Motivationally Guided Attention. J Neurosci 34:13757–13767.

Peck CJ, Salzman CD (2014b) Amygdala neural activity reflects spatial attention towards stimuli promising reward or threatening punishment. eLife 3:e04478.

Pessoa L, Adolphs R (2010) Emotion processing and the amygdala: from a "low road" to "many roads" of evaluating biological significance. Nature Reviews Neuroscience 11:773–782.

Pourtois G, Schettino A, Vuilleumier P (2013) Brain mechanisms for emotional influences on perception and attention: What is magic and what is not. Biological Psychology 92:492–512.

Prévost C, McNamee D, Jessup RK, Bossaerts P, O'Doherty JP (2013) Evidence for Model-based Computations in the Human Amygdala during Pavlovian Conditioning. PLOS Computational Biology 9:e1002918.

Rossi-Goldthorpe RA, Leong YC, Leptourgos P, Corlett PR (2021) Paranoia, self-deception and overconfidence. PLOS Computational Biology 17:e1009453.

Rutishauser U, Mamelak AN, Adolphs R (2015) The primate amygdala in social perception – insights from electrophysiological recordings and stimulation. Trends in Neurosciences 38:295–306.

Serences JT (2008) Value-Based Modulations in Human Visual Cortex. Neuron 60:1169–1181.

Shadlen MN, Kiani R (2013) Decision Making as a Window on Cognition. Neuron 80:791–806.

Shanahan LK, Bhutani S, Kahnt T (2021) Olfactory perceptual decision-making is biased by motivational state. PLOS Biology 19:e3001374.

Shirer WR, Ryali S, Rykhlevskaia E, Menon V, Greicius MD (2012) Decoding subject-driven cognitive states with whole-brain connectivity patterns. Cereb Cortex 22:158–165.

Summerfield C, Koechlin E (2010) Economic Value Biases Uncertain Perceptual Choices in the Parietal and Prefrontal Cortices. Frontiers in Human Neuroscience 4:208.

Venables WN, Ripley BD (2013) Modern Applied Statistics with S-PLUS. Springer Science & Business Media.

Visser RM, Bathelt J, Scholte HS, Kindt M (2021) Robust BOLD responses to faces but not to conditioned threat: challenging the amygdala's reputation in human fear and extinction learning. J Neurosci.

Voss A, Rothermund K, Brandtstädter J (2008) Interpreting ambiguous stimuli: Separating perceptual and judgmental biases. Journal of Experimental Social Psychology 44:1048–1056.

Watanabe N, Bhanji JP, Ohira H, Delgado MR (2019) Reward-Driven Arousal Impacts Preparation to Perform a Task via Amygdala–Caudate Mechanisms. Cerebral Cortex 29:3010–3022.

White CN, Mumford JA, Poldrack RA (2012) Perceptual Criteria in the Human Brain. J Neurosci 32:16716–16724.

Wiecki TV, Sofer I, Frank MJ (2013) HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. Front Neuroinform 7:14