



Published in final edited form as:

J Neurosci. 2013 May 8; 33(19): 8541–8548. doi:10.1523/JNEUROSCI.5754-12.2013.

Frontal Theta Overrides Pavlovian Learning Biases

James F. Cavanagh¹, Ian Eisenberg¹, Marc Guitart-Masip^{2,4}, Quentin Huys^{3,4}, and Michael J. Frank^{1,5,6}

¹Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI, USA ²Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, London, UK ³Translational Neuroimaging Unit, Department of Biological Engineering, ETH and University of Zurich, Switzerland ⁴Department of Psychiatry, Psychotherapy and Psychosomatics, University Hospital of Psychiatry Zurich, Switzerland ⁵Department of Psychiatry, Brown University, Providence, RI, USA ⁶Brown Institute for Brain Science, Brown University, Providence, RI, USA

Abstract

Pavlovian biases influence learning and decision making by intricately coupling reward seeking with action invigoration and punishment avoidance with action suppression. This bias is not always adaptive; it can oftentimes interfere with instrumental requirements. The prefrontal cortex is thought to help resolve such conflict between motivational systems, but the nature of this control process remains unknown. EEG recordings of mid-frontal theta band power are sensitive to conflict and predictive of adaptive control over behavior, but it is not clear whether this signal would reflect control over conflict between motivational systems. Here we utilized a task that orthogonalized action requirements and outcome valence while recording concurrent EEG in human participants. By applying a computational model of task performance, we derived parameters reflective of the latent influence of Pavlovian bias and how it was modulated by mid-frontal theta power during motivational conflict. Between subjects, individuals who performed better under Pavlovian conflict exhibited higher mid-frontal theta power. Within subjects, trial-to-trial variance in theta power was predictive of ability to overcome the influence of the Pavlovian bias, and this effect was most pronounced in individuals with higher mid-frontal theta to conflict. These findings demonstrate that mid-frontal theta is not only a sensitive index of prefrontal control, but it can also reflect the application of top-down control over instrumental processes.

Our prefrontal cortices allow us to deliberately overcome habitual or prepotent biases. Yet some biases can exist even in novel or seemingly impartial situations. For example, innate Pavlovian biases facilitate reward-induced vigor and punishment-induced inhibition, yet these biases can actually hinder learning in instrumental conditions (Guitart-Masip et al., 2011, 2012b). Here we asked whether prefrontal cortex can detect and resolve such conflict between separate motivational systems.

Motivated action selection is informed by at least three major processes that learn to associate stimuli with responses. Learning is primarily associated with two different instrumental processes. One which reinforces rewarded actions and suppresses punished actions, resulting in stimulus-response behavior (Frank, 2005). A second, more sophisticated prefrontal cortical operation involves an understanding of the consequences of actions, and leads to goal-directed choices (Hampton et al., 2006). These systems are respectively

referred to as *model-free* and *model-based*, referring to absence or presence of a ‘model’ of the environment. Logically, organisms with these systems should be able to learn any type of stimulus-action pairing, yet in reality they can fail in reliable and predictable ways. This failure can be due to a third process: influence of hard-coded Pavlovian responses to particular outcome expectations. Under Pavlovian influence, action selection and outcome valence are intricately interwoven.

While in many circumstances Pavlovian influences may facilitate instrumental learning, they can also impair it by stifling the pairing of conflicting action and valence requirements (Dayan and Balleine, 2002). This is most clearly shown with Hershberger’s (1986) chickens, which were unable to learn to move away from food in order to obtain it, and Holland’s (1979) rats, which approached a light predictive of food even when that led to food omission. In these cases, Pavlovian biases did not simply interfere with correct performance; they actually led to the opposite behavior than the task required.

Nevertheless, unlike chickens, people can eventually learn conflicting instrumental contingencies. Whereas previous research has highlighted the neural mechanisms of Pavlovian influences on instrumental performance (Cardinal et al., 2002; Talmi et al., 2008; Prévost et al., 2012), little is known about the systems that overcome Pavlovian conflict. To investigate this topic, we utilized mid-frontal electroencephalographic (EEG) signals associated with conflict, learning, and control. Mid-frontal theta-band power has been associated with the transient application of cognitive control to prevent impulsive responses (Cohen et al., 2009; Cavanagh et al., 2011). Thus, frontal theta may reflect the recruitment of control, but only if a model-based system recognizes the conflict between motivational systems.

However, other evidence indicates that mid-frontal theta covaries with reinforcement prediction error signals and can be predictive of subsequent learning and behavioral adjustments (Cavanagh et al., 2010, 2012a; van de Vijver et al., 2011). Thus this EEG signal appears to be influenced by both control (which could counter Pavlovian conflict) and value (which would covary with Pavlovian biases). By orthogonalizing action and valence, the current study was able to dissociate the influence of model-free (reflecting Pavlovian bias) and model-based systems (reflecting control over Pavlovian conflict) as indicated by frontal theta.

Materials and Methods

Participants

A total of 34 adults were recruited from the Brown University undergraduate subject pool and Providence community to complete the experiment (20 male; age $M=24$, $SD=4.44$, range=18–34). All participants had normal or corrected-to-normal vision, no history of neurological, psychiatric, or any other relevant medical problem, and were free from current psychoactive medication use. Participants were compensated \$15 for completing the task, with a monetary bonus depending on total accuracy levels (\$5 for achieving over 50% total accuracy and \$10 for achieving over 66% accuracy).

Task: Learning

The task was adapted from Guitart-Masip et al. (2011). Each trial consisted of three events: a cue, a target detection task, and an outcome (Fig. 1a). Before the experiment, participants were informed that each cue would either lead to reward or to punishment based on their response, and that no cue would lead to both. Participants were encouraged to explore both response options to best learn how to achieve the best outcome from each cue: they understood that they could either respond to the target (‘go’) or they could withhold a

response ('nogo'). There were four different cues that predicted unique optimal combinations of action and outcome: Go-to-Win, Go-To-Avoid, NoGo-To-Win, and NoGo-to-Avoid, each was presented 40 times across two blocks with a break between.

Trials began with the display of a colored shape cue (1000 ms). Each cue was 70% predictive of the correct action to take (whether to go or no-go) to gain the optimal outcome (reward or avoid punishment), whereas 30% or incorrect actions were reinforced. After a variable interval (250–2500 ms), the target detection stimulus appeared, which consisted of a white circle in the middle of the screen indicating the subject could respond (go) or not (no-go). The circle disappeared if participants pressed the button, or after 1000 ms. At 1000ms following the offset of the circle, feedback was presented (2000 ms) indicating reward (green +\$), punishment (red -\$), or neutral (a yellow bar) which was used to indicate either no reward or punishment avoidance, depending on the condition. The inter-trial interval consisted of a fixation cross for a variable interval (750–1500 ms).

In addition to these four probabilistic valenced conditions, there were two deterministic neutral conditions. Deterministic trials consisted of pictures of a hand indicating whether or not to press the response button. On these trials, participants were explicitly informed how to respond to target detection task and they were informed that there would be no outcome for these actions. These deterministic trials were intended to be used as motor action and inhibition contrasts, but they were not used for the present analysis. Note that the task structure was shorter (40 vs. 60 trials per condition) and harder (70% vs. 80% reinforcing) than previous versions of this learning task (Guitart-Masip et al., 2012b).

Learning Performance Measures

Individual differences in performance styles were defined in relation to the hardest NoGo-to-Win condition in the second time block. If subjects successfully inhibited action on more than 65% of trials in this second block, they were considered 'Learners' (N=17); the rest of the subjects were labeled 'Non-Learners' (N=17). This categorical assignment facilitates an intuitive display of performance patterns, but we utilized continuous measures of learning for all important statistical analyses.

To condense the specific types of Pavlovian biases underlying performance differences across conditions, we devised the following reinforcement responsiveness metrics. To summarize across all blocks, measures of reward-based invigoration ($(\text{Go on Go-to-Win} + \text{NoGo-to-Win}) / \text{Total Go}$) and punishment-based suppression ($(\text{NoGo on Go-to-Avoid} + \text{NoGo-to-Avoid}) / \text{Total NoGo}$) were averaged into a single measure of Pavlovian Performance Bias, see Fig 2b. While approach and avoidance conditions are not motivationally identical, this measure effectively merges comparable instrumental outcome-action adaptations. Thus if a participant were to learn the conditions perfectly, their Pavlovian Performance Bias measure would be 50% since half of the conditions facilitates and the other half contradicts Pavlovian response styles. Higher scores on the measure of Pavlovian Performance Bias therefore reflect a greater dependence on Pavlovian biases during decision making.

Task: Post-Learning Transfer Phase

After the learning phase, participants also completed a novel forced-choice transfer phase following the task (not included in earlier studies using this task, but similar to the transfer phase of Frank et al., 2004, see Fig. 1b). Data from this transfer phase was not available for one participant (a Learner) who had to leave early. In this transfer phase, each of the predictive cues was paired with each of the others in a two alternative forced choice scenario, and participants were told to select which cue was "more rewarding". No feedback

was presented, and each pairing was presented eight times. These choices were used to indicate relative preferences/valuations of the different cues to indicate their learned values outside of the instrumental learning environment. We reasoned that choice preferences in this phase would be indicative of Pavlovian biases in learned value, such that participants may assign a higher reward value for Go-to-Win than NoGo-to-Win cues. This hypothesized pattern of choices would reveal whether participants who do eventually learn the conflicting instrumental contingencies (e.g. Learners) nevertheless exhibit Pavlovian influences on value in their inherent preferences. Such a Pavlovian influence over value-related forced choice may reveal whether the source of individual differences in Pavlovian biases resides in: 1) the mechanisms giving rise to the bias itself (Learners – Non-Learners), or 2) whether individuals exhibit similar bias mechanisms but simply override them in the task conditions for which they are detrimental to performance (Learners = Non-Learners).

EEG Recording and Preprocessing

EEG was recorded using a 128 channel EGI system. EEG was recorded continuously with hardware filters set from .1 to 100 Hz, a sampling rate of 250 Hz, and an online vertex reference. Continuous EEG was epoched around the cues (-1500 ms to 5500 ms). Data were then visually inspected to identify bad channels to be interpolated and bad epochs to be rejected. Blinks were removed using independent component analysis from EEGLab (Delorme and Makeig, 2004). The vertex site was reconstructed; data were then converted to Current Source Density (CSD)(Kayser and Tenke, 2006). Broad band ERPs were filtered from .5 to 20 Hz.

Time-frequency calculations were computed using custom-written Matlab routines (Cavanagh et al., 2009). For condition-specific activities, time-frequency measures were computed by multiplying the fast Fourier transformed (FFT) power spectrum of single trial EEG data with the FFT power spectrum of a set of complex Morlet wavelets (defined as a Gaussian-windowed complex sine wave: $e^{i2\pi ft} e^{-t^2/(2\sigma^2)}$, where t is time, f is frequency (which increased from 1 to 50Hz in 50 logarithmically spaced steps), and defines the width (or “cycles”) of each frequency band, set according to $4/(2\pi f)$), and taking the inverse FFT. The end result of this process is identical to time-domain signal convolution, and it resulted estimates of instantaneous power (the magnitude of the analytic signal), defined as $Z[t]$ (power time series: $p(t) = \text{real}[z(t)]^2 + \text{imag}[z(t)]^2$). Each epoch was then cut in length (-500 to +1000 ms). Power was normalized by conversion to a decibel (dB) scale ($10 \cdot \log_{10}[\text{power}(t)/\text{power}(\text{baseline})]$), allowing a direct comparison of effects across frequency bands. For trial-to-trial analyses, EEG data was filtered from 4 to 8 Hz and Hilbert transformed to derive the single trial theta power envelopes. The baseline for each frequency consisted of the average power from 300 to 200 ms prior to the onset of the cues.

Based on previous literature (Cavanagh et al., 2012b), the stimulus-locked theta band power burst over mid-frontal sites (4–8 Hz, 175–350 ms) was *a priori* hypothesized to be the Region of Interest (ROI) involved in conflict and control. To verify this temporal, frequency, and spatial ROI using data-driven statistical tests, non-parametric Spearman’s correlations of time-frequency space were used with behavioral or model-based parameters. To diminish the influence of outliers, trial-by-trial theta power values were sigmoid transformed prior to use in the computational model.

Computational Modeling

An existing model of this task was utilized and refined to examine latent parameters thought to underlie individual differences in behavioral performance (Guitart-Masip et al., 2012b) and the degree to which these parameters were modified as a function of frontal theta. As in that study, models with increasing complexity (here labeled M1–M5) were assessed to

determine whether they capture additional variance and provide better fits to the data (penalizing for additional complexity). Here we implement the novel advancement of investigating the influence of trial-by-trial theta power on action selection using three competing models (M6a,b,c).

In all models, action values were estimated for each condition and a softmax choice function was used to predict the most likely action on each trial. The simplest model (M1) included two free parameters for scaling feedback sensitivity (ρ) and learning rate (ϵ). Reinforcements (r) took the form of $(-1,0,1)$ depending on the condition,

$$r \in (-1, 0, 1)$$

State-action values (Q values) were updated according to the delta learning rule with feedback sensitivity (ρ) scaling the reinforcement value and the learning rate (ϵ) scaling the update term.

$$Q_t(a_t|s_t) = Q_{t-1}(a_t|s_t) + \epsilon (\rho * r_t - Q_{t-1}(a_t|s_t))$$

Ensuing models included sequential additions, beginning with a third parameter in M2 to allow for irreducible noise (ξ) in action selection (to account for the possibility that some proportion of trials were not selected according to the model).

$$p(a|s) = \xi * \left[\frac{\exp(Q(a_t|s_t))}{\sum \exp(Q(q|s))} \right] + \frac{1-\xi}{2}$$

a fourth parameter of an overall bias to ‘go’ (b), regardless of valence, in M3,

$$Q_t(Go|s_t) = Q_t(Go|s_t) + b$$

a fifth parameter that allowed for potentially different sensitivities to reward vs. punishment (ρ_{rew} and ρ_{pun}) in M4. The critical sixth parameter in M5 specified the Pavlovian bias (π), which was the degree to which behavior is invigorated in response to stimuli that had positive learned value and is suppressed in response to stimuli that had negative learned value. In this model, the value V of each stimulus is learned as a function of reward history, then added to bias the action value $Q(Go)$ in proportion to the Pavlovian bias:

$$\begin{aligned} V_t(S_t) &= V_{t-1}(S_t) + \epsilon (\rho * r - V_{t-1}(S_t)) \\ Q_t(Go|s_t) &= Q_t(Go|s_t) + b + \pi * V_t(S_t) \end{aligned}$$

Finally, to investigate whether mid-frontal theta mitigated against Pavlovian bias, we investigated whether an Effect of Theta parameter (β) effectively weighted the trial-by-trial EEG theta power (Θ_t) to alter the balance between the instrumental controller (Q) and the Pavlovian controller (V). Previous studies have indicated that the influence of mid-frontal theta on cognitive control is primarily evident in conflict trials (Cavanagh et al., 2011), therefore this modulation was only modeled during conflict trials. We tested three models of this influence: M6a-c.

M6a tested if there was evidence for direct modulation of the Pavlovian influence (V) by theta power:

$$Q_t(Go|s_t) = \begin{cases} Q_t(Go|s_t) + b + (\pi + \beta * \theta_t) * V_t(S_t) & \text{if conflict} \\ Q_t(Go|s_t) + b + \pi * V_t(S_t) & \text{otherwise} \end{cases}$$

M6b tested if there was evidence for a direct modulation of the instrumental contribution (Q) by theta power:

$$Q_t(Go|s_t) = \begin{cases} (1 - \beta * \theta_t) * Q_t(Go|s_t) + \pi * V_t(S_t) + b & \text{if conflict} \\ Q_t(Go|s_t) + b + \pi * V_t(S_t) & \text{otherwise} \end{cases}$$

$$Q_t(NoGo|s_t) = \begin{cases} (1 - \beta * \theta_t) * Q_t(NoGo|s_t) & \text{if conflict} \\ Q_t(NoGo|s_t) & \text{otherwise} \end{cases}$$

M6c tested whether there was evidence that theta power shifted control from the Pavlovian influence (V) towards an instrumental controller (Q) using an instrumental-Pavlovian tradeoff parameter w instead of the Pavlovian influence π :

$$Q_t(Go|s_t) = \begin{cases} (1 - (w + \beta * \theta_t)) * Q_t(Go|s_t) + (w + \beta * \theta_t) * V_t(S_t) + b & \text{if conflict} \\ (1 - w) * Q_t(Go|s_t) + w * V_t(S_t) + b & \text{otherwise} \end{cases}$$

$$Q_t(NoGo|s_t) = \begin{cases} (1 - (w + \beta * \theta_t)) * Q_t(NoGo|s_t) & \text{if conflict} \\ (1 - w) * Q_t(NoGo|s_t) & \text{otherwise} \end{cases}$$

In all these models, if β was positive, it would indicate that increasing theta power was associated with greater expression of Pavlovian biases, whereas if β was negative, it would indicate that theta was a marker for the relative suppression of Pavlovian biases. In M6a, this is the consequence of $\beta * \theta_t$ directly modulating the Pavlovian influence. In M6b, reductions in the expression of Pavlovian contingencies are due to $\beta * \theta_t$ strengthening the instrumental component, while in M6c they are the result of $\beta * \theta_t$ modulating the competition between the Pavlovian and the instrumental component.

As in previous publications of this model, an Expectation-Maximization (EM) procedure was used for hierarchical model estimation of group and individual subject parameters (Huys et al., 2011; Guitart-Masip et al., 2012b). EM recursively iterates model fitting to inform the group distribution for each model parameter, which is used as a prior for parameter maximization of each individual subject. Recursion finishes when consecutive iterations converge to near-identical parameter values. Model comparison used the integrated Bayesian Information Criterion (iBIC). Whereas the BIC provides an estimate of the penalized *individual-level* likelihood of the data given a set of parameters, the iBIC estimates the penalized *group-level* likelihoods across the estimated *distribution* of the group-level hyper-parameters. Lower iBIC values indicate a model that fits the data better, with a difference of 4–12 iBIC values suggesting positive evidence, 12–20 suggesting strong evidence, and above 20 units suggesting very strong evidence (Kass and Raftery, 1995). As in previous uses of this model, (Guitart-Masip et al., 2012b) feedback sensitivities and the Pavlovian bias were constrained to be positive, and learning rates, instrumental-Pavlovian tradeoff parameter w and softmax noise were constrained to be between 0 and 1. All other parameters were unconstrained.

Results

Performance

Average performance accuracies followed qualitatively similar patterns as previous studies with this task (Guitart-Masip et al., 2012b), with good performance on Go-to-Win (accuracy $M=.88$, $SD=.13$), somewhat equivalent performance on Go-to-Avoid ($M=.68$, $SD=.20$) and NoGo-to-Avoid ($M=.68$, $SD=.21$) and poorest performance on NoGo-to-Win ($M=.48$, $SD=.34$). Figure 2a shows the group averages of the individual running accuracies in each condition. On Pavlovian congruent conditions, it is clear that all participants performed well. However, there was tremendous variance in performance on Pavlovian conflict conditions. Given that performance on Go-to-Avoid and NoGo-to-Win correlated with each other ($r_{(34)}=.43$, $p=.01$) and not with the congruent conditions (all p 's $>.16$), it is clear Non-Learners were not simply idiosyncratically bad in some conditions; rather individuals appeared to have reliable tendencies to rely on Pavlovian Biases. The summary measure of Pavlovian Performance Bias (Fig 2b) differed between groups ($t_{(32)}=4.62$, $p<.001$), and as expected, correlated with NoGo-to-Win accuracy ($r_{(34)}=-.81$, $p<.001$) and Go-to-Avoid accuracy ($r_{(34)}=-.67$, $p<.001$) effectively summarizing individual differences in the reliance on Pavlovian bias during the entire task.

Critically, both groups showed evidence for some Pavlovian influence over value learning in the post-task transfer phase (Fig 2c). These findings reveal that all individuals displayed a clear *win > go > nogo > avoid* hierarchy of explicit preferences. While the preference for *win > avoid* reflects the transfer phase instructions to select the “most rewarding” stimulus, a more subtle bias of *go > nogo* was also revealed in the pattern of choices. This finding is consistent with the idea that reward prediction errors invigorate action selection. There were no significantly different patterns between groups for any condition (t 's <1.6). Thus, although Learners were able to successfully suppress Pavlovian biases during learning of the conflict conditions, they nevertheless exhibited the same choice preference for Go-to-Win over NoGo-to-Win and Go-to-Avoid over NoGo-to-Avoid as did the Non-Learners, despite the fact that these cues were similarly predictive of reward. These findings suggest that the difference between Learners and Non-Learners may not reside in the mechanisms giving rise to Pavlovian influences, but instead may reflect a differential ability to override such biases.

EEG

To investigate the influence of frontal theta power on Pavlovian conflict, we correlated the Pavlovian Performance Bias measure with the theta power difference between Pavlovian conflict and congruent conditions. Importantly, this theta power contrast is orthogonal to action and reinforcement requirements (both conflict and congruent groupings involve one condition with a go action and one with a nogo action, and one condition with rewards and one with losses), providing a relatively clean measure of EEG activities associated with Pavlovian conflict rather than action or valence *per se*. There were no significant main effects of valence (win or avoid) or action (go vs. nogo) in the ROI.

Figure 3a shows the topography of electrodes with a significant correlation between theta band power and Pavlovian Performance Bias scores. Major findings occurred in three broad electrode clusters over frontal cortex, referred to as Mid-Frontal, Right-Mid, and Left Lateral. These regions are collapsed in Figure 3b to show the pixel-wise correlation of this performance measure with spectral differences. Significant effects were observed around the core temporal and frequency range of the ROI; these effects were replicated within each cluster independently as well. The inset shows the scatterplot of theta power with performance, revealing that the non-parametric correlations were significant with ($\rho_{(34)}=$

-.55, $p < .01$) or without ($\rho_{(32)} = -.50$, $p < .01$) outliers. Notably, this ROI time range converges with the mid-frontal P2-N2 complex of ERP components (Fig 3c), which are characterized by a strong theta band spectral dynamic (Cavanagh et al., 2012b). Although average theta power differences varied across frontal clusters (Fig 3d), the correlation values were similarly strong within the ROI time windows (Fig 3e). Interestingly, these correlations were maximal in the waxing of the theta power response.

In sum, convergent spatio-temporal-frequency findings revealed that individuals with greater frontal theta power during early cue processing in response to Pavlovian conflict were less compromised by a Pavlovian Performance Bias. We next assessed whether trial-to-trial variations of theta *within* an individual were related to varying abilities to override Pavlovian biases. This question is most straightforwardly addressed in the context of the computational model fits to behavior by investigating whether trial-specific theta power from the same ROI influenced the expression of Pavlovian bias, the recruitment of instrumental contingencies, or both.

Computational Modeling

Table 1 reveals that the stepwise addition of parameters in each model M1–M5 yielded increasingly better fits as measured by iBIC, and how the novel model M6a provided the strongest improvement upon the fit of the data from next most complex model with only a Pavlovian Bias parameter (M5). Importantly, M6a had highly similar parameters to M5; only the addition of the Effect of Theta parameter provided a better fit to the data. Figure 4a/b shows that within this best model (M6a), the Pavlovian Performance Bias (used in Figs 2 & 3) was correlated with both the Pavlovian Bias parameter ($\rho_{(34)} = .60$, $p < .01$) and the Effect of Theta parameter ($\rho_{(34)} = .37$, $p < .05$). Moreover, the Pavlovian Bias parameter was highly similar between M5 and M6 ($\rho = .71$, $p < .01$), and was uncorrelated with these Effect of Theta parameter ($p = .85$), highlighting the fact that the Effect of Theta parameter accounted for unique variance in the improved model fit in M6.

The mean value for the Effect of Theta parameter in M6a was significantly negative across the entire group (Mean = $-.67$, $t_{(33)} = 5.82$, $p < .01$), implying that trial-to-trial variations in theta negatively influenced within-subject Pavlovian biases. The Effect of Theta parameter was more negative in Learners (Mean = $-.94$) than in Non-Learners (Mean = $-.40$); these groups were significantly different from each other ($t_{(32)} = 2.55$, $p < .05$). Figure 4c reveals that inter-individual differences in theta power increases in response to Pavlovian conflict (across participants) correlated with intra-individual abilities to use theta (across trials) to overcome Pavlovian biases. Table 1 also shows that alternative formulations, where theta promoted instrumental contingencies or explicitly modulated the trade-off between instrumental and Pavlovian influences, provided inferior accounts of the data.

Thus, individuals who were more likely to detect presence of Pavlovian conflict exhibited higher mid-frontal theta power when conflict was high, and were thus better at overcoming Pavlovian biases in those trials for which theta was particularly evident. The strong negative coefficient implies that on trials in which theta power was high, there was a diminished Pavlovian bias. This result converges with the findings from the behavioral transfer phase which implied that even Learners exhibit a Pavlovian bias in their valuations, but are simply able to suppress that bias when they detect that it conflicts with the instrumental requirements of the task.

Discussion

This investigation revealed that conflict-induced mid-frontal theta power is indicative of the ability to overcome Pavlovian biases when they conflict with instrumental requirements.

This effect was observed both inter- and intra-individually, where greater theta to conflict was associated with increased top-down adaptive control. These results are consistent with prior studies showing that bilateral inferior frontal gyri are involved when overcoming Pavlovian bias (Guitart-Masip et al., 2012b), and that conflict-related mid-frontal theta influences the ability to prevent impulsive responding (Cavanagh et al., 2011).

The nature of Pavlovian biases

The results reported here replicate findings of a pervasive Pavlovian influence over instrumental performance (Talmi et al., 2008; Guitart-Masip et al., 2011, 2012a), including the finding that individuals vary widely in the expression of this influence (Guitart-Masip et al., 2012b). Yet it remains unknown if some individuals simply have a diminished Pavlovian influence over behavior, or if these individuals actively overcome this bias through effort-based and goal-directed cognitive control mechanisms. Convergent evidence suggests the latter case.

All individuals demonstrated evidence for a learned coupling between action and valence (e.g. they preferred Go-to-Win over NoGo-to-Win) in the post-task transfer phase. Thus, it appears that although Learners were able to successfully suppress Pavlovian biases during learning of the conflict conditions, they nevertheless exhibited the same forced choice preference as did the Non-Learners. This finding is particularly notable since Learners had more experience with positive prediction errors in NoGo-to-Win than did Non-Learners, yet still showed the same action-biased preferences. Furthermore, within any individual Learner, trials with lower mid-frontal theta responses to Pavlovian conflict during learning were associated with a greater propensity for Pavlovian biases on behavior. Together, these findings suggest that Learners do not differ from Non-Learners in terms of the mechanisms giving rise to such biases (putatively related to *model-free* cortico-striatal function), but rather in their ability to detect when these biases conflict with the rules of the task and need to be suppressed (putatively by *model-based* top-down prefrontal control), see also Guitart-Masip et al. (2012b).

The role of theta

The mid-frontal effect occurred during the P2-N2 time range of the ERP, a temporally specific window known to be affected by conflict-induced cognitive control and expectation-induced mismatch, and known to have a strong spectral signature in the theta band (Hanslmayr et al., 2008; Cavanagh et al., 2012b) and a presumed generator in mid-cingulate cortex (Van Veen and Carter, 2002; Yeung et al., 2004; Hanslmayr et al., 2008). This is the same approximate time period in the ERP which Holroyd et al. (2011) recently suggested was specifically sensitive to reward prediction errors, yet the current findings suggest that this previous effect may reflect general salience instead. Indeed, in this and other studies, mid-frontal theta power appears to reflect a generic signal of the need for top-down control, not an axiomatic reward prediction error (Oliveira et al., 2007; Cavanagh et al., 2012a). Yet very recent evidence has suggested that non-phase locked power (as used here) may preferentially reflect violations of probability, whereas phase-locked ERP amplitude (as in Holroyd et al. (2011)) may be primarily sensitive to valence (Hajihosseini and Holroyd, 2013). Clearly, more explicit hypothesis testing is needed to define the information content reflected within the ERP and constituent frequency bands, and how information may be differentially reflected in power vs. phase activities.

It has previously been shown that mid-frontal theta predicts behavioral slowing (Cavanagh et al., 2010) and switching (Cohen and Ranganath, 2007; van de Vijver et al., 2011) following prediction error; however it was not known if this effect relied the operation of a model-free controller for generic slowing/switching, or a model-based controller for

adaptive behavioral adjustment. The findings from this current study suggest the latter case, given that conflict-induced mid frontal theta was implicated in the ability to overcome Pavlovian biases through both invigoration and inhibition of action. Comparison of the three novel models clearly favored an account whereby trial-to-trial theta power suppressed Pavlovian influence (M6a), rather than promoting instrumental contingencies (M6b) or balancing the relative activity between the two (M6c). Based on theoretical and empirical findings, one candidate mechanism for this effect could be communication from the medial frontal cortex to the subthalamic nucleus via the hyperdirect pathway, which could raise the decision threshold to temporarily prevent the influence of striatal valuation signals on behavior (Frank, 2006; Cavanagh et al., 2011; Ratcliff and Frank, 2012; Zaghoul et al., 2012).

Remaining questions on the interplay between motivational systems

Pavlovian influence was modeled here and elsewhere as a bias in action invigoration/inhibition during response selection (Huys et al., 2011; Guitart-Masip et al., 2012b), but it remains possible that this effect additionally reflects biased learning. The architecture of the basal ganglia is structured to facilitate both learning and action selection in a manner biased by Pavlovian influences. Dopamine bursts to positive prediction errors facilitate plasticity along the D1-receptor mediated direct pathway while suppressing the D2-receptor mediated indirect pathway, while the opposite is true for negative prediction errors (Frank, 2005; Gerfen and Surmeier, 2011; Kravitz et al., 2012). Thus actions are more likely to be invigorated and reinforced following positive prediction errors, and more likely to be suppressed following negative prediction errors. As such, action biases could be accounted for by aberrant learning of action-values, or by a motivational alteration at the time of choice, or by an interaction of the two (e.g. Beeler et al., 2012). Given that these influences over motivation and learning are latent and possibly overlapping, it is difficult to parse the true nature of Pavlovian biases during action learning.

Guitart-Masip and colleagues recently demonstrated that increased activity in the striatum and substantia nigra/ventral tegmental area was associated with action invigoration and (2011 and (2012a) and trial-by-trial action values (2012b). When the effects of valence on action learning were tested, there were no neural correlates of state predictive value or a Pavlovian interaction that could account for the observed Pavlovian biases during learning. However, this is possibly due to a high correlation between action and state values in this task, making it difficult to tease apart with fMRI. In addition, it is unknown whether biases in the post-task transfer phase of the current study reflect motivational, learning, or other influences that boost the apparent value of action over omission.

While much of the evidence thus far favors a motivational account, it is likely that a more specific task and computational model will be required to test separable Pavlovian influences over motivation vs. learning. For example, although the post-training transfer phase data described here are generally supportive of a coupling between action and valence in value learning, the present model may not account for the full spectrum of choices. Specifically, participants reliably preferred Go-to-Avoid over NoGo-to-Avoid, which could not be explained by greater valuation *per se* but may require models that impose a greater degree of learning from positive prediction errors following Go actions than NoGo actions. Regardless of the specific mechanism of action/valence coupling, a similar top-down signal may be required to diminish it, suggesting that findings of increased frontal theta and bilateral inferior frontal gyrus (Guitart-Masip et al., 2012b) remain effective descriptions of the nature of model-based instrumental control.

Conclusion

Individual performances were characterized by a varied mixture of Pavlovian bias, instrumental learning, and model-based control. Through an innovative mixture of cognitive neuroscience and computational modeling we were able to determine the degree of Pavlovian bias over instrumental learning, and the nature of prefrontal control that was applied to ameliorate this bias. The evidence in this report suggests that mid-frontal theta is a sensitive index of model-based prefrontal control over behavior.

Acknowledgments

The authors thank Jerome Sanes for use of the EGI system. This project was supported by NIH 5T32MH019118-21, NIH RO1 MH080066-01, and NSF 1125788.

References

- Beeler, JA.; Frank, MJ.; McDaid, J.; Alexander, E.; Turkson, S.; Sol Bernandez, M.; McGehee, DS.; Zhuang, X. [Accessed December 14, 2012] A Role for Dopamine-Mediated Learning in the Pathophysiology and Treatment of Parkinson's Disease; *Cell Reports*. 2012. p. 1-15. Available at: <http://linkinghub.elsevier.com/retrieve/pii/S2211124712004111>
- Cardinal RN, Parkinson JA, Hall J, Everitt BJ. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience and biobehavioral reviews*. 2002; 26:321–352. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/12034134>. [PubMed: 12034134]
- Cavanagh JF, Cohen MX, Allen JJB. Prelude to and resolution of an error: EEG phase synchrony reveals cognitive control dynamics during action monitoring. *The Journal of neuroscience*; the official journal of the Society for Neuroscience. 2009; 29:98–105. Available at: <http://www.jneurosci.org/content/29/1/98.short>. [PubMed: 19129388]
- Cavanagh JF, Figueroa CM, Cohen MX, Frank MJ. Frontal Theta Reflects Uncertainty and Unexpectedness during Exploration and Exploitation. *Cerebral Cortex*. 2012a; 22:2575–2586. Available at: <http://cercor.oxfordjournals.org/content/22/11/2575>. [PubMed: 22120491]
- Cavanagh JF, Frank MJ, Klein TJ, Allen JJB. Frontal theta links prediction errors to behavioral adaptation in reinforcement learning. *NeuroImage*. 2010; 49:3198–3209. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2818688&tool=pmcentrez&rendertype=abstract>. [PubMed: 19969093]
- Cavanagh JF, Wiecki TV, Cohen MX, Figueroa CM, Samanta J, Sherman SJ, Frank MJ. Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature neuroscience*. 2011; 14:1462–1467. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/21946325>.
- Cavanagh JF, Zambrano-Vazquez L, Allen JJB. Theta lingua franca: a common mid-frontal substrate for action monitoring processes. *Psychophysiology*. 2012b; 49:220–238. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/22091878>. [PubMed: 22091878]
- Chase, HW.; Swanson, R.; Durham, L.; Benham, L.; Cools, R. Feedback-related Negativity Codes Prediction Error but Not Behavioral Adjustment during Probabilistic Reversal Learning. *J Cogn Neurosci*. 2010. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=20146610
- Cohen MX, Ranganath C. Reinforcement learning signals predict future decisions. *J Neurosci*. 2007; 27:371–378. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=17215398. [PubMed: 17215398]
- Cohen MX, Van Gaal S, Ridderinkhof KR, Lamme VA. Unconscious errors enhance prefrontal-occipital oscillatory synchrony. *Front Hum Neurosci*. 2009; 3:54. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=19956401. [PubMed: 19956401]
- Dayan P, Balleine BW. Reward, Motivation, and Reinforcement Learning. *Neuron*. 2002; 36:285–298. Available at: [http://dx.doi.org/10.1016/S0896-6273\(02\)00963-7](http://dx.doi.org/10.1016/S0896-6273(02)00963-7). [PubMed: 12383782]
- Delorme A, Makeig S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods*. 2004; 134:9–21. Available at:

http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15102499. [PubMed: 15102499]

- Frank MJ. Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism. *J Cogn Neurosci*. 2005; 17:51–72. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15701239. [PubMed: 15701239]
- Frank MJ. Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Netw*. 2006; 19:1120–1136. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=16945502. [PubMed: 16945502]
- Frank MJ, Seeberger LC, O'Reilly RC. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*. 2004; 306:1940–1943. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15528409. [PubMed: 15528409]
- Gerfen CR, Surmeier DJ. Modulation of striatal projection systems by dopamine. *Annual review of neuroscience*. 2011; 34:441–466. Available at: <http://www.annualreviews.org/doi/abs/10.1146/annurev-neuro-061010-113641>.
- Guitart-Masip M, Chowdhury R, Sharot T, Dayan P, Duzel E, Dolan RJ. Action controls dopaminergic enhancement of reward representations. *Proceedings of the National Academy of Sciences of the United States of America*. 2012a; 109:7511–7516. Available at: <http://www.pnas.org/content/109/19/7511.short>. [PubMed: 22529363]
- Guitart-Masip M, Fuentemilla L, Bach DR, Huys QJM, Dayan P, Dolan RJ, Duzel E. Action dominates valence in anticipatory representations in the human striatum and dopaminergic midbrain. *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 2011; 31:7867–7875. Available at: <http://www.jneurosci.org/content/31/21/7867.short>. [PubMed: 21613500]
- Guitart-Masip M, Huys QJM, Fuentemilla L, Dayan P, Duzel E, Dolan RJ. Go and no-go learning in reward and punishment: interactions between affect and effect. *NeuroImage*. 2012b; 62:154–166. Available at: <http://dx.doi.org/10.1016/j.neuroimage.2012.04.024>. [PubMed: 22548809]
- Hajihosseini A, Holroyd CB. Frontal midline theta and N200 amplitude reflect complementary information about expectancy and outcome evaluation. *Psychophysiology*. 2013
- Hampton AN, Bossaerts P, O'Doherty JP. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *J Neurosci*. 2006; 26:8360–8367. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=16899731. [PubMed: 16899731]
- Hanslmayr S, Pastotter B, Bauml KH, Gruber S, Wimber M, Klimesch W. The electrophysiological dynamics of interference during the Stroop task. *J Cogn Neurosci*. 2008; 20:215–225. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=18275330. [PubMed: 18275330]
- Hershberger WA. An approach through the looking-glass. *Animal Learning & Behavior*. 1986; 14:443–451. Available at: <http://www.springerlink.com/content/8g8ur4221688613m/>.
- Holland PC. Differential effects of omission contingencies on various components of Pavlovian appetitive conditioned responding in rats. *Journal of Experimental Psychology: Animal Behavior Processes*. 1979; 5:178–193.
- Holroyd CB, Krigolson OE, Lee S. Reward positivity elicited by predictive cues. *Neuroreport*. 2011; 22:249–252. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/21386699>. [PubMed: 21386699]
- Huys, QJM.; Cools, R.; Gölzer, M.; Friedel, E.; Heinz, A.; Dolan, R.J.; Dayan, P. Rangel, A., editor. [Accessed November 8, 2012] Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding; *PLoS computational biology*. 2011. p. e1002028 Available at: <http://dx.plos.org/10.1371/journal.pcbi.1002028>
- Kass RE, Raftery AE. Bayes Factors. *Journal of the American Statistical Association*. 1995; 90:773–795.
- Kayser J, Tenke CE. Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clinical*

neurophysiology_: official journal of the International Federation of Clinical Neurophysiology. 2006; 117:348–368. Available at: <http://dx.doi.org/10.1016/j.clinph.2005.08.034>. [PubMed: 16356767]

- Kravitz AV, Tye LD, Kreitzer AC. Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nature neuroscience*. 2012; 15:816–818. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/22544310>.
- Oliveira FT, McDonald JJ, Goodman D. Performance monitoring in the anterior cingulate is not all error related: expectancy deviation and the representation of action-outcome associations. *J Cogn Neurosci*. 2007; 19:1994–2004. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=17892382. [PubMed: 17892382]
- Prévost C, Liljeholm M, Tyszka JM, O’Doherty JP. Neural correlates of specific and general Pavlovian-to-Instrumental Transfer within human amygdalar subregions: a high-resolution fMRI study. *The Journal of neuroscience_: the official journal of the Society for Neuroscience*. 2012; 32:8383–8390. Available at: <http://www.jneurosci.org/content/32/24/8383.short>. [PubMed: 22699918]
- Ratcliff R, Frank MJ. Reinforcement-based decision making in corticostriatal circuits: mutual constraints by neurocomputational and diffusion models. *Neural computation*. 2012; 24:1186–1229. Available at: http://www.mitpressjournals.org/doi/abs/10.1162/NECO_a_00270. [PubMed: 22295983]
- Talmi D, Seymour B, Dayan P, Dolan RJ. Human pavlovian-instrumental transfer. *The Journal of neuroscience_: the official journal of the Society for Neuroscience*. 2008; 28:360–368. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2636904&tool=pmcentrez&rendertype=abstract>. [PubMed: 18184778]
- Van de Vijver I, Ridderinkhof KR, Cohen MX. Frontal oscillatory dynamics predict feedback learning and action adjustment. *Journal of cognitive neuroscience*. 2011; 23:4106–4121. Available at: http://www.mitpressjournals.org/doi/abs/10.1162/jocn_a_00110. [PubMed: 21812570]
- Van Veen V, Carter CS. The timing of action-monitoring processes in the anterior cingulate cortex. *J Cogn Neurosci*. 2002; 14:593–602. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=12126500. [PubMed: 12126500]
- Yeung N, Botvinick MM, Cohen JD. The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol Rev*. 2004; 111:931–959. Available at: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15482068. [PubMed: 15482068]
- Zaghloul KA, Weidemann CT, Lega BC, Jaggi JL, Baltuch GH, Kahana MJ. Neuronal activity in the human subthalamic nucleus encodes decision conflict during action selection. *The Journal of neuroscience_: the official journal of the Society for Neuroscience*. 2012; 32:2453–2460. Available at: <http://www.jneurosci.org/content/32/7/2453.short>. [PubMed: 22396419]

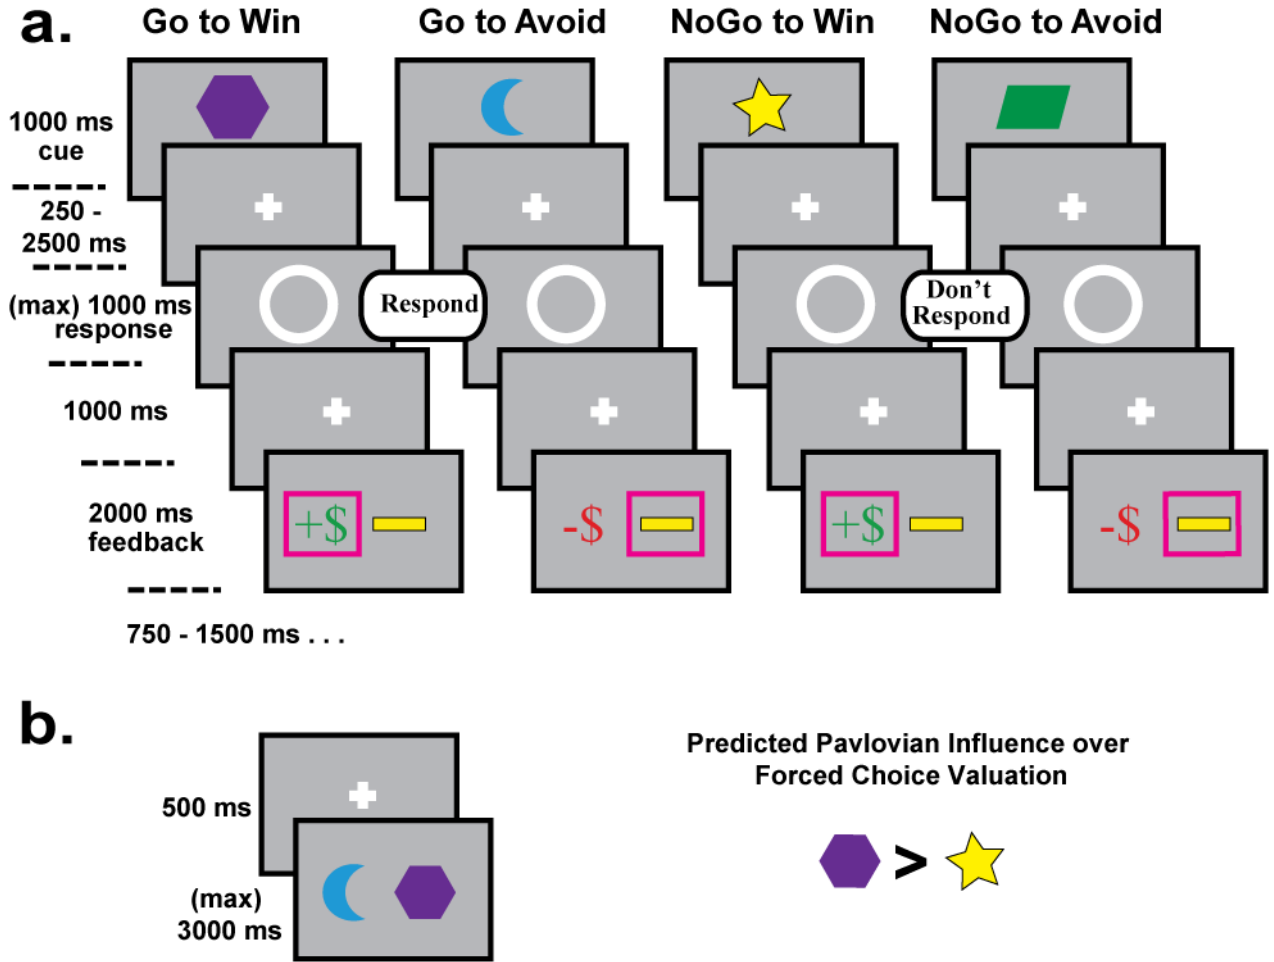


Figure 1. Tasks. **(a)** Training phase. Four different cues were 70% predictive of unique optimal combinations of action and outcome: Go-to-Win, Go-To-Avoid, NoGo-To-Win, and NoGo-to-Avoid. In the context of a Pavlovian bias over performance, conditions with congruent action-outcome pairings should be easy to learn (Go-to-Win, NoGo-to-Avoid), whereas conditions with conflicting action-outcome pairings should be hard to learn (Go-to-Avoid, NoGo-to-Win). **(b)** Post-training transfer phase. Each cue was paired together in a two-alternative forced-choice testing phase; participants were told to select the most rewarding stimulus. It was hypothesized that participants would show a Pavlovian bias during this phase by selecting Go-to-Win > NoGo-to-Win cues.

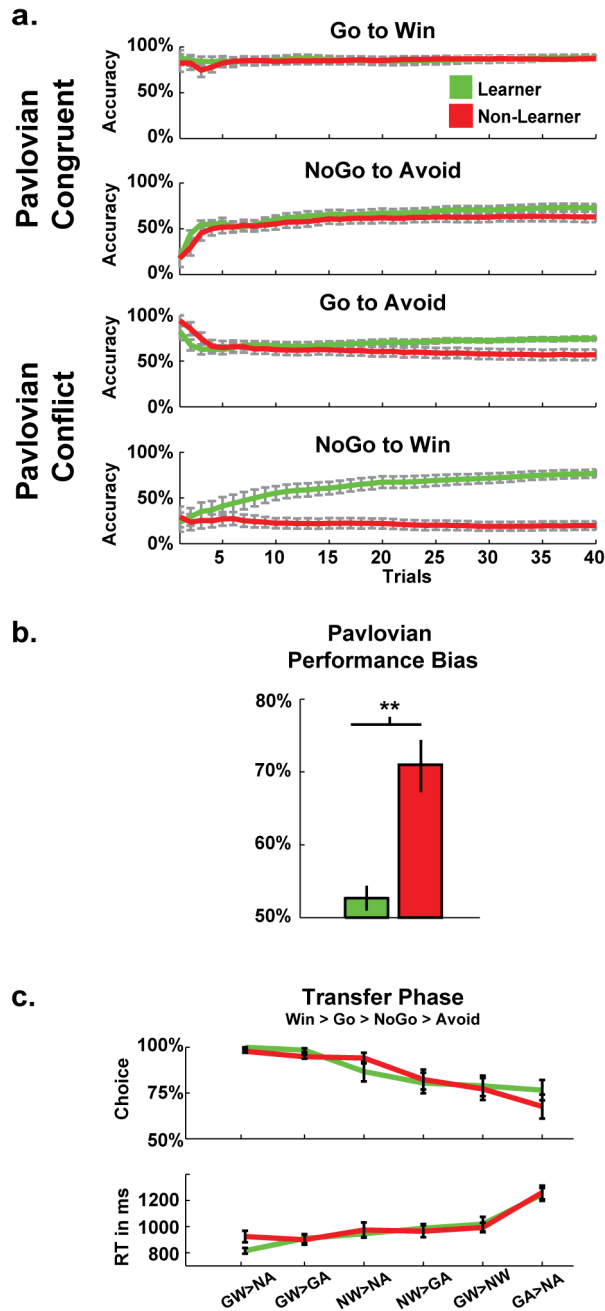
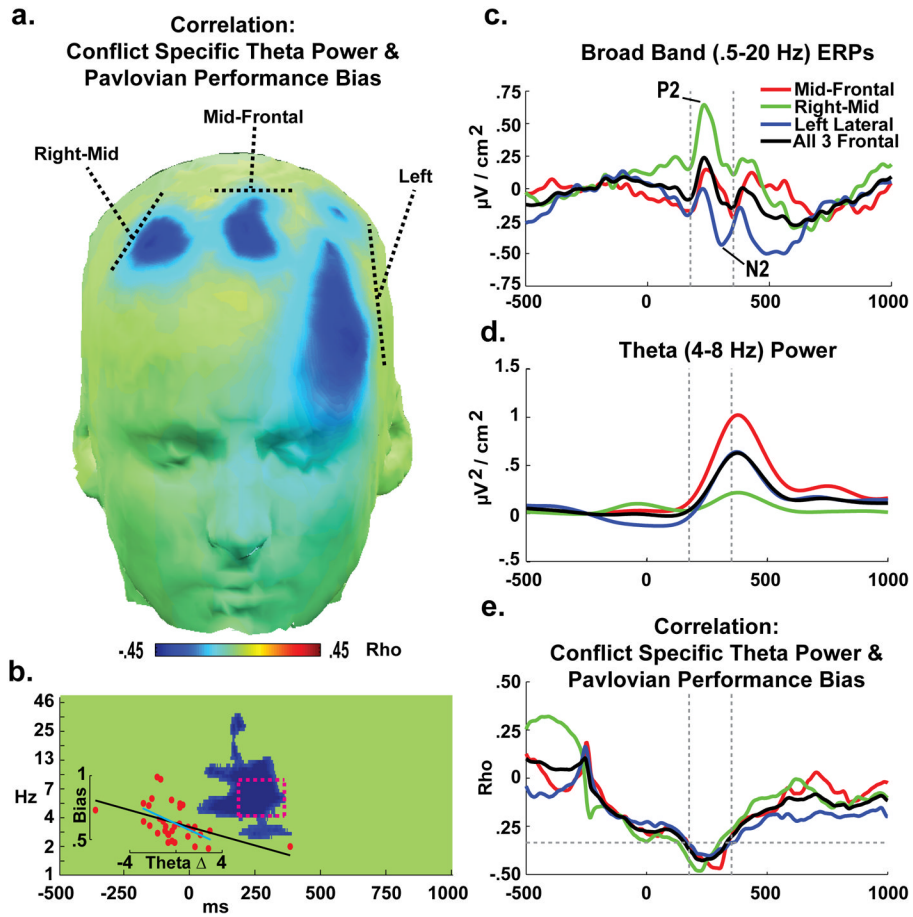


Figure 2. Task performance for Learners and Non-Learners, as defined by accuracy in the latter half of the hardest NoGo-to-Win condition. **(a)** Participant performance during training was similar on easy, congruent conditions, but there were large differences in harder conditions characterized by Pavlovian conflict (Go-to-Avoid & NoGo-to-Win). **(b)** Measures of total invigoration on reward conditions and suppression on avoidance conditions were combined to create a single measure of Pavlovian Performance Bias, capturing an individual’s aggregate tendency to commit an action in the presence of a reward-predictive cue and withhold an action in the presence of a punishment-predictive cue across the entire experiment. **(c)** In a post-task forced choice transfer phase, individuals displayed a *win > go*

> *nogo* > *avoid* ordering of preferences. This finding suggests that both groups had a Pavlovian influence during learning, but the Learner group was somehow able to overcome this bias in the training phase. (GW=Go-to-Win, GA=Go-to-Avoid, NW=NoGo-to-Win, NA=NoGo-to-Avoid).

**Figure 3.**

EEG activities to condition-specific effects, and relationship with performance. (a) There were significant inverse correlations across a range of frontal sites between the cue-locked difference in theta (4–8 Hz) power for conflicting-congruent conditions and Pavlovian Performance Bias. This finding indicates that individuals with greater frontal theta to motivational conflict were characterized by a smaller Pavlovian bias, and thus better performance in conflict trials. (b) Pixel-wise correlations reveal that effects were prevalent around the boxed theta band Region of Interest (ROI: 175–350 ms, 4–8 Hz) over these frontal clusters. The inset scatterplot shows significant non-parametric correlations with (black fit line: $\rho_{(34)} = -.55, p < .01$) or without (cyan fit line: $\rho_{(32)} = -.50, p < .01$) outliers. (c) ERPs demonstrate that the ROI time window occurs during the stimulus-locked P2-N2 complex. (d) This time window specifically captures the waxing of the stimulus-locked theta band power burst. (e) The correlation between conflict-congruent theta power and the Pavlovian Performance Bias was highly similar between sites, suggesting that these separate clusters reflect a common process. The discontinuity in the correlation at –250 ms is an artifact of baseline correction procedures.

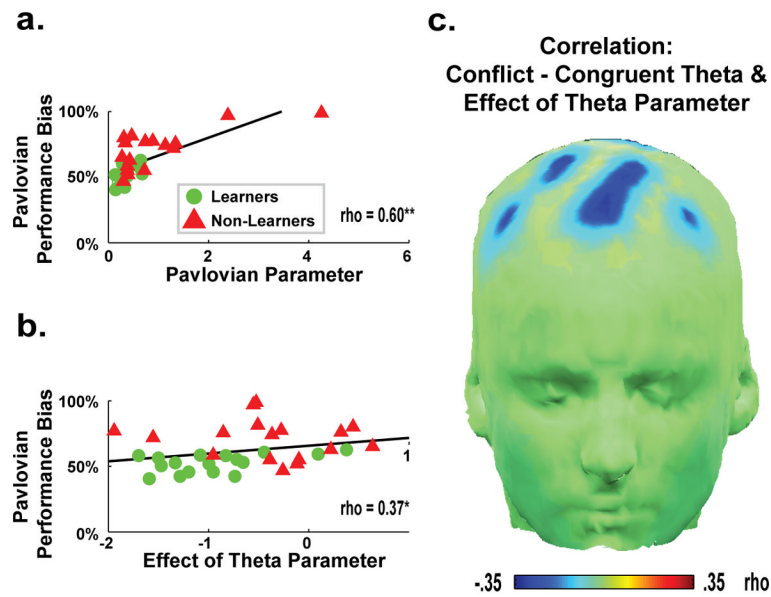


Figure 4.

Parameters from the best fitting computational model (M6a). **(a)** The Pavlovian Bias parameter estimated from the model significantly correlated with the Pavlovian Performance Bias metric used in Figures 2–3. **(b)** The parameter scaling the Effect of Theta in the model also significantly correlated with the Pavlovian Performance Bias metric, indicating that trial-to-trial variations in theta power diminished Pavlovian biases. **(c)** The Effect of Theta parameter had a more adaptive influence (was more negative) in those individuals who had larger conflict-induced theta power.

Table 1

Integrated Bayesian Information Criterion (iBIC) and parameters means (SDs) for each model M1-M5 and novel extensions M6a-c that examined the influence of trial-to-trial EEG theta power.

	M1	M2	M3	M4	M5	M6a	M6b	M6c
iBIC	5613	5615	5379	5141	4926	4857	4947	4893
Feedback Sensitivity (ρ)	4.05 (2.49)	4.67 (3.25)	4.95 (3.04)					
Reward Sensitivity (ρ_{rew})				12.82 (22.27)	6.78 (4.67)	6.86 (4.35)	7.85 (6.40)	9.81 (8.84)
Punishment Sensitivity (ρ_{pun})				3.76 (2.83)	4.81 (3.76)	6.26 (5.67)	4.77 (3.62)	9.36 (10.23)
Learning Rate (e)	.28 (.18)	.29 (.19)	.29 (.19)	.29 (.17)	.28 (.15)	.23 (.15)	.32 (.13)	.27 (.17)
Irreducible Noise (ξ)		.97 (.02)	.94 (.08)	.97 (.02)	.97 (.01)	.96 (.03)	.99 (.01)	.96 (.03)
Go Bias (b)			.38 (.62)	.13 (.85)	.50 (.60)	.58 (.61)	.49 (.59)	.58 (.63)
Pavlovian Bias (π)					0.48 (.77)	0.77 (.75)	.34 (.64)	
Effect of Theta (β)						-0.67 (.67)	.48 (.73)	-.32 (.58)
Tradeoff Parameter (w)								.31 (.14)