

Biologically plausible mechanisms underlying motor response correction during reward-based decision-making

Jung Hoon Lee¹

¹ Allen Institute for Brain Science, Seattle WA

Abstract

Our propensity to acclimate to new surroundings and choose a goal-directed behavior for maximal reward (i.e., optimal outcome) is natural, for it affects our survival. A line of studies suggested that anterior cingulate cortex (ACC) is a potential hub for regulating adaptive behaviors. For instance, an experimental study noted ACC contribution to selecting motor response for maximal reward; it found 1) that ACC neurons were selectively activated when reward was reduced and 2) that the suppression of ACC activity impaired monkeys' ability to change motor response to gain maximal reward. Then, how does ACC modulate motor responses? To address this question, we sought biologically-plausible mechanisms that can account for experimental findings mentioned above. Here, we built a computational model that can replicate the observed ACC activity patterns. Our simulation results raise the possibility that ACC can correct behavioral response by reading out and updating motor plans (guiding future motor responses) stored in prefrontal cortex (PFC).

Introduction

We are able to choose the best action to obtain the optimal outcome by evaluating available choices in terms of (expected) rewards; see (Soltani and Wang, 2008; Sugrue et al., 2005) for reviews. The exact value of action may be uncertain and even evolve over time, and thus the strategy to gain maximal rewards needs to be updated constantly. Then, how does the brain accomplish adaptive decision-making? In principle, the brain would need a precise mapping among at least three distinct factors, sensory inputs, (expected) rewards, and available choices. Notably, anterior cingulate cortex (ACC) has access to crucial information necessary for generating such mapping. First, dopamine, which has been thought to mediate reward signals (Berridge and Robinson, 1998;

Roesch et al., 2007; Sugrue et al., 2005), projects to ACC (Berger et al., 1991; Seamans, 2007). Second, ACC interacts with PFC (Paus, 2001; Sallet et al., 2011) known to participate in executive functions such as working memory (Fuster, 2001; Lara and Wallis, 2015; Miller and Cohen, 2001; Tanji and Hoshi, 2008). Consistent with these findings, ACC has been suggested to detect errors (Brown and Braver, 2005; Carter and Cohen, 1998; Ito et al., 2003; Ullsperger and von Cramon, 2003; Yeung et al., 2004) and select motor response for maximal reward (Bush et al., 2002; Holroyd and Coles, 2002, 2008; Shima and Tanji, 1998). Here, we aim to better understand the contribution of ACC to adaptive decision-making.

A detailed account of ACC activity during reward-based decision-making was reported by Shima and Tanji (Shima and Tanji, 1998), who trained monkeys to ‘push’ or ‘turn’ a lever when a visual cue was given. Monkeys received reward when they selected the reward-choice, which was randomly chosen. After a block of trials, the reward-choice was changed from one to another; to monkeys, this transition was signaled by reward reduction. When reward was reduced, monkeys needed to switch from previously selected motor response to another to obtain maximum reward. During reward-based decision-making, reward reduction depolarized ACC neurons, and the suppression of ACC neuron activity impaired monkeys’ ability to switch motor response accurately (Shima and Tanji, 1998). They further found that most ACC neurons, which responded to reward reduction, were sensitive to the sequence of prior-response (before reward reduction) and post-response (after reward reduction). For instance, a subset of ACC neurons was activated when monkeys changed their response from push to turn, but they remained quiescent when monkeys’ response sequence was reversed. The results indicate that ACC neurons can access information regarding the identity of prior-response and reward. However, a more recent study (Walton et al., 2006) questioned the accuracy of the central role of ACC in error correction (Shima and Tanji, 1998) by reporting that the lesion in ACC impaired monkeys’ ability to maintain the right choice rather than their ability to correct the choice (Walton et al., 2006).

This inconsistency is puzzling, as the same experimental protocol was employed in the two studies (Shima and Tanji, 1998; Walton et al., 2006). To gain insights into ACC function in adaptive decision-making, we sought neural mechanism that can account for these seemingly conflicting observations. Specifically, utilizing a biologically-plausible computational model, we asked if a

single network can account for both experimental findings (Shima and Tanji, 1998; Walton et al., 2006). Inspired by earlier studies suggesting that prefrontal cortex (PFC) subserve working memory (Baddeley, 2003; Miller and Cohen, 2001; Wang et al., 2013; Wilson et al., 1993) and that PFC and ACC are connected reciprocally (Paus, 2001; Sallet et al., 2011), we hypothesized 1) that PFC can hold motor plans in working memory to guide monkeys' response which is required at the onset of visual cues and 2) that ACC can read out and update motor plans stored in PFC. To address these hypotheses, we built a computational model consisting of ACC, PFC, and motor cortex (MC), which we included to describe the full process of motor response (i.e., lever movement) shift elicited by reward reduction.

In this study, we mimicked the experimental protocol used to measure ACC activity (Shima and Tanji, 1998). To model the effects of dopamine in ACC, we simulated the effects of D₂ receptors known to induce inhibition in ACC (Darvish-Ghane et al., 2016) by providing inhibitory currents to ACC neurons. In our model, when reward is reduced, inhibitory currents are removed. Our simulation results suggest 1) that the interplay between ACC and PFC can account for the experimentally observed ACC activity patterns (Shima and Tanji, 1998), 2) that ACC can read out and update the motor plan in PFC and 3) that the removal of a ACC neuron population driven by internal excitation within ACC can destabilize the motor plan in PFC. The results support our hypothesis and further suggest that a single network architecture can account for the two seemingly inconsistent experimental observations (Shima and Tanji, 1998; Walton et al., 2006). This model can provide a starting point from which we can develop biologically-plausible model of ACC that can account for diverse functions of ACC in a single framework (Ebitz and Hayden, 2016; Paus, 2001; Vassena et al., 2017; Yarkoni et al., 2011).

Results

Our model consists of three distinct cortical areas, ACC, PFC and MC; see Fig. 1A. Each area consists of multiple populations of leaky integrate-and-fire neurons (Methods). All neurons are identical in terms of internal dynamics (i.e., subthreshold dynamics and spike generation), but they receive distinct external background and recurrent synaptic inputs; see Methods. In ACC, three excitatory populations (PT, TP and NS) and two inhibitory populations (PT_i and TP_i) interact with one another via connectivity listed in Table 1 (Fig. 1). NS population does not receive any external

background inputs and is internally driven by other ACC neurons. PT_i and TP_i provide feedforward inhibition onto PT and TP, respectively. PT and TP project to NS, which sends excitation to both PT_i and TP_i ; that is, NS can project an indirect feedback inhibition onto PT and TP. Recurrent connections are established within the populations (Table 1); they are not illustrated in Fig. 1 for clarity. ACC is also connected with PFC and MC (Methods), both of which include two excitatory (P and T) and two inhibitory populations (P_i and T_i). In PFC, the excitatory recurrent connections within the excitatory population are set to be strong enough for PFC populations to generate persistent activity (Table 1), which are well-studied neural correlates of working memory. In MC and PFC, T and P represent ‘turn’ and ‘push’, respectively. In MC, they represent motor responses.

We simulated the experimental protocol with 3200 msec-long simulations by considering the effects of dopamine D_2 receptors which are known to induce inhibition in ACC (Darvish-Ghane et al., 2016); Fig. 1B illustrates the simulation protocol. In the simulation, the reward reduction is simulated by removing inhibitory currents to ACC (-250 pA). Each simulation consists of 3 blocks starting with a visual cue (Fig. 1B). As visual cues were used to signal monkeys to move the lever in the experiment (Shima and Tanji, 1998), they must not carry any information regarding reward-response. Therefore, in our model, we simulated visual cues with afferent inputs projecting to both populations in MC. As seen in Fig. 1B, visual cues are introduced to MC at 200, 1200 and 2200 msec after the onset of simulations, and their durations are 200 msec. To test reward effects on network activity, we performed simulations in two different conditions, continuous-reward (CW) and reduced-reward (RW). In the CW condition, the reward level remains constant, whereas in the RW condition, it is reduced at 2000 msec after the second visual cue. The initial motor plan is always ‘turn’ unless stated otherwise. As we assumed that PFC store motor plan as persistent activity, we provided 100 msec-long excitatory inputs to T population in PFC to initiate persistent activity within it before the first visual cue.

PFC can hold motor plan and guide MC response accordingly

We first tested model responses in the CW condition (i.e., the high level of reward through simulations). As seen in Fig. 2A, T population in PFC, stimulated by external inputs before the first visual cue, fires continuously at a high firing rate due to strong recurrent connections in PFC (Table 1). That is, ‘turn’ is retained as the motor plan during the entire simulation period. With

this motor plan in PFC, T population in MC responds to all three visual cues (Fig. 2B); ACC neurons remain quiescent. In the model, MC neurons are set to fire only when they receive afferent inputs from both PFC and visual cues. This makes MC neurons fire only during the visual cue, even though PFC neurons continuously project afferent inputs to MC T population. To examine if this result is robust over random connections and background inputs, we instantiated 100 independent networks. Each network has randomly chosen connections which are drawn from the same connectivity rule (Table 1), and we repeated the same simulation with independent sets of background Poisson spike trains (Table 2). In each simulation, we compared the average firing rates of P and T populations in MC during the presentation of third visual cue (2200-2400 ms). As seen in Fig. 2C, MC responses to the third are ‘turn’ in all 100 simulations, indicating that this network model sustains the same responses when reward is maintained at a constant level.

ACC neurons update motor plan in PFC when they are depolarized by reward reduction

We next tested responses of the model in the RW condition, in which reward is reduced at 2000 ms (after the second visual cue). As in the CW condition, T population in MC responds to the first two visual cues (200-400 and 1200-1400 msec), but P population responds to the third visual cue (2200-2400 msec) (Fig. 3A). This means motor response is updated successfully by the reduction of reward. Then, how does the reward shift motor response? In the model, when the reward is reduced, inhibition onto ACC neurons is lifted, making ACC neurons more responsive to afferent inputs. Among TP and PT populations in ACC, TP population is activated selectively (Fig. 3B) due to the excitation from T population in PFC (Fig. 3C).

We note that the activation of TP population in ACC initiates a series of changes in the model. First, the persistent activity of T population is turned off (Fig. 3C). Once TP population is active, it increases excitatory inputs to T_i in PFC, which in turn inhibits T population in PFC. Second, P population in PFC is turned on (Fig. 3C). As soon as T population in PFC is turned off, P population is liberated from lateral inhibition induced indirectly by T population (Fig. 1A) and starts firing due to external background inputs (Table 2); background inputs alone can drive PFC neurons to fire. Third, NS population in ACC becomes active due to excitation from TP population (Fig. 3B), which increases inhibition onto PT and TP by innervating PT_i and TP_i . This delayed feedback inhibition suppresses TP activity after active firing, meaning the activation of NS

population makes TP population transient, as observed by Shima and Tanji (Shima and Tanji, 1998).

With these changes, P population in PFC is active, when the third visual cue is given (2200-2400 msec), and P population in MC becomes active consequently (Fig. 3A). After the switch in this response is instantiated properly, we assumed that the reward level is restored and re-introduced inhibitory inputs to ACC, making ACC neurons quiescent again (Fig. 3B). These results suggest that ACC can shift motor response indirectly by updating the motor plan in PFC. As in the CW condition, we repeated simulations using 100 independently instantiated networks. To examine when the motor plan is switched in all these simulations, we calculated the selective index (*SI*) defined by Eq. 2 in each simulation. *SI* becomes positive when T population produces stronger responses in the selected 50 ms-time bin than P population. They are normalized to be between -1 and 1; see Methods. The simulation results are summarized in Fig. 4. TP neurons in ACC generate stronger spiking activity than PT neurons do (Fig. 4A), the persistent activity in PFC is successfully transferred from T to P after the reward reduction (Fig. 4B), and motor responses are properly switched in all 100 simulations (Fig. 4C).

We further tested the importance of ACC in shifting motor response in the model by removing inter-areal connections from ACC to PFC. As seen in Figs. 5A and B, T population in MC responds to the third visual cue when inter-areal connections from ACC to PFC are removed, confirming that it is ACC that corrects motor response via interactions with PFC. Finally, we tested whether these results depend on the initial motor plan by changing it from ‘turn’ to ‘push’. As seen in Fig. 6, we observed the equivalent results even when the initial motor plan was switched.

NS population is necessary for keeping motor plan in PFC stable

NS population in ACC shows delayed activity (Fig. 3B), and this population makes TP (PT) population activity transient. These transient and delayed responses are consistent with ACC activity patterns observed during reward-based decision-making (Shima and Tanji, 1998), suggesting that NS population is necessary for replicating ACC activity patterns experimentally observed. To gain better understanding of NS population function, we removed NS from ACC and repeated simulations in the RW condition. We found that MC response to the third population

became erroneous (Fig. 7A). In 6 out of 100 independent simulations, T population produces more responses than P population; that is, motor responses are not properly switched in 6 simulations. Interestingly, we note that both TP and PT populations become active when NS population is removed (Fig. 7B). PT population also becomes active because P population provides excitation to PT population once the motor plan is shifted to ‘push’ (Fig. 7C). With TP and PT populations active, both P and T populations in PFC receive inhibition, disrupting the persistent activity in both populations (Fig 7C). After the quiescent period, the persistent activity can emerge either in P or T population (Fig. 7D).

These observations led us to hypothesize that NS population may be responsible for stabilizing the motor plan stored in PFC. To test this hypothesis, we prolonged each simulation to have 6 visual cues. Specifically, 3 more visual cues are added, and simulations before the fourth visual cue are indeed identical to those with 3 visual cues discussed above. Figure 8A shows how PFC activity evolves over time in the RW condition without NS population. As seen in the figure, the persistent activity moves back and forth after reward reduction. This unstable PFC activity makes MC response to the last four visual cues become erroneous; T population, instead of the correct P population, sometimes responds to visual cues (Figs. 8B-E). We also note that incorrect answers to the first cue (after the reward reduction) are less frequent than those to other three cues (Fig. 8B-E). These results are consistent the observation that the lesion in ACC disrupted monkeys’ ability to maintain the correct choice (Walton et al., 2006), while it did not strongly degrade the accuracy of response to the first cue (after the reward reduction).

Discussion

To gain insights into ACC’s function in regulating motor response during adaptive decision-making, we sought neural circuits capable of 1) replicating ACC neuron activity patterns observed during a simple reward-based decision-making (Shima and Tanji, 1998) and 2) explaining how ACC neurons change motor response. In our model, there are two types of excitatory ACC neurons, PT/TP and NS. They show transient and delayed responses consistent with those of ACC neurons, which were experimentally observed (Shima and Tanji, 1998). Our simulation results suggest that the two neuron types have distinct functions. TP/PT neurons can read out and update motor plan when reward is reduced, which is consistent with correlated activity in the two areas

(Oemisch et al., 2015; Paus, 2001) and the suggested central role of PFC in error detection (Gehring and Knight, 2000). In contrast, NS neurons, internally driven by PT/TP neurons, do not participate in correcting responses directly. Instead, they enable MC and PFC to maintain the right choice by suppressing TP/PT neuron activity. Thus, our model raises the possibility that ACC can update and stabilize motor response, which, in turn, can explain the two seemingly conflicting experimental findings reported by earlier studies (Shima and Tanji, 1998; Walton et al., 2006) with a single mechanism.

Below, we discuss the potential links to the earlier models of ACC; see (Vassena et al., 2017) for a comprehensive review.

ACC as a motor control filter

Our model predicts that ACC neurons can switch motor response via interacting with PFC. That is, ACC in the model works as a type of 'motor control filter' proposed by (Holroyd and Coles, 2002, 2008). In the original motor control filter model, multiple motor controllers produce answers, and ACC selects the best controller for the given task. In our model, PFC is the only considered motor controller instead of multiple controllers, and ACC can dictate motor plan to be stored in PFC. However, our model does not exclude the possibility that ACC can interact with multiple motor controllers, as the motor control filter theory proposes. On the contrary, it proposes a potential mechanism by which ACC can select one motor controller over others. In our model, ACC suppresses the incorrect motor plan by providing di-synaptic inhibition to the PFC population representing it. We underline that this mechanism can also be used to inhibit all other unnecessary controllers. We note that some ACC neurons show sustained activity (Shima and Tanji, 1998), and they can disable unselected motor controllers by inducing constant (tonic) inhibition in them.

ACC as a conflicting monitor

In our model, populations in ACC have exclusive counter parts in PFC; for instance, PT (TP) population in ACC does not interact with T (P) population in PFC (Fig. 1A). These exclusive connections are desirable to implement the strategy to ensure maximal reward in the task used to study ACC activity during reward-based decision-making (Shima and Tanji, 1998). In the brain, however, the exclusive connections may be too restrictive, and the weak cross connections from PFC to ACC (e.g., P in PFC projects to TP in ACC) exist likely. With these cross connections, our model can effectively account for conflicting monitoring attributed to ACC (Botvinick et al., 2001;

Yeung et al., 2004), if PFC can encode multiple aspects of sensory evidence (e.g., color and identity of word). With PFC neurons responding to both the color and the identity of words, a word 'red' written in blue color will innervate more PFC neurons than a word 'red' in written in red color. That is, the incongruent word will enhance ACC activity more than the congruent one, as the conflicting monitoring theories suggested (Botvinick et al., 2001; Yeung et al., 2004), which is supported by PFC's dominant role in action monitoring occurring in ACC (Gehring and Knight, 2000).

Extension to unified framework of ACC function

It is well understood that ACC activity is elicited during various types of cognitive tasks (Ebitz and Hayden, 2016; Paus, 2001; Vassena et al., 2017; Yarkoni et al., 2011). For instance, ACC activity is reported to be correlated with expected values of outcomes (Shenhav et al., 2013; Walton et al., 2006), and it can also encode the efforts for given tasks or choice difficulty (Aarts et al., 2008; Botvinick, 2007; Hauber and Sommer, 2009; Schweimer and Hauber, 2006; Vassena et al., 2014). Although the diverse hypothetical functions of ACC pose great challenge when developing a unified theory of ACC function, recent theoretical studies provide potential mechanisms to account for diverse ACC functions in a single framework; see (Vassena et al., 2017) for a review. Predicted outcome model (PRO) and reward value prediction model (RVPM) assume that ACC neurons predict the value or the likelihood of future outcome via reinforcement learning (Alexander and Brown, 2011; Silvetti et al., 2011; Vassena et al., 2017).

Based on these studies, we may assume that behavioral tasks with stochastic rewards (i.e., uncertainty of rewards) are necessary to fully understand ACC function. As our goal is to elucidate biologically-plausible mechanisms to account for a simple deterministic behavioral task, in which components of the task have no uncertainty, we did not consider ACC functions associated with stochastic rewards in this study. In the future, we will extend our model by considering 1) more complex tasks with stochastic rewards (i.e., uncertainty) and 2) D₁ receptor which was suggested to underlie ACC's contribution to effort-based decision-making (Schweimer and Hauber, 2006). The extended study could bring us a step further to a single biologically-plausible theory to account for diverse functions of ACC.

It should be noted that biologically-plausible models can simulate multimodal neural signals including functional magnetic resonance image, electroencephalogram and single unit activity. As

these multimodal signals can be used to test a theory against abundant sets of data, we expect the extended version of our biologically-plausible model to provide rigorous ways to test multiple theories regarding ACC functions. For instance, it will produce multiple scale neural responses derived from PRO and RVPM models, and such predictions will help us choose the model that describes ACC function more accurately.

Methods

Our model consists of three cortical areas interacting with one another via inter-areal connections shown in Fig. 1 and Table 1. In each area of the model, neurons are split into several populations and have distinct connectivity. In addition, external background inputs (Poisson spike trains) are introduced to them to regulate their excitability and make their activity stochastic. The rates of Poisson spike trains are specific to neuron types (Table 2). We used NEST, the peer-reviewed simulator, to build our network model (Gewaltig and Diesmann, 2007).

Neuron and synapse models

All neurons in the model are implemented by leaky integrate-and-fire (LIF) neurons. Specifically, the subthreshold membrane potentials obey Eq. 1, and the crossing of the spike threshold (-55 mV) induces a spike. After the spikes are registered in the simulators, the membrane potentials are reset to the resting membrane potential (-70 mV), and no spike was allowed during the absolute refractory period (3 ms). The synapses in the model induce instantaneous voltage jumps in target neurons and decayed exponentially according to Eq. 1.

$$\frac{dV_i(t)}{dt} = -\frac{V_i(t)}{\tau_m} + \frac{1}{C} \sum_{j,k} w_{ij} e^{-(t-t_k)/\tau_{syn}} \quad (1)$$

, where V is the membrane potentials; where $\tau_m=10$ ms and $\tau_{syn}=2$ ms are the time constants of membrane potentials and synaptic events, respectively; where t_k is the spike time of presynaptic neurons; where w_{ij} are synaptic strengths listed in Table 1; where $C=250$ pF is the capacitance of the membrane. The NEST's native neuron model 'iaf_psc_exp' is used to implement this neuron (Gewaltig and Diesmann, 2007); see also (Potjans and Diesmann, 2014) for additional information. The synaptic weights are fixed during the entire simulation; that is, we ignored the plasticity of synapses. The external background spike trains are mediated by synapses whose strengths are 200 pA.

Model structure

Our model implements three cortical areas and their interactions to address our hypothesis. Each cortical area consists of identical LIF neurons and has multiple populations of excitatory and inhibitory neurons, which was inspired by cortical network models proposed earlier; see (Ardid et al., 2007; Compte et al., 2000) for instance. Here we assign specific functions to the populations in each area. All these three areas interact with one another via connections illustrated in Fig. 1. Specifically, ACC and PFC are reciprocally connected (Paus, 2001), and MC projects nonspecific inputs to ACC. In all three areas, excitatory populations consist of 400 neurons, whereas inhibitory populations consist of 100 neurons; for clarity, the raster plots illustrated in Results show 10% of neurons. All connections are drawn randomly using the connection probabilities and synaptic weights listed in Table 1.

The model ACC consists of three excitatory and two inhibitory populations. The first two excitatory populations named PT and TP aim to simulate ACC neurons observed in the experiment, which were sensitive to the sequence of motor responses prior to and post reward reduction (Shima and Tanji, 1998); TP and PT stands for ‘Turn to Push’ and ‘Push to Turn’, respectively. The two inhibitory populations TP_i and PT_i are exclusively connected to either TP or PT and provide feedforward inhibition onto the target population (Fig. 1). The third NS (standing for non-selective) population receives excitatory inputs from both PT and TP and sends excitation to TP_i and PT_i , which means that NS provides indirect feedback inhibition to PT and TP. ACC is known to have dopamine receptors, and we considered the effects of D_2 receptors in ACC. As D_2 receptors in ACC induce inhibition (Darvish-Ghane et al., 2016), we simulated the effects of dopamine by providing inhibitory currents to all ACC neurons (-250 pA). These inhibitory neurons are removed when the reward reduction is assumed.

The model PFC consists of two excitatory and two inhibitory populations. The two populations P and T are assumed to represent motor plans ‘Push’ and ‘Turn’, respectively. These neurons are connected with strong recurrent connections within P and T (Table 1) so that they could generate persistent activity which is believed to underlie working memory (Baddeley, 2003; Miller and

Cohen, 2001; Wang et al., 2013; Wilson et al., 1993). P and T also interact with each other and with inhibitory populations (Fig. 1).

The model MC has a similar structure as that of PFC. As seen in Fig. 1, the two excitatory populations (P and T) interact with two inhibitory populations. P and T populations receive afferent inputs from P and T population in PFC inspired by the influence of PFC on motor areas (Tanji and Hoshi, 2008). P and T are assumed to represent actual motor response, and the activation of P (T) population is interpreted as Push (Turn) response. In addition, we assumed that these two excitatory populations receive afferent inputs elicited by visual cues; see (Salinas and Romo, 1998) for the effects of sensory inputs onto motor cortex.

Simulation protocol

We mimicked the protocol used in the experiment (Fig. 1). Specifically, we simulated two trials prior to the onset of reward reduction and a single post trial. We performed simulations in two simulation conditions. In the reduced reward (RW) condition, the three visual cues are presented at 200, 1200 and 2200 msec, and the reward is assumed to be reduced at 2000 msec. Specifically, we removed inhibitory inputs to ACC neurons. We also performed the continuous reward (CW) as the control experiment, in which reward level is kept constant during the entire simulation period.

To examine the robustness of our results to stochasticity, 100 independent (but statistically equivalent) networks were instantiated using the same connectivity rule (Table 1), and 100 independent simulations were conducted. In each simulation, the background inputs to neurons in the model are independently generated using the rate given in Table 2.

Time course of PFC activity

To track the change in motor plans in PFC, we compared the outputs between T and P in PFC. Specifically, we split the spikes of PFC neurons into the 50 ms-long bins, which do not overlap one another. Then, we calculated the selective index (*SI*) using the binned firing rates (P_n , T_n) of P and T populations using Eq. 2

$$SI(t_n) = \frac{T_n - P_n}{\max(\max(T_n), \max(P_n))} \quad (2)$$

, where t_n is the center of n^{th} bin. When SI values are positive, T population produces stronger responses than P population. In contrast, SI values become negative, P population produces stronger responses. In the simulations, P and T populations are activated exclusively with one another due to the di-synaptic mutual inhibition (Table 1); once one of the populations is activated, its activity is persistent.

Acknowledgments

We wish to thank the Allen Institute founders, Paul G. Allen and Jody Allen, for their vision, encouragement and support.

References

- Aarts, E., Roelofs, A., and van Turenout, M. (2008). Anticipatory Activity in Anterior Cingulate Cortex Can Be Independent of Conflict and Error Likelihood. *J. Neurosci.* 28, 4671–4678.
- Alexander, W.H., and Brown, J.W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nat. Neurosci.* 14, 1338–1344.
- Ardid, S., Wang, X.-J., and Compte, A. (2007). An integrated microcircuit model of attentional processing in the neocortex. *J. Neurosci.* 27, 8486–8495.
- Baddeley, A. (2003). Working memory: looking back and looking forward. *Nat. Rev. Neurosci.* 4, 829–839.
- Berger, B., Gaspar, P., and Verney, C. (1991). Dopaminergic innervation of the cerebral cortex: Unexpected differences between rodents and primates. *Trends in Neurosci* 14, 21–27.
- Berridge, K.C., and Robinson, T.E. (1998). What is the role of dopamine in reward: Hedonics, learning, or incentive salience? *Brain Res. Rev.* 28, 308–67.
- Botvinick, M.M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cogn. Affect. Behav. Neurosci.* 7, 356–366.
- Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S., and Cohen, J.D. (2001). Conflict monitoring and cognitive control. *Psychol Rev.* 108, 624–652.

- Brown, J.W., and Braver, T.S. (2005). Learned Predictions of Error Likelihood in the Anterior Cingulate Cortex. *Science* (80-.). *307*, 1118–1121.
- Bush, G., Vogt, B.A., Holmes, J., Dale, A.M., Greve, D., Jenike, M.A., and Rosen, B.R. (2002). Dorsal anterior cingulate cortex: a role in reward-based decision making. *Proc. Natl. Acad. Sci.* *99*, 523–528.
- Carter, C.S., and Cohen, J. (1998). Anterior Cingulate Cortex, Error Detection, and the Online Monitoring of Performance. *Science* (80-.). *280*, 747–749.
- Compte, A., Brunel, N., Goldman-Rakic, P.S., and Wang, X.J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* *10*, 910–923.
- Darvish-Ghane, S., Yamanaka, M., and Zhuo, M. (2016). Dopaminergic Modulation of Excitatory Transmission in the Anterior Cingulate Cortex of Adult Mice. *Mol. Pain* *12*, 1–14.
- Ebitz, R.B., and Hayden, B.Y. (2016). Dorsal anterior cingulate: a Rorschach test for cognitive neuroscience. *Nat. Neurosci.* *19*, 1278–1279.
- Fuster, J.M. (2001). The prefrontal cortex - An update: Time is of the essence. *Neuron* *30*, 319–333.
- Gehring, W.J., and Knight, R.T. (2000). Prefrontal-cingulate interactions in action monitoring. *Nat. Neurosci.* *3*, 516–520.
- Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia* *2*, 1430.
- Hauber, W., and Sommer, S. (2009). Prefrontostriatal circuitry regulates effort-related decision making. *Cereb. Cortex* *19*, 2240–2247.
- Holroyd, C.B., and Coles, M.G.H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* *109*, 679–709.
- Holroyd, C.B., and Coles, M.G.H. (2008). Dorsal anterior cingulate cortex integrates reinforcement history to guide voluntary behavior. *Cortex* *44*, 548–559.
- Ito, S., Stuphorn, V., Brown, J.W., and Schall, J.D. (2003). Performance Monitoring by the Anterior Cingulate Cortex During Saccade Countermanding. *Science* (80-.). *302*, 120–122.
- Lara, A.H., and Wallis, J.D. (2015). The Role of Prefrontal Cortex in Working Memory: A Mini Review. *Front. Syst. Neurosci.* *9*, 1–7.
- Miller, E.K., and Cohen, J.D. (2001). An Integrative Theory of Prefrontal Cortex Function. *167–*

202.

Oemisch, M., Westendorff, S., Everling, S., and Womelsdorf, T. (2015). Interareal Spike-Train Correlations of Anterior Cingulate and Dorsal Prefrontal Cortex during Attention Shifts. *J. Neurosci.* *35*, 13076–13089.

Paus, T. (2001). Primate anterior cingulate cortex: where motor control, drive and cognition interface. *Nat. Rev. Neurosci.*

Potjans, T.C., and Diesmann, M. (2014). The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model. *Cereb. Cortex* *24*, 785–806.

Roesch, M., Calu, D., and Schoenbaum, G. (2007). Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *December 10*, 1615–1624.

Salinas, E., and Romo, R. (1998). Conversion of sensory signals into motor commands in primary motor cortex. *J. Neurosci.* *18*, 499–511.

Sallet, J., Mars, R.B., Quilodran, R., Procyk, E., Petrides, M., and Rushworth, M.F.S. (2011). *Neuroanatomical Basis of Motivational and Cognitive Control* (MIT Press).

Schweimer, J., and Hauber, W. (2006). Dopamine D1 receptors in the anterior cingulate cortex regulate effort-based decision making. *Learn. Mem.* *777–782*.

Seamans, J. (2007). Dopamine anatomy. *Scholarpedia* *2*, 3737.

Shenhav, A., Botvinick, M., and Cohen, J. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* *79*, 217–240.

Shima, K., and Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* *282*, 1335–1338.

Silvetti, M., Seurinck, R., and Verguts, T. (2011). Value and Prediction Error in Medial Frontal Cortex: Integrating the Single-Unit and Systems Levels of Analysis. *Front. Hum. Neurosci.* *5*, 1–15.

Soltani, A., and Wang, X.J. (2008). From biophysics to cognition: reward-dependent adaptive choice behavior. *Curr. Opin. Neurobiol.* *18*, 209–216.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* *6*, 363–375.

Tanji, J., and Hoshi, E. (2008). Role of the lateral prefrontal cortex in executive behavioral control. *Physiol. Rev.* *88*, 37–57.

Ullsperger, M., and von Cramon, D.Y. (2003). Error monitoring using external feedback:

specific roles of the habenular complex, the reward system, and the cingulate motor area revealed by functional magnetic resonance imaging. *J. Neurosci.* *23*, 4308–4314.

Vassena, E., Silvetti, M., Boehler, C.N., Achten, E., Fias, W., and Verguts, T. (2014). Overlapping neural systems represent cognitive effort and reward anticipation. *PLoS One* *9*.

Vassena, E., Holroyd, C.B., and Alexander, W.H. (2017). Computational models of anterior cingulate cortex: At the crossroads between prediction and effort. *Front. Neurosci.* *11*, 1–9.

Walton, M.E., Behrens, T.E.J., Buckley, M.J., Rushworth, M.F.S., and Kennerley, S.W. (2006). Optimal decision making and the anterior cingulate cortex. *Nat. Neurosci.* *9*, 940–947.

Wang, M., Yang, Y., Wang, C.J., Gamo, N.J., Jin, L.E., Mazer, J.A., Morrison, J.H., Wang, X.J., and Arnsten, A.F.T. (2013). NMDA Receptors Subserve Persistent Neuronal Firing during Working Memory in Dorsolateral Prefrontal Cortex. *Neuron* *77*, 736–749.

Wilson, F., Scaldidhe, S., and Goldman-Rakic, P. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science* (80-.). *260*, 1955–1958.

Yarkoni, T., Poldrack, R., and Nichols, T. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* *8*, 665–670.

Yeung, N., Botvinick, M.M., and Cohen, J.D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol. Rev.* *111*, 931–959.

Figures

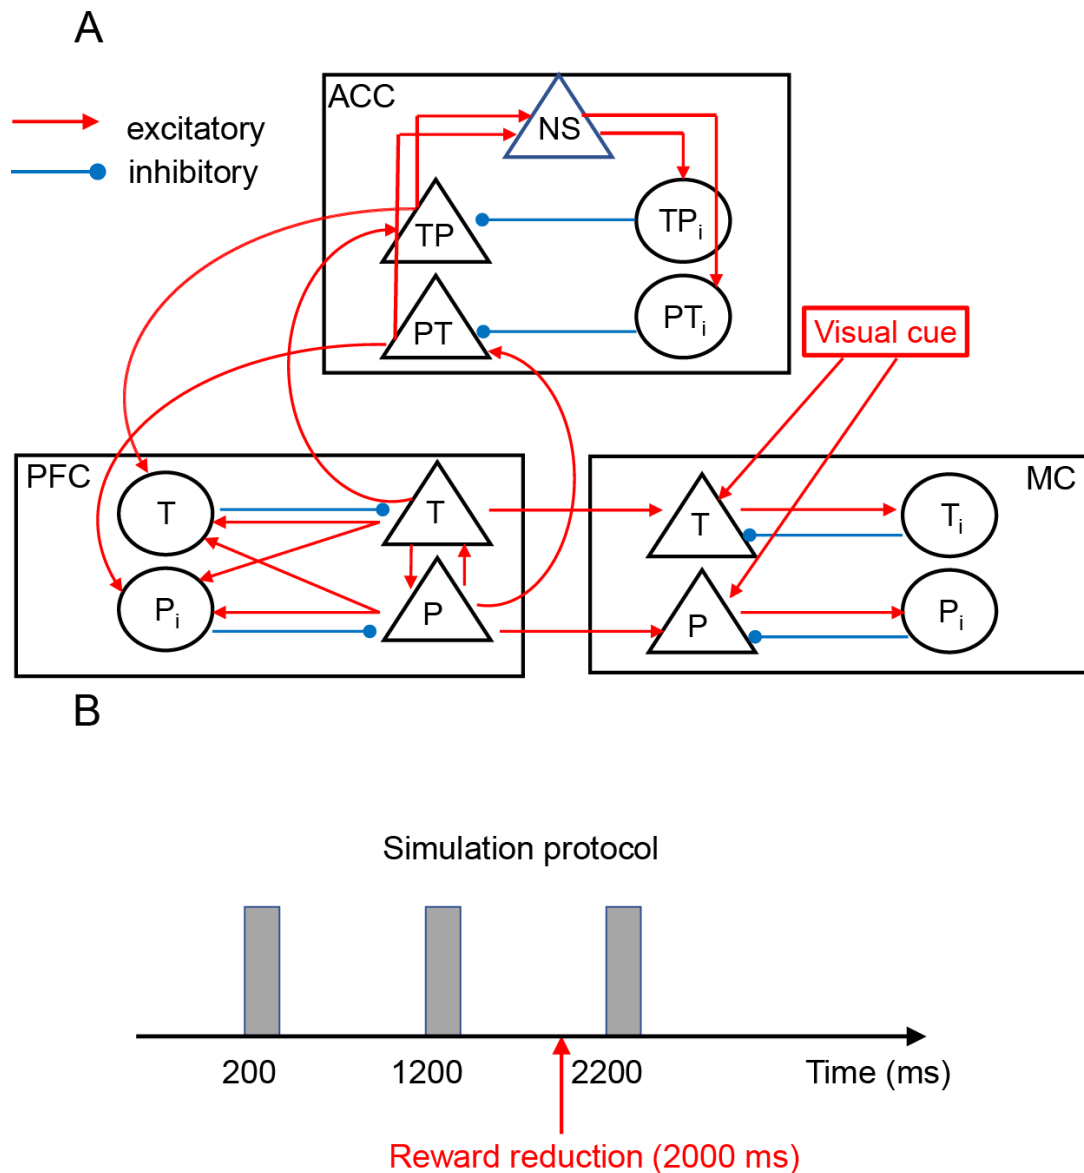


Figure 1: Structure of the model. (A), Schematics of the model consisting of three cortical areas. Triangles and circles represent excitatory and inhibitory neuron populations, respectively. Red and blue arrows indicate excitatory and inhibitory connections. For the sake of clarity, feedback connections from MC to ACC and recurrent connections within the same populations are not shown in the figure; see Table 1 for the full connectivity. (B), Protocol of simulations, three 200 ms-long visual cues are introduced at 200, 1200 and 2200 ms, and the reward is reduced at 2000 ms.

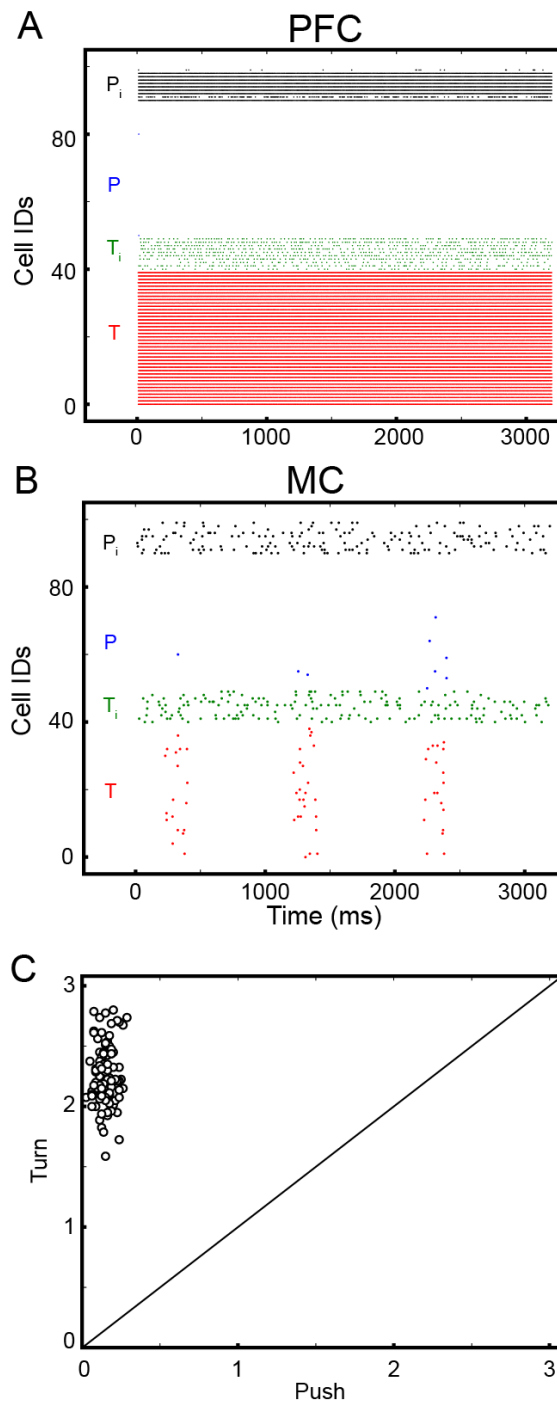


Figure 2: Model responses in the CW condition. For clarity, we show spikes from 10% of neurons. **(A)**, Spikes from T, P, T_i and P_i populations of PFC are shown in red, blue, green and black, respectively. **(B)**, Spikes from T, P, T_i and P_i populations of MC are shown in red, blue, green and black **(C)**, Comparison between T and P populations of MC in 100 independent simulations. x - and y -axis represent the firing rates (Hz) of P and T populations during the presentation of third visual cue; T and P represent ‘turn’ and ‘push’. In each simulation, we calculated the average firing rate of T and P populations in MC. Each dot represents comparison between them.

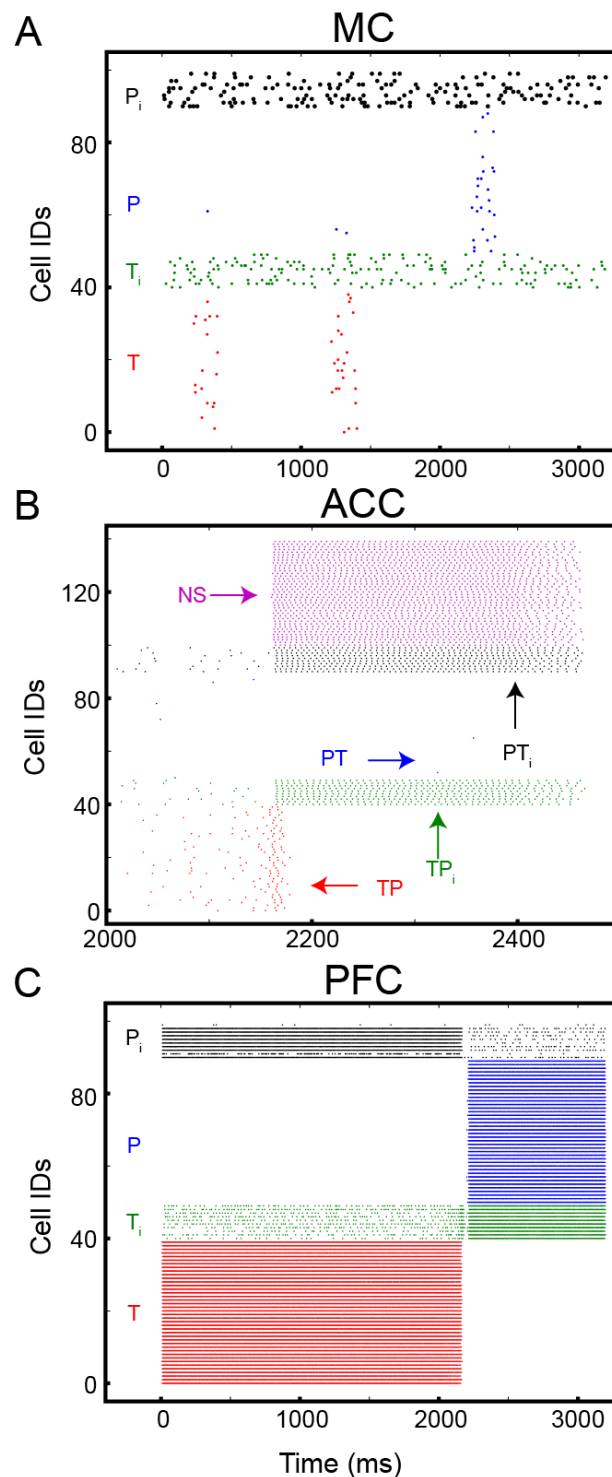


Figure 3: Model responses in the RW condition. For clarity, we show spikes from 10% of neurons. (A), Spikes from T, P, T_i and P_i populations of MC are shown in red, blue, green and black (B), TP, PT, TP_i, PT_i and NS populations of ACC are shown in red, blue, green, black and magenta. (C), T, P, T_i, and P_i of PFC are shown in red, blue, green, black.

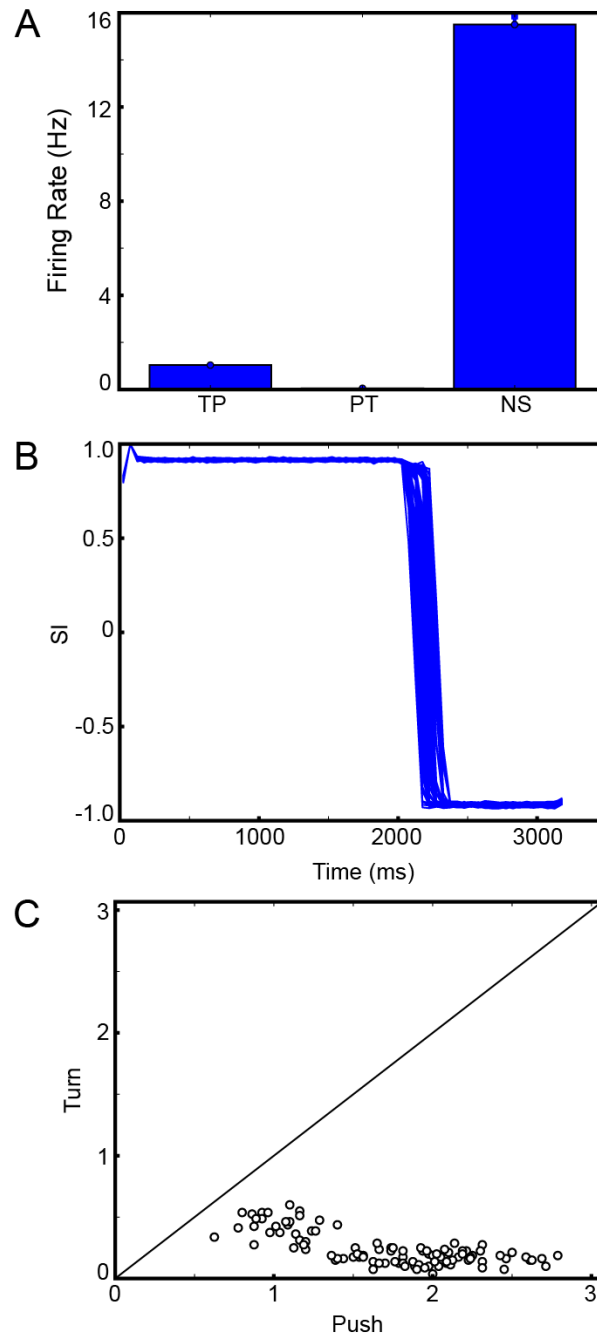


Figure 4: Responses from 100 independent simulations in the CW condition. (A), Mean values and standard errors of firing rates of TP, PT and NS populations in ACC neurons from 100 independent simulations; the spikes are collected from all ACC neurons during the entire simulation. In each simulation, we calculate the average firing rates of the three populations and calculated the mean and standard errors over 100 simulations. (B), SI values of PFC from 100 populations. (C), Comparison between T and P populations in MC. x - and y -axis represent the firing rates during the presentation of third visual cue. In 100 independent simulations, we calculated the average firing rates of T and P populations of MC. Each dot shows the firing rate in each simulation.

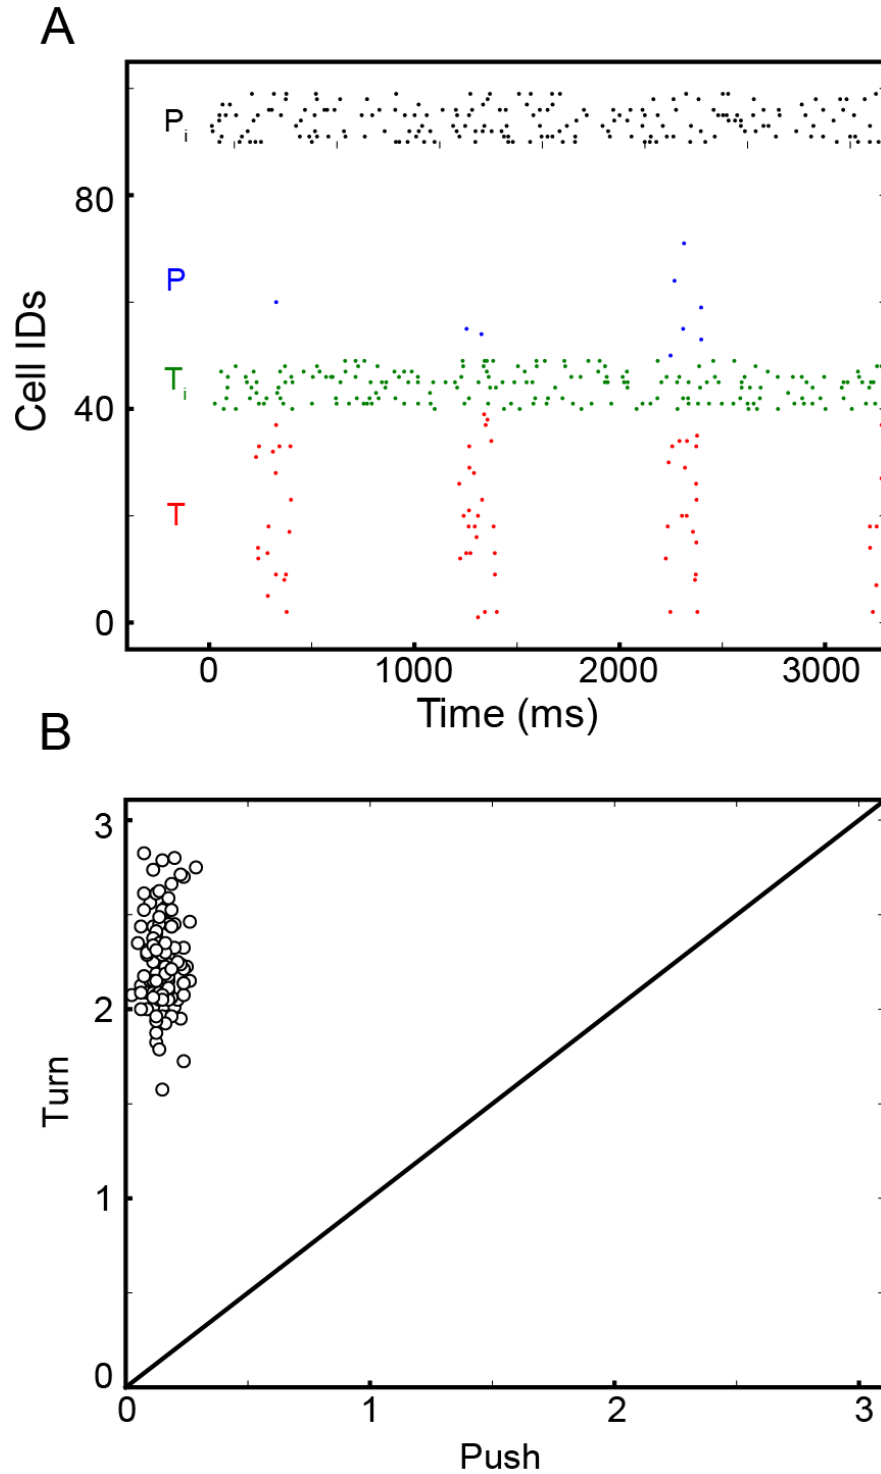


Figure 5: Model responses without afferent inputs from ACC to PFC. (A), Spikes from 10% of MC neurons. **(B)**, Comparison between MC responses in the 100 independent simulations.

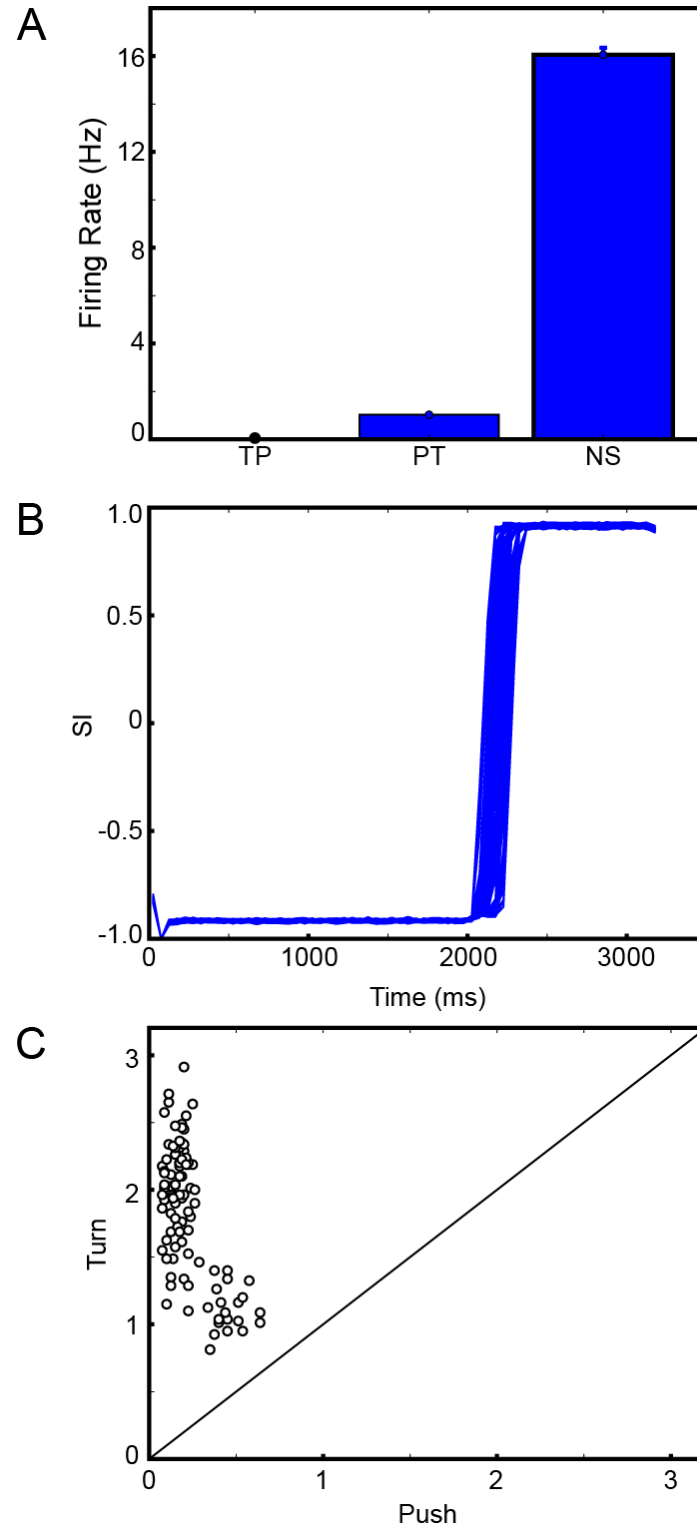


Figure 6: Model response with 'Push' as the initial motor plan. The same as in Figure 5, but a 'push' as the initial motor plan.

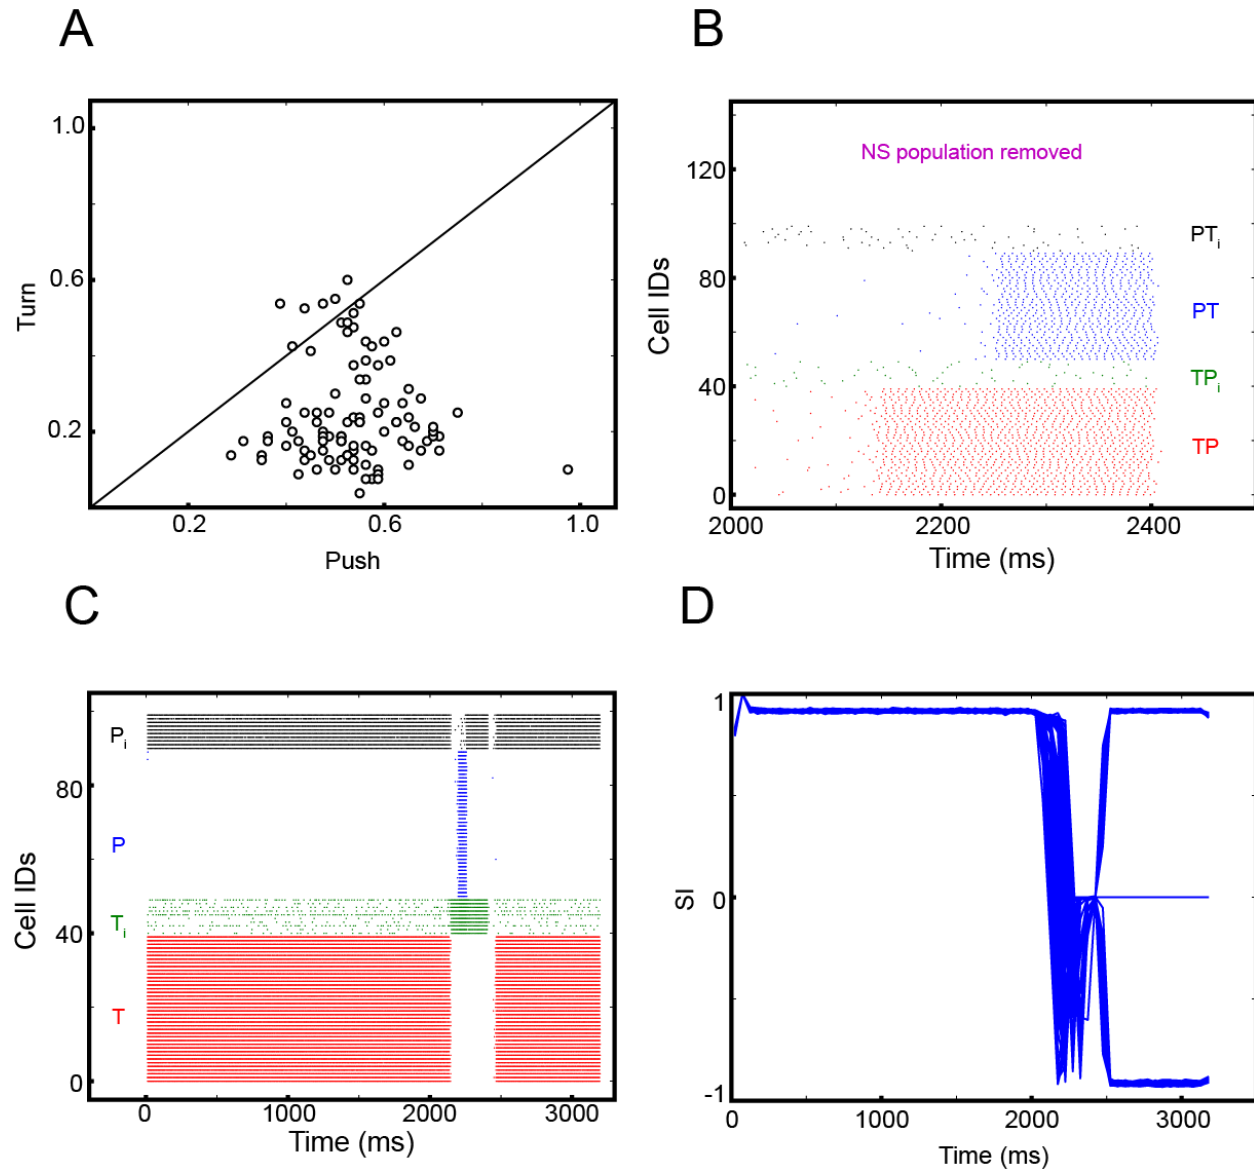


Figure 7: Model responses without NS population. (A), Comparison between MC responses in the 100 independent simulations. (B), Spikes from 10% of ACC neurons. TP, PT, TP_i, PT_i and NS populations of ACC are shown in red, blue, green, black and magenta. (C), Spikes from 10% of PFC neurons. T, P, T_i, and P_i of PFC are shown in red, blue, green, black. (D), Time course of SI.

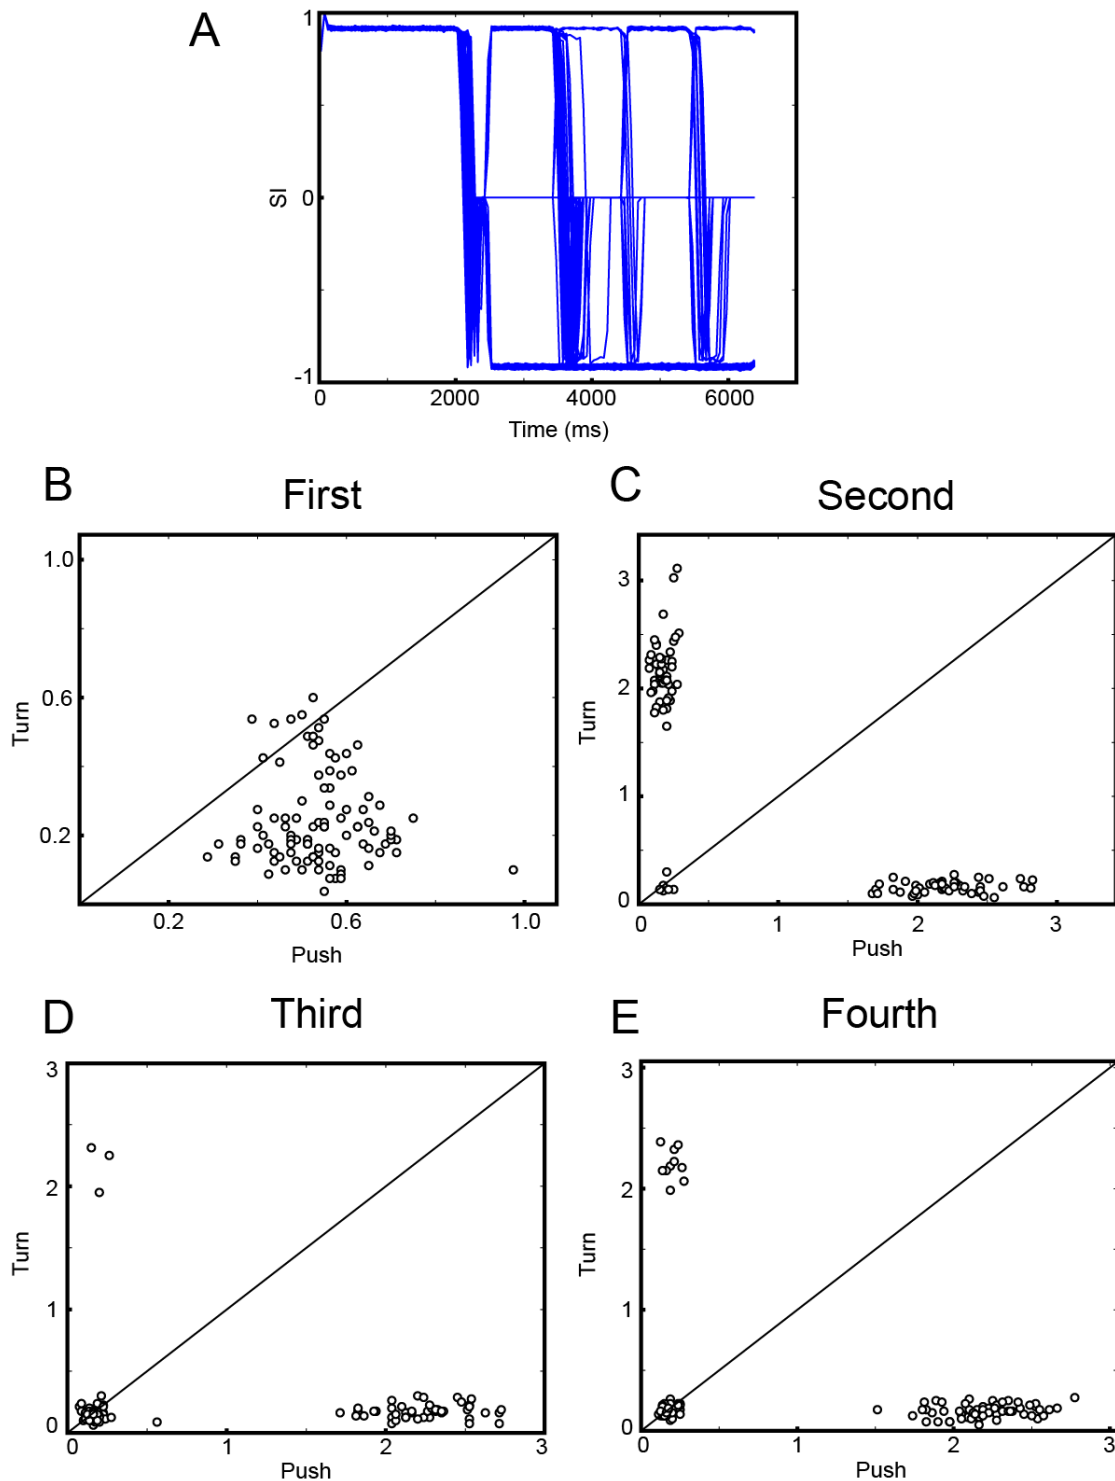


Figure 8: Model responses to 6 visual cues without NS population. (A), Time course of SI . (B)-(E), Comparisons between MC responses to the last four visual cues. The panels 'First', 'Second', 'Third', and 'Fourth' above indicate the third, fourth, fifth and sixth visual cues in the simulation, respectively.

Tables

Table 1: Connectivity pattern in the model. For each pair of pre- and post-synaptic neurons, connections are drawn randomly using the connection probability and strength listed below. Each connection type was separated by semicolon in the table and has a conduction delay. The connections among neurons in the same area have 2.0 ms delay, whereas the connections across neurons areas have 5 ms delay.

Connection type (presynaptic → postsynaptic)	Connection probability	Synaptic weights (pA)
Within ACC		
PT→PT; TP→TP	0.1	80
PT, TP→NS	0.1	100
NS→NS	0.2	60
NS→PT _i , TP _i	0.1	100
PT _i →PT; TP _i →TP	0.05	-100
PT _i →PT _i ; TP _i →TP _i	0.05	-50
Within PFC		
T→T; P→P	0.1	150
T→P; P→T	0.05	50
T→T _i ; P→P _i	0.1	20
T→P _i ; P→T _i	0.1	200
T _i →T _i , T; P _i →P _i , P	0.2	-400
Within MC		
T→T, T _i ; P→P, P _i	0.1	50
T _i →T, T _i ; T _i →P, P _i	0.3	-400
ACC→PFC		
TP→T _i ; PT→P _i	0.3	100
PFC→ACC		
P→PT; T→TP	0.1	3
PFC→MC		
T→T; P→P	0.05	10
MC→ACC		
T→TP, PT; P→TP, PT	0.2	100

Table 2: The external inputs to the model. In the model, three different types of external inputs are used. First, background inputs are introduced to all neurons except NS neuron in ACC during the entire simulation periods. This table lists the rate of Poisson spike trains selected for neuron types. Second, visual cues are introduced to P and T populations in MC. Third, inhibitory currents are introduced to all ACC neurons to simulate the effect of D2 dopamine receptors (Darvish-Ghane et al., 2016).

Background	
Neuron type	Rate
ACC	
PT, TP	500 Hz
PTi, TPi	800 Hz
PFC	
T, P, Ti, Pi	800 Hz
MC	
T, P	300 Hz
Ti, Pi	900 Hz
Visual cues	
Neuron type	Rate
T, P in MC	300 Hz
Inhibitory currents in ACC via dopamine D ₂	
TP, PT, NS, TPi, PTi	-250 pA