activity that anticipates future actions? Of specific interest are the means by which sensory information is integrated by layer 5 ALM neurons and converted into a motor plan, especially with respect to the roles of specific projection classes. Recordings of these neurons and their inputs throughout learning may help to shed light on how this is established. (2) What is the function, if any, of the ALM-dependent global activation of the cortex? It is especially pertinent because inactivation of brain regions other than ALM did not lead to measurable behavioral deficits (Allen et al., 2017; Makino et al., 2017), adding to the growing evidence that neural responses reflecting an animal's decision may not be causally involved in carrying out that decision (Katz et al., 2016). (3) Is ALM a universal hub for learned movement? This is especially pertinent given the fact that ALM activity is both necessary and sufficient for tongue protrusions (Guo et al., 2014; Komiyama et al., 2010), and all the mice in these studies were required to lick as part of the task. It is possible that other distinct premotor regions may exist for behaviors that do not involve orofacial movements. If this is the case, do all such premotor hubs share a common organizing principle based on their inputs and outputs?

Overall, these studies represent a major step forward toward a detailed understanding of the large-scale dynamics involved in motor preparation. By taking a more global view of activity, as is routine for simpler systems (Ahrens et al., 2012), we can begin to track—and one day to understand—processes that require complex interactions across brain regions.

REFERENCES

Ahrens, M.B., Li, J.M., Orger, M.B., Robson, D.N., Schier, A.F., Engert, F., and Portugues, R. (2012). Nature *485*, 471–477.

Allen, W.E., Kauvar, I.V., Chen, M.Z., Richman, E.B., Yang, S.J., Chan, K., Gradinaru, V., Deverman, B.E., Luo, L., and Deisseroth, K. (2017). Neuron *94*, this issue, 891–907.

Chen, T.W., Li, N., Daie, K., and Svoboda, K. (2017). Neuron *94*, this issue, 866–879.

Erlich, J.C., Bialek, M., and Brody, C.D. (2011). Neuron *72*, 330–343.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014). Neuron *81*, 179–194.

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Nature *535*, 285–288.

Komiyama, T., Sato, T.R., O'Connor, D.H., Zhang, Y.X., Huber, D., Hooks, B.M., Gabitto, M., and Svoboda, K. (2010). Nature *464*, 1182–1186.

Li, N., Chen, T.W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). Nature *519*, 51–56.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Nature *532*, 459–464.

Makino, H., Ren, H., Liu, H., Kim, A.N., Kondapaneni, N., Liu, X., Kuzum, D., and Komiyama, T. (2017). Neuron *94*, this issue, 880–890.

Rat mPFC and M2 Play a Waiting Game (at Different Timescales)

Angela J. Langdon,[1,*] Andrew M. Wikenheiser,[2] and Geoffrey Schoenbaum[2,*]
[1]Princeton Neuroscience Institute, Princeton University, Princeton, NJ, 08544, USA
[2]National Institute on Drug Abuse, Baltimore, MD, 21224, USA
*Correspondence: alangdon@princeton.edu (A.J.L.), geoffrey.schoenbaum@nih.gov (G.S.)
http://dx.doi.org/10.1016/j.neuron.2017.05.006

In this issue of *Neuron*, Murakami et al. (2017) relate neural activity in frontal cortex to stochastic and deterministic components of waiting behavior in rats; they find that mPFC biases waiting time, while M2 is ultimately responsible for trial-to-trial variability in decisions about how long to wait.

Good things come to those who wait—or so the proverb tells us. Yet animals and humans frequently choose to give up waiting for delayed reward, even when patience will ultimately maximize reward over the long run. The predictable (or, *deterministic*) element of the decision to abort waiting is thought to reflect learned expectations regarding the likely timing of outcomes and their value, including an inherent tendency to discount future reward. But like other choice behaviors, decisions about when to give up waiting are also highly variable, reflecting an apparently random (or, *stochastic*) element in the choice process that limits the ability to control both the selection and timing of actions. While the random element in choice is thought to be adaptive (for example, by promoting the exploration of previously unchosen actions), and comprises a key feature of many theoretical accounts of choice behavior, it remains a challenge to parse out the neural origin and mechanisms of variability in decisions about when to act.

Frontal regions of the brain, including medial prefrontal cortex (mPFC), are critically important for complex decision making generally (Dalley et al., 2004; Rushworth et al., 2011), including in decisions to wait for delayed reward (Narayanan and Laubach, 2006). Other frontal areas more proximal to primary motor output, such as M2 in rodents (or pre-SMA in primates), are also implicated in

decisions about when to act (Murakami et al., 2014). Despite this apparent overlap in function, and direct reciprocal connections between mPFC and M2, mPFC has been more extensively studied in interaction with brain areas such as the striatum and the amygdala, nodes that are generally thought to support associative learning (Cardinal et al., 2002; Kable and Glimcher, 2009). This focus has left open the question of how mPFC might interact with motor areas such as M2, particularly in the timing of actions and in deciding whether to wait for reward.

In a detailed study published in this issue of *Neuron*, Murakami et al. (2017) provide a clear, and perhaps surprising, answer to this question. To test the precise relationship between neural activity in frontal circuits and decisions about the timing of actions, the authors recorded and analyzed neural activity in mPFC and M2 of rats as they performed a waiting task. By patiently waiting with their snout in a port for an unpredictable duration, signaled by the delay between two tones, the rats could obtain a large water reward in a separate port. Deciding to abort waiting, by leaving the port after the first tone but before the second, resulted in a much smaller water reward. Throughout the task, the total duration of each trial was fixed, such that to maximize reward, the rats had only to wait for the second tone. Yet, despite this, the rats displayed a broad distribution of waiting times on this task, frequently aborting after apparently deciding to wait for the larger delayed reward.

Murakami et al. (2017) showed that the high degree of variability in how long the rats were willing to wait reflected the recent trial history of waiting times and reward combined with an apparent random process. They proposed a two-stage decision model in which stochastic influences are added to a learned deterministic wait time bias to produce the actual decision about how long to wait on any given trial. After pharmacologically inactivating mPFC and M2 to confirm their relevance for waiting behavior in their task, they posed the question: how does neural activity in mPFC and M2 relate to the two sources of variability governing the actual wait time on a given trial?

Strikingly, Murakami et al. (2017) found that activity in M2, but not in mPFC, represented the stochastic variation in how long the rats actually waited for reward on a given trial. Single units in M2 displayed patterns of spiking activity significantly correlated with trial-by-trial variation in wait times, and low-dimensional patterns in M2 population activity also correlated with actual wait times. In contrast, both M2 and mPFC contained single-unit and population-level activity correlated with the deterministic wait time bias, related to the recent history of wait times and reward experienced during the task. Even more interestingly, the effect of the deterministic bias on neural activity was dramatically different in the two areas. In M2, this information was represented briefly, in phasic, event-related activity inside the trial, whereas in mPFC, its representation extended throughout the trial, and into the inter-trial interval.

What does this apparent regional and temporal dissociation between stochastic and deterministic decision-related activity mean for the influence of frontal circuits on decisions to wait? One clear interpretation made by Murakami et al. (2017) is a model in which mPFC sets a general propensity to wait, based on recent experience, and then delegates more immediate and direct control of the timing of action to pre-motor areas such as M2. This implies that prefrontal areas are more important for monitoring and making longer-term behavioral adjustments, in line with their role in learning and representing the statistical structure of a task, including tracking changes in the contingency between actions and reward, while areas such as M2 that are closer to the actual control of the actions have a more direct influence on behavior in real time, including perhaps generating stochastic variance in behaviors that facilitate exploration and learning. One interesting prediction of this account is that the stochastic variability in the behavior and also in the neural activity in M2 might thus wane as the response becomes more ingrained, less exploratory, and more automated.

This interpretation also raises interesting questions for how we understand mPFC to control behavior. First, is it the case that the influence of mPFC is limited to this long-term biasing function, or might a task with different requirements promote a more immediate influence over the timing of actions? For example, given that prefrontal regions are crucial for abstract, rule-guided behavior, might one see a more immediate and punctate influence of mPFC in the current task if there were a stop signal or some other immediate cue that required the rat to redirect the default behavioral strategy? Such a cue would alter the rules governing the otherwise more automated behavior being implemented on each trial by downstream regions like M2. In theory, mPFC—at least parts of it—might be recruited by such a manipulation.

Second, is mPFC's involvement dependent on the precise statistical relationship between events in a task? Exponentially distributed events, such as the signaled wait times in this task, are characterized only by their timescale, rendering individual events largely unpredictable. An interesting extension to the present study would be to test how these behavioral and neural findings generalize to more structured distributions of signaled wait times. A bimodal distribution, for instance, might be hypothesized to promote the formation of an alternate, perhaps less variable, behavioral strategy: reliably wait for the shorter durations, then abort if these durations elapse with no completion signal. Indeed, in humans performing a similar waiting task with a bimodal distribution of reward times, mPFC was found to track the dynamic re-evaluation of whether to continue waiting, based on the elapsed time during a trial and learned expectations about the temporal distribution of reward (McGuire and Kable, 2015).

Finally, and relatedly, direct prefrontal control of action timing might be preferentially engaged in tasks in which there is more predictable structure in action-outcome contingencies to be behaviorally exploited. Such conditions are thought to distinguish goal-directed behavior, which promotes responding on the basis of response-outcome associations, from habitual or automatic responding, which is based on stimulus-response associations (Balleine and Dickinson, 1998). Importantly, it is known that unpredictable relationships between actions and reward favor the development of habitual responding. Indeed, habits are typically trained through variable interval reinforcement schedules with a similar distribution to the exponentially distributed wait times in the present study. Different regions

within mPFC, namely prelimbic (PL) and infralimbic (IL) cortex, are thought to have contrasting roles in habitual versus goal-directed action (Killcross and Coutureau, 2003). These different forms of behavioral control might bring out differential patterns of neural activity in PL versus IL, not apparent in the current task, that would be informative as to the general roles of mPFC subregions in behavior.

The findings in this study are thought provoking and stand as a counterpoint to the assumption of fine-grained behavioral control residing exclusively in prefrontal cortex. They show that at least in some cases, direct influence over behaviors typically associated with cognitive control and prefrontal function is delegated to brain regions more proximal to action execution itself. In fact, rather than being a mere effector of top-down signals from mPFC, when it comes to deciding exactly how long to wait, M2 seems to be in control.

ACKNOWLEDGMENTS

This work was supported by the Intramural Research Program at the National Institute on Drug Abuse, grant number ZIA-DA000587.

REFERENCES

Balleine, B.W., and Dickinson, A. (1998). Neuropharmacology *37*, 407–419.

Cardinal, R.N., Parkinson, J.A., Hall, J., and Everitt, B.J. (2002). Neurosci. Biobehav. Rev. *26*, 321–352.

Dalley, J.W., Cardinal, R.N., and Robbins, T.W. (2004). Neurosci. Biobehav. Rev. *28*, 771–784.

Kable, J.W., and Glimcher, P.W. (2009). Neuron *63*, 733–745.

Killcross, S., and Coutureau, E. (2003). Cereb. Cortex *13*, 400–408.

McGuire, J.T., and Kable, J.W. (2015). Nat. Neurosci. *18*, 760–766.

Murakami, M., Vicente, M.I., Costa, G.M., and Mainen, Z.F. (2014). Nat. Neurosci. *17*, 1574–1582.

Murakami, M., Shteingart, H., Loewenstein, Y., and Mainen, Z.F. (2017). Neuron *94*, this issue, 908–919.

Narayanan, N.S., and Laubach, M. (2006). Neuron *52*, 921–931.

Rushworth, M.F., Noonan, M.P., Boorman, E.D., Walton, M.E., and Behrens, T.E. (2011). Neuron *70*, 1054–1069.