

# Confidence and Corrections: How We Make and Un-Make Up Our Minds

John Pearson<sup>1,\*</sup> and Michael L. Platt<sup>1,\*</sup>

<sup>1</sup>Center for Cognitive Neuroscience and Department of Neurobiology, Duke University, Durham, NC 27708, USA

\*Correspondence: [pearson@neuro.duke.edu](mailto:pearson@neuro.duke.edu) (J.P.), [platt@neuro.duke.edu](mailto:platt@neuro.duke.edu) (M.L.P.)

DOI 10.1016/j.neuron.2009.09.011

Single neurons in several brain areas intervening between sensation and action signal the accumulation of sensory evidence favoring a particular behavioral response. Two new studies show that these same neurons encode decision confidence and that decision makers continue to process relevant sensory information even after a choice has been made.

The standard paradigm for probing sensory decision making, the two-alternative forced choice task (2AFC), bears striking resemblance to the problem confronted by a soccer goalkeeper defending a penalty kick—albeit without the legions of noisy fans. The average penalty kick clocks in at a speed of around 20 m/s, leaving the average goalkeeper, stationed 11 m away, somewhere in the neighborhood of half a second to decide whether to dive left or right. Before the ball reaches the goal line, the keeper must use sensory information to estimate its speed and trajectory, project its likely location by the time it reaches him, and plan a dive that will intercept it in time. Intuitively, the more obvious the kicker's movements, the faster and more accurately the keeper will react. And, indeed, psychophysical, computational, and neurophysiological studies have endorsed this intuition by showing that the neural processes responsible for rendering such a decision integrate sensory evidence over time, initiating action when a threshold is reached.

As informative as such studies have been, they tend to overlook more nuanced aspects of decision making. For example, because the keeper's reaction time is nearly as long as the travel time, he must often act on the basis of highly ambiguous information, even though the result may be no better than a guess. Even worse, a kick may bend mid-air on its way toward the goal, requiring a nearly impossible mid-air correction. No doubt the speed of the ball and the ability of the kicker to disguise his shot will influence both the keeper's confidence in his decision and his ability

to change his dive on the fly to meet the ball.

Two recent papers examined these important aspects of decision making that are often obscured in studies utilizing the standard 2AFC paradigm (Kiani and Shadlen, 2009; Resulaj et al., 2009). Both studies used modified versions of a now-classic 2AFC task used to study visual motion discrimination (Shadlen and Newsome, 2001). Typically, subjects are asked to judge the direction of motion of a field of dots moving stochastically on a computer monitor. On average, the dots move either right or left, though the coherence—the percentage of dots moving the same direction at any given time—may vary from just above 0% (random) to just below 100% (uniform motion). Clearly, the longer a subject is allowed to watch the display, the more information he obtains, and the more coherent the stimulus, the less time is needed to reach a conclusion.

Neurophysiological experiments in monkeys, who reported the direction of motion by shifting gaze to one of two response targets (Gold and Shadlen, 2002; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001), have focused on the responses of neurons in brain areas linked to either decoding the direction of motion (e.g., the middle temporal area, MT) or planning the eye movement response (e.g., the lateral intraparietal area, LIP). Neurons in MT, for example, tend to fire in direct proportion to the coherence of dot motion in their preferred direction. LIP neurons (as well as neurons in other areas within the visual orienting system), by contrast, show a ramp up in firing for motion favoring eye movements

in their preferred directions. Importantly, the slope of this ramp in firing varies as a direct function of dot coherence—implying integration of sensory evidence with respect to time. Moreover, firing rates tend to saturate at a fixed level that strongly predicts the time it takes the monkey to make a choice. In what appears to be a profound case of kismet, this pattern of firing is precisely what is predicted by computational models of sensory decision making developed solely to explain variation in decision accuracy and reaction time by human observers (Gold and Shadlen, 2007; Ratcliff, 1978; Shadlen et al., 2006; Smith and Ratcliff, 2004).

However, just as the goalkeeper trying to stop the penalty kick must dive either left or right, regardless of the quality of his evidence, 2AFC tasks that allow only a fixed amount of time for sensory integration may obscure much of the information encoded in the underlying neural signal. That is, while a 500 ms observation of dot motion may result in a saccade to the right for both 5% and 75% motion strengths, the animal's underlying confidence about the correctness of each response may be drastically different. In fact, in a study utilizing a time-limited version of the dot-motion task, error trials were typically preceded by periods in which the accumulated firing rates in LIP neurons remained low (Roitman and Shadlen, 2002). If, as claimed, these firing rates represented the ratio of evidence favoring a saccade in the neuron's preferred direction, low firing would imply roughly equal evidence in favor of either alternative, suggesting that most errors corresponded to decisions made with low confidence.

In a recent paper, Kiani and Shadlen tested this hypothesis by offering monkeys a third option in the time-limited version of the dot motion task while recording from neurons in LIP (Kiani and Shadlen, 2009). On a select subset of trials, a third target appeared following the stimulus presentation; selecting this target always resulted in a reward, albeit a smaller one than for a correct answer. This is the equivalent of giving the keeper a pass when he feels uncertain about where the kick will go. Intuitively, monkeys will select the opt-out target when they are less confident about what they have seen. As expected, monkeys selected this “sure” target more often for weaker motion strengths and shorter stimulus durations, resulting in improved choice accuracy across all other trials, implying that opt-out choices were made when the evidence was weaker and confidence lower. Moreover, by allowing the animals to opt-out of the discrimination task, the researchers could determine which patterns of firing were associated with low-confidence responses. Consistent with the temporal integration hypothesis, firing rates on opt-out trials remained intermediate between very high (strong evidence favoring eye movements in the neuron’s preferred direction) and very low (strong evidence favoring eye movements opposite the neuron’s preferred direction). Moreover, firing rates were highest when monkeys shifted gaze in the preferred direction of the neuron to select the sure target, again consistent with the idea that LIP neurons signal the likelihood of that choice.

Though the three-option task elicited only an implicit behavioral measure of confidence, the Kiani and Shadlen results confirm the predictions of the evidence accumulation model and add to a growing body of data on confidence-related neuronal activity elsewhere in cortex (Hayden et al., 2008; Kepecs et al., 2008). More importantly, they are the first to demonstrate how both a decision and its attendant confidence judgment can be encoded in a single data stream—precisely what would be expected from a normative, Bayesian decision making mechanism. How such a code might relate to the subjective experience of confidence remains an open question, but this and similar studies inform our understanding

of how the brain organizes the information available to it in the moments preceding a decision.

Just as crucial as the moments preceding a decision, however, are the moments following, in which organisms can sometimes adjust their actions in light of new information. In contrast with forced choice tasks, most decisions in real environments are not made in isolation, but unfold as part of complex foraging or mating policies that allow for adjustment, update, and learning (Stephens and Krebs, 1986). For example, a gazelle moving through open grassland can choose continually between attentive listening and eating as she scouts for predators. Although the keeper can often only look on helplessly in mid-air as the ball curves toward the other side of the net, sometimes, after committing to a course of action, the rapid accumulation of new information permits him to alter course mid-stride, changing action plans in response to updated evidence.

Resulaj and collaborators studied a striking example of this phenomenon in human subjects performing the dot motion discrimination task (Resulaj et al., 2009). As with monkeys, subjects were given up to 2 s of stimulus observation, but unlike the monkeys, they registered their decisions by moving a handle to fixed locations to the left or right in front of them. By requiring subjects to use the comparatively slower arm movements, obscuring these movements from view, and turning off the stimulus as soon as subjects began to reach, the experimenters created a situation in which movement trajectories might be altered by unprocessed visual information still making its way through the visual stream at the time of decision.

And indeed, this is exactly what Resulaj and collaborators observed. On most trials, subjects’ arm movements followed trajectories curved toward their eventual choice targets, but on a small subset, the arcs began heading in one direction before suddenly veering off in the opposite direction—suggesting a change of mind. This process, they showed, increased the overall accuracy of final decisions across all motion strengths and was correlated with stimulus information present up to 200 ms before subjects started to move. These data suggest that,

in the time between decision and action, the brain’s decision system continues to process stimulus information entering through the visual stream and integrates this new evidence during the course of movement.

To explain these postdecision corrections, the researchers proposed the existence of a second, postdecision “change-of-mind” bound. After accumulating evidence reached the threshold to initiate a movement, they argued, subjects continued to process further stimulus information, with actions reversing course if the subsequent evidence signal changed more than a certain percentage in favor of the alternative. Interestingly, fits of this model to the behavioral data also suggested a change-of-mind deadline, a point of no return past which new evidence could not alter subjects’ actions. More broadly, the data imply that the brain does not immediately dispose of accumulated information after a decision is made but maintains and updates these signals online for the purpose of altering ongoing behavior. Such findings not only place decision making within a more open-ended, naturalistic context, but also emphasize its connection with motor planning circuitry and control theory (Shadmehr and Wise, 2005).

Together, these studies contribute to an emerging picture of decision making in which the decision itself takes place amid more general processes of information gathering and online control. By offering insights into the brain’s capacities for graded judgments and self-correction, they may shed new light on the origins of phenomena such as delusions and fixed ideas that attend disorders such as schizophrenia and obsessive-compulsive disorder.

#### REFERENCES

- Gold, J., and Shadlen, M. (2002). *Neuron* 36, 299–308.
- Gold, J., and Shadlen, M. (2007). *Annu. Rev. Neurosci.* 30, 535–574.
- Hayden, B., Nair, A., McCoy, A., and Platt, M. (2008). *Neuron* 60, 19–25.
- Kepecs, A., Uchida, N., Zariwala, H., and Mainen, Z. (2008). *Nature* 455, 227–231.
- Kiani, R., and Shadlen, M. (2009). *Science* 324, 759.
- Ratcliff, R. (1978). *Psychol. Rev.* 85, 59–108.

Resulaj, A., Kiani, R., Wolpert, D., and Shadlen, M. (2009) *Nature*, in press. Published online August 19, 2009. 10.1038/nature08275.

Roitman, J., and Shadlen, M. (2002). *J. Neurosci.* 22, 9475–9489.

Shadlen, M., and Newsome, W. (2001). *J. Neurophysiol.* 86, 1916–1936.

Shadlen, M., Hanks, T., Churchland, A., Kiani, R., and Yang, T. (2006). *Bayesian Brain: Probabilistic Approaches to Neural Coding*, K. Doya, S. Ishii, A. Pouget, and R.P.N. Rao, eds. (Cambridge, MA: MIT Press), pp. 209–237.

Shadmehr, R., and Wise, S. (2005). *The Computational Neurobiology of Reaching and Pointing: A*

*Foundation for Motor Learning* (Cambridge, MA: MIT Press, Bradford Books).

Smith, P., and Ratcliff, R. (2004). *Trends Neurosci.* 27, 161–168.

Stephens, D., and Krebs, J. (1986). *Foraging Theory* (Princeton, NJ: Princeton University Press).