

The uncertainty of it all

Adam Kepecs

Two studies in this issue report the discovery of different types of uncertainty signals in little-studied, but critical, regions of the forebrain: decision confidence in the pulvinar and reward risk in the anterodorsal septum.

The world is fraught with uncertainty. External events are often unpredictable and our own brain processes may be noisy, generating additional uncertainty. Consequently, knowing one's degree of uncertainty confers benefits for a broad range of activities from the sophisticated to the mundane: managing a stock portfolio or deciding whether to carry an umbrella. Conversely, the pathological mis-evaluation of uncertainty contributes to a wide range of neuropsychiatric conditions, including anxiety, obsessive-compulsive disorder and addiction. Psychologists, behavioral ecologists and economists have long studied how we estimate uncertainty and deploy these estimates to guide behavior. For instance, when one is uncertain, learning rates should be boosted and attention enhanced to gather more information and thereby reduce the uncertainty¹. Recently, neuroscientists have joined the quest to study uncertainty and are starting to explore the contributions of different brain areas. In this issue of *Nature Neuroscience*, two studies report that distinct brain regions represent different facets of uncertainty: risk and confidence. Monosov and Hikosaka report the seminal discovery that most neurons in a hitherto little-examined forebrain nucleus, the anterodorsal septum (ADS), selectively signal reward uncertainty². And, in an elegant and impressive piece of work, Komura and colleagues find that confidence in a perceptual choice is represented in the pulvinar³, a higher order thalamic nucleus.

The concepts of 'risk' and 'confidence' can be defined in the context of human subjective experience and have been well studied using imaging approaches^{4–6}. How can these be studied in animals to dissect their underlying neural circuitry? First, such studies require behavioral tasks that incentivize an animal to evaluate and act on its own certainty. Second, they require computational models to link neural activity to hidden behavioral variables: unobservable internal variables such as uncertainty or attention. Traditional approaches rely on behavioral correlates—trial-by-trial correlations of neural activity

with observable aspects of behavior, such as choice or reaction time—but these are often difficult to interpret because they can arise for several reasons⁷. Recent work has risen to this challenge with quantitative models that describe trial-to-trial variation in animal behavior from observable variables in such a way that the internal variables arising from a model can be used as proxies for the underlying unobservable decision variables (for example, ref. 5). These methodological developments have enabled researchers to study the hidden variables underlying decisions, such as the utility or value of rewards, and are increasingly used to tackle other decision variables. The advantage of formal models is that they yield concrete predictions that are testable using traditional neurophysiological and behavioral tools without having to resort to semantic definitions or psychological labels. Applying computational models to processes such as uncertainty enables us to bridge the gap between the mental processes we seek to explain and the underlying computations realized at the level of neuronal circuits.

This approach has already gained considerable traction and has been used to demonstrate that some neurons signal the variance of reward payoff, a form of uncertainty that is also technically known as risk^{8–10}, and the degree of belief in the correctness of a decision, called confidence^{11–13}. But many of these neurons have also been observed to carry other signals, and it remains unclear whether any neuron encodes (or computes) a pure representation of these variables.

Monosov and Hikosaka² trained monkeys to watch a video screen and learn to associate different visual cues with variable probabilities of either appetitive or aversive outcomes (Fig. 1a,b). In one block of trials, the authors provided five distinct visual cues, each of which predicted the upcoming delivery of a drop of apple juice reward at different probabilities ranging from 0 to 1. In addition to the rewarded block, monkeys encountered an aversive block in which five distinct visual cues predicted different probabilities of an air puff directed at the monkey's face.

Initial experiments focused on the frontal part of the macaque brain. The authors systematically recorded neurons over a large area

during behavior. This turned up a surprise in the dorsal septum, a region that is directly connected to the hypothalamus, amygdala and cingulate gyrus, and has been thought to be involved in appetitive processes, but has remained largely *terra incognita*. The authors stumbled onto a subregion, the ADS, in which neurons responded specifically to cues that predicted uncertain rewards (25, 50 or 75%) and remained entirely silent when presented with the sure cues (0 or 100%). These neurons again remained silent in the aversive condition, establishing that ADS signals uncertainty about positive outcomes, but not about negative ones.

Uncertainty in this situation may be defined in different ways, such as the Shannon entropy, $p_{\text{reward}} \log p_{\text{reward}}$, or economic risk⁴, defined as the variance of the reward payoff distribution, in this case $p_{\text{reward}}(1 - p_{\text{reward}})$ (Fig. 1c). These definitions yield very similar predictions, with zero uncertainty at 0 or 100% cues and peak uncertainty at 50%, which is precisely what Monosov and Hikosaka² observed in the ADS (Fig. 1d). Can we tease apart which variety of uncertainty the neurons represent? Uncertainty about outcomes, as measured by Shannon entropy, depends only on the probability of reward, whereas risk-type uncertainty in economics depends on the variance of the reward distribution. Thus, Monosov and Hikosaka² introduced a new cue that predicted two drops of juice and compared it to a cue predicting a single drop with the same probability. They found that ADS neurons increased their firing for double rewards that were uncertain (Fig. 1d). In other words, these neurons carried representations not only about the reward probability, but also about reward size, consistent with the economic definition of risk, the variance of the reward distribution. More work will be required to determine whether this is true quantitatively and under a broader range of circumstances. Finally, the authors showed that ADS neurons can rapidly learn to predict reward risk as novel visual cues are introduced, each predicting distinct degrees of uncertainty.

The concept of risk discussed above is related to the prediction of uncertain outcomes that depend on external events that can be computed from the recent history

Adam Kepecs is at Cold Spring Harbor Laboratory, Cold Spring Harbor, New York.
e-mail: kepecs@cshl.edu

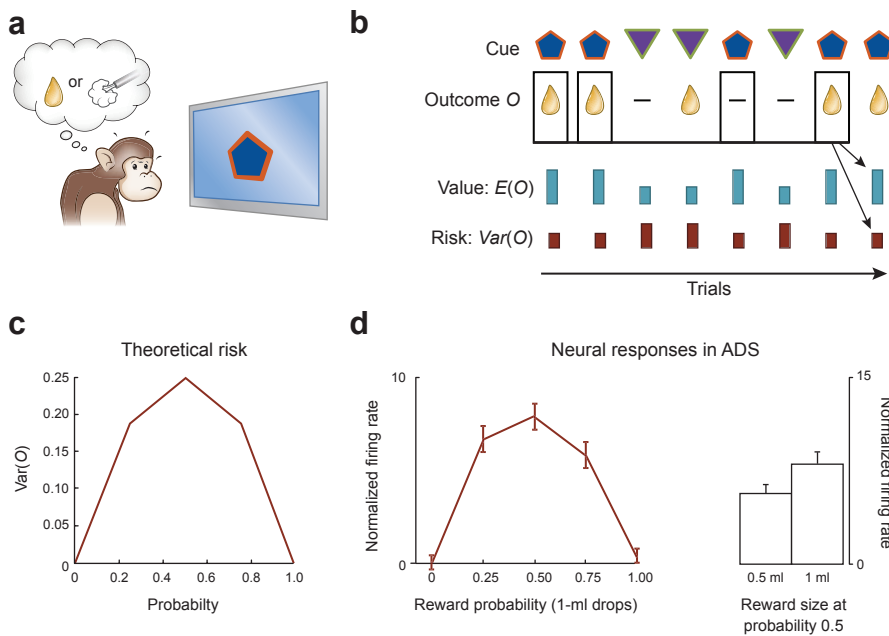


Figure 1 Reward risk-type uncertainty is represented in the ADS. (a) Monkeys were trained in a classical conditioning task to anticipate either appetitive or aversive outcomes (a drop of juice or an air puff) of different probabilities and amounts on the basis of visual cues. (b) The value and risk of outcomes O can be estimated by computing the mean, $E(O)$, and variance, $Var(O)$, of the recent history of outcomes associated with a given cue. For instance, if the pentagon-shaped cue is associated with a drop of juice in three out of four presentations, then $E(O) = 0.75$ and $Var(O) = 0.1875$. (c) Economic risk is defined as the variance of the reward outcome distribution. $Var(O)$ is highest for probability $p = 0.5$ and lowest when $p = 0$ or 1 . (d) ADS neurons fire in proportion to reward risk. For a fixed reward size, they fire as an inverted U-shaped function of reward probability (left), and for a given probability (0.5), they fire in proportion to reward size (right). Error bars indicate s.e.m. Panel **d** adapted from ref. 2 with permission.

of reward outcomes (Fig. 1b). Another type of uncertainty is decision confidence, a subjective estimate or belief about the outcome of a decision. Because decision confidence reflects noise internal to the decision process, even identical stimuli can generate different estimates about the uncertainty of decision outcomes.

Komura *et al.*³ tackled this issue, the neural representation of confidence in a perceptual choice. Confidence judgments are usually studied by means of explicit self-reports in humans^{6,14}. How can such a subjective state be studied rigorously in animals, where verbal confidence reports are unavailable? First, such studies require a suitable behavioral task that calls for the use of confidence information^{11,12}. Second, they require a model that can link confidence with observables, such as stimuli and outcomes.

The authors trained monkeys on a motion categorization task (Fig. 2a). On each trial, the monkeys viewed a cloud of dots, each either red or green, that were moving up or down. The monkeys had to report the predominant direction of the target color that was cued. When all of the dots were moving in one direction, the choice was easy, whereas when the mixture of upward and downward moving dots was nearly balanced, the choice was more

difficult. Correct choices were rewarded with a drop of juice. To make use of confidence information, the monkeys also had a third choice: to opt out of the categorization task and receive a smaller, but certain, reward¹². The idea was that if a monkey is confident, he should accept the categorization challenge and make a risky choice, whereas if he is less confident, he should choose the safe option. And indeed, the monkeys did choose the safe option only when presented with difficult stimuli.

Next, the authors used a theoretical model to relate the firing of their neurons to confidence. Signal detection theory provides a language for analyzing perceptual decisions under uncertainty and has been used to derive predictions for decision confidence¹¹. In this model, an observer must discriminate between a stimulus, s , and a category boundary, b . In each trial, the observer perceives s drawn from a normal distribution (Fig. 2b,c). The choice can be made by comparing whether $s < b$ or $b > s$. Confidence, d , is a function of the distance between these variables: $d = |s - b|$. This simple model for confidence predicts a distinctive and counterintuitive pattern: confidence decreases with increasing stimulus difficulty for correct choices and increases for error choices (Fig. 2c, bottom left). Examining

the stimulus and boundary configurations that could lead to a given choice offers an intuition behind this (Fig. 2c, top). For correct choices, the distance between stimulus distribution and the category boundary increases as the stimulus becomes easier. For error choices, however (which happen when a stimulus is perceived to be on the wrong side of the boundary), the distance between sampled stimulus and categorization boundary tends to be smaller for easy stimuli because the overlapping area of the two distributions becomes smaller. In other words, for easier stimuli, errors are rare, but in those cases in which they do occur, the decision maker cannot have been very confident.

Komura *et al.*³ decided to record neurons from the pulvinar, a higher order visual thalamic nucleus that has been implicated in attention, but whose function has remained largely mysterious. Notably, most pulvinar neurons precisely followed the predictions of this model for confidence (Fig. 2d, left). Can this model also account for opt-out choices? Confidence derived from this model on each trial is, as intuitively expected, a good predictor of average choice accuracy¹⁴. Thus, by setting an appropriate confidence threshold, this estimate can be used to decide whether it is more valuable to opt out or to make a perceptual choice (Fig. 2c, bottom right). Indeed, the firing rates of pulvinar neurons predicted choice accuracy and the probability of opting out, again consistent with this model (Fig. 2d, right).

Finally Komura *et al.*³ asked whether activity in the pulvinar contributes to opt-out choices. When they reversibly inactivated the lateral geniculate nucleus, the primary thalamic relay for visual information, they found that perceptual performance dropped and the number of opt-out choices increased, consistent with the idea that this pathway is responsible for relaying the perceptual content in a visual discrimination. In contrast, when they inactivated the pulvinar, perceptual decisions were unaffected, but the number of opt-out choices increased, consistent with the idea that this selectively affected decision confidence. The theoretical model cannot, however, rule out an alternative interpretation: that the opt-out threshold changed with this manipulation (Fig. 2c, bottom right) and the confidence reporting ability remained intact. For instance, if pulvinar inactivation either lowered risk-taking propensity or reduced attention, both could lead to an increased rate of opting out. Indeed, opt-out tasks are not ideal for studying confidence reporting because each trial only provides either a perceptual choice or an opt-out choice, making it difficult to draw firm behavioral conclusions. This highlights the

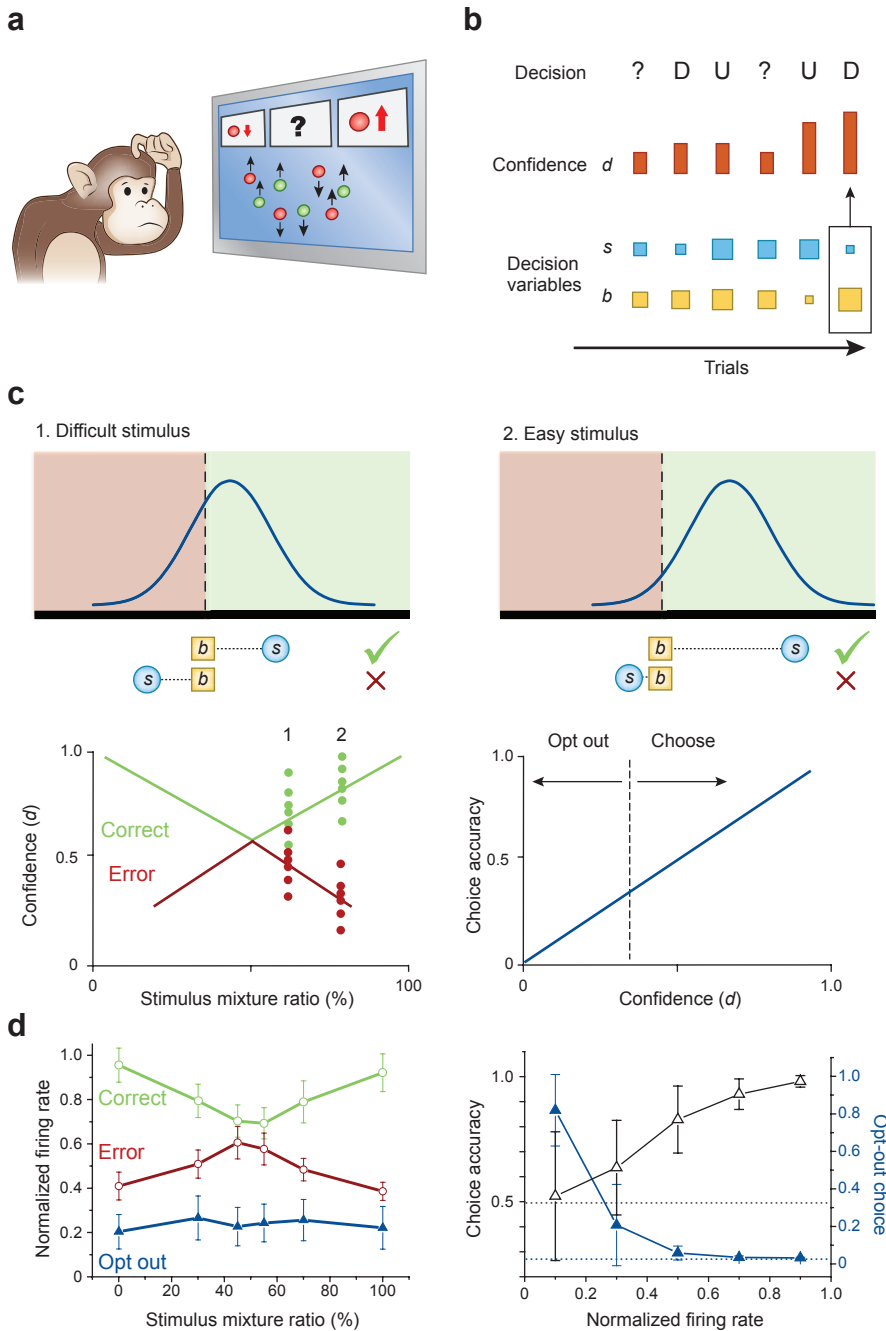


Figure 2 Decision confidence is represented in the pulvinar. **(a)** A three-alternative, forced-choice visual discrimination task, with two categorical choices based on the stimulus (whether more red dots in the stimulus mixture are moving upward or downward) and an unsure opt-out option that offers a certain, but small, reward. **(b)** Confidence can be computed as a single-trial estimate on the basis of the decision variables of the current trial. Decisions can be downward (D), upward (U) or opt out (?). On a given trial, decision confidence d can be estimated by computing the distance between the perceived stimulus, s , and the category boundary, b . **(c)** Computing decision confidence¹¹. Top, we assume the memorized value of b (dashed line) to be constant. Error choices occur when s is to the left of b (where the stimulus distribution extends into the red shaded region). When the stimulus distribution is easier (right), the red region under the curve shrinks and the green region expands. Thus, the maximum distance between s and b for error choices is lower and the maximum for correct choices is higher, so confidence estimates average lower for the rare easy-stimulus errors than for difficult-stimulus errors. Bottom left, confidence d estimated in each trial can be used to predict the mean choice accuracy. Bottom right, confidence threshold (dashed line) can be used to decide whether to opt out or make a category decision. **(d)** Neural activity in the pulvinar follows confidence (compare **c**, left) and predicts choice accuracy and opt-out choices (compare **c**, right). Error bars indicate s.d. across 72 neurons. Panel **d** adapted from ref. 3 with permission.

issue that behavioral task design remains critical even when using model-based interpretations. Despite these caveats, this is a groundbreaking study that opens a new avenue for understanding the functions of the pulvinar.

These two studies provide much-needed information about the neural processing of uncertainty. As with most exciting discoveries, they raise more questions than they answer. The ADS is well connected to brain regions that are involved in anxiety and learning. Thus, the finding that the ADS conveys information about reward risk suggests that it may be a neural substrate for how uncertainty induces anxiety and boosts learning. In contrast, the dorsal pulvinar is connected to a wide range of areas, including higher order visual, prefrontal and parietal cortices, and reports decision confidence, a subjective form of uncertainty reflecting belief in a decision. The functions of the pulvinar have remained largely unknown, although it has often been implicated in attention¹⁵. Some forms of attention appear to be directly driven by uncertainty¹, and hence confidence signals in the pulvinar could support an attentional function. More generally, one of the key uses of estimating one's own uncertainty is to drive different types of information-seeking behaviors, exploration, learning and attention, so as to reduce uncertainty. These distinct representations of uncertainty in the ADS and pulvinar raise the question of how and where these representations are integrated in a coherent way to guide behavior.

COMPETING FINANCIAL INTERESTS

The author declares no competing financial interests.

- Dayan, P., Kakade, S. & Montague, P.R. *Nat. Neurosci.* **3**, 1218–1223 (2000).
- Monosov, I.E. & Hikosaka, O. *Nat. Neurosci.* **16**, 749–755 (2013).
- Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T. & Miyamoto, A. *Nat. Neurosci.* **16**, 756–762 (2013).
- Huetzel, S.A., Stowe, C.J., Gordon, E.M., Warner, B.T. & Platt, M.L. *Neuron* **49**, 765–775 (2006).
- Sugrue, L.P., Corrado, G.S. & Newsome, W.T. *Science* **304**, 1782–1787 (2004).
- De Martino, B., Fleming, S.M., Garrett, N. & Dolan, R.J. *Nat. Neurosci.* **16**, 105–110 (2013).
- Maunsell, J.H. *Trends Cogn. Sci.* **8**, 261–265 (2004).
- Fiorillo, C.D., Tobler, P.N. & Schultz, W. *Science* **299**, 1898–1902 (2003).
- McCoy, A.N. & Platt, M.L. *Nat. Neurosci.* **8**, 1220–1227 (2005).
- O'Neill, M. & Schultz, W. *Neuron* **68**, 789–800 (2010).
- Kepecs, A., Uchida, N., Zariwala, H.A. & Mainen, Z.F. *Nature* **455**, 227–231 (2008).
- Kiani, R., & Shadlen, M.N. *Science* **324**, 759–764 (2009).
- Middlebrooks, P.G. & Sommer, M.A. *Neuron* **75**, 517–530 (2012).
- Rahnev, D.A., Maniscalco, B., Luber, B., Lau, H. & Lisanby, S.H. *J. Neurophysiol.* **107**, 1556–1563 (2012).
- Desimone, R., Wessinger, M., Thomas, L. & Schneider, W. *Cold Spring Harb. Symp. Quant. Biol.* **55**, 963–971 (1990).

