

which is proportional to the number of photons in the cavity. Such a scheme was proposed a few years ago by the ENS group⁴, but it has not been implemented yet, because the fact that the atom–field interaction is non-resonant makes it very weak. In the present experiment¹, a different trick is used: the interaction is fully resonant and thus much larger. An energy exchange does occur, but the parameters are chosen so that this energy exchange is fully reversible.

To understand how this exchange happens, we need to look briefly at concepts developed in the field of cavity QED. To cut a long story short, cavity QED uses optical or microwave resonators designed in such a way that the ‘electric field per photon’ is extremely large. Consequently, the coupling rate between an atom and the field is also very large, much larger than the dissipative effects caused by spontaneous emission from the atom, or the electric field leaking out of the cavity. So when the atom is within the cavity, the system can be described by only considering the coherent coupling between the atom and the cavity under various conditions. If there is no photon in the cavity, nothing happens; if there is one photon in the cavity, it is coherently absorbed and re-emitted by the atom before it leaves the cavity. In the latter case, the net energy exchange is zero, but it can be shown that because of this ‘Rabi cycle’ a phase shift occurs in the atomic wavefunction. Finally, if there is more than one photon, more Rabi cycles will occur, but this case can be ignored by the experimenters.

If one considers only the cases where the cavity is initially in an arbitrary quantum superposition, or mixture, of zero-photon and one-photon states, it is clear that this set-up will behave in a QND way. The phase information is extracted from the ‘meter’ atom using an interference effect, which transforms the phase shift into a detectable change in the atomic level. In fact, several schemes have been used by the ENS team¹: either a first atom is used to deposit one photon within the cavity and a second atom detects it; or both atoms perform two successive measurements of a weak initial field, repeatedly allowing the experimenter to know whether the cavity contains zero photons or one photon.

The ENS format¹ cannot easily be generalized to higher photon numbers (this would require using the non-resonant interaction described above), but it is worth pointing out that the techniques that yield this single-photon QND measurement may also generate other interesting effects. The fact that one photon can significantly shift the atomic phase is the key to building a quantum logic gate, suggesting a route to quantum computing⁵. It can be expected that this experiment will open the way to many other ones, which will explore further the very peculiar rules

for writing in and reading out the information encoded in quantum objects. □

Philippe Grangier is at the Institut d’Optique, BP 147-F91403, Orsay, France.

e-mail: philippe.grangier@iota.u-psud.fr

1. Nogues, G. *et al.* *Nature* **400**, 239–242 (1999).

2. Braginsky, V. B., Vorontsov, Y. L. & Thorne, K. S. *Science* **209**, 547–557 (1980).
 3. Grangier, P., Levenson, J. A. & Poizat, J.-P. *Nature* **396**, 537–542 (1998).
 4. Brune, M., Haroche, S., Lefèvre, V., Raimond, J. M. & Zagury, N. *Phys. Rev. Lett.* **65**, 976–979 (1990).
 5. Deutsch, D. & Ekert, A. *Physics World* **47–52** March (1998).

Neurobiology

Monkeys play the odds

M. James Nichols and William T. Newsome

On a sunny, spring morning, an experienced fisherman furrows his brow and weighs the merits of one fishing hole against another further upstream. At both sites he sees promising signs, and a novice could not distinguish between the two. But the veteran knows that, on mornings like this, bigger fish bite more often upstream — and he selects the upstream spot.

Like the fisherman we all make countless day-to-day decisions — selecting restaurants, betting in office football pools, buying used cars, and so on. We base these decisions not just on current information, but also on accumulated experience and expectations about the likelihood and size of rewards. Experimental evidence indicates

that animals also shape their behaviour based on the expected size and probability of rewards. Indeed, economists^{1–4}, psychologists^{5–8} and ecological biologists^{9,10} place these critical variables at the centre of models about human and animal decision-making.

Platt and Glimcher (reporting on page 233 of this issue)¹¹ have now applied decision theory to combined behavioural and neurophysiological experiments in rhesus monkeys, providing new insights into how decision variables are represented in the central nervous system. Despite the importance of reward expectations in shaping behaviour, neurophysiologists have largely ignored them in studies of brain function (but see Gallistel⁵). Instead, they have adopted a framework in which particular patterns of

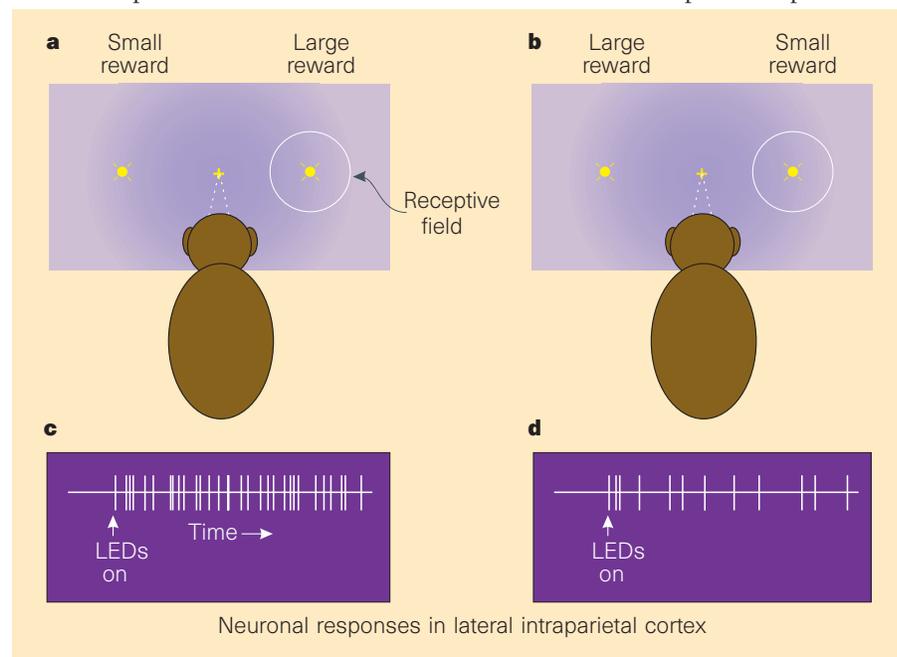


Figure 1 Assessing neural correlates of decision making. In a simple decision-making task devised by Platt and Glimcher¹¹, a monkey sits facing a viewing screen. On each trial, he first looks at a central fixation point (yellow cross). Then, two light-emitting diodes (LEDs; yellow circles) are illuminated at different locations on the viewing screen, one inside the receptive field (white circle) of the neuron under study, the other outside. After a delay, the fixation point is extinguished and the monkey receives a juice reward for looking to either of the two LEDs. a, In some blocks of trials, the LED within the receptive field is associated with a larger reward. b, In other blocks, the same LED is associated with a smaller reward. c, d, Schematized neuronal responses to the onset of the LED in the receptive field for the two reward conditions depicted in panels a and b. Vertical hatch marks indicate action potentials after illumination of the LEDs (vertical arrow). This hypothetical neuron discharges more vigorously for the same visual stimulus in its receptive field when the stimulus is associated with a larger juice reward.

sensory information are reflexively transformed into appropriate actions, with little regard for reward expectations or other variables. Although this simplified framework has been extremely powerful in early studies of the transformation from sensation to action, it cannot encompass the richness of the decision process.

In a key experiment (Fig. 1), Platt and Glimcher¹¹ placed a monkey in front of a viewing screen and repeatedly presented him with two identical visual stimuli — light-emitting diodes (LEDs). For each presentation of the two LEDs (called a ‘trial’), the monkey was rewarded with a squirt of juice if he directed his gaze from an initial target (a yellow cross called a ‘fixation point’) to either of the two LEDs. In other words, the animal, like the fisherman, was required to choose between two visually indistinguishable alternatives. Then, in blocks of roughly 100 trials, the authors manipulated the size of the juice reward associated with each target. In one block of trials, for example, the reward associated with one LED might be twice the size of the reward for the other one (Fig. 1a). In the next block of trials, these reward sizes might be reversed (Fig. 1b). Platt and Glimcher found that, over several trials at the

beginning of each block, the monkey learned the reward contingencies and adjusted his target selection accordingly. Specifically, he preferentially selected one target over the other in proportion to the relative reward sizes for the two targets.

While the monkey performed this simple decision-making task, the authors recorded neural activity in the lateral intraparietal cortex (LIP)¹². This area receives inputs from several parts of the brain involved in processing visual information, and it projects directly to the areas that control eye movements. Neurons in LIP discharge action potentials in response to visual stimuli that lie in restricted regions of visual space (their ‘receptive fields’). These neurons also discharge action potentials before and during gaze shifts to those visual stimuli. But neurons in LIP do not respond to stimuli that fall outside their receptive fields. Thus, LIP occupies an important intermediate stage in the transformation of visual signals into neural commands to move the eyes, and offers a point at which reward expectations could influence this transformation.

In each recording session, Platt and Glimcher placed one LED inside the receptive field (dashed circle in Fig. 1a and b) of the

neuron under study, and the other LED well outside it (although presumably within the receptive field of neurons in a different part of LIP). Consistent with previous results, when the two LEDs appeared, the one inside the receptive field elicited a response from the neuron. Furthermore, when the monkey shifted its gaze towards the LED in the receptive field, the neuron discharged before and during the eye movement, again consistent with previous findings. But the authors additionally found that when larger rewards were given for the LED in the receptive field than for the LED outside it (Fig. 1a), the neuron discharged more action potentials (vertical hatch marks in Fig. 1c) than when the same LED was associated with smaller rewards (Fig. 1d). So, given exactly the same visual display and exactly the same gaze shift, neurons in LIP respond differently depending on the animal’s expectation of reward size — an expectation that is evident in the pattern of decisions.

Many questions remain. Where in the brain are different rewards evaluated? Where are they associated with particular sensory stimuli? Where are reward expectations calculated and stored? Which areas of the brain are involved in more complex decisions? Platt and Glimcher have focused on LIP, but these processes are probably distributed across many regions of the brain. Furthermore, although monkeys estimate reward size efficiently in the simple task described here, humans estimate risk and reward quite poorly under some circumstances. A better understanding of the neural mechanisms underlying efficient estimation of reward may yield insights into why these mechanisms sometimes fail. In applying decision models to combined neurophysiological and behavioural experiments, Platt and Glimcher have begun to develop an important new framework for studying these issues in future experiments. □

M. James Nichols and William T. Newsome are at the Howard Hughes Medical Institute and Department of Neurobiology, Stanford School of Medicine, Stanford, California 94305-5125, USA. e-mails: jamie@monkeybiz.stanford.edu bill@monkeybiz.stanford.edu

Physiology

Crossing over to the dark side

Some locust species have a Jekyll-and-Hyde streak. Normally, they are a cryptic shade of green — shown in the photograph of *Schistocerca gregaria*, top — and lead solitary, unobtrusive lives. But at high population densities, their colour darkens and they become the gregarious, mobile animals that lay waste to crops. Now Amer Tawfik *et al.* have identified the hormone that triggers one aspect of this transmutation (*Proc. Natl Acad. Sci. USA* 96, 7083–7087; 1999).

Tawfik *et al.* used an albino mutant of the plague locust *Locusta migratoria* (middle), which did not darken when crowded. They found that implanting either the brain or the corpus cardiacum — an organ close to the brain — from normal locusts into the albino caused the production of melanin in the insect’s cuticle, and the normal darkening response. The hormone is probably made in the brain, then transferred to the corpus cardiacum for storage and release.

Next, Tawfik *et al.* used the albinos to identify the messenger itself. They injected albino locust nymphs with purified extracts from normal locusts: any that darkened had received the active ingredient. It turns out that this is a short peptide, 11 amino acids long, called [His⁷] corazonin. The bottom photograph shows an albino nymph of *L. migratoria* after



injection with corazonin (the effect on the green *S. gregaria* is similar). This hormone had already been found in another locust species, but its function was not known.

The hormone, christened ‘dark-colour-inducing neurohormone’, produces its effects even in locusts that are not crowded — the group is going on to investigate its effects on behaviour. Intriguingly, its amino-acid sequence is similar to that of the melanophore-stimulating hormone (MSH) found in mammals, and injecting MSH into albino locusts also causes their colour to darken. **John Whitfield**

1. von Neumann, J. & Morgenstern, O. *The Theory of Games and Economic Behavior* (Princeton Univ. Press, 1944).
2. Arnauld, A. & Nichole, P. *The Art of Thinking: Port-Royal Logic* (Bobbs-Merrill, Indianapolis, 1964).
3. Bernoulli, D. *The Works* (Birkhauser, Boston, 1982).
4. Kahneman, D., Slovic, P. & Tversky, A. *Judgement Under Uncertainty* (Cambridge Univ. Press, New York, 1982).
5. Gallistel, C. R. *Cognition* 50, 151–170 (1994).
6. Herrnstein, R. J. *J. Exp. Anal. Behav.* 4, 267–272 (1961).
7. DeVilliers, P. A. & Herrnstein, R. J. *Psych. Bull.* 83, 1131–1153 (1976).
8. Staddon, J. E. R. & Mothoral, S. *Psych. Rev.* 85, 436–444 (1978).
9. Pyke, G. H., Pulliam, H. R. & Charnov, E. L. *Q. Rev. Biol.* 52, 137–154 (1979).
10. Stephens, D. W. & Krebs, J. R. *Foraging Theory* (Princeton Univ. Press, 1986).
11. Platt, M. L. & Glimcher, P. W. *Nature* 400, 233–238 (1999).
12. Colby, C. L. & Goldberg, M. E. *Annu. Rev. Neurosci.* 22, 319–349 (1999).