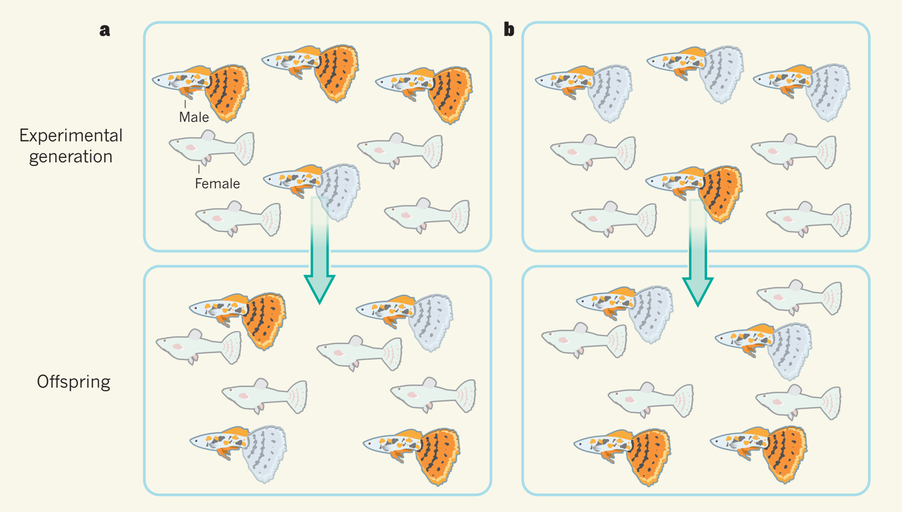

Figure 1 | The rare-male effect in a natural population. Hughes *et al.*[3] adjusted the frequency of male guppies in stream pools such that males with coloured tails were common and males with uncoloured tails were uncommon (**a**), or the reverse (**b**). The frequency of females was kept constant. In both types of pool, the authors found that the rare males obtained more matings with females and, as a result, contributed proportionally more offspring to the next generation.

in others. Then, to evaluate the reproductive success of the males, they undertook the daunting task of genotyping more than 1,400 offspring of pregnant females collected from the pools at the end of the experiment. Working across three stream systems and with an impressive level of replication, the researchers observed a strong pattern of elevated mating and reproductive success for whichever male type was rare.

Although this is an important study, it is not without limitations. One potential issue stems from the fact that, in guppies, fertilization is internal, and females can store sperm for weeks or months. When Hughes *et al.* collected females from their experimental pools and allowed the fish to give birth in the lab, they found strong evidence for a rare-male effect among the fathers of the first round of offspring. But when the females were obliged to use stored sperm for a second brood, the rare-male effect among the fathers of those offspring was not statistically significant. The authors argue that, in nature, female guppies mate frequently and rarely rely on stored sperm, so that all natural broods are essentially equivalent to the first broods in their experiment. However, the data reviewed by Hughes and colleagues in support of this claim are not extensive, and more work on this point would be helpful.

There was also a technical limitation in the genotyping technique used by the authors: relying on variable genomic regions known as microsatellites, they could confidently assign paternity to only a minority of individual offspring. Although there is no reason to expect this to create any bias in the data, it might be worthwhile revisiting the samples as more powerful methods for assigning paternity become available.

Hughes and colleagues' evidence for the rare-male effect in nature raises intriguing questions. The most obvious is, why do female guppies prefer males with uncommon colour patterns? Here we have more hypotheses than data sets. One theoretical analysis suggests that greater survival of rare male types, which has been found in guppies[2], might contribute to the evolution of such a preference, even if it is costly[9]. Interestingly, the same analysis points out that the negative feedback inherent in sexual selection for rarity (such selection causes rare types to become common) can ultimately prevent preferences for rare males from becoming ubiquitous. It remains to be determined whether preference for rarity varies among guppy females in experimental populations such as those used in this study. Hughes *et al.* point out that the rare-male effect might also emerge if females avoid mating with males that share the colour patterns of their previous mates, to increase the genetic diversity of their offspring. Alternatively, preferences for rarity could have evolved because they lead to lower rates of inbreeding if populations are small and intermittently isolated. It has even been suggested[10] that mating preferences for rare types provide no benefit for females and may be a manifestation of a preference for novelty that has evolved in other contexts.

The clever manipulations of natural populations used by Hughes *et al.* may, as the authors suggest, be a promising avenue for exploring the maintenance of variation by frequency-dependent selection in other systems. It might well be the case that, when put to the test, preferences for rarity might themselves not prove to be so rare. ■

Jeffrey S. McKinnon *is in the Department of Biology and Center for Biodiversity, East Carolina University, Greenville, North Carolina 27858, USA.* **Maria R. Servedio** *is in the Department of Biology, University of North Carolina Chapel Hill, Chapel Hill, North Carolina 27599, USA.*
e-mails: mckinnonj@ecu.edu; servedio@email.unc.edu

1. Kotiaho, J. S., LeBas, N. R., Puurtinen, M. & Tomkins, J. L. *Trends Ecol. Evol.* **23,** 1–3 (2008).
2. Olendorf, R. *et al. Nature* **441,** 633–636 (2006).
3. Hughes, K. A., Houde, A. E., Price, A. C. & Rodd, F. H. *Nature* **503,** 108–110 (2013).
4. Partridge, L. *Phil. Trans. R. Soc. Lond. B* **319,** 525–539 (1988).
5. Ajuria Ibarra, H. & Reader, T. *J. Zool.* **290,** 81–95 (2013).
6. Gray, S. M. & McKinnon, J. S. *Trends Ecol. Evol.* **22,** 71–79 (2007).
7. Mariette, M. M., Zajitschek, S. R. K., Garcia, C. M. & Brooks, R. C. *J. Evol. Biol.* **23,** 1772–1782 (2010).
8. Bleay, C., Comendant, T. & Sinervo, B. *Proc. R. Soc. B* **274,** 2019–2025 (2007).
9. Kokko, H., Jennions, M. D. & Houde, A. *Proc. R. Soc. B* **274,** 1317–1324 (2007).
10. Hughes, K. A., Du, L., Rodd, F. H. & Reznick, D. N. *Anim. Behav.* **58,** 907–916 (1999).

This article was published online on 30 October 2013.

NEUROSCIENCE

What to do and how

Decisions can differ depending on the context that surrounds them. Analyses of the prefrontal cortex region of the monkey brain indicate that a dynamical process at the neuronal population level controls this behaviour. **See Article p.78**

JEFFREY C. ERLICH & CARLOS D. BRODY

Neurons typically receive inputs from many thousands of others. This dense connectivity suggests that brain functions such as perception, cognition and motor control result from the concerted activity of populations of neurons, rather than from single-cell activity. Nevertheless, neuroscientists have long made discoveries by recording from, and thinking about, one neuron at a time. But more-recent developments in electrophysiological recording[1] and imaging[2] of the dynamics of hundreds of neurons simultaneously, alongside advances in data-analytic and theoretical tools[3,4], highlight the importance of understanding the relationship between the dynamics of large neural populations and the computations those dynamics embody[5–11]. In a tour de force reported on page 78 of this issue, Mante *et al.*[12] use electrophysiology together with sophisticated behavioural and neural modelling to add a population-dynamics view of the computations underlying flexible, context-dependent behaviour.

How you respond to a glance up a busy street should be very different depending on whether you intend to cross the road or hail a taxi. In the first case, you should pay attention to the motion of the cars, in the second case to their colour. Such stimulus-feature selection is an example of cognitive control[13] — our brains' ability to flexibly adapt how we process information from moment to moment, depending on current goals. Mante and colleagues wanted to know how the brain decides which aspect of a stimulus should guide a behaviour.

The authors designed an elegant task to explore this flexibility (Fig. 1a). They trained monkeys to look at a screen showing randomly moving dots, each dot randomly coloured either red or green. In trials in the colour context, the monkeys had to decide whether there were more red dots or more green dots, regardless of the dots' direction of motion. In the motion context, the animals decided whether most of the dots were moving left or right, regardless of the dots' colour. In both contexts, the same set of stimuli, containing motion and colour signals, was used.

The team then recorded from the monkeys' prefrontal cortex (PFC), a brain region thought to be crucial for flexible, context-dependent behaviour and which also controls the motor activity that animals use to report their decisions. Consistent with previous work[14,15], the activity of individual PFC neurons depended in complex, time-varying ways on multiple aspects of the task, including the motion and colour signals in the stimulus, the context and the monkey's decision. In what initially seemed a bewildering array of responses, these dependencies varied greatly across neurons.

To look at the data from a population viewpoint, the authors put together all the individual recordings for a given stimulus, and defined the neural population's response as the dynamical trajectory followed in a high-dimensional space in which each dimension represents the activity of one of the recorded neurons. They then used linear regression to find the axes (corresponding to a direction) in this space that best represented responses to the motion and the colour in the stimulus, and the monkeys' decision (Fig. 1b). They could therefore estimate separately how each signal's representation in the PFC developed over time, and how it depended on the context.

Which features of the data explained the context dependence of the behaviour? Not changes in the decision axis, which pointed in the same direction in both contexts. Nor could the behaviour be explained by irrelevant inputs being blocked from reaching the PFC, or by the inputs being changed in some context-dependent way, because the stimulus axes and the stimulus signals (motion and colour) on those axes remained remarkably constant across the two contexts. Yet, somehow, the decision signal was driven by the motion signal in the motion context, and by the colour signal in the colour context.

To investigate how the switch in response to the different contexts came about, Mante *et al.* used modelling. They took model neurons with nonlinear input–output functions, and densely connected them in a recurrent network. They provided the network with inputs representing the motion and colour signals and, separately, the current context signal. The activity of one of the model neurons was used to indicate the decision of the network.

The team then 'trained' the network by finding the strengths of the neuronal connections required to produce the appropriate, context-dependent decision in each trial. Starting the training from different initial random connection strengths led to different connection patterns that solved the problem. But cutting-edge analyses of how the model networks were solving the problem revealed that all the solutions shared two dynamical principles.

First, a line of closely spaced, stable, fixed points lay along the decision axis, with the two possible outcomes at opposite ends of the line. During each trial, stimulus-driven evidence for or against a decision accumulated along that line by gradually nudging the system from one fixed point to the next. Second, in dynamical systems, some inputs can cause a long-lasting perturbation in the system's trajectory, whereas other inputs can be short-lived, and rapidly decay away. In Mante and co-workers' model networks, the critical feature that depended on context was which of the inputs was short-lived. In the colour context, for example, signals from the irrelevant input (motion) were rapidly quenched and had no lasting impact on the neural trajectory, whereas colour signals persisted to drive the decision (Fig. 1c).

In other words, the context dependence was not achieved by controlling or gating the flow of information from one neuronal population to another. Instead, it was achieved within a single, densely interconnected cell population by controlling the directions along which the population dynamics quenched the neural activity. This dynamical principle can only be

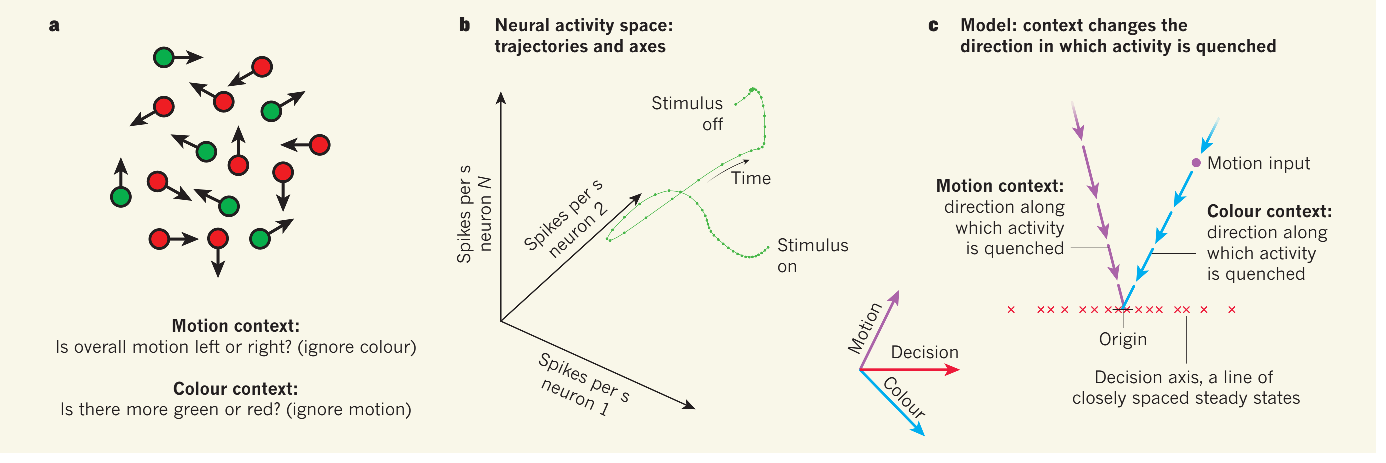

Figure 1 | The experiment and its results. a, Schematic view of the visual stimulus used by Mante *et al.*[12] in their study and the context-dependent decision their subjects had to make on each trial. **b**, In the neural activity space, each dimension corresponds to the activity of one recorded neuron. Together they give rise to the dynamical trajectory followed by the neuronal population in response to a particular stimulus (green). The thick coloured arrows illustrate the motion, colour and decision axes. **c**, Dynamical principle for context-dependent feature selection. In the colour context, one of the directions in neural space along which activity is quenched points in the same direction as the motion-input axis. Motion signals therefore decay rapidly in the colour context. In the motion context, the rapidly quenched direction points elsewhere, allowing the motion input to nudge the system from one fixed point to the next along the decision axis (red crosses), driving the decision. Neither the stimulus inputs nor the decision axes change across the two contexts; only the quenching direction changes.

understood — in fact, only exists — at the population level. And because the neural trajectories in the model closely resembled the neural trajectories in the data, Mante *et al.* propose this principle as the neural basis of flexible, context-dependent processing.

This fascinating study raises some intriguing questions. First, how do we go beyond correlations and causally test the authors' proposed principle? With existing perturbation methods, we cannot align manipulations with the critical axes that define the population's dynamical behaviour. Second, is the region of PFC that the authors recorded from even necessary for the context-dependent behaviour investigated here? The behavioural task used is new, so this question has never been addressed. Third, the conclusions on the behaviour of neural populations are based on data from many separate single-neuron recordings. Will the same principles hold when multi-neuron recordings are used to directly see the population neural trajectories in single trials? Finally, what is the specific mechanism that the network uses to change its quenching directions? These questions aside, Mante *et al.* have led the field with novel and deep ideas that should drive much discussion and thinking. ■

Jeffrey C. Erlich *and* **Carlos D. Brody** *are at the Princeton Neuroscience Institute and Department of Molecular Biology, Princeton University, Princeton, New Jersey 08544, USA.* **C.D.B.** *is also at the Howard Hughes Medical Institute.*
e-mail: brody@princeton.edu

1. Pfeiffer, B. E. & Foster, D. J. *Nature* **497,** 74–79 (2013).
2. Dombeck, D. A., Khabbaz, A. N., Collman, F., Adelman, T. L. & Tank, D. W. *Neuron* **56,** 43–57 (2007).
3. Yu, B. M. *et al. J. Neurophysiol.* **102,** 614–635 (2009).
4. Goldman, M. S. *Neuron* **61,** 621–634 (2009).
5. Laurent, G. *Nature Rev. Neurosci.* **3,** 884–895 (2002).
6. Briggman, K. L., Abarbanel, H. D. I. & Kristan, W. B. Jr *Science* **307,** 896–901 (2005).
7. Machens, C. K., Romo, R. & Brody, C. D. *J. Neurosci.* **30,** 350–360 (2010).
8. Harvey, C. D., Coen, P. & Tank, D. W. *Nature* **484,** 62–68 (2012).
9. Churchland, M. M. *et al. Nature* **487,** 51–56 (2012).
10. Shenoy, K. V., Sahani, M. & Churchland, M. M. *Annu. Rev. Neurosci.* **36,** 337–359 (2013).
11. Rigotti, M. *et al. Nature* **497,** 585–590 (2013).
12. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. *Nature* **503,** 78–84 (2013).
13. Miller, E. K. & Cohen, J. D. *Annu. Rev. Neurosci.* **24,** 167–202 (2001).
14. Roy, J. E., Riesenhuber, M., Poggio, T. & Miller, E. K. *J. Neurosci.* **30,** 8519–8528 (2010).
15. Clark, K. L., Noudoost, B. & Moore, T. *J. Neurosci.* **32,** 10907–10914 (2012).

CLIMATE SCIENCE

Uncertain then, irrelevant now

Uncertainty in estimates of the effects of aerosols on climate stems from poor knowledge of the past, pristine atmosphere — so getting a better understanding of these effects might not be as useful as was thought. SEE ARTICLE P.67

BJORN STEVENS

Aerosol particles in the atmosphere have long been feared to be the joker in the climate system's pack of cards. Even before the greenhouse effect became a household word, scientists had begun to consider the possibility that increasing concentrations of particles in the atmosphere's lowest layer (the troposphere) were acting to cool the planet, and thus masking what would otherwise be much larger temperature changes caused by rising concentrations of atmospheric carbon dioxide[1]. Reporting on page 67 of this issue, Carslaw *et al.*[2] use an innovative approach to demonstrate that most of the uncertainty in estimates of aerosol forcing — the aerosol perturbation caused by humans that induces a global surface-temperature change — that is associated with cloud-affecting aerosols can be attributed to uncertainty in the emission rate of aerosol precursors from natural sources. Their work suggests that, if there is an aerosol joker, it was probably played a century ago, and has become irrelevant to understanding present and future changes in global climate.

The aerosol, like the Earth system itself, is fascinatingly complex. How it interacts with clouds and radiation depends not only on its composition and concentration, but also on how these properties are distributed across particles of different sizes. All of these factors depend on uncertain microphysical processes and poorly understood aerosol sources, which in turn are integrated by the vagaries of the wind. Consequently, estimates of the strength of aerosol forcing derived from models that attempt to represent these poorly understood factors are controversial[3]. Carslaw *et al.* had the brilliant insight to recognize that something could be learned from this uncertainty.

Using a comprehensive model of global aerosols to train a simple statistical model, the researchers characterized the 28-dimensional