




Figure 1 | A proxy for ancient carbon dioxide levels. Coccolithophores are marine algae characterized by an exoskeleton of overlapping plates (coccoliths) composed of calcium carbonate. Bolton and Stoll¹ report that differences in the carbon-isotope composition of large and small coccoliths preserved in the geological record provide information about atmospheric carbon dioxide levels in ancient times.

of the small-celled alga *Emiliania huxleyi*, but Bolton and Stoll's results highlight the need to examine species of a range of sizes. The greater CO₂ sensitivity of large coccolithophores reported by the authors might affect competition among different-sized species, causing small species to outcompete large species at low CO₂ concentrations, and large coccolithophores to proliferate as CO₂ levels rise. The size-dependent concentration of inorganic carbon in coccolithophores could also influence the vertical flux of particulate organic and inorganic carbon in the sea, because cells and coccoliths from large species sink faster than those from smaller species.

Central to Bolton and Stoll's results is the conclusion that, as the concentration of CO₂ in sea water declines, a larger proportion of calcification is supported by CO₂ that enters the cell by diffusion. This CO₂ is converted to bicarbonate inside the cell to compensate for the diversion of pumped bicarbonate to photosynthesis. Coccolithophores might therefore produce less acid during calcification as the concentration of CO₂ rises, because a larger proportion of calcification is fed by bicarbonate (which produces 1 mole of acid per mole of calcium carbonate precipitated) than by CO₂ (which produces 2 moles of acid per mole of calcium carbonate precipitated). As a result, and because coccolithophores account for a large fraction of total calcification in the ocean, the currently expected decrease in the surface ocean's pH may be partially offset as CO₂ levels rise.

Perhaps of greater consequence to the global carbon cycle is how rising CO₂ concentrations will affect calcium carbonate precipitation by

coccolithophores overall. At low concentrations of 200–400 p.p.m., which occurred during past glacial periods and pertain today, calcification seems to decrease as atmospheric CO₂ levels rise⁹. Over a higher concentration range (400–750 p.p.m.), such as that expected during the next 100 years, the evidence relating to calcification trends is inconclusive^{8–10}.

NEUROSCIENCE

Dopamine ramps up

We thought we had figured out dopamine, a neuromodulator involved in everything from learning to addiction. But the finding that dopamine levels ramp up as rats navigate to a reward may overthrow current theories. [SEE LETTER P.575](#)

Yael Niv

Scientific findings typically come in two flavours: explanations for things we already knew occurred but had no idea why, or new phenomena that are clearly important but still mysterious. Howe and colleagues' finding¹, on page 575 of this issue, is of the latter kind — even if we don't yet know what it means, it stands to alter the way we think about dopamine.

Dopamine is a molecule that is broadcast throughout the brain and is involved in processes ranging from decision-making to schizophrenia, as well as most forms of addiction. The authors measured levels of dopamine in the

If, as Bolton and Stoll's findings suggest, the proportion of pumped bicarbonate used for calcification increases at higher concentrations of CO₂, this may partially counteract any suppression of calcification associated with future ocean acidification.

As the concentration of CO₂ in the atmosphere increases over the next 50 to 100 years, the 3-million-year transition from high to low CO₂ levels that Bolton and Stoll conclude occurred some 7 Myr ago will play back in reverse, but 30,000 to 60,000 times faster. Although the effects of this rapid carbonation of Earth's oceans on marine ecology and on the ocean's ability to absorb atmospheric CO₂ are uncertain, Bolton and Stoll provide insight into how a crucial component of marine phytoplankton communities worldwide may respond. ■

John Reinfelder is in the Department of Environmental Sciences, Rutgers University, New Brunswick, New Jersey 08901, USA. e-mail: reinfelder@envsci.rutgers.edu

1. Bolton, C. T. & Stoll, H. M. *Nature* **500**, 558–562 (2013).
2. Beerling, D. J. & Royer, D. L. *Nature Geosci.* **4**, 418–420 (2011).
3. Pagani, M., Freeman, K. H. & Arthur, M. A. *Science* **285**, 876–879 (1999).
4. Pagani, M., Liu, Z., LaRiviere, J. & Ravelo, A. C. *Nature Geosci.* **3**, 27–30 (2010).
5. LaRiviere, J. P. *et al. Nature* **486**, 97–100 (2012).
6. Bidigare, R. R. *et al. Glob. Biogeochem. Cycles* **11**, 279–292 (1997).
7. <http://keelingcurve.ucsd.edu>
8. Bach, L. T. *et al. New Phytol.* **199**, 121–134 (2013).
9. Beaufort, L. *et al. Nature* **476**, 80–83 (2011).
10. Iglesias-Rodriguez, M. D. *et al. Science* **320**, 336–340 (2008).

striatum of rats while the animals ran through mazes for food rewards. The striatum (Fig. 1a) is the area that contains the highest dopamine concentration in the brain. It is involved in action selection at all levels, from choosing which limb to move to selecting a goal to work towards. In a series of elegant experiments, Howe *et al.* established that dopamine concentration gradually ramps up as rats run towards a reward, and that the slope of the ramps relates to the amount of anticipated reward and the effort required to obtain it.

Why are these dopamine ramps so relevant? Dopamine-secreting (dopaminergic) neurons are special because they are thought to fire in unison, broadcasting a single all-important




Figure 1 | Rewards and the brain. **a**, Dopaminergic neurons in the midbrain project to all brain areas, most prominently to the striatum (black arrows). These cells fire at a constant rate of 3–5 spikes per second, with occasional phasic bursts or pauses on the occurrence of positive reward prediction errors (discovering that the local supermarket now supplies your favourite coffee beans, **b**) or negative reward prediction errors (sipping your coffee and finding that the milk has gone sour, **c**), respectively. The background (tonic) level of dopamine fluctuates slowly, possibly tracking the average rate of rewards (not shown). **d**, By measuring dopamine concentrations in the striatum of rats navigating mazes, Howe *et al.*¹ reveal a third mode of dopaminergic signalling: when a prolonged series of actions must be completed to obtain a reward (for instance, all the steps it takes to make a cup of coffee), dopamine concentration ramps up gradually, at each point in time signalling the predicted distance from the goal.

message widely. As such, and because dopamine is implicated in a bewildering number of disorders, neuroscientists have long been keen to understand the dopamine signal in the intact brain and how it can be restored when things go wrong. These efforts have led to a well-established theory, which may unfortunately be at odds with these dopamine ramps.

The central theory of dopaminergic firing comes from theoretical neuroscientists who noticed that the firing patterns of dopaminergic neurons bear an eerie resemblance to a key component in computational algorithms of reinforcement learning² called a reward prediction error³. Reward prediction errors quantify ‘surprise’ — the difference between the rewards we expect and those we get in reality⁴.

Imagine you drink coffee routinely. One day, while shopping, you find coffee beans you particularly like, which were not previously available in your town. According to the theory, and verified experimentally in laboratory settings^{5–7}, your dopaminergic neurons will fire to signal a positive reward prediction error due to the increase in your future expected (coffee) rewards (Fig. 1b). Such bursts of dopamine release affect learning in the striatum⁸, strengthening actions that preceded a positive prediction error — you will now be more likely to return to this shop. The flip side, a negative prediction error, occurs when a reward is below expectation — for instance, if you sip your coffee and find

that the milk has gone sour (Fig. 1c).

To the brain, new information that causes you to change your predictions and an actual reward that is at odds with your predictions are equivalent. In both cases, bursts or pauses in dopaminergic firing will notify the brain of the prediction error. Consequently, through dopamine-dependent learning, future predictions will become more accurate, and actions that led to better-than-expected outcomes will become more common.

The prediction-error theory is compelling because it is normative — it explains the role of dopamine within a prescriptive framework of how one should adapt behaviour to earn more rewards. However, it also has shortcomings. For instance, it fails to address dopamine’s effects on action vigour: Parkinson’s disease, caused by the death of dopaminergic neurons, results in slowing down of all actions, and the nickname for amphetamine (which mimics elevated dopamine levels) is ‘speed’.

Luckily, a straightforward extension of the theory fills this gap, suggesting that the background concentration of dopamine (termed tonic dopamine, to differentiate it from the phasic bursts and pauses that signal prediction errors) indicates the overall rate of rewards. Thus, tonic levels of dopamine quantify the cost of time, and should affect how much time we spend on each action⁹. According to this theory, phasic and tonic dopaminergic signalling

convey distinct but related information, affecting learning and vigour, respectively.

So where do the ramps Howe *et al.* observe fit in? They don’t — at least, not in any straightforward way. The ramps do not correspond to anything surprising and so cannot be signalling prediction errors. Because the ramps are predictable, the computational theory of reinforcement learning² dictates that they should disappear through learning. Instead, they persist even when rats have learned to predict the reward. If anything, they then become more informative, with the slope of the ramp differentiating whether the rat expects a small or large reward. Moreover, the ramps do not fit in with ideas about tonic levels of dopamine, because the ramps are too short-lived to be considered tonic and because the rate of reward does not increase as the rat traverses the maze.

What, then, do the ramps tell us about dopaminergic signalling in the brain? Mostly, that we have to go back to the drawing board. Obviously, before radically rethinking the current theory, it should be determined how general these findings are. Others have measured dopamine in tasks requiring a series of actions to obtain rewards, but have not seen reliable evidence of ramps¹⁰. Is this because ramping signals are confined to spatial navigation? Are there no dopamine ramps in our brains when we go about making coffee (Fig. 1d), unless this involves a walk to the nearest café?

Another question pertains to the origin of the ramps. Howe *et al.* used advanced, albeit indirect, methods to measure dopamine concentration in the striatum. By contrast, many other studies track the firing of dopaminergic neurons by recording electrical activity in the midbrain, where the cell bodies lie (Fig. 1a). Such recordings from rats running through mazes have yet to be reported. If similar ramps in neural firing rate were found, this would strengthen the current results. Yet, not finding them would also be revealing, because the relationship between firing of dopaminergic neurons and dopamine release in target areas is complex¹¹. For instance, the ramping activity of nearby striatal neurons (which has been recorded in some^{12,13}, but not all, areas of the striatum) may instigate dopamine release even without firing of the dopaminergic cells¹⁴. Whether this mode of dopamine release exists is not known. Regardless of this, once the boundaries of Howe and colleagues’ finding are established, it will once again be up to the theorists to explain what the ramps actually mean. ■

Yael Niv is at the Princeton Neuroscience Institute and the Psychology Department, Princeton University, Princeton, New Jersey 08544, USA.

e-mail: yael@princeton.edu

1. Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips,

- P. E. M. & Graybiel, A. M. *Nature* **500**, 575–579 (2013).
2. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press, 1998).
 3. Montague, P. R., Dayan, P. & Sejnowski, T. J. *J. Neurosci.* **16**, 1936–1947 (1996).
 4. Niv, Y. & Schoenbaum, G. *Trends Cogn. Sci.* **12**, 265–272 (2008).
 5. Tobler, P. N., Fiorillo, C. D. & Schultz, W. *Science* **307**, 1642–1645 (2005).
 6. Tobler, P. N., Dickinson, A. & Schultz, W. *J. Neurosci.* **23**, 10402–10410 (2003).
 7. Morris, G., Nevet, A., Arkadir, D., Vaadia, E. & Bergman, H. *Nature Neurosci.* **9**, 1057–1063 (2006).
 8. Reynolds, J. N. J., Hyland, B. I. & Wickens, J. R. *Nature* **413**, 67–70 (2001).
 9. Niv, Y., Daw, N. D., Joel, D. & Dayan, P. *Psychopharmacology* **191**, 507–520 (2007).
 10. Wanat, M. J., Kuhnen, C. M. & Phillips, P. E. M. *J. Neurosci.* **30**, 12020–12027 (2010).
 11. Montague, P. R. et al. *J. Neurosci.* **24**, 1754–1759 (2004).
 12. van der Meer, M. A. A. & Redish, A. D. *J. Neurosci.* **31**, 2843–2854 (2011).
 13. van der Meer, M. A. A., Johnson, A., Schmitzer-Torbert, N. C. & Redish, A. D. *Neuron* **67**, 25–32 (2010).
 14. Threlfell, S. et al. *Neuron* **75**, 58–64 (2012).

ASTROPHYSICS

Radioactive glow as a smoking gun

The observation of infrared emission following a short γ -ray burst lends support to the hypothesis that mergers of compact binary systems cause such bursts and produce the heaviest nuclei in the cosmos. [SEE LETTER P.547](#)

STEPHAN ROSSWOG

On 3 June 2013, NASA's Swift satellite detected an intense flash of γ -rays known as a short γ -ray burst. Follow-up observations by the Hubble Space Telescope revealed infrared emission that was present about 9 days after the burst, but had faded away after approximately 30 days. In this issue, Tanvir and colleagues¹ (page 547) propose that this emission was powered by the radioactive decay of heavy elements that were freshly synthesized in the merger of a compact binary system, consisting of either two neutron stars or a neutron star and a black hole. If this interpretation is correct, the observation could have profound consequences for high-energy astrophysics, cosmic nucleosynthesis and the detection of gravitational waves.

There are two flavours of γ -ray burst (GRB), each of different duration. Long bursts (lasting more than about 2 seconds) are produced by the death of a rare breed of massive star, whereas short bursts (less than 2 s) are thought

to result from compact-binary mergers. So far, we know of ten systems containing two neutron stars — extremely densely packed objects that are more massive than the Sun, but only about 12 kilometres in radius, and that consist predominantly of neutrons. As the stars orbit around each other, they emit gravitational waves and slowly spiral inwards until they merge (Fig. 1). Such orbital decays have been observed² and agree remarkably well with the predictions from Einstein's theory of general relativity.

The last moments of the spiral and the final merger are the prime targets for ground-based gravitational-wave laboratories such as LIGO³, VIRGO⁴ or KAGRA⁵. The gravitational-wave signal not only contains information about the physical parameters of the merging system (for example, the masses and spins of the neutron stars), but also carries the imprint of relativistic gravity and of the properties of matter at densities that exceed those inside an atomic nucleus. However, such a signal leaves us essentially blind as to the astronomical environment (for

example, the type of host galaxy) in which the merger occurs. Therefore, accompanying electromagnetic emission that can provide such extra information is highly welcome.

Compact-binary mergers are also interesting from a nucleosynthesis point of view: they eject roughly 1% of a solar mass of neutron-rich matter^{6–8}, and calculations show that this material is entirely made of heavy 'rapid neutron capture' elements^{9–11}. Such elements form through the rapid (as compared with radioactive decays) capture of neutrons on existing atomic nuclei, a mechanism known as the r-process. Generally, adding new neutrons to existing nuclei will not produce stable configurations, thus leading to a competition between further neutron captures and radioactive decays. If a lot of neutrons are present, as in the case of disrupted neutron-star material, very heavy and neutron-rich nuclei can be produced that are far from the 'valley of stability' — the location at which all the stable nuclei reside in the nuclear chart. In the most extreme cases, the neutron capture stops only near the 'neutron dripline', the point at which neutrons are no longer bound and can just 'drip out' of nuclei. Once all available neutrons have been consumed, the unstable nuclei decay towards the valley of stability with a wide variety of radioactive half-lives.

If this rapid neutron capture occurs inside the ejected material and sufficient radioactivity is left when the ejected matter becomes transparent, this should cause a radioactively powered, electromagnetic transient^{12–14}. This phenomenon is often called a macronova or kilonova in the astronomical literature. The delay time between the merger and the peak of the transient emission and its wavelength range depend sensitively on the opacity of the ejected material. Unfortunately, the opacity of such heavy r-process matter is not well known, but recent calculations¹⁵ indicate that it is orders of magnitude larger than that of iron-group-like nuclei. Therefore, macronovae should peak a few days after a merger at infrared wavelengths. The transient reported by Tanvir et al. may be the first event of this kind ever to be observed.

If this interpretation is correct, the observation represents a major advance linking



Figure 1 | Supercomputer simulation of the merger of two neutron stars¹⁹. The two stars merge within a few milliseconds and eject around 1% of a solar mass of matter into space. The radioactive decay of this material is likely to have been observed as an infrared transient by Tanvir and colleagues¹. The colour scale denotes the logarithm of temperature in megaelectronvolts. A temperature unit of 1 MeV corresponds to 1.16×10^{10} kelvins.