

Highly energetic wave conditions often destroy the complex three-dimensional surf zone structures (3), resulting in one or more uniform linear bars along the shore (see the right panel of the figure). Although it has not been shown by any of the modeling efforts described above, this is generally considered to be a forced response. What has been shown to be a forced response is the offshore migration of the linear bars when energetic conditions continue to prevail (4, 5). The sediment transport formulation in all prevailing models is based on near-bed velocities derived from experiment or theory (14, 15). Applying a similar approach during mild wave conditions typically results in incorrect onshore motion of the bar, or even the failure of onshore motion (16).

Hoefel and Elgar (1) now show that this onshore motion can also be predicted with a deterministically forced sediment transport model, provided that the sediment transport induced by flow acceleration within the waves is included. They extend a formulation by Drake and Calantoni (17) based on detailed numerical modeling of particle-fluid interactions in a sediment layer to the case of random waves, appropriate for field conditions. Hoefel and Elgar (1) demonstrate improved prediction of near-shore bar motion over a 45-day period. This implies a substantial extension of prediction horizons of deterministic forced models.

The introduction of wave acceleration or, more generally, temporal and spatial pressure gradients (17–19) is an important paradigm shift in describing and modeling sediment transport. Prevailing concepts are based on shear stress or work exerted by fluid velocities. Introducing pressure gradients will give a new boost to understanding sediment transport near the shore.

Breaking waves. ARGUS stations (Coastal Imaging Lab, Oregon State University) overlooking part of a beach are used to obtain photographic images of incident wave breaking. Waves prefer to break over shallow bars, where the foam of the breaking waves shows up as an area of high light intensity. By averaging over a large number of images (equivalent to a photographic time exposure), a stable estimate of light intensity is obtained (3). The sharp contrast in light intensity between areas of breaking and nonbreaking waves reflects the position of shallow bar areas and deeper channels and troughs. (Left) A time exposure of Palm Beach, Australia, with superposed bottom contours (sea-floor depth in meters with respect to mean sea level) displays a complex pattern of shallow shoals (light areas) cut by deep rip channels (dark areas). The beach is located on the left. (Right) A time exposure at the same beach shows intense wave breaking on a linear bar and additional breaking at the shoreline.

Do these findings imply that surf zone bar structures are deterministically forced and thus deterministically predictable? The answer has to be negative. Many studies of seabed and land geomorphology (9, 20) show that self-organization processes lead to emergent properties, which cannot be predicted from the physics of fundamental particles.

What is shown, though, is that reductionist studies can still yield new insights and that the fashionable self-organizational research approach is not the only route to increased understanding.

What hampers reductionist research progress is the enormous effort required to unravel physical processes from first principles. Hoefel and Elgar show that it pays off. Their study is a first step toward an operational new sediment transport approach. Developing this model further will require new physical insights and capabilities for modeling such phenomena as nonlinear wave kinematics in the surf zone.

References

1. F. Hoefel, S. Elgar, *Science* **299**, 1885 (2003).
2. L. D. Wright, A. D. Short, *Mar. Geol.* **56**, 93 (1984).
3. T. C. Lippmann, R. A. Holman, *J. Geophys. Res.* **97**, 11575 (1990).
4. J. A. Roelvink, M. J. F. Stive, *J. Geophys. Res.* **94**, 4785 (1989).
5. E. B. Thornton, R. Humiston, W. Birkemeier, *J. Geophys. Res.* **101**, 12097 (1996).
6. A. Ashton, A. B. Murray, O. Arnault, *Nature* **414**, 296 (2001).
7. R. Deigaard, N. Drønen, J. Fredsøe, J. Jensen, M. Jørgensen, *Coastal Eng.* **36**, 171 (1999).
8. A. Falques, G. Coco, D. A. Huntley, *J. Geophys. Res.* **105**, 24701 (2000).
9. B. T. Werner, T. Fink, *Science* **260**, 968 (1993).
10. G. Coco, D. A. Huntley, T. J. O'Hare, *J. Geophys. Res.* **105**, 21991 (2000).
11. R. T. Guza, D. L. Inmann, *J. Geophys. Res.* **80**, 2997 (1975).
12. H. N. Southgate, I. Moller, *J. Geophys. Res.* **105**, 11489 (2000).
13. S. Elgar, *J. Geophys. Res.* **106**, 4625 (2001).
14. A. J. Bowen, in *The Coastline of Canada*, S. McCann, Ed., Paper 10-80 (Geol. Surv. Canada, Ottawa, 1980), pp. 1–11.
15. J. A. Bailard, *J. Geophys. Res.* **86**, 10938 (1981).
16. E. Gallagher, S. Elgar, R. T. Guza, *J. Geophys. Res.* **103**, 3203 (1998).
17. T. G. Drake, J. Calantoni, *J. Geophys. Res.* **106**, 19859 (2001).
18. S. Elgar, R. T. Guza, M. Freilich, *J. Geophys. Res.* **93**, 9261 (1988).
19. S. Elgar, E. Gallagher, R. T. Guza, *J. Geophys. Res.* **106**, 11623 (2001).
20. M. A. Kessler, B. T. Werner, *Science* **299**, 380 (2003).

NEUROSCIENCE

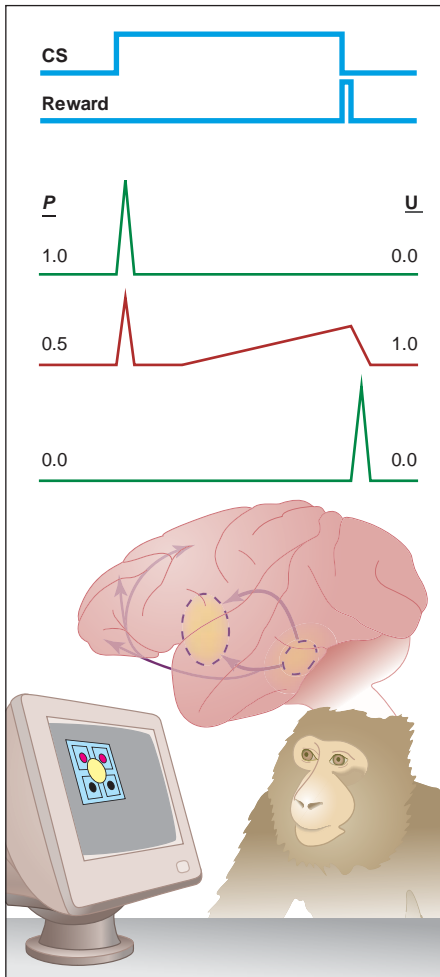
Gambling on Dopamine

Peter Shizgal and Andreas Arvanitogiannis

With bated breath, a player at a roulette table stares intently at the spinning wheel. As the ball comes to rest in one of the numbered slots, a smile crosses the gambler's face. This success strengthens his misguided belief in his ability to overcome the house advantage, and he prepares to wager again. The gambler's ability to detect the slot where the ball has settled depends on point-to-point connections between nerve cells at multiple levels of the visual system. The accompanying changes in emotion, attention, learning, and action depend on neurons with a very different pat-

tern of connectivity. Such neurons include midbrain dopamine neurons, which have cell bodies in the substantia nigra and ventral tegmental area of the midbrain, and highly divergent projections that connect with the frontal cortex, dorsal and ventral striatum, and other forebrain regions. Midbrain dopamine neurons go awry in Parkinson's disease, schizophrenia, and drug addiction. Data from both human and animal research implicate this small but widely connected neuronal population in motor control, motivation, effort, reward, analgesia, stress, learning, attention, and cognition. On page 1898 of this issue, Fiorillo *et al.* (1) report a new response mode for midbrain dopamine neurons and speculate how this new mode might contribute to the allure of gambling.

The authors are at the Center for Studies in Behavioral Neurobiology, Concordia University, Montréal, Quebec, H3G 1M8, Canada. E-mail: shizgal@csbn.concordia.ca; andreas@csbn.concordia.ca



Anticipating a reward. Electrophysiological responses of single midbrain dopamine neurons (yellow) were recorded while monkeys viewed a computer monitor. Unique visual stimuli were associated with different probabilities of a reward (a drop of syrup). Rewards were presented with different probabilities (P) when the conditioned stimulus was switched off (CS offset). The uncertainty (U) of the reward varied as an inverted U-shaped function of probability. When $P = 0$ or $P = 1$, the monkey is certain that reward delivery will or will not accompany CS offset. In contrast, when $P = 0.5$, the onset of the CS provides no information about whether reward will or will not occur, but it does predict the potential time of reward delivery. The interval between successive CSs varied unpredictably (not shown), and thus the onset of the CS is the earliest reliable predictor of the occurrence and/or the potential time of reward delivery. Once the meaning of the stimuli has been learned, the population of dopamine neurons responds to CS onset with a brief increase in activity when $P = 1.0$, but receipt of the expected reward does not provoke a strong change in firing. When $P = 0$, CS onset produced little response; if, however, the investigator violated expectation by delivering a reward, the dopamine cells responded with a brief, vigorous increase in firing. When $P = 0.5$ (and uncertainty about reward occurrence is maximal), a slow, steady increase in firing is seen prior to the time of potential reward delivery.

In a resting animal, midbrain dopamine neurons show a slow, steady (“tonic”) rate of firing; certain meaningful stimuli provoke brief, abrupt (“phasic”) changes in discharge. In previous work, Schultz and co-workers (2, 3) recorded the activity of single dopamine neurons in awake monkeys with microelectrodes. When the monkey received an unexpected reward, such as a drop of juice, most dopamine cells responded with a burst of firing. However, if the monkey learned that a stimulus, such as a particular pattern on a computer monitor, always preceded delivery of the reward, the dopamine neurons no longer responded to the reward but fired instead in response to the predictive (“conditioned”) stimulus. Omission of a predicted reward caused slowing or cessation of dopamine firing around the time that the reward was expected.

The phasic responses of midbrain dopamine neurons resemble a key signal in computer models based on animal learning (4–7). Through adjustment of connection weights in a neural network (“reinforcement”), these models are able to predict the achievement of a goal state (“reward”) and result in optimization of actions. The incremental improvement in predictions is driven by a “prediction error”: the difference between expected and experienced rewards. The new findings of Fiorillo *et al.* both strengthen and challenge the reinforcement-learning notion of midbrain dopamine neuronal activity while raising many fascinating new questions.

The novelty of the Fiorillo *et al.* study lies in their systematic variation of the proportion of conditioned stimuli that were followed by a reward. The conditioned stimulus associated with each reward probability (0.0, 0.25, 0.50, 0.75, 1.0) consisted of a unique visual pattern displayed on a computer monitor for 2 seconds (see the figure). Delivery of the reward, a drop of syrup, coincided with the offset of the conditioned stimulus. The dopamine neurons responded to the conditioned stimuli with phasic increases in firing that correlated positively with reward probability. In contrast, responses to reward delivery showed a strong negative correlation with reward probability. On trials when an expected reward was omitted, the firing of dopamine neurons tended to decrease at the time of potential reward delivery, and the magnitude of this dip tended to increase with the probability of a reward. The systematic variation in the strength of the phasic responses to conditioned stimuli and to reward delivery or omission supports and extends previous findings obtained at the two extreme probabilities (0.0 and 1.0): The higher the likelihood of reward, the

stronger the firing to the conditioned stimulus, the larger the decrease in firing to reward omission, and the weaker the firing to reward delivery.

Their most provocative results concern the activity of the dopamine cells before the time of potential reward delivery. In previous work carried out at the two extreme probabilities, the firing rate was stable. By exploring intermediate probabilities, Fiorillo *et al.* reveal a striking new pattern: The activity of dopamine neurons increased before the potential delivery of an uncertain reward. In contrast to the brief upswings in firing triggered by reward-predicting stimuli and unexpected rewards, the population firing rate rose steadily throughout most of the 2-second presentation of the conditioned stimulus when the probability of reward was 0.5, attaining a higher rate than when the reward probability was 0.25 or 0.75.

The authors propose that the sustained firing preceding the time of potential reward delivery tracked the uncertainty of the reward. The onset of the conditioned stimuli signaling probabilities of either 0.0 or 1.0 provided the monkey with definitive information as to whether the reward would be delivered; if the meaning of the stimuli had been fully learned, then uncertainty about reward occurrence following stimulus onset would be zero. In contrast, a reward was equally likely to be delivered or omitted when the probability was 0.5, and the monkey should have been maximally uncertain about reward occurrence. When the reward probability was 0.25 or 0.75, uncertainty was intermediate; if the monkey had bet on the occurrence of the reward following stimulus onset, it could have won, on average, three trials out of every four.

Regarding reinforcement learning models, the phasic response of the dopamine neurons encodes a prediction error that is used as a “teaching signal” to improve future predictions. The incremental adjustment of the weights causes the teaching signal to move backward in time toward the onset of the earliest stimulus that reliably predicts the occurrence of a reward (and/or its potential time of delivery). How the sustained response preceding potential reward delivery could be incorporated into such models is hard to see. If the dopamine signal serves as the teacher, and the sustained component is not filtered out, how could the sustained component remain stationary in time and amplitude over many trials? This problem is a knotty one because the sustained and phasic signals do not appear to be carried by independent populations of neurons. Thus, postsynaptic elements are likely to register the combined impact of

both the phasic and sustained signals. Given the slow time course of dopaminergic transmission, the authors are not confident that the two signals could be demultiplexed by postsynaptic elements.

Fiorillo *et al.* show that the sustained activity grows with the magnitude of the potential reward (the volume of syrup delivered). They propose that the dopamine responses multiplicatively combine reward magnitude with probability, thus encoding the expected reward value. This may be so in the case of the phasic response, which varies monotonically with reward probability (although data in addition to those shown are required to prove multiplicative combination). However, the sustained response increases as probability rises from 0.0 to 0.5 and then decreases toward zero as probability grows further to 1.0. This bitonic mapping of probability seems incompatible with the calculation of expected value; for example, this mapping yields a value at or near zero for a certain payoff. Given that many dopamine neurons show both phasic and sustained responses, the problem of extracting an uncorrupted expected-value signal is analogous (or identical) to the as-yet unsolved problem of extracting an uncorrupted prediction error.

When reward uncertainty was maximal, close to 30% of the dopamine neurons showed increased firing before potential reward delivery, and the firing rate of the population as a whole rose steadily during most of the 2-second presentation of the conditioned stimulus. Increasing the payoff (the volume of the syrup drops) strengthened sustained firing. Thus, a substantial release of dopamine might occur during prolonged periods of uncertainty about large payoffs. The authors propose that such increases in dopamine output may contribute to the rewarding properties of gambling. Given the importance of dopamine in addiction and that experimentally induced increases in dopamine tone can produce rewarding effects (8), this is not a suggestion to be taken lightly. Indeed, repeated intermittent increases in dopamine release can lead to sensitization, which increases appetitive motivation and the persistence of behaviors that lead to additional dopamine release (9). Results of neuroimaging studies in humans suggest elevated dopamine release during game-playing and gambling-like tasks (10–12).

The proposal that dopamine release driven by uncertainty about large payoffs could contribute to gambling comes at a moment when decision theorists are striving to integrate emotional and cognitive influences on choice (13, 14). The

roulette player faces a house advantage of 5.26% and may well be aware of these odds. Nonetheless, he harbors very common misperceptions about the statistical independence of successive outputs produced by randomizing devices such as roulette wheels (15); such misperceptions may coexist with awareness of the house edge and may contribute substantially to the gambler's misplaced confidence that he can beat the odds. The effectiveness of cognitive therapy aimed at correcting such misperceptions has been demonstrated after 6 to 12 months (16). However, emotions powerfully influence decision-making in real time (13, 14) and may provide a necessary complement to cognitive errors. The Fiorillo *et al.* proposal links such emotion-driven processes to the uncertainty-related activity of dopamine neurons.

These investigators ponder what a neural signal, which proves maladaptive in the artificial environment of a casino, does in the natural environment. They propose that such a signal may drive risk seeking, helping an animal beset with uncertainty to find better predictors of consequential events. Conditions that promote either risk aversion or risk seeking have been studied extensively. Generally, these studies detail influences unrelated to the reward value of uncertainty (17). Nonetheless, uncertainty could contribute, and the proposal by Fiorillo and colleagues merits investigation.

The functional argument advanced by the authors links the sustained firing of midbrain dopamine neurons to increased allocation of attention, thus addressing the question of why this signal firing precedes the time of potential reward delivery. Adapting some ideas from animal learning (18), they propose that increased attention during uncertainty may promote the learning of better predictors and actions. This aspect of their hypothesis is also likely to stimulate much empirical and modeling work. It will be fascinating to discern how phasic and sustained dopamine signals contribute to selective (19) and nonselective (20) aspects of attention.

The authors appeal to the inherent variability of interval timing to account for the time course of the sustained response. If interval timing were error-free, then uncertainty about the imminence of reward would be zero until just before the offset of the conditioned stimulus, even if the probability of reward occurrence were 0.5. However, the interval-timing mechanism is known to be noisy (21). Thus, during presentation of the conditioned stimulus, the monkey's confidence that reward is imminent should rise more gradually (22). Could this account for the slow in-

crease in population firing when probability equals 0.5? Clearly, uncertainty about the imminence of reward delivery does not suffice to produce sustained firing, as no such firing precedes the offset of the conditioned stimulus when reward probability equals 1.0. Why the contribution of uncertainty about time should depend on the uncertainty about occurrence is not clear. It is also perplexing that less sustained activity was seen when the conditioned stimulus was turned off midway between the time of its onset and the time of potential reward delivery. Perhaps this reflects the fact that dopamine modulates the impact of motivationally charged stimuli (23).

Fiorillo *et al.* have engaged in some theoretical risk taking that may reap scientific rewards. They propose that the activation of dopamine neurons by uncertainty mobilizes attention, motivates risk seeking, and promotes learning about relationships between external stimuli and consequential events. Their striking results and provocative ideas will increase uncertainty about some seemingly well-established ideas, attract the attention of researchers in many disciplines, motivate the search for new data and relationships, and promote learning about a neural system of great importance to understanding both normal and pathological behavior.

References and Notes

1. C. D. Fiorillo, P. N. Tobler, W. Schultz, *Science* **299**, 1898 (2003).
2. W. Schultz, *J. Neurophysiol.* **80**, 1 (1998).
3. _____, *Neuron* **36**, 241 (2002).
4. R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 1998).
5. P. R. Montague *et al.*, *Nature* **377**, 725 (1995).
6. P. R. Montague *et al.*, *J. Neurosci.* **16**, 1936 (1996).
7. W. Schultz *et al.*, *Science* **275**, 1593 (1997).
8. T. M. Tzschentke, *Prog. Neurobiol.* **56**, 613 (1998).
9. P. Vezina *et al.*, *J. Neurosci.* **22**, 4654 (2002).
10. M. J. Koepf *et al.*, *Nature* **393**, 266 (1998).
11. H. C. Breiter *et al.*, *Neuron* **30**, 619 (2001).
12. B. Knutson *et al.*, *J. Neurosci.* **21**, RC159 (2001).
13. G. F. Loewenstein *et al.*, *Psychol. Bull.* **127**, 267 (2001).
14. P. Slovic, M. Finucane, E. Peters, D. G. MacGregor, in *Intuitive Judgment: Heuristics and Biases*, T. Gilovich, D. Griffin, D. Kahneman, Eds. (Cambridge Univ. Press, Cambridge, UK), in press.
15. R. Ladouceur *et al.*, *J. Soc. Behav. Pers.* **10**, 473 (1995).
16. R. Ladouceur *et al.*, *J. Nerv. Ment. Dis.* **189**, 774 (2001).
17. M. Bateson, A. Kacelnik, in *Cognitive Ecology: The Evolutionary Ecology of Information Processing and Decision Making*, R. Dukas, Ed. (Univ. of Chicago Press, Chicago, IL, 1998), pp. 297–341.
18. J. M. Pearce, G. Hall, *Psychol. Rev.* **87**, 532 (1980).
19. P. Dayan *et al.*, *Nature Neurosci.* **3** (suppl.), 1218 (2000).
20. P. Redgrave *et al.*, *Neuroscience* **89**, 1009 (1999).
21. C. R. Gallistel, J. Gibbon, *Psychol. Rev.* **107**, 289 (2000).
22. C. R. Gallistel, in *Frequency Processing and Cognition*, P. Sedlmeier, T. Betsch, Eds. (Oxford Univ. Press, Oxford), in press.
23. K. C. Berridge, T. E. Robinson, *Brain Res. Brain Res. Rev.* **28**, 309 (1998).
24. We thank P. Dayan, C. R. Gallistel, R. Hastie, D. Kahneman, and J. McClelland.