

Discrete Coding of Reward Probability and Uncertainty by Dopamine Neurons

Christopher D. Fiorillo,* Philippe N. Tobler, Wolfram Schultz

Uncertainty is critical in the measure of information and in assessing the accuracy of predictions. It is determined by probability P , being maximal at $P = 0.5$ and decreasing at higher and lower probabilities. Using distinct stimuli to indicate the probability of reward, we found that the phasic activation of dopamine neurons varied monotonically across the full range of probabilities, supporting past claims that this response codes the discrepancy between predicted and actual reward. In contrast, a previously unobserved response covaried with uncertainty and consisted of a gradual increase in activity until the potential time of reward. The coding of uncertainty suggests a possible role for dopamine signals in attention-based learning and risk-taking behavior.

The brain continuously makes predictions and compares outcomes (or inputs) with those predictions (1–4). Predictions are fundamentally concerned with the probability that an event will occur within a specified time period. It is only through a rich representation of probabilities that an animal can infer the structure of its environment and form associations between correlated events (4–7). Substantial evidence indicates that do-

pamine neurons of the primate ventral mid-brain code errors in the prediction of reward (8–10). In the simplified case in which reward magnitude and timing are held constant, prediction error is the discrepancy between the probability P with which reward is predicted and the actual outcome (reward or no reward). Thus, if dopamine neurons code reward prediction error, their activation after reward should decline monotonically as the

probability of reward increases. However, in varying probability across its full range ($P = 0$ to 1), a fundamentally distinct parameter is introduced. Uncertainty is maximal at $P = 0.5$ but absent at the two extremes ($P = 0$ and 1) and is critical in assessing the accuracy of a prediction. We examined the influence of reward probability and uncertainty on the activity of primate dopamine neurons.

Two monkeys were conditioned in a Pavlovian procedure with distinct visual stimuli indicating the probability ($P = 0, 0.25, 0.5, 0.75,$ and 1.0) of liquid reward being delivered after a 2-s delay (11). Anticipatory licking responses during the interval between stimulus and reward increased with the probability of reward (Fig. 1), indicating that the animals discriminated the stimuli behaviorally. However, at none of the intermediate probabilities was there a difference in the amount of anticipatory licking between rewarded and unrewarded trials (fig. S1). This suggests that the expectation of reward did not fluctuate significantly on a trial-by-trial basis as a result of the monkey learning the reward schedule (11).

Dopamine neurons of ventral midbrain areas A8, A9, and A10 (fig. S2) were identified solely on the basis of previously described electrophysiological characteristics, particularly the long waveform of their impulses (1.5 to 5.0 ms) (11). The analyses presented here are for the entire population of dopamine neurons sampled, without selection

Institute of Physiology, University of Fribourg, CH-1700 Fribourg, Switzerland, and Department of Anatomy, University of Cambridge, Downing Street, Cambridge CB2 3DY, UK.

*To whom correspondence should be addressed. E-mail: cdf28@cam.ac.uk

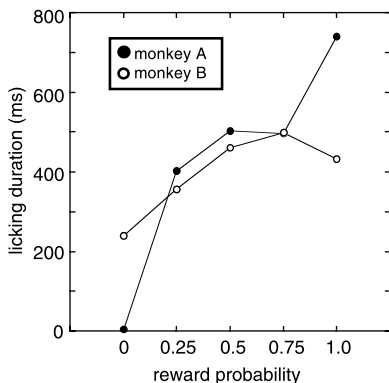


Fig. 1. Conditioned licking behavior increased with reward probability. The ordinate displays the mean duration of licking during the 2-s period from conditioned stimulus onset to potential reward. Each point represents the mean (\pm SEM) of 2322 to 6668 trials. Standard errors are too small to be visible. The behavioral data shown were collected between the first and last day of recordings and include data collected in the absence of physiological recordings.

for the presence of any event-related response. Dopamine neurons ($n = 188$) showed little or no response to fully predicted reward ($P = 1.0$), but they displayed the typical phasic activations (8–10) when reward was delivered with $P < 1.0$, even after extensive training (Fig. 2, A and B). The magnitude of the reward responses increased as probability decreased, as illustrated by linear regression analyses (correlation coefficient $r^2 = 0.97$, $P = 0.002$ and $r^2 = 0.92$, $P = 0.01$ in monkeys A and B, respectively) (Fig. 2C and fig. S3A) (12). Although dopamine neurons

discriminated the full range of probabilities effectively as a population, in contrast to Fig. 2A, many single neurons appeared not to discriminate across the full range (13). For trials in which reward was predicted with intermediate probabilities ($P = 0.25$ to 0.75) but did not occur, neuronal activity was significantly suppressed. The amount of suppression tended to increase with probability ($r^2 = 0.65$, $P = 0.20$ and $r^2 = 0.80$, $P = 0.10$ in monkeys A and B, respectively) (Fig. 2, B and D) although the quantification of suppression may have been limited by the

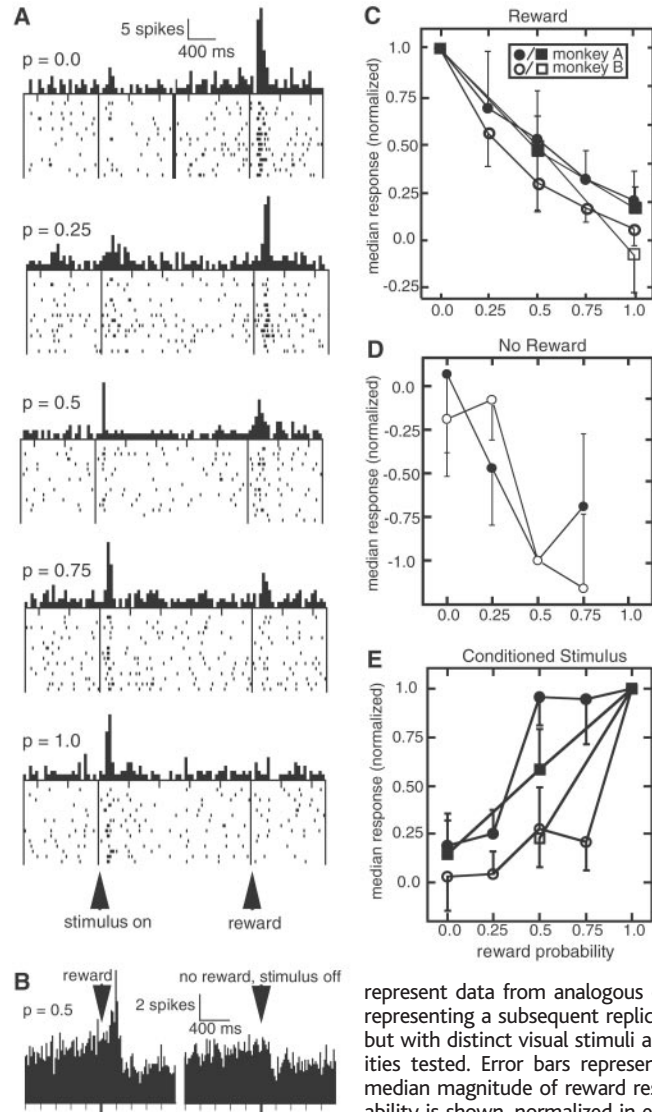


Fig. 2. Dependence of phasic neuronal responses on reward probability. (A) Rasters and histograms of activity in a single cell, illustrating responses to the conditioned stimuli and reward at various reward probabilities, increasing from top to bottom. The thick vertical line in the middle of the top panel ($P = 0$) indicates that the conditioned stimulus response to the left and the reward response to the right were not from a single trial type as in other panels but were spliced together. Reward at $P = 0.0$ was given in the absence of any explicit stimulus at a rate constant of 0.02 per 100 ms and thus presumably occurred with a low subjective probability (11). Only rewarded trials are shown at intermediate probabilities. Bin width = 20 ms. (B) Population histograms of rewarded (left) and unrewarded (right) trials at $P = 0.5$ ($n = 39$, monkey A, set 1). Bin width = 10 ms. (C to E) The median response ($n = 34$ to 62) measured in fixed standard windows, along with symmetric 95% confidence intervals (bars) (11). Circles and squares

represent data from analogous experiments, with the squares representing a subsequent replication of the prior "circle" data but with distinct visual stimuli and only two or three probabilities tested. Error bars represent standard errors. In (C), the median magnitude of reward responses as a function of probability is shown, normalized in each neuron to the response to unpredicted reward. Unpredicted reward caused a median increase in activity that ranged from 76 to 270% above baseline for the four picture sets. Analogous to (C), fig. S3A shows means (\pm SEM) for a subset of responsive neurons (17). In (D), the median magnitude of responses to no reward as a function of probability is shown, normalized in each neuron to the response at $P = 0.5$ ranged from -22 to -55% below baseline. Symbols represent picture sets as shown in (C). At reward probability $P = 0$ for monkey B, a neutral visual stimulus was predicted ($P = 0.5$) by the conditioned stimulus. The data point shows the response after the neutral stimulus failed to occur. In (E), responses to conditioned stimuli are shown, normalized in each neuron to the response to the stimulus predicting reward at $P = 1.0$. The median response to this stimulus ranged from 67 to 194% above baseline. Symbols represent picture sets as shown in (C). The stimuli with $P = 0$ for monkey A, set 2, and for monkey B, set 1, predicted the subsequent occurrence of a neutral visual stimulus with $P = 0.5$.

low spontaneous activity levels. Conditioned stimuli elicited the typical phasic activations (8–10), with their magnitude increasing with increasing reward probability ($r^2 = 0.80$, $P = 0.04$ and $r^2 = 0.69$, $P = 0.08$ in monkeys A and B, respectively) (Figs. 2, A and E, and 3, A and B). In summary, the phasic activations varied monotonically with reward probability, although further conclusions about the quantitative relations are not warranted (13).

The present work revealed an additional, previously unreported activation of dopamine neurons. There was a sustained increase in activity that grew from the onset of the conditioned stimulus to the expected time of reward (Fig. 3, A and B). At $P = 0.5$, 29% of 188 neurons showed significant increases in activity before potential reward, whereas 3% showed decreases ($P < 0.05$, Wilcoxon test). By contrast, at $P = 1.0$, only 9% showed significant increases, and 5% showed significant decreases. For the population response, the sustained activation was maximal at $P = 0.5$, less pronounced at $P = 0.25$ and 0.75 , and absent at $P = 0.0$ and 1.0 (Fig. 3C and fig. S3B). Statistical analysis revealed a significant effect of uncertainty on the population response ($P < 0.005$ in each of four data sets) (11), indicating that the sustained activation codes uncertainty (14). Furthermore, the peak of the sustained activation occurs at the time of potential reward, which corresponds to the moment of greatest uncertainty (15). The particular function of uncertainty signaled by dopamine neurons is not known (13), but we note that common measures of uncertainty (variance, standard deviation, and entropy) are all maximal at $P = 0.5$ and have highly nonlinear relations to probability, being very sensitive to small changes in probability near the extremes ($P = 0$ or 1).

The phasic and sustained activations differed not only in timing and relation to reward probability, but also in their occurrence in single neurons. In Fig. 3D, the magnitude of the phasic and sustained activation is shown for each neuron ($n = 241$). First, a substantial number of neurons had little or no response of either type (13); however, the magnitudes of each type of response fell along a continuum, with no evidence for subpopulations among dopamine neurons. Second, the magnitude of the sustained activation showed no consistent relation to the magnitude of phasic activation across neurons. This was the case both for the phasic response to conditioned stimuli ($r = 0.095$, $P > 0.10$) and for the response to unpredicted reward ($r = -0.024$) (Fig. 3D). In contrast, there was a significant positive correlation of phasic responses between conditioned stimuli and reward ($r = 0.196$, $P < 0.01$) (fig. S4). Thus, the phasic and sustained activations appear to occur inde-

pendently and within a single population of dopamine neurons.

Although the sustained activation occurs in response to reward uncertainty, it is important to know whether it is specific to motivationally relevant stimuli or generalizes to all uncertain events. We conditioned two visual stimuli in a series, with the second following the first in only half of the trials ($P = 0.5$). The stimuli were distinct but entirely analogous. Dopamine neurons showed neither sustained (Figs. 3C and 4A) nor phasic responses (Fig. 2, D and E) to either the first or second of these stimuli. Thus, the sustained activation seems to be related to uncertainty about motivationally relevant stimuli.

If the sustained dopamine activation is related to the motivational properties of uncertain

rewards, it should vary with reward magnitude. We used distinct visual stimuli to predict the magnitude of potential reward at $P = 0.5$ and found that the sustained activation of dopamine neurons increased with increasing reward magnitude ($n = 84$, $P < 0.02$ in each monkey) (Fig. 4A) (11). The sustained activation could reflect the discrepancy in potential reward rather than absolute reward magnitude. To address this issue, we performed an additional experiment (53 neurons in monkey B) in which reward was delivered in each trial but varied between two magnitudes at $P = 0.5$. One stimulus predicted a small or medium reward, another predicted a small or large reward, and a third predicted a medium or large reward. The sustained activation was maximal after the stimulus predicting the largest variation (small versus large reward) ($P < 0.01$) (Fig.

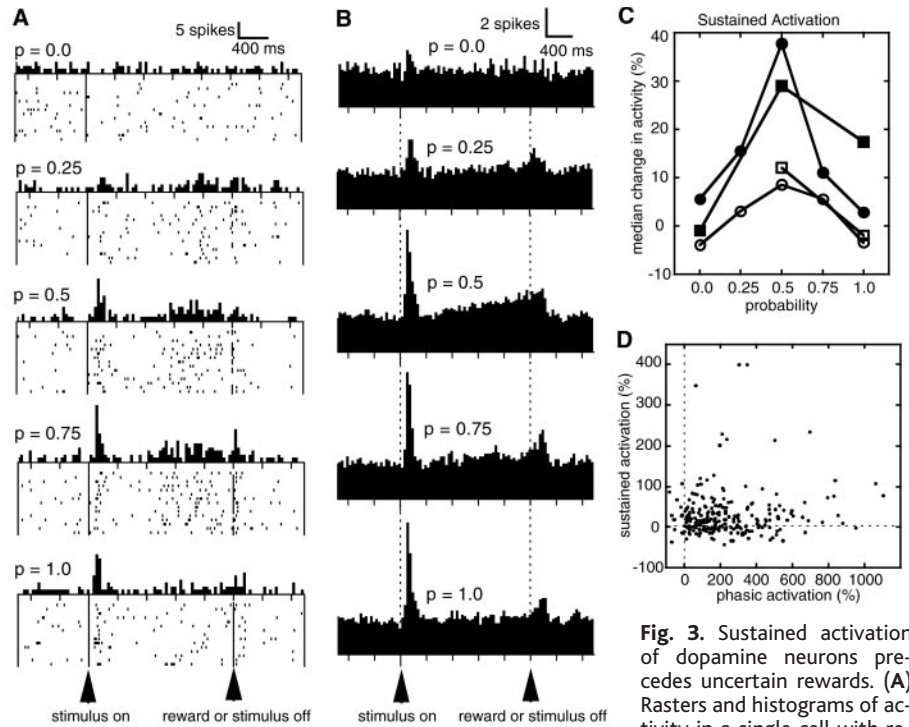


Fig. 3. Sustained activation of dopamine neurons precedes uncertain rewards. (A) Rasters and histograms of activity in a single cell with reward probabilities ranging from 0.0 (top) to 1.0 (bottom). This neuron showed sustained activation before potential reward at all three intermediate probabilities. Both rewarded and unrewarded trials are shown at intermediate probabilities; the longer vertical marks in the rasters indicate the occurrence of reward. Bin width = 20 ms. (B) Population histograms at reward probabilities ranging from 0.0 (top) to 1.0 (bottom). Histograms were constructed from every trial in each neuron in the first picture set in monkey A (35 to 44 neurons per stimulus type; 638 total trials at $P = 0$ and 1200 to 1700 trials for all other probabilities). Both rewarded and unrewarded trials are included at intermediate probabilities. At $P = 0.5$, the mean (\pm SD) rate of basal activity in this population was 2.5 ± 1.4 impulses per second before stimulus onset and 3.9 ± 2.7 in the 500 ms before potential reward. (C) Median sustained activation of dopamine neurons as a function of reward probability. In analogy, means (\pm SEM) are shown in fig. S3B for a subset of responsive neurons (17). Symbols have the same meaning as in Fig. 2C. For monkey A, set 1, the points at $P = 0.25$ and 0.75 may underestimate the amount of sustained activation, as 11 cells with unusually high levels of sustained activity at $P = 0.5$ (median activation of 72%) were not tested at $P = 0.25$ or 0.75 . This was because, at the time of those experiments, the novel form of activation cast doubt on the dopaminergic identity of the neurons. For $P = 0$ in monkey A, set 2, and in monkey B, set 1, there was a 50% chance of a neutral stimulus following the conditioned stimulus. (D) Sustained responses (at $P = 0.5$) plotted against phasic responses to unpredicted reward ($P = 0$) for all neurons recorded in both monkeys (188 neurons, with an additional 53 neurons tested with different reward magnitudes as in Fig. 4B; five outlying neurons, in both dimensions, are not shown).

4B). These data indicate that the amount of sustained activation by reward uncertainty in dopamine neurons increases with the discrepancy between potential rewards.

The present results demonstrate two distinct response types in dopamine neurons. Brief, phasic activations changed monotonically with increasing reward probability, whereas slower, more sustained activations developed with increasing reward uncertainty. These sustained activations were not observed in previous studies in which predictions had low uncertainty. Thus, the activity of dopamine neurons carries information about two intimately related but fundamentally distinct statistical parameters of reward. A potentially analogous coding scheme was identified in neurons of the fly visual system, in which the visual stimulus and uncertainty about that visual stimulus appeared to be coded independently in single neurons (16).

By systematically varying reward probability, we show that the phasic activity of dopamine neurons matches the quantitative definition of reward prediction error. Re-

sponses to reward decreased with increasing reward probability, and, conversely, responses to the predictive stimulus increased. Furthermore, reward always elicited responses when it occurred at $P < 1$, even after thousands of pairings between stimulus and reward. By always coding prediction error over the full range of probabilities, dopamine neurons could provide a teaching signal in accord with the principles of learning originally described by Rescorla and Wagner (17–19).

In addition to those principles described by Rescorla-Wagner, other basic intuitive principles of associative learning have been described, focusing in particular on the importance of attention (20, 21). It is generally accepted that no single principle alone is sufficient to explain all observations of animal learning, and the various theories are thus considered to be complementary (6, 7). The Pearce-Hall theory proposes that attention (and thus learning) is proportional to uncertainty about reinforcers (21, 22). As dopamine neurons are activated by reward uncertainty, dopamine could facilitate attention and learning in accord with the Pearce-Hall theory. This raises the possibility that two fundamental principles of learning are embodied by two distinct types of response in dopamine neurons (23).

The link between uncertainty, attention, and learning has two related aspects [another aspect is given in (24)]. The goal of learning can be seen as finding accurate predictors for motivationally significant events. Subjective uncertainty indicates that the animal lacks an accurate predictor and thus indicates the utility of identifying a more accurate predictor (25). Similarly, and as indicated by mathematical principles of information (26), only in the presence of uncertainty is it anticipated that there will be information available in the outcome. If reward ($P = 1$) or no reward ($P = 0$) occurs exactly as predicted, that event contains no information beyond that already given by the conditioned stimulus; that is, it is redundant. However, when the prediction of reward is uncertain, the outcome (reward or no reward) always contains information. The outcome at $P = 0.5$ contains, on average, the maximal amount of information (one bit) of any probability. The processing of this reward information is demonstrated by the fact that prediction error signals are always generated in dopamine neurons when reward outcomes occur under conditions of uncertainty. Thus, subjective reward uncertainty corresponds both to the utility of identifying more accurate predictors and to the expectation of reward information. Through its widespread influence, dopamine could control a nonselective form of attention or arousal, which is dependent on uncertainty and designed to aid the learning of predictive stimuli and actions.

Although dopaminergic signals may promote a particular form of attention, an extensive literature has already established the critical importance of dopamine in reward and reinforcement. Whereas the phasic response of dopamine neurons to reward prediction error fits remarkably well with dopamine's presumed role in appetitive reinforcement (10, 17, 18), the activation by reward uncertainty may appear inconsistent with a reinforcing function. This apparent discrepancy would be resolved to the extent that postsynaptic neurons can discriminate the two forms of activity. However, it seems unlikely that the two patterns of activity can be discriminated perfectly, especially given the slow time course of dopamine transmission. Rather than arguing against a role for the activity of dopamine neurons in reinforcement, one might ask whether reward uncertainty itself has rewarding and reinforcing properties. Indeed, gambling behavior is defined by reward uncertainty and is prevalent throughout many cultures. Animals display a potentially related behavior, preferring variable over fixed reward schedules [for discussion, see (27) and (28)]. The present results suggest that dopamine is elevated during gambling in a manner that is dependent on both the probability and magnitude of potential reward. This uncertainty-induced increase in dopamine could contribute to the rewarding properties of gambling, which are not readily explained by overall monetary gain or dopamine's corresponding role in prediction error (as losses tend to outnumber gains) (29). The question arises as to why a reward signal would be produced by reward uncertainty. Although risk-taking behavior may be maladaptive in a laboratory or casino, where the probabilities are fixed and there is nothing useful to learn, it could be advantageous in natural settings, where it would be expected to promote learning of stimuli or actions that are accurate predictors of reward (25). Thus, the sustained, uncertainty-induced increase in dopamine could act to reinforce risk-taking behavior and its consequent reward information, whereas the phasic response after prediction error could mediate the more dominant reinforcement of reward itself.

References and Notes

1. R. P. N. Rao, D. H. Ballard, *Nature Neurosci.* **2**, 79 (1999).
2. D. M. Wolpert, Z. Ghahramani, *Nature Neurosci.* **3** (suppl.), 1212 (2000).
3. E. K. Engel, P. Fries, W. Singer, *Nature Rev. Neurosci.* **2**, 704 (2001).
4. R. P. N. Rao, B. A. Olshausen, M. S. Lewicki, Eds., *Probabilistic Models of the Brain* (MIT Press, Cambridge, MA, 2002).
5. C. R. Gallistel, *The Organization of Learning* (MIT Press, Cambridge, MA, 1990).
6. A. Dickinson, *Contemporary Animal Learning Theory* (Cambridge Univ. Press, Cambridge, 1980).
7. J. M. Pearce, *An Introduction to Animal Cognition* (Lawrence Erlbaum, Hove, UK, 1987).
8. W. Schultz, P. Apicella, T. Ljungberg, *J. Neurosci.* **13**, 900 (1993).
9. P. Waelti, A. Dickinson, W. Schultz, *Nature* **412**, 43 (2001).

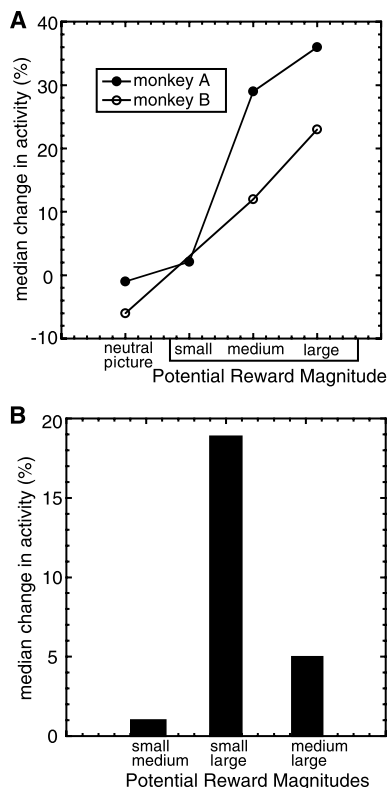


Fig. 4. Sustained activation is dependent on the discrepancy in potential reward magnitude. **(A)** All stimuli predicted potential reward (0.05, 0.15, or 0.5 ml of liquid) or a neutral picture at $P = 0.5$. Data are from 35 cells in monkey A and 49 cells in monkey B. **(B)** Each stimulus predicted that reward would be one of two potential magnitudes, each at $P = 0.5$, as indicated on the abscissa. Every trial was rewarded with one of the two potential reward magnitudes. Data are from 53 cells in monkey B.

10. W. Schultz, *J. Neurophysiol.* **80**, 1 (1998).
11. Materials and methods are available as supporting material on *Science Online*.
12. Simple linear regression coefficients for each type of phasic response were calculated for each set of data for which all probabilities were tested ($P = 0.0, 0.25, 0.5, 0.75$, and 1.0 in Fig. 2, C and E; $P = 0.0, 0.25, 0.5$, and 0.75 in Fig. 2D). This was done only as an approximation and does not imply linearity in the response functions. In addition to the nonlinear factors discussed in (13), there is imprecision in the subjective timing of the 2-s interval between stimulus onset and potential reward (15). This probably accounts for the small but significant activation to "fully" predicted reward in monkey A (Figs. 2C and 3B).
13. Unpublished data (30), as well as Figs. 3 and 4, suggest that the responses of dopamine neurons multiplicatively combine the probability and magnitude of reward. Thus, it is not necessarily the case that the maximal responses observed in this study for a given reward magnitude (those at $P = 0.0, 0.5$, or 1.0 , depending on the type of response) are actually the maximal evoked responses of a given neuron. One would expect that, like other neurons coding the intensity of a signal, dopamine neurons have a stimulus-response function that is sigmoid, being insensitive to values above or below a particular range. The likelihood that individual neurons have distinct thresholds has critical implications for understanding the shape of the probability-response functions presented in Figs. 2 and 3 and could explain why many neurons shown in Fig. 3D appear to be unresponsive. The shape of the probability functions that we measured would depend on the range of values to which most of the neurons are sensitive. Because these ranges are unknown, the only interpretation that should be given to the data at this time is that dopamine neuronal responses follow probability or uncertainty in a monotonic fashion.
14. The present experiments were performed with a standard delay conditioning procedure, meaning that the conditioned stimulus remained on for the full 2-s delay until the potential time of reward. In a separate experiment, a smaller number of neurons ($n = 22$) were tested with trace conditioning in which the conditioned stimulus indicating the probability of reward was on for 1 s, and potential reward occurred following an additional 1-s interval after stimulus offset. Although there may have been some sustained activation in the trace condition at $P = 0.5$ ($P < 0.1$), the activity preceding potential reward (during either 250- or 500-ms periods) was significantly less than that in experiments with delay conditioning ($P < 0.05$, Mann-Whitney test). Furthermore, a distinct behavioral pattern emerged with trace conditioning; the likelihood of licking increased before stimulus offset, decreased subsequently, and then increased again before reward. The explanation for the apparent discrepancy between trace and delay conditioning is unclear, but it could be related to the presence of temporal information provided by the continued presence of the delay stimulus; that is, as long as the delay stimulus is present, the time of potential reward must not have passed, and this information could suppress incoming inhibitory signals that are (imprecisely) timed to coincide with potential reward (15).
15. Objectively, potential reward always occurred after a 2-s delay. However, it is known that subjective timing is imprecise. Thus, the time course of the slowly developing sustained activation could reflect the increasing likelihood that the interval is nearing completion. Unpublished data (30) on the phasic activation of dopamine neurons to the delivery of reward earlier or later than predicted suggest a similar degree of temporal imprecision in the prediction. It is therefore reasonable to hypothesize that dopamine neurons code the uncertainty in reward in the subsequent moment (the very near future).
16. A. L. Fairhall, G. D. Lewin, W. Bialek, R. R. de Ruyter van Steveninck, *Nature* **412**, 787 (2001).
17. R. R. Montague, P. Dayan, T. J. Sejnowski, *J. Neurosci.* **16**, 1936 (1996).
18. W. Schultz, P. Dayan, R. R. Montague, *Science* **275**, 1593 (1997).
19. R. A. Rescorla, A. R. Wagner, in *Classical Conditioning II: Current Research and Theory*, A. H. Black, W. S. Prokasy, Eds. (Appleton-Century-Crofts, New York, 1972), pp. 64–69.
20. N. J. Mackintosh, *Psychol. Rev.* **82**, 276 (1975).
21. J. M. Pearce, G. A. Hall, *Psychol. Rev.* **87**, 532 (1980).
22. H. Kaye, J. M. Pearce, *J. Exp. Psychol. Anim. Behav. Process.* **10**, 90 (1984).
23. The fact that there are two distinct dopamine signals, each with unique properties, suggests two distinct functions for dopamine. However, this does not necessarily imply that the two signals must be processed independently. Thus, each signal may contribute to the performance of two or more functions. Furthermore, questions concerning the functions of dopamine in target areas (such as reinforcement and attention) are distinct from questions about the qualitative nature of stimuli (rewarding versus attention-inducing) that contribute to the activation of dopamine neurons (31).
24. P. Dayan, S. Kakade, P. R. Montague, *Nature Neurosci.* **3** (suppl.), 1218 (2000).
25. In the artificial, impoverished conditions of a laboratory setting or a casino, the probabilities associated with particular stimuli or actions are fixed, and there is nothing else useful to be learned. However, the natural environment contains a high degree of correlation between a multitude of events; this is implicit in the adaptive utility of associative learning. Thus, an animal should not assume that uncertainty signals the objective absence of accurate predictors but rather that it is ignorant of those predictors. Although accurate predictors of reward may not always be present in the environment, one would not expect the learning machinery of the brain to assume their absence. In fact, there is a period of uncertainty about all rewards before accurate predictors are found. If subjective uncertainty is assumed to result from ignorance of predictors rather than absence of predictors, then it would be appropriate for subjective uncertainty to have attention-inducing and reinforcing properties that would ultimately enhance learning and reduce uncertainty.
26. C. E. Shannon, *Bell Syst. Tech. J.* **27**, 379 (1948).
27. C. R. Gallistel, J. Gibbon, *Psychol. Rev.* **107**, 289 (2000).
28. J. E. Mazur, *Psychol. Rev.* **108**, 96 (2001).
29. Alternative attempts to explain gambling behavior focus on the fact that people (particularly those with prefrontal deficits) may misperceive reward probabilities or magnitudes or combine them in an inappropriate manner. This "cognitive" hypothesis fails to explain why gambling is appealing (and sometimes addictive) to a large number of otherwise healthy people, most of whom are aware that the odds are against them and that they have lost and will continue to lose money. In addition, any attempt to explain gambling behavior must address the fact that gambling is common at all probabilities (except $P = 0$ or 1 , by definition) and all reward magnitudes. The present work suggests that activation of dopamine neurons may occur to a comparable extent during the expectation of a small reward at intermediate probabilities or a large reward at low probabilities. Thus, dopamine could contribute to the appeal of gambling in general. A behavior as prevalent as gambling must be explained in terms that are consistent with natural selection. The present hypothesis does so by pointing out that risk-taking promotes learning in natural environments.
30. C. D. Fiorillo, P. N. Tobler, W. Schultz, unpublished data.
31. P. Redgrave, T. J. Prescott, K. Gurney, *Trends Neurosci.* **22**, 146 (1999).
32. We thank A. Dickinson, S. Baker, R. Moreno, K. Tsutsui, I. Hernadi, P. Dayan, and two anonymous reviewers for helpful comments on the manuscript. Funding was provided by the Human Frontiers Science Program (C.D.F.), Swiss National Science Funds (W.S. and P.N.T.), and Wellcome Trust (W.S.).