SUPPORTING ONLINE MATERIAL

Material and Methods

The individual animals, the basic design of the experiments and the electrophysiological techniques for extracellularly recording from dopamine neurons were identical to those previously reported (10). All procedures were performed in Fribourg, complied with the Swiss Animal Protection Law and were supervised by the Fribourg Cantonal Veterinary Office.

Experimental design. Two adult female *Macaca fascicularis* monkeys were mildly fluid deprived. They were trained in a Pavlovian procedure in which distinct visual stimuli predicted specific amounts of sweetened liquid (0.00 ml, 0.05 ml over 40 ms, 0.15 ml over 100 ms, or 0.50 ml over 240 ms) with specific probabilities (P = 0.0, 0.5, or 1.0) (Fig. S1). We used not more than two rewards per stimulus, which allowed us to explore several stimuli with different reward conditions during the limited testing period with each neuron. We assume the frequency and amount of liquid to provide reasonable approximations of the animals' estimates of the probability and magnitude of reward. Stimuli were chosen to have similar physical salience but to be easily discriminated. To aid discrimination, each stimulus was presented at a unique location on the computer monitor. Liquid was delivered via a computer-controlled solenoid valve from a spout in front of the animal's mouth. The onset of liquid delivery occurred 2 s after the onset of visual stimuli, and offsets of visual stimuli and liquid flow coincided. Licking behavior was monitored with an infrared detector. 'Unpredicted' liquid, not signaled by any immediately preceding stimulus, was delivered to each neuron in a separate block of trials. The inter trial interval

(from reward to next conditioned stimulus or reward) averaged 9 s, consisting of a fixed 4 s plus an exponentially distributed interval with a mean of 5 s.

The computer that controlled behavior did not deliver liquid in a completely random manner. To prevent long streaks in which a stimulus was repeatedly followed by the same reward outcome, the program insured that the actual frequencies would precisely match the assigned probabilities after 8 consecutive trials of a specific visual stimulus. The 'counter' was reset if the experimenter interrupted the recording for more than a few seconds. Although it would seem to be difficult given the intermixed trial types, it would by possible in principle for an animal to learn this structure and thereby reduce its uncertainty about reward. Previously published analysis of behavior and neural data suggests that the animals did not learn to take advantage of this structure (10).

Training consisted of 100–200 trials of each stimulus per day, five days per week, for about five weeks. Recordings began only after substantial pretraining (5-8 days and 600–1500 trials of each type) and emergence of discriminative conditioned licking responses during the stimulus and preceding the time of reward.

Electrophysiological Recordings. As previously described (S1, 8-11), dopamine neurons in the substantia nigra and ventral tegmental area were identified solely by their discharge characteristics, including low basal firing rates (0.1 – 8.0 Hz) and long duration, initially negative or positive waveforms (1.5 – 5.0 ms, high-pass filtered at 100 Hz and -3 dB). Prior studies in primates have shown that ventral midbrain neurons having these properties are antidromically activated by stimulation of the striatum, and their firing is suppressed by systemic administration of the dopamine D2 receptor agonist apomorphine

(S1). These characteristics are similar to those of identified dopaminergic neurons in other mammalian species (e.g. S2, S3, S4).

Recording sites. Recording sites were marked with small electrolytic lesions and reconstructed from 40 μm thick, stereotaxically oriented coronal brain sections, stained with cresyl violet or antibodies to tyrosine hydroxylase. Recording sites overlapped substantially with those described in a previous report which shows plots of neuronal positions relative to regions of dense tyrosine hydroxylase staining (10). Planes of recorded neurons ranged from 5.5 to 10.5 mm anterior to the interaural line.

Data analysis. Statistical analysis of neural activity followed our previously described methods (8, 10). Typically, at least 15 trials of each trial type were performed per neuron; the minimum accepted trial number for analysis was 7. Average firing rates were measured in standard time windows (see below) and divided by the average rate in a 1 s control period immediately preceding event onset to calculate the percent change in impulse rate. These values were normalized by dividing them by the response to an analogous event (either a visual stimulus or liquid delivery) recorded in the same neuron. Normalized percent changes were used for both statistical analysis and graphical display. The 95% confidence intervals were calculated in the same manner as in the preceding report (10), multiplying the appropriate t-value by the interquartile range and dividing by 1.075 times the square root of the number of observations (S5). Activity in the standard time windows was compared to the 1 s control activity using a Wilcoxon matched-pairs, signed rank test on normalized counts in each trial with each neuron ($p<0.01$). We employed the Mann-Whitney test for assessing the discrimination between different trial types within single neurons ($p<0.01$) and the Wilcoxon test for comparing responses within

populations of neurons. The Bonferroni method was used to correct for multiple comparisons.

Standard time windows were fixed across trial types and across neurons, and were chosen so as to capture most of the period in which neural activity changed. Following onset of visual stimuli, the windows were 90-180 ms for monkey A and 110-240 ms for monkey B. For responses following liquid onset, or visual stimulus offset in the case of no reward, the window was 120-320 ms in both monkeys. Peak dopamine responses are typically delayed by about 150 – 200 ms after an error event. A single window was chosen to capture both the periods of suppression and excitation.

A particular time window of 250–400 ms was employed for the specific experiment shown in figure 3A, B, because responses were spread over a longer duration due to prolonged liquid flow with unexpectedly higher volumes. In many past experiments in our laboratory, the animals were able to predict that at a particular moment in time, a drop of a known volume of liquid either would or would not be delivered. A particular volume of liquid always corresponds to a particular duration of liquid flow, so that if a particular volume is expected, then the onset of liquid flow can be used to predict its overall duration. Thus the prediction error, and the dopamine response, is time locked to the onset and does not continue for the duration of the liquid flow. In some of the present experiments however, and particularly that shown in figure 3A, B, both the theoretical prediction error and the dopamine response are spread out over time. In figure 3A, the activation can be seen to be particularly sustained in response to 0.5 ml of liquid flowing for 240 ms. Most other neurons tested in this experiment showed similarly long-lasting responses. In principle, the positive error signal in this case would begin only after 120 ms, since the

expected liquid volume lasts only for 120 ms, and would continue until 240 ms when liquid flow stops. The negative error signal to the small reward (0.05 ml over 40 ms) in this experiment would not be expected to begin until 40 ms.

Additional analysis of data shown in Figure 4

The sensitivity or gain of the neural responses as a function of liquid volume adapted according to the prediction made by the visual stimulus, so that responses appeared to be equivalent regardless of their absolute magnitude (Fig. 4). We considered two hypotheses concerning what aspect of the prediction evoked the adaptation. First, the adaptation in sensitivity may have consisted of normalization to some measure of the discrepancy between likely outcomes, such as the range or standard deviation. Alternatively, normalization could have occurred to the expected value. The experiments were not originally designed to discriminate between these two possibilities, and in the experiment depicted in figure 4C left, expected value and range perfectly covaried. However, in the experiment of figure 4C right, the two varied in a partially independent manner across visual stimuli, and therefore this data set provided an opportunity to compare the two hypotheses. The neural responses of figure 4C right were replotted after normalizing the abscissa by either the difference (range) in potential volumes (Fig. S2 top) or by expected liquid volume (mean) (Fig. S2 bottom). The observation that neural responses in all three conditions appeared to be identical could be explained by the fact that all pairs of reward outcomes were exactly one range apart (Fig. S2 top). By contrast, when liquid volume is expressed in units of the mean, the difference between pairs of reward outcomes ranged from 1.00 to 1.64 means (Fig. S2 bottom), and yet neural responses appeared insensitive to this discrepancy. This did not appear to be due to saturation of the response, since

responses to unpredicted volumes of 0.15 ml in the same neurons were about twice as large (Fig. S2). In order to statistically compare the two normalization procedures, we compared the slopes for each pair of reward outcomes (Fig. S2). The slopes did not differ from one another after normalizing by the range (Fig. S2 top) ($p > 0.2$ for all three comparisons, Wilcoxon paired sample test, $n = 53$), but the slope corresponding to liquid volumes of 0.05 and 0.50 was significantly less than either of other two after normalizing by the mean (Fig. S2 bottom) ($p < 0.001$). To directly compare the effect of normalization by range versus mean on the slopes, the difference between the slope for the 0.05–0.50 ml pair and the mean of the other two slopes was divided by the mean slope. This ratio was calculated in each neuron after normalization to the mean, and again after normalization to the range, and was significantly greater after normalization by the mean ($p<0.0001$, $n = 53$, Wilcoxon paired sample test). This analysis suggests that normalization by the range could account for the identical responses, whereas normalization by the mean or expected value would not in itself appear to be fully sufficient to account for the identical responses. Although the present evidence on this point is limited, it suggests that normalization by the range provides the more parsimonious explanation. As the range perfectly covaried with the standard deviation in all the present experiments, the observed adaptation appeared to occur relative to the standard deviation, which is an accepted measure of uncertainty. Furthermore, past experiments indicate that the sustained, delay-period activity of dopamine neurons may represent the standard deviation or some other measure of uncertainty. Studies on motion-sensitive neurons of the fly suggest that they possess information about the standard deviation and use it for normalization in a manner analogous to what we observe in dopamine neurons (21, 22).

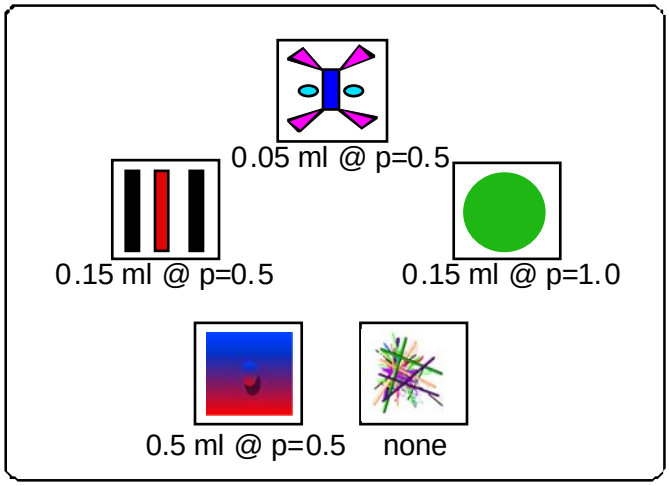

Fig. S1 Visual stimuli indicated probabilities of various liquid volumes. One stimulus was presented in each trial on a computer monitor directly in front of the animal. Each stimulus was always presented in the same unique location. The particular stimuli illustrated here were used in animal A in the experiments illustrated in figures 1 and 4. A particular image was never used in more than one experiment in an individual animal. Different images were used in Animal B.

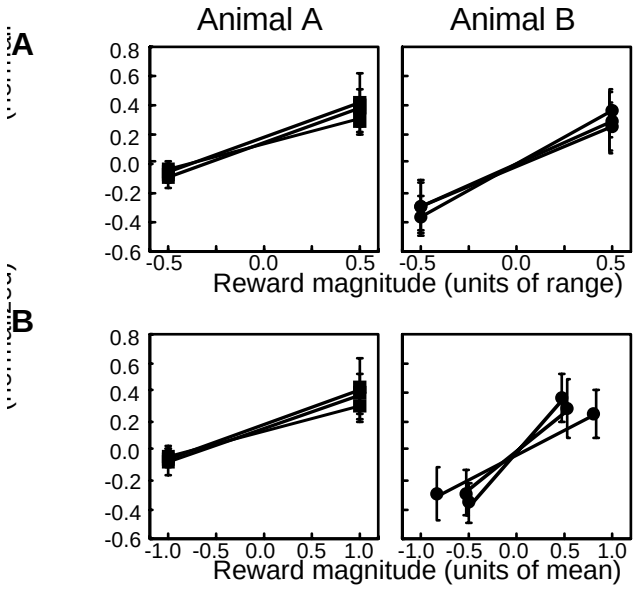

Fig. S2. Adaptation of neural sensitivity to liquid volume following reward-predicting stimuli. Same data as in figure 4C, but replotted after normalizing the abscissa by either the range of potential liquid volumes predicted by a visual stimulus (top), or by the expected value (mean) indicated by a visual stimulus (bottom). Each line connects a pair of points representing the two potential reward outcomes predicted by a distinct visual stimulus. Each point represents the median response (±95% confidence intervals) of the population taken after normalizing to the response following unpredicted reward recorded in the same neuron (0.15 ml; median activation of 266% in animal A, n = 57, and 97% in animal B, n = 53). The lesser variation of the slopes in panel A suggests that dopamine neurons or their inputs may normalize the liquid volumes by range or standard deviation rather than expected value or mean.

References and Notes

S1. W. Schultz, R. Romo, *J. Neurophysiol.* **57**, 201 (1987).

S2. A.A. Grace, B.S. Bunney, *Neuroscience* **10**, 301 (1983).

S3. B.S. Bunney, Aghajanian, R.H. Roth, *Nature New Biology* **245**, 123 (1973).

S4. A. Studer, E. Sundström, G. Jonsson, W. Schultz, *Neuropharmacology* **27**, 923 (1988).

S5. B. Iglewicz, in *Understanding Robust and Exploratory Data Analysis*. D.C. Hoaglin, F. Mosteller, J.W. Tukey (Wiley, New York, 1983), pp. 404-430.