You might find this additional information useful...

This article cites 33 articles, 10 of which you can access free at:
http://jn.physiology.org/cgi/content/full/98/3/1428#BIBL

Updated information and services including high-resolution figures, can be found at:
http://jn.physiology.org/cgi/content/full/98/3/1428

Additional material and information about *Journal of Neurophysiology* can be found at:
http://www.the-aps.org/publications/jn

This information is current as of September 11, 2007 .

Statistics of Midbrain Dopamine Neuron Spike Trains in the Awake Primate

Hannah M. Bayer,[1] **Brian Lau,**[1,2] **and Paul W. Glimcher**[1]

[1]*Center for Neural Science, New York University; and* [2]*Center for Neurobiology and Behavior, Columbia University, New York, New York*

Bayer HM, Lau B, Glimcher PW. Statistics of midbrain dopamine
neuron spike trains in the awake primate. *J Neurophysiol* 98: 1428–1439,
2007. First published July 5, 2007; doi:10.1152/jn.01140.2006. Work in
behaving primates indicates that midbrain dopamine neurons encode a
prediction error, the difference between an obtained reward and the
reward expected. Studies of dopamine action potential timing in the alert
and anesthetized rat indicate that dopamine neurons respond in tonic and
phasic modes, a distinction that has been less well characterized in the
primates. We used spike train models to examine the relationship be-
tween the tonic and burst modes of activity in dopamine neurons while
monkeys were performing a reinforced visuo-saccadic movement task.
We studied spiking activity during four task-related intervals; two of
these were intervals during which no task-related events occurred,
whereas two were periods marked by task-related phasic activity. We
found that dopamine neuron spike trains during the intervals when no
events occurred were well described as tonic. Action potentials appeared
to be independent, to occur at low frequency, and to be almost equally
well described by Gaussian and Poisson-like (gamma) processes. Unlike
in the rat, interspike intervals as low as 20 ms were often observed during
these presumptively tonic epochs. Having identified these periods of
presumptively tonic activity, we were able to quantitatively define phasic
modulations (both increases and decreases in activity) during the intervals
in which task-related events occurred. This analysis revealed that the
phasic modulations of these neurons include both bursting, as has been
described previously, and pausing. Together bursts and pauses seemed to
provide a continuous, although nonlinear, representation of the theoreti-
cally defined reward prediction error of reinforcement learning.

INTRODUCTION

Over the past three decades, a significant amount of data has
been gathered about the spiking properties of midbrain dopa-
mine neurons. Intracellular recordings made from dopamine
neurons in rodent slice (Johnson et al. 1992; Kita et al. 1986),
studies of the intact anesthetized rodent (Grace and Bunney
1984; Kitai et al. 1999; Tepper et al. 1995), and studies of
awake behaving rodents (Freeman et al. 1985; Hyland et al.
2002) have all contributed to our understanding of the biophys-
ical properties of this system during both tonic and phasic
modes of activation (Grace 1991).

Recently, systems level studies of the activity of these
neurons in the awake behaving primate have begun to indicate
that phasic activity after a reward is systematically related to
the difference between the magnitude of behavioral reinforce-
ment received by the primate and the magnitude of the rein-
forcement that the primate is presumed to expect (Schultz
1998). At this level of analysis, there has been growing
evidence that transient increases in spike rate, the phasic bursts
that have been widely observed in the rodent preparation,
appear to continuously encode positively valued differences

between the expected and obtained reward, the reward predic-
tion error (Bayer and Glimcher 2005; Morris et al. 2004;
Waelti et al. 2001). In a similar way, current evidence suggests
that reductions in baseline activity may be related to negatively
valued reward prediction errors (Hollerman and Schultz 1998;
Ljungberg et al. 1992), a feature consistent with temporal-
difference models of reinforcement learning (Montague et al.
1997; Schultz et al. 1997).

We were interested in using an existing database of dopa-
mine spiking patterns (Bayer and Glimcher 2005) to examine
the statistics of dopamine firing rates in the awake behaving
primate to ask three questions. *1*) How are the phasic responses
(both increases and decreases in activity) of these cells related
to the theoretically defined reward prediction error? *2*) How do
the phasic modulations of these neurons by rewards relate to
the tonic and burst modes of activation observed in the rodent?
3) Do the statistics of action potential generation in the alert
animal support the notion that these cells can serve as pace-
makers as has been proposed in the rodent (Meck and Benson
2002)?

We therefore examined the spiking properties of midbrain
dopamine neurons while monkeys were learning, by trial-and-
error, when to make an eye movement to receive a fluid
reward. Within that context, we examined the properties of
dopamine spike trains under four conditions: after an auditory
tone that initiated each trial, while the animal was waiting to
make the eye movement, after the delivery of the reward, and
during an epoch measured between trials when no stimuli or
rewards were presented. The second and fourth conditions
were associated with continuous average levels of activity that
might be expected to reflect the tonic mode of activation that
has been observed previously in the rodent. The first and third
conditions were associated both with the existence of reward
prediction errors at a systems level and the generation of phasic
responses at a physiological level.

We began by quantitatively examining the interspike inter-
vals (ISIs) of dopamine (DA) neuron activity when no task-
related activation was expected, focusing our analysis on the
epoch between trials. We found that, during these intervals, the
neurons fired at a low continuous rate with moderate variabil-
ity. We saw little evidence either of strongly periodic behavior
(of the type observed in other known pacemaker circuits) or of
phasic modulations during these intervals. Our analysis of the
distribution of ISIs during tonic activity was more closely
matched to the irregular dopaminergic spike activity found in
anesthetized animals than the dopaminergic pacemaker activity
observed predominantly in vitro. Although DA neurons in the
awake primate do spike at reasonably regular intervals under

---

these conditions, their ISIs were marginally better described as Poisson-like (Gamma) than as Gaussian in their distribution. Based on our analysis of this presumptively tonic pattern of activation, we defined phasic modulations in activity as ISIs above or below the 95% CIs observed during tonic activity. In this way, we identified bursts of activity similar to those seen in the awake behaving rat (Hyland et al. 2002). The frequency and duration of these bursts was correlated with the magnitude of a mathematically defined positive reward prediction error as might be expected from previous work in the primate (Mirenowicz and Schultz 1994; Schultz et al. 1997; Waelti et al. 2001). We also, however, observed a pause in activity that occurred on some trials that has been observed previously in the primate (Hollerman and Schultz 1998; Ljungberg et al. 1992) but has received less attention than the burst response. The durations of pauses in activity, periods during which the neurons were completely silent, were also systematically correlated with a mathematically defined reward prediction error (RPE), but in this case with negatively valued reward prediction errors. These results suggest that the afferent inputs to the DA neurons, which are known to be correlated with D1 receptor activity and that are known to control phasic bursting (Floresco et al. 2003; Goto and Grace 2005; Grace 1991) may also initiate pauses under some circumstances. Thus while we observed both tonic and phasic patterns of activity, the phasic pattern we observed included modulations of the durations of pauses in activity, which is not a feature of existing reinforcement learning models of dopaminergic activity.

METHODS

Two male rhesus macaques (*Macaca mulatta*) were used as subjects. All animal procedures were developed in association with the University Veterinarian, approved by the New York University Institutional Care and Use Committee, and designed and conducted in compliance with the Public Health Service's *Guide for the Care and Use of Laboratory Animals*. All surgical and training procedures were performed using standard protocols that have been described in detail previously (Handel and Glimcher 1997). The database of spike patterns analyzed here with regard to the statistics of neuronal action potentials served as the subject of a previous report on reinforcement learning (Bayer and Glimcher 2005).

Recording protocol

We used ultrasonography to place guide tubes and electrodes in the ventral midbrain (Glimcher et al. 2001). Neurons at, or caudal to, the anatomical location of the substantia nigra pars compacts (SNc), as determined by ultrasonography, were classified as dopaminergic based on three criteria: they had relatively long action potentials (typically ~2 ms), their baseline firing rates were relatively low (mean: 5.3 ± 1.5 spikes/s), and they had a phasic response to unpredicted fluid rewards. A subset of these neurons, which were typical of the population, were localized histologically to the SNc and the ventral tegmental area (VTA).

To ensure that we had successfully isolated single neurons, we visually assessed spike waveforms for identity before beginning data collection and throughout the recording process. After all data were collected, we created ISI histograms (ISIHs) and examined, for each cell, the occurence of intervals below a conservative estimate (2 ms) of the biophysical refractory period for these cells. Fifteen of the neurons in our population had <0.1% of their observed ISIs shorter than this interval. Because there were typically ~1,000–2,000 ISIs per cell, this meant that in all likelihood there were less than one or

two recorded action potentials of dubious provenance for each unit in this group. We compared the ISI distributions for these very well-isolated units to the ISI distributions for our entire population. For the entire population, we observed <1% of the ISIs were <2 ms in any unit studied, and on further analysis, we found no detectable difference in the ISI distributions for the population as a whole compared to the 15 best isolated neurons.

Task

Monkeys were trained to perform a saccade timing task in which they had the opportunity to learn, by trial-and-error, when to initiate a saccade to an eccentric target without an external "go"-cue. Saccade timing trials (Fig. 1) began after an intertrial interval of unpredictable duration followed by an audible beep. Three hundred milliseconds later, a central yellow light emitting diode (LED) was illuminated, and the subject was required to align gaze with this stimulus (±3°) within 1,000 ms. Three hundred milliseconds after gaze was aligned with this central LED, it turned red, and a single red eccentric LED was illuminated at 10° of vertical elevation (the location of the target was identical during all experiments). During the next 4 s, the subject could initiate a saccade to the eccentric target at any time. After gaze was shifted into alignment with the eccentric LED, the subject was required to maintain gaze for another 250 ms. Both LEDs were extinguished, and the subject either received a reward or not. However, a new trial would not begin until the 4-s interval was complete.

During each trial, the subject received a reward if he executed the saccade during an unsignalled temporal window that was embedded in the 4-s interval. From the beginning to the end of the window, the volume of liquid reward that the animal could earn increased linearly; thus the animal could maximize his reward by learning where the end of this unsignalled window occurred and choosing to make his saccade during that interval. The duration of the window was scaled logarithmically as it was moved later in the trial. The temporal position of the interval was shifted between blocks of trials in an uncued manner (see Bayer and Glimcher 2005 for more details).

Data analysis

For each behavioral trial on which the animal made a saccade to the eccentric target, we measured how long the animal waited to make the saccade, the interval during which the saccade would be rewarded, the volume of liquid reward that the animal received, and the times at which action potentials occurred during four intervals within the task: two intervals of presumptively tonic activity and two intervals that included phasic activity (Fig. 1). The tonic intervals were a 1,500-ms baseline period starting 800 ms after the reward from the previous trial was delivered, and a variable length wait period, starting with the onset of the eccentric target, and extending until the saccade was initiated. The length of
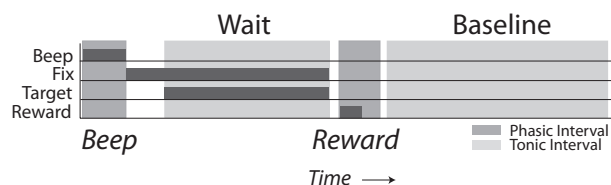

FIG. 1. Saccade timing task events and intervals presented as a function of time. Each trial begins with a 300-ms warning tone marked beep. A fixation target is illuminated for 300 ms after which a single eccentric target is also illuminated for ≤4,000 ms. The monkey is free to make a saccade shifting gaze to the eccentric target at any time during this interval. Depending on when the monkey makes his eye movement, a reward of variable magnitude is delivered after the eye movement is complete. An intertrial interval (variable length, ≤4,000 ms) marked baseline ensues. Light gray bars mark intervals during which tonic-related activity was measured. Dark gray bars indicate intervals during which phasic activity was analyzed.

the wait interval ranged from 300 to 3,000 ms, depending on how long the animal waited to make the saccade. The phasic intervals were a 500-ms beep interval starting at the time when an auditory beep signaled the onset of the trial and a 500-ms reward interval starting at the time when the eccentric target was extinguished (on rewarded trials, this coincided with the onset of reward delivery).

To examine the statistical properties of DA neuron spike trains during epochs of tonic firing, we first computed the time elapsed between each spike collected during both the baseline and wait intervals for each neuron. We constructed two ISIHs for each neuron: one for each of these presumptively tonic intervals. To quantify these distributions, we also computed the coefficient of variation (CV) during both intervals, which was the ratio of the SD to the mean of the ISI distribution. The CV provides a single-parameter estimate of the variability of the neuronal spike train. A CV near 1 indicates Poisson-like variability. A CV near 0.35 indicates highly regular spike trains characteristic of previously described pacemaking systems.

We also characterized the underlying distribution from which the ISIs appeared to have been drawn. To do this, we fit each empirical ISI distribution with two different models: *1*) a two-parameter Gamma distribution $[f(x|\alpha,\beta) = x^{\alpha-1}e^{-x/\beta}/(\beta^{\alpha}\Gamma(\alpha))]$, and *2*) a two-parameter Gaussian distribution $[f(x|\mu,\sigma) = e^{-(x-\mu)^2/(2\sigma^2)}/(\sigma\sqrt{2\pi})]$. The Gamma distribution is commonly used to model ISI distributions (Brown et al. 2003), and it includes the exponential distribution, which characterizes a stationary Poisson process (a classic early model of the spike generator), as a special case ($\alpha = 1$). Unlike the exponential distribution, the Gamma and the Gaussian distributions allow for varying periods of inactivity in the neuron immediately after a spike. The parameters for each model were fit using the method of maximum likelihood (gamfit and normfit in Matlab), which yielded parameter estimates and the log-likelihood evaluated at the model parameters.

We used Akaike's information criterion (AIC; AIC = 2 × log-likelihood − 2*k*, where *k* is the number of parameters) to compare model fits (Brown et al. 2003; Burnham and Anderson 1998). We also used the variance accounted for (VAF) by the models as an additional, more intuitive, measure to compare the model fits [(Total variance in the data − Residual variance not accounted for by the function)/Total variance in the data].

We quantified regularities in the firing patterns of DA neurons using autocorrelation functions. Previous studies have identified a series of regular multiple peaks in the autocorrelation functions for dopaminergic neurons in other preparations (Hyland et al. 2002; Paladini and Tepper 1999; Shepard and German 1988; Tepper et al. 1995), evidence for clear periodicity in those spike trains. For each cell in our database, we therefore computed an autocorrelation function and performed the following two analyses on those functions. First, we averaged together the autocorrelation functions from individual trials to determine whether we could observe consistent changes in the likelihood that an action potential would occur immediately after an action potential had been generated (for each neuron). This revealed, as we expected, a "quiet period" after action potentials; a time during which spikes were unlikely to occur. To quantify this quiet period, we measured the time it took for each cell to return, after action potential generation, halfway to the maximum probability of spike generation (observed for that neuron). To do this, we smoothed the autocorrelation functions by averaging them with a 25-ms sliding window that yielded a unique "half-maximum" time for each function, an approach modeled after a similar measure used in a previous report on DA neurons (Wilson et al. 1977). Second, we used the best-fitting model for the ISI distribution of each neuron to generate a predicted autocorrelation function under the assumption that the spikes were generated as a renewal process; that is, ISIs were sampled independently and identically from the best-fitting ISI distribution. For a stationary renewal process, the autocorrelation can be calculated directly from the ISI distribution (Cox and Lewis 1966) for comparison with the observed ISI distribution for that particular cell.

To better understand the properties of dopaminergic spike generation, we also assessed the goodness-of-fit of the Gamma and Gaussian models using a test developed by Brown et al. (2001). This allowed us to quantitatively determine whether the tonic activity from DA neurons could be characterized as a renewal process with ISIs that were either Gamma or Gaussian distributed. Briefly, this approach to analyzing spike trains begins by noting that a critical problem with traditional approaches (like the VAF described above) is that those approaches are based on the notion that the underlying variables are continuous measures—which is not the case for spike trains. To engage this problem, Brown et al. turned to the time-rescaling theorem, which states that any point process, such as a Poisson process or a renewal process with Gamma or Gaussian distributed ISIs, can be transformed through its conditional intensity function into a realization of a Poisson process with unit rate, meaning that the ISI distribution of this rescaled process is exponentially distributed.

Second, and the critical step for our purposes, is that the theorem allowed us to assess the goodness-of-fit of the Gamma and Gaussian models for DA neurons. For each neuron, we used the best-fitting Gamma (including the special case of an exponential density for a Poisson process) and Gaussian models of the ISI distribution to estimate a conditional intensity function assuming that the observed spike times were generated from a fixed (stationary) renewal process. We used the time-rescaling theorem to transform the observed spike times through the estimated intensity functions of the Gamma and Gaussian models. The ISI distribution of these rescaled spike times will be exponentially distributed with unit rate if the model fits the data. The critical step is to ask how closely each of the estimated intensity functions comes to achieving this exponential distribution. If any of the proposed models can achieve that goal, the spike train under study can be well described as a renewal process of that type. We can assess whether the Gamma or Gaussian models achieved this by comparing the ISI distribution of the rescaled spike times with an exponential distribution with unit rate. In practice, we did this by first transforming the rescaled ISIs such that, if they were exponentially distributed, they would be now uniformly distributed over the range from 0 to 1. This means that we can measure the goodness-of-fit by comparing the transformed model of the observed data to a uniform distribution. We did this graphically by plotting the sorted, transformed ISIs against the theoretically defined cumulative distribution function of a uniform density. If a model is correct, the points from this plot will lie on the main diagonal of the graph, and deviations of this line from the main diagonal indicate deviations of the observed spike train from the best-fitting model of that type. Brown et al. referred to these graphs as Kolmogorov-Smirnov (KS)-plots. We summarized these KS-plots using the KS statistic (Press et al. 1992), which measures the difference between two distributions and ranges from 0 to 1: 0 indicating a perfect fit of the model to the observed point process, and 1 indicating no fit.

Once we had characterized the patterns of activity observed during the baseline period in each of these ways, we could ask whether significant deviations from this pattern occurred at other times during each trial, indicating the onset of what we defined as a phasic modulation. We searched for phasic responses in each cell using the distributional model that accounted for the most variance in that cell, (in practice either a Gamma or a Gaussian distribution). To do this, we set thresholds for ISIs that represented the upper and lower 95% CIs of this distribution. We defined the onset of a burst as the occurrence of two successive action potentials that were separated by an ISI shorter than the lower threshold (mean of 33.3 ms and an SD of 27.3

ms across our population).[1] We defined the offset of a burst as the last spike that was preceded by an ISI shorter than the threshold. We also used this method to identify pauses in the tonic activity of dopamine neurons. Pauses were defined as two sequential action potentials separated by an interval that was longer than the upper threshold (mean of 369.0 ms and an SD of 103.3 across our population). Once we had identified bursts and pauses during the beep and reward intervals that by definition included behaviorally relevant events, we could correlate these phasic modulations with regard to the relevant behavioral events.

Previously published results indicated that the average firing rate of DA neurons after the delivery of a reward may reflect the difference between the magnitude of the reward the animal has just received and a weighted average of the magnitudes of the preceding rewards (Bayer and Glimcher 2005; Schultz et al. 1997). More formally, this suggests that the phasic modulations of DA neurons can be predicted from an equation having the form

Average Firing Rate

$$= R_{(\text{most recent})} - [(R_{-1} \times w_{-1}) + (R_{-2} \times w_{-2}) + \ldots + (R_{-n} \times w_{-n})]$$

where $R_{(\text{most recent})}$ is the value of the most recently received reward, R_{-1} through R_{-n} are the rewards on a set of previous trials, and w_{-1} through w_{-n} are the weights used to average these previous rewards. Most theoretical models that have been used to account for the phasic modulation (Rescorla and Wagner 1972; Schultz et al. 1997; Sutton and Barto 1981) go a step further and suggest that the magnitude of the weights should decline exponentially if these neurons participate in reinforcement learning.

As in our previous study (Bayer and Glimcher 2005), we computed the empirical reward prediction error function for the neurons, which is the empirical function that predicts average DA firing rate during the postreward interval from the history of reward magnitudes over the past 10 trials, using a linear regression on recent rewards to predict firing rate during the first 500 ms of the reward interval.[2] The linear regression thus provided a set of weights taking the following form

$$\beta_0 \times (R_t) + \beta_1 \times (R_{t-1}) + \beta_2 \times (R_{t-2}) + \ldots + \beta_{10} \times (R_{t-10})$$

Where R_t is the amount of fluid reward provided on the current trial and takes a positive value, R_{t-1} is the amount of reward obtained on the previous trial taking a negative value, and so on with all remaining coefficients (the β-weights) taking a negative value. Note that the regression does not require that these terms take values having these particular signs, but if the coefficients construct a reward prediction error they must do so. In practice, the negative sum of $\beta_1 \times (R_{t-1}) + \beta_2 \times (R_{t-2}) + \ldots + \beta_{10} \times (R_{t-10})$ is found to be equal to $\beta_0 \times (R_t)$ for the neurons we have studied. The regression thus yielded a set of β values defining the best linear rule for predicting the firing rate of the DA neurons from the recent history of rewards. We have previously shown that the weighting function derived in this way almost perfectly approximates the exponentially weighted average of the theoretically defined reward prediction error and does so without making any other assumptions than linearity (for more details on this approach to the reward prediction error, see Bayer and Glimcher 2005).

We constructed ISIHs during the beep and reward intervals that were segregated by reward prediction error to confirm that there were differences in the general distribution of ISIs resulting from the differences in average firing rate observed previously. Finally, we computed three different characteristics of phasic activity on all identified pauses and bursts: the number of spikes per burst, burst latency, and pause duration. This allowed us to examine the relationship between these variables and reward prediction error.

Population level analysis

To compare the characteristics of phasic activity across the entire population of neurons, we normalized the results for each cell according to the range of responses observed in that cell. Relative burst size was computed by dividing the actual burst size by the mean burst size observed for all bursts recorded from the same neuron. Relative pause duration and relative burst latency were computed the same way; the temporal interval observed on each trial was divided by the mean pause duration or burst latency (respectively) observed for that cell. This allowed us to determine whether there were changes in the characteristics of phasic activity across the entire population of neurons.

RESULTS

Tonic activity

To ensure that the baseline and wait intervals did, in fact, represent periods when there were no significant task-related phasic modulations in neuronal activity, we constructed perievent time histograms for each cell. For each neuron, four histograms were generated: two during the baseline period, one aligned to the time at which the reward was delivered and the other to the end of the trial, and two during the wait period, one aligned to the illumination of the eccentric target and the other to the offset of the target. Figure 2 plots these histograms for a single typical neuron. This pattern of results was observed in all of our neurons, suggesting that the average firing rate during these intervals did not include phasic modulations of the neurons triggered by afferent input linked to task events. This was a prerequisite for the following analysis of spiking that presumes the activity during these periods was largely stationary in nature.

Figure 3, *A* and *B*, shows the distributions of ISIs we observed for a single neuron during the baseline and wait periods, respectively. Both distributions have a central peak at ~144 ms and drop off at the same rate for larger and smaller ISIs. Note that the distribution includes many ISIs <80 ms, the classical threshold beneath which a pair of spikes can be considered for categorization as a burst in the rodent (Grace and Bunney 1984). The distribution of ISIs during the baseline and wait intervals for a second cell are shown in Fig. 3, *C* and *D*. To characterize the variability of ISIs observed across our population, Fig. 3, *E* and *F*, plots the observed CVs for each neuron during the baseline and wait intervals, respectively. To further quantify the distribution of ISIs observed during baseline firing in our neurons, we fit the baseline ISIH of each neuron with both a Gaussian density and a Gamma density.[3] Figure 4, *A* and *B*, shows the ISIHs of the example cells from Fig. 3, *A* and *B*, with both the best fit Gaussian function and the best fit Gamma function. The distribution of ISIs for cell 1 was slightly better described by a Gaussian function, although the difference in terms of variance accounted for was only ~3%. Figure 4*B* shows the ISIH and best-fitting functions based on the data from the second example cell from the previous

---

[1] Previous studies in the rat (Grace and Bunney 1984) have often defined a burst as two or more spikes having an ISI of <80 ms and in which the amplitude of subsequent spikes declines. Our goal here was not to precisely replicate that measure but rather to identify phasic modulations that reflected a transient and statistically significant deviation from the baseline tonic activity of each neuron. In practice, the bursts we describe here would all have been identified using this more classical measure, although many more bursts would have been defined with the classical 80-ms measure derived from study of the rodent than with our approach in these primate data.

[2] We selected 500 ms for consistency with our previous work and because a 500-ms period encompasses the longest pauses we observed. In practice, the exponentially declining beta weights we observed are fairly robust to measured interval duration as indicated by the analysis of longer reward intervals presented in the results section.

[3] An analysis of the wait interval, not shown here, yielded a similar result.

Baseline - aligned to reward
(starting 500 ms after reward onset)

Baseline - aligned to end of trial

Wait - aligned to target onset

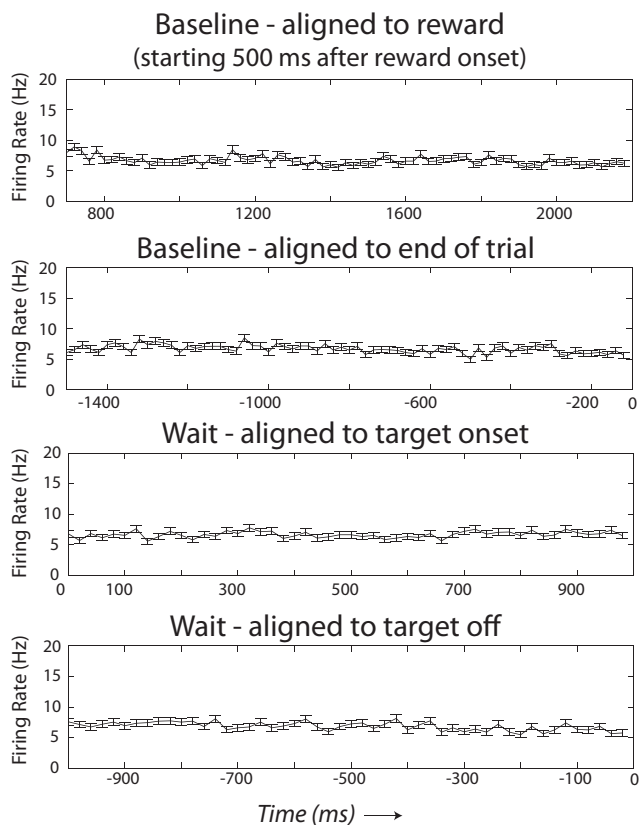## Wait - aligned to target off

Time (ms) ⟶

FIG. 2.   Perievent time histograms aligned to start and end of baseline and wait intervals for a single neuron. Mean and SE are plotted for each 20-ms bin; 1,024 trials were used to compute these histograms. Flat histograms indicate that no task-related phasic activity occurred during these intervals.

figure. This cell was better described by a Gamma function, although again, the difference in variance accounted for was quite small, only ~2%.

To examine whether the Gamma function accounted for significantly more variance than the Gaussian function for every cell in our population, we computed the difference in variance accounted for by the two fits as a function of the total variance accounted for by both fits and used AIC to determine which model better described the data. Figure 4C plots a histogram of the difference between the AIC for the Gamma distribution minus the AIC for the Gaussian distribution. Note that most cells were better described by the Gamma distribution. To further examine this difference, Fig. 4D plots the variance accounted for by each model, for each cell. Gray points identify neurons better fit by the Gamma distribution according the AIC used in Fig. 4C. Note that these two measures largely agree, but much more important is the observation that both models fit all cells quite well. In summary, these data suggest that, in the awake-behaving monkey, more cells were fit better by a Gamma function than by a Gaussian function, but this difference was very small.

Figure 5, A and B, shows the distribution of parameters for the Gaussian models that best fit the data. The average mean was 175 ± 34 ms, and the average SD was surprisingly broad at 107 ± 36 ms. Figure 5, C and D, shows the parameters of the Gamma functions that best fit the neuronal data.

We also examined the patterns of sequential spike generation in our neurons by computing autocorrelation functions during the

baseline period. Figure 6, A and B, shows the autocorrelation functions for two example cells in gray. These functions show a very low probability of spike generation immediately after a spike, which gradually increased until it reached a maximum probability of spike generation about 100–140 ms later. The black lines show analytically derived autocorrelation functions for the best-fitting Gamma and Gaussian functions for these two neurons. Note that, as in Fig. 4, cell 1 appears better described by a renewal process with Gaussian distributed ISIs and cell 2 appears better described by a renewal process with Gamma distributed ISIs. Note that neither of these autocorrelation functions shows the oscillatory pattern of correlations that are evidence of an underlying periodic process as has been observed in some rodent dopamine neurons (see Fig. 3C of Hyland et al. 2002 for an example of this). We saw no evidence in our population of regular firing cells with strong multiple peaks in the autocorrelation functions. Figure 6C plots the time to half-maximum for each neuron in our population. The time to half-maximum is the interval required after a spike for the autocorrelation to rise to one half its maximum value. This is the interval required for the probability of spike generation to rise to 50% of maximal probability. The average across the population is 106 ms. Figure 6D plots, for the best-fitting model of each cell (Gamma or Gaussian, by AIC), the correlation between the observed and analytically derived autocorrelation functions. Again note that the models do
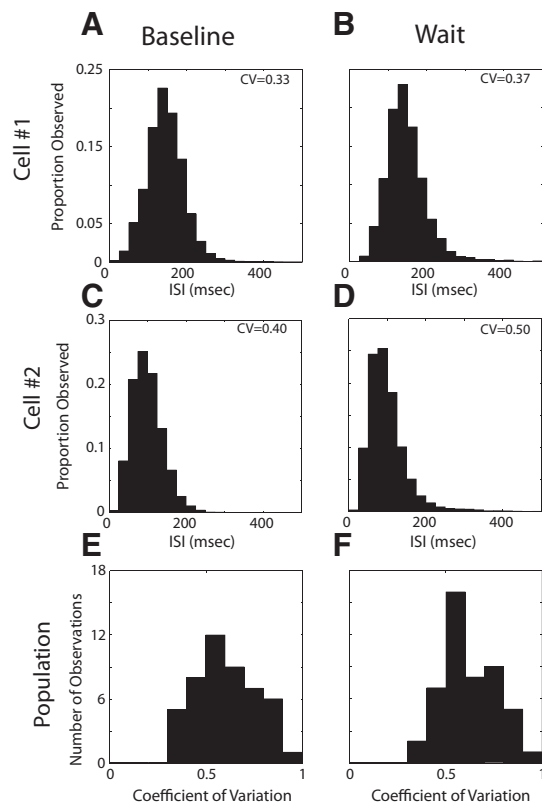

FIG. 3.   Interspike interval (ISI) histograms for the baseline and wait intervals for 2 individual neurons and the distribution of CVs for population. A: ISI histogram for example neuron 1 during baseline interval. B: ISI histogram for example neuron 1 during wait interval. C: ISI histogram for example neuron 2 during baseline interval. D: ISI histogram for example neuron 2 during wait interval. E: distribution of CV for all cells during baseline interval (mean = 0.61, SD = 0.16). F: distribution of CV for all cells during wait interval (mean = 0.62, SD = 0.15).
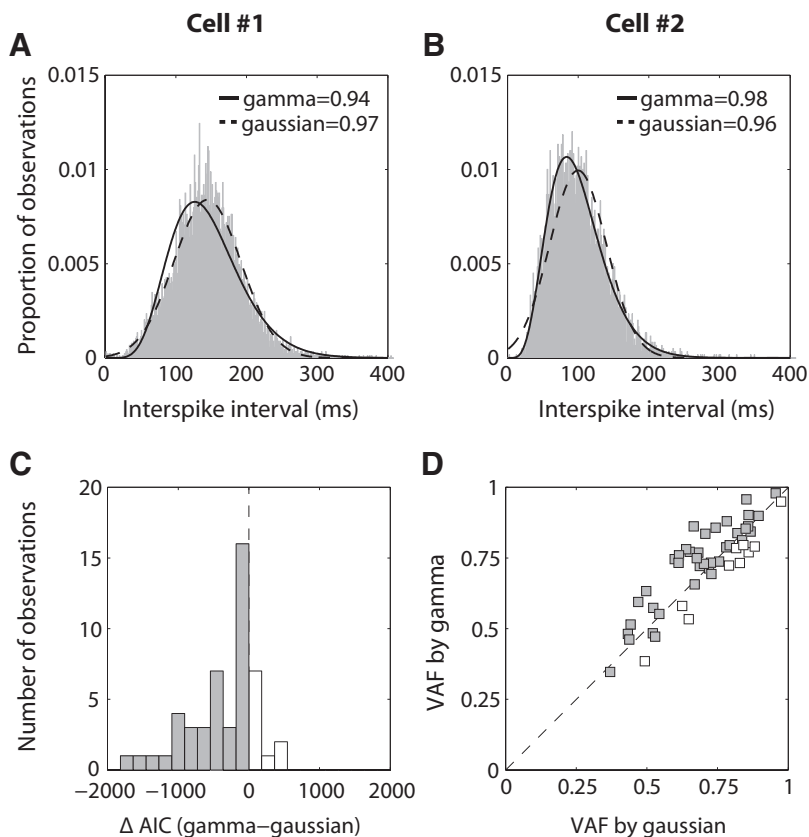
FIG. 4. Fitting a model to the ISI histogram (ISIH). *A*: ISIH for example neuron 1 (from Fig. 3) plotted in gray, with best-fitting Gaussian and Gamma distributions plotted on top in dashed and solid lines, respectively. The numbers give the VAF for each model. *B*: ISIH for example neuron 2 plotted in gray, with best-fitting Gaussian and Gamma functions plotted on top in dashed and solid lines, respectively. *C*: histogram of differences in Akaike's information criterion (AIC) scores across population. *D*: variance-accounted-for by each model. Neurons for which the Gamma model was a better fit (smaller AIC score) are shaded gray. Gaussian and Gamma models both fit the data reasonably well.

a reasonable job of accounting for the spiking behavior of the neurons.

To further assess the accuracy of the Gaussian and Gamma models for these neurons, we plotted KS plots of our neurons using the methods of Brown et al. (2001; see METHODS). Figure 7, *A* and *B*, shows these plots for two example cells. We used the KS plots to measure the accuracy with which a renewal process with exponential, Gamma, or Gaussian distributed ISIs
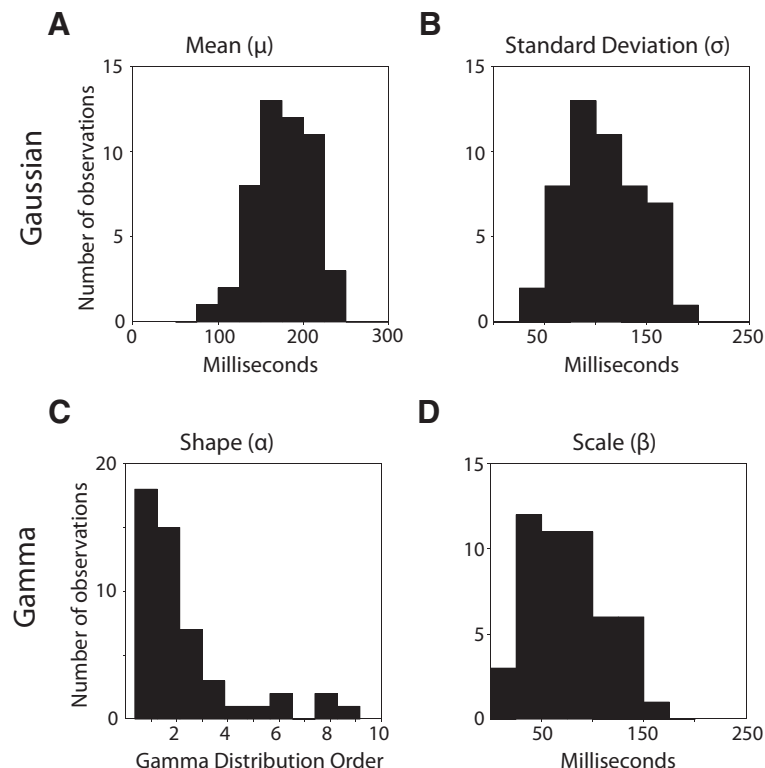

FIG. 5. Distributions of parameters for best-fitting models of activity during the baseline period. *A*: histogram of means from best-fitting Gaussian functions (mean = 175, SD = 34). *B*: histogram of SD from best-fitting Gaussian functions (mean = 107, SD = 36). *C*: histogram of α parameters from best-fitting Gamma functions (mean = 3.0, SD = 1.8). *D*: histogram of β parameters from best-fitting Gamma functions (mean = 75, SD = 38).
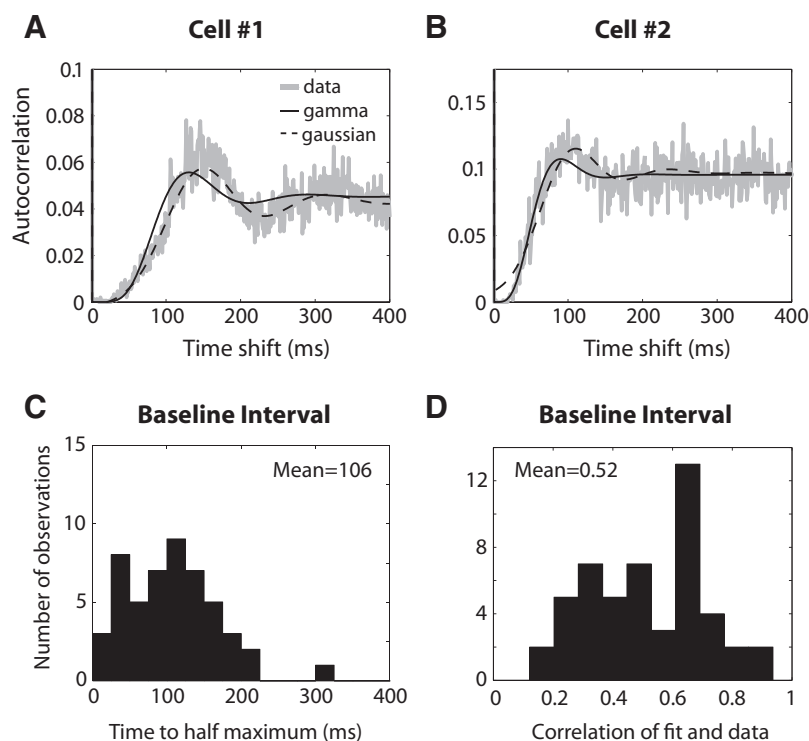
FIG. 6. Autocorrelation analysis. *A*: gray line is autocorrelation function for example neuron 1 (from Fig. 3) estimated from data during baseline interval (1-ms binning). Dashed and solid black lines are the autocorrelation functions predicted by a renewal process with either a Gaussian or Gamma ISI distribution obtained by fitting these distributions to observed ISIs. For this neuron, the simple renewal model with a Gaussian ISI distribution produces an autocorrelation function very similar to that observed in the actual neuron. *B*: autocorrelation function for example neuron 2 during baseline interval. *C*: histogram of time to half-maximum autocorrelation values for all neurons during baseline period (mean = 106, SD = 59). *D*: for each neuron, we compared observed autocorrelation with the autocorrelation function predicted by a renewal process with best-fitting ISI distribution for each neuron. These comparisons were summarized using a correlation coefficient obtained from time shifts 1–400, excluding *time 0*. This figure plots the distribution of correlation coefficients we observed for all neurons in our population during the baseline interval.

characterized the spike trains of individual neurons. The closeness with which the line for a given model approximates the main diagonal is a measure of the goodness of fit for that model. Note that as in our previous analyses, cell 1 is better described as Gaussian and cell 2 as Gamma. Figure 7*C* quantifies this across the population using a KS statistic. Values of 0 indicate a perfect fit between the model and the data. Note that across our population, we once again observed a better fit of the Gamma functions, but Gaussian functions also do a good job of describing the data for some neurons.

Phasic activity

The activity we observed during the baseline and wait intervals thus appeared to be reasonably well described by simple stochastic models in which all spikes were independent events and were not associated with afferent activity triggered by task events. These are properties we take to be characteristic of the tonic activity studied by other researchers in rodent DA neurons. To search for phasic events during our reward and beep intervals, we therefore used a simple statistical criterion. Any consecutive ISIs that lay below the lower 95% CIs of the modeled distribution for that neuron were labeled bursts, and any interval that lay above the 95% CI was labeled a pause.

Figure 8*A* shows the results of our burst/pause detection algorithm for a single neuron. Plotted in black are the times at which individual spikes occurred with respect to the time at which a task-associated reward was delivered, and plotted on top of those in gray are the intervals that were identified as bursts (plotted in thick gray bars) or pauses (plotted as thin black bars). Trials are sorted by the magnitude of the reward prediction error, as measured on that trial, based on an analysis of the firing rate and the reward history associated with that particular neuronal recording session. At very negative reward prediction errors, the neuron clearly paused on most of the trials we recorded, whereas at very

positive reward prediction errors, the cell exhibited a burst of action potentials. Stars mark trials in the upper half of the figure where the neuron paused despite a positive reward prediction error. Stars mark trials in the lower half of the plot where bursts occurred under conditions of a negative reward prediction error. These data indicate that this neuron shows two forms of phasic modulation correlated with the empirically defined reward prediction error: both bursts and pauses.

Figure 8, *B–F*, shows the distribution of ISIs observed during the bursts that followed the trial initiating beep and the reward. After the beep, there was a shift in the ISI toward the left, with an average ISI of 48 ms and a CV of 0.26. During the postreward bursts, there was also a leftward shift in the ISI distribution, which was correlated with the reward prediction error. The mean ISI duration for bursts associated with large positive reward prediction errors was only 28 ms, whereas the mean ISI duration associated with smaller reward prediction errors was 49 ms (comparable with the responses after the beep). Unexpectedly, the responses associated with slightly negative reward prediction errors (Fig. 8*E*) showed a biphasic distribution of ISIs, as the neuron sometimes responded with a very short burst (mean of short ISIs = 48 ms) and sometimes with a short pause (mean of long ISIs = 284 ms).[4] For very negative reward prediction errors, however, the neuron always responded with a pause in activity (mean ISI = 339 ms).

---

[4] Before drawing any conclusions from this bimodal distribution about the underlying neuronal representation of predicted reward magnitude, it should be bourne in mind that this bimodal distribution may simply reflect a limitation of our algorithm for estimating the reward prediction error encoded by firing rate of the dopamine neurons. If positive and negative reward prediction errors are encoded differently (nonlinearly) by the dopamine neurons, a misidentification of the zero-point in the reward prediction scale might result. In fact, one might expect precisely such a misidentification from our linear regression if the true encoding of reward prediction error by dopamine neurons shows a nonlinearity that favors positive reward prediction errors.
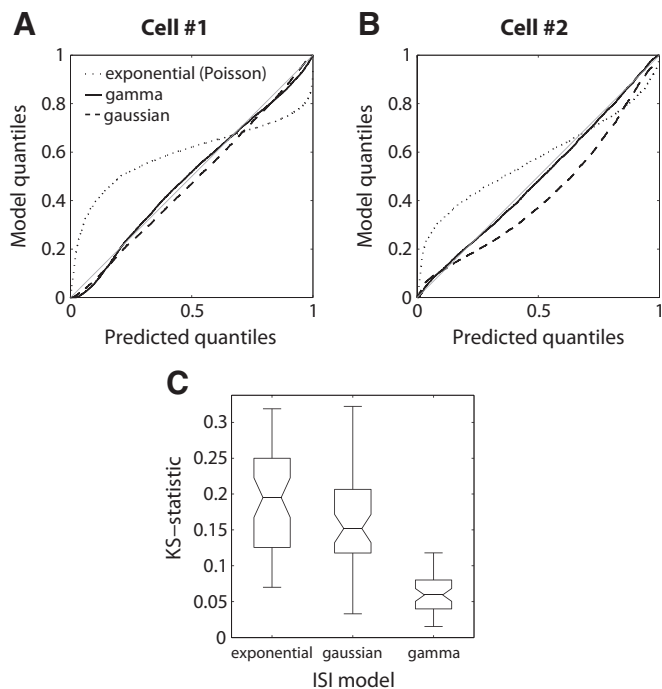
FIG. 7. Goodness-of-fit. We assessed how well different models fit data using the method proposed by Brown et al. (2001). *A*: model quantiles plotted against predicted quantiles for example cell 1 (from Fig. 3) from data during baseline interval. Thin gray line is the unity line; any empirical comparison that runs along this line indicates a perfect fit of that model to data. Dotted, dashed, and solid black lines represent model quantiles predicted by stationary renewal models using Exponential, Gaussian, and Gamma ISI distributions, respectively. Kolmogorov-Smirnov (KS)-statistics are Exponential = 0.303, Gaussian = 0.033, and Gamma = 0.041. *B*: model quantiles plotted against predicted quantiles for example cell 2. KS-statistics are Exponential = 0.224, Gaussian = 0.140, and Gamma = 0.022. *C*: boxplots of the K-S statistic for the population. KS-statistic encapsulates how well each model fits data by essentially reporting maximum vertical distance of empirical curves in *A* from main diagonal. KS-statistic of 0 indicates a perfect fit. Horizontal notches indicate medians for each model, extent of boxes represents upper and lower quartiles of distributions, and whiskers give range of statistic.

To better quantify the relationship between both types of phasic neuronal activity and the reward prediction error, Fig. 9*A* plots the number of spikes in the burst as a function of reward prediction error for our two example cells. For both cells, there is a relatively linear relationship between the number of spikes in the burst and reward prediction error over the limited range of reward prediction error for which bursts occurred. We plotted the duration of the pause during trials when the received reward was less than the expected reward (Fig. 9*B*). We found that the pauses in activity were longer for more negative reward prediction errors, with a range of ∼100 ms and this relationship was largely linear over the limited range at which bursts did not occur. We also plotted the latency to burst onset as a function of reward prediction error for both of these individual neurons (Fig. 9*C*) and found that larger reward prediction errors were generally associated with earlier burst onsets, but only with about a 30-ms difference between the earliest and the latest burst and with a function that has a step-like quality for these neurons.

Figure 10 shows the same measurements plotted for the entire population. Figure 10*A* plots the relationship between relative burst size and reward prediction error. As observed in the example neurons, the increases in reward prediction error

are generally correlated with increases in burst size (Spearman rank correlation, $r_s = 0.30$; $P < 0.001$). Figure 10*B* indicates that pause duration was also associated with changes in reward prediction error. More negative reward prediction errors elicited longer pauses ($r_s = -0.21$; $P < 0.001$). Finally, the most positive reward prediction errors were not associated, at the population level, with a decrease in the delay before the burst ($P > 0.10$; Fig. 10*C*). Taken together, these results suggest that
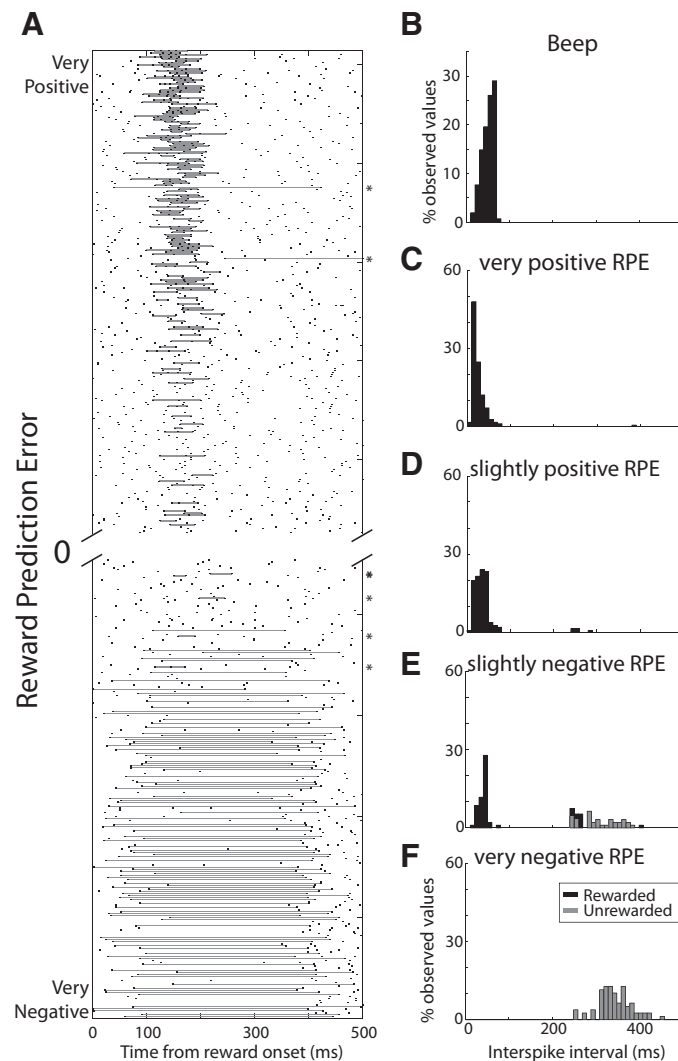

FIG. 8. Phasic activity. *A*: rasters for individual trials from example neuron 1 (from Fig. 3) aligned to time of reward delivery, sorted from most negative reward prediction error to most positive reward prediction error using best-fit reward predictor. Black tick marks indicate timing of individual action potentials. Horizontal bars mark periods of phasic activity detected by our statistical analysis. Bursts are highlighted using thicker gray bars, and pauses are highlighted using thinner black bars. Majority of bursts plotted occur for positive prediction errors (top half), whereas majority of pauses occurred for negative prediction errors (bottom half). Rare pauses or bursts that occurred for positive or negative prediction errors, respectively, are marked with an asterisk at the *right* of plot. *B*: ISI histograms for intervals identified as phasic activity after trial-initiating beep. According to most models, these intervals should be associated with positive reward prediction errors. *C–F*: ISI histograms for phasic modulations during reward intervals sorted by magnitude of measured reward prediction errors. Very positive reward prediction errors (RPEs) correspond to trials with RPEs > 0.1 (73 trials). Slightly postive RPEs correspond to trials with RPEs ≥ 0 and <0.1 (359 trials). Slightly negative RPEs correspond to trials with RPEs ≥ 0.1 and <0 (444 trials). Very negative RPEs correspond to trials with RPEs < −0.1 (131 trials).
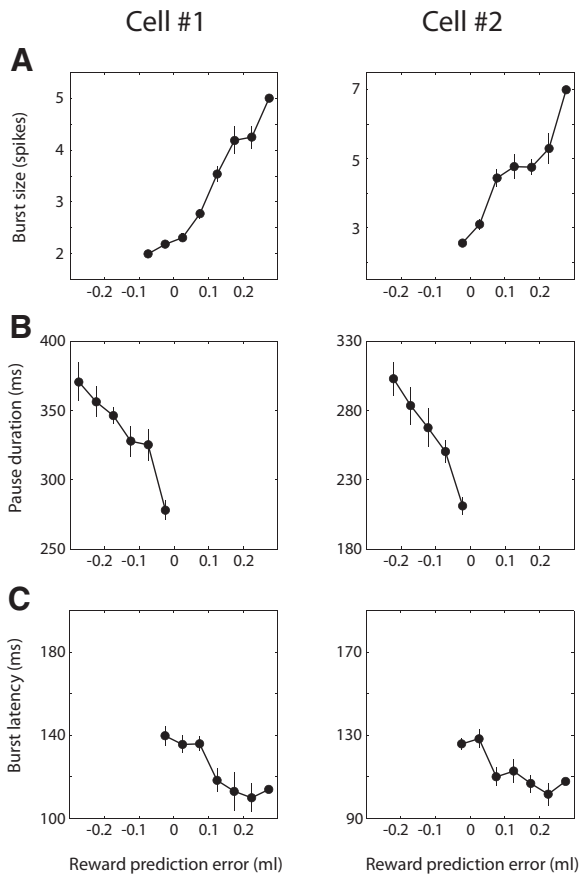
FIG. 9. Burst characteristics plotted as a function of reward prediction error for two example neurons, on the *left* for example neuron 1 and on the *right* for example neuron 2 from Fig. 3. *A*: mean burst size (number of spikes per burst) plotted as function of reward prediction error with SE bars (*left*, example neuron 1; *right*, example neuron 2). *B*: mean pause duration plotted as a function of reward prediction error with SE bars (*left*, neuron 1; *right*, neuron 2). *C*: mean burst latency plotted as a function of reward prediction with SE bars (*left*, neuron 1; *right*, neuron 2).

both classes of phasic activity, bursts and pauses, are correlated with the history of rewards encountered by the animal, a signal presumably carried by afferent inputs to these neurons.

To further explore the relationship between bursts and pauses, we plotted average firing rate (which would reflect both increases and decreases in activity) as a function of the reward prediction error during postreward intervals of different length (Fig. 11). If we set the interval to the length of the longest burst we observed, this function should approximate the steepest possible slope for positive reward prediction errors. If the interval was set to the longest pause we observed, the function should approximate the steepest possible slope for negative reward prediction errors. To perform this analysis, we therefore selected postreward intervals of 150, 200, 400, and 600 ms and empirically derived reward prediction error functions for each duration (as in Bayer and Glimcher 2005). Figure 11*A* plots the relationship between the reward prediction error and firing rate for each of these interval durations. Figure 11*B* plots the weighting function derived from an analysis of that interval. Note that the relationship between firing rate and the reward prediction error remains nonlinear regardless of the duration over which spikes are averaged. This suggests a discontinuity between bursts and pauses with regard to the linear encoding of

previous rewards by spike rate during any fixed postreward interval, the standard measure used in most models of the reward prediction error term.

Finally, we also hoped to determine whether the characteristics of the phasic modulations observed during the reward interval were related to the tonic firing rates of the cells before the delivery of the reward on that trial. This serves as a test to examine the degree to which the underlying state of the neuron influenced ongoing spike rate. Figure 12 plots the mean and SD of firing rate during the baseline interval of each trial (for all neurons) as a function of mean firing rate during the reward interval on that same trial. Note that there is a significant, but
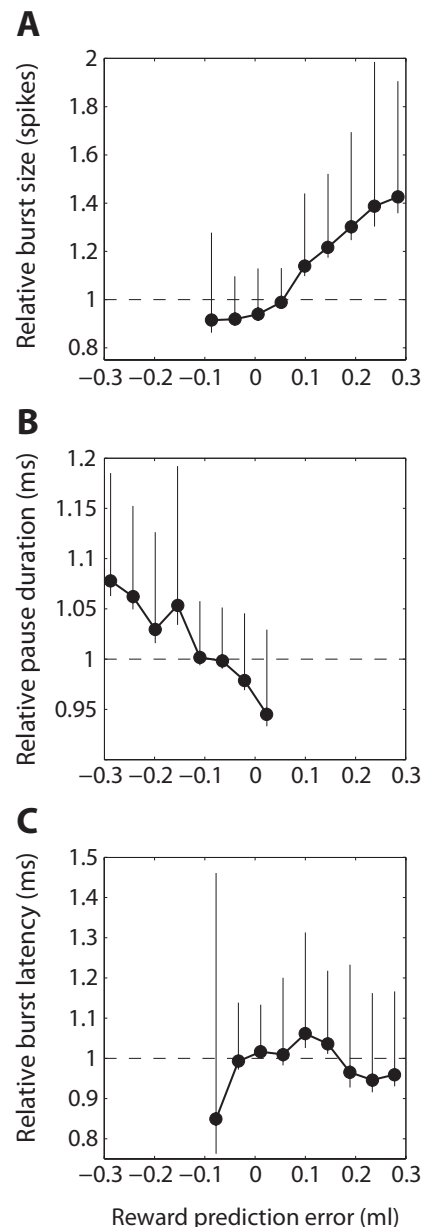

FIG. 10. Phasic response characteristics plotted as a function of reward prediction error for population. *A*: relative burst size plotted as function of reward prediction error. Top error bars indicate SD across population; bottom error bars identify SE estimated from population. *B*: relative pause duration plotted as a function of reward prediction error. Top error bars, SD; bottom error bars, SE. *C*: relative burst latency plotted as a function of reward prediction error. Top error bars, SD; bottom error bars, SE.

A   Firing Rate vs. Reward Prediction Error   ## B   Linear Weighting Function
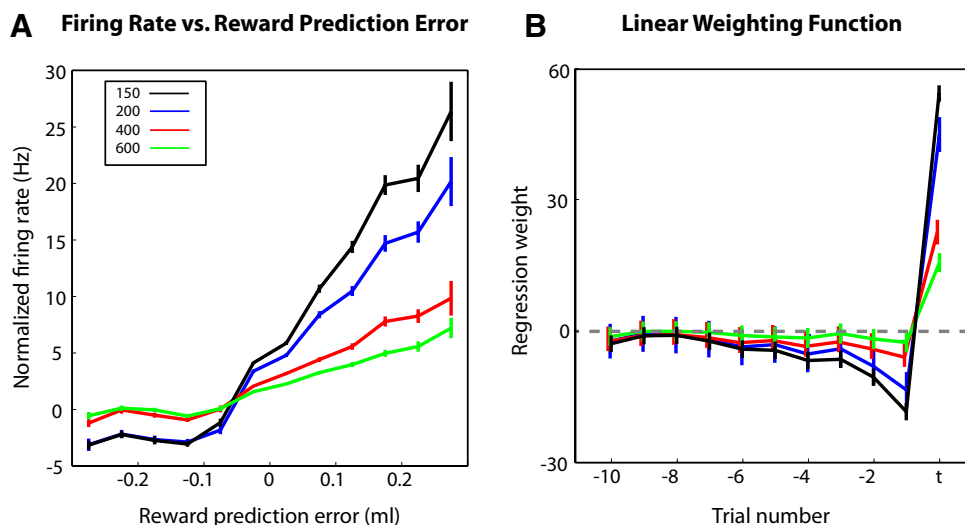

FIG. 11.   *A*: average firing rate as a function of reward prediction error during postreward intervals of different lengths. *B*: best-fit weighting functions from which reward prediction error used on the lower axis of *A* was derived. See Bayer and Glimcher (2005) for more details on this function.

very small, relationship between these variables. These data suggest that the prior tonic state of the neurons, their ongoing level of excitability, contributes weakly, but significantly, to the phasic responses of these cells in this task.

DISCUSSION

We studied midbrain DA neurons during two intervals that might be expected to contain tonic activity: one occurring between trials of an experimenter-controlled task and the other occurring while the animal was waiting to make an eye movement for a juice reward. We examined the statistics of sequential ISIs during these epochs by computing the CV. We also used autocorrelation analyses to measure the level of rhythmicity in the sequential process of spike generation. For all of our neurons, we found that
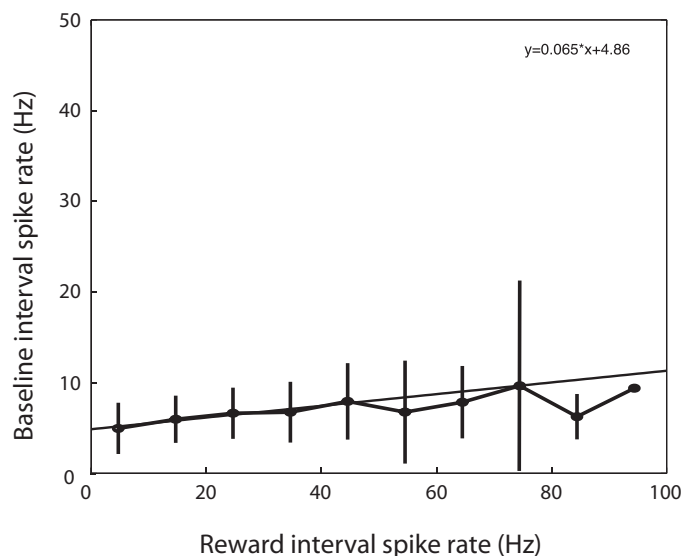

FIG. 12.   Relationship between phasic modulation and spike rate during the immediately preceding baseline interval for all trials and all neurons. Mean and SD of firing rates are plotted, on a trial-by-trial basis, during sequential baseline and reward (150-ms duration interval from Fig. 11) intervals. $R^2 = 0.05$, $P < 0.0001$, $n = 15641$. Best-fitting line: $y = 0.065x + 4.86$.

during both epochs, the CV was significantly <1 (averaging ~0.6) and that, although the autocorrelograms showed a reduced likelihood of firing during an extended interval (on the order of 100 ms) after the occurrence of each action potential, there was no significant evidence for nonindependence of ISIs during these epochs in any of the analyses that we performed. Our results are compatible with the conclusion that activity during these intervals, which have been previously studied in the monkey (Bayer and Glimcher 2005; Hollerman and Schultz 1998; Ljungberg et al. 1992; Morris et al. 2004; Waelti et al. 2001), largely represents the tonic activity that has been observed in the rodent system. It should, however, be noted that we did not observe the highly rhythmic pacemaker mode that has occasionally been observed in that system (Grace and Bunney 1984; Hyland et al. 2002; Silva and Bunney 1988).

We were able to use the distributional analyses of this presumptively tonic activity to define phasic activity as any group of spikes separated by an ISI that was significantly shorter or longer than average, using a 95% CI as we observed it in this system. We found that DA neurons both bursted and paused (with ISIs of ≤350 ms) in a manner correlated with the subject's recent reward history. Furthermore, the firing rates we defined as phasic occurred along a continuum that ranged from long pauses to short pauses to short bursts to long bursts. The neurons thus appeared to deviate from tonic levels of activity in a continuous fashion either by increasing or decreasing spike generation rates as has been previously suggested (Waelti et al. 2001).

The relationship between bursts and pauses was, however, nonlinear with regard to average spike rates during a fixed postreward interval and the reward prediction error. Average firing rates during fixed intervals short enough to capture bursts could not be described as linear with regard to both positive and negative reward prediction errors simultaneously, although the encoding of reward prediction errors by bursts or pauses individually could be described as largely linear. Thus the onset, duration, and magnitude of the phasic modulation of dopamine neurons are all correlated with reward prediction error, a result that expands significantly on previous findings (Nakahara et al. 2004; Satoh et al. 2003; Schultz et al. 1997; Waelti et al. 2001).

Reward prediction errors and phasic activity

Central to our understanding of the function of midbrain DA is that these neurons appear to operate in two modes: a burst, or phasic, mode and a tonic mode. This distinction arises from a broad array of studies that have identified pharmacological, biophysical, and anatomical distinctions between these two modes and which posit different roles for these two modes in the control of behavior. Previous studies in the rodent, for example, have associated bursts with D1 receptor family activation (Goto and Grace 2005), the stimulation of NMDA receptors (Johnson et al. 1992), postsynaptic effects on synaptic strength (Goto and Grace 2005; Lisman and Grace 2005), and modulation of hippocampal activity (Goto and Grace 2005). In contrast, tonic activity in these neurons is believed to reflect D2 receptor activation, activity in the prefrontal cortex, and a presynaptic mechanism for the modification of synaptic strength (Lisman and Grace 2005). These data and others have led to the proposal of a phasic/tonic model for dopaminergic activity (Grace 1991). In this model, prefrontal mechanisms regulate the baseline of tonic activity in these neurons which controls, through homeostatic mechanisms, the overall pharmacologic sensitivity of the dopaminergic targets. Other systems linked to areas like the hippocampus are proposed to govern the phasic activation of this system in a way that regulates learning.

In comparing the statistics of dopaminergic phasic activity in awake primates with the activity of dopaminergic neurons in rodents, we found many similarities that support the importance of these distinctions in the awake-behaving primate. However, we also found that there were some apparent differences between the phasic mode that has been previously reported in the rodent and the reward-related phasic modulations that both we and others working in the monkey have observed. We observed both pauses and bursts during phasic modulations, and both pauses and bursts were related to the recent reward history of the animal. When the most recent reward received by the animal was more than the average of recent rewards, the neurons responded with a burst of action potentials. When the recent reward was smaller than this average, the neurons paused. Our results thus support the hypothesis that phasic modulations in DA firing rate are driven by reward-sensitive afferents, but extend that hypothesis to include brief pauses in activity.

Previous results, however, have suggested that the average spike rates of DA neurons during a postreward interval of fixed length may encode a wider range of positive reward prediction errors than negative reward prediction errors (Bayer and Glimcher 2005; Satoh et al. 2003). This effect emerges largely from the fact that the baseline firing rates of the neurons are so low. These dopamine neurons can increase their firing rate by a factor of 10 or more but can only decrease their firing rate by a few hertz before reaching a rate of 0 during fixed postreward intervals of limited length. Our finding that pause duration is correlated with negative reward prediction errors suggests the possibility that negative reward prediction errors may be encoded by these neurons for lower values of the reward prediction error than had been previously suspected. In interpreting this finding, it is critical to note, however, that the reward prediction errors encoded by bursts and pauses are not linearly related by average firing rate during fixed intervals short enough to detect postreward bursts. With regard to firing rate during fixed intervals of the type usually described in the literature, positive reward prediction errors are encoded with a much steeper slope than are negative reward prediction error and the inflection in this slope seems to occur at or near the zero point at which predicted and obtained reward are identical. This is an observation that may support the suggestion of Daw et al. (2002) that DA is only one of two neural systems carrying information about reward prediction errors. On the other hand, the observation that the postreward pause duration of these neurons does encode the negative reward prediction error may mitigate against this conclusion.

Our data, however, do not seem to support the notion that these neurons serve as pacemakers in the awake-behaving primate. An analysis of the ISIHs, the CVs, and the autocorrelation functions for this population of neurons indicates that this is unlikely. Our data cannot rule out the possibility that other primate DA neurons have these properties but neither does our small sample provide evidence for the existence of pacemaking neurons in this species.

Summary

The results reported in this paper suggest a specific relationship between the phasic modulations of DA cells and reward prediction errors. Bursts in activity appear to encode positive prediction errors and pauses in activity of ≤350 ms appear to encode negative errors. Both of these classes of modulation may well be compatible with existing models of phasic activity. The tonic activity we observed, while broadly similar to that observed in the rodent, did show some differences. If, as we propose here, the baseline and wait period activity we measured constitutes tonic activity, the range of frequencies that constitute tonic activity in the monkey may be broader than that observed in the rat. Our ISIHs indicate that, under these conditions, ISIs <80 ms were extremely common and occurred stochastically throughout our tonic interval, despite the fact that ISIs <80 ms in duration have been used, in combination with other criteria, to define bursting activity in many rodent preparations. Finally, we found no evidence for a repetitive firing pacemaker mode in any of the neurons we examined.

REFERENCES

Bayer HM, Glimcher PW. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47: 1–13, 2005.
Brown EN, Barbieri R, Eden UT, Frank LM. Likelihood methods for neural data analysis. In: *Computational Neuroscience: A Comprehensive Approach*, edited by Feng J. London: CRC, 2003, p. 253–286.
Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput* 14: 325–346, 2001.
Burnham KP, Anderson DR. *Model Selection and Inference: A Practical Information-Theoretic Approach.* New York: Springer-Verlag, 1998.
Cox DR, Lewis PAW. *The Statistical Analysis of Series of Events.* London: Methuen, 1966.

Daw ND, Kakade S, Dayan P. Opponent interactions between serotonin and dopamine. *Neural Netw* 15: 603–616, 2002.

Floresco SB, West AR, Ash B, Moore H, Grace AA. Afferent modulation of dopamine firing differentially regulates tonic and phasic dopaminergic transmission. *Nat Neurosci* 6: 968–973, 2003.

Freeman AS, Meltzer LT, Bunney BS. Firing properties of substantia nigra dopamine neurons in freely moving rats. *Life Sci* 36: 1983–1994, 1985.

Glimcher PW, Ciaramitaro VM, Platt ML, Bayer HM, Brown MA, Handel A. Application of neurosonography to experimental physiology. *J Neurosci Methods* 108: 131–144, 2001.

Goto Y, Grace AA. Dopaminergic modulation of limbic and cortical drive of nucleus accumbens in goal-directed behavior. *Nat Neurosci* 8: 805–812, 2005.

Grace AA. Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: a hypothesis for the etiology of schizophrenia. *Neuroscience* 41: 1–24, 1991.

Grace AA, Bunney BS. The control of firing pattern in nigral dopamine neurons: burst firing. *J Neurosci* 4: 2877–2890, 1984.

Handel A, Glimcher PW. Response properties of saccade-related burst neurons in the central mesencephalic reticular formation. *J Neurophysiol* 78: 2164–2175, 1997.

Hoffman RE, Shi WX, Bunney BS. Nonlinear sequence-dependent structure of nigral dopamine neuron interspike interval firing patterns. *Biophys J* 69: 128–137, 1995.

Hollerman JR, Schultz W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat Neurosci* 1: 304–309, 1998.

Hyland BI, Reynolds JNJ, Hay J, Perk CG, Miller R. Firing modes of midbrain dopamine cells in the freely moving rat. *Neuroscience* 114: 475–492, 2002.

Johnson SW, Seutin V, North RA. Burst firing in dopamine neurons induced by N-methyl-D-aspartate: role of electrogenic sodium pump. *Science* 258: 665–667, 1992.

Kita T, Kita H, Kitai ST. Electrical membrane properties of rat substantia nigra compacta neurons in an in vitro slice preparation. *Brain Res* 372: 21–30, 1986.

Kitai ST, Shepard PD, Callaway JC, Scroggs R. Afferent modulation of dopamine neuron firing patterns. *Curr Opin Neurobiol* 9: 690–697, 1999.

Lisman JE, Grace AA. The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron* 46: 703–713, 2005.

Ljungberg T, Apicella P, Schultz W. Responses of monkey dopamine neurons during learning of behavioral reactions. *J Neurophysiol* 67: 145–163, 1992.

Meck WH, Benson AM. Dissecting the brain's internal clock: how frontal-striatal circuitry keeps time and shifts attention. *Brain Cogn* 48: 195–211, 2002.

Mirenowicz J, Schultz W. Importance of unpredictability for reward responses in primate dopamine neurons. *J Neurophysiol* 72: 1024–1027, 1994.

Montague PR, Dayan P, Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci* 16: 1936–1947, 1997.

Morris G, Arkadir D, Nevet A, Vaadia E, Bergman H. Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron* 43: 133–143, 2004.

Nakahara H, Itoh H, Kawagoe R, Takikawa Y, Hikosaka O. Dopamine neurons can represent context-dependent prediction error. *Neuron* 41: 269–280, 2004.

Paladini CA, Tepper JM. GABA(A) and GABA(B) antagonists differentially affect the firing pattern of substantia nigra dopaminergic neurons in vivo. *Synapse* 32: 165–176, 1999.

Press WH, Teukolsky SA, Vetterling WT, Flannery BP. *Numerical Recipes in C.* Cambridge, MA: Cambridge University Press, 1992.

Rescorla RA, Wagner AR. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: *Classical Conditioning: Current Research and Theory,* edited by Black AH and Prokasy WF. New York: Appleton-Century-Crofts, 1972, vol. 2, p. 64–99.

Satoh T, Nakai S, Sato T, Kimura M. Correlated coding of motivation and outcome of decision by dopamine neurons. *J Neurosci* 23: 9913–9923, 2003.

Schultz W. Predictive reward signal of dopamine neurons. *J Neurophsyiol* 80: 1–27, 1998.

Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science* 275: 1593–1599, 1997.

Shepard PD, German DC. Electrophysiological and pharmacological evidence for the existence of distinct subpopulations of nigrostriatal dopaminergic neuron in the rat. *Neuroscience* 27: 537–546, 1988.

Silva NL, Bunney BS. Intracellular studies of dopamine neurons in vitro: pacemakers modulated by dopamine. *Eur J Pharmacol* 149: 307–315, 1988.

Sutton RS, Barto AG. *Reinforcement Learning: An Introduction.* Cambridge, MA: MIT Press, 1998.

Tepper JM, Martin LP, Anderson DR. GABAA receptor-mediated inhibition of rat substantia nigra dopaminergic neurons by pars reticulata projection neurons. *J Neurosci* 15: 3092–3103, 1995.

Waelti P, Dickinson A, Schultz W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 412: 43–48, 2001.

Wilson CJ, Young SJ, Groves PM. Statistical properties of neuronal spike trains in the substantia nigra: cell types and their interactions. *Brain Res* 136: 243–260, 1977.

Young ED, Robert JM, Shofner WP. Regularity and latency of units in ventral cochlear nucleus: implications for unit classification and generation of response properties. *J Neurophysiol* 60: 1–29, 1988.