Supplemental Information

Supplemental Figures

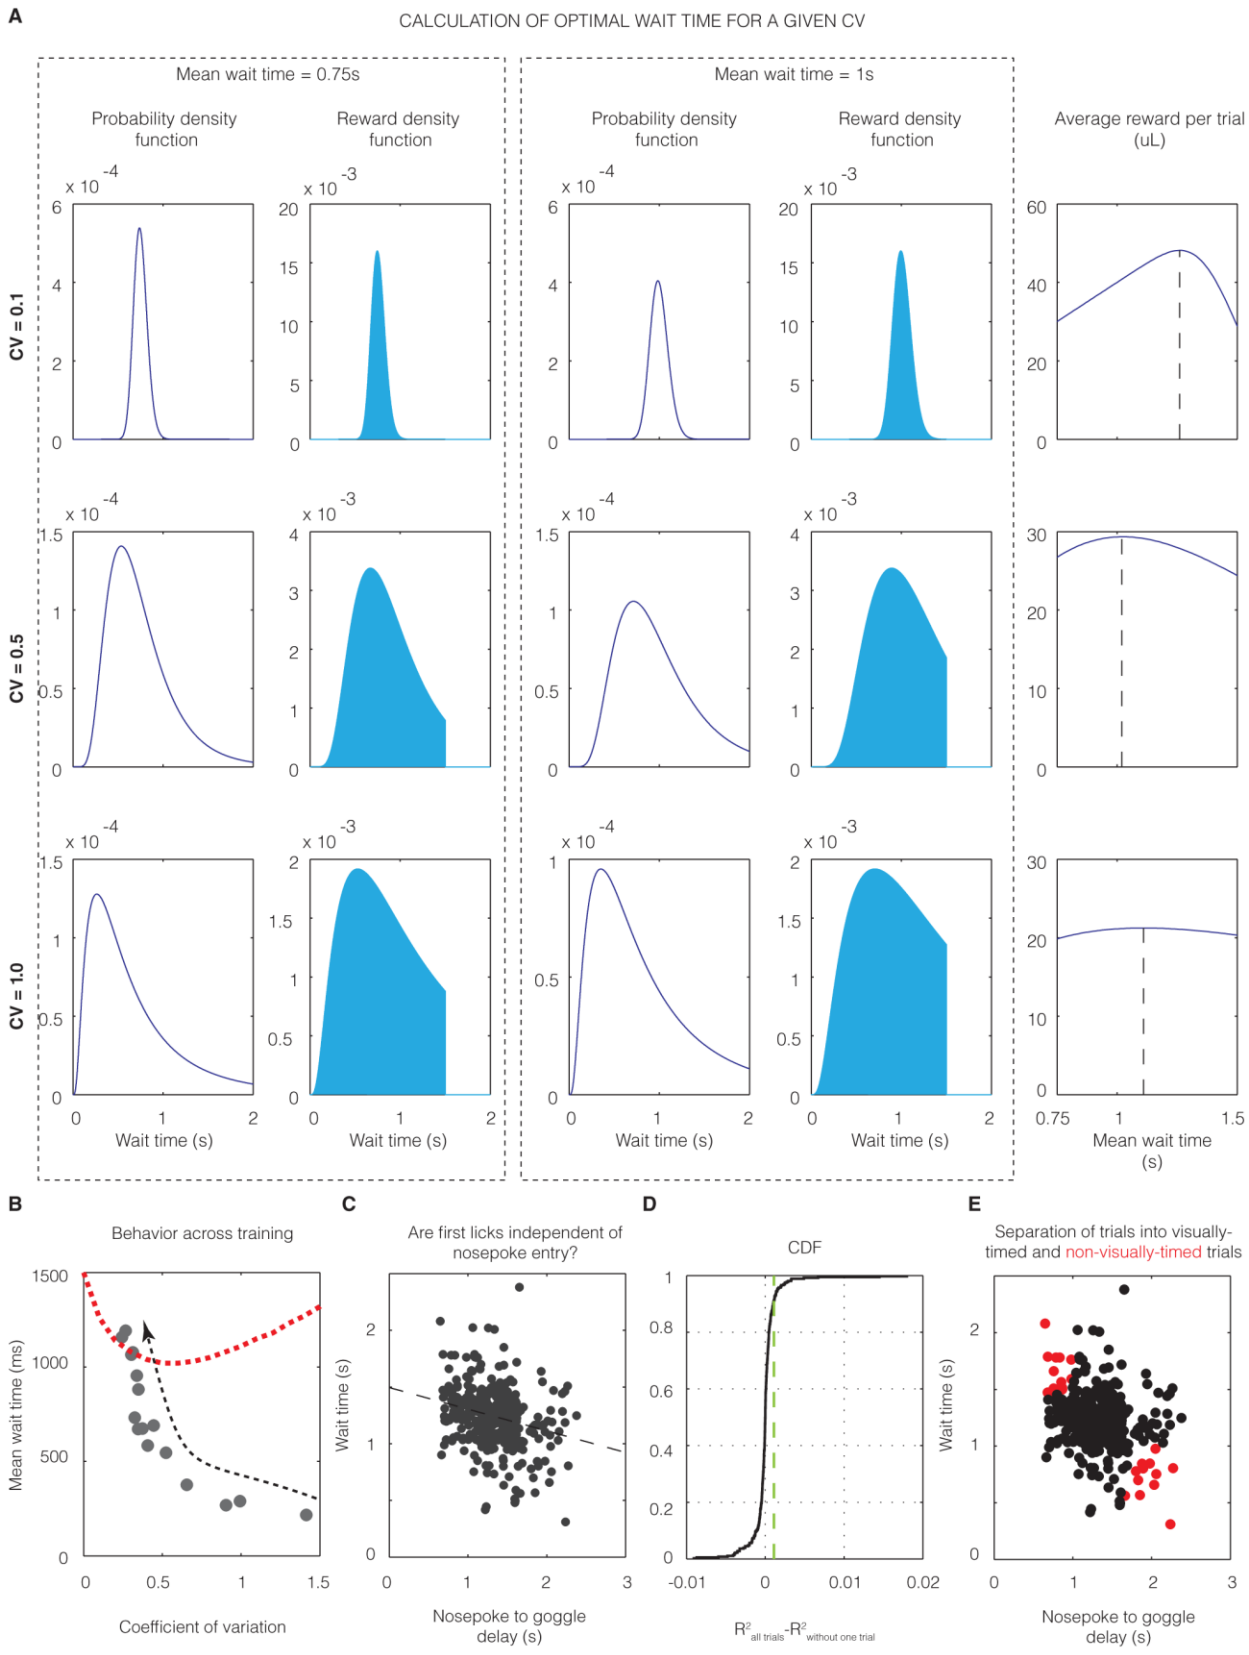
A
CALCULATION OF OPTIMAL WAIT TIME FOR A GIVEN CV
Mean wait time = 0.75s
Mean wait time = 1s
Probability density function
Reward density function
Average reward per trial (uL)
CV = 0.1
CV = 0.5
CV = 1.0
Wait time (s)
Mean wait time (s)
B
Behavior across training
Mean wait time (ms)
Coefficient of variation
C
Are first licks independent of nosepoke entry?
Wait time (s)
Nosepoke to goggle delay (s)
D
CDF
E
Separation of trials into visually-timed and non-visually-timed trials

Figure S1. Related to Figure 1; **A.** Calculation of optimal wait time given a coefficient of variation (CV) is shown. For a given value of CV and mean wait time, the probability density function of licking and the corresponding reward density function (reward obtained at a wait time multiplied by the probability density at that wait time) are plotted, assuming a lognormal distribution of wait times. Thus, the average reward obtained for a given CV and mean wait time is shown by the blue shaded area. Combinations of three different CV's and mean wait times are shown. For a given value of CV, this procedure is repeated for mean wait times ranging from 0.75 s to 1.5 s (last column) so as to calculate the mean wait time that leads to the maximum average reward on a trial (black dashed line). The resultant curve showing optimal wait times for different CV's is shown in Figure 1D. **B.** The behavior of an example animal is shown from the first day of training to the last, showing a constant progression towards optimality. **C-E.** The procedure for separating non-visually-timed trials from visually-timed trials for the example session shown in Figure1E-G. **D.** The CDF of the change in R^2 upon the removal of every individual trial is shown. Every trial that showed a positive change of at least 1.5 times the interquartile change from the median change is judged to be a non-visually-timed trial. This threshold is shown by the dashed green line.

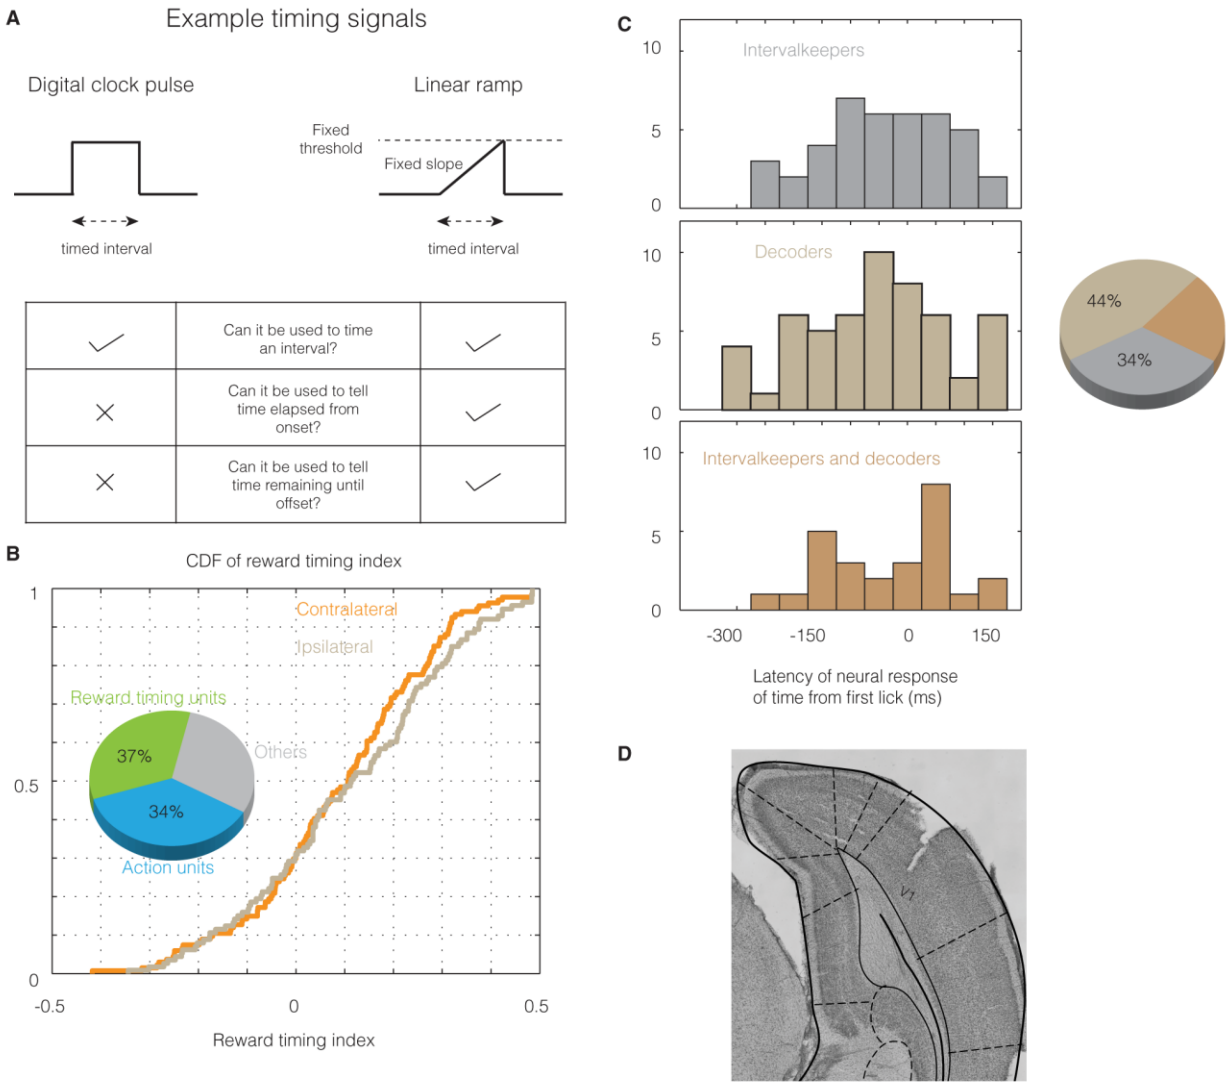

Figure S2. Related to Figure 2; **A.** Two possible example timing signals are shown. One is a digital clock pulse, whereas the other is a linear ramp with constant slope and threshold. Both can be used to time an interval, even though the digital clock pulse cannot be used to continually represent the passage of time. Hence, even neural signals that are time-invariant during an interval (other than at onset and offset) can represent temporal intervals. Note though, that we do not assert that our observed neural signals need abide by either form. **B.** The units that did not show trial-by-trial correlations (n=241) were analyzed to test for reward timing (Chubykin et al., 2013; Shuler and Bear, 2006). The reward timing analysis is explained in *Supplemental Experimental Procedures*. A total of 37% units were classified as reward timing units. The CDF of reward timing index for stimulations of contralateral and ipsilateral eyes (with respect to the hemisphere from which the unit was recorded) are shown. The median reward timing index of 0.11 implies that for a median wait time of one second (close to the asymptotic wait time of animals), the median neural report of time is 1.2 seconds. This is consistent with prior observations (Chubykin et al., 2013; Shuler and

Bear, 2006). **C.** All 122 action units (showing trial-by-trial correlations with action) from trained animals were manually classified (see *Supplemental Experimental Procedures*) as either intervalkeepers or decoders or both. The latency of neural response of time and the proportions are shown for each category. **D.** Histology showing an electrode track in V1. The targeted location for implantation of electrodes was 1.5 mm anterior and 4.2 mm lateral from lambda, at a depth of 1.0 mm.

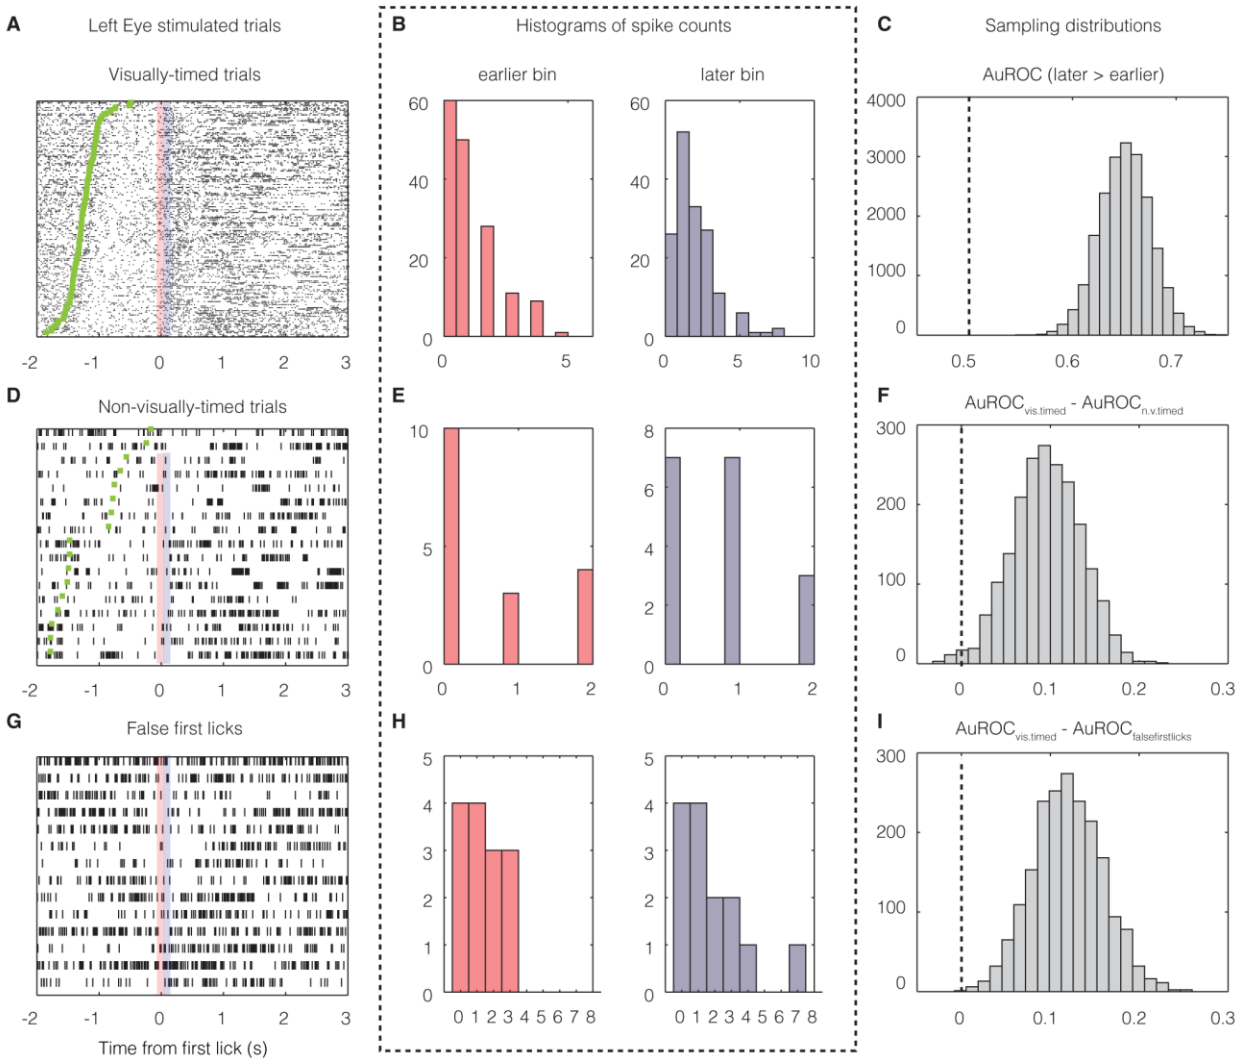

Figure S3. Related to Figure 3; The analysis method is shown for the example single unit shown in Figure 2J. **A.** The latency at which the maximum change in firing rate for this unit (see Experimental Procedures and Figure S7D) happened in visually-timed trials was 50 ms after the first lick. The quantification of this firing rate change was done by analyzing the spike count difference in a 100 ms bin before (red) and after (blue) the first lick. **B.** The corresponding spike count histogram is shown**. C.** The boostrapped sampling distribution of AuROC of spike counts in these two bins. **D & E**. The corresponding firing properties are shown for non-visually-timed trials. **F.** The sampling distribution for the difference in AuROC between visually-timed and non-visually-timed trials. **G & H.** The corresponding firing properties for false first licks (see Experimental Procedures). **I.** The sampling distribution of difference in AuROC between visually-timed and false first licks.

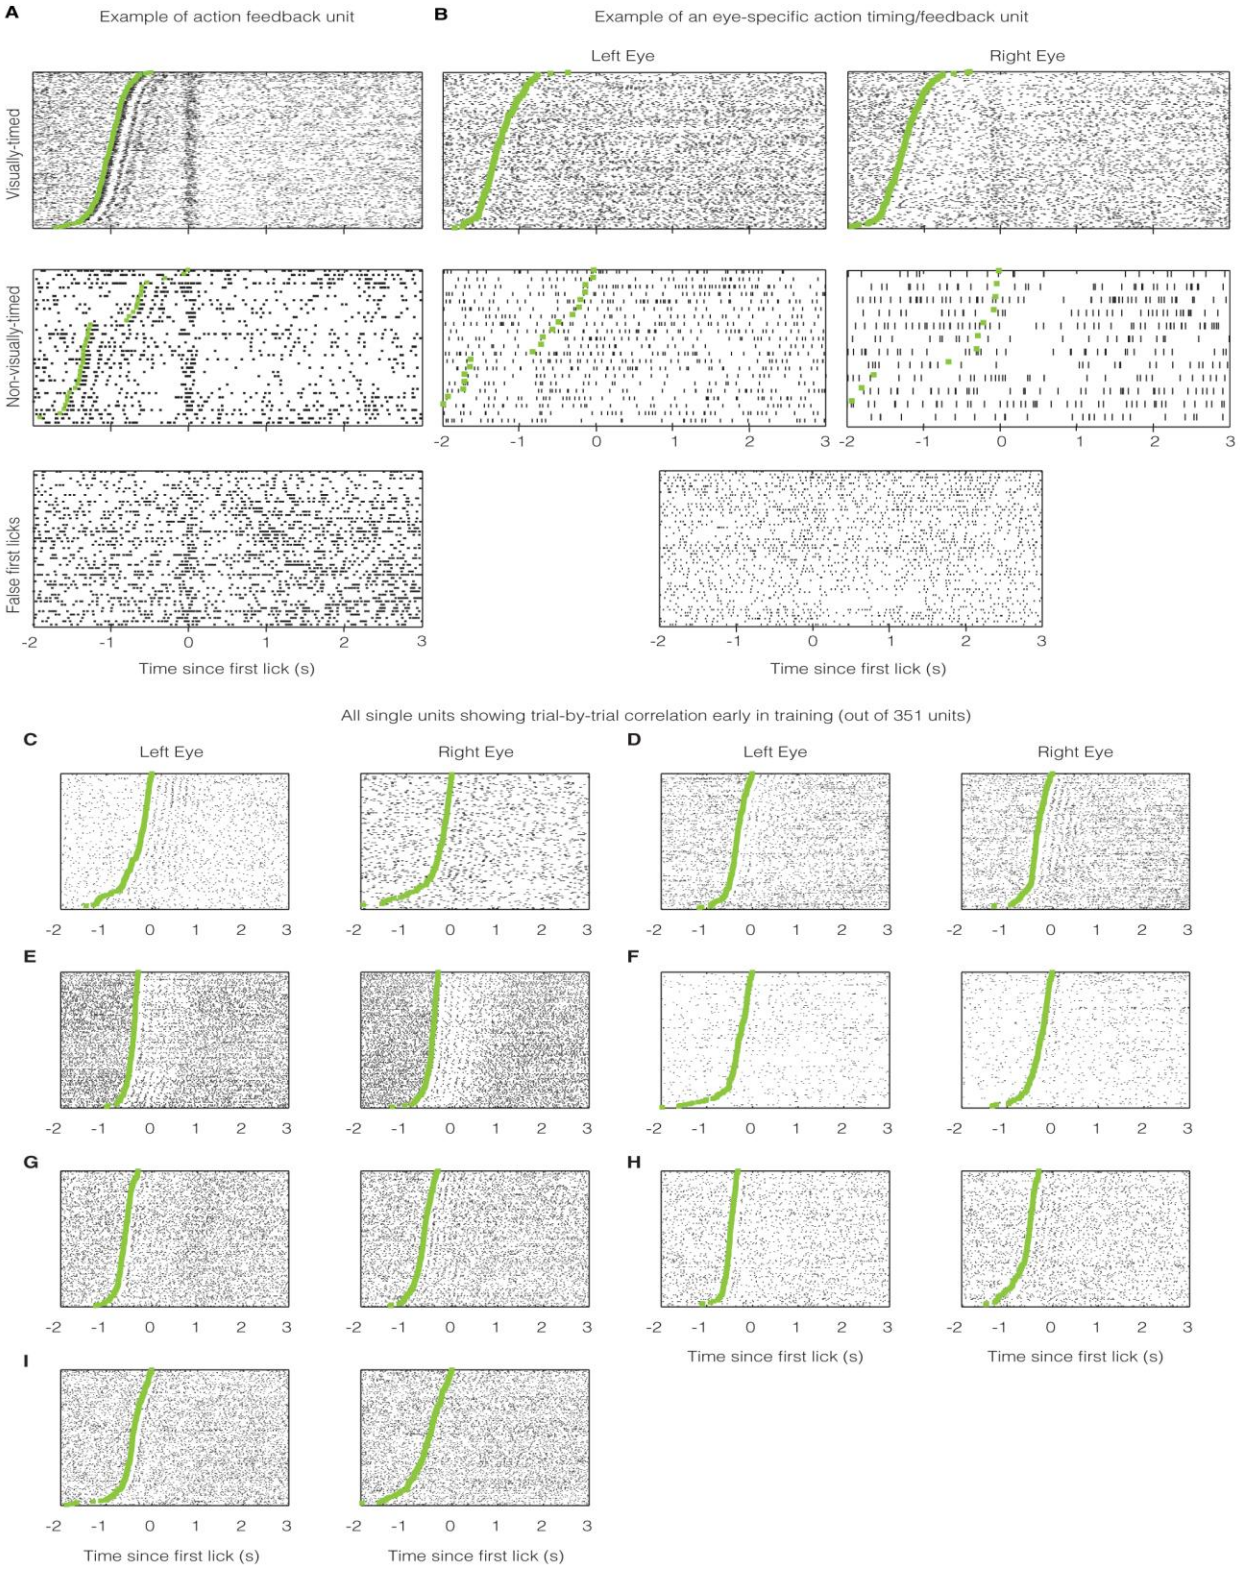
A
Example of action feedback unit
B
Example of an eye-specific action timing/feedback unit
Left Eye
Right Eye
Visually-timed
Non-visually-timed
False first licks
-2 -1 0 1 2 3
Time since first lick (s)
All single units showing trial-by-trial correlation early in training (out of 351 units)
C
D
E
F
G
H
I

Figure S4. Related to Figure 4; **A.** An example feedback unit with no significant difference between visually-timed and non-visually-timed licks (top and middle panel) that also shows a significant trial-by-trial correlation with the first lick on non-visually-timed licks (middle panel) is shown. This unit also shows a significant trial-by-trial correlation with false first licks (bottom panel). **B.** An example of an eye-specific action timing/feedback unit from the population data shown in Figure 4D. Even though this unit shows no obvious trial-by-trial correlation with the first lick in nosepoke-time right eye trials, this difference was not statistically significant due to only 5 trials being available for analysis in the non-visually-timed trials. The unit does show a significant difference ($p<0.001$) between visually-timed and false first licks. However, we required that both the visually-timed to non-visually-timed comparison *and* the visually-timed to false first lick comparison show significant differences. Hence, this unit was classified as an action timing/feedback unit, even though it is unlikely to be a simple reflection of the first lick (due to eye-specificity and significant differences in visually-timed and false first licks). **C-I.** All seven single units out of 353 that were recorded early in training showing trial-by-trial correlations with the first lick as determined by our analysis are shown. The false positive rate of our two-leveled (each level's false positive rate = 5%) nested analysis is lower than at least 5%. Hence, the observed rate of trial-by-trial correlations of 7/353*100 = 2% likely reflects false positives.

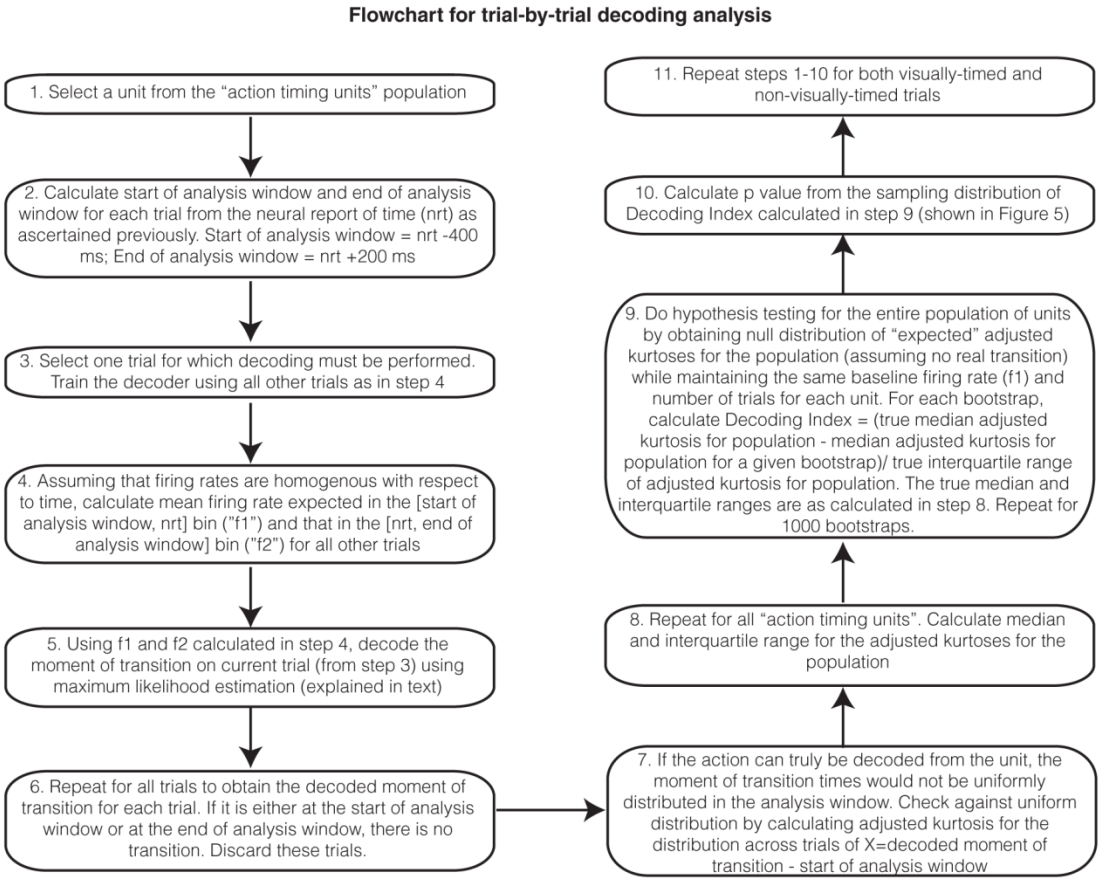

Figure S5. Related to Figure 5. This flowchart explains the procedure used for the trial-by-trial decoding analysis shown in Figure 5. A full description of this analysis can be found in *Supplemental Experimental Procedures*.

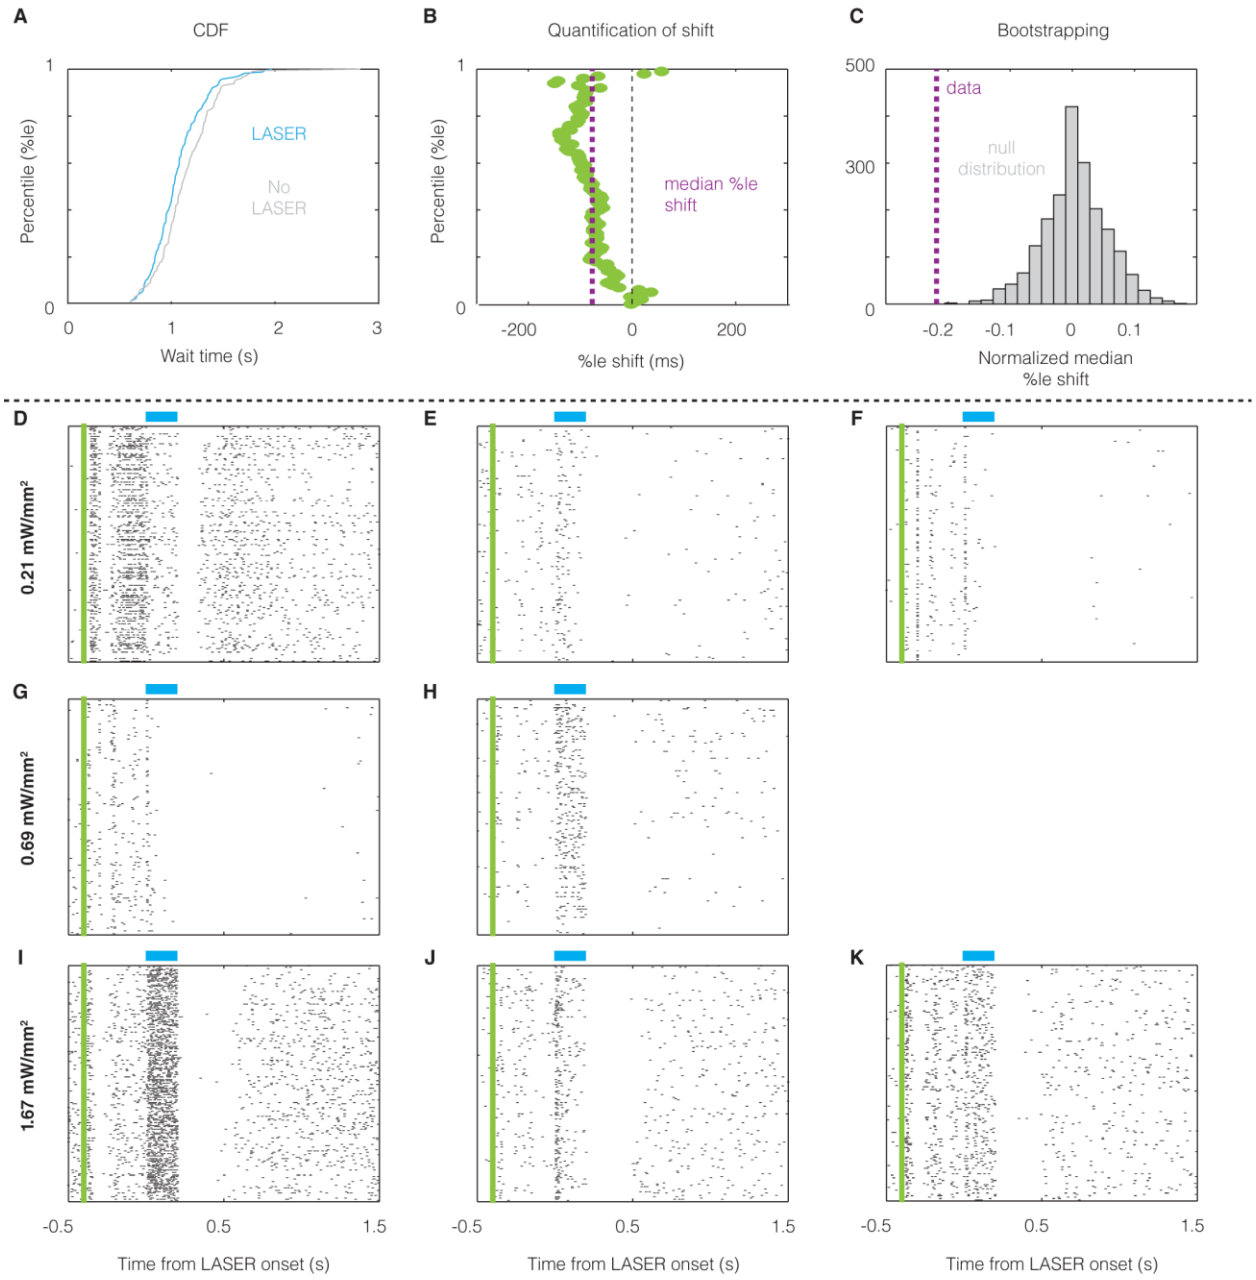

Figure S6. Related to Figure 6; **A-C.** Quantification of shift in wait times induced by the presentation of laser (see Experimental Procedures). **A** is reproduced from Figure 5B. Our aim was to quantify the consistent shift observed throughout the CDF, across different percentiles. To this end, we measured the shift between the two distributions at each percentile and quantified the median percentile shift, as shown in **B**. **C.** The null distribution of the normalized median percentile shift (normalized by the interquartile range of the wait time distribution in no-laser trials) was obtained using bootstrapping by resampling with replacement from the wait time distribution of no-laser trials. The p value of the observed median percentile shift for this session was calculated using a two-tailed measure ($p<0.001$). **D-K.** Single unit recordings from

one animal performed to confirm perturbations of neural activity indeed showed that the laser (blue bar) affected neural activity across all intensities used.

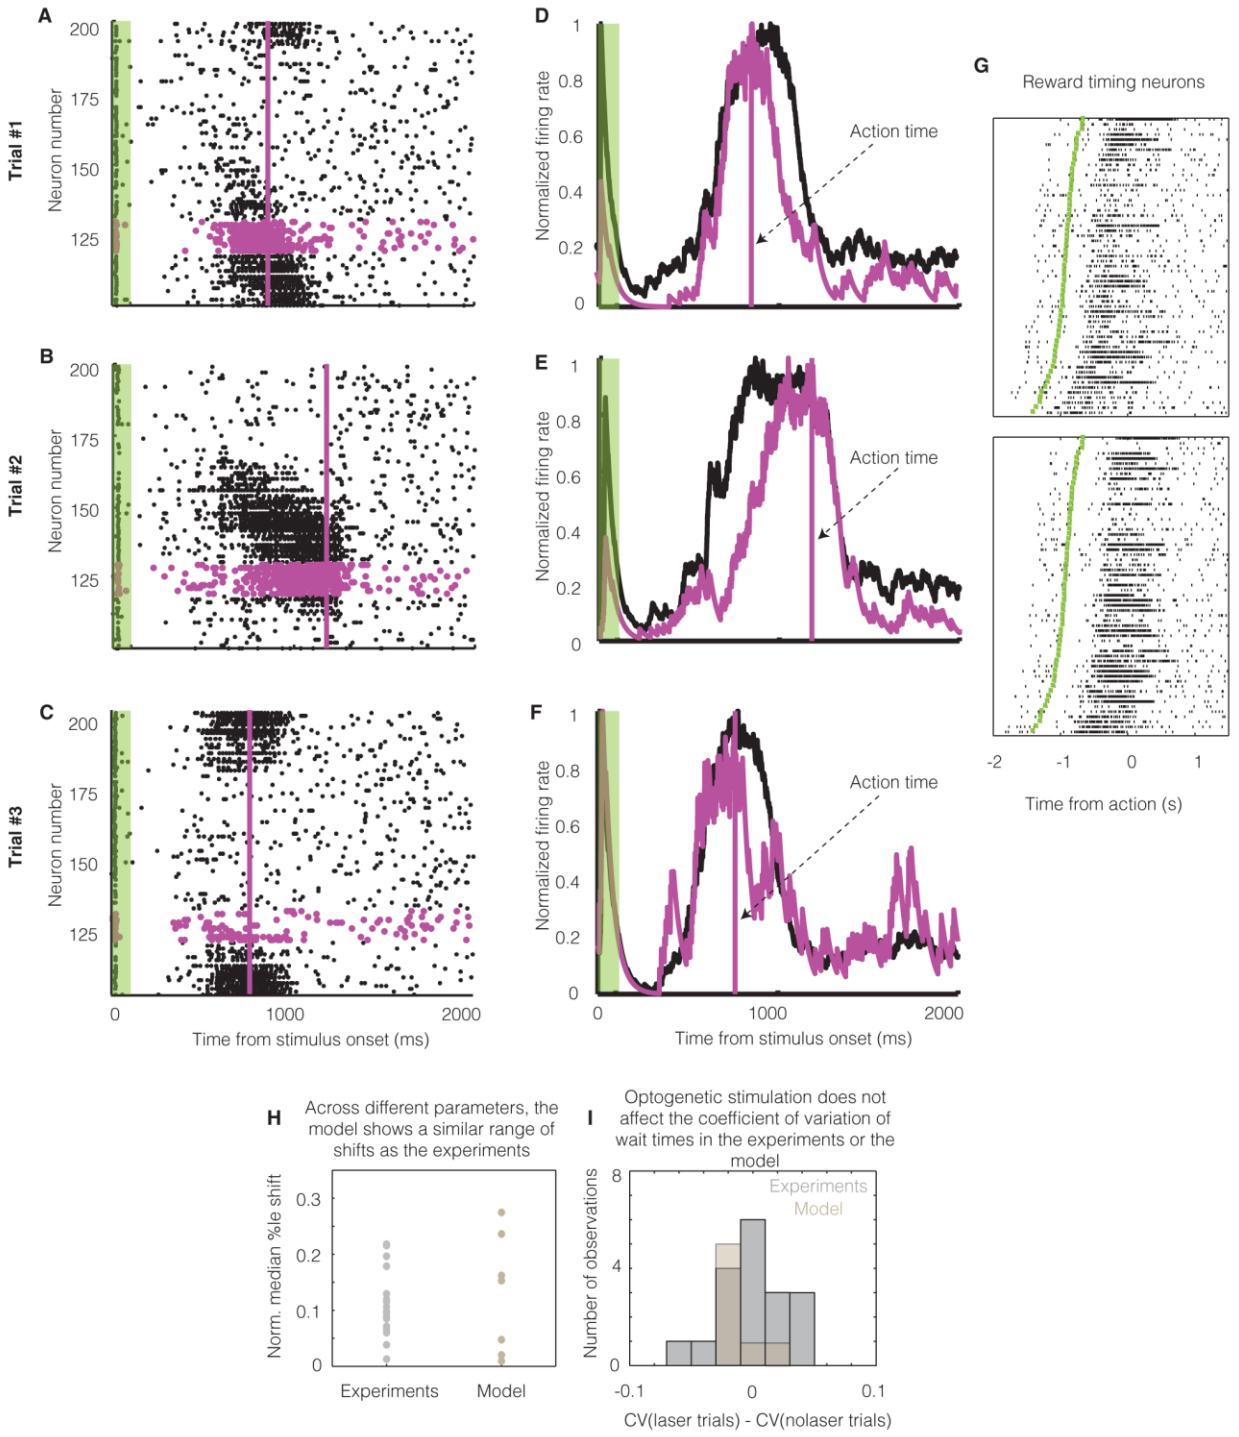

Figure S7. Related to Figure 7; **A-C.** The activity of each unit in the P population of the spiking neuronal network model (as shown in Figure 6A) is shown for three different trials, showing action generation (purple line) following visual stimulation (green bar) (see Experimental Procedures). The neurons controlling the action are shown in purple. **D-F.** The summed population activity is shown for both the action generating

subpopulation (purple) and all others (black). The action is generated at the peak of this summed activity, but even individual units show trial-by-trial correlations with the action (as shown in Figure 6B). Importantly, the population activity of the other neurons does not correlate with the action time, but reflects the median reward time. **G.** The activity of single units that do not control the action does not show trial-by-trial correlations with the action. Instead, these neurons represent the median reward time. Neuron 156 is shown in upper panel and neuron 146 in the lower. **H.** The range of normalized median percentile shifts observed for the model is similar to the experimentally observed shifts for across different parameters. For the experiments, the shifts from all 18 sessions are shown. For the model, the fraction of excitatory and inhibitory neurons activated by the laser in ascending order of observed shifts are: (10%, 20%), (5%, 10%), (8%, 10%), (9%, 20%), (10%, 10%), (6%, 10%) and (4%, 10%). **I.** Individual sessions in both the experiment and the model show non-significant differences in coefficient of variation due to the optogenetic perturbation ($p>0.05$ with Holm-Bonferroni correction for multiple comparisons, two-tailed bootstrapping). The population of shifts in CV is also not significantly different from zero for either the experimental data ($p=0.85$, two-tailed Wilcoxon signed-rank test, $W_{18}=81$, $z = -0.1960$) or the model ($p=0.19$, two-tailed Wilcoxon signed-rank test, $W_7=3$).

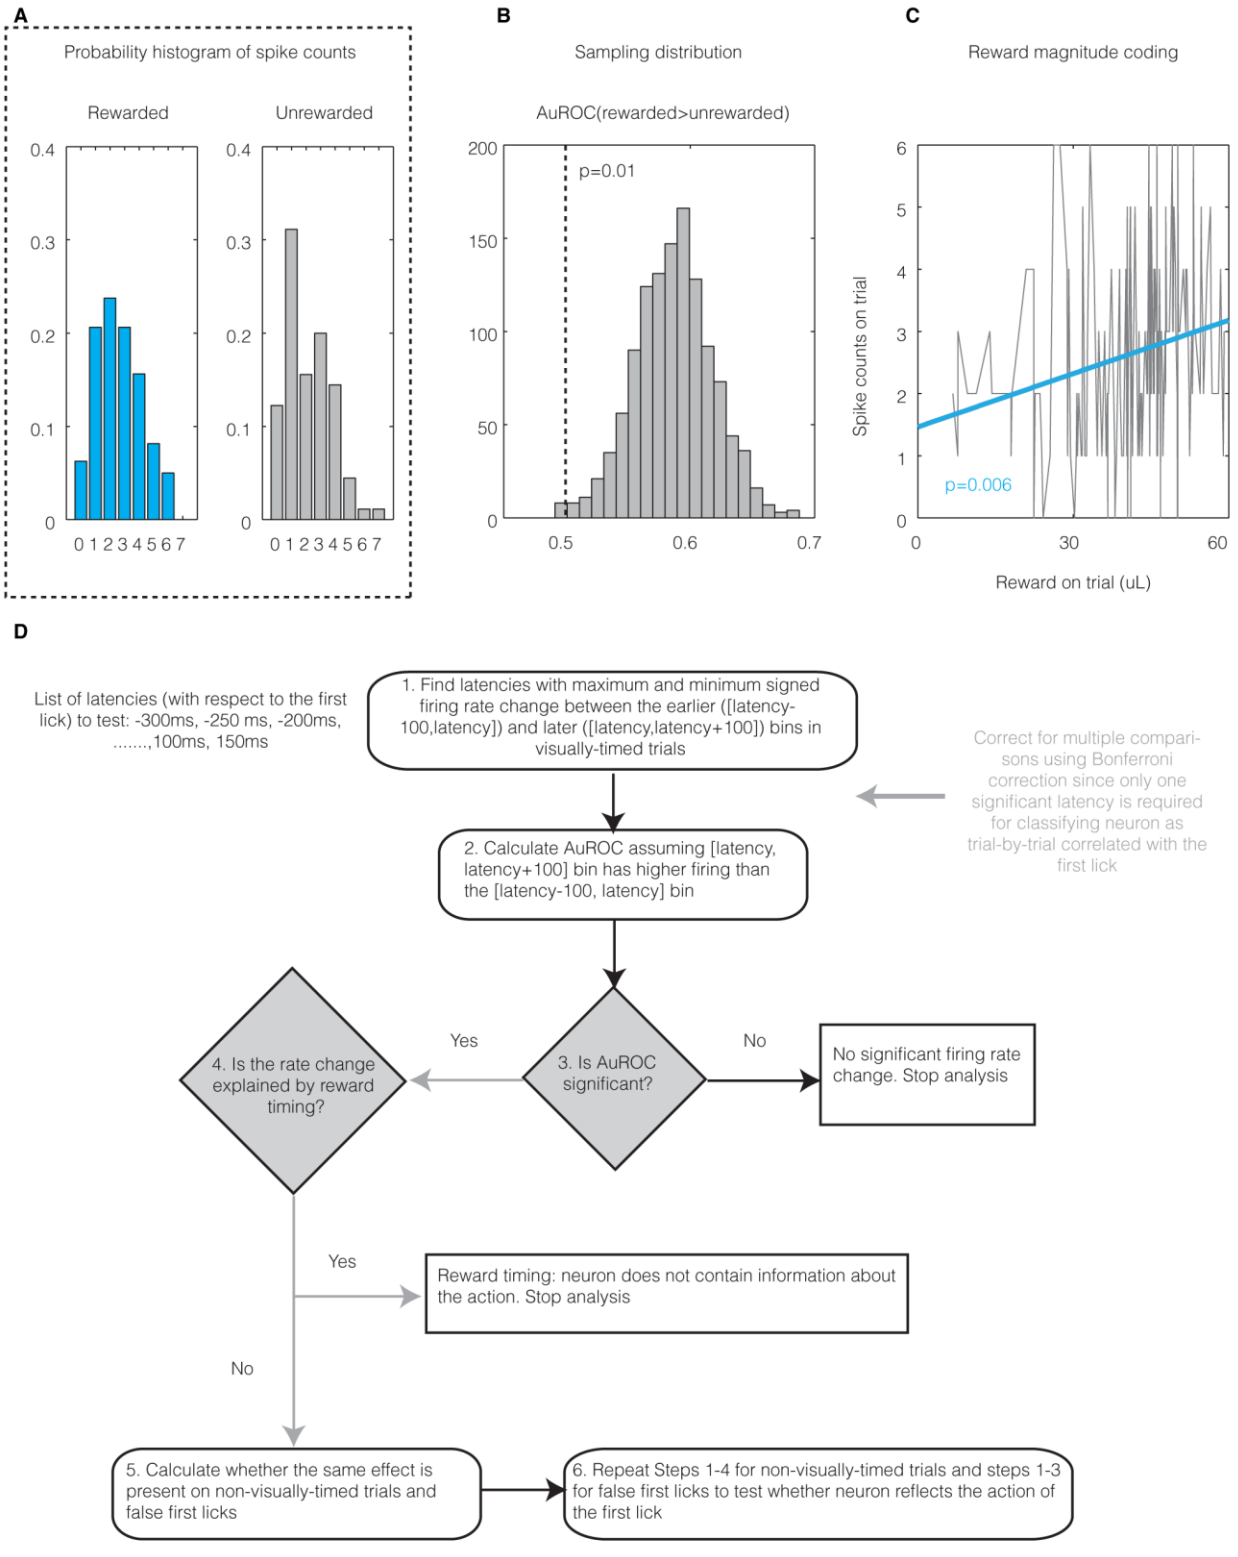

Figure S8. Related to Figure 8 (A-C) and to Figures 2-4 (D); **A-C.** Analysis for testing reward magnitude information in V1 for the example unit shown in Figure 8A. **A.** The

probability histogram of spike counts (normalized for number of trials) shows a higher firing rate in rewarded trials. **B.** This difference is quantified as an AuROC and tested for significance by generating a sampling distribution using bootstrapping (see Experimental Procedures). **C.** If the above comparison shows a significant difference ($p<0.05$) between rewarded and unrewarded trials, we tested whether the spike counts actually reflect the magnitude of reward obtained. The grey line shows the unsmoothed spike counts on a trial with a given obtained reward magnitude; the cyan line is the best fit linear correlation. The unit is deemed as representing reward magnitude information only if both of these tests are significant, and if the difference in spike counts between rewarded and unrewarded trials does not result from the difference in lick density (see *Supplemental Experimental Procedures*). **D.** Flow chart showing the analysis of neural data for timing (see Experimental Procedures). The latencies with maximum and minimum signed firing rate change are detected. An ROC analysis detects whether a linear classified can assert the firing rates as changing at that latency. If the firing rate change is significant, it is tested to see whether it can be explained by reward timing, i.e. firing profile representing an average interval from the visual stimulus. If it cannot be explained by reward timing, the activity is analyzed on non-visually-timed trials and false first licks. If the neuron's activity cannot be explained by the mere presence of an action, it is classified as action timing and if it can be explained by the action alone, it is classified as action feedback.

Supplemental Experimental Procedures

1. **Surgical techniques:**
   The electrode implant surgeries were performed once the animals learned the task (at least three consecutive sessions with median wait times greater than a second), except for three animals for which data was acquired early in training (Figure 4E). During the surgery, animals were anesthetized using isoflurane, had scalps shaved, were given dexamethasone (.2mg/mL), and were placed into a stereotaxic apparatus atop a heating pad. After swabbing the area to be incised with betadine and ethanol, a subcutaneous injection of lidocaine (0.1mL) was administered. An incision was made and connective tissue was retracted to expose the skull. The surface of the skull was dried by applying 3% hydrogen peroxide, following which six screws were uniformly placed across the exposed skull. Bilateral craniotomies and durotomies were then made overlying the primary visual cortex to implant custom-made movable 16-channel electrode arrays (targeted stereotaxic coordinates of 1.5 mm anterior, 4.2 mm lateral to lambda at a depth of 1.0 mm). Following the insertion of electrodes, the ground wires were wrapped around the screws, and a petrochemical gel was applied on the exposed surface of the brain. The entire skull was then covered with dental cement. The animal was given a subcutaneous drug cocktail containing an antibiotic (gentamicin, 1mg/mL) and an NSAID (flunixin, 4mg/mL).

2. **Histology and immunohistochemisty:**
   General perfusion and histology: Animals were deeply anesthetized with pentobarbital and transcardially perfused with 1x PBS followed by 4% PFA. Brains were then harvested and

stored in 4% PFA for two hours until being transferred to a solution of 30% sucrose in PBS until time of sectioning. Brains were sectioned into 60µm slices and mounted on gelatin-coated slides (FisherBrand). Slices were then imaged using an inverted microscope (Zeiss Instruments). Electrode tracks were verified to be found in V1 (see Figure S2D). Based on the distance moved by electrodes and the histology, the cortical layer for each unit was estimated. The different classes of units were found to be distributed across all targeted layers (4, 5 and 6). It must be noted that these numbers reflect a sampling bias towards the initial depth at which the electrodes were implanted. For the 122 "action units", 62% were from layer 4, 25% from layer 5 and 13% from layer 6. For the 38 "action timing units", 63% were from layer 4, 26% from layer 5 and 11% from layer 6. For the 7 "action feedback units", 4 were from layer 4, 1 from layer 5 and 2 from layer 6.

Immunohistochemistry for channelrhodopsin-infected animals: Animals were perfused and brains sectioned as described above. After sectioning, using a mini-orbital shaker, slices were washed in PBS for ten minutes three times before being blocked in 10%NDS/0.5%triton-X/PBS solution for one hour. Slices were then incubated with a primary anti-GFP antibody (ab290, abcam, 1:5000) overnight at 4C. After three 10m PBS washes, slices were incubated with the secondary donkey anti-rabbit Alexa Fluor 488 antibody (Life Technologies, A21206, 1:1000) to verify protein expression (see Figure 5D).

3. **Behavioral task:**
Prior to behavioral training, goggle posts were mounted using dental cement on the skull of animals under sterile surgical conditions. Following a recovery period of 7-10 days, animals were water deprived but given ad-lib access to food. Animals began behavioral training after two days of water deprivation. During this period, they obtained water only within the behavioral task in addition to brief scheduled access periods every day and a night of free access every week.

The behavioral task was a visually-cued timing task, performed in a standard operant chamber (Med Associates) containing a port (nosepoke) that housed a lick tube. A behavioral session consisted of 360 trials. A nosepoke entry was required to initiate a trial. After the nosepoke, a random delay of 0.6-1.6s expired during which the animal is allowed to lick. At the expiry of this random delay, if the animal had not licked in the previous 0.2-1.2s (randomly set on each nosepoke entry), the visual stimulus was presented to either the left or the right eye (Figure 1A). If instead the animal had licked, 0.2-1.2s had to elapse without a lick for the visual stimulus to be presented. The visual stimulus was a full field green LED-powered flash of 100ms presented using detachable head-mounted goggles covering each eye. Entries into the nosepoke and licks were registered using two IR emitter-receiver pairs (Med Associates). If, without having exited the nosepoke, the rat subsequently licked following the offset of the visual stimulus, it immediately received a water reward with a probability of 5/6. Rewarded and unrewarded trials were randomly interleaved such that no more than 4 unrewarded trials could be presented in a row. The reward volume was chosen depending on the wait time as shown in Figure 1B. The reward delivery system used a solenoid (custom modified to eliminate any operational noise) which was controlled by a microcontroller generating sub-millisecond precision for open times. A trial ended when the rat exited the nosepoke. An obligatory intertrial interval (ITI) of 500ms was enforced after

every trial such that the animal had to wait outside the nosepoke until its expiry. Entries into the nosepoke within the ITI did not result in a visual stimulus presentation and caused a reset of the ITI. After every 12 trials, animals were presented with a "NoStim" trial (total of 30, in addition to 360 visual stimulus presentations). In this trial type, the animal had to wait in the nosepoke as it would on a visual stimulus trial but instead of a light flash, no stimulus was presented. These NoStim trials were included so as to determine whether licking behavior was driven by and timed from the visual stimulus, as well as to assess neural activity on trials in which a lick was registered following this NoStim event. As mentioned in the text, licks after a NoStim or false entry to the nosepoke (during the mandatory ITI) were combined to form the "false first licks" as shown in Figure 3A,B and Figure S3 and Figure S4.

A slightly different procedure was used for behavioral shaping early in training. On the first day of training, the animals were given a free reward on every trial 500 ms after the corresponding nosepoke entry. If the animals did not stay in the nosepoke until visual stimulus presentation, no further free reward was delivered on the next nosepoke entry. The lick-free pause (shown in Figure 1A) required to deliver a visual stimulus was set between 200-400 ms. Free rewards were then subsequently delivered only on the first 200, 100 and 20 trials respectively in the next 3 sessions. On the 4^{th} session, the lick-free pause requirement was increased to 200-600 ms, followed consecutively by 200-800 ms, 200-1000 ms and finally, 200-1200 ms, as they progressed in training. This resulted in the animals licking prior to visual stimulus presentation only on a very small fraction of trials (<10%) by the final stage. In some cases, to speed the rate at which the animals increased their wait time, the ramp duration (wait time at which maximum reward was obtained) was increased to 2 seconds. Once the median wait time increased beyond 1s, the ramp duration was held at 1.5 seconds.

4. **Separation of non-visually-timed trials from visually-timed trials:**
When parsing non-visually-timed and visually-timed trials, any trial with a wait time less than 300 ms and a wait time greater than 3000 ms was also grouped into the category of non-visually-timed trials (we do not necessarily know if these were timed from the nosepoke, but it is unlikely that they were timed from the visual stimulus, and since we were only interested in testing visually-timed trials against non-visually timed trials, these were included along with the non-visually-timed trials). Further, if the coefficient of variation of the wait time distribution was greater than 0.6 (very low precision), it was deemed that the animal had not yet learned to time its licks from the visual stimulus. This happened only early in training and neural recordings were obtained from such sessions only in the three animals shown in Figure 3. Nevertheless, the neural analysis as discussed below was carried out to test for trial-by-trial correlations with the action, even though, at this early stage, we could not test whether visually-timed and non-visually-timed trials showed different neural responses.

5. **Neural recordings and spike sorting:**
Once the animals were deemed behaviorally-trained (at least three consecutive sessions with median wait times greater than one second), custom-built 16-channel electrode arrays were implanted bilaterally in V1 (targeted stereotaxic coordinates of 1.5mm anterior and 4.2 mm lateral to lambda at a depth of 1mm) following sterile surgical procedures. The same parameters were used for recovery and subsequent water deprivation prior to recordings

during behavior. Signal from each electrode array was sent to headstage amplifiers, and, bandpass filtered (1-10 kHz) neural signals were sampled at 33 kHz by commercially available hardware (Neuralynx). The electrode arrays were moved down by a distance of approximately 15-60 µm on a day, if necessary, so as to increase the yield of single unit recordings. Commercially available software packages were used to manually sort the waveforms into single units (Offline Sorter, Plexon) and to visualize the timestamps (Neuroexplorer, Nex Technologies) of these sorted units. The manual sorting was done by tracing the boundaries of the clusters in a 3D space formed by three of the following variables: principal components 1 and 2, valley timing, non-linear energy and peak minus valley amplitude. Units with unstable waveforms or greater than 3% refractory period (1.5 ms) violations were excluded, along with multi-unit recordings. A unit was deemed as different from any previous recording if one of the following criteria were met: 1) no unit was observed from the same channel on the previous session; units recorded from the same channel on the previous day showed 2) different waveforms, or 3) different visually-evoked response; or, 4) the electrode array was moved more than 60 µm between the two sessions. If none were met, then the unit was deemed as a repeat. In this case, only the first instance was counted. Experiments were continued only in the animals (n=5) that produced usable neural recordings. No blinding or randomization of subjects was performed as these experiments used a within-subject design. See #2 above for information on histology.

6. **Analysis of units recorded early in training:**
The animals from which the neural recordings early in training were obtained were implanted prior to any behavioral training and were also run on a reward profile that ramped to 1 second instead of 1.5 seconds. The lowest coefficient of variation of the wait times from a session with at least one single unit recorded was 0.45 across these animals, showing that the recordings were all made in an early stage of learning.

   Early in training, the earliest latency tested for trial-by-trial correlations with the action depended on the wait time distribution (this was true also for neurons later in training, but the following procedure later in training always resulted in the earliest latency being -300 ms). This was done to avoid the short-latency visually-evoked activity from directly affecting action responses. If the 25^{th} percentile of the wait time distribution was greater than 500 ms, the earliest latency tested was -300 ms from the action. Instead, if it fell between 400 and 500 ms, the earliest latency tested was -200 ms, and if it was between 300 and 400 ms, the earliest latency tested was -100 ms and if it was between 200 and 300 ms, the earliest latency tested was 0 ms. If the 25^{th} percentile fell between 100 and 200 ms, the earliest latency tested was 100 ms with respect to the action and if it was less than 100 ms, the units from that session were not analyzed for action timing.

7. **Testing whether there is a firing rate change locked to the first lick:**
The significance testing for AuROC was done by estimating the sampling distribution using bootstrapping (Figure S8D). This was done by sampling with replacement 18000 (= number of latencies tested×2000; so p=0 implies p<0.001 after Bonferroni correction) times from the available pool of trials, i.e. individual trials were resampled, including both the earlier and later bin. The p value was tested by reading off the percentile at which chance level of 0.5 for the AuROC lay, using a two-tailed test (Obuchowski and Lieber, 1998). If this p value was

found to be less than 0.05, it means that there was a significant change in firing rate at the corresponding latency. Further action-timing analysis was performed only if this was true. The next step, as mentioned above, was to test whether the observed change in firing rate could be explained by visually-evoked activity. This was done by generating fake first lick events (not to be confused with false first licks) for each trial, by randomly permuting the wait times. This ensured that the mean interval between the visual stimulus and the fake first lick is preserved, but the relationship to individual trials, being shuffled, was randomized. If such shuffling did not affect the AuROC at the corresponding latency with respect to the new fake first licks, then the change in firing rate was judged to have resulted from visually-evoked activity, and not timing of, or execution of, the action. Such a procedure was repeated 18000 times to generate the sampling distribution of $AuROC_{observed}$ - $AuROC_{fakefirstlicks}$. The p value was read off as the percentile at which chance level of zero lay. If this p value was also found to be less than 0.05, the neuron was deemed to show trial-by-trial correlations with the action ("action unit").

8. **Testing for specificity of action timing with respect to eye of origin:**
A significant ($p<0.05$) eye difference for the action timing response on action timing/feedback units (Figure 4D) was tested by checking whether the AuROC for the action response was different following the stimulation of both eyes. This difference ($|AuROC_{lefteye} - AuROC_{righteye}|$) is labeled as "eye-specificity index" in Figure 4D.

9. **Trial-by-trial decoding analysis:**
This analysis assumes that the action is predicted by a sudden change in firing rate of a unit. It also assumes that the activity of a unit can be described as a Poisson point process. That the firing rates of units show a transition locked to the action has already been confirmed in Figures 2 and 3. The aim of the trial-by-trial decoding analysis is to test whether the moment of first lick on a given trial can be decoded from the firing rate transition on that trial. To this end, for a given single unit, a decoded transition time is calculated on every trial. This operation defines the decoder. To calculate the transition, the decoder assumes that the firing rate prior to transition and following transition is fixed and time-independent. To maximize the likelihood that this assumption (that firing rates before and after transition are time-independent) is correct, the analysis window is restricted to start 400ms prior to the expected transition time and end 200 ms after the expected transition time. Further, in order to avoid transient visual response, if the start of the analysis window is less than 300ms away from the visual stimulus, that trial is ignored for the analysis. The expected transition time for a given unit is taken as the neural report of time calculated previously (shown in Figure 4C). The decoder is then trained to learn the expected firing rates before and after transition. This training is done by estimating these corresponding firing rates from all other trials excluding the trial of interest. Once the firing rates before (f1) and after (f2) transition are known, the moment of transition is calculated using Maximum Likelihood Analysis, as explained in (West and Ogden, 1997). This procedure is repeated for every trial to calculate the decoded moment of transition for every trial. If the decoded moment of transition is either at the start of the analysis window or at the end of the analysis window, it means that there is no transition on that trial (i.e. firing rate is constant throughout the analysis window). Thus, these trials are discarded.

A seemingly intuitive first step to test whether the decoded moment of transition accords with the action is to test for a linear relationship between both. However, such an analysis is incorrect since the decoded moment of transition must lie within the analysis window, which is locked to the action. Thus, there would always be a linear correlation between the decoded moment of transition and the action (seen in Figure 5A for both example units). If there really was no transition that predicted the action on a trial-by-trial basis, the decoded moment of transition would be distributed uniformly with respect to the action in the analysis window. Needless to say, such detected transitions (when there is no real transition) are obtained by chance, since on every trial there is a non-zero probability that the firing pattern behaves as if there is a transition. Thus, the null hypothesis for this test is that the decoded moment of action for any given unit is distributed uniformly in the analysis window (Example 2 in Figure 5). On the other hand, if there truly is a transition that predicts the action, the distribution of decoded moments of transition would be peaked (Example 1 in Figure 5). The deviation from the null distribution can then be calculated by the excess kurtosis of the distribution of decoded moments of transition from that of a uniform distribution (=1.8).

Since doing the hypothesis testing on individual units does not test whether the population of "action timing units" predicts the action time as a whole, we do the significance testing for the entire population. To do so, we first calculate the adjusted kurtosis (excess kurtosis over uniform = kurtosis – 1.8) for each unit within the population. The distribution of adjusted kurtosis across all units forms the observed distribution of kurtosis. The expected null distribution of kurtosis can be determined by calculating the decoded moments of transition for every trial of each unit, assuming that the firing rate does not show any transition (i.e. remains at f1 for the entire analysis window). Since the expected variability in kurtosis depends strongly on the number of trials, to generate the null distribution of kurtosis, the number of trials for each unit is held constant. For a given iteration of the bootstrapped null distribution, in order to normalize for the expected variability in kurtosis due to the number of trials, a statistic called the Decoding Index was defined (akin to a z-score) as:

Decoding Index = (true median adjusted kurtosis for the population – median adjusted kurtosis for null distribution)/true inter-quartile range of adjusted kurtosis for the population.

If Decoding Index = 0, the population predicts the action only at chance levels. Thus, the null hypothesis can be tested against by assessing whether the sampling distribution (obtained over 1000 iterations of bootstrap) of the Decoding Index is significantly greater than zero. The one-tailed p value is the percentile at which the sampling distribution of Decoding Index crosses zero (Figure 5B).

This entire analysis is repeated for both visually-timed and non-visually-timed trials and the results are as shown in Figure 5B.

10. **Classification into intervalkeepers and decoders:**
Action units (showing trial-by-trial correlations with the action) were manually classified into intervalkeepers or decoders for the purpose of Figure S2C using the following criteria: if the unit showed a sustained response from visual stimulus to action, it was called a

intervalkeeper. If instead, it only showed a peak modulation around the action, it was called a decoder. If it showed both, it was called a intervalkeeper and a decoder.

11. **Reward timing analysis:**
In order to test whether the firing rate showed any long-latency modulations, we compared long-latency responses evoked by the stimulation of either eye to the baseline firing rate observed just prior to the stimulus presentation. A window of 100 ms was used to obtain the baseline spike counts. The activity in the first 300 ms following visual stimulus onset was considered as a transient response, and hence, was not analyzed. We used a sliding window of 100 ms, sliding by 50 ms, up to a maximum of 2 seconds, to test for significant modulations of firing with respect to baseline, using an ROC analysis (quantified by AuROC). Significance testing was done using bootstrapping, similar to the action timing analysis. The longest streak of consecutive significant modulation was identified, and if this streak lasted less than 300 ms, the unit was deemed as a non-reward timing unit. Instead, if the unit showed significant modulation in more than 6 consecutive bins (spanning 300 ms), the unit was deemed as a reward timing unit. In order to determine the neural report of time, the firing profile was first classified as a sustained excited, sustained inhibited or a peak excited or peak inhibited profile. If the longest streak of modulation did not begin for at least 500 ms after the visual stimulus offset, the response was classified as a peak; otherwise, it was classified as a sustained response. The sign of the modulation determined whether it was excited or inhibited. For sustained responses, the neural report of time was defined as the last bin that showed significant modulation compared to the baseline, whereas for peaks, it was defined as the bin showing maximum modulation. The reward timing index for a unit was defined as the ratio of the difference between the neural response of time and the median wait time to the sum of these two quantities. For instance, if the neural report of time was 1.22 seconds for a session with median wait time equaling one second, the reward timing index will be ~0.1. This is the median reward timing observed for the population, as shown in Figure S2B.

12. **Optogenetic surgeries:**
Three craniotomies were made overlying primary visual cortex (V1; 1-2mm anterior, 4-4.5mm lateral from Lambda) in order to allow for cortical infections at two depths (0.9-1.0mm and 0.5mm ventral from dura). A nanoliter injector (Drummond Scientific) was backfilled with mineral oil as well as either the virus containing the channelrhodopsin (ChR2) construct (UPenn Vector Core: AAV9.CAG.ChR2-Venus.W.SV40[49]) or saline. The injector was then lowered into the most ventral depth and KwikCast (World Precision Instruments) was applied to the surrounding area and allowed to cure before infusing. The pump introduced 32.2nL of infusant every 10s over 2.5 minutes for a total volume of 483nL after which the tip was kept in place for 5 minutes to ensure that all infusant diffused away from the tip. The injector was then raised to the more dorsal infection site and another injection was similarly carried out. This process was repeated for each craniotomy for a total of 12 infections delivering roughly 5.8µL of infusant throughout cortex. Following infusions, Kwikcast was placed atop each hemisphere's three craniotomies, and, after curing, the skin was sutured together.

13. **Rationale for using median percentile shift as the shift statistic:**

One possible way to quantify whether the optogenetic perturbation caused any significant change in wait time is to do a Kolmogorov-Smirnov test to check for significant differences in distributions of wait times between laser and no-laser trials. However, our aim was to test whether there was a consistent shift introduced by the presence of the laser, rather than merely testing for differences in distribution (for instance, KS test also returns significance for significant change in variance). To this end, we defined the shift statistic as the absolute median percentile shift (see below) and generated bootstrapped sampling distributions for significance testing. This is also why we avoided standard ordinal tests that simply check for differences in median.

14. **Testing whether the aggregate effect of the laser was present in all animals:**
In order to test whether the significant effect of the laser in the aggregate as shown in Figure 6E was consistently present in all animals, we tested individual experimental animals against the controls. Even under greatly reduced statistical power, we found that two animals showed significant differences with the control group (p values of 0.01 and 0.03) with the remaining animal missing significance (p=0.11, one tailed, Welch's t test). In this last animal, three individual sessions showed significant shifts in wait time distribution between laser and no-laser trials. In order to test whether the optogenetic results were significant merely due to the duplication of data at each laser power, we also tested the difference between the experimental and saline groups after averaging the observed shifts at each intensity within an animal. The comparison is still significant (p=0.0061, one-tailed, t=2.92, Welch's t test).

15. **Testing whether laser presentation led to an increase in licking on NoStim trials:**
To test whether the laser presentation led to an observable increase in the probability of licking in individual experimental animals, we tested for differences in the probability of licking in the presence and absence of laser using bootstrapping with 10,000 re-samples with replacement. To correct for the multiple comparisons, we used Holm-Bonferroni correction (since it is more powerful than Bonferroni correction) and found no significant effect in any animal ($p>0.05$).

16. **Spiking neuronal model:**
Neurons in each population were modeled as leaky integrate-and-fire units. The equation for the membrane potential for neuron i is described by the differential equation

$$C dV_i/dt = g_L (E_L - V_i) + g_{E,i} (E_E - V_i) + g_{I,i} (E_I - V_i),$$

where C is the membrane capacitance, V_i is the membrane potential; g_L, $g_{E,i}$ and $g_{I,i}$ are the conductance of the leak current, total excitatory and total inhibitory synaptic currents, respectively, and E_L, E_E and E_I are the corresponding reversal potential of each current.

To simulate the effect of optogenetic activation of neurons, we included an additional current of the form $I_{ChR2} = g_{ChR2,i} (E_{ChR2} - V_i)$, where $g_{ChR2,i}$ and E_{ChR2} are the conductance and reversal potential of the membrane current associated with ChR2.

Except for the leak conductance, the total synaptic conductance takes the form (e.g. for excitation) $g_{E,i} = g_E \Sigma W_{ij} s_j$, where g_E is the maximum conductance of the current, s_j is the

synaptic efficacy of presynaptic synapse j and W_{ij} is the synaptic strength. Each s_j is modeled as an exponentially decaying function that obeys the differential equation

$$ds_k/dt = - (1/\tau_s)\, s_k + \rho\, (1-s_k)\, \Sigma\, \delta(t-t_{pre}),$$

where τ_s is the synaptic decay time constant, ρ is the fractional change in synaptic efficacy due to the arrival of a presynaptic spike that occurs at time t_{pre} and the sum is over all pre-synaptic spikes; δ is the Dirac delta function.

We computed the response of the network to an initial presentation of a 100 ms long full-field visual stimulation, represented as a train of action potentials with a constant firing rate. The initial values of the synaptic weights W_{EE}, W_{PP} and W_{PE} were obtained after training[27,28] so that the activity of the excitatory population E exhibited an initial rise in its firing rate in response to the visual stimulation followed by a slow decay due to reverberations within the population that decayed to baseline 1000 ms after the stimulus on average. Changes in synaptic weights were calculated using the Reward-Dependent Expression (RDE) rule[27]. The output from the E population, and consequently from the I population onto P, led to a response that was qualitatively similar to the previously observed peak responses. Once trained, the values of these synaptic weights remained constant for the rest of the simulations.

To simulate the spontaneous activity of neurons, all three populations received feed-forward excitatory input in the form of random Poisson distributed spike trains at a constant frequency of 5 Hz. The addition of noise induced trial-by-trial variability in the response of the network. Furthermore, to represent a transient visual input, excitatory populations E and P received additional excitatory feed-forward input at a higher mean rate, representing drive from the visual thalamus. In this network, the E population represents "intervalkeepers" since they represent the target interval in the duration of sustained responses triggered by the visual stimulus. The P population represents "decoders" since they represent the moment of target interval expiry. Action neurons were selected from the P population and consisted of 10 neurons that were spatially close in synaptic connectivity, and thus showed a high degree of correlation. The action was signaled when the summed firing rate of this subpopulation exceeded 15 Hz and reached its peak value (Figure S7). Other neurons in this population, synaptically-distant from these action neurons, were thus less correlated to the action neurons and (consequently) the time of action. Hence, they did not represent the action time, but represented the median reward time (Figure S7).

For simulating the optogenetic perturbation, there were two parameters of interest, viz. the fraction of excitatory and inhibitory neurons activated by the laser. Since the fraction of V1 neurons affected by the laser is likely small, we modeled the above fractions to be in the range of 5-20%. Further, we modeled the activation as affecting more inhibitory neurons than excitatory ones. This was for two reasons. First, based on the firing patterns of the neurons recorded during optogenetic stimulation, the predominant effect in the network is the inhibition of activity, as is observed immediately following laser offset (Figure S6D-K) (the enhancement of activity while the laser is on does not indicate whether the laser has a net excitatory or inhibitory effect on the network as neurons expressing ChR2 will be activated by the laser regardless of whether they are excitatory or inhibitory). Second, the laser power

decreases rapidly over depth of the cortex and hence, the strongest activation would be in the superficial layers. In layer 1, ~90% neurons are inhibitory(Fitzpatrick et al., 1987; Hendry et al., 1987; Meinecke and Peters, 1987). Within the excitatory populations (E and P), we assumed an equal split of neurons getting activated. The identity of the neurons getting activated within each subpopulation was randomized. We found that if inhibitory activation was twice that of excitatory activation, the network was nearly balanced with the resulting shift in wait times being minimal. Deviations around this ratio caused shifts in wait times similar to those observed experimentally (see Figure S7H).

17. **Analysis of reward magnitude responses:**
Reward responses were quantified by testing the spike count distribution of one of three possible bins, all starting 200 ms after the action (reward onset) and ending at one of 300, 400 or 500 ms. All p values were corrected for multiple comparisons using Bonferroni correction. For each bin, the spike count distribution on rewarded trials was tested against that on unrewarded trials using AuROC. If this was found to be significant, the response was subsequently tested to determine if there was a significant correlation between the spike counts on a trial and the reward obtained. The neuron was deemed to convey reward magnitude information only if both the above tests were significant. Further, we also tested whether the observed difference between rewarded and unrewarded trials resulted from a difference in licking behavior for these trial types in the corresponding bins. To this end, we quantified the lick count and spike count on each trial and tested whether the change in spike counts between rewarded and unrewarded trials was linearly correlated with the corresponding change in lick counts. This was done using bootstrapping by resampling with replacement from the available trials and generating a sampling distribution of the slope in linear correlation. If the sampling distribution was significantly different from zero, it was deemed that the difference in neural activity between rewarded and unrewarded trials resulted from the difference in lick rates on the corresponding trials. No unit tested showed such a relationship ($p>0.5$).

Supplemental References

Fitzpatrick, D., Lund, J.S., Schmechel, D.E., and Towles, a C. (1987). Distribution of GABAergic neurons and axon terminals in the macaque striate cortex. J. Comp. Neurol. *264*, 73–91.

Hendry, S.H.C., Schwark, H.D., and Jones, E.G. (1987). Numbers and Proportions of GABA-lmmunoreactive Different Areas of Monkey Cerebral Cortex Neurons in. J. Neurosci. *7*, 1503–1519.

Meinecke, D.L., and Peters, a (1987). GABA immunoreactive neurons in rat visual cortex. J. Comp. Neurol. *261*, 388–404.