

Published in final edited form as:

J Neurosci. 2012 May 23; 32(21): 7311–7315. doi:10.1523/JNEUROSCI.0027-12.2012.

Single-neuron responses in the human nucleus accumbens during a financial decision-making task

Shaun R. Patel^{1,3,*}, Sameer A. Sheth^{1,*}, Matt K. Mian¹, John T. Gale¹, Benjamin D. Greenberg⁴, Darin D. Dougherty², and Emad N. Eskandar¹

¹Nayef Al-Rodhan Laboratories, Department of Neurosurgery, Massachusetts General Hospital, Harvard Medical School, Boston, MA

²Department of Psychiatry, Massachusetts General Hospital, Harvard Medical School, Boston, MA

³Department of Anatomy and Neurobiology, Boston University School of Medicine, Boston, MA

⁴Department of Psychiatry, Butler Hospital and Brown University, Providence, RI

Abstract

Linking values to actions and evaluating expectations relative to outcomes are both central to reinforcement learning and are thought to underlie financial decision-making. However, neurophysiology studies of these processes in humans remain limited. Here, we recorded the activity of single human nucleus accumbens neurons while subjects performed a gambling task. We show that the nucleus accumbens encodes two signals related to subject behavior. First, we find that under relatively predictable conditions single neuronal activity predicts future financial decisions on a trial-by-trial basis. Interestingly, we show that this activity continues to predict decisions even under conditions of uncertainty (e.g. when the probability of winning and losing is 50/50 and no particular financial choice predicts a rewarding outcome). Furthermore, we find that this activity occurs, on average, two seconds before the subjects physically manifests their decision. Second, we find that the nucleus accumbens encodes the difference between expected and realized outcomes, consistent with a prediction error signal. We show this activity occurs immediately after the subject has realized the outcome of the trial and is present on both the individual and population neuron level. These results provide human single neuronal evidence that the nucleus accumbens is integral in making financial decisions.

INTRODUCTION

Many studies have implicated the midbrain dopaminergic system in encoding a prediction error signal that identifies differences between expectations and outcomes (Montague et al., 1996; Schultz et al., 1997; Pagnoni et al., 2002; O'Doherty et al., 2004; Pessiglione et al., 2006; Day et al., 2007; Zaghoul et al., 2009). These findings have supported the development of reinforcement learning models that maintain historical information about rewarding actions and predictions on future rewarding states (for review see, ((Montague et al., 2004))). Together, these systems are thought to promote learning by selecting motor behaviors based on current sensory information and evaluating outcomes relative to internal goal states. As a result, behaviors that are rewarded are favored over those that lead to no

Corresponding Author: Emad N. Eskandar, Massachusetts General Hospital, 55 Fruit St. White 502, Boston, MA 02114, eeskandar@partners.org.

*These authors contributed equally to this work.

Conflict of Interest: The authors report no conflict of interest.

reward. The nucleus accumbens (NAc) is often thought of as playing a critical role in learning and motivation because of its rich connectivity with midbrain dopaminergic neurons, prefrontal, and limbic areas (Ikemoto and Panksepp, 1999; Joel and Weiner, 2000; Schultz, 2000; Graybiel, 2005; Flagel et al., 2011). As an extension, dysfunctions of the NAc are implicated in conditions such as major depressive disorder (MDD), obsessive-compulsive disorder (OCD), addiction, and others (Gao et al., 2003; Giacobbe and Kennedy, 2006; Greenberg et al., 2006).

Here, we recorded single-neuronal responses from the NAc of eight human subjects undergoing planned deep brain stimulation surgery for the treatment of OCD or MDD. Microelectrode recordings constitute a routine part of deep brain stimulation surgery. Experiments were approved by the Massachusetts General Hospital Institutional Review Board, bore no connection to clinical decisions regarding the appropriateness or type of surgery, and posed no additional risk. Subjects could stop participating at any time before or during the procedure. To examine the role of the NAc in evaluating financial decisions, we designed a behavioral task allowing us to explore two critical components: (i) binding of predicted stimulus value to action, and (ii) evaluating differences between expectation and outcome.

MATERIALS AND METHODS

Study subjects

We recruited eight subjects undergoing planned deep brain stimulation surgery for the treatment of MDD (5 subjects) or OCD (3 subjects) for participation in this study. Six subjects were male and two were female, with mean ages of 47 and 37, respectively. Each individual was evaluated and considered for surgery by a multidisciplinary team of neurologists, neurosurgeons, and psychiatrists. Once approved and scheduled for surgery, an independent member of the research team approached each patient to describe this research study. At that time, the risks and benefits were clearly explained to the subject. All study subjects enrolled voluntarily and provided informed consent under guidelines approved by the Massachusetts General Hospital Institutional Review Board. All subjects were free to withdraw from the study at any time, including during surgery, without consequence to their clinical care.

Microelectrode recordings are performed routinely during deep brain stimulation surgery. Microelectrode wires are lowered to target brain regions prior to implantation of the permanent stimulating electrode. This procedure allows the neurosurgeon to verify the target brain region based on the physiological properties of individual neurons, in order to optimally position the permanent stimulating electrode. Thus, the only modification to the surgical procedure related to this study was the addition of the behavioral paradigm (Gale et al., 2011).

Task Presentation

The behavioral task was performed intraoperatively during microelectrode recordings. A computer monitor was affixed to an adjustable arm mounted to the operating table and positioned within comfortable viewing distance of the subject. A button box was similarly mounted near the subject's right hand. Subjects were in a comfortable semi-recumbent position as is standard for these procedures. The behavioral task was presented using custom written software in Matlab (Mathworks, Natick, MA), Monkeylogic (Asaad and Eskandar, 2008a; 2008b).

The behavioral task was analogous to the classic card game 'War' (Fig 1A). The subject and computer opponent were each dealt a single card; the player with the highest card won. To

simplify the game, the deck was limited to five cards – even cards from 2 through 10 of one suit. Cards were drawn randomly with replacement, and duplicates were permitted. The rules were carefully explained to each subject preoperatively, and each was allowed to practice the task for a short duration to demonstrate comprehension. The task required the subject to evaluate his card, determine its value, and place a \$5 or \$20 wager with the goal of maximizing profit. Declining to wager was not permitted. Thus the logical wager when dealt a 10-card was \$20, as that hand would win 80% of the time and draw 20% of the time. Similarly, the logical wager when dealt a 2-card was \$5, as that hand would either lose or draw.

Each trial began with a central fixation point presented for 350 ms, which cued the subject that the trial was about to begin. Next, the subject's card and the back of the opponent's card was displayed for 1000 ms. Subsequently, two red circles appeared, indicating the mapping of each button (left and right) to its respective wager (\$5 and \$20). The button map was oriented randomly such that the \$5 and \$20 wagers were assigned to the left and right buttons equally. The appearance of the button map served as the go-cue, indicating permission to register a wager. The time interval between appearance of the go-cue and button press was considered the reaction time. A maximum of 5000 ms was allotted for the subject to make a wager. Following the button push, there was a randomized delay period of 250–500 ms, which was immediately followed by the appearance of the subject's and revealed opponent's card for 1000–1250 ms. The outcome of the hand was thus first realized at this time. Lastly, feedback was given by explicitly displaying the amount they won or lost, or "Draw".

Electrophysiology

Intraoperative microelectrode recordings were performed using tungsten microelectrodes (FHC, Bowdoin, ME, or Alpha-Omega Engineering, Nazareth, Israel) with typical impedances of approximately 1 M Ω . The electrodes were advanced into the ventral striatum (Fig. 1B) using a motorized microdrive (Alpha-Omega Engineering, Nazareth, Israel) with 0.01 mm precision. Analog data were band-pass filtered between 300 Hz and 6.5 kHz, recorded at 20 kHz by a PowerLinc 1401 acquisition system (CED, Cambridge, England), and stored for post-hoc analysis. The neurophysiology data were sorted into individual neuronal clusters using waveform principal component analysis (Offline Sorter, Plexon, Houston, TX). Spiking data were filtered for isolation quality, stability, and signal-to-noise; of the eight subjects, data from seven subjects met the criteria and were used for analysis.

Single-unit and population responses

All analyses were performed using custom written software in the MATLAB programming environment using standardized mathematical and signal processing toolboxes.

For the purpose of visualizing the neuronal time course, continuous time firing rate histograms were computed for each neuron by applying a 500 ms sliding window with 10 ms steps. All individual responses were computed using the raw spike data (trials with mean firing rates below 0.5Hz were not used). To allow for comparisons between cells, we normalized the neuronal activity during each trial by subtracting the average activity during a 500 ms pre-fixation window. Statistical differences in neuronal activity between trial types were assessed using a two-tailed t-test ($p < 0.05$) for single-unit and population responses in one of two a priori defined windows: 0–500 ms for go-cue and 250–750 ms for feedback analyses. All results are given with their mean and standard error of the mean.

Receiver operating characteristic and bootstrap

A receiver operating characteristic (ROC) analysis was performed on pooled data from all neurons. ROC analysis is a system for performing binary classifications and provides a method for quantifying and visualizing the true-positive vs. false-positive rate. Put another way, it is a method to quantify the extent to which an ideal observer can predict a particular outcome given a neuronal response. We performed a ROC analysis for predicting financial decisions when considering all cards and just the 6-card alone. To compute the ROC, we pooled the activity within a 500 ms time window (following the go-cue) from all neurons during a specific trial type (e.g. all cards or 6-card alone) and divided them into trials in which the subject bet high and bet low. We then rank-ordered the firing rates and plotted the values. A predictive signal would deviate from unity and have an area under the curve (a.u.c) different from 0.5 (Swets, 1996). To examine statistical significance, a bootstrap randomization process was applied by shuffling the association between the trial type and the neuronal data, while maintaining the overall distribution of each trial type, and recomputing the a.u.c value 1000 times. An a.u.c. value from the original computation that fell outside the 95% confidence interval of the permuted set was considered statistically significant.

RESULTS

On average patients performed 2.4 sessions with 118 trials per session. The average reaction time was $1.3 \pm .15$ seconds (Mean \pm S.E.M.; Fig 1C), and subjects bet high in 45% of the trials (Fig 1D). We isolated 19 neurons from 7 subjects with an average of 1.1 neurons/session and a mean firing rate of 9 ± 1.2 Hz (Mean \pm S.E.M.). Each subject contributed between one and four neurons. Given their phasic response pattern and predominance in the striatum (Heimer et al., 1997), it is most likely that we were recording from inhibitory medium-sized spiny projection neurons.

To examine the role of the NAc in binding stimulus value to action, we investigated whether neuronal activity could predict wager direction (high vs. low). We normalized each trial to its baseline activity, pooled all trials, and performed a ROC analysis. We found that during the 500 ms interval following the go-cue, NAc activity significantly predicted whether the subject would bet high or low (a.u.c. = 0.52; $p = 0.03$, randomization test). A single-neuronal example is presented in Fig 2. Furthermore, we found this activity occurred on average $1.9 \pm .02$ seconds (Mean \pm S.E.M.) before the subjects physically manifested their bet. To confirm that this signal was not the result of an underlying change in baseline firing rate, we computed the average firing rate over the entire trial for bet high and bet low trials for all card presentations for each subject and performed a two-tailed t-test. We found no changes beyond those related to high or low bets. Additionally, the number of high and low bet trials were equally represented for all cards (χ^2 test, $p = 0.62$). Furthermore, we know this effect was not simply the result of the motor movement, because the button cues (\$5 and \$20) were randomly mapped to the left and right button. Together, these results describe a neural signal that predicts future financial decisions.

We next investigated how the NAc responds when the stimulus does not predict outcome (i.e. there is no logical basis for making financial decisions), as is often the case in real world financial decisions. To explore this situation, we normalized and pooled all the “6” card trials and performed a ROC analysis (Fig 3A). We chose the 6-card because it does not predict outcome; that is to say, the subjects have an equal chance of winning or losing, and thus their behavior is not based on an obvious expected reward value. We again found that during the 500 ms interval following the go-cue, NAc activity significantly predicted whether the subject would bet high or low (ROC, a.u.c. = 0.62; $p = 0.003$, randomization test; Fig 3B). We also found that this activity occurred well before ($2 \pm .04$ seconds; Mean \pm

S.E.M) the subjects expressed their bet. The number of high and low bet trials were equally represented (χ^2 test, $p = 0.34$). Additionally, the distribution of outcomes (wins and losses) was equal on previous trials for the 6-card trials in which the bet was high (χ^2 test, $p = 0.76$) and low (χ^2 test, $p = 0.21$), demonstrating that the behavioral signal was not a simple reflection of the outcome on the previous trial. These results demonstrate a predictive behavioral signal under conditions of uncertainty.

To examine the role of the NAc in encoding a prediction error (PE) signal, we divided trials based on subject expectation (unexpected or expected) and the actual outcome (positive or negative; Fig 4A). For each neuron, we computed continuous firing rate histograms for unexpected positive and negative trials within the outcome epoch (250–750 ms after outcome). This is the time when the opponent's card is revealed and the first time that the subject knows the outcome of his hand. We found that 21% of the neurons significantly modulated (two-tailed t-test, $p < 0.05$) for unexpected positive and negative trials. Specifically, we found that NAc activity was potentiated for unexpected positive trials ($PE > 0$) and attenuated for unexpected negative trials ($PE < 0$). To examine population level changes, we computed a population average by normalizing trials from each neuron to its baseline activity. The effect was also found on the population level over the same interval (two-tailed t-test, $p = 0.03$; Fig 4B). We found no significant change in activity for expected positive or negative trials ($PE = 0$) over the same interval (Fig 4C). These data provide evidence for a prediction error signal that may drive adaptive behavior by identifying differences in expectations and outcomes.

DISCUSSION

To the best of our knowledge, these data represent the first single-unit recordings from the human nucleus accumbens during a financial decision-making task. Previously, much of the work exploring the role of the NAc in computing financial decisions has been done using functional imaging (Knutson et al., 2001; 2005; Knutson and Bossaerts, 2007). Using event-related imaging methods, researchers have been able to more closely study the 'where' and 'when' of financial computations. Our study attempts to extend the resolution even further by examining individual neuronal responses during these computations. In this study, we report three important contributions to the understanding of financial decision-making.

First, NAc activity predicts future financial decisions on a trial-by-trial basis. Possible interpretations of the observed predictive signal include encoding of reward expectancy, risk-taking, or probability estimation. Because subjects' behavior was fairly stereotyped (they tended to bet high for 8 and 10 cards and low for 2 and 4 cards) it is conceivable that the observed signal reflects a reward expectancy (i.e. prediction of upcoming reward) rather than future financial decisions. To investigate this possibility, we analyzed data restricted to trials wherein the subject was dealt a 6-card, which provides no expectation regarding the outcome, as there is a 50/50 chance of winning or losing. In this situation, we found that neuronal activity continued to signal future financial decisions, even when there is no clear expectation of reward. The same argument applies for probability estimation, which in this task was essentially a proxy for reward expectation. There is little evidence to support the idea that the observed signal reflects risk-taking; subjects made appropriate choices, betting low for low cards, high for high cards, and evenly for 6-cards.

Second, we characterized the temporal evolution of the predictive behavioral signal. Interestingly, we found that the activity occurred roughly two seconds before the decision was physically manifested--a fairly large latency for neurophysiological signals. The early onset of this signal during the go-cue is consistent with the motivational role of the NAc, but may also reflect the influence of cortical information streams from frontal networks.

Third, we found single neuronal evidence for a prediction error signal in the NAc. Our findings parallel previously reported prediction error signals in the human substantia nigra (Zaghloul et al., 2009). Zaghloul et al. reported peak prediction error activity roughly 250 ms into the outcome period in the substantia nigra, whereas we found peak activity at 450 ms in the NAc. This difference in latency is consistent with the time required for transmission of dopamine from midbrain dopaminergic neurons to the ventral striatum.

One notable characteristic of the PE signal was that the increase in firing for positive PE responses was proportionally greater than the decrease in firing for negative PE responses. This feature may reflect a basic property of the NAc: that it is more responsive to rewards than losses. Alternatively, it may be related to the known dynamic range of dopaminergic neurons. The baseline firing rate of these neurons is about 5–7 Hz. Phasic firing rates can rise as high as 30–40 Hz, but these neurons can only reduce their firing rate to zero. Hence this floor effect is another possible mechanism for the observed asymmetry in NAc PE responses.

A fundamental limitation of this and other human single unit physiology studies is that we are restricted to recording neuronal activity from a structure implicated in the subjects' pathology. Nonetheless, subjects' behavior was rational in this simple task, and the results are consistent with the neurophysiological literature in animal models and functional imaging literature in humans. Therefore, if any difference exists between these subjects and healthy controls it is likely to be more of a small quantitative difference rather than a prominent qualitative difference. These findings serve as a testament not only to the commonality of these responses across species, but also to the importance of these computations in the human brain, since they can be seen even under pathological conditions.

Acknowledgments

This work was supported by grants from the National Science Foundation (IOB 0645886), the National Institutes of Health (NEI 1R01EY017658-01A1, NIDA 1R01NS063249, NIMH Conte Award MH086400, R25NS065743), the Klingenstein Foundation and the Howard Hughes Medical Institute. We thank W. F. Asaad and L. J. Toth for their helpful comments and discussion.

REFERENCES

- Asaad WF, Eskandar EN. A flexible software tool for temporally-precise behavioral control in Matlab. *J Neurosci Methods*. 2008a; 174:245–258. [PubMed: 18706928]
- Asaad WF, Eskandar EN. Achieving behavioral control with millisecond resolution in a high-level programming environment. *J Neurosci Methods*. 2008b; 173:235–240. [PubMed: 18606188]
- Day JJ, Roitman MF, Wightman RM, Carelli RM. Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. *Nat Neurosci*. 2007; 10:1020–1028. [PubMed: 17603481]
- Flagel SB, Clark JJ, Robinson TE, Mayo L, Czuj A, Willuhn I, Akers CA, Clinton SM, Phillips PEM, Akil H. A selective role for dopamine in stimulus-reward learning. *Nature*. 2011; 469:53–57. [PubMed: 21150898]
- Gale JT, Martinez-Rubio C, Sheth SA, Eskandar EN. Intra-operative behavioral tasks in awake humans undergoing deep brain stimulation surgery. *J Vis Exp*. 2011
- Gao G, Wang X, He S, Li W, Wang Q, Liang Q, Zhao Y, Hou F, Chen L, Li A. Clinical Study for Alleviating Opiate Drug Psychological Dependence by a Method of Ablating the Nucleus accumbens with Stereotactic Surgery. *Stereotact Funct Neurosurg*. 2003; 81:96–104. [PubMed: 14742971]
- Giacobbe P, Kennedy SH. Deep brain stimulation for treatment-resistant depression: a psychiatric perspective. *Curr Psychiatry Rep*. 2006; 8:437–444. [PubMed: 17094923]

- Graybiel AM. The basal ganglia: learning new tricks and loving it. *Curr Opin Neurobiol.* 2005; 15:638–644. [PubMed: 16271465]
- Greenberg BD, Malone DA, Friehs GM, Rezai AR, Kubu CS, Malloy PF, Salloway SP, Okun MS, Goodman WK, Rasmussen SA. Three-year outcomes in deep brain stimulation for highly resistant obsessive-compulsive disorder. *Neuropsychopharmacol.* 2006; 31:2384–2393.
- Heimer L, Alheid G, de Olmos J, Groenewegen H, Haber SN, Harlan R, Zahm D. The accumbens: beyond the core-shell dichotomy. *J Neuropsych Clin N.* 1997; 9:354–381.
- Ikemoto S, Panksepp J. The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking. *Brain Res Rev.* 1999; 31:6–41. [PubMed: 10611493]
- Joel D, Weiner I. The connections of the dopaminergic system with the striatum in rats and primates: an analysis with respect to the functional and compartmental organization of the striatum. *Neuroscience.* 2000; 96:451–474. [PubMed: 10717427]
- Knutson B, Adams CM, Fong GW, Hommer D. Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J Neurosci.* 2001; 21:RC159. [PubMed: 11459880]
- Knutson B, Bossaerts P. Neural antecedents of financial decisions. *J Neurosci.* 2007; 27:8174–8177. [PubMed: 17670962]
- Knutson B, Taylor J, Kaufman M, Peterson R, Glover G. Distributed neural representation of expected value. *J Neurosci.* 2005; 25:4806–4812. [PubMed: 15888656]
- Montague PR, Dayan P, Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci.* 1996; 16:1936–1947. [PubMed: 8774460]
- Montague PR, Hyman SE, Cohen JD. Computational roles for dopamine in behavioural control. *Nature.* 2004; 431:760–767. [PubMed: 15483596]
- O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science.* 2004; 304:452–454. [PubMed: 15087550]
- Pagnoni G, Zink CF, Montague PR, Berns GS. Activity in human ventral striatum locked to errors of reward prediction. *Nat Neurosci.* 2002; 5:97–98. [PubMed: 11802175]
- Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature.* 2006; 442:1042–1045. [PubMed: 16929307]
- Schultz W. Multiple reward signals in the brain. *Nat Rev Neurosci.* 2000; 1:199–207. [PubMed: 11257908]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science.* 1997; 275:1593–1599. [PubMed: 9054347]
- Swets, JA. Signal detection theory and ROC analysis in psychology and diagnostics: Collected papers. Lawrence Erlbaum Associates, Inc.; 1996.
- Zaghloul KA, Blanco JA, Weidemann CT, McGill K, Jaggi JL, Baltuch GH, Kahana MJ. Human substantia nigra neurons encode unexpected financial rewards. *Science.* 2009; 323:1496–1499. [PubMed: 19286561]

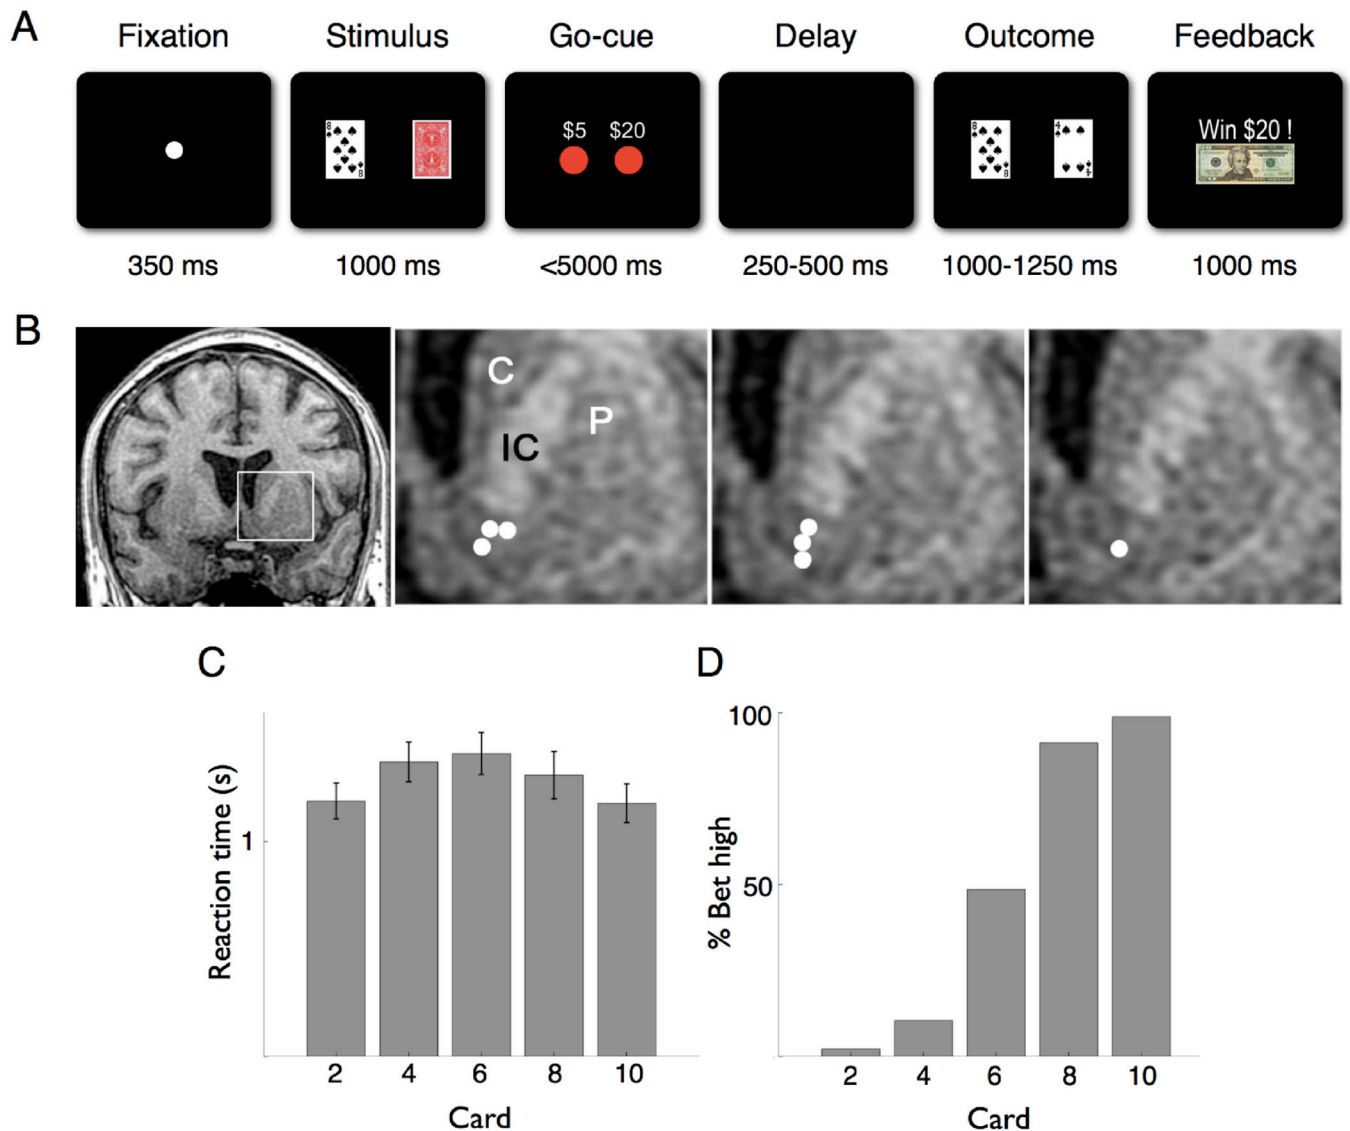


Fig 1. Behavioral task, recording location, and subject performance. **(A)** The Fixation screen was followed by the Stimulus screen showing the subject's card and the back of the opponent's card. Appearance of the wager choices (\$5 or \$20) signified the Go-Cue. After a variable delay period, the Outcome screen showed both the subject's and revealed opponent's card. The Feedback screen then explicitly displayed how much was won or lost, or whether the hand was a draw. **(B)** Coronal T1-weighted MRI showing the recording sites from the seven subjects in whom individual neurons were isolated. The first panel shows a slice 2 mm anterior to the anterior commissure (AC). The following three panels show the inset area (left striatum) 2, 3, and 4 mm anterior to the AC, respectively, with recording sites in the ventral striatum denoted by white circles. C, caudate; P, putamen; IC, anterior limb of internal capsule. **(C)** Average reaction time (in seconds) of all subjects by card value. Error bars are S.E.M. **(D)** Percentage of high wager (\$20) trials by card value.

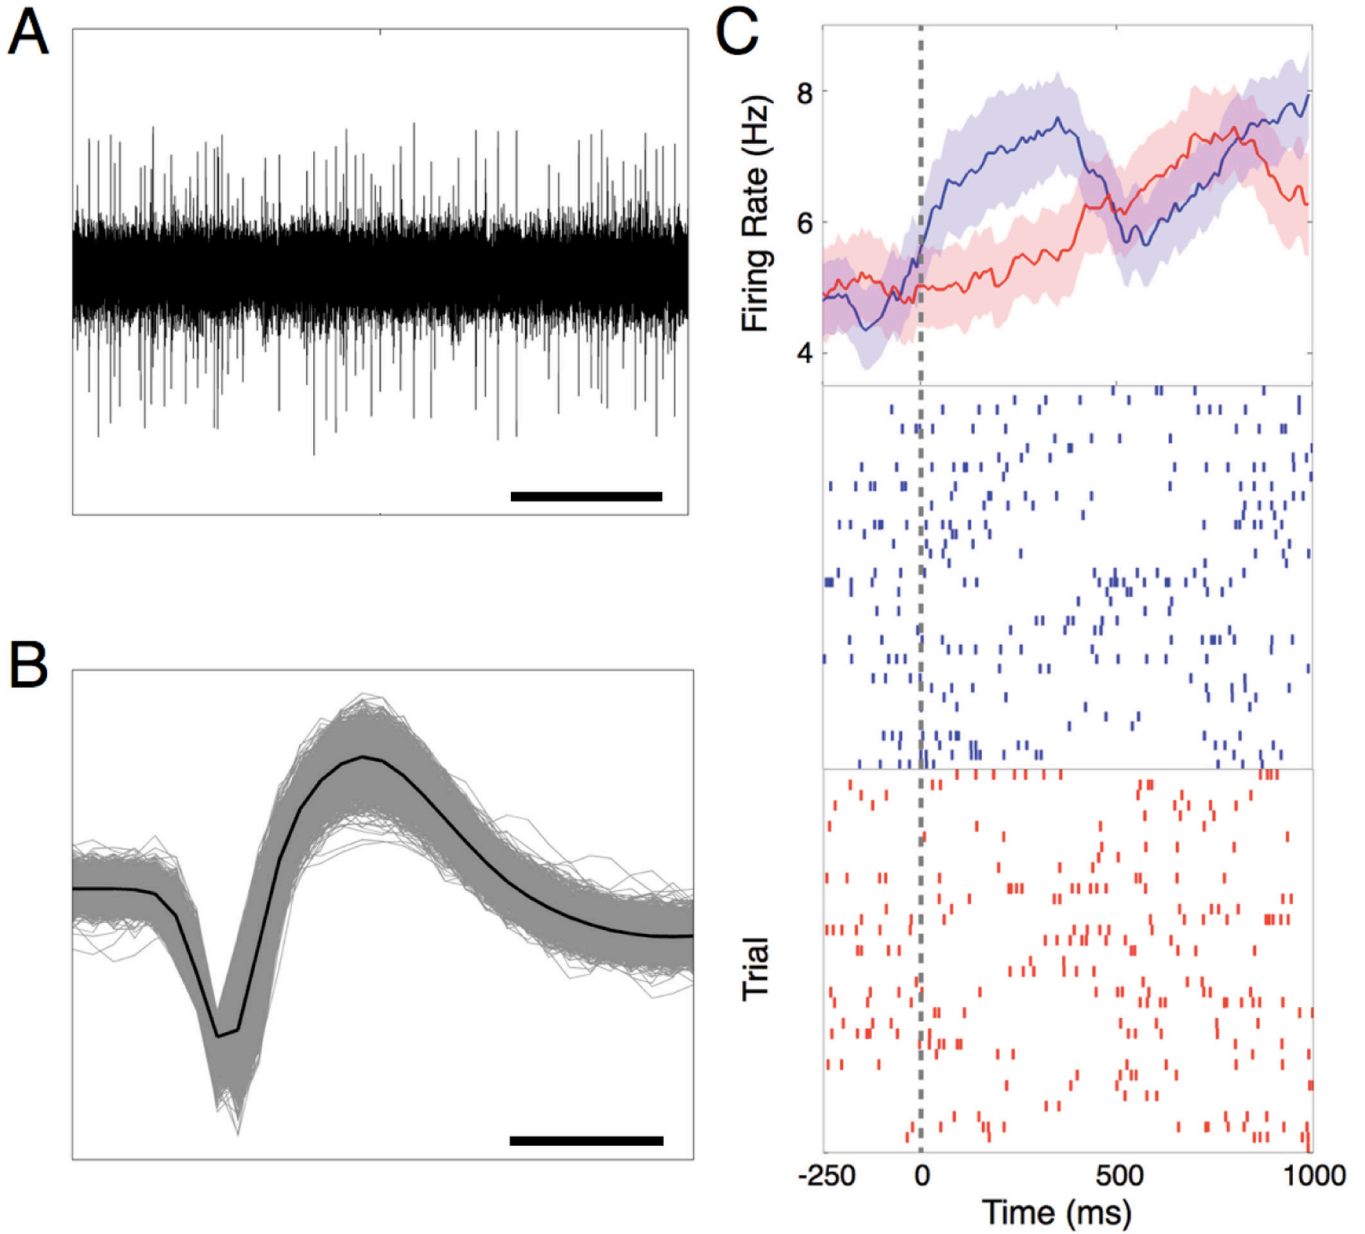


Fig 2. (A) Example of a raw analog voltage trace from one recording session (bar depicts 0.5 seconds). (B) Example of waveform from one isolated neuron (bar depicts 0.5 ms). (C) Single-neuron example of nucleus accumbens activity predicting wager at the go-cue. Continuous firing rate histogram (top panel) and raster plots (middle and lower panels) of all trials in which the subject bet high (blue) and bet low (red) relative to the go-cue (gray dashed line). For each plot, the y-axis represents trial number while the x-axis represents time. Each hashmark represents a single action potential.

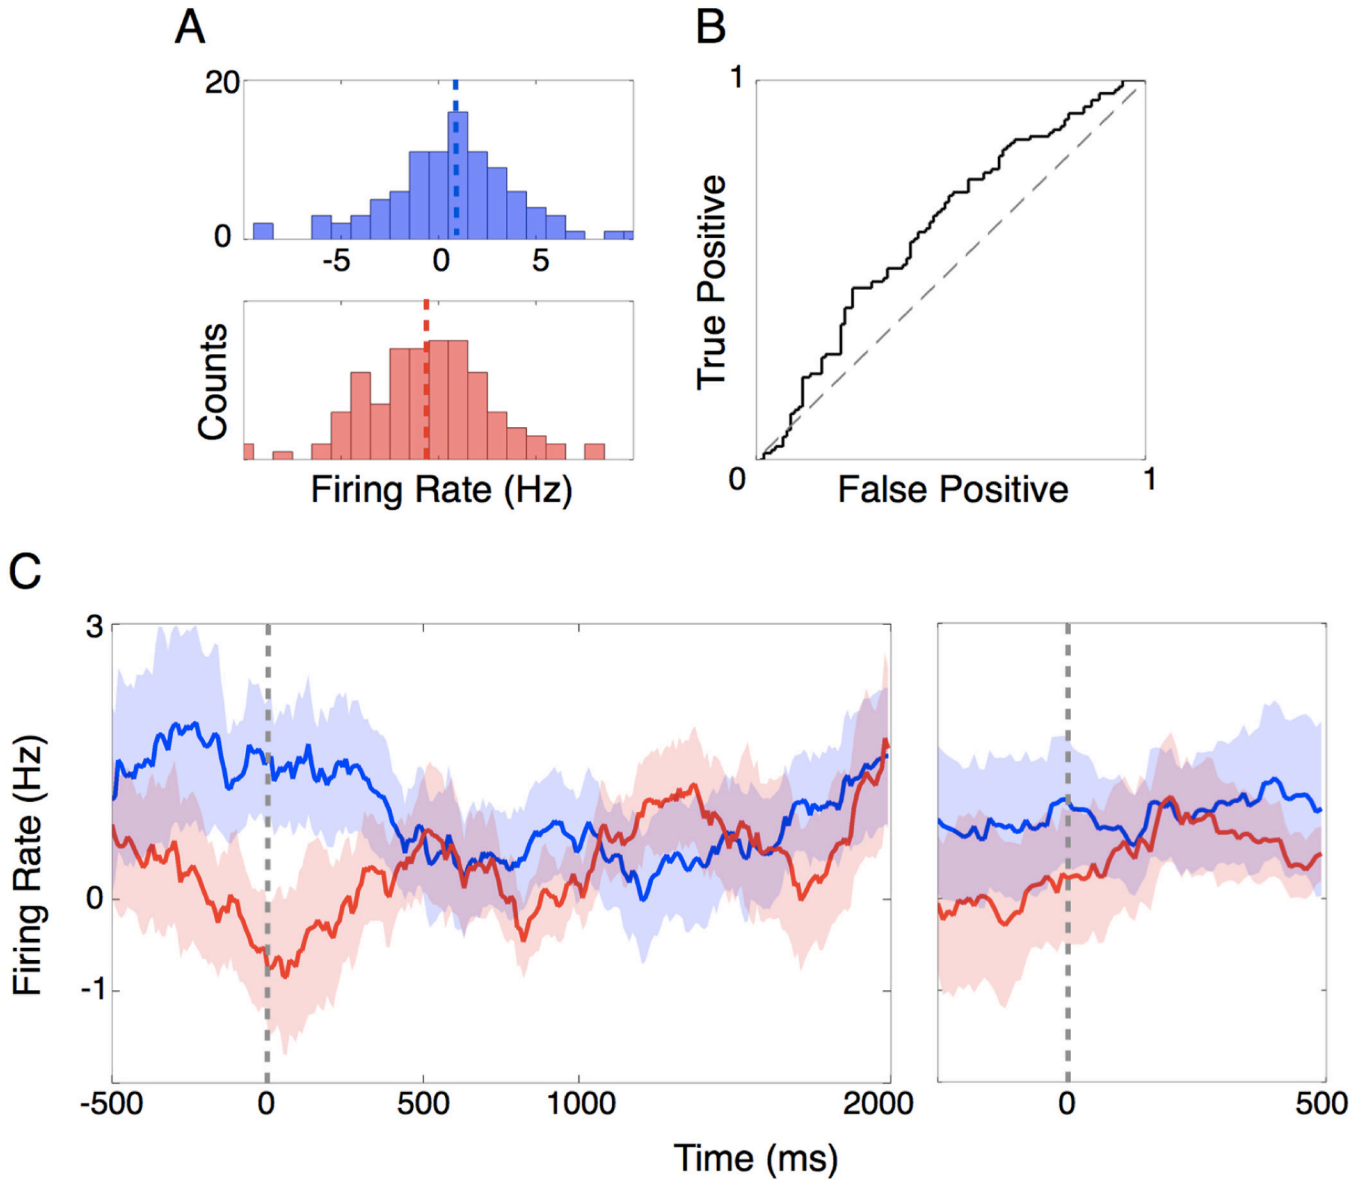


Fig 3. Population response of NAc activity predicts wager at the go-cue for 6-card trials. Each trial was normalized to its pre-fixation baseline activity. **(A)** Histogram of average nucleus accumbens activity during the first 500 ms following the go-cue for trials in which the subject bet high (blue) or low (red). The mean activities for bet high and low trials are represented by the blue and red dashed lines, respectively. Axis limits are the same for the top and bottom panels. **(B)** ROC curve for bet high and bet low trials during the first 500 ms following the go-cue (ROC, a.u.c. = 0.62; $p = 0.003$, randomization test). **(C)** Continuous firing rate histograms for bet high (blue) and bet low (red) trials relative to go-cue (left panel) and choice (right panel).

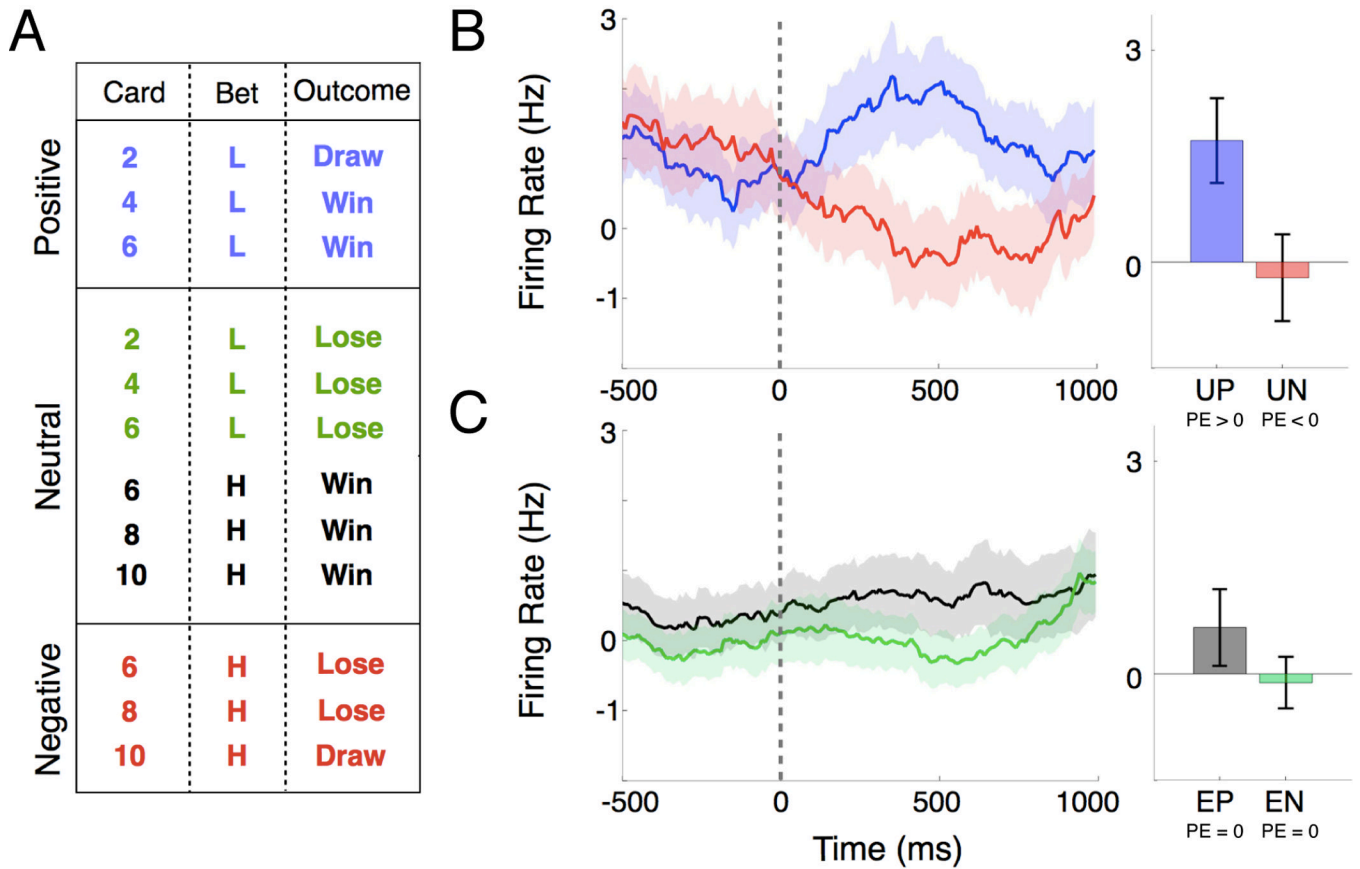


Fig 4. Population nucleus accumbens activity encodes a prediction error signal. Each trial was normalized to its pre-fixation baseline activity. **(A)** Chart of trials composing positive, negative, and neutral prediction error, based on card identity (2, 4, 6, 8, 10), bet (H = high, L = low), and outcome (Win, Lose, Draw). **(B)** Continuous firing rate histogram (left panel) relative to the outcome epoch for PE>0 (unexpected positive; blue) and PE<0 (unexpected negative; red) trials. Bar plot (right panel) depicting mean activity over the 250 to 750 ms interval following the outcome for PE>0 (unexpected positive; UP; blue) and PE<0 (unexpected negative; UN; red) trials (two-tailed t-test, $p = 0.03$). **(C)** Continuous firing rate histogram (left panel) relative to the outcome for the two situations in which PE=0: expected positive (green) and expected negative (brown) trials. Bar plot (right panel) depicting mean activity over the 250 to 750 ms interval following the outcome for expected positive (EP; green; PE=0) and negative (EN; brown; PE=0) trials (two-tailed t-test, $p = 0.24$). Error bars indicate S.E.M.