

Multi-areal neural dynamics encode human decision making

Author list: Jaqueline A. Overton^{1,2}, Karen Moxon³, Matthew P. Stickle³, Logan M. Peters³, Jack J. Lin¹, Edward F. Chang⁴, Robert T. Knight^{5,6}, Ming Hsu⁵, Ignacio Saez^{2,7,8}.

¹Dept. of Neurological Surgery, UC Davis; ²Dept. of Neuroscience, Icahn School of Medicine at Mount Sinai; ³Dept. of Biomedical Engineering, UC Davis; ⁴Dept. of Neurosurgery, UCSF; ⁵Helen Wills Neuroscience Institute, UC Berkeley; ⁶Department of Psychology, UC Berkeley; ⁷Dept. of Neurosurgery, Icahn School of Medicine at Mount Sinai; ⁸Dept. of Neurology, Icahn School of Medicine at Mount Sinai.

Abstract

Value-based decision-making involves multiple cortical and subcortical brain areas, but the exact nature of neural activity underlying choice behavior has been difficult to parse out in the human brain. Here, we use intracranial recordings from neurosurgical patients to decode trial-by-trial choice, showing that information related to expressed choices is contained in high frequency bands (gamma, high frequency activity) and is distributed across multiple brain regions, suggesting that distributed processes underlie human decisions under uncertainty. Furthermore, our results show that reward-based choices can be robustly decoded on a trial-by-trial basis from anatomically distributed iEEG signals.

Main

Decision-making involves the coordinated activity of multiple brain areas. Earlier accounts proposed compartmentalized and sequential roles to individual brain regions¹⁻³, such as valuation and comparison in orbitofrontal cortex (OFC)⁴⁻⁶, and action selection in lateral prefrontal cortex (LPFC)⁷. Animal studies however support the notion that neural activity supporting decisions appears at similar timescales across multiple prefrontal and other regions^{8,9}, indicating that coordinated activity within and across regions is implicated in generating choices. Thus, recent accounts describe decision-related neural activity as distributed, circuit-level processes^{10,11}. However, whether neurophysiological activity related to choice arises simultaneously or sequentially in the human brain has been difficult to ascertain in the human brain, partially given the difficulty of obtaining simultaneous, localized recordings of neural activity across the multiple implicated brain regions at a sufficient temporal resolution to study fast decision processes^{2,3,12}. Here, we leverage human intracranial electroencephalography (iEEG) recordings to analyze distributed neural activity during a decision-making game. iEEG provides high-quality (high signal-to-noise, high temporal resolution) neurophysiological recordings simultaneously from several cortical and subcortical areas¹³.

We recorded intracranial electroencephalography (iEEG) data from 36 medication refractory epilepsy patients while they played a gambling task¹⁴, 20 of which resulted in data of sufficient quality and were included in final analyses (see Methods). Patients made trial-by-trial choices between a safe prize and a risky gamble with varying win probabilities (see Fig. 1a). Patients gambled more often in trials with higher win probability, as expected in reward-maximizing behavior^{15,16} (Fig. 1b and Extended Data Fig. 1). Patients underwent either electrocorticography, providing subdural coverage predominantly in frontoparietal

regions (ECoG, 9/20 patients, Fig. 1c right), or stereotactic EEG (sEEG), predominantly in deep temporal lobe regions (amygdala, hippocampus, insula; 11/20 patients; Fig. 1c left and center and Extended Data Fig. 2). We analyzed electrophysiological recordings from electrodes located in regions involved in reward-related behavior: orbitofrontal cortex (OFC), lateral prefrontal cortex (LPFC), cingulate cortex (CC), precentral gyrus (PrG), postcentral gyrus (PoG), parietal cortex (PC), amygdala (Amy), hippocampus (Hipp) and insula (Ins) ($n=1085$ electrodes total; Fig. 1c and Extended Data Table 1).

We analyzed local field potential (LFP) data from all electrodes located in grey matter in the above locations of interest and characterized neural activity in each neural frequency band (delta/ δ [1-4Hz], theta/ θ [4-8Hz], alpha/ α [8-12Hz], beta/ β [12-30Hz], gamma/ γ [30-70Hz], and high-frequency activity/HFA [70-200Hz]). First, for each electrode we determined whether significant power modulation existed in any frequency band during choice by comparing the deliberation epoch (-1 to 0s before choice button press) to baseline (-0.5 to 0s prior to stimulus presentation): 76.8% of electrodes showed significant modulation (either increase or decrease) in at least one frequency band and were termed task-active ($p<0.01$, paired t-test; Fig. 2a, Extended Data Tables 1, 2 and 3). Electrodes were often task-active in multiple frequencies (mean = 2.09 ± 1.70 frequency bands; Extended Data Table 4), with similar proportions of task-active electrodes across frequency bands (mean = $34.9\%\pm 0.05$, Fig. 2b and Extended Data Table 5). The proportion of task-active electrodes was homogeneous across anatomical regions (mean = $77.0\%\pm 0.07$, Fig. 2c, Extended Data Tables 2 and 7), indicating widespread power modulation distributed across frequency bands and brain regions during deliberation.

Task-active electrodes in individual brain regions could show increases or decreases. Therefore, we next examined the proportion of electrodes showing significant increases or decreases in power separately. Overall, 50.5% of electrodes showed power increases and 55.6% showed decreases, with 29.3% of electrodes showing a combination of increases and decreases in separate frequency bands (e.g. an increase in delta accompanied by a decrease in beta; Extended Data Fig. 4). There was significant heterogeneity in the proportion of electrodes that showed power increases/decreases across regions (Fig. 2d and Extended Data Fig. 4). For example, electrodes in Amy/HC predominantly showed power increases, whereas power decreases were most common in PC/PrG (Fig. 2d, 2f, and Extended Data Table 8). Grouping activations by frequency band instead of region shows a complementary depiction of power modulations. For example, power modulation in δ/θ consisted predominantly of power increases, whereas β/γ modulations consisted predominantly of power decreases (Fig. 2g, Extended Data Fig. 4 and Table 9). Finally, HFA was unique in that it showed almost exclusively bidirectional modulation, with individual electrodes within a region showing either increases or decreases (Fig. 2d, Fig 2g, and Extended Data Fig. 4).

Therefore, we observed a rich pattern of region-frequency specific patterns of task-related power modulation, with similar patterns of power modulation across sets of regions. For example, PoG/PrG/PC electrodes predominantly showed power decreases across most frequency bands ($\delta/\theta/\alpha/\beta/\gamma$; average decrease = 35.5%); Amy/HC showed the opposite pattern, with widespread power increases across power bands ($\delta/\theta/\alpha/\beta/\gamma$ /HFA; average increase = 27%), whereas frontal regions (OFC/LPFC/CC), showed concomitant power increases in lower frequencies ($\delta/\theta/\alpha$) and power decreases in higher frequencies (β/γ /HFA). This analysis revealed three sets of regions with similar modulation patterns: prefrontal (OFC/LPFC/CC), frontoparietal (PrG/PoG/PC) and limbic (Amy/HC/Ins) (Fig. 2e). To assess the similarities of power modulation patterns across these hypothesized sets of regions, we sought to classify regions according to their power modulation patterns in an unsupervised way. Specifically, we carried out

dimensionality reduction followed by unsupervised clustering on power modulation patterns by characterizing the percentage of electrodes showing power increases/decreases for each region in each patient, separately for each frequency band (n=104 datapoints, 12 dimensions [6 frequencies x increase/decrease]; see Methods). This analysis significantly separated the regions into three hypothesized functional clusters, corresponding to prefrontal (OFC/LPFC/CC), frontoparietal (PrG/PoG/PC) and limbic (Amy/HC/Ins) circuits, (bootstrapped $p < 0.05$; Fig 2e, Extended Data Table 10), suggesting that patterns of power modulation are distinct across these proposed sub-circuits. In summary, the direction of power modulation (increase/decrease) depended on both frequency and region, with three sets of regions (prefrontal [OFC/LPFC/CC], frontoparietal [PrG/PoG/PC and limbic [Amy/HC/Ins]) showing distinct power modulation patterns.

Power modulations are thought to reflect multiple cognitive processes underlying choice behavior. Next, we investigated which neural features across regions and frequency bands were associated with choices. Specifically, we sought to identify which neural features showed significant differences between safe bet and gamble trials across all electrodes (choice-encoding electrodes; see Methods). Overall, 42.7% of electrodes encoded choice in one or multiple frequency bands (average = 0.55 ± 0.73 frequency bands per electrode; Fig. 3a; Extended Data Tables 11, 12). Choice encoding was more common in higher frequencies (β/γ /HFA, mean = 6.05%/13.26%/19.72%, respectively) than in lower frequencies ($\delta/\theta/\alpha$, mean = 2.76%/2.14%/3.24%, respectively; $p < 10^{-6}$ between average percent encoding in low and high frequencies; Fig. 3b, Extended Data Tables 13, 14). In contrast, the overall proportion of choice-encoding electrodes was largely homogeneous across anatomical regions (average = $37 \pm 7.7\%$), but greatest in frontoparietal regions (52.2%), followed by prefrontal regions (40.5%), and lastly limbic regions (33.1%) (see Fig. 3c). Next, we examined the pattern of choice encoding across anatomical regions of interest and frequency bands (Fig. 3d). We found that the frequency-specific patterns of choice encoding were similar across anatomical regions. Specifically, we observed more choice-encoding electrodes in high-frequencies compared to low frequencies across all regions (Fig. 3d), suggesting a homogeneous and widespread involvement of γ and HFA in choice behavior. These results indicate broadly distributed encoding of choice information across brain regions, particularly in higher frequencies.

Finally, to confirm choice-encoding information in γ and HFA, we sought to build a model to decode trial-by-trial choice from the neural activity in those frequency bands. Specifically, we focused on activity during deliberation (-1 s to 0 prior to button press) across all regions. To reduce the number of dimensions input to the decoder, we tested two dimensionality reduction strategies (principal component analysis [PCA] alone and linear dynamical systems [LDS] after PCA) followed by two classifier strategies (a simple Euclidean distance classifier [ED]¹⁷ or dynamic time warping [DTW]). PCA and ED are standard methods for dimensionality reduction and classification, respectively; we chose to compare them to LDS and DTW which are potentially better suited to neural time series and may result in better decoding accuracy. LDS characterizes ongoing neural dynamics by identifying the underlying latent variables (LVs) that capture time-varying neural activity, and DTW is a distance metric that specifically accounts for temporal variation in neural activity (see Methods). Using a leave-one-out validation strategy, we found that dimensionality reduction using LDS combined with classification using DTW outperformed ED approaches (ANOVA $F(4,19) = 39.1$, $p < 10^{-17}$; Fig. 4a). The optimal decoding strategy, LDS followed by DTW, accurately classified choice $74.3 \pm 3.4\%$ across all patients and produced a maximal trial-by-trial single patient decoding accuracy of 79.8%. Therefore, a low dimensional, linear dynamical model followed by a classifier that accounted for variation in temporal dynamics across trials was the optimal strategy for choice decoding.

Performance of the LDS+DTW was close to the optimal performance across all models (best accuracy = $79.7\% \pm 3.36\%$, $p > 0.01$ Bonferroni corrected, paired t-test). Importantly, LDS+DTW achieved above-chance decoding in all patients, with a minimum decoding accuracy of 65.4%, indicating that LDS+DTW decoder was robust to variation in electrode number and anatomical localization across patients (Fig. 4a, Extended Data Table 15). In addition, we found that decoder performance depended on the win probability of the gamble trials, with the decoder performance falling to chance levels when win probability was either 0% or 100% (i.e. when there was no risk associated with the gamble choice), yet for all risky choices (win probability between 10% and 90%), the decoder performance was similar (Fig. 4b, Extended Data Table 16). Therefore, we excluded trials with no uncertainty (gamble win probability = 0% or 100%) from decoder results (see Extended Tables 16a and 16b). Electrodes with <50 trials after artifact rejection were excluded from analyses.

Each dimension of the LDS model defines a latent variable (LV) that describes how the dynamics unfold, thus, we examined the single-trial neural trajectories identified by the LDS model (Fig. 4c) within the manifolds defined by 3 latent variables (LVs; Fig. 4d) to further probe the underlying dynamics of choice processing. Some of these dimensions capture the choice-related dynamics while others capture dynamics unrelated to choice that may reflect stochastic processes. Selecting the three LVs whose representations during gamble trials were most different from safe bet trials, we observe choice-related dynamics separating during the deliberation phase (Figure 4c). Next, by plotting the state space represented by these 3 LVs (Figure 4d), we can observe that the neural trajectories define a 3-D manifold within which the dynamics traverse during choice. Early in deliberation (-1 s to -500 ms pre-choice), single-trial trajectories overlapped (Fig. 4c and Extended Data Movie 1). As the patient approached their selection (500 ms to 0 s pre-choice), neural trajectories progressed to non-overlapping regions of the state-space representing their final choice, suggesting the existence of separate attractors for gamble and safe bets (Fig. 4d). This settling of the neural dynamics to distinct attractors in the state space presumably reflects the independent information used by the decoder to predict choices.

In summary, by leveraging multi-region sEEG recordings, we show that individual electrodes across a variety of cortical and subcortical brain regions present distinct patterns of power modulation that defined sets of regions (prefrontal, limbic, frontoparietal). Despite heterogeneity in task-related power modulation, all regions similarly encoded choice information in high frequency bands (primarily γ /HFA). These neural features can be leveraged to build decoders capable of classifying trial-by-trial choices in every subject in our dataset, despite differences in anatomical coverage and surgical strategy.

Our results show that multi-frequency band, multi-region power is modulated during deliberation, indicating that these processes occur simultaneously and generate complex patterns of power modulation that can be detected with distributed recordings obtained with sEEG. Sensory processing, motor output, attentional and decision processes are cognitive sub-processes that occur during value-based decision-making and have been previously shown to be reflected in oscillatory processes¹⁸⁻²⁰. For example, we observed widespread beta-band power decreases in frontoparietal contacts (Fig. 2d), possibly reflecting beta desynchronization associated with pre-motor processes²¹. Delta-theta increases were widespread in limbic and prefrontal electrodes, an observation consistent with the appearance of theta-band oscillations in the hippocampus and prefrontal cortex during goal-directed behavior, navigation and memory formation^{22,23}. Therefore, it is likely that widespread low-frequency power modulations reflect the neural basis of these and other cognitive processes. Alternatively, similar patterns of low-frequency power modulation may reflect the establishment of functional communication across regions, in a pattern

reminiscent of “functional fingerprinting” in which frequency-specific activations reflect computations associated with specific cognitive processes²⁴.

Although our task was not designed to parse out all underlying component cognitive processes, we were able to show that power modulation in higher frequencies (γ /HFA) reflected choices. This larger impact of higher frequency bands in choice-encoding is consistent with the notion that broadband gamma activity reflects local neuronal activation related to value-based decision-making^{5,13,25–27}. HFA modulation was bidirectional, reflecting heterogeneity in local spike rate encoding^{25,26,28}. Contrary to the region-specific patterns of frequency power modulation, γ /HFA choice encoding was present in all regions examined (Fig. 3). This observation is consistent with the notion that value-based decision-making is a distributed process engaging multiple brain areas, likely simultaneously, a notion supported by an increasing amount of evidence^{11,29,30} as well as the bi-directional nature of the connections between reward-related brain areas³¹. In our dataset, all regions showed some level of choice selectivity, although the proportion of choice-selective sites was higher in frontoparietal than in medial temporal sites (Fig 3), consistent with the importance of frontoparietal regions in decision-making³². Importantly, we identified an impact of activity in these frequency bands on final choice but did not disambiguate the specific nature of the underlying computations, which could vary by region. For example, OFC activity reflects a map of the relative values of options under consideration^{4,33,34,34}, which is likely to impact choice; conversely other brain regions (e.g. LPFC) could reflect computations that are closer to the actual nature of the choice than to an abstract value space^{5,9}.

Finally, we were able to decode choices from multi-frequency, multi-region neural activity. We found that a combined LDS+DTW strategy produced the best decoding accuracy (Fig. 4). LDS models generated low-dimensional representations (latent variables) of neural activity, which traversed the state space during deliberation before settling on one of two attractors, for either safe bet or gamble decisions (Fig. 4). Interestingly, individual LVs showed rapid switching during deliberation (Fig. 4b), suggesting that both safe bet/gamble attractors were visited multiple times before a choice was made, as the dynamics moved back and forth between states representing gamble and safe choice. This is similar to activity patterns observed during deliberation in multi-electrode OFC decoding in non-human primates that reflect the fast alternative evaluation of binary choices as in our task^{5,35}. However, the trajectory of higher-dimensional representations incorporating multiple LVs showed a progressive separation of gamble and safe bet trajectories (Fig. 4d), suggesting that deliberation in this higher-dimensional space may evolve progressively until a decision-threshold is reached^{36,37}, potentially bridging observations of evidence accumulation and fast alternative switching. Finally, DTW algorithms resulted in the best average decoding performance, indicating decision-encoding across sites and patients had consistent yet time-varying temporal dynamics that impacts the ability of models to decode choices. Intriguingly, we show that the decoder fails to classify choices when win probability was either 0% or 100% (Fig. 4d), suggesting that these choices, which do not imply risk, may have a different neural signature, and that the encoding of choice intent was similar regardless of the difficulty of the choice. The widespread pattern of activity underlying behavior choice-encoding is likely to be crucial for the consistent performance of the choice-decoding algorithms. Indeed, we show that our final LDS+DTW model successfully decoded choice in all patients, despite great differences in coverage (see Fig. 1 and Extended Data Fig. S2), highlighting the importance of leveraging distributed recordings rather than a predetermined set of ROIs.

In summary, we show that choice information is anatomically distributed, present in high (γ /HFA) frequency activity in a set of cortical and subcortical regions, and that decoding of trial-by-trial choices is

robust to variation in electrode placement across patients. The combination of invasive recordings, economic probes of decision-making, and ML approaches will open the door to characterizing circuit-wide activity underlying decision-making behavior. Future efforts will be tasked with development of more general brain-state decoders for other cognitive (i.e. memory, attention) and translational (i.e. pathological vs healthy brain states) contexts.

Acknowledgements

We would like to thank L. Nuñez, C. Meikle, and C. Foreman for help with data collection. We would especially like to thank the patients for their willingness to participate in this research. The project described was supported by the National Institute of Mental Health through grant number K01MH108815, and the National Center for Advancing Translational Sciences, National Institutes of Health, through grant number UL1 TR001860 and linked award TL1 TR001861. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

References

1. Rangel, A., Camerer, C. & Montague, P. R. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience* **9**, 545–556 (2008).
2. Padoa-Schioppa, C. Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci* **34**, 333–359 (2011).
3. Rushworth, M. F. S., Mars, R. B. & Summerfield, C. General mechanisms for making decisions? *Curr Opin Neurobiol* **19**, 75–83 (2009).
4. Padoa-Schioppa, C. & Assad, J. A. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
5. Rich, E. L. & Wallis, J. D. Decoding subjective decisions from orbitofrontal cortex. *Nat Neurosci* **19**, 973–980 (2016).
6. Peters, J. & D’Esposito, M. The drift diffusion model as the choice rule in inter-temporal and risky choice: A case study in medial orbitofrontal cortex lesion patients and controls. *PLoS Comput Biol* **16**, e1007615 (2020).
7. Seo, M., Lee, E. & Averbeck, B. B. Action selection and action value in frontal-striatal circuits. *Neuron* **74**, 947–960 (2012).

8. Kennerley, S. W., Dahmubed, A. F., Lara, A. H. & Wallis, J. D. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* **21**, 1162–1178 (2009).
9. Wallis, J. D. & Miller, E. K. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *European Journal of Neuroscience* **18**, 2069–2081 (2003).
10. Hunt, L. T. *et al.* Triple dissociation of attention and decision computations across prefrontal cortex. *Nat Neurosci* **21**, 1471–1481 (2018).
11. Hunt, L. T. & Hayden, B. Y. A distributed, hierarchical and recurrent framework for reward-based choice. *Nat Rev Neurosci* **18**, 172–182 (2017).
12. Rangel, A., Camerer, C. & Montague, P. R. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience* **9**, 545–556 (2008).
13. Parvizi, J. & Kastner, S. Promises and limitations of human intracranial electroencephalography. *Nat Neurosci* **21**, 474–483 (2018).
14. Saez, I. *et al.* Encoding of Multiple Reward-Related Computations in Transient and Sustained High-Frequency Activity in Human OFC. *Current Biology* **28**, 2889-2899.e3 (2018).
15. Niv, Y. Reinforcement learning in the brain. *Journal of Mathematical Psychology* **53**, 139–154 (2009).
16. Doll, B. B., Simon, D. A. & Daw, N. D. The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology* **22**, 1075–1081 (2012).
17. Foffani, G. & Moxon, K. A. PSTH-based classification of sensory stimuli using ensembles of single neurons. *J Neurosci Methods* **135**, 107–120 (2004).
18. Helfrich, R. F. *et al.* Neural Mechanisms of Sustained Attention Are Rhythmic. *Neuron* **99**, 854-865.e5 (2018).
19. Khanna, P. & Carmena, J. M. Neural oscillations: beta band activity across motor networks. *Current Opinion in Neurobiology* **32**, 60–67 (2015).

20. Schroeder, C. E. & Lakatos, P. Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci* **32**, 9–18 (2009).
21. Kühn, A. A. *et al.* Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. *Brain* **127**, 735–746 (2004).
22. Ekstrom, A. D. *et al.* Human hippocampal theta activity during virtual navigation. *Hippocampus* **15**, 881–889 (2005).
23. Buzsáki, G. & Moser, E. I. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat Neurosci* **16**, 130–138 (2013).
24. Siegel, M., Donner, T. H. & Engel, A. K. Spectral fingerprints of large-scale neuronal interactions. *Nat Rev Neurosci* **13**, 121–134 (2012).
25. Buzsáki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience* **13**, 407–420 (2012).
26. Buzsáki, G. & Wang, X.-J. Mechanisms of Gamma Oscillations. *Annual Review of Neuroscience* **35**, 203–225 (2012).
27. Lopez-Persem, A. *et al.* Four core properties of the human brain valuation system demonstrated in intracranial signals. *Nat Neurosci* **23**, 664–675 (2020).
28. Ray, S. & Maunsell, J. H. R. Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biol* **9**, e1000610 (2011).
29. Smith, E. H. *et al.* Widespread temporal coding of cognitive control in the human prefrontal cortex. *Nat Neurosci* **22**, 1883–1891 (2019).
30. Spitman, M., Seo, H., Lee, D. & Soltani, A. Multiple timescales of neural dynamics and integration of task-relevant signals across cortex. *Proceedings of the National Academy of Sciences* **117**, 22522–22531 (2020).

31. Haber, S. N. & Knutson, B. The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology* **35**, 4–26 (2010).
32. Andersen, R. A. & Cui, H. Intention, Action Planning, and Decision Making in Parietal-Frontal Circuits. *Neuron* **63**, 568–583 (2009).
33. Padoa-Schioppa, C. Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex. *J. Neurosci.* **29**, 14004–14014 (2009).
34. Schuck, N. W., Cai, M. B., Wilson, R. C. & Niv, Y. Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron* **91**, 1402–1412 (2016).
35. Peixoto, D. *et al.* Decoding and perturbing decision states in real time. *Nature* **591**, 604–609 (2021).
36. Krajbich, I. & Rangel, A. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences* **108**, 13852–13857 (2011).
37. Pedersen, M. L., Frank, M. J. & Biele, G. The drift diffusion model as the choice rule in reinforcement learning. *Psychon Bull Rev* **24**, 1234–1251 (2017).

Figures

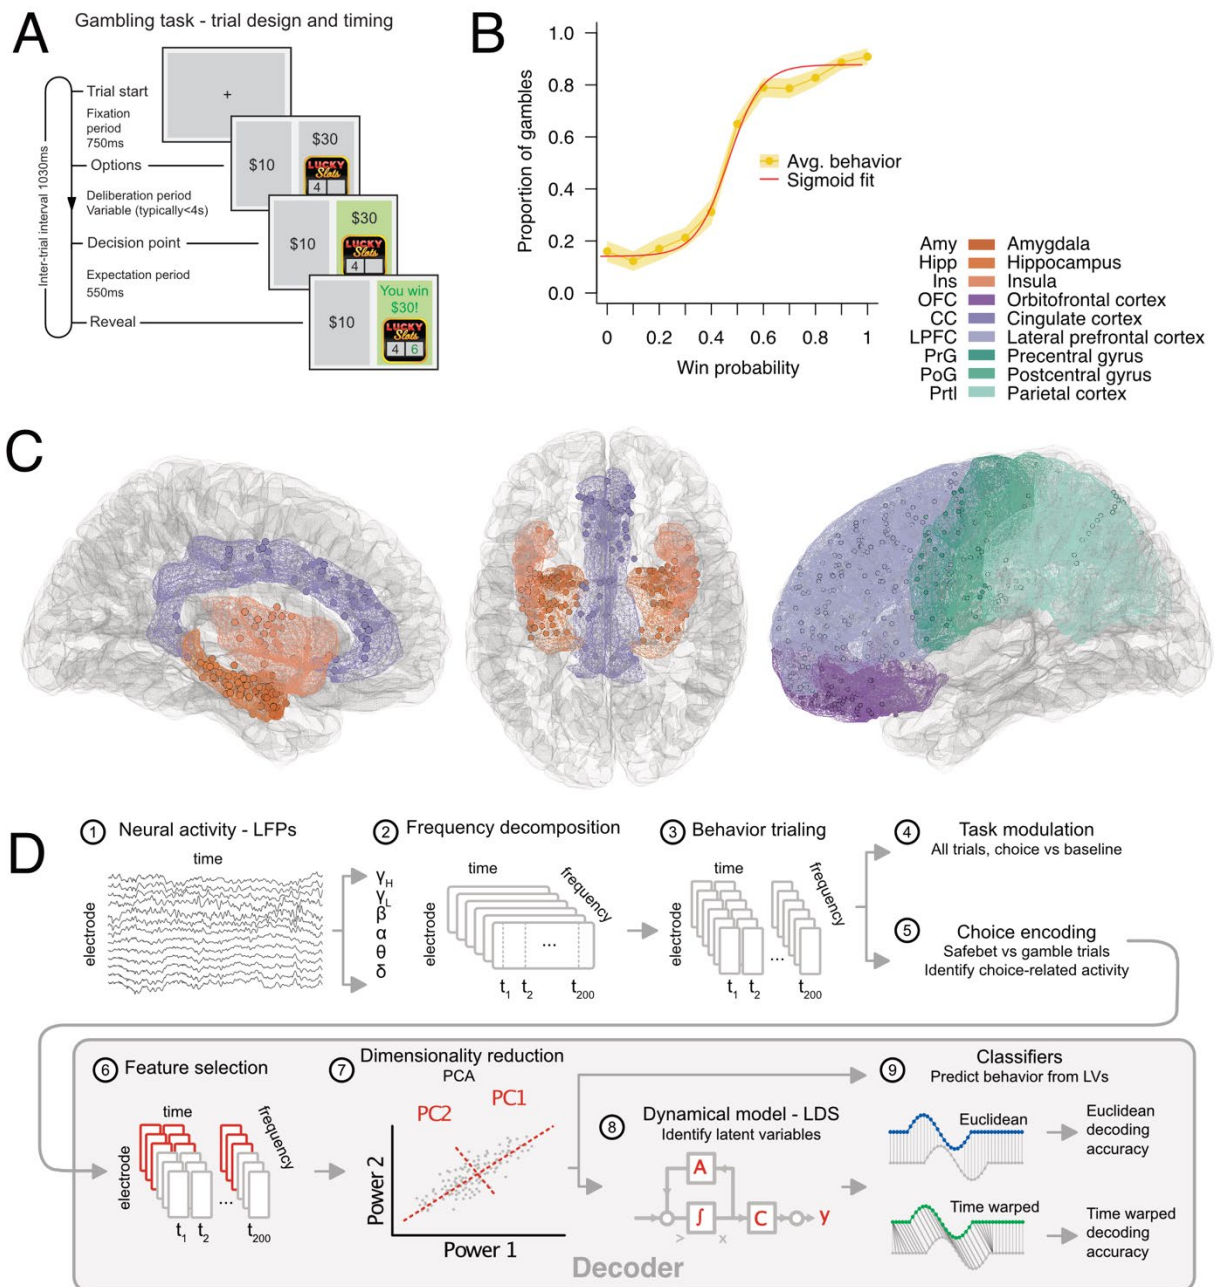


Figure 1. Experimental design and analytical strategy. **a**, Patients played a gambling game in which they chose between a safe \$10 reward and a risky gamble for a higher reward. Gamble win probability varied parametrically on a trial-by-trial basis (0-100%). **b**, Average patient behavior (yellow line, shadow indicates SEM), shown as the proportion of risky choices (gambles) for each win probability, averaged in 10% increments. Patients gambled more often as the win probability increased. Red line represents a sigmoidal fit to the average behavior. **c**, Anatomical intracranial EEG (iEEG) coverage. We collected data from patients that underwent either electrocorticography (n=9, ECoG) or stereotactic EEG (n=11, sEEG) during epilepsy surgery (n=20 patients and 1042 electrodes total). We focused our analyses on electrodes located

in grey matter in regions involved in reward-related behavior: orbitofrontal cortex (OFC), lateral prefrontal cortex (LPFC), precentral gyrus (PrG), cingulate cortex (CC), precentral gyrus (PrG), postcentral gyrus (PoG), parietal cortex (PC), amygdala (Amy), hippocampus (Hipp) and insula (Ins). **d**, Analytical strategy. Local field potentials (LFPs) from each grey matter electrode were cleaned and preprocessed (1), and subsequently decomposed into power across frequency bands (delta [1-4Hz], theta [4-8Hz], alpha [8-12Hz], beta [12-30Hz], gamma [30-70Hz] and high-frequency activity [HFA; 70-200Hz], 2). After trialing around behavioral events of interest (patient choice, 3), we determined which electrodes presented task modulation (significant power change versus baseline, 4) and choice encoding (significantly different power between gamble and safe bet trials, 5) in any frequency band. Choice-encoding features (electrodes and frequencies) were selected (6) and subject to dimensionality reduction (7), followed by dynamical modeling and classification (8-9) or direct classification (9). We assessed model quality by examining the classification accuracy using a leave-one-out procedure.

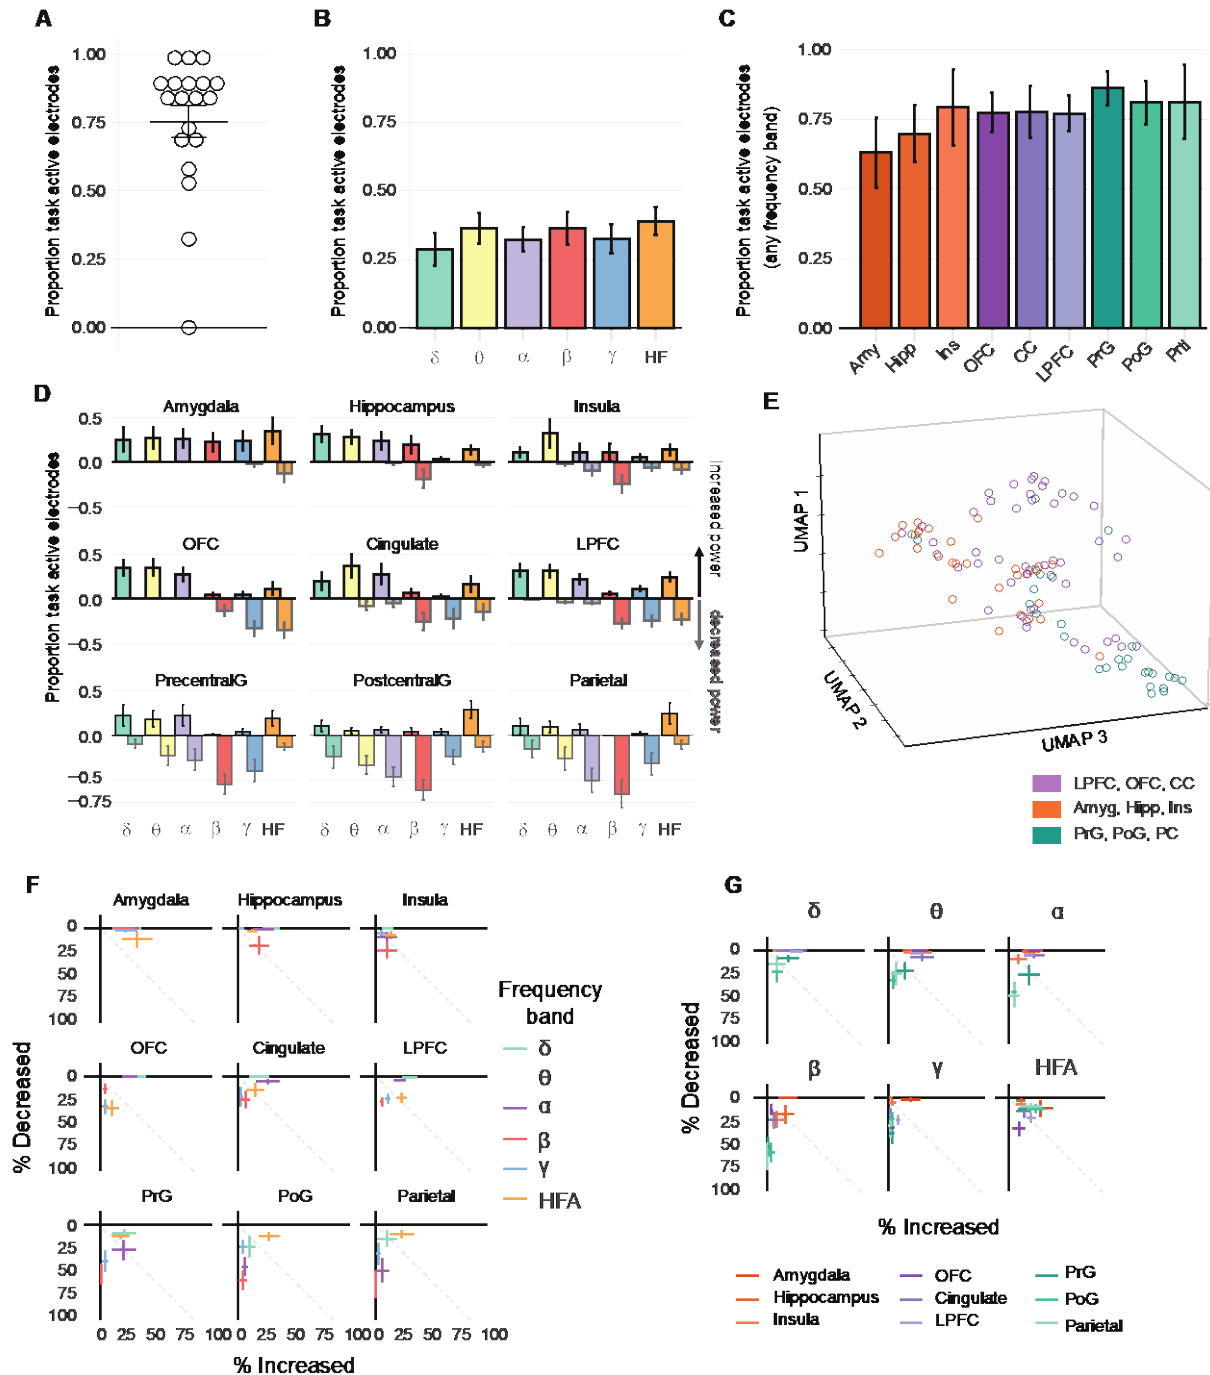


Figure 2. Proportion of task-active electrodes showing significant power modulation during deliberation. Power was compared between deliberation (-1s to 0s pre-choice) and baseline (fixation cross) epochs. **a**, Proportion of task active electrodes per patient. Horizontal lines show mean +/- SEM. **b-c**, Mean proportion of task-active electrodes across patients, grouped by frequency (across all regions, **b**) and anatomical region (across all frequencies, **c**). **d**, Mean proportion of task-active electrodes across patients, grouped by both region and frequency band. The proportion of electrodes that showed significant increases or decreases relative to baseline are shown as positive and negative bars on the

vertical axis, respectively. **e**, Clustering of regions into functional groups: limbic (Amy, HC, Ins), prefrontal (OFC, LPFC, CC) and frontoparietal (PrG, PoG, PC). **f**, same data as in (c), but showing the patient-averaged proportion of electrodes showing increases (y-axis) or decreases (x-axis) separated by regions (panels) and frequency bands (points). **(g)** Same data as in (c) but showing the patient-averaged proportion of electrodes showing increases (y-axis) or decreases (x-axis) separated by frequency bands (panels) and regions (points).

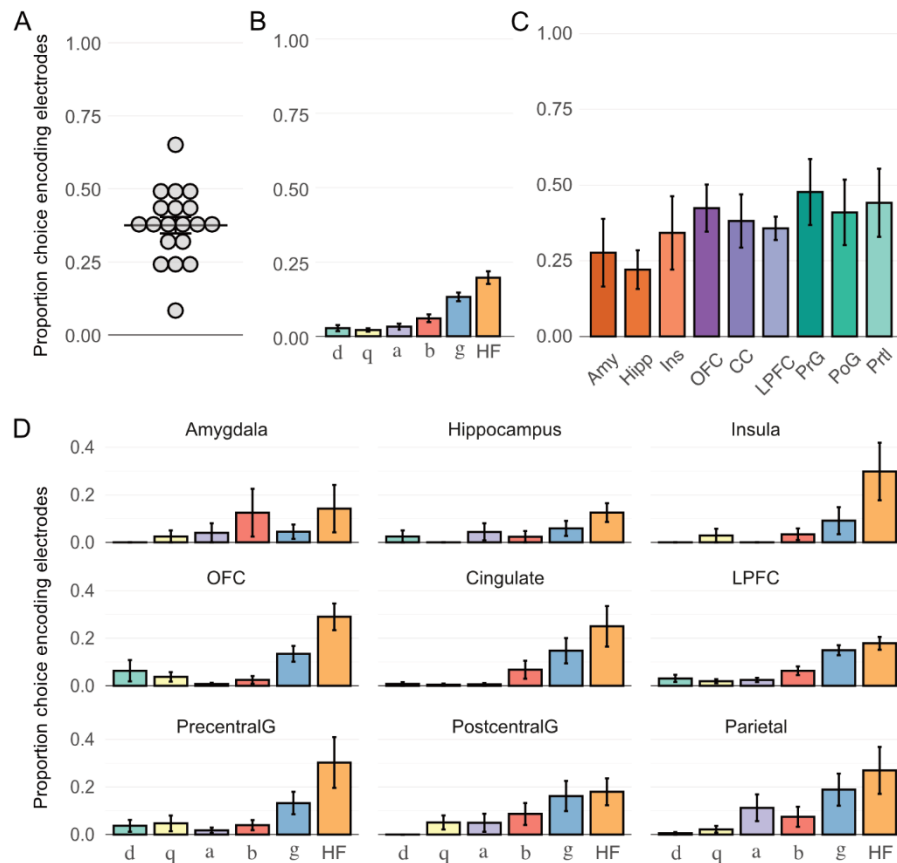


Figure 3. Proportion of choice encoding electrodes showing a significant difference in power between safe bet and gamble trials. Power was compared between safe bet/gamble trials during deliberation (-1s to 0s pre-choice) using an analytical strategy that allowed differences in timing of differences (see Methods). **a**, Proportion of choice encoding electrodes per patient. Horizontal lines show mean \pm SEM. **b-c**, Mean proportion of choice encoding electrodes across patients, grouped by frequency band (across all regions, **b**) and anatomical region (across all frequencies, **c**). An electrode was considered choice-encoding if there was a significant activation between gamble and safe bet trials in any frequency band (see Methods). **(c)** Mean proportion of choice encoding electrodes for any frequency band across subjects grouped by region. **(d)** Mean proportion of choice encoding electrodes across subjects, grouped by both region and frequency.

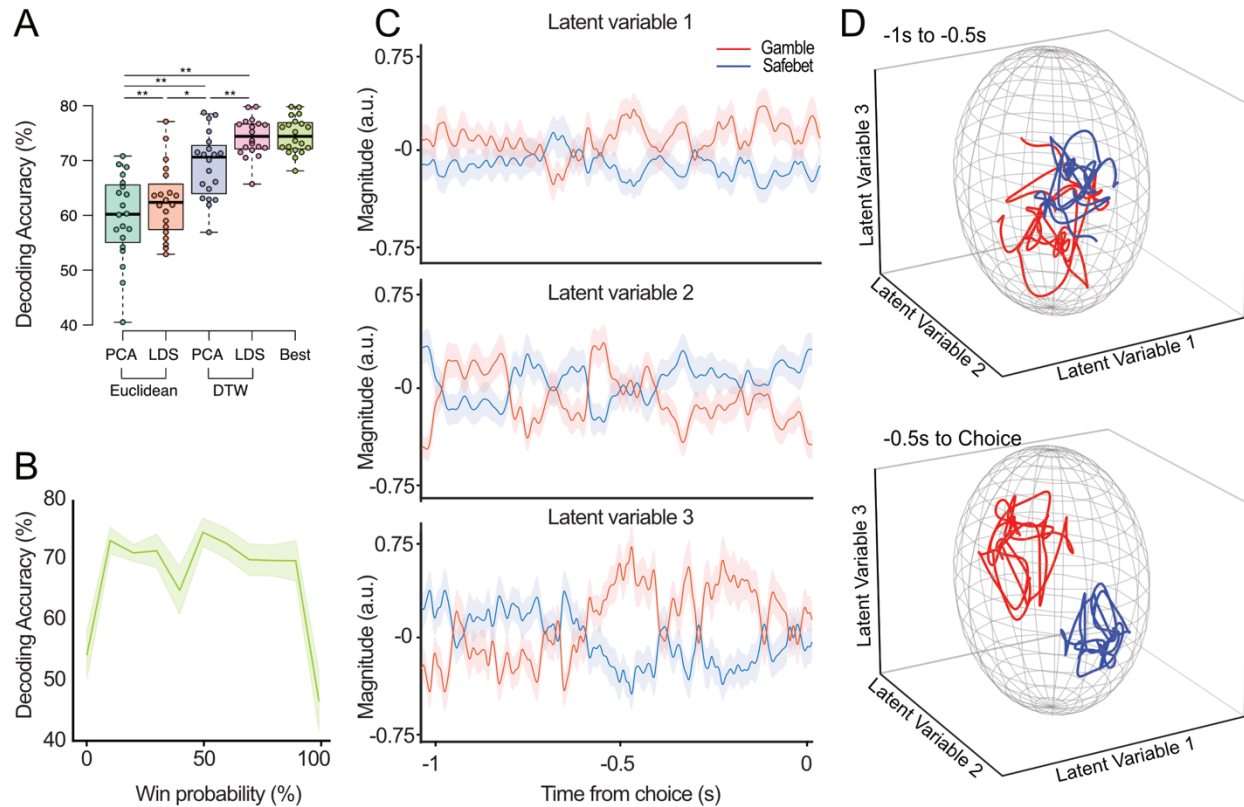


Figure 4. Choice decoding models and results. **a**, Classification performance for each decoding strategy, across all patients. We used two different dimensionality reduction methods, PCA and LDS, and two classification strategies, ED and DTW (see Methods). Results are shown for each combination of dimensionality reduction and classification. The “best” category represents the highest classification accuracy achieved for a given patient, across all strategies. **b**, Decoding accuracy for the LDS-DTW model as a function of trial win probability, averaged across all patients. **c**, Temporal evolution of LDS latent variables in an example patient. The plot shows the mean \pm SE of the average latent variable trajectories for gamble (red) and safe bet (blue) trials. **d**, Latent variable trajectories in 3-dimensional space, separated in early (-1s to -0.5s pre-choice, top) and late (-0.5 to 0s pre-choice, bottom) deliberation period.

Methods

Subjects

Data were collected from 34 patients (19 female) with intractable epilepsy who were implanted with chronic subdural grid or strip electrodes (electrocorticography, ECoG) or stereotactic EEG (sEEG) electrodes as part of a procedure to localize the epileptogenic focus. Electrode placement was based solely on the clinical needs of each patient. Data were recorded at five hospitals: The University of California (UC), San Francisco Hospital (n=3), the Stanford School of Medicine (n=3), UC Irvine Medical Center (n=23), and UC Davis Medical Center (n=5). As part of the clinical observation procedure, patients were off anti-epileptic medication during these experiments. All subjects gave written informed consent to participate in the study in accordance with the University of California, Davis or University of California,

Berkeley Institutional Review Board. Patients understood that they could decline participation at any time, and verbal assent was reaffirmed prior to each experimental task.

Electrophysiological data acquisition

ECoG and sEEG activity was recorded, deidentified, and stored at the same time as behavioral data. Data was collected using Tucker-Davis Technologies, Nihon-Kohden, or Natus systems. Data processing was identical across all sites: channels were amplified x10000, analog filtered (0.01-1000 Hz) with > 1kHz digitization rate, re-referenced to a common average offline, high-pass filtered at 1.0 Hz with a symmetrical (phase true) finite impulse response (FIR) filter (~35 dB/octave roll-off). Behavioral data was simultaneously collected using a PC laptop running Python (v.2.7) and PsychoPy (v.1.85.2) and synchronized with a timed visual stimulus (trial start) recorded by a photodiode through an analog input to the electrophysiological system.

Behavioral task

We probed risk-reward tradeoffs using a simple gambling task described previously³⁸. Briefly, patients chose between a safe bet (\$10, fixed) or a gamble for potential higher winnings (between \$15 and \$30). Gamble win probability varied per trial based on an integer between 0-10 shown at game presentation. At the time of outcome ($t=550\text{ms}$ post-choice, Fig 1a), a second number (also 0-10) is revealed. The gamble results in a win if the second number is greater than the first, and ties were not allowed, therefore, a shown '2' had a win probability of 20% and an '8' had a win probability of 80%. Both numbers were randomly generated using a uniform distribution. Location of safe bet and gamble options (left/right) were randomized across trials. Patients played 10 practice trials, repeated as many times as necessary, to ensure they had full knowledge of the (fair) structure of the task prior to game play (200 trials). Timing is summarized in Figure 1a. Trials started with a fixation cross ($t=0$), followed by a game presentation screen ($t=750\text{ms}$). Patients had up to 8s to respond (mean reaction time = 1.4s), and gamble outcome presentation appeared 550ms after button press (choice). A new round started 1s after outcome reveal. The experimental task typically lasted 12-15min. This gambling task minimized other cognitive demands (working memory, learning, etc.) on our participants while allowing us to test for decision-making under risk. Behavioral performance was assessed by examining the proportion of trials in which the patient chose to gamble as a function of win probability; the proportion of risky trials was calculated for each win probability value (0-100% in 10% increments) and fit with a logistic curve (Fig. 1B, Extended Data Fig. 1). As a control for behavioral data quality, we excluded patients in which a logistic function did not appropriately fit the relationship between percentage of gambles and win probability ($p<0.05$, logistic fit). Results from a subset of these patients playing the same task were published previously³⁸. Fourteen (14) patients did not show a significant fit and were removed from further analysis, leaving twenty ($n=20$) subjects.

Anatomical analyses

Electrode localization was based strictly on clinical criteria for each patient, 9/20 had electrocorticography (ECoG) grids, predominantly in orbitofrontal, lateral prefrontal, and parietal regions, whereas 11/20 had stereotactic EEG (sEEG) coverage, predominantly of deep temporal lobe regions (amygdala, hippocampus) (Extended Data Fig 2). For each patient, we collected a pre-operative anatomical MRI (T1) image and a post-implantation CT scan. The CT scan allows identification of individual electrodes but offers poor anatomical resolution, making it difficult to determine their anatomical location. Therefore, the CT scan was realigned to the pre-operative MRI scan following a previously described procedure³⁹. Briefly, both

the MRI and CT images were aligned to a common coordinate system and fused with each other using a rigid body transformation. Following CT-MR co-registration, we compensated for brain shift, an inward sinking and shrinking of brain tissue caused by the implantation surgery. A hull of the patient brain was generated using the FreeSurfer analysis suite, and each grid and strip was realigned independently onto the hull. This step was necessary to avoid localization errors of several millimeters common in ECoG patients. Subsequently, each patient's brain and the corresponding electrode locations were normalized to a template using a volume-based normalization technique and snapped to the cortical surface³⁹. Finally, the electrode coordinates are cross-referenced with labeled anatomical atlases (Brainnetome atlas) to obtain the gross anatomical location of the electrodes, verified by visual confirmation of electrode location based on surgical notes. We selected all grey matter electrodes across a broad set of regions known to be involved in reward-related behavior for analysis (Fig. 1C): lateral prefrontal cortex (LPFC; 391 electrodes from n=19 patients), orbitofrontal cortex (OFC; 193, n=18), cingulate cortex (CC, 84, n=13), hippocampus (HC; 65, n=13), amygdala (Amy; 32, n=11), insula (Ins; 46, n=8), precentral gyrus (PrG; 108, n=11), postcentral gyrus (PoG; 88, n=9), and parietal cortex (PC; 78, n=8) (Fig. 1; see Extended Data Table 1 for a complete account of electrode numbers across regions and patients).

Electrophysiological analyses

Quality control and preprocessing

Epileptogenic channels and channels with excessive noise (low signal-to-noise ratio, 60 Hz line interference, electromagnetic equipment noise, amplifier saturation, poor contact with cortical surface) were identified and deleted. Out of 1194 electrodes localized to regions of interest, 1085 were artifact-free and included in subsequent analyses. Additionally, all channels were visually inspected to exclude epochs of aberrant or noisy activity (typically <1% of datapoints). Data analysis was carried using custom scripts written in MATLAB and Fieldtrip toolbox⁴⁰. Data for each channel was downsampled to 1KHz. each channel was lowpass filtered (200Hz), highpass filtered (1Hz), and notch filtered (60Hz and harmonics) to remove line noise, and downsampled to 1 kHz if necessary. Electrode channels were re-referenced to a common average reference of all electrodes in each strip/grid. Even though bipolar derivations or white matter referencing are often used for sEEG electrodes, we opted to use a single (CAR) re-referencing strategy for both ECoG and sEEG electrodes for analytical consistency. Trials were epoched to the time of decision using a [-4,3]s window around events of interest (options presentation, and patient choice), and the leading and trailing 1s of data were discarded to remove edge effects. Time-frequency representations (DPSS taper method) were plotted for each region and patient (averaged across electrodes and trials) and visually inspected for artifacts.

Time-frequency representation of neural activity (bandpass estimates)

To examine the role of individual oscillatory bands, we decomposed the neural activity into canonical, discrete activity bands: (delta, δ [1-4Hz]; theta, θ [4-8Hz]; alpha, α [8-12Hz]; beta, β [12-30Hz]; gamma, γ [30-70Hz]; high frequency activity, HFA [70-200Hz]) for each grey matter electrode for each patient using the Filter-Hilbert method. Power in the 6 bands was calculated by applying a Butterworth bandpass filter (order 3 for delta and order 4 for all other power bands) and Hilbert transform and multiplying the resultant complex signal by its complex conjugate⁴¹. As before, one second is removed from the beginning and end of the data to reduce edge effects. Prior to dimensionality reduction for classification, power data for each trial and channel was smoothed and downsampled using a 50ms sliding window with 10ms step increments³⁸ and z-scored over the time dimension within each band to correct for the 1/f profile of neural activity.

Task-Active Electrodes

To examine patterns of electrodes that showed significant task-related power modulation, we compared average power estimates during deliberation period (-1 to 0s pre-choice) with baseline power estimates (-0.5 to 0 pre-stimulus onset) for each electrode and each frequency band independently (paired t-test across trials, $\alpha = 0.01$). Data were analyzed and plotted using custom scripts in Matlab and R. To investigate patterns across patients, regions, and powerbands, we summarized and plotted the proportion of task-active electrodes for each patient (Fig. 2A, Extended Data Tables 3), and then calculated the mean proportion of task-active electrodes across patients for each power band (Fig. 2B, Extended Data Table 5) and region (Fig 2C, Extended Data Table 7). Mean patient proportions and standard errors provide a depiction of population activity across patients that would be obscured in the aggregate. For comparison, proportions of task modulated electrodes overall ($n=1042$ total electrodes) are summarized separately (Extended Data Tables 2,4,6). Finally, in order to probe the apparent homogeneity of the results, we plotted proportions of task active electrodes that increased or decreased power per power band and regions (Fig. 2D,F,G, Extended Data Figure 3, and Tables 8,9) which revealed patterns that were quantified with clustering methods below.

Region Clustering

Analysis of task-active electrodes that increased or decreased power (above) revealed an apparent pattern of similarities among three groups of regions (Fig. 2D): prefrontal (OFC/LPFC/CC), frontoparietal (PrG/PoG/PC), and limbic (Amy/HC/Ins). To test this hypothesis, we applied an unsupervised clustering algorithm (k-nearest neighbor) to evaluate whether the observed similarities could define functional groups that underwent similar patterns of power modulations. We parameterized activity in each brain region by estimating the proportion of electrodes that showed increases or decreases in power modulation during the deliberation period, separately for each patient, region, and frequency band (see Task Modulation, above). Therefore, each patient-region combination (104 total data points, see Extended Data Table 1) is initially represented by a single point in 12D space (increase or decrease in power x 6 frequency bands). To reduce noise or redundancy, we performed dimensionality reduction using Uniform Manifold Approximation and Projection⁴². UMAP converts the data into a k-neighbor graph and then identifies a projection into a lower dimensional space by minimizing the cross-entropy between the two representations while maintaining the fundamental characteristics of the original graph. Once the data were projected into a lower dimension we applied the k-nearest-neighbor algorithm to sort the data into a specified number of clusters. This method effectively separated the regions into three hypothesized functional groups (bootstrapped $p < .05$) (Extended Data Table 9).

Choice-Encoding Electrodes

Next, we defined choice-encoding electrodes as those that showed a significant period of activation during deliberation (-1 to 0s pre-choice) between trials where a gamble bet was chosen compared to safe bet. This was determined by permutation test. Trials were separated by subsequent choice (gamble or safe bet) and then a two-sample t-test was applied to every time point (1 ms resolution) in the 1-second deliberation epoch. MATLAB's `bwconnect` function was used to find contiguous suprathreshold clusters of points ($\alpha < 0.05$). An electrode was categorized as choice-encoding by comparing the sum of the T-statistic for the largest cluster to a trial-shuffled null distribution at α of 0.001 (2-tailed). To determine the null distribution, trial labels were shuffled 10,000 times and the sum of the T-statistic of the largest suprathreshold cluster was recorded on each iteration. This was repeated for each electrode, for each power band, for each patient. Similar to task-active electrodes, we plotted the mean proportion of choice-

encoding electrodes for each patient (Fig. 3A, Extended Data Table 11), and summarized mean proportions of choice-encoding electrodes across patients by frequency band (Fig 3B, Extended Data Table 14) and region (Fig 3C, Extended Data Table 13), as well as separated by region and frequency (Fig 3D). Regions and frequency bands with the greatest proportions of choice encoding were selected as features for subsequent choice decoding (below).

Decoding

Identifying optimal window of interest

To identify the optimal window of time before choice, we tested performance of a Euclidean distance classifier (ED, see below) on the first 50 ms before choice and then repeatedly increased the size of the window by 500 ms increments up to 2 s prior to choice. This was done separately for each subject, using all available electrodes for all six power bands and all available trials with leave-one-out (LOO) cross-validation. Classifier performance peaked at approximately 1 s prior to choice and for all subsequent analyses, we used a 50 ms resolution³⁸ and 1 s window prior to choice. Results were similar if we started 2 s before choice and added 50 ms windows to reach the time of choice (Supplementary Methods Figure 1). The resulting classification had modest performance, correctly classifying above chance in 18/20 patients (mean classification accuracy = $59.3 \pm 0.06\%$, $p < 0.001$ vs bootstrapped performance), reaching classification accuracies as high as 72.3% in one patient.

Feature Selection

Choice encoding results (Fig 3) indicated that choice-related information was distributed across all regions analyzed and was most likely to be carried by higher frequency bands (HFA, gamma, and possibly beta). We tested the contribution of beta to classification performance using linear dynamical systems (LDS) modeling and dynamic time warping (DTW) (see below) with gamma/HFA, with and without beta; classification results were nearly identical ($R^2 = 0.75$) in both cases and not significantly different ($p = 0.60$) (Supplementary Methods Table 1). Therefore, we selected gamma and HFA power bands in all regions for subsequent choice decoding.

Dimensionality Reduction

To reduce the dimensionality of data used for classification, we used two different dimensionality reduction techniques: principal component analysis (PCA) and linear dynamical systems modeling (LDS).

Principal component analysis (PCA)

To account for the variation in anatomical coverage and number of electrodes across patients, we performed PCA separately for each region in each patient. For each region, the original neural features were power in the gamma/HFA frequency bands for all electrodes. The original dataset for each patient thus contained a variable number of electrodes across a variable set of regions, and two power estimates (gamma and HFA). Through PCA, we reduced the number of features to the (varying) number of regions for each subject times the number of bands used (2, gamma/HFA).

Linear dynamical system (LDS) model

As an alternative to PCA, we employed LDS modeling. LDS models allow characterization of the development of neural activity through time, and thus may be better suited for identifying how neural computations related to the on-going task unfold⁴³. As with PCA, LDS reduces the dimensionality of the data⁴⁴ but additionally uses an expectation-maximization likelihood function to solve for the parameters of the model.

Equation 1:

$$x_{t+1} = Ax_t + w_t$$

Equation 2:

$$y_t = Cx_t + v_t$$

Equation 1 shows how the next state, x_{t+1} , is calculated based on the previous state x_t and the state noise w_t . Equation 2 calculates the output, y_t , based on the state x_t and output noise v_t . Both A and C are learned using expectation maximization. The LDS seeks to determine a set of low-dimensional representations of neural activity that have maximal predictive power in their temporal trajectories. As above, the first PC for each region-powerband combination was used in LDS modeling. The appropriate dimensionality of the model (number of latent variables, LVs) was determined empirically by re-running the model for a range of dimensions (1-15) and examining the predictability of neural dynamics. When increasing the dimension size does not improve the predictability of the neural dynamics, the optimal dimensionality has been reached. The dimensionality value ranged from 8 to 16 across patients. In an effort to use a consistent number of dimensions for all patients, we selected a value of 7 LVs. Each dimension of the LDS model defines a latent variable that describes how the dynamics unfold and these LVs can be exploited to further probe the underlying dynamics of choice processing.

Classifiers

Once the dimensionality of the neural data was reduced using either PCA or LDS, we employed classification techniques to attempt to decode trial-by-trial patient choices from the neural data alone. Again, we employed two different strategies: a simple Euclidean Distance classifier and a Dynamic Time Warping approach.

Euclidean Distance Classifier (ED)

We used a simple Euclidean distance (ED) classifier as this method is computationally efficient and shown to be as effective as linear discriminate analysis (LDA) at classifying neural activity in response to different events¹⁷. To estimate the decoding accuracy, we used a leave-one-out procedure. Briefly, we created separate neural trajectory templates for safe bets and gambles by averaging the dimensionality-reduced neural features (PCs or LVs, of size n ROIs \times n frequencies (2) \times n time bins) across trials of either type for each electrode, on a per patient basis. A single test trial of similar dimensions was left out from both the dimensionality reduction and this averaging procedure. After the templates were calculated, the overall Euclidean distance between this trial and the average safe bet/gamble templates was calculated. The predicted choice was then assigned to the closest template in Euclidean space. If the predicted choice matched the true, behaviorally expressed choice (either safe bet or gamble), the trial was correctly classified. This procedure was repeated for all trials in our sample. Performance was defined as the percentage of correctly classified trials, which is calculated on a patient-by-patient basis.

Dynamic Time Warping (DTW)

While an efficient classification method, the ED classifier's fixed template matching and temporal comparison may not be the best approach for decoding complex neural processes. We attempted to improve on the ED classifier by using dynamic time-warping (DTW), a distance metric that is capable of accounting for variation in time of neural activation through flexible stretching or shrinking of the time

dimension. Briefly, DTW seeks to minimize differences between two temporal sequences (in this case, the individual trial and each of the two templates) while allowing some flexibility in the matching. Specifically, instead of matching point-by-point, temporal warping that respects the sequence of points in each time series is allowed. DTW does this by using dynamic programming to calculate whether one or neither of the time series should be stretched in order to minimize the Euclidean distance between the two time series for each successive point in time, we used the Matlab implementation of dynamic time warping (`dtw`). We hypothesized this additional temporal flexibility would result in classification improvement if the timing of neural activity supporting patient choices varies from trial to trial, but the underlying computations are the same. For DTW we used the same leave-one-out procedure as the ED-based classifier with the average for both safe and gamble trials separately warped with the left-out trial before calculating the ED.

Bootstrapping

To identify bias in the neural data that may account for performance above the expected 50% chance performance, we randomly swapped the labels (gamble or safe bet) for each trial and assessed performance on the PCA+LDS+DTW model repeatedly, 1000 times (e.g. bootstrap performance). The distribution of performance provides an empirical assessment of chance performance for each subject. This is especially important here because in this task, subjects are more likely to select a gamble than a safe bet (see Results) and this will bias the performance of the classifier. The bootstrap performance across subjects was almost exactly 50% as expected ($50.1 \pm 0.002\%$), while the actual performance was significantly better than chance ($74.3 \pm 3.4\%$).

Latent variable trajectories and manifolds

We examined the latent variables (LVs) from the LDS model in order to probe the underlying choice dynamics. For plotting purposes, three of seven latent variables from one patient were selected to maximize differences between gamble and safe bet trials. The patient with the best decoder performance was chosen (p06, 77.0%, see Extended Methods Table 16). A manifold was defined by an ellipsoid that contained the average gamble and safe bet trajectories for the selected patient (min and max for each LV dimension plus an additional 70% (LV₁) or 60% (LV₂, LV₃)).

References

38. Saez, I. *et al.* Encoding of Multiple Reward-Related Computations in Transient and Sustained High-Frequency Activity in Human OFC. *Current Biology* **28**, 2889-2899.e3 (2018).
39. Stolk, A. *et al.* Integrated analysis of anatomical and electrophysiological human intracranial data. *Nature Protocols* **13**, 1699 (2018).
40. Oostenveld, R., Fries, P., Maris, E. & Schoffelen, J.-M. FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience* **2011**, e156869 (2010).
41. Cohen, M. X. *Analyzing Neural Time Series Data: Theory and Practice*. (MIT Press, 2014).

42. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. Preprint at <http://arxiv.org/abs/1802.03426> (2020).
43. Macke, J. H. *et al.* Modelling low-dimensional dynamics in recorded spiking populations. in (2011).
44. Ghahramani, Z. & Hinton, G. E. *Parameter Estimation for Linear Dynamical Systems*. (1996).