

RESEARCH ARTICLE | *Higher Neural Functions and Behavior*

Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles

Tommy C. Blanchard,¹ Steven T. Piantadosi,¹ and Benjamin Y. Hayden²

¹Department of Brain and Cognitive Sciences, Center for Visual Science, and Center for the Origins of Cognition, University of Rochester, Rochester, New York; and ²Department of Neuroscience and Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, Minnesota

Submitted 7 November 2017; accepted in final form 1 December 2017

Blanchard TC, Piantadosi ST, Hayden BY. Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *J Neurophysiol* 119: 1305–1318, 2018. First published December 6, 2017; doi:10.1152/jn.00808.2017.—Classification of neurons into clusters based on their response properties is an important tool for gaining insight into neural computations. However, it remains unclear to what extent neurons fall naturally into discrete functional categories. We developed a Bayesian method that models the tuning properties of neural populations as a mixture of multiple types of task-relevant response patterns. We applied this method to data from several cortical and striatal regions in economic choice tasks. In all cases, neurons fell into only two clusters: one multiple-selectivity cluster containing all cells driven by task variables of interest and another of no selectivity for those variables. The single cluster of task-sensitive cells argues against robust categorical tuning in these areas. The no-selectivity cluster was unanticipated and raises important questions about what distinguishes these neurons and what role they play. Moreover, the ability to formally identify these non-selective cells allows for more accurate measurement of ensemble effects by excluding or appropriately down-weighting them in analysis. Our findings provide a valuable tool for analysis of neural data, challenge simple categorization schemes previously proposed for these regions, and place useful constraints on neurocomputational models of economic choice and control.

NEW & NOTEWORTHY We present a Bayesian method for formally detecting whether a population of neurons can be naturally classified into clusters based on their response tuning properties. We then examine several data sets of reward system neurons for variables and find in all cases that neurons can be classified into only two categories: a functional class and a non-task-driven class. These results provide important constraints for neural models of the reward system.

clustering; functional subtypes; mixed selectivity; prefrontal cortex; reward

INTRODUCTION

A neuron can be characterized based on its response (generally, its tuning function) to various task variables. Tuning properties for a set of neurons can then be used to categorize them into discrete subgroups with distinct functional proper-

ties. For example, one could imagine that neurons may be classified based on whether they respond to value, to space, or to both. While such groupings may be heuristically useful, they do not necessarily identify natural clustering that, to use Plato's phrase, "carves nature at its joints." That is, a categorization procedure either can detect discrete categories or it can artificially impose groupings on what is actually a continuum of responses. Thus, for example, a population of neurons could have selectivities for reward and space that are drawn from two uniform and independent distributions; selecting the neurons that are most strongly reward sensitive and putting them into a category would impose a cluster on what is a actually a continuum.

Standard clustering methods like k-means do not distinguish between these two sorts (i.e., natural and externally imposed) of categories. Instead, they produce best clusterings without providing any indication about whether the underlying distribution is more clustered than would be expected by chance. Another approach with the same limitation is to use a cutoff like $P < 0.05$ to divide cells into categories and use the resulting cell classifications to make inferences about category function. This method can be susceptible to both Type I and Type II errors (Maxwell and Delaney 1993). Three aspects of this approach are particularly problematic. First, it uses categorical classification (significant vs. nonsignificant) on continuous variables and thus throws away a great deal of information (a Type II error problem). Second, whether a neuron reaches significance will depend on the number of trials collected; this arbitrary factor will increase variance in the estimate of clusterings (another Type II problem). Third, due to the mathematics of joint probability, when data are noise limited (i.e., almost always), fewer neurons will reach significance for two variables than for either one. This bias will lead to a systematic overestimate of coding disjunctions—that is, neurons will appear more purely categorical than they really are (a Type I problem).

Understanding the distribution of tuning functions is important for several reasons. First, delineation of cell types is an important basic description of the nature of responses in a brain area. Classification provides a useful way to simplify what are often complex and heterogeneous response patterns. For example, comparisons of response properties between brain regions gives insight into the core differences in their functions

Address for reprint requests and other correspondence: B. Y. Hayden, University of Minnesota, Minneapolis, MN 55455 (e-mail: benhayden@gmail.com).

(e.g., Kennerley et al. 2009; Seo and Lee 2009; Sleezer et al. 2017; Strait et al. 2015; Tsujimoto et al. 2011; Wallis et al. 2001). Second, categorization provides important constraints on neural models. For example, two-pool models of economic choice rely on classification of cells into two types based on selectivity for individual stimuli in economic choice tasks (Chau et al. 2014; Hayden and Moreno-Bote 2017; Hunt et al. 2015; Louie et al. 2011; Rustichini and Padoa-Schioppa 2015; Soltani et al. 2006). Finally, categorization of neurons provides a possible baseline for establishing functional implications of anatomical and genetic methods that delineate specific cell types (e.g., Kvitsiani et al. 2013).

The question of whether there are natural functional categories of cells in most brain regions remains unanswered. This problem is particularly salient now. Recent successes of population analysis methods have demonstrated the power of ensemble coding and distributed computation (Blanchard et al. 2015b; Churchland et al. 2012; Hunt and Hayden 2017; Hyman et al. 2012; Kristan and Shaw 1997; Ma et al. 2014; Mante et al. 2013; Stokes et al. 2013). Moreover, theoretical advances confirm the flexibility and power of mixed encoding schemes (Barak et al. 2013; Fusi et al. 2016; Ganguli and Sompolinsky 2012; Pouget and Sejnowski 1997; Rigotti et al. 2010, 2013). Finally, a recent high-profile study specifically investigating the question of tuning categories reported no evidence of categorical organization in rodent parietal cortex (Raposo et al. 2014). This finding raises the possibility that category-free selectivity may be a general property of neural systems.

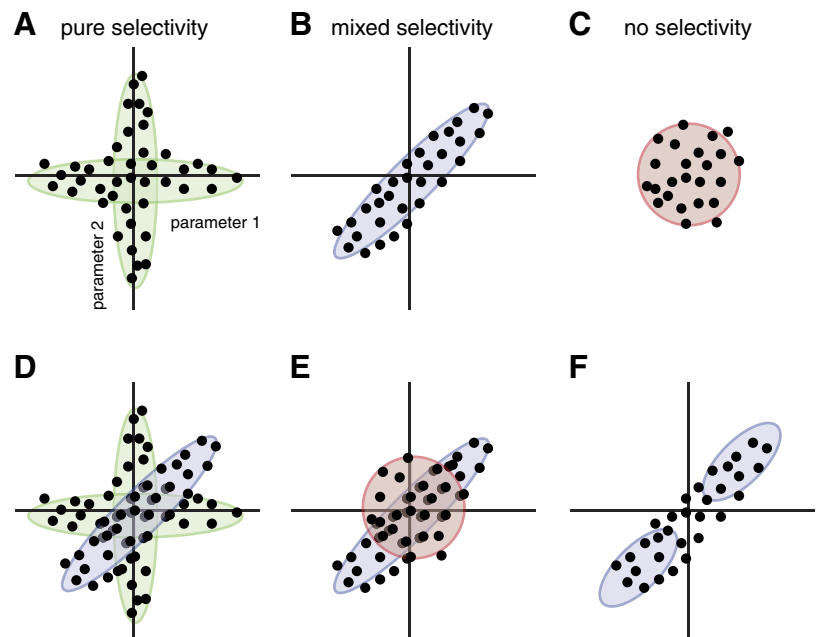
We developed a novel analysis approach by using a generative Bayesian model. While these problems can be solved using conventional frequentist statistics, they are more naturally solved using Bayesian approaches, wherein the standard equations take into account estimates of the uncertainty in the encoding of variables in an optimal manner. This approach also allows us to place a confidence estimate on the likelihood that null results stem from lack of sufficient power or whether they reflect true nulls. Our model makes statistical assumptions about the components likely to have given rise to the observed

firing patterns. It identifies three potential components in any data set: 1) A *no-selectivity component* that generates neurons that can confidently be said to have no tuning to any task variables (that is, it did not simply fail to achieve statistical significance), 2) a set of *pure-selectivity components* that generates neurons that have tuning to only one task variable, and 3) a *multiple-selectivity component* that generates neurons that have tuning to multiple task variables. The statistical analysis assumes that the observed data come from a weighted mixture of these components and infers the weights from the neural firing patterns. Thus, it might discover that most neurons are purely selective (strong categorical organization, Fig. 1A) or come from one large category of multiple selectivity cells (Fig. 1B), or even that they are nonselective (not task relevant, Fig. 1C). It could also show a mixture of these sets, involving multiple discrete categories (Fig. 1, D and E). Alternatively, and in contrast to clustering methods, the analysis might show that the data are not highly informative about the classification, telling us that additional techniques or larger data sets are required.

To identify these components, we use Bayesian tools [specifically, Markov chain Monte Carlo (MCMC) sampling] in the programming language Stan to infer likely values of the mixture weights and parameters of these components (Carpenter et al. 2015). Our analysis starts from relatively unbiased assumptions about what the parameter values may be (see MATERIALS AND METHODS). Several choices in our modeling setup are designed to make this analysis robust, in particular to outliers and the presence of no-selectivity neurons. These include deliberate inclusion of a noise component to remove the influence of no-selectivity neurons, use of *t*-distributions as mixture components due to their robustness to outliers (Peel and McLachlan 2000), and inclusion of variance and covariance from the underlying regression to correctly handle uncertainty in each neuron's response profile. Our method is publicly and freely available (Blanchard 2016).

To validate this method, we applied it to synthetic data sets with known category structure and show that it accurately

Fig. 1. Cartoon illustrating possible effects in a data set. Consider a set of neurons in which cells' tuning properties for two parameters are independently assessed and then *z*-score transformed. Scatter plots are cartoon versions; each dot corresponds to a neuron; the two dimensions are two tuning dimensions. *A*: neurons may fall into pure-selectivity clusters. In the example, neurons are strongly tuned for parameter 1 or 2 but not both. *B*: another possibility is that neurons fall into a single larger cluster selective for both variables. In the presence of noise, it is often difficult to distinguish pure from multiple selectivity. *C*: a third possibility is that neurons will have no selectivity for either variable and will be best classified as nonselective. Some cells may be significantly tuned, but thus number is no greater than the false positive rate expected by chance. *D*: a population of neurons can also contain a combination of subsets that are pure and mixed, as for example when a population contains subpopulations correspond to each of two variables and a third that integrates them. *E*: neuronal populations can also contain other mixtures of these. In particular, our data suggest that a combination of multiple and no selectivity populations are common, perhaps even universal, in reward regions. *F*: conventional clustering methods make it difficult to judge categorical structure of populations. For example, a method like *k*-means clustering will divide cells into two categories (two blue ovals) even if they are from a single distribution.



infers their parameters. We then examined functional properties of neurons in four reward-sensitive frontostriatal brain areas in a variety of economic decision-making tasks [orbitofrontal cortex (OFC, Blanchard et al. 2015a; Wang and Hayden 2017); dorsal anterior cingulate cortex (dACC, Azab and Hayden 2016, 2017; Blanchard et al. 2014, 2015b); ventromedial prefrontal cortex (vmPFC, Strait et al. 2014); and ventral striatum (VS, Strait et al. 2015)]. For all 20 sets of variables we examined (ones relevant to economic choice and executive control) we found the same pattern: neurons fell into only two categories, one category containing all task-sensitive cells, and another one containing cells best classified as pure no-selectivity. We found no evidence for distinct categorical tuning in any of the data sets we analyzed.

MATERIALS AND METHODS

Model

The model was implemented in the programming language Stan (Carpenter et al. 2015). Code is available on Github and distributed under a GPL3 license address: <https://github.com/TommyBlanchard/StanNeuronModelling>.

Formally, the estimated x, y points in a scatter plot (each point corresponding to a neuron's estimated tuning parameters) are assumed to be generated from a weighted mixture of two-dimensional t -distributions. The multivariate t -distributions used a fixed 50 degrees of freedom to approximate a normal distribution while being relatively robust to outliers. The t -distributions were all centered at $(0, 0)$, and each are defined by a covariance matrix that captures their key theoretical properties for each type of mixture component. The pure-tuned distributions allow variation only along one axis, with variance coming from a half-cauchy prior:

$$\sigma_{\text{pure}_x} \sim \text{cauchy}(0, 5)$$

$$\sigma_{\text{pure}_y} \sim \text{cauchy}(0, 5)$$

The multiple selectivity distribution is defined by a correlation matrix (r) with an LKJ prior and a scaling parameter (s):

$$s \sim \text{cauchy}(0, 5)$$

$$r \sim \text{lkj_corr}(1)$$

The LKJ prior selects correlation matrices in proportion to the determinant of the correlation matrix raised to the power of a parameter minus 1 (Lewandowski et al. 2009):

$$P(r) = \det(r)^{n-1}$$

With the parameter set to 1, this yields a uniform distribution over all correlation matrices.

The no-tuning component is defined as having no variance along either axis, thus is a covariance matrix of 0s. This gives us the covariance matrix for all of the components:

$$q_x = [\sigma_x, 0]$$

$$q_y = [0, \sigma_y]$$

$$q_{\text{mixed}} = s * r * s$$

$$q_{\text{no_tuning}} = [0, 0]$$

where q_x , q_y , q_{mixed} , and $q_{\text{no_tuning}}$ are the covariance matrices for the x -aligned pure-tuned, the y -aligned pure-tuned, mixed, and no-tuning distributions, respectively.

Note that each of these generative models specifies that "true" locations for each point (which is itself a set of beta weights for a neuron's response to each of two dimensions of input), and then each point is additionally assumed to be measured with Gaussian noise.

This noise from the GLM part of the analysis is simply added to a component's covariance matrix when estimating the probability of the data point given a particular component:

$$P(X_{n,n} | q_n) \sim \text{multi_student_t}(X_n, 50, q_n + \hat{\sigma}_n)$$

where X_n is the location of the n th neuron in the two-dimensional space and $\hat{\sigma}_n$ is the noise in the estimate of the location of that point as estimated by the GLM component of the analysis.

To compute the likelihood of each data point, we marginalize over which cluster it came from. The weights defining this cluster assignment all come from a Dirichlet distribution or beta distributions with priors that would be Jeffreys' priors for multinomial outcomes:

$$\theta \sim \text{dirichlet}(0.5, 0.5)$$

$$\zeta \sim \text{dirichlet}(0.5, 0.5)$$

$$\eta \sim \text{dirichlet}(0.5, 0.5)$$

Here, θ , ζ , and η are the axes, mixed-tuning, and no-tuning weights, respectively. The weights determine the probabilities of a data point coming from each of the components.

The probability of a data point X_n given the model parameters is therefore given by summing up all of the ways that the data point could have been generated, as shown by the multiple routes in Figure 2. In particular, a data point could be generated as noise, pure component, or multiple component, each of which happens with their weighted probabilities, which are summed in this equation. Conditioned on the model parameters, each data point is treated independently, meaning that the probability of the data is found as a product of the probabilities for each individual point, assuming conditional independence between measured points.

For running inference, we used default parameters from the Stan package. Specifically, we ran the No-U-Turn sampler, an extension to the Hamiltonian Monte Carlo algorithm (Hoffman and Gelman 2014; Neal 2011). We ran this MCMC algorithm for 5,000 steps and 5 chains. Convergence was assessed using visual analysis of trace plots and confirmed by the Gelman and Rubin convergence diagnostic (Gelman and Rubin 1992). The Gelman-Rubin convergence diagnostic (which specifically measures the ratio of variance across different simulations) was below 1.05 for all variables for all data shown, indicating that the numerical methods used in the analysis accurately recovered the true posterior parameter values.

Generation of Physiological Data

All data come from previously published studies; detailed methods are given in the appropriate citation. The following is a brief summary.

Surgical procedures. All animal procedures were approved by the University Committee on Animal Resources at the University of Rochester and were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals. Male rhesus macaques (*Macaca mulatta*) served as subjects.

Recording sites. We approached vmPFC, VS, OFC, and dACC through standard recording grids (Crist Instruments). Images of the recording sites and tasks can be found in Fig. 3. We defined vmPFC as Area 14, lying within the coronal planes situated between 29 and 44 mm rostral to the interaural plane, the horizontal planes situated between 0 and 9 mm from the brain's ventral surface, and the sagittal planes between 0 and 8 mm from the medial wall. We defined VS as NAcc core lying within the coronal planes situated between 28.02 and 20.66 mm rostral to interaural plane, the horizontal planes situated between 0 to 8.01 mm from ventral surface of striatum, and the sagittal planes between 0 to 8.69 mm from medial wall. We defined OFC as Area 13, lying within the coronal planes situated between 29.50 and 35.50 mm rostral to interaural plane, the horizontal planes

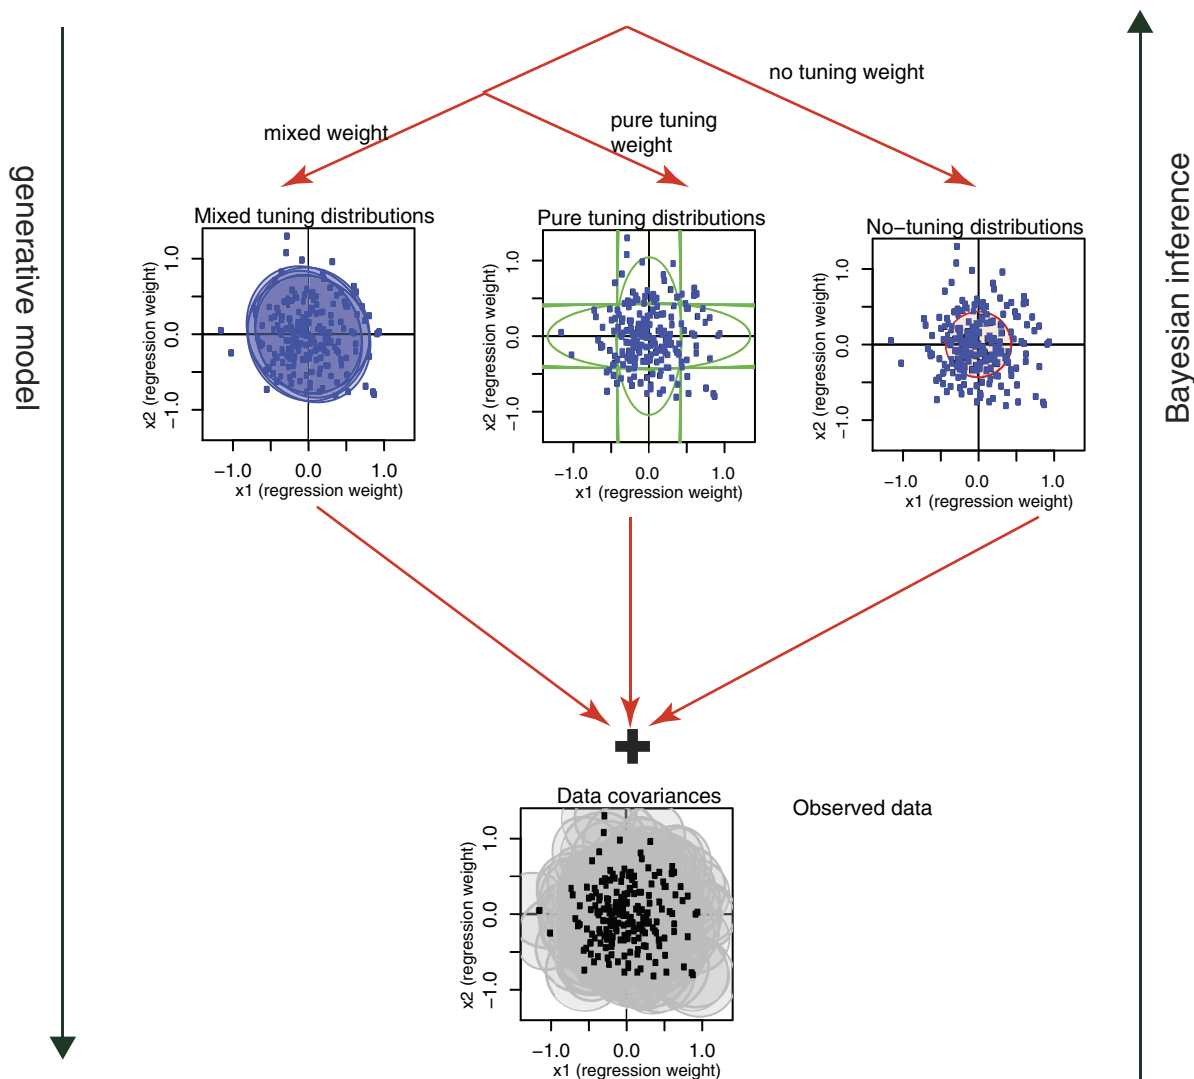


Fig. 2. Illustration of the approach used by our model. This visualization uses a simulated population with neurons coming from all three distributions. Fit distributions are colored based on which distribution they are most likely to have come from according to the model. Blue, multiple selectivity; green, pure tuning; red, no tuning. The row of three scatter plot panels shows the shape of each component of the model. The shading is proportional to the amount of weight on that component. *Bottom* panel shows the observed data, with the circles indicating the variance on the estimate of their location.

situated between 0 to 6.00 mm from the brain's ventral surface, and the sagittal planes between 6.54 to 13.14 mm from medial wall. We defined dACC as Area 24 (for a discussion of our decision to use this terminology, see Heilbronner and Hayden 2016), lying within the coronal planes situated between 29.50 and 34.50 mm rostral to interaural plane, the horizontal planes situated between 4.12 to 7.52 mm from the brain's dorsal surface, and the sagittal planes between 0 and 5.24 mm from medial wall.

Electrophysiological techniques. Single electrodes (Frederick Haer, impedance range 0.8 to 4 M Ω) were lowered using a microdrive (NAN Instruments) until waveforms between 1 and 3 neuron(s) were isolated. Individual action potentials were isolated on a Plexon system (Plexon, Dallas, TX). Neurons were selected for study solely on the basis of the quality of isolation; we never preselected based on task-related response properties. No neurons that surpassed our isolation criteria were excluded from analysis.

Eye-tracking and reward delivery. Eye position was sampled at 1,000 Hz by an infrared eye-monitoring camera system (SR Research). Stimuli were controlled by a computer running MATLAB (MathWorks) with Psychtoolbox and Eyelink Toolbox. Visual stimuli were colored rectangles on a computer monitor placed 57 cm from the

animal and centered on its eyes. A standard solenoid valve controlled the duration of juice delivery. The relationship between solenoid open time and juice volume was established and confirmed before, during, and after recording.

Behavioral tasks. Subjects performed in four different tasks with the same basic structure and one task with a somewhat different structure (Fig. 3). For the neuronal recordings in vmPFC, *subjects B* and *H* performed the *risky choice task*; for VS, *subjects B* and *C* performed the *risky choice task*; for *OFC study 1*, *subjects B* and *J* performed the *curiosity gambling task*; for *OFC study 2*, *subjects B* and *H* performed the *riskless choice task*; for *dACC study 1*, *subjects B* and *J* performed the *token risky choice task*; for *dACC study 2*, *subjects B* and *J* performed the *diet selection task*. All tasks made use of rectangles indicating reward amount and either probability or (in the diet selection task) delay. This method produces reliable communication of abstract concepts such as reward, probability, and delay to monkeys (see corresponding references, and also Blanchard et al. 2015b; Blanchard and Hayden 2015; Hayden et al. 2010a; Pearson et al. 2010, for quantitative analyses demonstrating the robustness of these behavioral methods). In all cases, variables were selected to be uncorrelated. Across the tens of thousands of trials used, observed

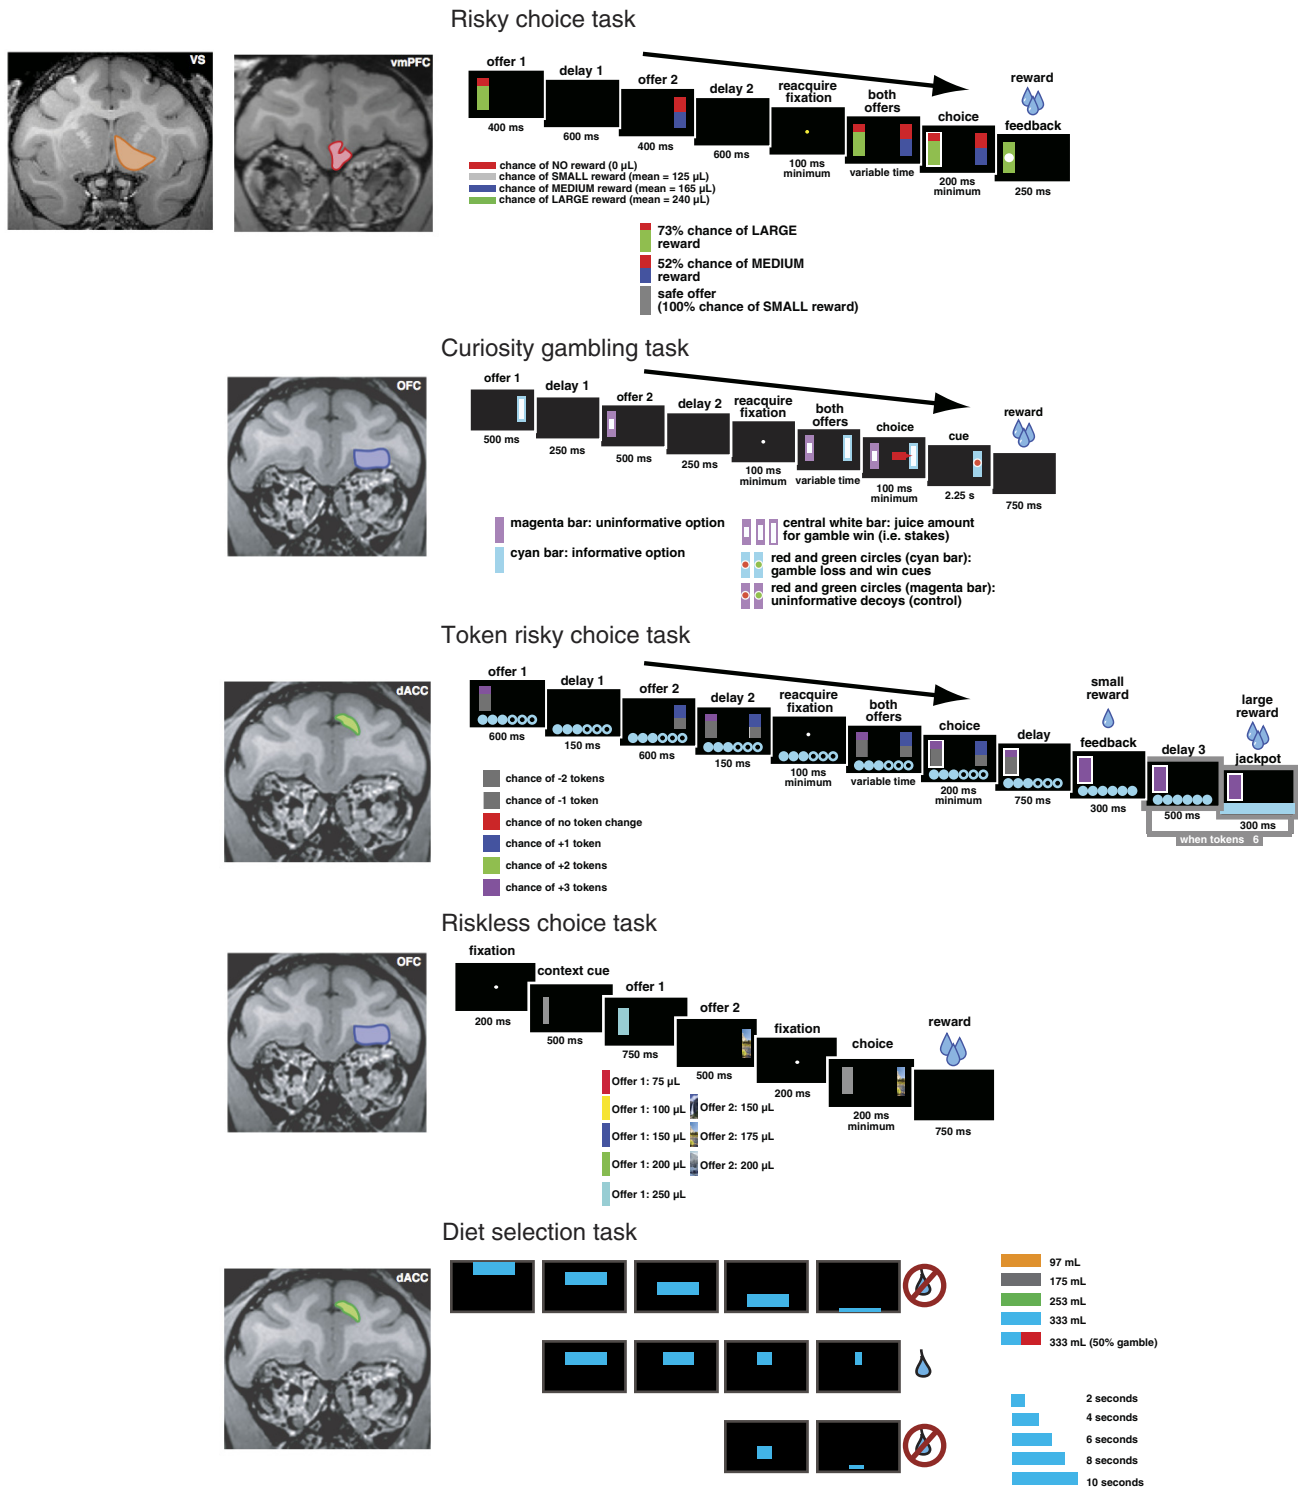


Fig. 3. Illustration of the brain areas and tasks used in our data sets analyzed here.

correlation coefficients were between -0.01 and $+0.01$ and did not significantly differ from zero for any pairs of variables used.

Risky choice task. Two offers, indicated by rectangles, were presented on each trial. Options offered a risky bet for liquid reward (there were safe offers as well, which were excluded from these analyses). Gamble offers were defined by two parameters, *reward size* and *probability*. The size of the blue or green portion of the rectangle signified the probability of winning a medium (mean 165 μ L) or large reward (mean 240 μ L), respectively. These probabili-

ties were drawn from a uniform distribution between 0 and 100%. The rest of the bar was colored red; the size of the red portion indicated the probability of no reward. On each trial, one offer appeared on the left side of the screen and the other appeared on the right. The sides of the first and second offer (left and right) were randomized by trial. Each offer appeared for 400 ms and was followed by a 600-ms blank period. After the offers were presented separately, a central fixation spot appeared and the monkey fixated on it for 100 ms. Following this, both offers appeared simultaneously and the

animal indicated its choice by shifting gaze to its preferred offer and maintaining fixation on it for 200 ms.

Curiosity gambling task. A similarly structured gambling task, where gambles always carried a reward probability of 50% and the size of a white bar in the center of each offer indicated reward size at 21 levels: from 75 to 375 μl in increments of 15 μl . Each trial, the monkey chose between an informative gamble (a cyan bar; if chosen a cue 2.25 s before the potential reward would indicate whether the monkey was about to win) and an uninformative gamble (a magenta bar; if chosen the cue was replaced with an uninformative decoy; Fig. 2B). The stakes of both options and the order and side on which the informative option appeared were all randomized on all trials. Critically, the information was not revealed during the presentation of the cues, but only after the choice was made. Thus neural responses to the offers were not themselves reflective of the information.

Riskless choice task. This task was structured similarly to the other tasks, with the following exceptions: all cues were 100% valid. The first offer was one of five values [50 100 150 200 250 μl]; the second offer was one of three values [150 175 200 μl]. Half of trials (randomly selected by trial) were experienced (offer was given and choice of the offer would repeat it) and half were described (offer was indicated with a valid visual cue).

Token risky choice task. Another similarly structured gambling task, where gambles each had two potential outcomes, wins or losses in terms of “tokens” displayed onscreen as cyan circles. A small reward (100 μl) was administered concurrently with gamble feedback on each trial, regardless of gamble outcome. Trials in which the monkey accumulated six or more tokens triggered an extra “jackpot” epoch in which a very large reward (300 μl) was administered (Fig. 2C).

Diet selection task. This task was based on the famous foraging task of the same name. On each trial, a rectangle glided from top to bottom of screen at a rate such that it would take one second to traverse the screen. Fixation of the rectangle caused it to pause and begin shrinking at a constant rate. Successful maintenance of fixation led to a juice reward corresponding to the rectangle’s color [75 135 212 293 293/0 μl]. (The symbol 293/0 indicates a gamble with 50% probability of 293 and 50% of zero). Widths indicated delay and were drawn from [2 4 6 8 10 s].

RESULTS

Description of the Model

Consider an experiment in which two variables are manipulated independently across several trials so that there is no correlation between their magnitudes. In a risky choice, for example, these variables could be reward amount and reward probability. Neural firing rates are recorded for a set of neurons over a few hundred trials as these variables vary randomly and independently from trial to trial and neurons are found to encode both variables. We wish to know whether the population of neurons encoding these two variables is naturally organized into discrete clusters. That is, are the amount and probability neurons the same neurons or separate classes of neurons? Or are the variables assigned at random to individual neurons?

Our model employs a two-part estimation scheme. First, we z -score the firing rates; then we analyze each neuron’s responses with a linear model and compute regression weights (i.e., beta coefficients or values) for each variable as well as the variance and covariance in these estimates. The beta values resulting from these computations summarize the response of each neuron to each variable. Our analysis uses a separate regression from the mixture model for reasons of computa-

tional efficiency, although it is possible in principle to integrate both aspects into a single hierarchical analysis.

Thus, the basic starting point for our analysis can be viewed as a scatter plot, with a point for each neuron located at its x -location the estimated effect for one task variable and its y -location the estimated effect on the neuron for another (see, for example, Fig. 2). The question we address is: what unobserved components drive the pattern in this scatter plot—are there underlying populations of neurons that are selective to both task variables, to one, or to none?

To answer this question, the collection of beta weights is modeled as a mixture of nonselective vs. task selective cells, and then as cell subtypes within the task-relevant categorization. The no selectivity cells are assumed to have beta weights drawn from zero. In practice, of course, the measured numerical value will be different than zero due to statistical noise in sampling and estimation; our model takes this into account. Cells that are unlikely to be no selectivity are modeled as either pure selectivity (beta is zero on one variable and nonzero on the others) or multiple selectivity (the betas come from a covariance matrix).

All of the parameters are jointly inferred from the firing patterns, meaning that we can provide estimates of the proportion of cells that are nonselective (alpha), the proportion that are purely selective (beta), the covariance matrix of cells that have multiple selectivity (1-beta), as well as the probability that any individual cell belongs to any of these categories.

The inferred distribution on covariance matrices for the multiple selectivity cells is particularly interesting because it tells us whether, for instance, cells that respond positively to reward probability likely also respond positively (positive correlation) or negatively (negative correlation) to increases in reward. Or, if this covariance matrix has zero covariance, that would indicate that individual cells have statistically reliable but independent responses to these two task variables.

The model also includes parameters for the “scale” of firing patterns, fit separately to each mixture component. These parameters allow the model to fit the measurement scale used empirically as well as the range of possible firing rates observed. This feature is important because it lets the model deal gracefully with large differences in tuning strengths in different data sets or across different neuron types.

Note that the inference scheme is not dependent on a specific procedure of fitting or clustering. Our procedure is rather a statistical technique that gives the optimal (relative to our assumptions) beliefs about what parameter values are likely, given the observed data. For example, instead of simply being able to state that 40% of cells are task relevant, we can say that there is a 95% chance that between 35% and 50% of the cells are task relevant. (We provide median estimates and credible intervals below). Thus, if the data are underinformative (e.g., due to a low number of cells or trials), our technique could tell us so—the parameter values will have a large range of possible values. Moreover, due to our use of Bayesian inference techniques, the uncertainty in each parameter appropriately adjusts for the uncertainty in the others, meaning that our estimates of variables like the covariance in task selective cells correctly adjusts for uncertainty in the categorization, as well as the uncertainty in the estimated regression coefficients.

Testing the Model on Simulated Data

To confirm that our model works as anticipated, we first tested it on simulated data. We therefore constructed ersatz neuronal data sets. For each neuron, we generated a “true” tuning parameter by sampling from a distribution that could be described as one of the three classes (multiple selectivity, pure selectivity, or no tuning). Separately for each simulated neuron, we generated simulated trials by generating a random value for each variable for each trial and generated neural firing as a product of the variables and the neuron’s tuning parameters (plus Gaussian noise).

We chose the strength of the tuning parameters and the amount of noise so that the regression coefficient values and variances roughly match the levels we saw in our real data sets (median absolute regression coefficient value in simulated data = 0.092, range in true data sets = [0.039 to 0.132]; median regression coefficient variance in simulated data set = 0.0019, range in true data sets = [0.0019 to 0.0107]). As described above, we regressed the firing rates onto their tuning variables, separately for each neuron, and then obtained estimates of regression coefficient and covariance. We then fit our Bayesian model to these outputs.

We first tested whether the model could correctly identify that a population of neurons composed of a single population with multiple selectivity (Fig. 2). The model correctly places all of its weight on the multiple-selectivity component (median multiple selectivity weight = 1, 95% credible interval = [0.93 to 1]).

We next extended our test to consider the case of a population of cells with neurons from each component (Fig. 4, A and B). In this constructed population, 25% of cells were purely selective for variable 1, 25% purely selective for variable 2, 25% had multiple selectivity for both variables, and 25% had no selectivity to either. Our model was able to accurately estimate the proportions of no-selectivity neurons (correct answer: 0.25, model median weight: 0.25, 95% credible interval: [0.15 to 0.36]). It also captured the proportion of the signal weight that was due to multiple selectivity (correct answer: 0.33, model median 0.4, 95% credible interval: [0.24 to 0.59]). Finally, it accurately estimated the proportion of the pure-selective neurons that were sensitive to each variable (correct answer = 0.5, model median weight = 0.53, 95% credible interval = [0.36 to 0.72]).

We then tested the model’s ability to detect correlations in the neurons’ tunings. In a population of neurons with tuning to two variables, there is often a fixed relationship between how they are tuned to each of the two variables. For example, neurons positively tuned to one variable may more frequently be positively tuned to a second variable. This relationship would be observed, for example, if neurons encode value rather than its components (e.g., reward amount and reward probability; Strait et al. 2014, 2015). Note that in this case, as above, the two regressors themselves will not be correlated with each other, even if the firing rate is correlated with both of them in the same way.

To see whether our model can detect this relationship, we tested two more simulated data sets, with the same proportions of neurons coming from each of the components as above, but with a positive correlation of 0.5 between the tunings for the two variables for the multiple-selectivity neurons (Fig. 4,

C–F). The model again was able to correctly estimate the weights for each component. It was also able to correctly estimate the correlation in the multiple-selectivity component (correct answer = 0.5, median for the population of only multiple-selectivity neurons = 0.45, 95% credible interval = [0.33 to 0.56]; median correlation estimated for population with all components = 0.54, 95% credible interval = [0.33 to 0.76]).

Because our model is designed to deal with extremely noisy data sets, we wanted to ensure that it converged to a single set of parameters as more data were added. We therefore tested the model’s behavior for a population of 50% multiple selectivity and 50% no-selectivity neurons. The weights assigned to these guesses rapidly converged as more neurons were added (Figs. 5A). We also examined a population of 50% pure-selectivity and 50% no-selectivity neurons (Figs. 5B). Finally, we tested a population of 25% pure-selectivity to variable 1, 25% pure-selectivity to variable 2, 25% multiple selectivity, and 25% with no selectivity (Fig. 5C). In all three cases, the model quickly converged to the correct weightings after observing data from between 50 and 100 neurons. Additional neurons produced greater convergence but did not qualitatively affect the data.

Strong Evidence Against Categorical Selectivity in Orbitofrontal Cortex

We next applied this analysis technique to real data. We started with a data set collected in the OFC (Blanchard et al. 2015a). In the *curiosity tradeoff task* monkeys chose between offers that differ in two discrete dimensions, reward (water amount) and informativeness (information about the upcoming outcome of a gamble), which is shown, by behavior, to be valued—presumably because it sates curiosity (Kidd and Hayden 2015). Both variables were selected at random from a uniform distribution on each trial and were therefore uncorrelated in the data set.

These two dimensions were both encoded by neurons in OFC during the offer period of the task. Our model can answer some key questions our earlier study could do only crudely. For example, are the neurons that encode reward different from the ones that encode informativeness? Or are these variables distributed in two sets of neurons—or at random across neurons? And are the neurons whose firing rates appeared to be unrelated to these variables simply too weakly tuned to detect an effect, or is there a way to confidently classify them as no-selectivity neurons?

The model produced clear results. We found, first, that reward-sensitive neurons and informativeness-sensitive neurons do not constitute different sets of cells (Fig. 6). Instead, they came from a single larger class of task-relevant cells with multiple tunings (median multiple tuning signal weight = 1, 95% credible interval = [0.96 to 1]). We also found that the correlation between the coefficients is not significant, indicating that the variables are not fully integrated into a single value-like variable (median r = 0.1, credible interval = [−0.22 to 0.41]). This finding confirms and provides more solid grounding for one of the central results of that earlier paper, which used simpler methods.

We did see a clear split between the task-selective cells and the no-selectivity cells (median no selectivity weight = 0.43; 95% credible interval = [0.23 to 0.63]). Thus it does not

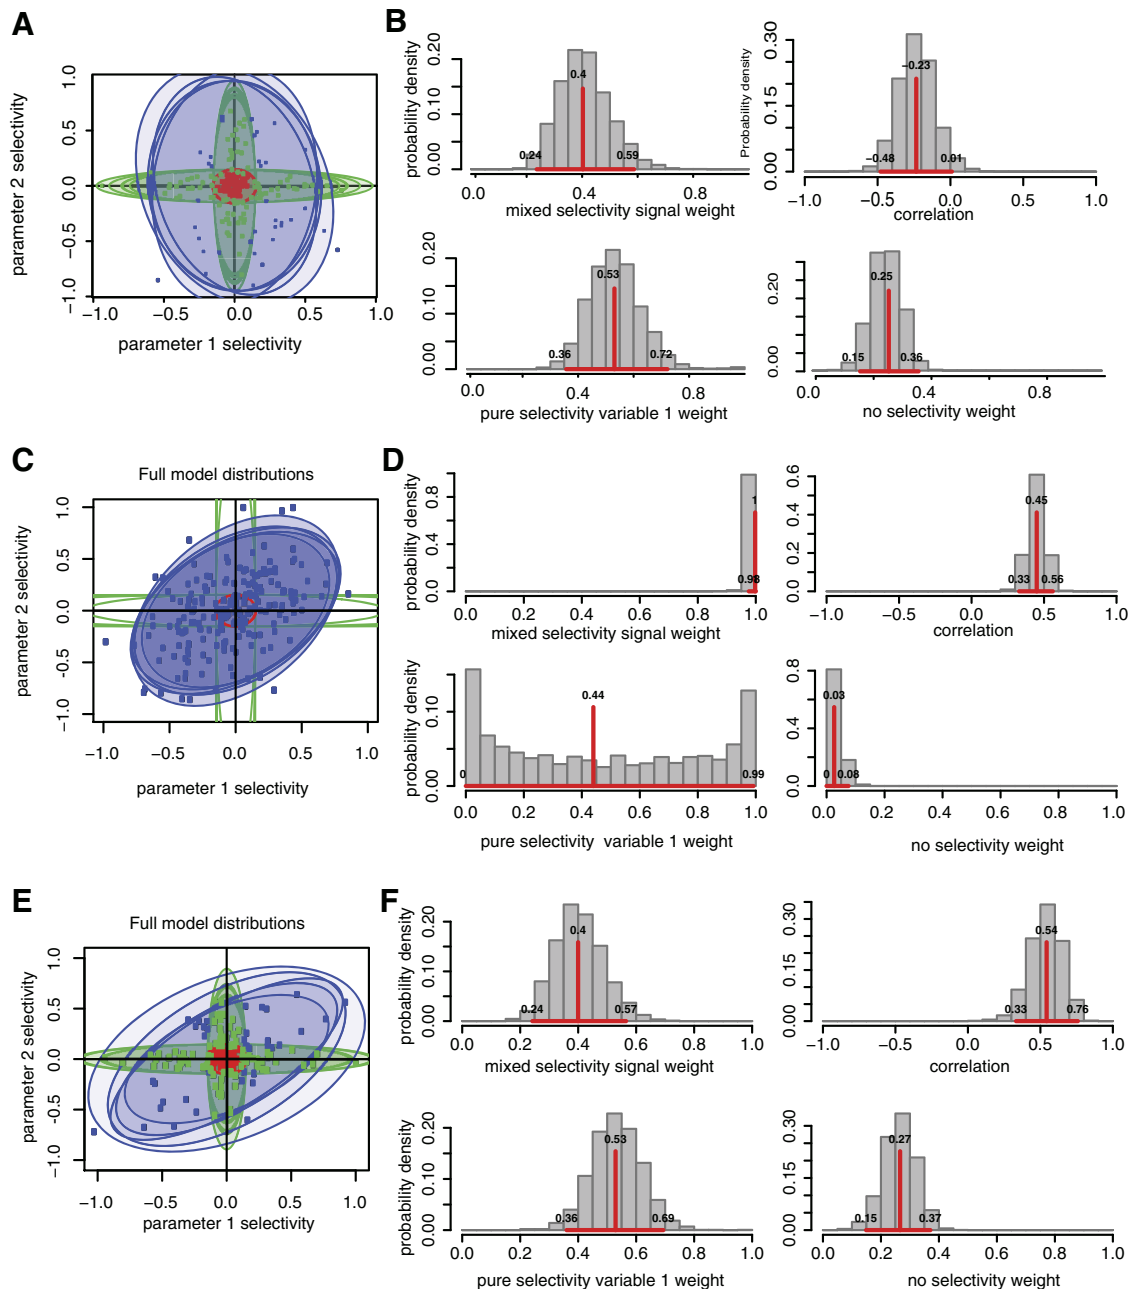


Fig. 4. Model fit to simulated data with correlated tunings. Model visualizations follow the same format as Fig. 2. *A*: model visualization for simulated data with a multiple tuned population with a 0.5 correlation between the tuning for X1 and X2. *B*: probability distribution functions of posteriors for *A*. *C* and *D*: same as *A* and *B*, but with a multiple and nonselective set. *E* and *F*: same as *A* and *B*, but with the population split between no tuning, pure tuning, and multiple tuning.

appear to be the case that untuned cells are simply tuned ones for which we did not collect enough data—the model tells us that we had enough data to see an effect had it been there, at least in a substantial number of these cells. This result was not reported in that paper (indeed, the methods we used could not detect it; Blanchard et al. 2015a).

Note that our new results, which rely on Bayesian statistics, are qualitatively different from those derived from conventional statistics because they permit inferences that take into account our uncertainty over whether or not a neuron is task relevant. In methods that attempt to classify neurons based on a statistical threshold, the evidence provided by many subthreshold neurons (e.g., those that do not differ significantly from zero) has no effect;

in this analysis, the joint inference of each neuron's classification and relevance uses as much information from the data as possible, making the analysis much more sensitive.

Another important question in OFC is whether neurons are selective for the identity of specific offers (a labeled line code), or whether they use a single consistent format to encode the values of each of the two offers. Labeled line codes are important in many models of economic choice (Chau et al. 2014; Hayden and Moreno-Bote 2017; Louie et al. 2011; Rustichini and Padoa-Schioppa 2015; Soltani et al. 2006) but recent work suggests they may not be an accurate description of neurons in decision-making structures like OFC/vmPFC (Blanchard et al. 2015a; Rich and Wallis 2016; Strait et al.

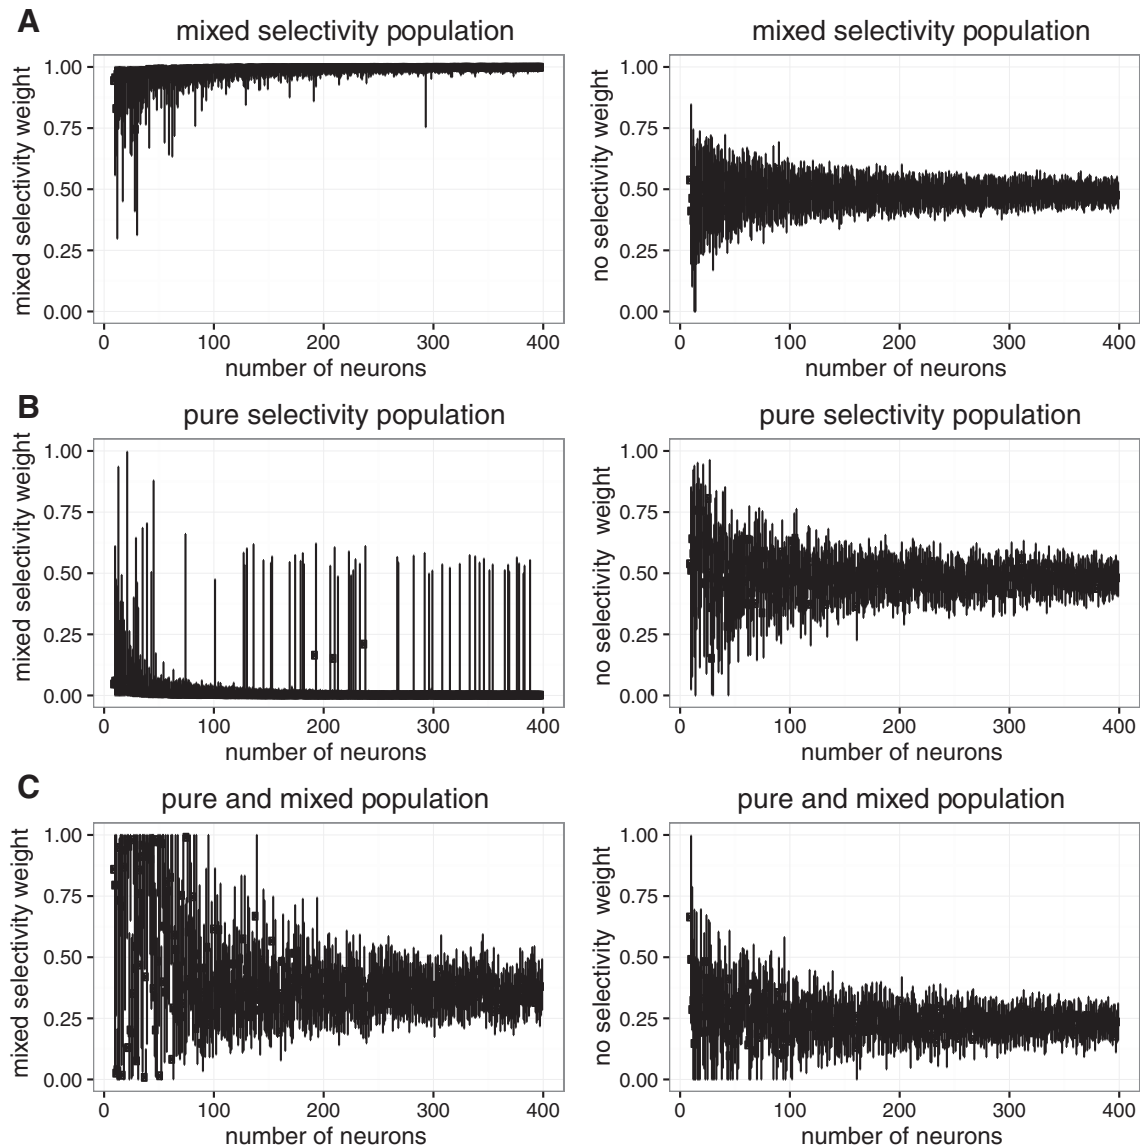


Fig. 5. Weights converge quickly to correct answer as more data is given. Multiple selectivity (*left*) and no selectivity (*right*) weights for a multiple/no selectivity population (A), a pure/no selectivity population (B), and a multiple/pure/no selectivity population (C).

2014;) and parietal cortex (Raposo et al. 2014). We first examined whether OFC neurons show categorical tuning based on side (neurons coding the left offer and neurons coding the right offer), as in LIP (Gold and Shadlen 2007; Platt and Glimcher 1999). Value here was defined based on the offered reward size (Blanchard et al. 2015a.)

Our analysis provides strong evidence against categorical selectivity (median multiple tuning weight = 1, 95% credible interval = [0.94 to 1]). Strength of tuning for left and for right offers is highly positively correlated (median correlation = 0.71, 95% credible interval = [0.37 to 0.95]). Another related possibility would be that neurons are selective for the value of the first and second offer, respectively. Our results argue strongly against categorization by side as well (median multiple tuning weight = 0.98, 95% credible interval = [0.87 to 1]). Instead, neurons that are selective for the value of the first offer are much more likely to be the ones selective for the value of the second offer. A different possibility would be neurons sensitive to first offer value and chosen value; we see no evidence

for such a distinction in OFC (median multiple tuning weight = 0.97, 95% credible interval = [0.92 to 1]). In other words, the neurons were consistently responding to the values of the offer currently presented whether it appeared first or second, the left or the right, or the offer or the chosen. The evidence points against a labeled-line coding scheme.

Strong Evidence Against Categorical Selectivity in Other Data Sets

We then considered a total of twenty pairs of variables coming from five data sets from four different brain regions (Fig. 8). We looked at firing rates of neurons in vmPFC and VS in a gambling task (Strait et al. 2014, 2015), dACC in a diet selection task (Blanchard and Hayden 2014; Blanchard et al. 2015b) and in a token gambling task (Azab and Hayden 2016, 2017; Strait et al. 2016), and OFC neurons in a riskless choice task (Wang and Hayden 2017). In all cases, monkeys made economic decisions based on combining different, orthogonal,

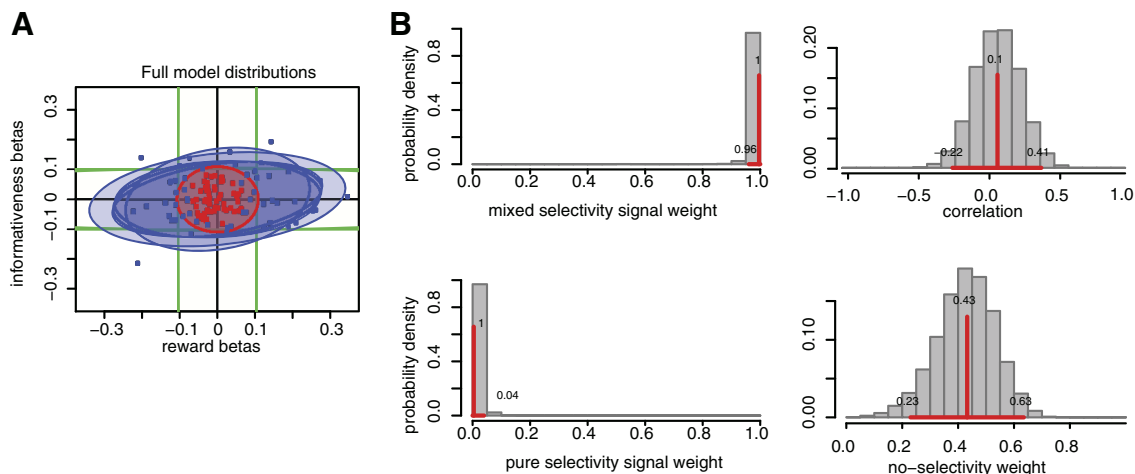


Fig. 6. Model fit to orbitofrontal cortex (OFC) data set 1 (curiosity tradeoff task). *A*: we fit regression weights for reward sensitivity (*x*-axis) and informativeness (*y*-axis). Cells were categorized either as multiple tuning (blue oval) or as nonselective (red oval). *B*: posteriors. Most weight was put onto the multiple-tuning class, indicating neurons selectivity for both variables. The correlation between these variables was not found to be significant; it overlapped with zero. Pure selectivity signal weight applied to few cells. Data fit to no-tuning category shows a discrete cluster of cells that were not sensitive to either variable.

dimensions. There were several differences in the tasks as well; see MATERIALS AND METHODS and Fig. 3 for details. However the structure was similar; in particular, all pairs of component variables were uncorrelated by the design of the task.

In all cases, we saw the same basic pattern: neurons are categorized into two sets, a single task-sensitive set and a

no-selectivity set (Fig. 7). The task-sensitive set consisted of neurons differentially sensitive to all tested task factors and the no-selectivity set was, as far as we could tell, not modulated by the variables we chose. We did not see any pairs of variables for which there are clear neuronal classes. Indeed, for all of our data sets values of over 0.05 for the pure-selective component




Fig. 7. Summary of population data. Plots of weights given by the model to each of 20 neural data sets (see Fig. 7). Bars indicate credible intervals. *A*: weights given to multiple selectivity are high and overlap with 1.0 in all cases. *B*: weights given to pure selectivity are weak and overlap with zero in all cases. *C*: weights given to no selectivity are positive and do not include 0 in most cases.

fell outside of the 95% credible interval, meaning that with $P < 0.05$ confidence we can reject more than 5% of neurons being purely selective (Fig. 7B).

Estimating the Proportion of No-Selectivity Neurons

One benefit of our method is that it allows us to estimate the proportion of no-selectivity neurons (that is, neurons whose responses are not selective for the chosen task variables) in our population. We know of no previous estimate for this parameter in frontostriatal regions. Conventional analysis methods make it impossible to know whether nonselective cells are truly not task related, or whether more trials would allow such neurons to pass a significance threshold. We report these proportions in Fig. 7C.

In analyses of physiological data, correlation is often used to detect relationships between tuning to two different variables. For example, correlations between regression weights can provide information about integrated vs. disjoint (i.e., multiplexed) coding schemes. One limitation of this approach is that it underestimates the true correlation, raising the possibility of Type II errors. Specifically, a correlation using the entire population will produce an estimate biased by neurons that are not task selective.

Our method allows us to more accurately compute the correlation between variables by weighting each neuron in the correlation by its probability of being part of the multiple tuning component. Thus, neurons with low likelihood of showing task-relevance can have a weaker influence on the computed correlation, in proportion to their likelihood. Because these neurons are not task driven, they can be expected to add noise to the correlation estimate. Thus, we expect correlations based on our method to be stronger than those estimated using conventional approaches (and also to be more accurate). Indeed, this is what we found. For all data sets with significant correlations (positive or negative), our method estimates a stronger correlation than the standard method (Fig. 8). A sign test on the magnitudes of the effects confirmed that this difference between the models is significant ($P < 0.001$). Of

the 20 individual cases, 15 showed a significant individual effect as well ($P < 0.05$, permutation test). The cases that showed a nonsignificant trend were EV of the first/chosen option (VS), EV of the first/chosen option (vmPFC), reward size and informativeness of the first option (OFC), reward size of first and second option (OFC), and reward size of first and chosen option (OFC).

Public Disposition of Our Code and Data

All code relevant for using our analysis procedure is on GitHub, along with instructions for use. All data analyzed here are also available there. The address is <https://github.com/TommyBlanchard/StanNeuronModelling>.

DISCUSSION

A population of cells can, in principle, be organized into functionally discrete clusters with categorically distinct response patterns. We developed a new statistical method that estimates the proportion of neurons that belong, respectively, to no-selectivity (not sensitive to chosen task parameters), pure-selectivity (firing selective for a single task parameter), and/or multiple-selectivity (selective for two task parameters simultaneously) categories. While these are relatively standard methods in statistics, they have not, so far as we know, been used to understand categorization of neuronal subtypes. Re-analysis of our own data from previous studies reveals the presence of two discrete categories consistently for several brain areas and comparisons: one category containing task-nonspecific cells and one containing neurons selective for task variables. These findings raise the possibility that, for many variables of interest to neuroscientists, populations of task-selective neurons may be category free (Raposo et al. 2014).

One unexpected finding of our analysis is that many neurons fall into the category of no-selectivity neurons—that is, the algorithm believes it has enough data to say with confidence that they are nonselective for task variables, and not simply lacking enough data. These neurons presumably have some

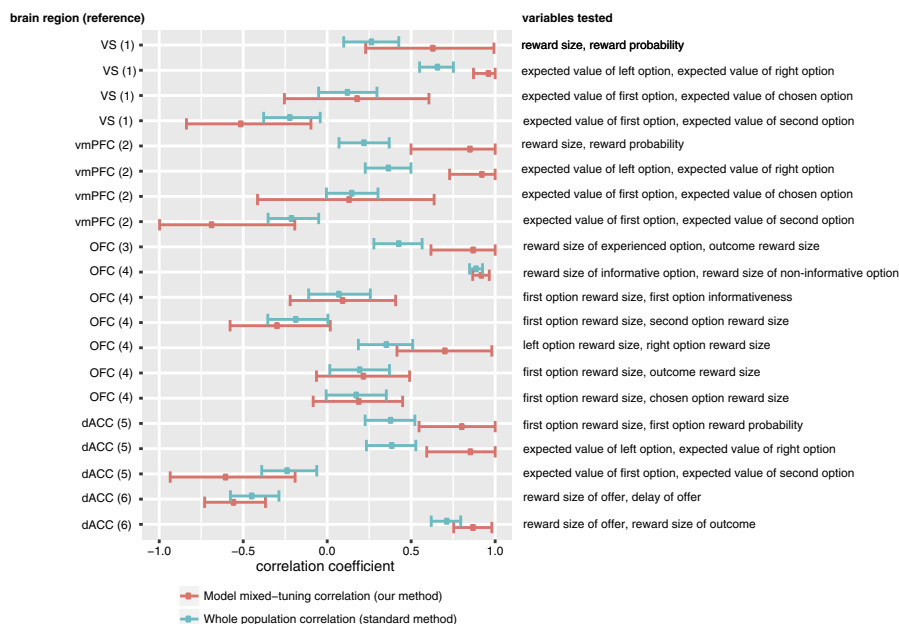


Fig. 8. Illustration of measured correlations and confidence intervals using standard methods (green bars) and using our methods described here. In all cases with significant effects, measured correlations were more extreme (farther from zero) using our method. Brain areas are indicated on the left as are references. Names of variables are indicated on the right. dACC, dorsal anterior cingulate cortex; vmPFC, ventromedial prefrontal cortex; VS, ventral striatum. References: (1) Strait et al. 2015; (2) Strait et al. 2014; (3) Wang and Hayden 2017; (4) Blanchard et al. 2015; (5) Strait et al. 2016; (6) Blanchard and Hayden 2014.

function that we did not elicit. For example, our tasks were all oculomotor; perhaps they were purely manual modality cells. In any case, identifying the functional role of these cells is an important goal for future studies. One immediate benefit of this finding is that, by assessing the amount of noise associated with a given neuron, we can more accurately estimate population parameters, such as the correlations between tuning parameters in neuronal data. In any case, the identification of such cells presents a great opportunity to improve estimates of parameter correlations among driven cells in a population. It also provides the opportunity to address questions about neurons in regions that are associated with weak driving, such as posterior cingulate cortex and ventromedial prefrontal cortex—it may be that we have simply not collected enough data in these regions, or it may be that they contain neurons that are simply not task selective (Hayden et al. 2010b).

Neurons are often selective for multiple variables at the same time; when these tuning functions interact, this property is known as mixed selectivity (Barak et al. 2013; Fusi et al. 2016; Rigotti et al. 2013). While mixed selectivity is not the same as category-free similarity—it allows for both categorical and noncategorical selectivities—it has a great deal in common. Mixed selectivity has many convenient features, including scalability, flexibility, robustness to damage, and the lack of need for precise wiring (Barak et al. 2013; Fusi et al. 2016; Raposo et al. 2014; Rigotti et al. 2013). Our results are also consistent with a view in which ensemble responses provide the best description of neuronal decision-making and control systems, one that is largely independent of the properties of elemental cells (Churchland et al. 2012; Hyman et al. 2012; Kristan and Shaw 1997; Ma et al. 2014; Mante et al. 2013; Stokes et al. 2013).

One method in the literature provides a formal text for categorically using frequentist approaches. The PAIRS test provides an important tool and important evidence for noncategorical tuning (Raposo et al. 2014). That study also demonstrated category-free selectivity in one region, the rodent parietal cortex. Our approach confirms and extends the ideas from that paper. First, in that work, which uses a null-hypothesis significance testing approach, the null hypothesis is that there are no categories. This means that a nonsignificant result, and a finding of “no categories” provides a somewhat limited conclusion, as failures to reject the null must. Our Bayesian approach gets around this limitation and allows us to assign a confidence to the hypothesis that there are no categories. Second, our method is flexibly able to deal with any relationship between the tuning properties of neurons to different variables. Thus it is more general and can deal with more variety of population structures. Third, our method identifies, with confidence levels, information about what neuron types compose a population. This in turn allows more accurate measures of correlations between variables, and potentially more robust classifications. These measures can be used to test hypotheses about the relationship between cell types and firing rates, for example. Finally, the PAIRS test does not deal in a principled way with the possibility that only a subset of neurons measured may be task relevant; that is, it ignores a possible no-selectivity cluster. Nonetheless, the PAIRS test represents an important and valuable tool for the field, one that is complementary to ours. Moreover, the results of that test in

rodent parietal cortex present an important challenge to the idea of categorical tuning, one that our results echo.

It is very common to perform statistical hypothesis testing and then do further analyses based on the results of such tests. Such a practice has several disadvantages relative to our approach. First, it uses a fixed statistical significance threshold (normally $\alpha = 0.05$) to categorize data as significant and nonsignificant. Doing so reduces a continuum of information into a binary and thus throws out a good deal of meaningful information and reduces sensitivity. Further analyses on such data necessarily produces Type II errors, and can, in some cases produce Type I errors (Maxwell and Delaney 1993; MacCallum et al. 2002). The standard approach likely also introduces a potential nonindependence error: by removing cells that do not achieve a given P value, the remaining cells are certain to be nonrepresentative of the population in an anticongervative direction: examination of only cells that pass the threshold can give results that appear, incorrectly, to be statistically strong. A similar problem is known in neuroimaging as “double-dipping,” which selects the brain region of interest with one statistical test and subsequently conducts nonindependent analysis on the selected region (Kriegeskorte et al. 2009). Our methods offer a new and direct way around this problem. Finally, in cases where the no-selectivity neurons are scientifically important, a conservative cutoff (like $P < 0.05$) is also likely to overestimate the proportion of no-selectivity cells in the population, which can in turn lead to false positives (Humphreys and Fleishman 1974; Maxwell et al. 1984).

It is important to describe several limitations of our modeling approach to define the boundaries of where our method is most usefully applied and to point toward further extensions of the general approach. First, the model assumes a specific form of mixture components (nonselective, pure selective, and multiple selective). If there were a population that did not fit into these categories—for instance, two distributions with different orientations—then the population would not be well modeled by our existing framework. Similarly, if the distributions did not have zero mean or were not unimodal, the present method would not be appropriate. (Note that our data were z -scored and thus forced to have a zero mean). This limitation may not be as bad as it seems. Any such data could be incorporated in principle into a mixture model like ours by mathematically defining and including the corresponding distribution with its associated free parameters. For instance, numerous distributions with free orientations and means could be used in situations where there are few specific hypotheses about how neural responses are likely to behave. In this way, our approach is an instantiation of a more general technique that is increasingly common, Bayesian data analysis, where a family of hypothesized generative models may be written down, and the data can be used to optimally infer how likely each model or parameter value is. This family of extensions of our model which focuses on modeling the underlying components of the data is likely to yield deep insights into the organization of neural systems. A third limitation is that our approach does not work with variables that are themselves correlated. For example, suppose that we were interested in coding of mathematical expected value (the product of probability and stakes) and probability—these two variables are

themselves strongly correlated with each other. Our analysis could not deal well with this condition.

There are also limitations on the neural data that we present. First, our data come from a relatively narrow range of tasks and from a relatively narrow range of brain regions. These limitations reflect the nature of the data we had available to us to use, and reflect the importance of the questions to our laboratory. Nonetheless, these are all association regions—neither sensory nor motor—as such our results do not provide information on whether such regions (or indeed other association regions) would provide categorical responses. A second limitation of the data we present is that we did not examine all possible variables that could affect firing. In particular we did not examine ostensibly primitive variables like sensory and motor ones. These basic variables may have a special and distinct representation in reward regions. Unfortunately, the experiments we used were not designed to disambiguate these variables from other confounded ones, especially related to value. Finally, a limitation of our neural data is that we cannot say for certain that we were effective in driving the neurons in our samples. It may be the case that the parameters we used as regressors were ones along which neurons were weakly selective. It is possible that other parameters would have produced much greater neuronal responses and perhaps would have shown categories.

ACKNOWLEDGMENTS

We thank Habiba Azab, Meghan Castagno, Giuliana Loconte, Marc Mancarella, Brianna Sleezer, Caleb Strait, and Maya Wang for help with data collection and organization.

GRANTS

This work was supported by an R01 (DA037229) from the National Institute on Drug Abuse (<https://www.drugabuse.gov>) to B. Y. Hayden.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

T.C.B., S.T.P., and B.Y.H. conceived and designed research; T.C.B. and S.T.P. performed experiments; T.C.B. and S.T.P. analyzed data; T.C.B., S.T.P., and B.Y.H. interpreted results of experiments; T.C.B., S.T.P., and B.Y.H. prepared figures; T.C.B., S.T.P., and B.Y.H. drafted manuscript; T.C.B., S.T.P., and B.Y.H. edited and revised manuscript; T.C.B., S.T.P., and B.Y.H. approved final version of manuscript.

REFERENCES

- Azab H, Hayden BY. Shared roles of dorsal and subgenual anterior cingulate cortices in economic decisions (Preprint). *bioRxiv* 074484, 2016. doi:10.1101/074484.
- Azab H, Hayden BY. Correlates of decisional dynamics in the dorsal anterior cingulate cortex. *PLoS Biol* 15: e2003091, 2017. doi:10.1371/journal.pbio.2003091.
- Barak O, Rigotti M, Fusi S. The sparseness of mixed selectivity neurons controls the generalization-discrimination trade-off. *J Neurosci* 33: 3844–3856, 2013. doi:10.1523/JNEUROSCI.2753-12.2013.
- Blanchard TC. A general method for estimating categorical composition of neural data. 2016. <https://github.com/TommyBlanchard/StanNeuronModelling>.
- Blanchard TC, Hayden BY. Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *J Neurosci* 34: 646–655, 2014. doi:10.1523/JNEUROSCI.3151-13.2014.
- Blanchard TC, Hayden BY. Monkeys are more patient in a foraging task than in a standard intertemporal choice task. *PLoS One* 10: e0117057, 2015. doi:10.1371/journal.pone.0117057.
- Blanchard TC, Hayden BY, Bromberg-Martin ES. Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* 85: 602–614, 2015a. doi:10.1016/j.neuron.2014.12.050.
- Blanchard TC, Pearson JM, Hayden BY. Postreward delays and systematic biases in measures of animal temporal discounting. *Proc Natl Acad Sci USA* 110: 15491–15496, 2013. doi:10.1073/pnas.1310446110.
- Blanchard TC, Strait CE, Hayden BY. Ramping ensemble activity in dorsal anterior cingulate neurons during persistent commitment to a decision. *J Neurophysiol* 114: 2439–2449, 2015b. doi:10.1152/jn.00711.2015.
- Blanchard TC, Wilke A, Hayden BY. Hot-hand bias in rhesus monkeys. *J Exp Psychol Anim Learn Cogn* 40: 280–286, 2014. doi:10.1037/xan0000033.
- Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A. Stan: a probabilistic programming language. *J Stat Softw* 76: 1–32, 2017. doi:10.18637/jss.v076.i01.
- Chau BKH, Kolling N, Hunt LT, Walton ME, Rushworth MFS. A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nat Neurosci* 17: 463–470, 2014. doi:10.1038/nn.3649.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. Neural population dynamics during reaching. *Nature* 487: 51–56, 2012. doi:10.1038/nature11129.
- Delaney H, Maxwell SE, Delaney HD. Bivariate median splits and spurious statistical significance bivariate median splits and spurious statistical significance. *Psychol Bull* 113: 181–190, 1993. doi:10.1037/0033-2909.113.1.181
- Fusi S, Miller EK, Rigotti M. Why neurons mix: high dimensionality for higher cognition. *Curr Opin Neurobiol* 37: 66–74, 2016. doi:10.1016/j.conb.2016.01.010.
- Ganguli S, Sompolinsky H. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annu Rev Neurosci* 35: 485–508, 2012. doi:10.1146/annurev-neuro-062111-150410.
- Gelman A, Rubin D. Inference from iterative simulation using multiple sequences. *Stat Sci* 7: 457–511, 1992. doi:10.1214/ss/1177011136.
- Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci* 30: 535–574, 2007. doi:10.1146/annurev.neuro.29.051605.113038.
- Hayden BY, Heilbronner SR, Platt ML. Ambiguity aversion in rhesus macaques. *Front Neurosci* 4: 166, 2010a. doi:10.3389/fnins.2010.00166.
- Hayden BY, Moreno-Bote R. A neuronal theory of sequential economic choice (Preprint). *bioRxiv* 221135, 2017. doi:10.1101/221135.
- Hayden BY, Smith DV, Platt ML. Cognitive control signals in posterior cingulate cortex. *Front Hum Neurosci* 4: 223, 2010b. doi:10.3389/fnhum.2010.00223.
- Heilbronner SR, Hayden BY. Dorsal anterior cingulate cortex: a bottom-up view. *Annu Rev Neurosci* 39: 149–170, 2016. doi:10.1146/annurev-neuro-070815-013952.
- Hoffman MD, Gelman A. The no-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J Mach Learn Res* 15: 1351–1381, 2014.
- Humphreys LG, Fleishman A. Pseudo-orthogonal and other analysis of variance designs involving individual-differences variables. *J Educ Psychol* 66: 464–472, 1974. doi:10.1037/h0036539.
- Hunt LT, Behrens TEJ, Hosokawa T, Wallis JD, Kennerley SW. Capturing the temporal evolution of choice across prefrontal cortex. *eLife* 4: 1–25, 2015. doi:10.7554/eLife.11945.
- Hunt LT, Hayden BY. A distributed, hierarchical and recurrent framework for reward-based choice. *Nat Rev Neurosci* 18: 172–182, 2017. doi:10.1038/nrn.2017.7.
- Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proc Natl Acad Sci USA* 109: 5086–5091, 2012. doi:10.1073/pnas.1114415109.
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21: 1162–1178, 2009. doi:10.1162/jocn.2009.21100.
- Kidd C, Hayden BY. The psychology and neuroscience of curiosity. *Neuron* 88: 449–460, 2015. doi:10.1016/j.neuron.2015.09.010.
- Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci* 12: 535–540, 2009. doi:10.1038/nn.2303.
- Kristan WB Jr, Shaw BK. Population coding and behavioral choice. *Curr Opin Neurobiol* 7: 826–831, 1997. doi:10.1016/S0959-4388(97)80142-0.

- Kvitsiani D, Ranade S, Hangya B, Taniguchi H, Huang JZ, Kepecs A.** Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature* 498: 363–366, 2013. doi:10.1038/nature12176.
- Lewandowski D, Kurowicka D, Joe H.** Generating random correlation matrices based on vines and extended onion method. *J Multivariate Anal* 100: 1989–2001, 2009. doi:10.1016/j.jmva.2009.04.008.
- Louie K, Grattan LE, Glimcher PW.** Reward value-based gain control: divisive normalization in parietal cortex. *J Neurosci* 31: 10627–10639, 2011. doi:10.1523/JNEUROSCI.1237-11.2011.
- Ma L, Hyman JM, Phillips AG, Seamans JK.** Tracking progress toward a goal in corticostriatal ensembles. *J Neurosci* 34: 2244–2253, 2014. doi:10.1523/JNEUROSCI.3834-13.2014.
- MacCallum RC, Zhang S, Preacher KJ, Rucker DD.** On the practice of dichotomization of quantitative variables. *Psychol Methods* 7: 19–40, 2002. doi:10.1037/1082-989X.7.1.19.
- Mante V, Sussillo D, Shenoy KV, Newsome WT.** Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503: 78–84, 2013. doi:10.1038/nature12742.
- Maxwell SE, Delaney HD.** Bivariate median splits and spurious statistical significance. *Psychol Bull* 113: 181–190, 1993. doi:10.1037/0033-2909.113.1.181.
- Maxwell SE, Delaney HD, Dill CA.** Another look at ANCOVA versus blocking. *Psychol Bull* 95: 136–147, 1984. doi:10.1037/0033-2909.95.1.136.
- Neal RM.** MCMC using Hamiltonian dynamics. In: *Handbook of Markov Chain Monte Carlo*, edited by Brooks S, Gelman A, Jones G, Meng S. New York: Chapman and Hall, 2011, chapt. 5, p. 113–162. doi:10.1201/b10905-6.
- Pearson JM, Hayden BY, Platt ML.** Explicit information reduces discounting behavior in monkeys. *Front Psychol* 1: 237, 2010. doi:10.3389/fpsyg.2010.00237.
- Peel D, McLachlan GJ.** Robust mixture modelling using the t distribution. *Stat Comput* 10: 339–348, 2000. doi:10.1023/A:1008981510081.
- Platt ML, Glimcher PW.** Neural correlates of decision variables in parietal cortex. *Nature* 400: 233–238, 1999. doi:10.1038/22268.
- Pouget A, Sejnowski TJ.** Spatial transformations in the parietal cortex using basis functions. *J Cogn Neurosci* 9: 222–237, 1997. doi:10.1162/jocn.1997.9.2.222.
- Raposo D, Kaufman MT, Churchland AK.** A category-free neural population supports evolving demands during decision-making. *Nat Neurosci* 17: 1784–1792, 2014. doi:10.1038/nn.3865.
- Rich EL, Wallis JD.** Decoding subjective decisions from orbitofrontal cortex. *Nat Neurosci* 19: 973–980, 2016. doi:10.1038/nn.4320.
- Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S.** The importance of mixed selectivity in complex cognitive tasks. *Nature* 497: 585–590, 2013. doi:10.1038/nature12160.
- Rigotti M, Ben Dayan Rubin D, Wang X-J, Fusi S.** Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Front Comput Neurosci* 4: 24, 2010. doi:10.3389/fncom.2010.00024.
- Rustichini A, Padoa-Schioppa C.** A neuro-computational model of economic decisions. *J Neurophysiol* 114: 1382–1398, 2015. doi:10.1152/jn.00184.2015.
- Seo H, Lee D.** Behavioral and neural changes after gains and losses of conditioned reinforcers. *J Neurosci* 29: 3627–3641, 2009. doi:10.1523/JNEUROSCI.4726-08.2009.
- Sleezer BJ, LoConte GA, Castagno MD, Hayden BY.** Neuronal responses support a role for orbitofrontal cortex in cognitive set reconfiguration. *Eur J Neurosci* 45: 940–951, 2017. doi:10.1111/ejn.13532.
- Soltani A, Lee D, Wang XJ.** Neural mechanism for stochastic behaviour during a competitive game. *Neural Netw* 19: 1075–1090, 2006. doi:10.1016/j.neunet.2006.05.044.
- Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J.** Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78: 364–375, 2013. doi:10.1016/j.neuron.2013.01.039.
- Strait CE, Blanchard TC, Hayden BY.** Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. *Neuron* 82: 1357–1366, 2014. doi:10.1016/j.neuron.2014.04.032.
- Strait CE, Sleezer BJ, Blanchard TC, Azab H, Castagno MD, Hayden BY.** Neuronal selectivity for spatial positions of offers and choices in five reward regions. *J Neurophysiol* 115: 1098–1111, 2016. doi:10.1152/jn.00325.2015.
- Strait CE, Sleezer BJ, Hayden BY.** Signatures of value comparison in ventral striatum neurons. *PLoS Biol* 13: e1002173, 2015. doi:10.1371/journal.pbio.1002173.
- Tsujimoto S, Genovesio A, Wise SP.** Comparison of strategy signals in the dorsolateral and orbital prefrontal cortex. *J Neurosci* 31: 4583–4592, 2011. doi:10.1523/JNEUROSCI.5816-10.2011.
- Wallis JD, Anderson KC, Miller EK.** Single neurons in prefrontal cortex encode abstract rules. *Nature* 411: 953–956, 2001. doi:10.1038/35082081.
- Wang MZ, Hayden BY.** Reactivation of associative structure specific outcome responses during prospective evaluation in reward-based choices. *Nat Commun* 8: 15821, 2017. doi:10.1038/ncomms15821.