

Research Articles: Systems/Circuits

Abstract value encoding in neural populations but not single neurons

<https://doi.org/10.1523/JNEUROSCI.1954-22.2023>

Cite as: J. Neurosci 2023; 10.1523/JNEUROSCI.1954-22.2023

Received: 18 October 2022

Revised: 7 May 2023

Accepted: 11 May 2023

This Early Release article has been peer-reviewed and accepted, but has not been through the composition and copyediting processes. The final version may differ slightly in style or formatting and will contain links to any extended data.

Alerts: Sign up at www.jneurosci.org/alerts to receive customized email alerts when the fully formatted version of this article is published.

1 **Title:** Abstract value encoding in neural populations but not single
2 neurons

3
4 **Abbreviated title:** Transformable subspaces in economic choice

5
6 **Authors:** Justin M. Fine, David J-N. Maissou, Seng Bum Michael Yoo, Tyler V. Cash-Padgett,
7 Maya Zhe Wang, Jan Zimmermann, and Benjamin Y. Hayden

8
9 **Affiliation**

10 Department of Neuroscience and Center for Magnetic Resonance Research
11 University of Minnesota, Minneapolis MN 55455

12
13 **Corresponding author**

14 Benjamin Yost Hayden
15 Email: benhayden@gmail.com
16 Current address: Department of Neurosurgery
17 Baylor College of Medicine
18 Houston, TX, 77040

19
20 **46 pages, 8 figures, 0 tables**

21 **Abstract:** 193 words

22 **Introduction:** 689 words

23 **Discussion:** 1392 words

24
25 **Conflict of interest:** None

26
27 **Acknowledgements**

28 We thank Geoffrey Schoenbaum for helpful discussions. We thank Marc Mancarella,
29 Caleb Strait and Tommy Blanchard for assistance with data collection, Sarah Heilbronner for
30 help with anatomy, and the rest of the Hayden and Zimmermann labs for valuable discussions.
31 Funding provided by National Institute on Drug Abuse grant P30 DA048742-01A1 (to BYH and
32 JZ), National Institute for Biomedical Imaging Grant P41 EB027061 (to BYH and JZ), and the
33 UMN AIRP award (to BYH and JZ).

34
35

36

ABSTRACT

37

38

39

40

41

42

43

44

45

46

47

48

49

50

An important open question in neuroeconomics is how the brain represents the value of offers in a way that is both abstract (allowing for comparison) and concrete (preserving the details of the factors that influence value). Here we examine neuronal responses to risky and safe options in five brain regions that putatively encode value in male macaques. Surprisingly, we find no detectable overlap in the neural codes used for risky and safe options, even when the options have identical subjective values (as revealed by preference) in any of the regions. Indeed, responses are not just uncorrelated but occupy distinct (semi-orthogonal) encoding subspaces. Notably, however, these subspaces are linked through a linear transform of their constituent encodings, a property that allows for comparison of dissimilar option types. This encoding scheme allows these regions to have their cake and eat it too: they can encode the detailed factors that influence offer value (here, risky and safety) but also directly compare dissimilar offer types. Together these results suggest a neuronal basis for the qualitatively different psychological properties of risky and safe options and highlight the power of population geometry to resolve outstanding problems in neural coding.

51

SIGNIFICANCE STATEMENT

52

53

54

55

56

57

58

59

To make economic choices, we must have some mechanism for comparing dissimilar offers. We propose that the brain uses distinct neural codes for risky and safe offers, but that these codes are linearly transformable. This encoding scheme has the twin advantages of allowing for comparison across offer types while preserving information about offer type, which in turn allows for flexibility in changing circumstances. We show that responses to risky and safe offers exhibit these predicted properties in five different reward-sensitive regions. Together, these results highlight the power of population coding principles for solving representation problems in economic choice.

INTRODUCTION

60

61 We are often faced with the need to choose between options that differ qualitatively. To
62 make such choices, it is helpful to have access to an abstract representation of the value of each
63 option. Much of neuroeconomics is predicated on the assumption that such representations must
64 exist and, in particular, that they exist in the form of specialized *abstract value neurons* (Platt
65 and Glimcher, 1999; Padoa-Schioppa and Assad, 2006; Kennerley et al., 2009; Lau and
66 Glimcher, 2008; Peters and Buchel, 2010; Kolling et al., 2016). Such neurons would, by
67 definition, have firing rates that covary monotonically with the values of offers regardless of
68 their other qualities. That is, two options with identical subjective values would elicit the same
69 firing rate response in an abstract value neuron even if they differed in other ways. For example,
70 if asked to evaluate a Ferris wheel ride and an equally valued cupcake, a value neuron will
71 necessarily have the same response to both. One advantage of the abstract neuron
72 representational system is that a downstream decoder can produce good choices simply by
73 identifying the neuron or neurons that encode the highest value and selecting the appropriate
74 choice. However, while such an encoding scheme has its advantages, it has a major weakness – it
75 is inflexible in situations where the relative importance of these features changes.

76 More broadly, the abstract value neuron idea is out of step with modern thinking on how
77 information is coded in populations of neurons. While the “neuron doctrine” is focused on the
78 idea that information explicitly encoded in the firing rates of single neurons, the “population
79 doctrine” emphasizes the flexible and expressive power of neuronal populations (Saxena and
80 Cunningham, 2019; Ebitz and Hayden, 2021). From the perspective of populations, abstraction
81 can come from geometry; specifically, from transformable subspaces (Elsayed et al., 2016; Tang
82 et al., 2020; Yoo and Hayden, 2020; Libby and Buschman, 2021). This way of generating

83 abstraction has the major advantage of allowing a population of neurons to simultaneously carry
84 detailed about the features of the offers and to have an abstract representation of value.

85 We hypothesized that neurons in core value regions use the population approach to
86 representing value. To test this hypothesis, we studied encoding of risky and safe options in a
87 two-option risky choice task. Psychologists have established that our minds treat risky and safe
88 options differently. We have a strong preference for safe options (Holt and Laury, 2002;
89 Kacelnik and Bateson, 2013; Heilbronner, 2017), risky options differentially activate
90 motivational and emotional factors (Loewenstein et al., 2001; Slovic et al., 2004 and 2007;
91 Lerner et al., 2015), and risk differs qualitatively from certainty among dimensions like the
92 prospect of learning and satisfaction of curiosity (Binde, 2013; Heilbronner and Hayden, 2013;
93 Wang and Hayden, 2021). Population geometric encoding would allow for value comparison
94 while also maintaining information that allows risk and surety to be treated differently. It is well
95 established that neural responses in core value regions scale with the probability (and thus the
96 subjective value) of risky offers (Kennerley et al., 2009; So and Stuphorn, 2010; Raghuraman
97 and Padoa-Schioppa, 2014; McCoy and Platt, 2005; Strait et al., 2014; Azab and Hayden, 2017).
98 Perhaps surprisingly, little of this research has addressed the critical question of whether *risky*
99 and *safe* offers use the same coding scale, and none of it has used a dense sampling of
100 probability space needed to retrospectively identify equally valued risky and safe offers and test
101 how their codes are related.

102 We examined responses of neurons in five core reward areas in a gambling task (**Figure**
103 **1A**). We found that responses of single neurons to safe offers are not just distinct but are
104 unrelated to responses evoked by risky offers. Instead, risky and safe offers are encoded by
105 overlapping sets of neurons using distinct and semi-orthogonal subspaces. At the same time,

106 these subspaces are mutually transformable. These results are consistent with the hypothesis that
107 the brain's core value regions strategically use subspace orthogonalization to flexibly partition
108 different offer types in a way that allows for value comparison. More generally, these results
109 endorse the utility of population perspectives in tackling classic problems in neuroscience.

110 **MATERIALS AND METHODS**

111 *Surgical procedures.* Subjects were male macaques. All procedures were approved by
112 either the University Committee on Animal Resources at the University of Rochester or the
113 IACUC at the University of Minnesota. Animal procedures were also designed and conducted in
114 compliance with the Public Health Service's Guide for the Care and Use of Animals. All of the
115 animals were handled according to approved institutional animal care and use committee
116 (IACUC) protocols (#2005-619 38127A) of the University of Minnesota. The protocol was
117 approved by the Committee on the Ethics of Animal Experiments of the University of Minnesota
118 (NIH permit number: A3456-01). All surgery was performed under controlled anesthesia. Six
119 male rhesus macaques (*Macaca mulatta*) served as subjects. A small prosthesis for head fixation
120 was used. Subjects were habituated to laboratory conditions and then trained to perform
121 oculomotor tasks for liquid rewards. We place a Cilux recording chamber (Crist Instruments)
122 over the area of interest (see Behavioral Tasks for breakdown). We verified positioning by
123 magnetic resonance imaging with the aid of a Brainsight system (Rogue Research). Animals
124 received appropriate analgesics and antibiotics after all procedures. Throughout both behavioral
125 and physiological recording sessions, we kept the chamber clean with regular antibiotic washes,
126 and we sealed them with sterile caps.

127 *Recording sites.* We approached our brain regions through standard recording grids (Crist
128 Instruments) guided by a micromanipulator (NAN Instruments). All recording sites were selected
129 based on the boundaries given in the Paxinos atlas (Paxinos et al., 2000). In all cases we sampled
130 evenly across the regions. Neuronal recordings in OFC were collected from *subjects P and S*
131 (Yoo and Hayden, 2020); recordings in vmPFC were collected from *subjects B and H* (Strait et
132 al., 2014); recordings in pgACC were collected from *subject B and V* (Maisson et al., 2021);

133 recordings from PCC were collected from *subject P and S* (Wang et al., 2021); and recording in
134 VS were collected from *subject B and C* (Strait et al., 2016). For details, see **Figure 1B**:

135 We defined **OFC 13** as lying within the coronal planes situated between 28.65 and 34.05
136 mm rostral to the interaural plane, the horizontal planes situated between 3 and 6.5 mm from the
137 brain's ventral surface, and the sagittal planes between 5 and 14 mm from the medial wall. The
138 coordinates correspond to area 13m in Paxinos et al. (2000).

139 We defined **vmPFC 14** as lying within the coronal planes situated between 29 and 44
140 mm rostral to the interaural plane, the horizontal planes situated between 0 and 9 mm from the
141 brain's ventral surface, and the sagittal planes between 0 and 8 mm from the medial wall. These
142 coordinates correspond to area 14m in Paxinos et al. (2000).

143 We defined **pgACC 32** as lying with the coronal planes situated between 30.90 and 40.10
144 mm rostral to the interaural plane, the horizontal planes situated between 7.30 and 15.50 mm
145 from the brain's dorsal surface, and the sagittal planes between 0 and 4.5 mm from the medial
146 wall (**Figure 1B**). Our recordings were made from central regions within these zones, which
147 correspond to area 32 in Paxinos et al. (2000).

148 We defined **PCC 29/31** as lying within the coronal planes situated between 2.88 mm
149 caudal and 15.6 mm rostral to the interaural plane, the horizontal planes situated between 16.5
150 and 22.5 mm from the brain's dorsal surface, and the sagittal planes between 0 and 6 mm from
151 the medial wall. The coordinates correspond to area 29/31 in Paxinos et al. (2000).

152 We defined **VS** as lying within the coronal planes situated between 20.66 and 28.02 mm
153 rostral to the interaural plane, the horizontal planes situated between 0 and 8.01 mm from the
154 ventral surface of the striatum, and the sagittal planes between 0 and 8.69 mm from the medial
155 wall. Note that our recording sites were targeted towards the nucleus accumbens core region of

156 the VS.

157 We confirmed recording location before each recording session using our Brainsight
158 system with structural magnetic resonance images taken before the experiment (0.5 voxels). We
159 confirmed recording locations by listening for characteristic sounds of white and gray matter
160 during recording, which in all cases matched the loci indicated by the Brainsight system with an
161 error of ~1 mm in the horizontal plane and ~2 mm in the z-direction.

162 *Electrophysiological techniques.* Either single (FHC) or multi-contact electrodes (V-
163 Probe, Plexon) were lowered using a microdrive (NAN Instruments). Individual action potentials
164 were isolated on a Plexon (Dallas, TX) or Ripple (Salt Lake City, UT). Neurons were selected
165 for study solely based on quality of isolation; we never preselected based on task-related
166 response properties. All collected neurons for which we managed to obtain at least 300 trials
167 were analyzed; no neurons that surpassed our isolation criteria were excluded from analysis.

168 *Eye-tracking and reward delivery.* Eye position was sampled at 1,000 Hz by an infrared
169 eye-monitoring camera system (SR Research). Stimuli were controlled by a computer running
170 Matlab (Mathworks) with Psychtoolbox and Eyelink Toolbox. Visual stimuli were colored
171 rectangles on a computer monitor placed 57 cm from the animal and centered on its eyes (Fig.
172 1A). A standard solenoid valve controlled the duration of juice delivery. Solenoid calibration was
173 performed daily.

174 *Behavioral tasks.* Six macaques performed in the risky choice task. Both tasks made use
175 of vertical rectangles indicating reward amount and probability. We have shown in a variety of
176 contexts that this method provides reliable communication of abstract concepts such as reward,
177 probability, delay, and rule to monkeys (Hayden et al., 2010; Blanchard et al., 2014).

178 *Risky choice task.* The task presented two offers on each trial. A rectangle 300 pixels tall

179 and 80 pixels wide represented each offer (11.35° of visual angle tall and 4.08° of visual angle
180 wide; Fig. 2A). Two parameters defined gamble offers, *stakes* and *probability*. Each gamble
181 rectangle was divided into two portions, one red and the other either gray, blue, or green. The
182 size of the color portions signified the probability of winning a small (125 μ L, gray), medium
183 (165 μ L, blue), or large reward (240 μ L, green), respectively. We used a uniform distribution
184 between 0 and 100% for probabilities. The size of the red portion indicated the probability of no
185 reward. Offer types were selected at random with a 43.75% probability of blue (medium-stakes)
186 gamble, a 43.75% probability of green (high-stakes) gambles, and a 12.5% probability of gray
187 options (safe offers).

188 On each trial, one offer appeared on the left side of the screen and the other appeared on
189 the right. We randomized the sides of the first and second offer. Both offers appeared for 400 ms
190 and were followed by a 600 ms blank period. After the offers were presented separately, a central
191 fixation spot appeared, and the subject fixated on it for 100 ms. Next, both offers appeared
192 simultaneously and the animal indicated its choice by shifting gaze to its preferred offer and
193 maintaining fixation on it for 200 ms. Failure to maintain gaze for 200 ms did not lead to the end
194 of the trial but instead returned the monkey to a choice state; thus subjects were free to change
195 their mind if they did so within 200 ms (although in our observations, they seldom did so).
196 Following a successful 200-ms fixation, the gamble was resolved, and the reward was delivered.
197 We defined trials that took > 7 sec as inattentive trials and we did not include them in the
198 analyses (this removed $\sim 1\%$ of trials). Outcomes that yielded rewards were accompanied by a
199 visual cue: a white circle in the center of the chosen offer. All trials were followed by an 800 ms
200 intertrial interval with a blank screen.

201 *Estimation of subjective value equivalence.* We calculated the indifference point between

202 safe and risky offers. For each subject, independently, we fitted a sigmoidal function to the
203 distribution of choices (safe or risky) across the full range of risky offer probabilities.

$$f(x) = \frac{a}{b + e^{bx}}$$

204 where x = the probability associated with the risky offer, and $f(x)$ = the likelihood of choosing
205 the safe offer; a (the maximum value of the curve) and b (the growth rate; steepness) are
206 coefficients of the function, estimated by the fitting procedure for maximizing R^2 . Using the
207 fitted sigmoidal function, we then estimated the value of x needed to produce a safe choice
208 likelihood of 0.5; that is, a risky offer choice is equally likely as a safe choice. This is called the
209 *indifference point*. We performed this analysis separately for medium- and high-stakes gambles,
210 and separately for each subject. Note that these curves are the result of a fitting procedure that
211 has greater error in fitting the asymptotes. This is due to the small number of variant trials at
212 these asymptotes that the algorithm is highly sensitive to outliers at these values.

213 *Statistical methods.* We constructed peristimulus time histograms by aligning spike
214 rasters to the presentation of the first offer and averaging firing rates across multiple trials. We
215 calculated firing rates in 20 ms bins, but we analyzed them in longer (500 ms) epochs. Some
216 statistical tests of neuronal activity were only appropriate when applied to single neurons
217 because of variations in response properties across the population. We have used this epoch in all
218 our past research on this and similar tasks (Strait et al., 2014, 2015, 2016; Azab and Hayden,
219 2017, 2018, 2020). We have found that this epoch provides a good characterization of functional
220 responses and allows for fair comparison across brain regions. We used it here for those reasons
221 and because adherence to a single pre-planned epoch of interest reduces the likelihood of
222 inadvertent “p-hacking”.

223 *Shuffle analysis:* to understand the values of correlations that we would obtain by chance,

224 in the absence of a true effect, is it useful to repeat our analysis on dummy data. To do so, we
225 performed a specific control procedure developed by Elsayed et al. (2016). In this analysis, we
226 wished to preserve the features of each neuron without removing that information, so we
227 shuffled across time, meaning that we randomly sorted the spikes to occur at random times, but
228 maintained their linkage to the neuron that generated them. This procedure then gave dummy
229 data that we performed our correlation analyses on. More specifically, we shuffled the data from
230 across all three of the previously computed response matrices (safe, matched-medium, and
231 matched-large) and computed the alignment index between pairs of shuffled sets. Because the
232 data are randomized, any alignment index at or below this dummy threshold would be
233 considered at least semi-orthogonal. We then, following standard logic of bootstrap analyses,
234 repeated this process over 1000 randomizations. Then, to test for significance, we computed the
235 99% confidence interval across iterations; a value outside of this range can therefore be said to
236 be significant at $p < 0.005$ (given that our t-test is assumed to be one-tailed).

237 *Population overlap.* We adapted a previously published procedure (Azab and Hayden,
238 2017). For each neuron, we isolated both medium- and high-stakes risky offers with subjectively
239 equal values to the respective subject. For each trial, we computed the average firing rate across
240 a 500 ms window, starting from 100 ms after the appearance of the first offer (see above). We
241 then regressed the firing rates on the corresponding offer probability on each trial to produce a
242 beta weight representing that neurons degree of linear encoding of offer probability. For safe
243 offers, we were limited by the fact that there is only a single probability value (100%) for safe
244 offers. Therefore, we also computed the average firing rate from a 500 ms window that preceded
245 the start of the trial by 100 ms. For each trial, we computed the difference between the offer
246 response and the pre-trial response, to represent the degree of linear encoding for the safe offer

247 type. Finally, we performed a correlation between the absolute value of the safe offer encoding
248 and the absolute value of the risky offer beta weights. We plotted the two encoding weights
249 against each other and computed the line of best fit, given the correlation coefficient.

250 Decoding analysis. We built a pseudo-population of pseudo-trials. First, we isolated
251 neural responses to the safe offers and the equivalent risky offers. Then, we collapsed the firing
252 rates for each trial into an average for the 500 ms offer 1 epoch. We randomly selected 1000
253 samples for each neuron, under both risk conditions, resulting in two $n \times 1000$ matrices (one for
254 each offer type; safe vs. risky), where n represents the number of neurons recorded from each
255 region. We used an elastic-net logistic GLM to decode safe versus risky trials. Elastic-net offers
256 regularization over the decoding weights, including both L1 and L2-norms of both Lasso and
257 Ridge regression, respectively. Elastic-net GLMs can provide parameter estimates with reduced
258 bias and variance when there is collinearity in the predictors (Zou and Hastie, 2003).
259 Specifically, the Lasso regularization has the effect of shrinking weakly contributing predictors
260 to 0. The ridge regularization has the effect of reducing the size of all predictors. The elastic-net
261 represents a mixture of both these costs in determining GLM parameters. We fit the GLM across
262 1000 within-sample bootstraps. On each iteration, a randomly selected subset, constituting half
263 of each original safe vs. risky offer type pseudo-population, was used to fit the model and the
264 other half was used for cross-validation. The resulting cross validation provided a measure of
265 accuracy with which the elastic-net logistic regression model was able to decode safe vs. risky
266 offer types from the cross-validation set of firing rates.

267 We performed this entire procedure on 1000 bootstrapped iterations of randomly
268 constructed pseudo-populations. We averaged the accuracy across these out-of-sample
269 bootstrapped cross-validations and calculated the standard error across iterations. This procedure

270 was performed independently for safe vs. medium-risky and safe vs. high-risky, and was
271 performed independently for each subject. Finally, the entire procedure was performed on
272 shuffled data, to confirm that the expected prediction accuracy for a binary decoder of
273 randomized data would be 50%. To do this, each original pseudo-population was instead
274 comprised of a randomly selected trial, regardless of the offer type. Additionally, each pseudo-
275 trial was randomly assigned a label as one of the possible offer types. We then compared the
276 average decoding accuracy between safe-risky and the shuffled data, using a standard two-
277 sample t-test, with Bonferoni corrections.

278 *Subspace alignment.* We followed the procedure described in Elsayed et al., 2016.
279 Specifically, for each structure, we separated offers and neural responses by their risk profile
280 (safe and equivalent risky offer of medium- and high-stakes), as described previously. For each
281 neuron, we identified two factors to incorporate into a single condition: time and offer position.
282 Time included the same 500 ms period following the onset of offer 1 and preceding the onset of
283 offer 2. Time was segmented into 20 ms bins. For each 20 ms bin, we computed the mean firing
284 rate across trials on which the offer was positioned on either the left or right of the screen. Thus,
285 we constructed a condition (time X offer position) X neuron matrix of mean firing rates; that is, a
286 50 X n-neurons matrix. One such matrix was constructed for safe offers, one for medium, and
287 one for high-stakes risky offers of equivalent value. Firing rates in prefrontal areas of macaques
288 tend to be sparse. So, we smoothed these matrices using a gaussian filter, with a sigma equal to
289 one. We smoothed across columns in the matrix; in other words, we smoothed each neurons
290 individually. We then normalized the smoothed matrices, by computing the z-score within each
291 cell, to account for differences in encoding scaling between neurons.

292 Next, we performed a principal component analysis, using eigenvalue decomposition, on

293 the safe response matrix, providing a transformation matrix into which we projected both the
 294 safe response matrix and each of the risky response matrices. We computed the explained
 295 variance due to each of the principal components. We performed the same process of
 296 dimensionality reduction for each of the risky offers, projecting both the safe response and
 297 corresponding risky response data into the resulting principal component spaces (medium- and
 298 high-stakes risky response each into their own principal component spaces). To determine if the
 299 subspaces were aligned, we computed an alignment index:

300

$$A_{idx} = \frac{\text{Tr}(D_{risky}^T C_{safe} D_{risky})}{\sum_{i=1}^{sel-dim} \sigma_{safe}(i)}$$

301 where $\text{Tr}()$ is the sum along the diagonal entry, $sel-dim$ = the number of selected principal
 302 components (or ten, in the current study), D_{risky} are the set of top $sel-dim$ eigenvectors, C_{safe} is the
 303 covariance matrix for the safe responses, $\sigma_{safe}(i)$ is i -th singular value of C_{safe} . Essentially, the
 304 variance explained in safe responses by the top ten principal components of the risky responses
 305 is normalized against the sum of the variance explained by the top ten principal components of
 306 the safe responses. Note that we also performed this calculation using both the top 4 and top 7
 307 principal components. This control did not change the results of the significance tests, and so
 308 they are not reported.

309 In the Results, we described the difficulty associated with interpreting alignment indices
 310 given the structure of the data, and the corresponding need to use shuffling procedures to
 311 determine the effective upper and lower limits of measured alignment. To implement this
 312 procedure, we tested whether the safe and risky subspaces were more or less orthogonal, relative
 313 to randomly sampling within the space of this fixed correlation structure. We concatenated all
 314 safe and risky offer data into a single matrix. We then computed the covariance matrix and

315 performed an eigenvalue decomposition for the covariance matrix. We then randomly sampled
316 subspaces that were aligned to the fixed correlation structure of the response space, as follows:

$$v_{align} = \frac{ortho(U\sqrt{Sv})}{\|U\sqrt{Sv}\|_2}$$

317 where U and S are the eigenvectors and eigenvalue matrices, respectively, of the
318 computed covariance matrix. A matrix (v) was drawn from a normal distribution with a mean =
319 0.0 and variance = 1.0. *Orth()* computes the orthonormal basis of the projected matrix. This
320 process essentially maintains the neuronal covariance structure of the original covariance matrix
321 used for the eigenvalue decomposition. We repeated this process across 1000 iterations and
322 computed the alignment index for each, according to the above description. We then calculated
323 the average alignment index and the 99% confidence intervals across the 1000 iterations.

324

RESULTS**325 Behavior**

326 On each trial of the *risky choice task*, a macaque (*Macaca mulatta*) subject chose
327 between two offers that varied in magnitude and probability (**Figure 1A**, Strait et al., 2014, see
328 **Methods**). Safe offers (12.5% of offers) provided a small volume of juice (125 μL) with 100%
329 certainty. Risky offers provided either a medium (165 μL , 43.75% of offers) or large (240 μL ,
330 43.75% of offers) volume of juice with a defined probability (0-100%, 1% increments). The
331 offer types for the two offers were generated at random and independently on each trial. Our
332 dataset consists of responses from six subjects in 315 recording sessions and consists of 211,884
333 trials (average 672.6 trials per session). Target regions for neural recordings are illustrated in
334 **Figure 1B** and anatomical boundaries are provided in the **Methods**.

335 Subjects consistently performed at a high level, were modestly risk-seeking, and did not
336 differ from each other qualitatively (**Figure 1C**). Details of species-typical behavior in this task
337 are given elsewhere (in most detail in Farashahi et al., 2018 and 2019). Results of these analyses
338 are not repeated here, except to confirm, as we have previously shown, that subjects' behavior is
339 quite stable and consistent within subjects, both across and within sessions, and across subjects
340 and sessions (**Figure 1D-F**). Indeed, all subjects showed the same behavioral patterns we have
341 observed using this task in past studies.

342 All analyses in this paper make use of subjective values instead of expected values. To
343 identify the relative values of safe offers, we computed the risky-safe indifference point. We
344 calculated, separately for each subject and separately for medium and large stakes offers, the
345 likelihood that the subject would choose the safe offer as a function of the probability associated
346 with the risky offer. We then calculated indifference using a standard approach in which we fit

347 the resulting data with a sigmoid curve and calculated the point at which the best-fitting curve
348 crossed the indifference line (see **Methods**).

349 As in all of our past studies using this task, all subjects were risk-seeking. The cross-
350 subject average indifference point for medium-stakes risky offers corresponded to an offer
351 probability of 0.33 ± 0.05 (standard deviation). A risk-neutral subject would have had an
352 indifference point at 0.76. The average indifference point for high-stakes risky offers
353 corresponded to an offer probability of 0.11 ± 0.04 . A risk-neutral subject would have had an
354 indifference point of 0.52. As we have observed many times, preferences were highly consistent
355 across many contexts. For example, indifference points are similar for the first and second offers
356 (offer 1: medium: 0.34; high: 0.11; offer 2: medium: 0.22; high = 0.11), for offers made early
357 and late in the session (early: medium: 0.29; high: 0.09; late: medium: 0.31; high: 0.13), and
358 when risky offers appear on the left or right (left: medium: 0.27, high: 0.11; right: medium: 0.36,
359 high: 0.12). Data for an example subject are shown in **Figure 1D-F**. Data for individual subjects
360 are shown in **Figure 2**.

361

362 **Lack of pupil size effects**

363 Because our central research goal is to ascertain the influence of offer type on neural
364 activity, we wanted to ensure that our effects were not due to attention or arousal. We therefore
365 examined the relationship between two proxies for attention, looking at time and pupil size in
366 four of our subjects (subjects B, H, P, and S; these are the subjects for which we had pupil size
367 recorded). We found no detectable relationship. Specifically, during the offer 1 epoch, subjects
368 fixated the first offer for 198.5 ms if it was risky and for 199.1 ms if it was safe. These two are
369 not different ($p=0.85$ for the group, p -values were the same or higher for all four subjects).

370 During the offer 2 epoch, subjects fixated the second offer for 222.4 ms if it was risky and for
371 220.9 ms if it was safe (again, these are not different, $p=0.73$). During the choice epoch, we
372 calculated the total looking time and, again, found no differences (302.7 ms for chosen risky and
373 302.5 for chosen safe, $p = 0.93$). Nor were any of the individual subjects' fixation time
374 differences statistically significant for any of these measures ($p > 0.05$ in all cases). These results
375 do not appear to be dependent on our choice of epoch; we used longer analysis (1 sec) epochs for
376 each of these analyses and found qualitatively matched results.

377 During the offer 1 epoch, the pupil size did not differ for risky and safe options in any of
378 the four subjects ($p=0.10$ for subject S and $p > 0.5$ for the other three subjects). Specifically,
379 relative to the baseline value (defined as 0), risky offer-evoked pupil size was -0.92 z-score units
380 (for all subjects averaged); the corresponding value for safe offers was 0.93 z-score units.
381 Likewise, during the offer 2 epoch, relative to the baseline value (again, defined as 0), risky
382 offer-evoked pupil size was -0.81 z-score units (for all subjects averaged); the corresponding
383 value for safe offers was 0.80 z-score units. These are not different for any subject ($p = 0.21$ for
384 subject H, $p = 0.39$ for subject P, $p > 0.5$ for the other two subjects). Finally, during the choice
385 epoch, relative to the baseline value of 0, risky choice-evoked pupil size was -0.79 z-score units
386 (for all subjects averaged); the corresponding value for safe offers was also 0.79 z-score units.
387 These are not different for any subject ($p > 0.5$ in all cases).

388

389 **No special subpopulations for risky-and safe-preferring neurons**

390 We recorded responses of 843 neurons in five brain regions while our subjects performed
391 the risky choice task: vmPFC (area 14, 156 neurons), OFC (area 13, 157 neurons), pgACC (area
392 32, 255 neurons), PCC (area 29/31, 151 neurons), and VS (nucleus accumbens, 124 neurons).

393 We recorded in two subjects for all areas, although different subjects were used for the different
394 areas (see **Methods**). Detailed analyses of responses to risky offers were reported previously for
395 vmPFC and VS (Strait et al., 2014 and 2015). We have never previously examined neural
396 responses to safe offers.

397 For these and subsequent analyses, we used each individual's subjective indifference
398 point. We then defined a range of probabilities ($\pm 2.5\%$, total range of 5.0%) and treated all
399 offers within that range as being subjectively equivalent to the safe value. We subsequently
400 checked for robustness by repeating the following analyses using a larger range ($\pm 5\%$, total
401 range of 10%) but because we found no qualitative differences, we do not report those results.
402 We also repeated all analyses using within-session estimates of subjective value (rather than
403 within-subject, across-session). Again, because we found no qualitative differences, we do not
404 report those results.

405 Example neurons *vmPFC.70* and *pgACC.17* (**Figure 3A and B**) showed responses to safe
406 offers whose magnitude could not be predicted from responses to risky offers. Additional sample
407 cells, from each targeted area, including those that showed positive and negative monotonic
408 tuning for risky values and, again, unrelated responses to safe offers (**Figure 3C-E**).

409 If neural responses to risky and safe offers differ, one reason may be that they make use
410 of different subpopulations of neurons, ones specialized for risky and safe values. To determine
411 whether the brain makes use of separate subpopulations of neurons specialized for encoding the
412 values of safe and risky options, we adapted the approach we developed previously (Azab and
413 Hayden, 2017 and 2018). Specifically, we have shown that if there are two categorically distinct
414 populations of neurons, defined by distinct tuned variables (e.g., selectivity for offer 1 and
415 selectivity for offer 2), then correlating the unsigned regression weights (betas) will produce a

416 negative (<0) correlation. Conversely, if the two variables are found in a single population, then
417 correlating their unsigned betas will produce a positive correlation (**Figure 4A**). If the variables
418 are distributed at random in the sample of neurons, then correlation will produce a null result.
419 Here, to perform this calculation, we regressed probability for each risky option against firing;
420 for safe options, we computed the response evoked by the presentation of the safe offer by
421 comparing the safe-evoked response to the pre-trial baseline firing rate (see **Methods**). We found
422 that, in the OFC, the degree to which the neural populations encode safe offers was positively
423 correlated with the extent to which they encode matched risky offers (medium-stakes: $r = 0.49$, p
424 < 0.001 ; high-stakes: $r = 0.29$, $p < 0.001$; Pearson's correlation; **Figure 4B**). We found similar
425 results in all 5 other areas ($p < 0.001$, in all cases, Pearson's correlation; **Figure 4C-F**).

426 Note that, due to the non-normality of the data, it is possible that these correlations are
427 unduly driven by the few outlier datapoints. We therefore repeated these correlations using
428 Spearman's correlation. All reported correlations are still significant. Specifically, for the safe vs.
429 medium, we find values of OFC: $\rho=0.433$; $p=0.007$; vmPFC: $\rho=0.501$; $p<0.0001$; pgACC:
430 $\rho=0.318$; $p=0.008$; PCC: $\rho=0.498$; $p<0.001$; VS: $\rho=0.544$; $p<0.001$. For the safe vs. high, we
431 find values of OFC: $\rho=0.288$; $p=0.01$; vmPFC: $\rho=0.401$; $p<0.008$; pgACC: $\rho=0.438$; $p=0.001$;
432 PCC: $\rho=0.543$; $p<0.001$; VS: $\rho=0.511$; $p<0.001$. These results are consistent with more detailed
433 studies from our lab using these and similar datasets showing no intrinsic categories in responses
434 (Blanchard et al., 2018).

435

436

437

438 **Offer type (safe vs. risky) is decodable in all five regions**

439 Our principal question is how the brain might be able to disambiguate two offers that are
440 behaviorally indistinguishable but constituted by qualitatively distinct risk features. Neurons in
441 these regions may – and likely are - tuned for multiple variables (Rigotti et al., 2013; Raposo et
442 al., 2014; Fusi et al., 2016; Blanchard et al., 2018; Johnston et al., 2020). This fact is important
443 because if we assume a neuron is only tuned to a single variable, we may under-weight or ignore
444 its other quite real tunings, which would introduce biases - likely strong ones - into our analysis.
445 To deal with this problem, we use an elastic-net regression, which is a regularization procedure
446 that modifies a generalized linear model (GLM) with two regularization terms: the L1 and L2
447 norms of both Lasso and Ridge regression, respectively. The Lasso term increases sparsity by
448 penalizing weakly contributing neurons, while the Ridge term increases variance by reducing the
449 size of all predictors. The elastic net GLM can provide parameter estimates with reduced bias
450 and variance when there is collinearity in the predictors (Zhou and Hastie, 2003). We used the
451 model as a logistic decoder to predict whether an offer was safe or risky based on the fitted
452 model for population firing rates, using a bootstrap approach (see **Methods**).

453 We found that in the OFC, the elastic-net was able to decode the safe vs. medium-risky
454 offer type from population activity with 94.9% accuracy from subject P and 90.4% accuracy in
455 subject S (safe vs. high risky: 91.2% accuracy in subject P; 87.9% accuracy in subject S). We
456 found similar results in the other 5 brain areas (**Figure 5**).

457 To determine if decoding accuracy was significant, we repeated the process after
458 shuffling the pseudo-populations by randomly selecting from each offer type and each offer
459 epoch, and randomly assigning a safe or risky label to the response (see **Methods**). In OFC, the
460 decoding accuracy from shuffled data was 49.8% for both subject P and subject S, which
461 constitutes the expected random 50/50 guess for a binary decoder. Decoding safe vs. medium-

462 risky offer type was significantly higher than chance (Students t-test, subject P: $t = -861.4$, $p <$
463 0.001 ; subject S: $t = -1058.8$, $p < 0.001$; even after Bonferoni correction). We found similar
464 results in all 5 other areas, across subjects and risky offer magnitudes ($p < 0.001$ in all cases).

465

466

467

468

469

470 **Orthogonal response subspaces for value-matched risky and safe offers**

471 The underlying connectivity of neuron ensembles can effectively constrain activity to be
472 correlated, rendering a low-dimensional subspace (Gallego et al., 2017; Ebitz and Hayden,
473 2021). Successfully decoding the safe vs. risky identity of equally valued offers could then be
474 accomplished through their organization into distinct subspaces. We adapted previously used
475 approaches to characterize the uniqueness of safe and matched-risky subspaces (Elsayed et al.,
476 2016; Yoo and Hayden, 2020). Specifically, we projected risky offer responses into the safe offer
477 response subspace and computed the explained variance for each (**Figure 6**).

478 When we project data from the risky subspaces (medium and high), they have lower
479 values than when we project the safe values into the same subspace (**Figure 6A**), indicating the
480 presence of coding orthogonality. If the brain used collinear codes for risky and safe offers, then
481 the two sets of lines would be the same height. Indeed, a plot of the difference between the
482 average of the two sets of times (**Figure 6B**) shows that in all cases, these are greater than zero –
483 if the risky and safe subspaces were collinear, the differences here would be precisely equal to
484 zero. Note that these data represent the difference between the measured value and the shuffle

485 alignment index; this means that data with higher shuffle alignment indices will appear lower as
486 a result. Indeed, these results should not be taken to imply that the safe vs. high alignments are
487 consistently greater than the safe vs. medium alignments, even though it appears this way.
488 However, these plots are made relative to the shuffled values, which have consistent differences
489 related to the idiosyncratic properties of the statistics of their spike trains. Instead, the only
490 strong conclusion that can be drawn from these data is that both medium and high risky offers
491 are consistently encoded in subspaces semi-orthogonal to those of safe offers. Note also that the
492 projection weights are not necessarily reflective of the amount of orthogonality in a linear sense,
493 meaning two indices that are close to each other but significantly different are not necessarily
494 quantitatively similar.

495 We used these projections to quantify the extent to which subspaces were aligned (A_{idx} ;
496 see **Methods**). Purely orthogonal subspaces would have an alignment of 0; collinear ones would
497 have an alignment of 1.0; semi-orthogonal subspaces would have an intermediate value between
498 0.0 and 1.0. Note that semi-orthogonal subspaces would be sufficient to produce separate
499 representations and would satisfy our hypotheses. The alignment analysis indicated that in OFC,
500 for example, comparisons of safe and medium-stakes risky response subspaces had an alignment
501 index of $A_{idx} = 0.231$. Safe and high-stakes subspaces had an $A_{idx} = 0.266$.

502 These numbers do not account for the measurement limits imposed the by structure of the
503 data. We performed two control procedures, one to determine the practical upper measurable
504 alignment index, and the other to determine the practical lower measurable alignment index. To
505 determine the upper limit, we performed a control procedure to measure a threshold below which
506 indices would be considered at least semi-orthogonal (see **Methods**; Elsayed et al., 2016). Any
507 alignment index at or. To do this, we shuffled the data from across all three response matrices

508 (safe, matched-medium, and matched-large) and computed the alignment index between pairs of
509 shuffled sets (see **Methods**). We repeated this process over 1000 randomizations. Then, to test
510 for significance, we computed the 99% confidence interval across iterations; a value outside of
511 this range can therefore be said to be significant at $p < 0.005$ (given that our t-test is assumed to
512 be one-tailed). We found that the average shuffled alignment index in OFC was $A_{idx} = 0.276$.
513 Both the safe-medium and safe-high alignment indexes were below the 99% confidence interval
514 (0.271 - 0.281) and thus both significant at $p < 0.005$ (or $p < 0.01$ with a two-tailed t-test).

515 To determine the lower limit, we repeated the shuffle procedure, but randomized across
516 all axes. This analysis approach, in effect, identifies the alignment index we would observe if the
517 data were entirely orthogonal. In OFC, the alignment index across totally shuffled data was A_{idx}
518 $= 0.091 \pm 0.009$ (99% confidence interval). The safe-medium and safe-high alignment indexes
519 are both quite a bit greater than this noise floor (thus both significant at $p < 0.005$). In other
520 words, response subspaces for safe and equally valued risky offers in OFC are less orthogonal
521 than random data, indicating that they are partially, albeit not completely, aligned. We found
522 similar results in all structures (below the 99% confidence interval, in all cases; $p < 0.005$). We
523 can conclude, then, that these responses have an intermediate level of alignment.

524 Note that it is worth reading these numbers with a good degree of caution. The space of
525 orthogonality measures is inherently non-linear – for example, a value of 0.8 is not in any
526 meaningful sense twice as collinear as a value of 0.4. Indeed, a quantitative interpretation of
527 these numbers would require strong assumptions about the decoding process used by
528 downstream structures. Nonetheless, measures of statistical significance here are interpretable
529 and meaningful. Thus, when we say that the measured values are intermediate, we do not mean
530 that they lie roughly halfway between orthogonal and collinear; instead, we mean that they are

531 neither at one extreme nor the other. Moreover, it is worth pausing to emphasize what can and
532 cannot be inferred about orthogonality of subspaces from these numbers. On one hand the results
533 of the significance tests can be interpreted in a conventional way – a significant difference
534 between an upper and lower bound can be taken as evidence of semi-orthogonality. On the
535 other hand, the magnitude of the effect comes with several caveats that make interpreting it,
536 beyond the results of significance tests, difficult. This magnitude depends on, for example, the
537 signal-to-noise ratio in that population of neurons, their baseline firing rate, and the particular
538 nature of their distribution.

539 Finally, we asked how neural responses to medium- and high-stakes risky offers related to
540 each other. If the qualitative difference in codes for risky and safe offers reflect, even in part, the
541 qualitative differences between risk and safety, then two different risk offers should differ less
542 than the risky and safe offers. We therefore repeated the above analyses on medium- and high-
543 stakes risky offers. In OFC, we found that the medium- and high-stakes risky responses had an
544 alignment index of $A_{idx} = 0.275$. This value is significantly greater than the value of the medium-
545 risky and high-risky indices (0.231 and 0.266, respectively, see above). Moreover, this value is
546 within the confidence interval for shuffled data (0.271 - 0.281) and is not significantly different
547 from chance ($p \approx 0.29$). We found similar effects for the other brain areas. Specifically, in
548 vmPFC, the difference in A_{idx} between the both medium-stakes and high-stakes and risky-safe is
549 positive (0.051 and 0.018 for medium and high; $p < 0.01$ for both; for this and the following, we
550 used the bootstrap test described above). In pgACC, these numbers were 0.046 and 0.017
551 ($p < 0.01$ for both); in PCC, the numbers were 0.041 and 0.009 ($p < 0.01$ for both). In VS, the
552 numbers were 0.038 and 0.007 ($p < 0.01$ for medium and $p < 0.05$ for high).

553

554

555 **Orthogonal risky and safe response subspaces can be transformed to be aligned**

556 Above, we demonstrated that the population subspaces for risky and safe offers are semi-
557 orthogonal (see cartoon in **Figure 7A**). A primary question, then, is to what extent can these
558 subspaces be transformed into a common subspace that would aid in their comparison? Such a
559 common subspace could provide a mechanism for comparing offers of equal value (Yoo and
560 Hayden, 2020). We reasoned that if the response subspaces could be transformed such that their
561 hyperplanes were collinear, then the axis along which they become aligned is most likely the axis
562 that describes their relative values (**Figure 7B**).

563 To investigate whether our data here obey these principles, we performed a canonical
564 correlation analysis on the subspace loadings (i.e., the principal component projection weights,
565 see **Methods**; Gallego et al., 2018; Susillo et al., 2015) for the first three principal components
566 from the safe and risky offer response matrices. We also randomly shuffled the subspace
567 loadings 1000 times and performed the same canonical correlation analysis. In OFC, we found
568 that safe and equally valued medium-stakes offer response subspaces could be significantly
569 aligned to correlation of $r = 0.53$ ($p = 0.023$, bootstrap rank test). Safe and matched-high stakes
570 risky offer response subspaces reached a maximal correlation of $r = 0.61$ ($p < 0.01$). We found
571 similar results across the remaining five brain areas in our dataset ($p < 0.05$ in all cases, **Figure**
572 **8**). Specifically, we found that in all five brain areas, subspaces are linked. As above, we
573 repeated these canonical correlation analyses, but for the two risky offers (medium and high
574 stakes). They are also transformable. Specifically, in OFC, we find a correlation of $r=0.62$ ($p =$
575 0.011 , bootstrap rank test). In vmPFC, we find a correlation of $r=0.69$ ($p = 0.008$, bootstrap rank
576 test); in pgACC, we find a correlation of $r=0.71$ ($p = 0.007$, bootstrap rank test). In PCC, we find

577 a correlation of $r=0.64$ ($p = 0.013$, bootstrap rank test); finally, in VS, we find a correlation of
578 $r=0.72$ ($p = 0.007$, bootstrap rank test).

579 One limitation of this analysis is that it is based on neural responses occurring during the
580 first offer period only. That is, our analysis assumes implicitly that neural responses to the first
581 offer will be identical to those for the second offer, and that cross-trial subspaces will allow
582 comparison. This assumption can be tested. We therefore performed an analysis using neural
583 responses to the second offers. Specifically, we performed a canonical correlation analysis on
584 safe (offer 1) and equally valued middle-risk (offer 2), on safe (offer 1) and equally valued high-
585 stakes (offer 2), as well as the reverse. In all cases, we found that the principle of
586 transformability was preserved. Specifically, in OFC, we find a correlation of $r=0.65$ ($p = 0.009$,
587 bootstrap rank test). In vmPFC, we find a correlation of $r=0.68$ ($p = 0.01$, bootstrap rank test); in
588 pgACC, we find a correlation of $r=0.70$ ($p = 0.006$, bootstrap rank test). In PCC, we find a
589 correlation of $r=0.62$ ($p = 0.016$, bootstrap rank test); finally, in VS, we find a correlation of
590 $r=0.67$ ($p = 0.010$, bootstrap rank test).

591

DISCUSSION

592

593

594

595

596

597

598

599

600

601

602

603

604

605

606

607

608

609

610

611

612

613

We find that neural responses to risky and safe offers use distinct codes in five core reward areas. Our task uses a single safe value and a dense sampling of 99 different risky values (1-99%), meaning we can use behavior to precisely infer equivalent risky and safe values. Indeed, we find that risky and safe options are encoded in distinct (semi-orthogonal) ensemble subspaces. The risky and safe subspaces, while different, are systematically related so that there is a ready transformation between them; this transformability means that a downstream area can straightforwardly compare the values of the two offers despite the differences in the codes used to represent them. We conjecture that this distinct but transformable code can allow the brain to achieve two conflicting goals, that is, to maintain a separate representation of different offer types, but also to allow for their direct comparison on a single scale.

Typically, neuroeconomic models are concerned with developing theories to explain how we compare disparate option types (Plassmann et al., 2007; Padoa-Schioppa, 2011; Levy and Glimcher, 2012). While abstraction is important, the more general problem faced by the brain is a bit more complex - it must use a system that simultaneously maintains options' features (so that they can be used to influence behavior if needed) while *also* allowing for comparison of dissimilar features. For example, circumstances may change rapidly so that surety is more or less valuable; the decision-maker must be able to selectively change the relative value of risky options. The use of linked but semi-orthogonal subspaces for qualitatively different feature types allows this goal to be accomplished quickly. Our results, then, support the notion that these ostensible value regions do contribute to evaluation and comparison processes in choice, but do so in a way that maintains information about the qualities of the options, in addition to their values.

614 Our results have some bearing on longstanding discussions about the nature of abstract
615 value representation and their role in facilitating comparison (Padoa-Schioppa, 2011; Montague
616 and Berns; Hayden and Niv, 2021; Platt and Glimcher, 1999; Padoa-Schioppa and Assad, 2006;
617 Kennerley et al., 2009; Peters and Buchel, 2010; Kolling et al., 2016; So and Stuphorn, 2010;
618 Raghuraman and Padoa-Schioppa, 2014; McCoy and Platt, 2005). Traditional neuroeconomics
619 proceeds from the ‘neuron doctrine,’ which, in neuroeconomics, implies that comparison across
620 offer types requires the existence of neurons whose responses are the same for equally valued
621 offers of different types. By contrast, the ‘population doctrine’ allows for other forms of
622 equivalence, including translatable but distinct subspaces (Elsayed et al., 2017; for reviews, see
623 Saxena and Cunningham, 2019; Ebitz and Hayden, 2021). Equally valued risky and safe offers
624 are a well-known example of offer types with qualitatively different features that can be readily
625 compared, and for this reason, have long been the focus of studies about value representation
626 (Levy et al., 2010; Tobler and Weber, 2014). Thus, our results both raise the possibility of and
627 add empirical support for the idea that qualitatively different values can be compared in a way
628 that occurs at the population level and does not require value neurons. To speculate a bit,
629 subspace orthogonalization may offer more encoding flexibility than single neuron encodings
630 because they can be rapidly adjusted to context without requiring rewiring or retraining; they
631 may also be more robust to certain forms of error or degradation.

632 The concept of neural subspaces originates in motor cortex and can be used to explain
633 how motor preparation is kept separate from action, so as to allow for fast responding while
634 preventing precocious movement (Kaufman et al., 2014; Elsayed et al., 2016). We have argued
635 that the same principles apply to core value regions, and that subspace orthogonalization allows
636 for the separation of evaluation and comparison processes (Yoo and Hayden, 2020). The present

637 results suggest a new use for subspace partitioning - keeping track of two qualitatively different
638 types of rewards in a format so that the details of their properties (here, risky vs. safe) are
639 maintained, while also allowing for ready comparison in a downstream structure.

640 Orthogonalized subspaces are an emergent property of neurons with mixed selectivity.
641 Mixed selectivity, or the simultaneous tuning to multiple features, is an important property of
642 neurons in prefrontal and associated regions (Barak et al., 2013; Rigotti et al., 2013; Fusi et al.,
643 2016). Mixed selectivity is the foundation of many useful properties, including feature binding,
644 learning, abstraction, and generalization, and code morphing (Bernardi et al., 2020; Parthasarathy
645 et al., 2017). Notably, a good deal of the analysis of physiological data assumes, either explicitly
646 or tacitly, that neurons have singular tuning, not mixed selectivity. For example, approaches that
647 divide neurons into specific categorical types tend to classify those neurons according to which
648 feature drives them most strongly. By doing so, these approaches zero out those neurons' tuning
649 for other parameters, and make the population appear more categorical than it really is. That
650 work, however, risks misinterpretation of neural data, precisely because of the remarkable power
651 of mixed selectivity (Fusi et al., 2016).

652 One highlight of our study is that we were able to directly compare responses in five
653 brain regions in the same task. This fact means that we can ask questions about the unique
654 functional properties of each region, such as whether different regions contribute differentially to
655 evaluation, comparison, and action selection elements of choice. Here, we replicate our past
656 findings showing that functional properties are largely qualitatively the same (Strait et al., 2015
657 and 2016; Maisson et al., 2021; Fine and Hayden, 2022). This does not necessarily prove these
658 areas are strictly overlapping in their functions, nor does it prove that economic functions are
659 purely distributed. Indeed, there is a quite a good deal of evidence of functional specialization in

660 the brain (Wilson et al., 2010; Rushworth et al., 2011; Passingham and Wise, 2012). Instead, we
661 propose that, to some extent, economic functions are a general feature of prefrontal cortex (and
662 related structures like PCC and VS), and that, for core economic functions, these regions
663 function as a hierarchy, rather than as a series of modules (Hunt and Hayden, 2017; Yoo and
664 Hayden, 2018 and 2021; Fine and Hayden, 2021; Maïsson et al., 2021). Moreover, our results
665 here extend these previous ideas and suggest that these regions not only have similar coding
666 repertoires but use similar computational algorithms (here, judicious control of subspace) to
667 implement them.

668 Warren Weaver (1982) and Lola Lopes (1987) have both argued, in their work on
669 decision-making and risk, that it is a mistake to assume that two options with the same expected
670 value (including subjective expected value) are or should be treated the same way by decision-
671 makers. There are many natural situations in which a risky prospect is quite different from a safe
672 one, even if they are matched for subjective value. To give some examples, risky options elicit
673 emotions, and anxieties, promote learning, reward curiosity, generate error signals, require
674 maintenance of an eligibility trace, may be associated with exploratory rather than exploitative
675 states, require additional brain computations, and may in some cases differentially elicit
676 executive control (e.g. Loewenstein et al., 2001; McCoy and Platt, 2005; Slovic et al., 2007; Platt
677 and Huettel, 2008; Barasëghyan et al., 2013; So and Stuphorn, 2016). For these reasons, risky
678 and safe options ought to elicit at least somewhat distinct brain responses, even when they are
679 matched for subjective value. Our results show that in five brain areas, they elicit strikingly
680 different response patterns. This does not mean that an abstract single-neuron value code does
681 not exist in the brain; it may be found, for example, in a hierarchically later area, such as dorsal
682 anterior cingulate cortex (dACC, Cai and Padoa-Schioppa, 2012), or supplementary motor areas

683 (So and Stuphorn, 2010). Another possibility, however, is that the brain makes do without an
684 abstract value subspace, and instead achieves comparison through alternative means (Vlaev et
685 al., 2012; Yoo and Hayden, 2018; Hayden and Niv, 2021; Walasek and Brown, 2021).

686

REFERENCES

687

1. Azab, H., & Hayden, B. Y. (2017). Correlates of decisional dynamics in the dorsal

688

anterior cingulate cortex. *PLoS biology*, 15(11), e2003091.

689

2. Azab, H., & Hayden, B. Y. (2018). Correlates of economic decisions in the dorsal and

690

subgenual anterior cingulate cortices. *European Journal of Neuroscience*, 47(8), 979-993.

691

3. Azab, H., & Hayden, B. Y. (2020). Partial integration of the components of value in

692

anterior cingulate cortex. *Behavioral Neuroscience*, 134(4), 296.

693

4. Barak, O., Rigotti, M., & Fusi, S. (2013). The sparseness of mixed selectivity neurons

694

controls the generalization–discrimination trade-off. *Journal of Neuroscience*, 33(9),

695

3844-3856.

696

5. Binde, P. (2013). Why people gamble: A model with five motivational dimensions.

697

International Gambling Studies, 13(1), 81-97.

698

6. Blanchard, T. C., Piantadosi, S. T., & Hayden, B. Y. (2018). Robust mixture modeling

699

reveals category-free selectivity in reward region neuronal ensembles. *Journal of*

700

neurophysiology, 119(4), 1305-1318.

701

7. Blanchard, T. C., Wolfe, L. S., Vlaev, I., Winston, J. S., & Hayden, B. Y. (2014). Biases

702

in preferences for sequences of outcomes in monkeys. *Cognition*, 130(3), 289-299.

703

8. Cai, X., & Padoa-Schioppa, C. (2012). Neuronal encoding of subjective value in dorsal

704

and ventral anterior cingulate cortex. *Journal of Neuroscience*, 32(11), 3791-3808

705

9. Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of

706

action choices. *Annual Review of Neuroscience*, 33, 269-298.

707

10. Ebitz, R. B., & Hayden, B. Y. (2021). The population doctrine in cognitive neuroscience.

708

Neuron.

- 709 11. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P.
710 (2016). Reorganization between preparatory and movement population responses in
711 motor cortex. *Nature communications*, 7(1), 1-15.
- 712 12. Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D., & Soltani, A. (2019). Flexible
713 combination of reward information across primates. *Nature human behaviour*, 3(11),
714 1215-1224.
- 715 13. Farashahi, S., Azab, H., Hayden, B., & Soltani, A. (2018). On the flexibility of basic risk
716 attitudes in monkeys. *Journal of Neuroscience*, 38(18), 4383-4398.
- 717 14. Fine, J. M., & Hayden, B. Y. (2022). The whole prefrontal cortex is premotor cortex.
718 *Philosophical Transactions of the Royal Society B*, 377(1844), 20200524.
- 719 15. Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: high dimensionality for
720 higher cognition. *Current opinion in neurobiology*, 37, 66-74.
- 721 16. Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural manifolds for
722 the control of movement. *Neuron*, 94(5), 978-984.
- 723 17. Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., & Miller, L. E.
724 (2018). Cortical population activity within a preserved neural manifold underlies multiple
725 motor behaviors. *Nature communications*, 9(1), 1-13.
- 726 18. Glimcher, P. W., & Fehr, E. (Eds.). (2013). *Neuroeconomics: Decision making and the*
727 *brain*. Academic Press.
- 728 19. Grupe, D. W., & Nitschke, J. B. (2013). Uncertainty and anticipation in anxiety: an
729 integrated neurobiological and psychological perspective. *Nature Reviews Neuroscience*,
730 14(7), 488-501.
- 731 20. Hayden, B., Heilbronner, S., & Platt, M. (2010). Ambiguity aversion in rhesus macaques.

- 732 Frontiers in neuroscience, 4, 166.
- 733 21. Hayden, B. Y., & Niv, Y. (2021). The case against economic values in the orbitofrontal
734 cortex (or anywhere else in the brain). *Behavioral Neuroscience*, 135(2), 192.
- 735 22. Heilbronner, S. R. (2017). Modeling risky decision-making in nonhuman animals: shared
736 core features. *Current opinion in behavioral sciences*, 16, 23-29.
- 737 23. Heilbronner, S., & Hayden, B. (2013). Contextual factors explain risk-seeking
738 preferences in rhesus monkeys. *Frontiers in neuroscience*, 7, 7.
- 739 24. Heilbronner, S., Hayden, B. Y., & Platt, M. (2011). Decision salience signals in posterior
740 cingulate cortex. *Frontiers in Neuroscience*, 5(55).
- 741 25. Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American*
742 *economic review*, 92(5), 1644-1655.
- 743 26. Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework
744 for reward-based choice. *Nature Reviews Neuroscience*, 18(3), 172-182.
- 745 27. Johnston, W. J., Palmer, S. E., & Freedman, D. J. (2020). Nonlinear mixed selectivity
746 supports reliable neural computation. *PLoS computational biology*, 16(2), e1007544.
- 747 28. Kacelnik, A., & Bateson, M. (1996). Risky theories—the effects of variance on foraging
748 decisions. *American Zoologist*, 36(4), 402-434.
- 749 29. Kahneman, T. & Tversky, A. (1979). Prospect Theory: An analysis of decisions under
750 risk. *Econometric*, 47(2).
- 751 30. Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical
752 activity in the null space: permitting preparation without movement. *Nature neuroscience*,
753 17(3), 440-448.
- 754 31. Kennerley, S. W., Dahmubed, A. F., Lara, A. H., & Wallis, J. D. (2009). Neurons in the

- 755 frontal lobe encode the value of multiple decision variables. *Journal of cognitive*
756 *neuroscience*, 21(6), 1162-1178.
- 757 32. Kolling, N., Wittmann, M. K., Behrens, T. E., Boorman, E. D., Mars, R. B., &
758 Rushworth, M. F. (2016). Value, search, persistence and model updating in anterior
759 cingulate cortex. *Nature Neuroscience*, 19(10), 1280-1285.
- 760 33. Lau, B., & Glimcher, P. W. (2008). Value representations in the primate striatum during
761 matching behavior. *Neuron*, 58(3), 451-463.
- 762 34. Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision
763 making. *Annual review of psychology*, 66, 799-823.
- 764 35. Levy, D. J., & Glimcher, P. W. (2012). The root of all value: a neural common currency
765 for choice. *Current opinion in neurobiology*, 22(6), 1027-1038.
- 766 36. Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural
767 representation of subjective value under risk and ambiguity. *Journal of*
768 *neurophysiology*, 103(2), 1036-1047.
- 769 37. Libby, A., & Buschman, T. J. (2021). Rotational dynamics reduce interference between
770 sensory and memory representations. *Nature neuroscience*, 24(5), 715-726.
- 771 38. Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings.
772 *Psychological bulletin*, 127(2), 267.
- 773 39. Lopes, L. L. Between hope and fear: The psychology of risk. *Adv. Exp. Soc. Psychol.* 20,
774 255–295 (1987).
- 775 40. Maisson, D. J. N., Cash-Padgett, T. V., Wang, M. Z., Hayden, B. Y., Heilbronner, S. R.,
776 & Zimmermann, J. (2021). Choice-relevant information transformation along a
777 ventrodorsal axis in the medial prefrontal cortex. *Nature communications*, 12(1), 1-14.

- 778 41. McCoy, A. N., & Platt, M. L. (2005). Risk-sensitive neurons in macaque posterior
779 cingulate cortex. *Nature neuroscience*, 8(9), 1220-1227.
- 780 42. Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode
781 economic value. *Nature*, 441(7090), 223-226.
- 782 43. Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model.
783 *Annual review of neuroscience*, 34, 333-359.
- 784 44. Parthasarathy, A., Herikstad, R., Bong, J. H., Medina, F. S., Libedinsky, C., & Yen, S. C.
785 (2017). Mixed selectivity morphs population codes in prefrontal cortex. *Nature*
786 *neuroscience*, 20(12), 1770-1779.
- 787 45. Passingham, R. E., & Wise, S. P. (2012). *The neurobiology of the prefrontal cortex:*
788 *anatomy, evolution, and the origin of insight* (No. 50). Oxford University Press.
- 789 46. Paulus, M. P., & Angela, J. Y. (2012). Emotion and decision-making: affect-driven belief
790 systems in anxiety and depression. *Trends in cognitive sciences*, 16(9), 476-483.
- 791 47. Paxinos, G., Huang, X. F., & Toga, A. W. (2000). *The rhesus monkey brain in stereotaxic*
792 *coordinates*.
- 793 48. Plassmann, H., O'doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes
794 willingness to pay in everyday economic transactions. *Journal of neuroscience*, 27(37),
795 9984-9988.
- 796 49. Platt, M. L., & Huettel, S. A. (2008). Risky business: the neuroeconomics of decision
797 making under uncertainty. *Nature neuroscience*, 11(4), 398-403.
- 798 50. Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal
799 cortex. *Nature*, 400(6741), 233-238.
- 800 51. Peters, J., & Büchel, C. (2010). Neural representations of subjective reward

- 801 value. *Behavioural Brain Research*, 213(2), 135-141.
- 802 52. Raghuraman, A. P., & Padoa-Schioppa, C. (2014). Integration of multiple determinants in
803 the neuronal computation of economic values. *Journal of Neuroscience*, 34(35), 11583-
804 11603.
- 805 53. Rappoport, D., Kaufman, M. T., & Churchland, A. K. (2014). A category-free neural
806 population supports evolving demands during decision-making. *Nature neuroscience*,
807 17(12), 1784-1792.
- 808 54. Reyna, V. F., & Huettel, S. A. (2014). Reward, representation, and impulsivity: A
809 theoretical framework for the neuroscience of risky decision making.
- 810 55. Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., & Fusi, S.
811 (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*,
812 497(7451), 585-590.
- 813 56. Rushworth, M. F., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E.
814 (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron*, 70(6),
815 1054-1069.
- 816 57. Saxena, S., & Cunningham, J. P. (2019). Towards the neural population doctrine. *Current*
817 *opinion in neurobiology*, 55, 103-111.
- 818 58. Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2007). The affect heuristic.
819 *European journal of operational research*, 177(3), 1333-1352.
- 820 59. Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2004). Risk as analysis and
821 risk as feelings: Some thoughts about affect, reason, risk, and rationality. *Risk Analysis:*
822 *An International Journal*, 24(2), 311-322.
- 823 60. So, N. Y., & Stuphorn, V. (2010). Supplementary eye field encodes option and action

- 824 value for saccades with variable reward. *Journal of Neurophysiology*, 104(5), 2634-2653.
- 825 61. Strait, C. E., Blanchard, T. C., & Hayden, B. Y. (2014). Reward value comparison via
826 mutual inhibition in ventromedial prefrontal cortex. *Neuron*, 82(6), 1357-1366.
- 827 62. Strait, C. E., Sleezer, B. J., & Hayden, B. Y. (2015). Signatures of value comparison in
828 ventral striatum neurons. *PLoS biology*, 13(6), e1002173.
- 829 63. Strait, C. E., Sleezer, B. J., Blanchard, T. C., Azab, H., Castagno, M. D., & Hayden, B.
830 Y. (2016). Neuronal selectivity for spatial positions of offers and choices in five reward
831 regions. *Journal of neurophysiology*, 115(3), 1098-1111.
- 832 64. Sussillo, D., Churchland, M. M., Kaufman, M. T., & Shenoy, K. V. (2015). A neural
833 network that finds a naturalistic solution for the production of muscle activity. *Nature*
834 *neuroscience*, 18(7), 1025-1033.
- 835 65. Tang, C., Herikstad, R., Parthasarathy, A., Libedinsky, C., & Yen, S. C. (2020).
836 Minimally dependent activity subspaces for working memory and motor preparation in
837 the lateral prefrontal cortex. *Elife*, 9, e58154.
- 838 66. Tobler, P. N., & Weber, E. U. (2014). Valuation for risky and uncertain choices.
839 In *Neuroeconomics* (pp. 149-172). Academic Press.
- 840 67. Yoo, S. B. M., Hayden, B. Y., & Pearson, J. M. (2021). Continuous decisions.
841 *Philosophical Transactions of the Royal Society B*, 376(1819), 20190664.
- 842 68. Yoo, S. B. M., & Hayden, B. Y. (2020). The transition from evaluation to selection
843 involves neural subspace reorganization in core reward regions. *Neuron*, 105(4), 712-724.
- 844 69. Yoo, S. B. M., & Hayden, B. Y. (2018). Economic choice as an untangling of options
845 into actions. *Neuron*, 99(3), 434-447.
- 846 70. Walasek, L., & Brown, G. (2021). Incomparability and Incommensurability in Choice:

- 847 No Common Currency of Value?. psyarXiv.
- 848 71. Wang, M. Z., & Hayden, B. Y. (2021). Latent learning, cognitive maps, and curiosity.
- 849 Current Opinion in Behavioral Sciences, 38, 1-7.
- 850 72. Wang, M. Z., Hayden, B. Y., & Heilbronner, S. R. (2022). A structural and functional
- 851 subdivision in central orbitofrontal cortex. Nature communications, 13(1), 3623.
- 852 73. Weaver, W. Lady luck: The theory of probability. (Dover Press, 1982).
- 853 74. Wilson, C. R., Gaffan, D., Browning, P. G., & Baxter, M. G. (2010). Functional
- 854 localization within the prefrontal cortex: missing the forest for the trees?. Trends in
- 855 neurosciences, 33(12), 533-540.
- 856 75. Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. (2011). Does the brain calculate
- 857 value?. Trends in cognitive sciences, 15(11), 546-554.
- 858

859

860

Author contributions:

861

Conceptualization: DM, JF, SY, JZ, BH

862

Methodology: DM, BH, JF

863

Investigation: TC, MW, BS, JF

864

Visualization: DM

865

Supervision: JZ, BH, JF

866

Writing – original draft: DM, JZ, BH

867

Writing – review & editing: DM, JF, SY, JZ, BH

868

869

870

Data and materials availability:

871

All analyses were performed using stock Matlab functions and no custom code was

872

generated. Data used for all reported analyses are available on Dryad

873

[\(<https://doi.org/10.5061/dryad.ttdz08kx9>\).](https://doi.org/10.5061/dryad.ttdz08kx9)

874

875

876

877

Figure Captions

878

879 **Figure 1.** *Task, targeted structures, behavior, and calculation of equivalent risky and safe*
880 *values. (A)* Structure of our Risky Choice Task (Strait et al., 2014). Task consists of 400 ms offer
881 1 presentation, a 600 ms blank period, a 400 ms offer 2 presentation, another 600 ms blank
882 period, then a fixation spot and, on fixation, reappearance of both offers and choice (indicated by
883 saccade). For each offer, the magnitude of the associated reward (stakes) is indicated by the
884 bottom color (green, high or blue, medium) of the stimulus. The probability of being rewarded is
885 indicated by its size. **(B)** Anatomical positions of our brain regions of interest: OFC (purple),
886 vmPFC (purple-pink), pgACC (pink), PCC (gold), and VS (grey). **(C)** Likelihood of choosing
887 the first offer as a function of its value relative to the second (specifically, for signed value
888 difference). Sigmoid fits of raw binary data shown (see **Methods**). Gray lines: individual
889 subjects; black line: group average. In this and subsequent panels, a horizontal dashed line
890 indicates the indifference point (the point at which choices are 50/50). **(D)** Likelihood of
891 choosing a safe option as a function of the probability of the risky option for medium (blue) and
892 high (green) stakes offers. All data were analyzed on a subject-by-subject basis, so only data for
893 one example subject (subject B) are shown. Other subjects showed similar patterns. Vertical
894 black lines indicate the probability used as the SV-equivalence point for the subject (the arrow
895 points to the indifference point for medium- (blue) and high-stakes (green) risky offers). **(E)**
896 Same as D, except data are separated for left and right offers. Side of presentation does not affect
897 choice much. **(F)** Same as D, except data are separated by trials that were in the first (early) or
898 second (late) half of a session.

899

900 **Figure 2.** Indifference points for each of our five subjects individually. Y-axis shows the
901 probability (from 0% to 100%) associated with a risky offer that is equivalent in subjective
902 value, as assessed by preference indifference. Data are separated by whether the offer is medium
903 (left) or high (right) stakes). Subjects generally prefer risky offers, as indicated by the fact that
904 points tend to be on the lower y-values of the graph. Preferences are lower for high stakes than
905 for medium stakes, indicating that subjects took account of stakes. Subject indifferent to stakes
906 would show no difference between the medium and high stakes conditions. Error bars are not
907 shown because, due to the large number of trials, they are smaller than the bars.

908

909 **Figure 3. Responses of single neurons.** This figure shows the average responses of sample
910 neurons to safe and risky offers of differing values, as well as the average response similarity
911 rates. **(A)** Peristimulus time histogram from mean firing rates of sample neuron vmPFC.70. Each
912 line indicates the average response across offers of a given risk profile (grey: all safe offers; blue:
913 all medium-stakes risky offers; green: all large stakes risky offers). The grey shaded box
914 indicates the 1-second period from which the 500-ms epoch 1 analysis window was extracted,
915 where the onset of the first offer is time-locked to zero seconds. **(B)** This is a plot of data
916 collected from a sample neuron pgACC.17, which showed a response to safe offers that was
917 statistically different from the response to equivalent risky offers. Depicted are the average
918 responses to medium-stakes (left; blue) and high-stakes (right; green), separated by probability
919 ranges of 0.05. The black point indicates the average response of the given neuron to safe offers
920 (error bars denote the SEM across responses to safe offers). The diagonal black line indicates a
921 fitted regression line, showing positive monotonic tuning. **(C-E)** Same as (B) but demonstrating

922 sample cell responses to an assortment of medium- and high-stakes offers from across all target
923 areas.

924

925 **Figure 4.** *Overlapping populations encode both safe and risky offers.* **(A)** Conceptual schematic
926 indicating the analytical approach for identifying functional subpopulations. The right panels
927 propose the slope of a best fit line describing the encoding formats for each population type. **(B)**
928 For OFC (area 13), left panel: scatter plots represent the absolute degree of encoding of safe
929 offers plotted against that of the medium-stakes risky offers for each neuron (each point
930 corresponds to data from one neuron). The black line indicates the line of best fit for the
931 correlation between the absolute degree of encoding of safe and risky offers across neurons. Red
932 bars indicate the 99% confidence interval for the line of best fit. The inlaid numbers indicate the
933 Pearson correlation coefficient (r) and the significance (p). Right panel: same as left, but
934 comparing safe offer encoding and high-stakes risky offer encoding. **(C-F)** Same as B, but for
935 vmPFC (area 14), pgACC (area 32), PCC (area 29/31), and ventral striatum, respectively.

936

937 **Figure 5.** *Safe and equivalently valued risky offers are readily decoded.* Decodability for safe
938 and risky offers is high in all areas for both medium- and high-stakes gambles. Shuffled data
939 refers to decodability of randomly assigned safe/risky labels to neural responses that are
940 completely shuffled across trials and cells. Bars represent the mean decoding accuracy across out
941 of sample cross-validations Error bars indicate the standard error across cross-validations. The
942 horizontal grey line indicates the chance decoding accuracy level.

943

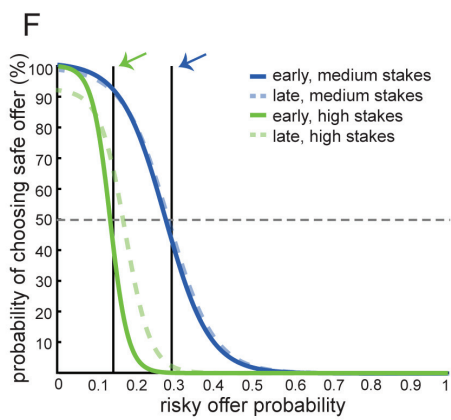
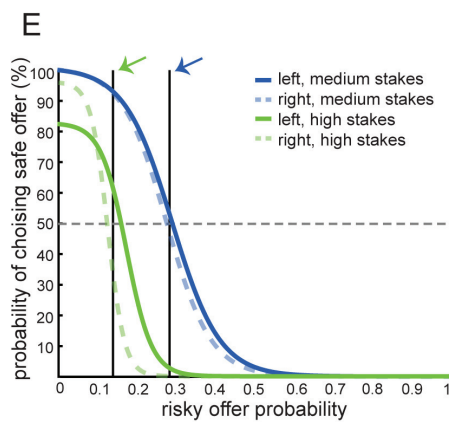

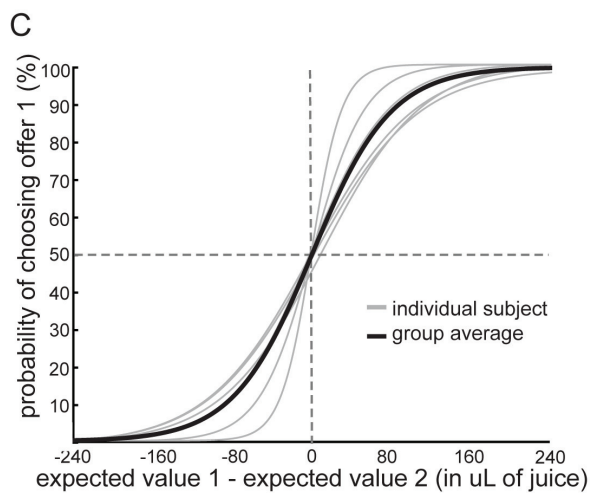

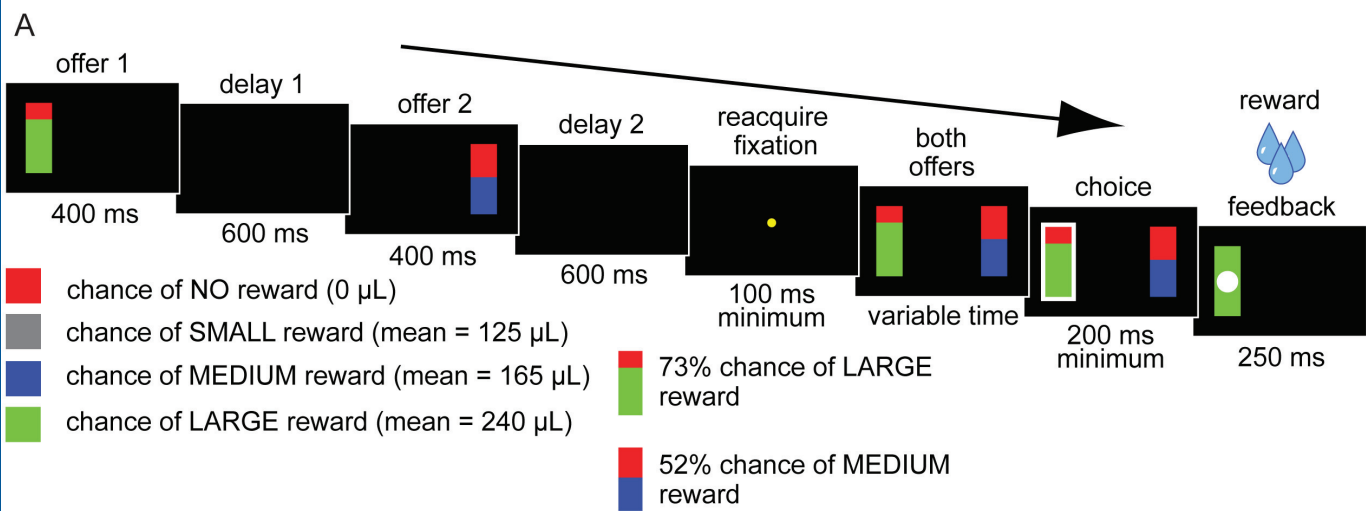
944 **Figure 6.** *Subspace Alignment between Safe and Equivalent Risky Offers.* **(A)** For each structure,

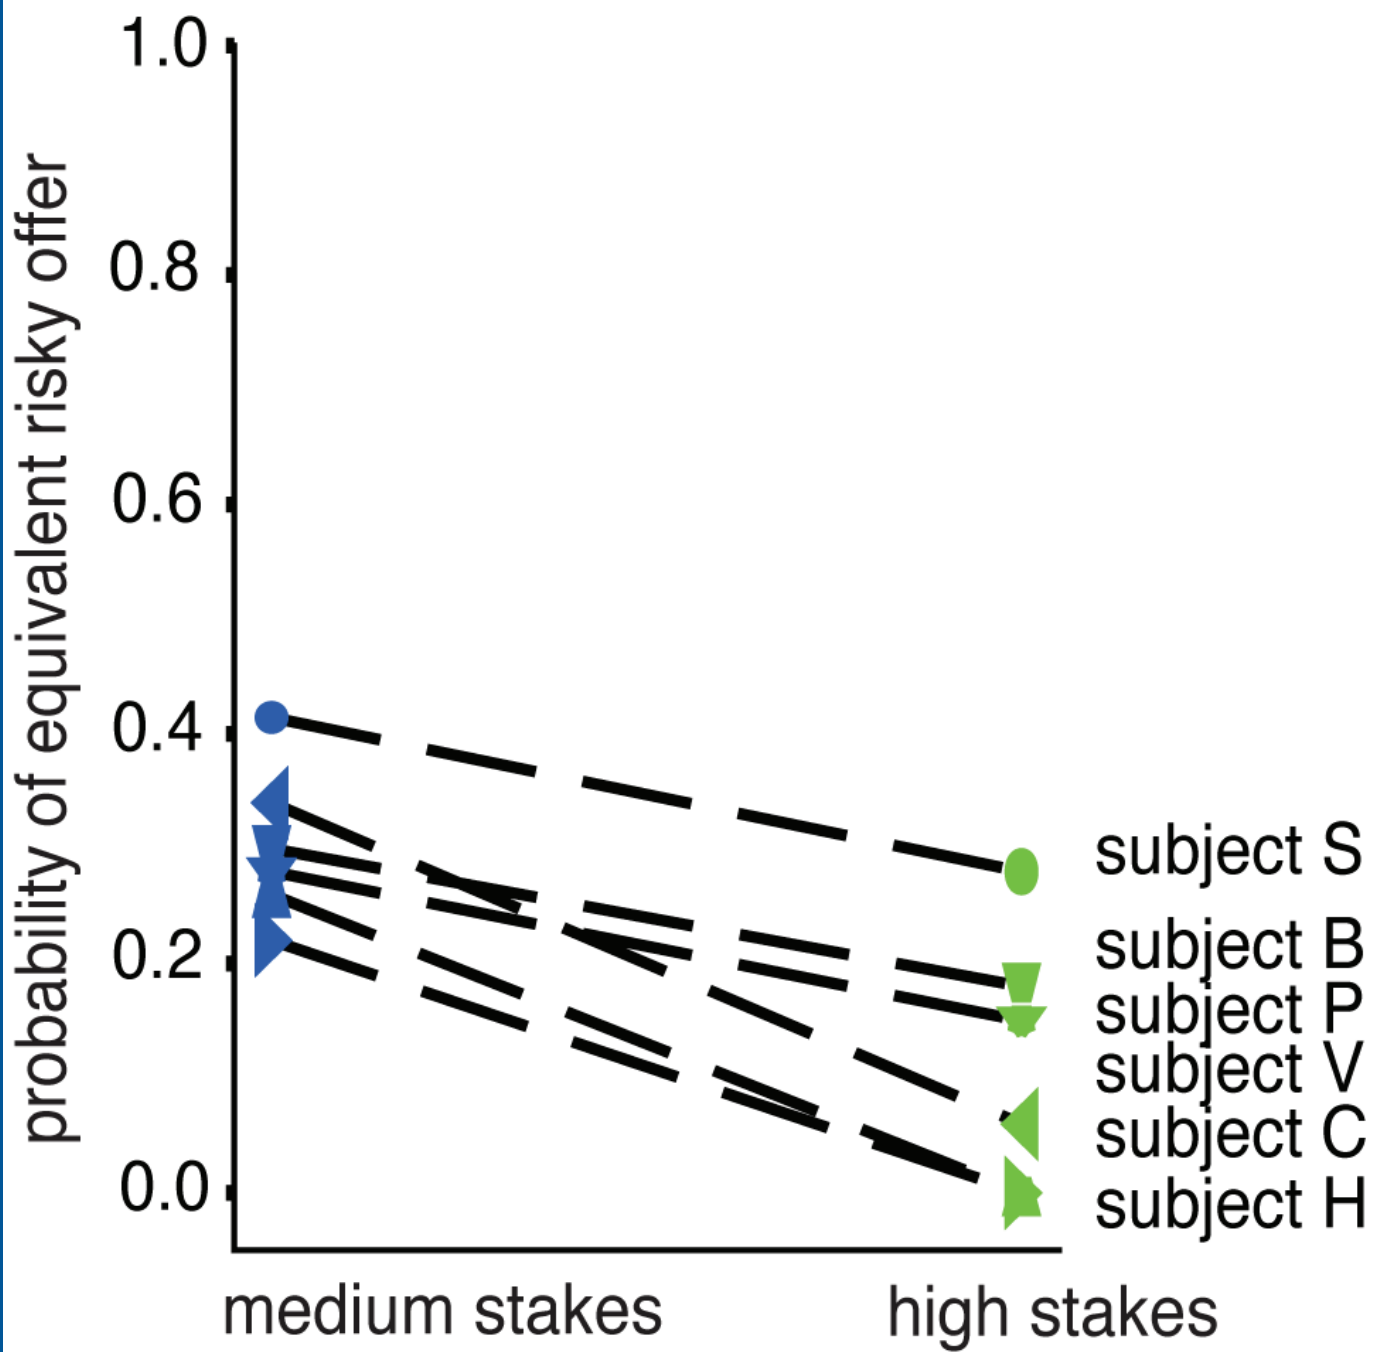
945 we performed a principal components analysis and projected the resulting dimensions for
946 medium (top row) and high (bottom row) stakes offers onto the corresponding ones for safe
947 offers. If subspaces for the risky and safe codes were the same (or collinear), then the projections
948 would match. Instead, the projections for the risky offers (lighter lines) are consistently lower
949 than the control values (the projections for the safe offers, darker lines). In other words, the fact
950 that the lighter (risky) lines are lower than the darker (safe) ones indicates that the two types of
951 offers are encoded in semi-orthogonal subspaces. **(B)** This finding is summarized in panel B,
952 which gives a summary of the difference between the average shuffle alignment index and either
953 the safe-medium (left bar) or the safe-high (right bar) alignment index. Error bars indicate the
954 standard error across computed differences. Note that these data represent the difference between
955 the measured value and the shuffle alignment index; this means that data with higher shuffle
956 alignment indices will appear lower as a result.

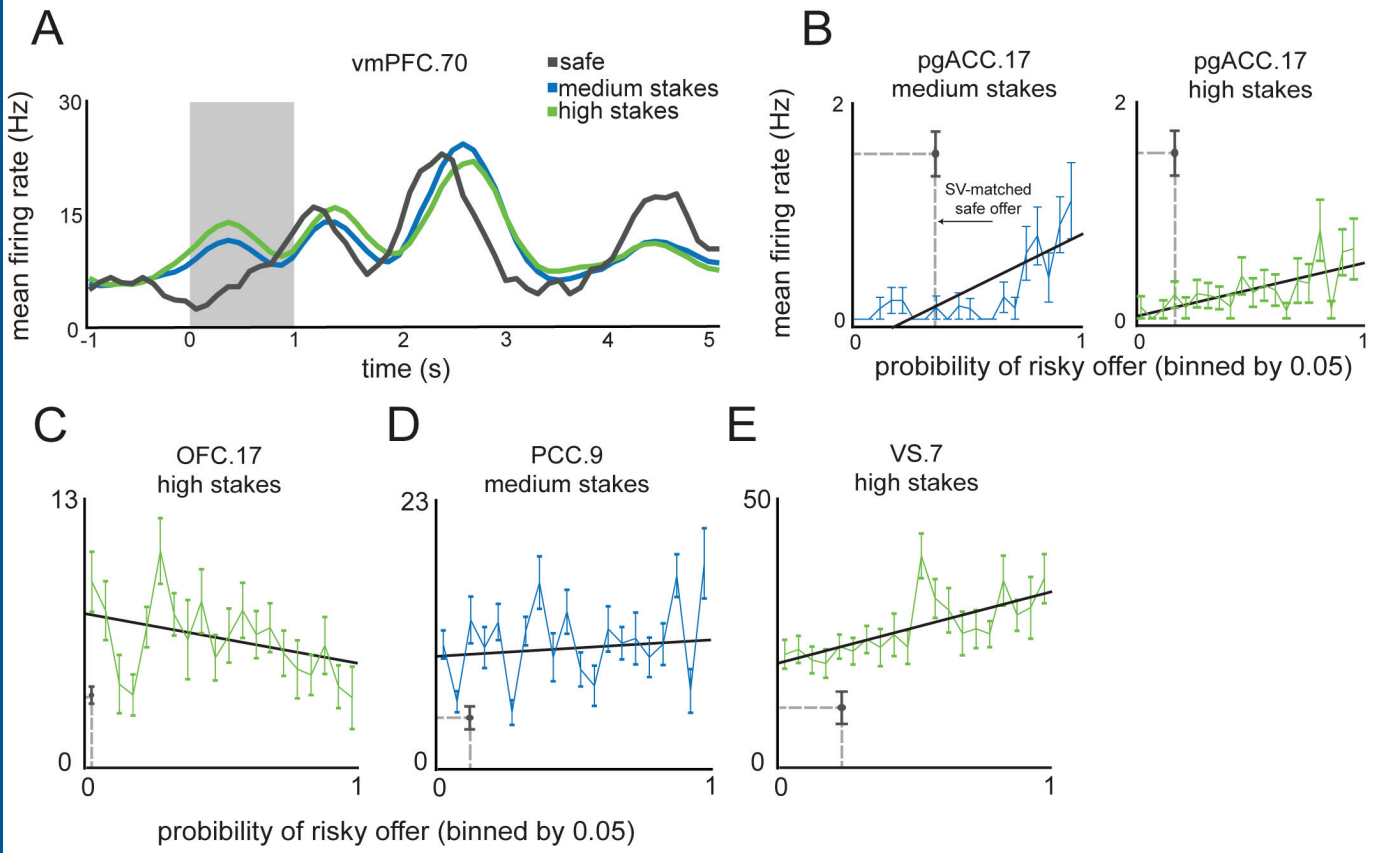
957

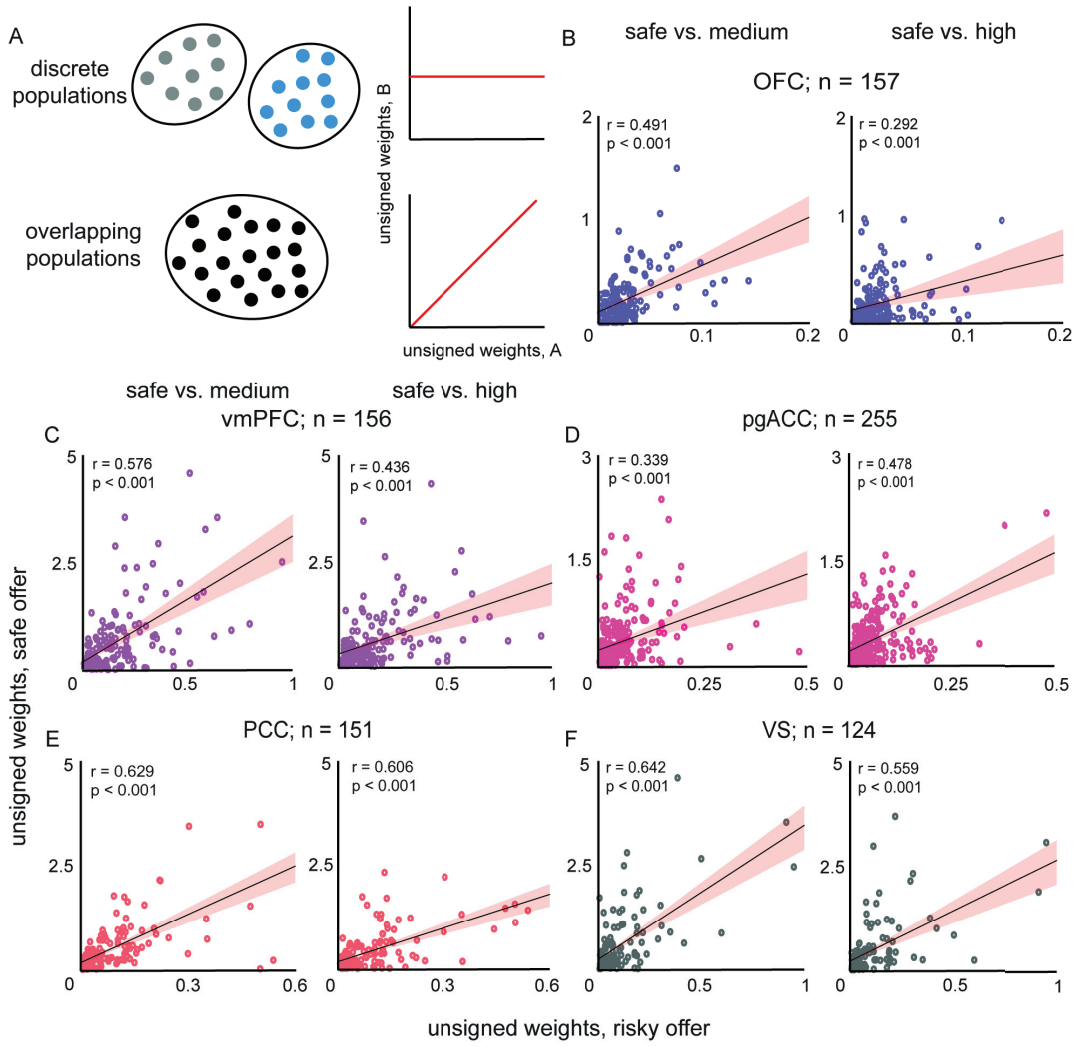
958 **Figure 7.** *Subspace rotational transformations and canonical correlations.* **(A)** A cartoon
959 demonstrating an example of orthogonal (left) and aligned (right) hyperplanes projected onto the
960 first three dimensions (cf. Yoo and Hayden, 2020). **(B)** A procedural schematic demonstrating
961 the process of linear algebraic rotational transformation, to align subspaces by maximizing the
962 Pearson's correlation coefficient via canonical correlation.


963 **Figure 8.** Scatter plots, for each structure, of projections of medium and safe offer responses (top
964 row) and high and safe offer responses (bottom row) projected into safe offer response
965 subspaces. Plotted are projections onto the first 2 principal component loadings for the original
966 (darker circles) and transformed (lighter pluses) projections. The inlaid coefficients are the
967 Pearson's correlation coefficients between projections on the first 2 principal component loading
968 for the original (r1) and transformed (r2) responses. The lines denote the best-fit line for the
969 corresponding correlation coefficients (r1: darker line; r2: lighter line).
970


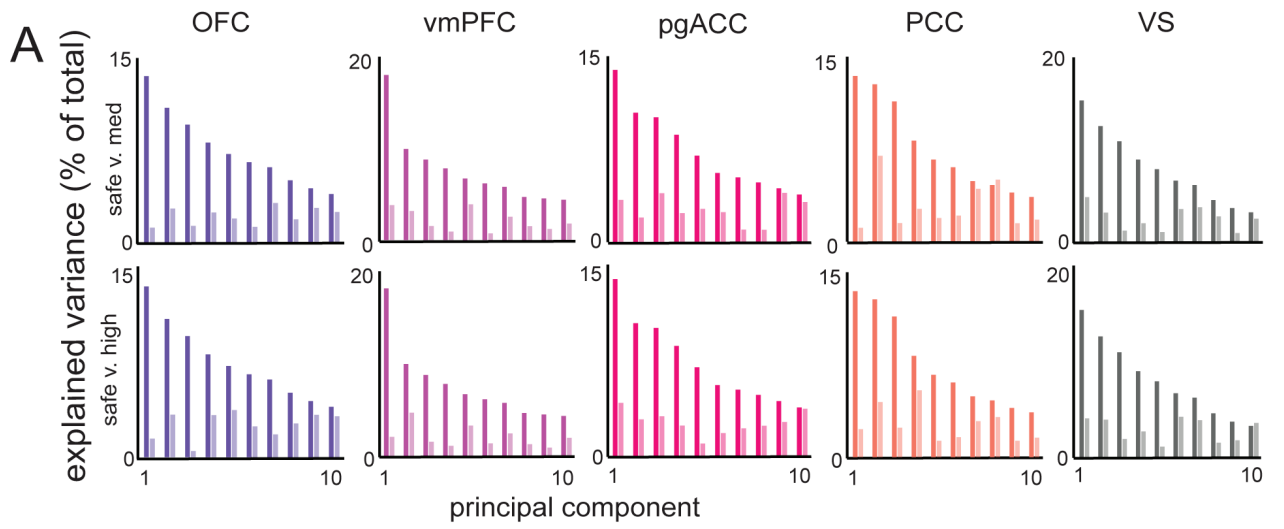




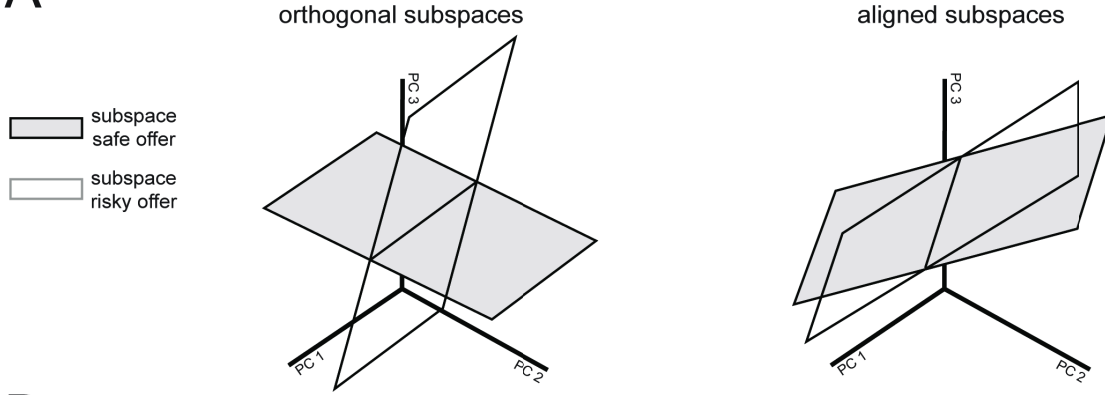








A



B

