Dorsal anterior cingulate cortex neurons only partially integrate determinants of value

Habiba Azab and Benjamin Hayden

Department of Neuroscience and Center for Magnetic Resonance Research
University of Minnesota, Minneapolis, Minnesota, United States of America

*Corresponding author:
    Habiba Azab
    Department of Neuroscience
    University of Minnesota
    Minneapolis, MN, 55455
    Email address: habiba.azab@gmail.com

CONFLICTS OF INTEREST
The authors declare no conflicts of interest, financial or otherwise.

ABSTRACT

37

38   Evaluation often involves integrating dissimilar determinants of value. A brain region can

39   therefore be placed either before or after a presumed evaluation stage by measuring whether

40   responses of its neurons depend simultaneously on multiple determinants of value. A brain

41   region could also, in principle, show partial integration, which would suggest that it occupies a

42   middle position between (pre-evaluative) non-integration and (post-evaluative) full integration.

43   Existing mathematical techniques cannot distinguish full from partial integration and therefore

44   risk misidentifying regional function. Here we use a new Bayesian regression-based approach to

45   analyze responses of neurons in dorsal anterior cingulate cortex (dACC) to risky offers. We find

46   that dACC neurons only partially integrate across outcome dimensions, indicating that dACC

47   cannot be assigned to a purely post-evaluative position. Neurons in dACC also show putative

48   signatures of value comparison, suggesting that comparison processes do not require evaluation

49   to be complete to proceed.

50

51 **INTRODUCTION**

52   To evaluate an option, we must consider all of the aspects of the option that influence

53 value and then combine them to generate an integrated value signal (Rangel et al., 2008; Kable &

54 Glimcher, 2009; Busemeyer et al., 2019; Tversky, 1969; Tversky & Kahneman, 1974). For

55 example, a risky prospect (a gamble) may involve two possible outcomes, both of which have

56 some influence on our likelihood of choosing the gamble itself. To evaluate the gamble, we must

57 decide how appealing each of those outcomes is and then somehow combine those component

58 evaluations. This could involve a process of multiplication and addition, as dictated by normative

59 approaches, or could be any of a number of satisficing heuristic approaches (Stewart et al., 2006;

60 Busemeyer and Townsend, 1998; Tversky and Kahneman, 1974; Farashahi et al., 2019). Clearly,

61 this integration process must be reified in the brain—we necessarily process dimensions

62 separately and our actions reflect only the integrated value, so combination has to occur

63 somewhere in the middle (Bowman et al., 2012; Roesch and Olson, 2003; Roesch and Olson,

64 2004). How it occurs remains an important problem in neuroeconomics.

65   One important goal in neuroeconomics is identifying where this integration occurs

66 (Rangel et al., 2009). Scientists therefore often seek to categorize brain regions as either

67 preceding or following the integration process. A standard approach is to probe the responses of

68 single neurons and test whether their responses are driven by multiple determinants of value

69 (Padoa-Schioppa and Assad, 2006; Kahnt et al., 2011; Hosokawa et al., 2013; Kennerley et al.

70 2009; Blanchard et al., 2015; So and Stuphorn, 2010; O'Neill & Schultz, 2018; Raghuraman &

71 Padoa-Schioppa, 2014). Simultaneous selectivity for multiple factors that determine value/utility

72 has therefore become a hallmark of cross-dimensional integration (Padoa-Schioppa, 2011; Hunt

73 et al., 2015; Fellows, 2006; Kim et al., 2008; O'Neill and Schultz, 2018; Strait et al., 2014).

74   However, this interpretation relies on the assumption that the process of evaluation is

75   anatomically discrete. Another possibility would be that it unfolds gradually across many

76   regions, much as complex form is gradually built up over multiple regions of the ventral visual

77   stream (Yoo and Hayden, 2018; Hunt and Hayden, 2017; Chen and Stuphorn, 2015; Cisek,

78   2012). That is, reward-sensitive regions may be aligned along a hierarchy, and the dimensions

79   may become gradually more united the further we travel along it. Critically, the standard

80   statistical approach of detecting dependence on multiple determinants of value does not

81   distinguish full (i.e. complete) from partial integration. We recently developed a method that

82   makes use of Bayesian regression and cross-validation that can distinguish partial from complete

83   correlations across dimensions in neural ensemble regressions, although we used it for a different

84   purpose (Azab and Hayden, 2017). Here we use that methods to address this question.

85   This approach also allows us to address a question about how evaluation relates to

86   comparison. Conventional approaches choice assume the completion of evaluation before the

87   start of comparison (Rangel et al., 2008, Padoa-Schioppa, 2011). For example, competitive

88   inhibition models, the dominant models of economic choice, involve separate pools of integrated

89   (i.e. post-evaluation) value neurons competing (Chau et al., 2014; Rustichini and Padoa-

90   Schioppa, 2015; Hunt et al., 2012; Hunt et al., 2015; Louie et al., 2011). However, there is no

91   special reason why comparison must wait until evaluation is complete (Yoo and Hayden, 2018).

92   In distributed choice systems, such as beehives, comparison processes can progress even in the

93   absence of full cross-dimensional integration. For example, individual bees that participate in a

94   swarm decision may sample only a subset of available determinant dimensions and nonetheless

95   contribution directly to comparison (Seeley, 2009; Seeley & Buhrman, 1999). It is not clear, in

96   the human brain, whether evaluation is complete before comparison can begin, or whether

4

97   comparison starts soon after evaluation does (Hunt et al., 2014; Cisek, 2012; Chen & Stuphorn,

98   2015; Strait et al., 2015; Azab & Hayden, 2018; Balasubramani et al., 2018). Cognitive models

99   of value-based decision-making have, in fact, often proposed that the evaluation and comparison

100   stages occur in parallel (Busemeyer et al., 2019; Noguchi & Stewart, 2018). Knowing whether

101   signatures of integration across dimensions is partial or complete can help shed light on this

102   question.

103   These questions are especially interesting to ask in the dorsal anterior cingulate cortex

104   (dACC). Among value-sensitive or economic regions, there is a broad consensus that this is an

105   anatomically or hierarchically late region (Heilbronner and Hayden, 2016; Rushworth et al.,

106   2011; Schall et al., 2002; Wunderlich et al., 2009). That is, it may serve as a repository for feed-

107   forward signals generated by other value-sensitive regions, such as orbitofrontal cortex and

108   ventromedial prefrontal cortex, and have a direct role in influencing motor and premotor

109   processing that other economic regions lack. As a consequence, it stands to reason that, among

110   value-sensitive regions it should be most likely to carry fully-integrated signals. Indeed, current

111   debates about the economic function of the dACC are generally split between (late) post-

112   comparison selection role and (middle) post-evaluative comparison role (Hare et al. 2011;

113   Wunderlich, Rangel & O'Doherty, 2009; Rangel & Hare, 2010; Azab & Hayden, 2017).

114   Here, we analyzed responses of neurons in dACC during a two-option token gambling

115   task (Azab & Hayden, 2017; Azab & Hayden, 2018). We find that while neurons encode both

116   large and small prospective outcomes of each gamble, their tuning functions for these outcomes

117   are similar but are demonstrably not identical. In other words, dACC responses reflect partial

118   integration across dimensions, not complete integration. We also observed putative signatures of

119   comparison in these neurons, suggest that comparison does not require evaluation to be complete

120   to proceed. These results suggest that dACC, even though it is normally placed late in the choice

121   hierarchy, does not follow the completion of integration. Moreover, we believe these results are

122   more consistent with distributed than with highly modular accounts of economic choice.

123
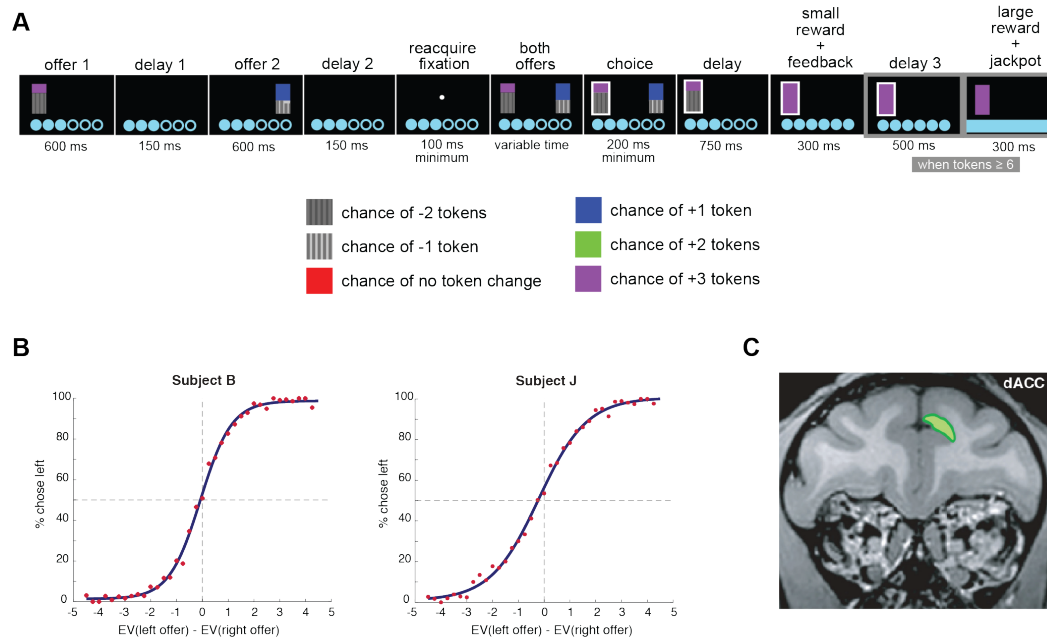
Figure 1: ***A:*** *Example trial from token-gambling task.* **B.** *Behavior for each subject, fit to a sigmoid function. Subjects choose the left option more often as its value increases, as would be expected given understanding of the task. EV: expected value of gamble.* ***C:*** *Regions of interest (for exact coordinates, see Methods). Figure adapted with permission from Azab & Hayden, 2018.*

RESULTS

Behavior is informed by integrated value

131
132
133

134   Details of the *token risky choice task* are provided in the **Methods** and are illustrated in

135 **Figure 1A** (see also Azab and Hayden, 2017 and 2018). On each trial, a macaque subject chose

136 between two gambles. Gambles offered the possible gain or loss of virtual tokens. Tokens

137 accumulated across trials; once subjects accrued six, they received a large water reward.

138 Gambles consisted of two possible outcomes (which we denote "large" and "small"—the small

139 one was most often negative or zero). Gambles appeared asynchronously; this staggered

140 presentation allowed us to examine neuronal responses to the first gamble independent of the

141 value of the second.

142
143 ***Gamble values influence subjects' choices***
144

	Subject B (n = 66 sessions)		Subject J (n = 74 sessions)	
Regressor	Mean regression coefficient	Wilcoxon sign-rank Z-statistic (p-value)	Mean regression coefficient	Wilcoxon sign-rank Z-statistic (p-value)
Left offer – win	1.40	7.06 (p = 1.64 x 10^{-12})	1.15	7.47 (p = 7.773 x 10^{-14})
Left offer – loss	0.486	6.97 (p = 3.26 x 10^{-12})	0.293	7.22 (p = 6.35 x 10^{-13})
Left offer – probability of win	5.12	7.06 (p = 1.64 x 10^{-12})	3.70	7.47 (p = 7.73 x 10^{-14})
Right offer – win	-1.33	-7.06 (p = 1.65 x 10^{-12})	-1.33	-7.47 (p = 1.67 x 10^{-12})
Right offer – loss	-0.412	-6.83 (p = 8.37 x 10^{-12})	-0.237	-7.06 (p = 1.67 x 10^{-12})

Right offer – probability of win	-4.85	-7.06 (p = 1.64 x 10⁻¹²)	-4.09	-7.47 (p = 7.73 x 10⁻¹⁴)

Table 1: *Subjects' choices were influenced by all values characterizing both gambles. Average regression coefficients from a multiple logistic regression model of choice (left=1, right=0) against the variables in the regressor column, for each behavioral session. The distribution of regression coefficients across sessions for every behavioral variable listed deviated significantly from zero. A 'loss' within a gamble was always less than or equal to the win outcome of that gamble, and may or may not be a negative value. The magnitude of the regression coefficient indicated the extent to which this variable influenced choice, while its sign indicated whether higher values of this variable (positive vs. negative) favored choice of the left vs. right gamble, respectively. Adapted with permission from Azab & Hayden, 2018.*

Details of the subjects' behavior are provided elsewhere (Azab and Hayden, 2017; Azab & Hayden, 2018; Farashahi et al., 2018). Briefly, both subjects preferred the option with the higher expected value (subject B: 80.3%; subject J: 75.1%; two-sided binomial test: both P < 0.0001). Behavior was significantly influenced by all three components (the two outcomes and the probability) of both gambles (see **Table 1**).

We first tested whether *choices* reflect the integrated values of the gambles or the components of each gamble independently. That is, whether the weighted product of outcome magnitudes and probabilities explained behavior over and above these gamble components individually. We fit choices to two multiple linear regression models, one including the three individual components of each gamble and another that included their interaction. For both subjects, inclusion of the integrated value predictor improved model fit (Wilcoxon sign-rank test between Akaike Information Criterion, AIC, subject B: z=7.06; subject J: z=7.47; p<0.0001 for both subjects). These results provide important evidence that our subjects behaviorally integrate the values. In other words, they indicate that the final common pathway (the oculomotor control neurons) encode integrated value, and therefore indicate that the integration must occur before or at that point (Roesch and Olson, 2003; Roesch and Olson, 2004).

9
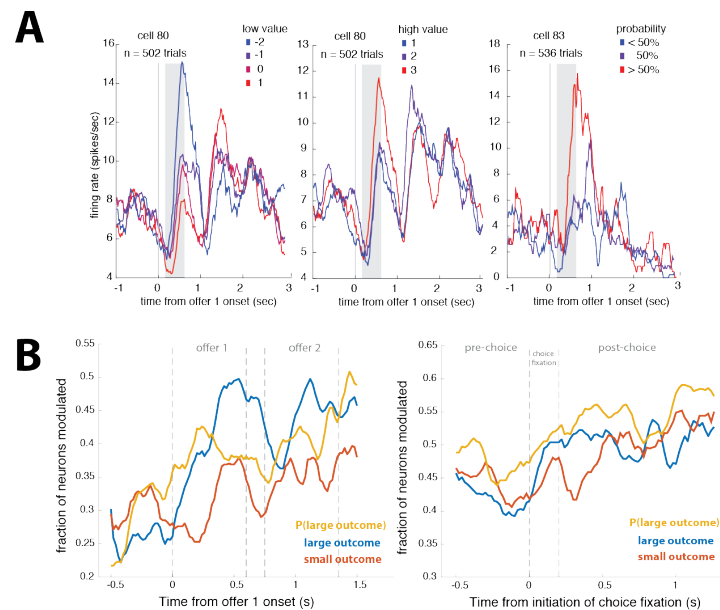
Figure 2: *Neural responses to gamble components.* **A:** *Responses of example neurons to the large outcome (left), probability of large outcome (middle), and small outcome (right) of the first gamble. Gray region: period of significant modulation.* **B:** *Proportion of neurons modulated by the components of the first gamble throughout the course of the trial. This fraction was computed by fitting a stepwise regression model to the normalized firing rate over a 500 ms sliding window against the components of the attended gamble (left) or the chosen gamble (right), along with other task-relevant variables (see Methods). Apparent effects before time zero reflect the width of the analysis window.*

Responses of dACC neurons carry value-relevant information

We recorded 129 neurons from dACC in two subjects (subject B: 55 neurons, subject J: 74 neurons; see **Figure 1B** for region of interest). Basics of modulation in response to task-relevant variables are presented in detail in previous manuscripts and are not repeated here (Azab & Hayden, 2017 and 2018). We use the same epoch we used in our previous papers, a 500 ms window starting 100 ms after events. Example cells are shown in **Figure 2**. For the first example cell, lower low values result in higher firing rates (t-test on firing rates, p<0.001). For the second example cell, higher high values result in higher firing rates (p<0.005). For the third example cell, firing depended on the probability of the large outcome (p<0.001).

Responses of 17.05% of neurons were affected by the size of the large outcome (n=22/129, one-sided binomial test: p<0.0001). Responses of 20.2% were affected by the

193   probability of that outcome (n=26/129, p<0.0001), and 8.53% were affected by the small

194   outcome (n=11/129, p=0.059) of the same gamble in the first epoch. Note that the lack of a

195   significant effect for the small outcome raises the possibility that dACC simply does not encode

196   this parameter. However, more sensitive tests do show an effect. Specifically, we performed a

197   stepwise linear regression model; a significant fraction of neurons showed dependence on the

198   small outcome in both epochs (epoch 1: 34.1%, n = 44/129, P < 0.0001; epoch 2: 55.0%, n =

199   71/129, two-sided binomial test: P < 0.0001).

200
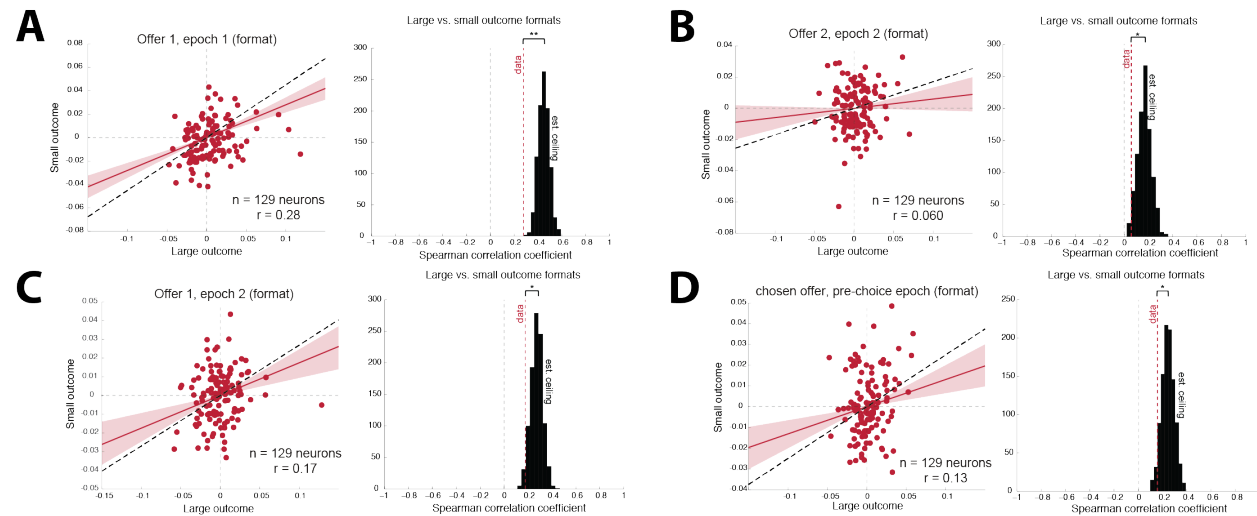
201
202   ***Figure 3****: Large and small outcomes are encoded in partially-separable formats in the*
203   *recorded population. **A**: Correlation between formats of large and small outcomes for the first*
204   *gamble in the first epoch. (left) Scatter-plot showing the correlation between regression*
205   *coefficients for large and small outcomes for each neuron. Red line indicates the positive*
206   *correlation between regression coefficients. Shaded red region indicates the 95% credible*
207   *interval on the data correlation. Black dashed line indicates the estimated ceiling correlation*
208   *between these variables in this epoch. (right) Results of permutation test. Black histogoram*
209   *indicates the distribution of the estimated ceiling correlation from permuted data. Red dashed*
210   *line indicates degree of correlation observed in the data. **B**: Correlation between variable*
211   *formats for the second gamble in the second epoch. **C.** Correlation between variable formats for*
212   *the first gamble in the second epoch. **D**: Correlation between variable formats for the chosen*
213   *gamble in the pre-choice epoch. *: p < 0.05, ** p < 0.01.*
214
215

216

217   **dACC neurons use partially integrated formats to encode gamble outcomes**

218   We have used the term *ensemble coding format* (or just *format*) to refer to the vector of

219   tuning coefficients for all recorded neurons (Strait et al., 2014; Wang and Hayden, 2017). We

220   then correlate formats across conditions to determine how they relate. Bayesian regression

221   allows us to draw samples from the likely format (rather than the single estimate conventional

222   regression provides, Gelman et al., 2013). All correlations below use Spearman (i.e. rank-order)

223   correlation, which is more robust to outliers than Pearson correlation, and is thus a more

224   conservative method. This analysis approach has a good deal of conceptual resemblance to

225   representational similarity analysis (RSA; Kriegeskorte et al., 2008); the main practical

226   difference is that it uses parameter tuning (specifically, regression weights) rather than raw

227   response (e.g. firing rates), making it more appropriate for the questions we want to ask here (see

228   **Discussion**).

229   We found that ensemble coding formats for large and small outcomes of each offer are

230   positively correlated (r=0.28, p<0.001, 1000 samples). That is, our population uses overlapping

231   (i.e. more similar than expected by chance) neural codes to signal the values of the large and

232   small outcomes; this is a signature of cross-dimensional integration (Strait et al., 2014;

233   Blanchard et al., 2015).

234   But does it reflect a *full* cross-dimensional integration? Mathematically speaking, why is

235   the measured correlation coefficient less than 1.0? It could reflect intermediate levels of

236   integration or it could reflect complete integration that appears intermediate due only to

237   measurement noise. (Recall that neurons are inherently noisy, so we can be certain that such

238   noise is a major factor). To differentiate these possibilities, we repeated the same analysis on a

239   shuffled dataset (where the values of the two variables of interest are shuffled *within* trial). This

240   tells us the maximum correlation observable between the encoding of two variables given the

241   noise in our data, which we call the *ceiling correlation*.

242   The ceiling correlation for large and small outcomes in the first epoch was r=0.45

243   (standard deviation = 0.047; **Figure 3A**). This result indicates that noise is high in our dataset (as

244   it is in most neural recordings) and that good deal of the explanation for the low correlation we

245   observed is simply noise. However, our observed correlation (r = 0.28, see above) is nonetheless

246   significantly lower than this value (p=0.001, bootstrap significance test, 1,000 permutations).

247   Thus, although large and small outcomes are encoded in similar formats, these formats are not

248   identical, as one would expect if dACC outputs reflected a completed value computation.

249   Mathematically speaking, the difference in formats necessarily reflects a qualitative (non-affine)

250   difference in tuning—a consistent gain shift or baseline shift would not be detected by our

251   method and would not generate a correlation significantly different from ceiling. In other words,

252   dACC neurons show evidence simultaneously of vectors corresponding integration and non-

253   integration, and these vectors are combined in our sample of neurons.

254   We found a similar pattern when considering the second gamble in the *second offer*

255   *epoch* (t=100 to 600 ms after the second offer appears). Specifically, although we did not find

256   evidence of integration across the two dimensions (r=0.059, p=0.060, compared to zero), we

257   found that whatever integration existed was less than the ceiling correlation value (r=0.17;

258   p=0.028; Figure 3B). Thus, it appears that dACC does not fully integrate either offer. (Note that,

259   because it is a null result, the lack of integration observed for the second offer does not prove no

260   integration occurred, although it is consistent with that possibility.)

261   We next asked whether full integration occurs when the information about the gamble is

262   transferred to working memory. It appears not. The coding of the first offer during the second

263   offer epoch (when it was no longer on screen) resulted in the same pattern: partial but not

264   complete integration (r=0.18, p<0.001 for comparison with zero. Ceiling correlation was r=0.27,

265   which was larger than the observed value, p=0.031; **Figure 3C**).

266

267   **dACC continues to show partial integration around the time of choice**

268   We next examined the 500 ms immediately preceding choice fixation (*pre-choice epoch*).

269   Whereas above we focused on offered gambles, here we focused on the *chosen gamble*, with the

270   reasoning that identifying an option for choice would potentially result in full integration. The

271   formats for integration of the two stakes is, indeed, positive (r=0.16, p<0.001). However, as like

272   in previous epochs, this correlation is less than ceiling (r=0.25, p=0.047; Figure 3D). These

273   results thus show that separation in large- and small-outcome formats persists at least until right

274   before choice is indicated.

275
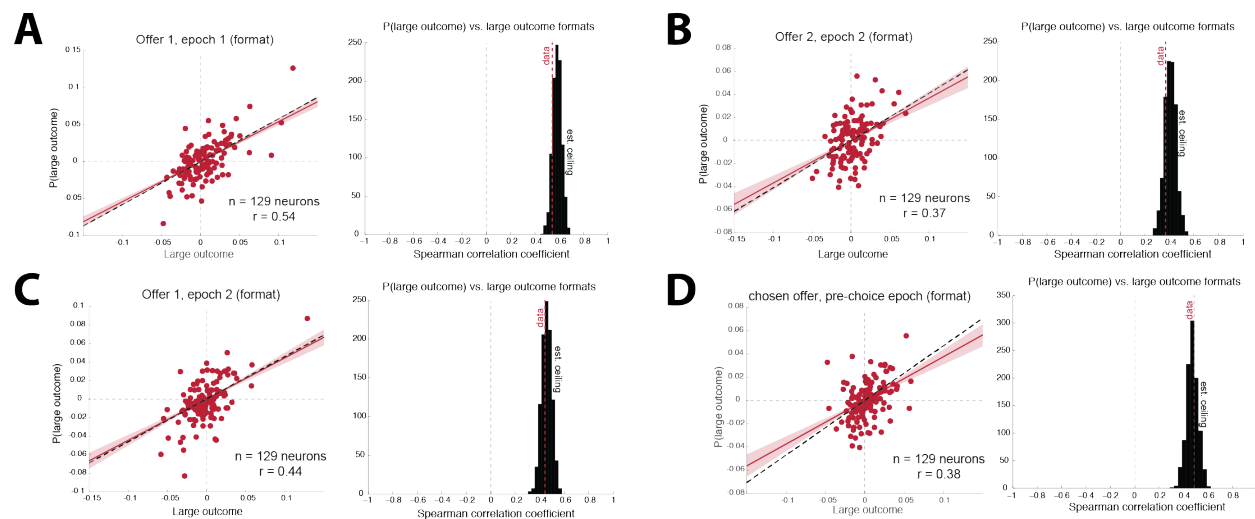
276
277   _**Figure 4:**_ *Large outcomes and their probabilities are encoded in integrated formats in*
278   *the recorded population. A: Correlation between formats of large outcomes and probabilities for*
279   *the first gamble in the first epoch. (left) Scatter-plot showing the correlation between regression*
280   *coefficients for large outcomes and probabilities for each neuron. Red line indicates the positive*
281   *correlation between regression coefficients. Shaded red region indicates the 95% credible*
282   *interval on the data correlation. Black dashed line indicates the estimated ceiling correlation*
283   *between these variables in this epoch. (right) Results of permutation test. Black histogoram*

284 *indicates the distribution of the estimated ceiling correlation from permuted data. Red dashed*
285 *line indicates degree of correlation observed in the data. **B:** Correlation between variable*
286 *formats for the second gamble in the second epoch. **C.** Correlation between variable formats for*
287 *the first gamble in the second epoch. **D:** Correlation between variable formats for the chosen*
288 *gamble in the pre-choice epoch.*
289
290

291 **Formats used to encode stakes and probabilities are consistent with full integration**

292 **for those dimensions**

293 For the first gamble in the first epoch, we find that the large outcome and its probability

294 were encoded in similar formats (r=0.54, p<0.001). This degree of correlation was not

295 significantly less than our estimated ceiling (ceiling r=0.58, p=0.135; Figure 4A). While we

296 cannot definitively conclude that there is full integration of these variables, the difference

297 between these results and the previous ones is striking. Indeed, for the first offer in the second

298 epoch (i.e. when it is in working memory), large outcomes and probabilities are also encoded in

299 similar formats (r=0.44, p<0.001) that are no different from ceiling (r=0.46, p=0.39; Figure 4B).

300 For the second offer, large outcomes and their probabilities are also encoded in similar formats

301 (r=0.37, p<0.001), and again no different from ceiling (r=0.41, p=0.22; Figure 4C). The chosen

302 gambles' large outcomes and probabilities are encoded in similar formats in the 500 ms

303 preceding choice (r=0.49, p<0.001), and again no different from ceiling (r=0.47, p=0.63; Figure

304 4D). Although the lack of a difference is not sufficient to prove identity, the matching patterns

305 for all four conditions, and their striking differences with the stakes analyses above, suggest that

306 large stakes and probability may be fully integrated in dACC.

307 These results are important because they indicate that the lack of integration observed for

308 large and small stakes is not a consequence of our task design, of the psychology of our subjects,

309 or of a limitation in our analysis approach. It appears that dACC is capable of fully integrating

15

310   some variables and maintaining some separation in others, and our methods are capable of

311   detecting this partial integration when it occurs.

312   This final point is further emphasized by our next analysis. In the first epoch, the small

313   outcomes and probabilities were encoded in similar formats ($r=0.11$, $p=0.002$), which are less

314   correlated than ceiling ($r=0.30$, $p=0.001$). Results are not conclusive for the same gamble in the

315   second epoch: there is no correlation between the formats for small outcomes and large-outcome

316   probabilities ($r=0.050$, $p=0.11$), but this correlation is no different from our estimated ceiling,

317   which is very low ($r=0.13$, $p=0.057$). The small outcome and large-outcome probability of the

318   second gamble are encoded in weakly-correlated formats in the second epoch ($r=0.063$,

319   $p=0.049$), and this correlation is no less than our estimated ceiling (also low, ceiling $r=0.13$,

320   $p=0.12$). The low stakes and probability components of the chosen gamble were encoded in

321   positively-correlated formats in the pre-choice epoch ($r=0.17$, $p<0.001$); no less than ceiling

322   ($r=0.15$, $p=0.65$). The lack of full integration between small outcomes and probabilities in the

323   first epoch (which is highly significant and survives Bonferroni correction) highlights the

324   qualitatively different way in which large and small offer outcomes are treated.

325

DISCUSSION

326

327   We examined the neural encoding of risky offers in dACC. We made use of a recently-

328   developed statistical method that can disambiguate variable encodings that are partially

329   correlated from ones that are fully correlated but different due solely to noise (Azab and Hayden,

330   2017). We find, first, that while dACC neurons encode both large and small prospective

331   outcomes of individual gambles on the screen and in working memory, the population tuning

332   formats for the two outcomes are only partially correlated—neither orthogonal nor collinear.

333   The "mid-evaluation" signal we observe is a type that is predicted by hierarchical

334   (sometimes called serial) models of value, which involve multiple small steps leading from a

335   space of options to one of actions rather than a single value estimation locus (Hunt & Hayden,

336   2017; Yoo & Hayden, 2018). In such models, reward and value processing occur along an

337   anatomical hierarchy and computation within that hierarchy implements a gradual transformation

338   from stimulus information to motor commands, rather than a series of categorical steps (Hunt et

339   al., 2012; Hunt & Hayden, 2017; Cisek, 2012; Chen & Stuphorn, 2015; Cisek & Kalaska, 2010).

340   More specifically, we have proposed that the relationship between different hierarchical levels

341   may be one of untangling—of rotating representations into a format more usable for effector

342   systems, while retaining information (Yoo and Hayden, 2018; DiCarlo et al.. 2012). In the

343   domain of economic choice the tangled format involves the two non-integrated dimensions and

344   the untangled format is an integrated value representation. The untangling theory, which is based

345   on how the visual stream works, proposes that each anatomical stage untangles information from

346   the earlier stage and allows for intermediate amounts of untangling.

347   The first major effect we observe suggests that dACC follows the initiation of value

348   integration but precedes its completion. Together these results suggest that dACC cannot be

17

349   assigned either a purely pre- or post-evaluation stage in choice. Moreover, in conjunction with

350   other results showing evidence for value comparison in dACC (e.g. Azab & Hayden, 2017; Azab

351   & Hayden, 2018; Klein-Flugge & Bestmann, 2012; Hare et al. 2011; Wunderlich et al., 2009;

352   Rangel & Hare, 2010), these findings indicate that value integration does not need to be

353   complete before value comparison begins (Hunt et al., 2014).

354   These results complement our recent findings using this dataset (Azab & Hayden, 2017).

355   In that study, we concluded that dACC occupies a middle role in the *comparison* of values. Our

356   complementary results here suggest that dACC occupies a middle role in the *computation* of

357   values. Juxtaposed, the two results have a larger implication: that the completion of integration is

358   not essential to the initiation of comparison. That is, it may be impossible, even in theory, to

359   draw a line, either in anatomical space or in time, between evaluation and comparison stages of

360   choice (Yoo & Hayden, 2018).

361   The identification of abstract amodal value signals is a desideratum of neuroeconomics

362   (Shizgal, 1997; Padoa-Schioppa, 2011; O'Doherty, 2014; Montague & Berns, 2002; Levy &

363   Glimcher, 2012). Our results suggest that, if these exist, they are most likely to occur only

364   downstream of dACC. We envision two major possibilities about this site of value

365   representation. First, value representations may only be complete in the premotor or motor

366   system (see Cisek, & Kalaska 2010, Cisek, 2007, and Thura & Cisek, 2016 for related

367   arguments). Second, that may be no single location at which value is represented abstractly and

368   completely (Chen & Stuphorn, 2015; Cisek, 2012; Yoo & Hayden, 2018; Hunt et al., 2015;

369   Balasubramani et al., 2018).

370   Ours is not the first study to consider representations of multi-attribute offers (O'Neill &

371   Schultz, 2018; Raghuraman & Padoa-Schioppa, 2014; Kennerley et al., 2009; Strait et al., 2014;

372   Blanchard et al., 2015; Padoa-Schioppa & Assad, 2006; Hunt et al., 2015). A broad finding from

373   these studies is that neurons in several regions show some overlap in the way they encode the

374   multiple dimensional determinants of value. The major difference is that we show that

375   integration of attributes into value representations is not an all-or-none phenomenon. The

376   distinction between attributes that are not integrated vs. somewhat integrated may not be

377   meaningful for readout, since this information is extractable in either case—future work is

378   needed to know what information is read out of dACC responses.

379   We have proposed that what we call a value may be better thought of as evidence in favor

380   of the proposition that the option will be chosen (Azab & Hayden, 2017; see also Hayden &

381   Moreno-Bote, 2018). From this perspective, there is no reason to expect that valuation must be

382   complete before comparison can begin. Each dimension can provide its own evidence in favor or

383   against the choice of an option. That evidence can be summed in one specific place or it can be

384   fed into a slowly evolving action plan without ever being integrated in one place. We and others

385   have previously argued that the analogy to bee swarm decision-making is helpful for

386   understanding choice (Hunt & Hayden, 2017; Eisenreich et al., 2017; Couzin, 2009; Passino et

387   al., 2008; Seeley et al., 2012; Pirrone et al., 2018). Individual bees evaluate hive sites in an

388   extremely noisy fashion (Seeley, 2009; Seeley & Buhrman, 1999). Indeed, an individual scout

389   may only assess one or two of the several dimensions along which hives vary. Even if no

390   individual bee performs a full cross-dimensional integration, the hive as a whole can still make a

391   good decision—without ever computing anything like value—through a quorum sensing

392   procedure. And the hive as a whole takes several days to select a hive site. In the middle of this

393   time, the hive may be said to have made a partially-complete evaluation process.

394   Similarly, some *cognitive* (rather than neural) models of value-based decision-making do

395   not make this assumption of discrete, serial stages; ordinal within-attribute comparisons

396   oftentimes contribute to the construction of each option's value. In this framework, comparison

397   would, in some sense, *precede* evaluation, and all value signals are entirely relative (Noguchi &

398   Stewart, 2018; Busemeyer et al., 2019). These models—also ones of value-based decision-

399   making, albeit at a different level of granularity—clearly demonstrate the importance of

400   integrating different levels of analysis when exploring the inner workings of a cognitive process,

401   and the many ways that a system (biological or otherwise) could potentially instantiate it.

402

REFERENCES

Azab, H., & Hayden, B. Y. (2017). Correlates of decisional dynamics in the dorsal anterior cingulate cortex. *PLoS biology*, *15*(11), e2003091.

Azab, H., & Hayden, B. Y. (2018). Correlates of economic decisions in the dorsal and subgenual anterior cingulate cortices. *European Journal of Neuroscience*, *47*(8), 979-993.

Balasubramani, P.P., Moreno-Bote, R., and Hayden, B.Y. (2018). Using a simple neural network to delineate some principles of distributed economic choice. Front. Comput. Neurosci. 12, 22.

Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, *85*(3), 602-614.

Blanchard, T. C., Piantadosi, S., & Hayden, B. Y. (2017). Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *Journal of neurophysiology*.

Blanchard, T. C., Strait, C. E., & Hayden, B. Y. (2015). Ramping ensemble activity in dorsal anterior cingulate neurons during persistent commitment to a decision. *Journal of neurophysiology*, *114*(4), 2439-2449.

Bowman, N. E., Kording, K. P., & Gottfried, J. A. (2012). Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. Neuron, 75(5), 916-927.

Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial vision*, *10*, 433-436.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. Psychological review, 100(3), 432.

Busemeyer, J. R., Gluth, S., Rieskamp, J., and Turner, B. M. (2019). Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-based Decisions. Trends in Cognitive Sciences 1873

Calhoun, A. J., & Hayden, B. Y. (2015). The foraging brain. Current Opinion in Behavioral Sciences, 5, 24-31.

Chau, B. K., Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. Nature neuroscience, 17(3), 463.

Chen, X., & Stuphorn, V. (2015). Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions. *Elife*, *4*, e09418.

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *362*(1485), 1585-1599.

Cisek, P. (2012). Making decisions through a distributed consensus. *Current opinion in neurobiology*, *22*(6), 927-936.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience*, *33*, 269-298.

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613-617.

Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in cognitive sciences*, *13*(1), 36-43.

DiCarlo, J.J., Zoccolan, D., and Rust, N.C. (2012). How does the brain solve visual object recognition? Neuron 73, 415–434.

Eisenreich, B. R., Akaishi, R., & Hayden, B. Y. (2017). Control without controllers: toward a distributed neuroscience of executive control. *Journal of cognitive neuroscience*, *29*(10), 1684-1698.

Farashahi, S., Azab, H., Hayden, B., & Soltani, A. (2018). On the flexibility of basic risk attitudes in monkeys. *Journal of Neuroscience*, 2260-17.

Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D., & Soltani, A. (2019). Flexible combination of reward information across primates. Nature human behaviour, 3(11), 1215-1224.

Fellows, L. K. (2006). Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. *Brain*, *129*(4), 944-952.

Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.

Hare, T. A., Schultz, W., Camerer, C. F., O'Doherty, J. P., & Rangel, A. (2011). Transformation of stimulus value signals into motor commands during simple choice. *Proceedings of the National Academy of Sciences*, 201109322.

Hayden, B., & Gallant, J. (2013). Working memory and decision processes in visual area v4. Frontiers in neuroscience, 7, 18.

Hayden, B. Y. (2018). Economic choice: the foraging perspective. Current Opinion in Behavioral Sciences, 24, 1-6.

Heilbronner, S. R., & Hayden, B. Y. (2016). Dorsal anterior cingulate cortex: a bottom-up view. *Annual review of neuroscience*, *39*, 149-170.

Heilbronner, S. R., & Hayden, B. Y. (2016). The description-experience gap in risky choice in nonhuman primates. Psychonomic bulletin & review, 23(2), 593-600.

Hunt, L. T., Behrens, T. E., Hosokawa, T., Wallis, J. D., & Kennerley, S. W. (2015). Capturing the temporal evolution of choice across prefrontal cortex. *Elife*, *4*, e11945.

Hunt, L. T., Dolan, R. J., & Behrens, T. E. (2014). Hierarchical competitions subserving multi-attribute choice. *Nature neuroscience*, *17*(11), 1613.

Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. *Nature Reviews Neuroscience*, *18*(3), 172.

Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F., & Behrens, T. E. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature neuroscience*, *15*(3), 470.

Hunt, L. T., Malalasekera, W. N., de Berker, A. O., Miranda, B., Farmer, S. F., Behrens, T. E., & Kennerley, S. W. (2018). Triple dissociation of attention and decision computations across prefrontal cortex. *Nature neuroscience*, *21*(10), 1471.

Hosokawa, T., Kennerley, S. W., Sloan, J., & Wallis, J. D. (2013). Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *Journal of Neuroscience*, *33*(44), 17385-17397.

Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron*, *63*(6), 733-745.

Kahnt, T., Heinzle, J., Park, S. Q., & Haynes, J. D. (2011). Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. *Neuroimage*, *56*(2), 709-715.

Kennerley, S. W., Dahmubed, A. F., Lara, A. H., & Wallis, J. D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. *Journal of cognitive neuroscience*, *21*(6), 1162-1178.

Kim, S., Hwang, J., & Lee, D. (2008). Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron*, *59*(1), 161-172.

539   Klein-Flügge, M. C., & Bestmann, S. (2012). Time-dependent changes in human corticospinal
540   excitability reveal value-based competition for action during decision processing.
541   *Journal of neuroscience*, *32*(24), 8373-8382.
542
543   Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-
544   connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, *2*,
545   4.
546
547   Levy, D. J., & Glimcher, P. W. (2012). The root of all value: a neural common currency for
548   choice. *Current opinion in neurobiology*, *22*(6), 1027-1038.
549
550   Louie, K., Grattan, L. E., & Glimcher, P. W. (2011). Reward value-based gain control: divisive
551   normalization in parietal cortex. Journal of Neuroscience, 31(29), 10627-10639.
552
553   Maxwell, S. E., & Delaney, H. D. (1993). Bivariate median splits and spurious statistical
554   significance. *Psychological bulletin*, *113*(1), 181.
555
556   Montague, P. R., & Berns, G. S. (2002). Neural economics and the biological substrates of
557   valuation. *Neuron*, *36*(2), 265-284.
558
559   Noguchi, T., & Stewart, N. (2018). Multialternative decision by sampling: A model of decision
560   making constrained by process data. *Psychological review*, *125*(4), 512.
561
562   O'Doherty, J. P. (2014). The problem with value. *Neuroscience & Biobehavioral Reviews*, *43*,
563   259-268.
564
565   O'Neill, M., & Schultz, W. (2018). Predictive coding of the statistical parameters of uncertain
566   rewards by orbitofrontal neurons. *Behavioural brain research*.
567
568   Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annual
569   review of neuroscience*, *34*, 333-359.
570
571   Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode
572   economic value. *Nature*, *441*(7090), 223.
573
574   Passino, K. M., Seeley, T. D., & Visscher, P. K. (2008). Swarm cognition in honey bees.
575   *Behavioral Ecology and Sociobiology*, *62*(3), 401-414.
576
577   Paxinos, G., Huang, X. F., & Toga, A. W. (2000). The rhesus monkey brain in stereotaxic
578   coordinates.
579
580   Pirrone, A., Azab, H., Hayden, B. Y., Stafford, T., & Marshall, J. A. (2018). Evidence for the
581   speed–value trade-off: Human and monkey decision making is magnitude sensitive.
582   Decision, 5(2), 129.
583

Raghuraman, A. P., & Padoa-Schioppa, C. (2014). Integration of multiple determinants in the neuronal computation of economic values. *Journal of Neuroscience*, *34*(35), 11583-11603.

Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature reviews neuroscience*, *9*(7), 545.

Rangel, A., & Hare, T. (2010). Neural computations associated with goal-directed choice. *Current opinion in neurobiology*, *20*(2), 262-270.

Roesch, M. R., & Olson, C. R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. Science, 304(5668), 307-310.

Roesch, M. R., & Olson, C. R. (2003). Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. Journal of Neurophysiology, 90(3), 1766-1789.

Rushworth, M. F., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron*, *70*(6), 1054-1069.

Rustichini, A., & Padoa-Schioppa, C. (2015). A neuro-computational model of economic decisions. Journal of neurophysiology, 114(3), 1382-1398.

Schall, J. D., Stuphorn, V., & Brown, J. W. (2002). Monitoring and control of action by the frontal lobes. *Neuron*, *36*(2), 309-322.

Seeley, T. D. (2009). *The wisdom of the hive: the social physiology of honey bee colonies*. Harvard University Press.

Seeley, T. D., & Buhrman, S. C. (1999). Group decision making in swarms of honey bees. *Behavioral Ecology and Sociobiology*, *45*(1), 19-31.

Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, *335*(6064), 108-111.

Shizgal, P. (1997). Neural basis of utility estimation. *Current opinion in neurobiology*, *7*(2), 198-208.

Sleezer, B. J., Castagno, M. D., & Hayden, B. Y. (2016). Rule encoding in orbitofrontal cortex and striatum guides selection. Journal of Neuroscience, 36(44), 11223-11237.

So, N. Y., & Stuphorn, V. (2010). Supplementary eye field encodes option and action value for saccades with variable reward. *Journal of Neurophysiology*, *104*(5), 2634-2653.

630 Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. Cognitive psychology,
631   53(1), 1-26.
632

633 Strait, C. E., Blanchard, T. C., & Hayden, B. Y. (2014). Reward value comparison via mutual
634   inhibition in ventromedial prefrontal cortex. *Neuron*, *82*(6), 1357-1366.
635

636 Strait, C. E., Sleezer, B. J., Blanchard, T. C., Azab, H., Castagno, M. D., & Hayden, B. Y.
637   (2015). Neuronal selectivity for spatial positions of offers and choices in five reward
638   regions. *Journal of neurophysiology*, *115*(3), 1098-1111.
639

640 Strait, C. E., Sleezer, B. J., & Hayden, B. Y. (2015). Signatures of value comparison in ventral
641   striatum neurons. *PLoS biology*, *13*(6), e1002173.
642

643 Thura, D., & Cisek, P. (2016). Modulation of premotor and primary motor cortical activity
644   during volitional adjustments of speed-accuracy trade-offs. *Journal of Neuroscience*,
645   *36*(3), 938-956.
646

647 Tversky, A. (1969). Intransitivity of preferences. *Psychological review*, *76*(1), 31.
648

649 Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases.
650   *science*, *185*(4157), 1124-1131.
651

652 Wang, M. Z., & Hayden, B. Y. (2017). Reactivation of associative structure specific outcome
653   responses during prospective evaluation in reward-based choices. Nature
654   communications, 8, 15821.
655

656 Wunderlich, K., Rangel, A., & O'Doherty, J. P. (2009). Neural computations underlying
657   action-based decision making in the human brain. *Proceedings of the National
658   Academy of Sciences*, *106*(40), 17199-17204.
659

660 Yoo, S. B. M., & Hayden, B. Y. (2018). Economic choice as an untangling of options into
661   actions. *Neuron*, *99*(3), 434-447.
662

663

664

MATERIALS AND METHODS

665
666
667   The data used here were analyzed and summarized in other manuscripts (Strait et al.,
668   2015; Azab & Hayden, 2017; Azab & Hayden, 2018; Farashahi et al., 2018).
669

Surgical Procedures

671   All procedures were approved by the University Committee on Animal Resources at the
672   University of Rochester and were designed and conducted in compliance with the Public Health
673   Service's Guide for the Care and Use of Animals. Two male rhesus macaques (*Macaca mulatta*:
674   subject B age 5y. 7mo.; subject J age 6y. 7mo. at the start of recording) served as subjects in both
675   studies. We used standard procedures, as described previously (Strait et al., 2014). A small
676   prosthesis for holding the head was used. Animals were habituated to laboratory conditions and
677   then trained to perform oculomotor tasks for liquid reward. A Cilux recording chamber (Crist
678   Instruments, Hagerstown, Maryland, USA) was placed over the dACC and attached to the
679   calvarium with ceramic screws. Appropriate anaesthesia was used at all times; induction was
680   performed with ketamine and isoflurane was used for maintenance. For surgical induction, we
681   used 10–15 mg/kg of ketamine, 0.25 mg/kg of midazolam, and 2– 4 mg/kg of propofol. For
682   maintenance, we used isoflurane, ad lib level, set depending on active monitoring procedure. For
683   systemic antibiotics, we used cefazolin and for topical application, we used standard veterinary
684   triple antibiotic. For analgesics, we used meloxicam, and, when judged necessary by veterinary
685   staff, buprenorphine. Post-operative care included close monitoring and restoration of fluid
686   intake. Animals received appropriate analgesics and antibiotics after all procedures. Position was
687   verified by magnetic resonance imaging with the aid of a Brainsight system (Rogue Research
688   Inc., Montreal, Quebec, Canada). Throughout both behavioral and physiological recording
689   sessions, the chamber was kept sterile with regular antibiotic washes and sealed with sterile caps.
690

Recording Site

692   We approached dACC through a standard recording grid (Crist Instruments). We defined
693   dACC according to the Paxinos atlas (Paxinos et al, 2000). Roughly, we recorded from a ROI
694   lying within the coronal planes situated between 29.50 and 34.50 mm rostral to interaural plane,
695   the horizontal planes situated between 4.12 to 7.52 mm from the brain's dorsal surface, and the
696   sagittal planes between 0 and 5.24 mm from medial wall (Figure 1B). Our recordings were made
697   from a central region within this zone. We confirmed recording location before each recording
698   session using our Brainsight system with structural magnetic resonance images taken before the
699   experiment. Neuroimaging was performed at the Rochester Center for Brain Imaging, on a
700   Siemens 3T MAGNETOM Trio Tim using 0.5 mm voxels. We confirmed recording locations by
701   listening for characteristic sounds of white and gray matter during recording, which in all cases
702   matched the loci indicated by the Brainsight system.
703

Electrophysiological Techniques

705   Single electrodes (Frederick Haer & Co., Bowdoin, Maine, USA; impedance range 0.8 to
706   4 MU) were lowered using a microdrive (NAN Instruments, Nazaret Illit, Israel) until waveforms
707   of between one and three neuron(s) were isolated. Individual action potentials were isolated on a
708   Plexon system (Plexon, Inc., Dallas, Texas, USA). Neurons were selected for study solely on the
709   basis of the quality of isolation; we never pre-selected based on task-related response properties.
710   All collected neurons for which we managed to obtain at least 250 trials were analyzed.

711
712   **Eye Tracking and Reward Delivery**
713   Eye position was sampled at 1,000 Hz by an infrared eye-monitoring camera system (SR
714   Research, Ottawa, Ontario, Canada). Stimuli were controlled by a computer running Matlab
715   (Mathworks) with Psychtoolbox (Brainard, 1997) and Eyelink Toolbox (Cornelissen, Peters &
716   Palmer, 2002). Visual stimuli were colored rectangles on a computer monitor placed 57 cm from
717   the animal and centered on its eyes. A standard solenoid valve controlled the duration of juice
718   delivery. The relationship between solenoid open time and juice volume was established and
719   confirmed before, during, and after recording.
720
721   **Behavioral Task**
722   Subjects performed a two-option gambling task. The task was conceptually similar to
723   risky choice tasks we have used previously (Strait et al., 2014), with two major differences: (1)
724   monkeys gambled for virtual tokens—rather than liquid—rewards, and thus (2) outcomes could
725   be losses as well as wins. Previous training history for these subjects included a set shifting task
726   (Sleezer et al., 2016), a simple choice task (Heilbronner and Hayden, 2016), an attentional task
727   (Hayden and Gallant, 2013), and several foraging and foraging-inspired tasks (see Calhoun and
728   Hayden, 2014; Hayden, 2018)
729   Two offers were presented on each trial. Each offer was represented by a rectangle 300
730   pixels tall and 80 pixels wide (11.35° of visual angle tall and 4.08° of visual angle wide). 20% of
731   options were safe (100% probability of either 0 or 1 token), while the remaining 80% were
732   gambles. Safe offers were entirely red (0 tokens) or blue (1 token). The size of each portion
733   indicated the probability of the respective reward. Each gamble rectangle was divided
734   horizontally into a top and bottom portion, each colored according to the token reward offered.
735   Gamble offers were thus defined by three parameters: two possible token outcomes, and
736   probability of the top outcome (the probability of the bottom was strictly determined by the
737   probability of the top). The top outcome was 10%, 30%, 50%, 70% or 90% likely on gamble
738   offers.
739   Six initially unfilled circles arranged horizontally at the bottom of the screen indicated
740   the number of tokens to be collected before the subject obtained a liquid reward. These circles
741   were filled appropriately at the end of each trial, according to the outcome of that trial. When 6
742   or more tokens were collected, the tokens were covered with a solid rectangle while a liquid
743   reward was delivered. Tokens beyond 6 did not carry over, nor could number of tokens fall
744   below zero.
745   On each trial, one offer appeared on the left side of the screen and the other appeared on
746   the right. Offers were separated from the fixation point by 550 pixels (27.53° of visual angle).
747   The side of the first offer (left and right) was randomized by trial. Each offer appeared for 600
748   ms and was followed by a 150 ms blank period. Monkeys were free to fixate upon the offers
749   when they appeared (and in our observations almost always did so). After the offers were
750   presented separately, a central fixation spot appeared and the monkey fixated on it for 100 ms.
751   Following this, both offers appeared simultaneously and the animal indicated its choice by
752   shifting gaze to its preferred offer and maintaining fixation on it for 200 ms. Failure to maintain
753   gaze for 200 ms did not lead to the end of the trial, but instead returned the monkey to a choice
754   state; thus, monkeys were free to change their mind if they did so within 200 ms (although in our
755   observations, they seldom did so). A successful 200 ms fixation was followed by a 750 ms delay,
756   after which the gamble was resolved and a small reward (100 μL) was delivered—regardless of

757   the outcome of the gamble—to sustain motivation. This small reward was delivered within a 300
758   ms window. If 6 tokens were collected, a further delay of 500 ms was followed by a large liquid
759   reward (300 μL) within a 300 ms window, followed by a random inter-trial interval (ITI)
760   between 0.5 and 1.5 s. If 6 tokens were not collected, subjects proceeded immediately to the ITI.
761   Each gamble included at least one positive or zero-outcome, ensuring that every gamble
762   carried the possibility of a win (or at least no change in tokens). This decreased the number of
763   trivial choices presented to subjects, and maintained motivation.
764
765   **Statistical Methods**
766   **Peri-stimulus time histograms (PSTHs)** were constructed by aligning spike rasters to
767   the presentation of the first offer and averaging firing rates across multiple trials. Firing rates
768   were calculated in 20 ms bins but were generally analyzed in longer epochs. For latency
769   analyses, we used PSTHs with firing rates calculated in 5 ms bins to grant us finer resolution to
770   detect shorter latencies.
771   Firing rates were **normalized** by subtracting the mean and dividing by the standard
772   deviation of the neuron's entire PSTH (i.e. z-scoring).
773   We tested for correlations between variables of interest, since these could produce
774   spurious results in our format and population correlation analyses. Although gamble values were
775   chosen independently, safe gambles (where the large and small outcome were of the same value)
776   introduced a spurious correlation between large and small outcomes. We thus exclude trials
777   where the gamble of interest was a safe gamble, as we have done in previous manuscripts (refs).
778   We tested for single-unit modulation using a **multiple generalized linear regression**
779   **model**, including the following task-relevant variables: large outcomes, small outcomes, and
780   probability of the large outcomes for gambles 1 (and two, if that gamble had been presented), the
781   number of tokens collected as of the beginning of the trial, and the side the first offer appeared
782   on. We also included the side of the chosen offer for analyses in the pre-choice and post-choice
783   epochs. We use the same variables to fit stepwise linear regression models to assess small-
784   outcome modulation.
785   **Analysis epochs** were chosen *a priori*, before data analysis began, to align with epochs
786   we have chosen for previous studies (Strait et al., 2014; Azab & Hayden, 2017; Azab & Hayden,
787   2018). The first and second offer epochs were defined as the 500 ms epoch beginning 100 ms
788   after the offer was presented. The pre-choice epoch was the 500 ms epoch before choice was
789   indicated using an express saccade. The post-choice epoch was the 500 ms epoch immediately
790   after choice fixation was complete, and before reward feedback was revealed. These epochs were
791   chosen before data collection began; indeed, they originally were selected for a previous study
792   on a different brain region, and to account for hypothesized delays in neural responding (Strait et
793   al., 2014). We maintained the same epochs here as a mechanism to reduce the likelihood of p-
794   hacking.
795   All **fractions of neurons** were tested for significance using a one-sided binomial test.
796   When testing for significant encoding in the population, we use an alpha level of 0.05 to indicate
797   chance (i.e. the number of neurons that would exhibit significant modulation at random).
798   **Format and population correlation analyses** were performed in the following manner.
799   We use **beta correlation analyses** to assess whether neurons represented two variables (or the
800   same variable at different time periods) using similar / orthogonal / opposing formats, in
801   overlapping / orthogonal / distinct populations. To do this, we first found the regression
802   coefficient associated with the variable of interest per neuron. We estimate this using a multiple

803 linear regression model including the same variables detailed above (for spatial location
804 analyses, we split trials by the side of offer appearance, and thus do not include that variable in
805 our regression model). We then combined the regression coefficients associated with the variable
806 of interest into a vector of the same length as the number of neurons in our sample. This vector
807 indicates the relative strength (after normalizing) and direction of modulation for each individual
808 neuron in the population, in response to a particular variable in a particular epoch. We call this
809 the population "format". We compared different formats by finding the Spearman correlation
810 coefficient between them.
811   Note that this approach provides a more sensitive method of examining population
812 properties than conventional approaches, which involve determining which cells cross a
813 significance threshold and then using those as focal cells for further analyses. By taking on all
814 cells regardless of their response, our method accomplishes two things. First, it uses all available
815 information, even that information that is not sufficient by itself to achieve significance. Second,
816 it avoids introducing a hard categorical boundary, which can introduce false positives in data
817 (Blanchard et al., 2017; Maxwell & Delaney, 1993).
818   We extend this method to account for the noise inherent in estimating each neuron's
819 encoding of each variable of interest, which existing methods do not account for (Azab &
820 Hayden, 2017). We used a Bayesian regression to obtain a probabilistic distribution over each
821 regression coefficient for each neuron (rather than an individual coefficient estimate per neuron;
822 this is akin to taking into account the confidence interval on each regression coefficient
823 estimate). We sampled 1,000 regression coefficients from this distribution for each neuron, to
824 obtain a probability distribution of (1,000) potential formats for the population, for each of two
825 task conditions. We then performed the correlation analyses on each of these samples across task
826 conditions, thus generating a probability distribution of (1,000) correlation coefficients. This is a
827 more robust estimate of the correlation between formats, as it takes into account the uncertainty
828 inherent in estimating any individual regression coefficient, and allows us to view the spread of
829 the distribution of this correlation when this significant source of noise is taken into account.
830   The Spearman correlation coefficient between signed regression coefficients indicated
831 whether variables were represented in a similar *format* i.e. directionality of tuning across the
832 population. (Note that we used Spearman instead of Pearson because it is more robust to outliers
833 and is therefore more statistically conservative). A positive correlation indicated a preservation
834 of directionality, while a negative correlation suggested variables were represented in opposing
835 directionality of firing rate modulation. No correlation suggests orthogonal formats, but we draw
836 no strong conclusions from these.
837   Similarly, the correlation coefficient between unsigned regression coefficients indicated
838 whether similar neuronal populations tended to be involved in encoding the two variables in
839 question, regardless of their direction of modulation. A positive correlation indicated overlapping
840 populations, while a negative correlation indicated separate ones. A lack of correlation suggests
841 orthogonal populations (i.e. encoding one variable does not affect the neuron's likelihood of
842 encoding the other variable), although given the lack of interpretability of null results, it does not
843 constitute evidence for lack of true correlation.
844   It is important to note that the analysis method we use here is rooted in previous analysis
845 techniques used by our lab and others. The closest relative is Representational Similarity
846 Analysis (RSA; Kriegeskorte, Mur & Bandettini, 2008; Hunt et al., 2018). In RSA, neural
847 representations (vectors of neural responses) associated with different behavioral conditions are
848 compared and their similarity characterized to facilitate comparison with different data sets,

849   sometimes collected using different means. Rather than comparing raw neural responses, though,

850   our method relies on comparing regression coefficients; this allows us to hone in on the signal

851   associated with a particular variable, allowing us to compare the neural responses to different

852   variables at the same point in time.

853

854