Partial Adaptation to the Value Range in the Macaque Orbitofrontal Cortex

Katherine E. Conen¹ and Camillo Padoa-Schioppa^{1,2,3}

Departments of ¹Neuroscience, ²Economics, and ³Biomedical Engineering, Washington University, St Louis, Missouri 63110

Values available for choice in different behavioral contexts can vary immensely. To compensate for this variability, neuronal circuits underlying economic decisions undergo adaptation. In orbitofrontal cortex (OFC), neurons encode the subjective value of offered and chosen goods in a quasilinear way. Previous experiments found that the gain of the encoding is lower when the value range is wider. However, the parameters OFC neurons adapted to remained unclear. Furthermore, previous studies did not examine additive changes in neuronal responses. Computational considerations indicate that these factors can directly impact choice behavior. Here we investigated how OFC neurons adapt to changes in the value range. We recorded from two male rhesus monkeys during a juice choice task. Each session was divided into two blocks of trials. In each block, juices were offered within a set range of values, and ranges changed between blocks. Across blocks, neuronal responses adapted to both the maximum and the minimum value, but only partially. As a result, the minimum neural activity was elevated in some value ranges relative to others. Through simulation of a linear decision model, we showed that increasing the minimum response increases choice variability, lowering the expected payoff. This effect is modulated by the balance between cells with positive and negative encoding. The presence of these two populations induces a non-monotonic relationship between the value range and choice efficacy, such that the expected payoff is highest for decisions in an intermediate value range.

Key words: decision-making; economic choice; neuronal plasticity; optimal coding; range adaptation; subjective value

Significance Statement

Economic decisions are thought to rely on the orbitofrontal cortex (OFC). The values available for choice vary enormously in different contexts. Previous work showed that neurons in OFC encode values in a linear way, and that the gain of encoding is inversely related to the range of available values. However, the specific parameters driving adaptation remained unclear. Here we show that OFC neurons adapt to both the maximum and minimum value in the current context. However, adaptation is partial, leading to contextual changes in the response offset. Interestingly, increasing the activity offset negatively affects choices in a simulated network. Partial adaptation may allow the circuit to maintain information about context value at the cost of slightly reduced payoff.

Introduction

Neuronal adaptation takes place throughout the brain. Although its function is not fully understood, in sensory systems adaptation may contribute to homeostatic regulation (Benucci et al., 2013; Hengen et al., 2013), efficient perceptual representation (Dan et

al., 1996; Lewicki, 2002; Gutnisky and Dragoi, 2008; Adibi et al., 2013), and sharper behavioral performance (Krekelberg et al., 2006; Liu et al., 2016). Context adaptation has also been observed in the neuronal representation of subjective values. Studies in nonhuman primates have found adaptive coding in several brain regions, including orbitofrontal cortex (OFC; Padoa-Schioppa, 2009; Kobayashi et al., 2010; Yamada et al., 2018), anterior cingulate cortex (Cai and Padoa-Schioppa, 2014), and the amygdala (Bermudez and Schultz, 2010; Saez et al., 2017). In humans, experiments measuring BOLD activity have shown context-adapting value signals in ventromedial prefrontal cortex (vmPFC), ventral striatum, and other brain areas (Elliott et al., 2008; Cox and Kable, 2014; Burke et al., 2016). More recent work has begun to explore the behavioral implications of value adaptation using a combination of experimental and theoretical approaches. One study found that adaptation in OFC reduces variability in value-based decisions, increasing the subject's ex-

Received Sept. 4, 2018; revised Jan. 21, 2019; accepted Feb. 13, 2019.

Author contributions: K.E.C. and C.P.-S. designed research; K.E.C. performed research; K.E.C. analyzed data; K.E.C. wrote the paper; C.P.-S. edited the paper.

This work was supported by the National Institutes of Health (Grants R01-MH104494 to C.P.-S. and F31-MH107111 to K.E.C.). We thank H. Schoknecht for help with animal training and S. Ballesta, W. Shi, and E. Bromberg-Martin for comments on earlier versions of the paper.

The authors declare no competing financial interests.

Correspondence should be addressed to Camillo Padoa-Schioppa at camillo@wustl.edu.

K. E. Conen's present address: Department of Collective Behaviour, Max Planck Institute for Ornithology, and Department of Biology, University of Konstanz, Universitätsstraße 10, 78464 Konstanz, Germany.

<https://doi.org/10.1523/JNEUROSCI.2279-18.2019>

Copyright © 2019 the authors

pected payoff (Rustichini et al., 2017). Other work suggests that value adaptation on a shorter time scale may induce irrational decision patterns (Soltani et al., 2012; Yamada et al., 2018).

Despite these advances, many aspects of range adaptation remain poorly understood. For example, previous studies did not distinguish between neurons adapting to the value range (i.e., the difference between maximum and minimum value) and neurons adapting to the maximum value (Padoa-Schioppa, 2009; Kobayashi et al., 2010; Cox and Kable, 2014). Also, previous work focused exclusively on gain adaptation (Padoa-Schioppa, 2009; Kobayashi et al., 2010; Cox and Kable, 2014; Rustichini et al., 2017) and did not examine changes in the intercept of the response function. To address these outstanding issues, we recorded from monkeys engaged in a juice choice task. We focused on OFC, an area engaged in value-based decisions (Fellows, 2011; Wallis, 2012; Rudebeck and Murray, 2014; Schultz, 2015; Padoa-Schioppa and Conen, 2017). Each neuron was recorded during consecutive blocks of trials. Within each trial block, offer values for each juice varied within a specified range. However, value ranges varied across blocks. Neurons adapted to both the maximum and the minimum values available in the current context. However, responses did not remap completely to the new value range (partial adaptation). In particular, the encoding slopes measured with wide value ranges were steeper than expected under full adaptation. Furthermore, the response to the minimum value increased as a function of the minimum value. Importantly, partial remapping reflected the final adapted state of neurons, not simply an incomplete temporal process.

We complemented these experimental results with a series of simulations. Using a linear decision model, we showed that increasing the minimum activity adds noise to the decision process, increasing choice variability and reducing the expected payoff. However, this theoretical loss is relatively minor compared with the effect of narrowing the dynamic range. Moreover, the presence of both positive and negative value encoding moderates this effect, keeping expected payoff more consistent across different ranges. Incomplete adaptation may allow the circuit to remain flexible and maintain information about the overall value of the context, at the cost of a slight decrease in expected payoff.

Materials and Methods

Experimental procedures. Two adult male macaques (*Macaca mulatta*; Monkey D, 11.5 kg; Monkey F, 11.0 kg) participated in this study. During the experiment, the animal sat in an electrically insulated enclosure (Crist Instruments) with its head fixed. Cues were displayed on a computer monitor placed 57 cm in front of the animal. Monkeys performed a variant of a juice choice task used in several previous studies (Padoa-Schioppa and Assad, 2006). The task was run on custom-written software (<http://www.monkeylogic.net/>) based on MATLAB (MathWorks). Eye position was monitored with an infrared video camera (EyeLink, SR Research).

In each session a monkey chose between two juices, A and B, offered in varying quantities. Juice A was defined as the preferred juice (i.e., 1A was generally chosen over 1B). In each trial, the monkey began by fixating on a central point. After 1 s, cues appeared on each side of the central fixation, indicating the current range of possible offers. The cues consisted of a set of filled and empty squares. The color of the squares indicated the juice type, the total number of squares represented the maximum possible offer for that juice in the current trial, and the filled squares represented the minimum possible offer in that trial (Fig. 1A). The cues remained on screen for 1 s and were then replaced by a set of solid squares denoting the offers on the current trial. After a randomly variable delay (1–2 s), the central fixation point disappeared and targets appeared next to each offer (go signal). The monkey indicated its choice with a saccade to one of the targets and, after 0.75 s, received the juice corresponding to

the chosen offer. If the monkey broke fixation before the go signal or if it failed to fixate the target for 0.75 s after the saccade, the trial was aborted and the monkey received no reward.

Each session consisted of 2–3 blocks, each lasting ~250 trials. The offered quantity varied pseudorandomly from trial to trial within a defined range. Within a block, the range of possible offers was kept constant for each juice. The monkey could either learn the value range implicitly through experience or explicitly using the range cues. We do not attempt to distinguish between these possibilities. Between blocks, the range of available offers for each juice changed. There were three possible ranges for each juice: “low” (0–3 uA or 0–6 uB), “high” (2–5 uA or 4–10 uB), and “wide” (0–5 uA or 0–10 uB). Most range transitions consisted of an increase/decrease in the minimum value ($V_{\min}$), whereas the maximum value ($V_{\max}$) either remained constant or shifted in conjunction with $V_{\min}$. When $V_{\min}$ and $V_{\max}$ changed together, the difference $V_{\max} - V_{\min}$ was kept constant. We counterbalanced the type of range transition across sessions. In a small subset of sessions, $V_{\max}$ increased/decreased while $V_{\min}$ was kept at zero. The ranges of Juice A and Juice B could change in the same direction or in different directions.

Before training, a head-restraint device and a recording chamber were implanted on the skull under general anesthesia. The recording chamber (main axes: 50 × 30 mm) was centered on inter-aural coordinates (A30, L0). Structural MRI scans were obtained before and after implantation and used to guide recordings. We recorded neuronal data from the central OFC, in a region approximately corresponding to area 13m (Ongür and Price, 2000; Monkey D: A 31:36, L –6:–10; Monkey F: A 31:37 L –6:–11 and 6:11). Recordings were conducted using tungsten electrodes (diameter: 125 μ m; FHC) or 16-channel silicon V-probes (diameter: 185 μ m, spacing between electrodes: 100 μ m; Plexon). Electrodes were lowered vertically into position each day using a custom-built microdrive (step size: 2.5 μ m). The recording depth was determined ahead of time based on structural MRI.

Electrical signals were amplified (gain: 10,000) and bandpass filtered (low-pass cutoff: 300 Hz, high-pass cutoff: 6 kHz; Lynx 8, Neuralynx). Action potentials were detected on-line by setting a threshold during recording, and waveforms crossing the threshold were saved (sampling rate: 40 kHz; Power 1401, Cambridge Electronic Design). Spike sorting was conducted off-line using standard software (Spike 2, Cambridge Electronic Design). Neurons were included in the analysis if they remained stable and well isolated for at least 120 trials in each of two blocks. Responses that were not stably isolated for the full session were only analyzed for the trials in which they were stable. In the V-probe recordings, spikes from the same neuron were occasionally picked up by two neighboring contacts. These instances were detected manually based on the consistent presence of simultaneous spikes. If units in neighboring channels shared >70% of spikes, they were considered duplicates, and one of the units was excluded from the analysis.

All experimental procedures conformed to the NIH *Guide for the Care and Use of Laboratory Animals* and were approved by the Institutional Animal Care and Use Committee at Washington University in St Louis.

Experimental design and statistical analyses. The study included two male rhesus monkeys. We recorded from several hundred cells in each animal (1262 neurons total). All analyses were performed in MATLAB (MathWorks). Pairwise comparisons were evaluated using the Wilcoxon signed rank test. Comparisons between two distributions were made using the Wilcoxon rank sum test. Correlations refer to Pearson's correlation unless otherwise specified. *P* values for sign rank and rank sum tests were calculated by normal approximation using MATLAB's built-in functions. Tests for bimodality were done using Hartigan's dip test, adapted from code written by Nicholas Price (Monash University, Victoria, Australia; <http://www.nicprice.net/diptest/>). *P* values for the dip test were calculated by bootstrap permutation with 1000 repetitions. *P* values are reported without adjustment for multiple comparisons. Additional details on the data analysis are provided in the following three sections. Data and code are available upon reasonable request.

Analysis of behavior. Choice behavior was analyzed separately in each trial block. We defined the choice pattern as the percentage of trials in which the animal chose Juice B as a function of the offer ratio q_B/q_A .

Figure 1. Task outline and behavioral results. **A**, Schematic of a session. The monkey chose between two juices, each associated with one color. The animal initiates each trial by fixating on a central point. After 1 s, range cues appear on either side of the central fixation. Filled squares indicate the minimum possible offer for a given juice and total squares indicate the maximum. In this trial, the cue informed the monkey that the quantity of Juice B offered would be between 4 and 10 drops, whereas the quantity of Juice A would be between 0 and 3 drops. After 1 s, the cues are replaced by two sets of filled squares representing the current offers. In this trial, the animal chose between 1 drop of grape juice and 4 drops of fruit punch. After a variable interval (1–2 s), the central fixation disappears, and targets appear next to each offer, cueing the monkey to indicate its choice. The monkey then makes a saccade to one of the targets and holds fixation for 0.75 s, after which it receives the chosen juice. Each session consists of 2–3 blocks, each with ~250 trials. Ranges remain constant across trials within each block and change between blocks. **B**, Each juice is offered in one of three ranges: low, high, or wide. The two juices can be offered in either the same type of range or different types of range within a block. **C–H**, Changes in choice behavior across sessions. **C**, **D**, Illustration of the behavioral response function for changing values of relative value ρ (**C**), or behavioral steepness η (**D**). Increased ρ corresponds to a decreased probability of choosing offer B for a given offer (q_B/q_A). Increased η corresponds to less variable choice behavior. **E–H**, Animal behavior across sessions. Each point represents the behavior from one pair of blocks in a session ($n = 205$). **E**, Relative value in the earlier versus later block of a session. Values of ρ were strongly correlated across pairs of blocks in a session ($r = 0.73$, $p = 5.5 \times 10^{-35}$, Pearson correlation) and slightly elevated in later blocks (median $\Delta\rho = 0.07$; $p = 0.011$, $W = 8.4 \times 10^4$). **F**, Values of η were correlated across blocks ($r = 0.24$, $p = 4.4 \times 10^{-4}$, Pearson correlation) but did not differ between the first and second blocks of a session ($p = 0.48$, $W = 1.1 \times 10^4$). **G**, Fractional difference in ρ across range types. The fractional difference is defined as the difference in parameter values divided by their sum. Median differences: $\rho_{\text{high}} - \rho_{\text{low}} = 0.040$ ($p = 2.3 \times 10^{-4}$, $W = 436$), $\rho_{\text{wide}} - \rho_{\text{low}} = 0.032$ ($p = 8.3 \times 10^{-5}$, $W = 1.5 \times 10^3$), $\rho_{\text{high}} - \rho_{\text{wide}} = 0.041$ ($p = 0.054$, $W = 479$). **H**, Fractional difference in η across range types. Median differences: $\eta_{\text{high}} - \eta_{\text{low}} = 0.11$ ($p = 7.7 \times 10^{-3}$, $W = 384$), $\eta_{\text{wide}} - \eta_{\text{low}} = 0.069$ ($p = 4.2 \times 10^{-3}$, $W = 1.4 \times 10^3$), $\eta_{\text{high}} - \eta_{\text{wide}} = 0.049$ ($p = 0.14$, $W = 449$). Black diamonds indicate median fractional differences, and asterisks indicate significant difference from 0 ($p < 0.05$). All p values, Wilcoxon signed rank test.

where q_X is the quantity of Juice X offered to the monkey. We fit the choice pattern to a sigmoid function using logistic regression:

$$P(\text{choice B}) = 1/(1 + e^{-x}),$$

$$X = a_0 + a_1 \log(q_B/q_A).$$

From this fit, we computed the relative value of the two juices (ρ) and the sigmoid steepness (η):

$$\rho = \exp(-a_0/a_1),$$

$$\eta = a_1.$$

We examined changes in ρ and η as a function of the range type. To do so, we compared data for all pairs of blocks within a session. We recorded during 107 sessions, each of which included 2 or 3 range conditions, yielding a total of 236 unique block pairs. Block pairs were excluded from the behavioral analysis if there were <2 offer types with choices split between the two juices (complete or quasi-complete separation; 31 block pairs excluded). This criterion was imposed because the logit regression did not fully converge, resulting in large errors of measure for η and ρ . For the remaining 205 block pairs, we computed a fractional difference for each parameter across different range types, where we defined the fractional difference as the difference divided by the sum.

Classification of neuronal responses. Because we were interested in the effects of adaptation at steady state, we discarded the first 16 trials of each block, thereby excluding trials where the monkey had not yet experienced the full range of values. We analyzed cell data in seven time windows following offer onset: post-offer (0.5 s after offer onset), late-delay (0.5–1.0 s after offer onset), pre-go (0.5 s before the go signal), reaction time (time from go cue to target acquisition, usually ~ 200 ms), post-juice (0.5 s after juice delivery), and post-juice 2 (0.5–1 s after juice delivery). Data were analyzed independently for each block. We defined a “trial type” as a set of two offers and the monkey’s choice between them. For example, if the monkey was offered 1A versus 6B and chose Juice B, the trial type would be [1A:6B; B]. A “neuronal response” was defined as the activity of one cell in one time window as a function of the trial type, across two blocks.

A response was considered task-related if it passed an ANOVA (factor: trial type; $p < 0.05$) in each block. To classify task-related responses, we regressed each response against variables *offer value A*, *offer value B*, *chosen value*, and *chosen juice*, separately in each block. Variables that provided a nonzero regression slope in both blocks ($p < 0.05$) were said to “explain” the response. When a response was explained by more than one variable, we assigned it to the variable providing the highest total R^2 (summed over the 2 blocks). Most subsequent analyses focused on *offer value* and *chosen value* responses. Because we were interested in the effects of changing the value distribution, we excluded responses from analysis if the value range differed by <0.5 units of Juice B (uB) between blocks (132 responses). We also excluded cells with dramatic changes in pretrial firing rate across blocks ($>1.6\times$ change during the fixation time window, 208 responses), because large variability in baseline activity could obscure effects on cell tuning. Including these responses in the analysis added noise but did not qualitatively alter the results.

Different neuronal responses encoded value with a positive slope (higher activity for higher value; 71%) or with a negative slope (lower activity for higher value; 29%). For our analyses, we rectified neuronal responses with a negative slope as follows:

$$s_{\text{rectified}} = -s_{\text{raw}} \\ b_{\text{rectified}} = b_{\text{raw}} + s_{\text{raw}}(V_{\min} + V_{\max}),$$

where s and b indicate the slope and y -intercept of the original linear regression (“raw”) and the rectified tuning function (“rectified”). $V_{\max}$ and $V_{\min}$ indicate the maximum and minimum value offered in the current condition. Rectification maintained the same dynamic range and same slope magnitude, but with a positive rather than negative sign. Hence, for a rectified response, the maximum evoked response corresponds to $V_{\max}$ rather than $V_{\min}$. Neuronal responses with positive and rectified negative slopes were pooled in our main analysis. However, we reproduced qualitatively similar results for positive and negative encoding responses examined separately.

Analysis of range adaptation. The analysis of range adaptation focused on *offer value A*, *offer value B*, and *chosen value* responses. We grouped responses into three types of range transition: change $V_{\max}$ only, change $V_{\min}$ only, and change both. Transition types could be divided further based on the direction of change (increase/decrease). For *offer value* responses, we controlled the value range so that each transition type was consistent across sessions. Thus, if we describe the offer value range as a fraction of the wide value range (ΔV_{wide}), the normalized ranges were 0–0.6 uV (low range), 0.4–1 uV (high range), and 0–1 uV (wide range) for all *offer value* responses. *Chosen value* ranges depended on the choice

pattern of the animal, and in particular on the relative value ρ , which varied across sessions even when the two juices were identical. For the purpose of this study, we considered the maximum/minimum chosen value changed if the difference between blocks was >0.5 uB.

For each response, we regressed the firing rate onto value separately in each block. We obtained the slope of encoding (s) from each fit. Slopes were compared directly across range types, and the relationship between slope and range was tested more precisely using adaptation ratios (see Results). We also used the regression to calculate the responses to the minimum and maximum values ($R_{\min}$ and $R_{\max}$):

$$R_{\min} = s * V_{\min} + b,$$

$$R_{\max} = s * V_{\max} + b,$$

where $V_{\min}$ and $V_{\max}$ are the minimum and maximum values in the current block and b is the y -intercept of the linear fit.

We computed the normalized difference for conditions where either $V_{\min}$ or $V_{\max}$ changed alone:

$$\Delta R_{\min} = (R_{\min, \text{wide}} - R_{\min, \text{narrow}}) / (R_{\max, \text{wide}} - R_{\min, \text{wide}}),$$

$$\Delta R_{\max} = (R_{\max, \text{wide}} - R_{\max, \text{narrow}}) / (R_{\max, \text{wide}} - R_{\min, \text{wide}}),$$

and for conditions where both $V_{\min}$ and $V_{\max}$ changed:

$$\Delta R_{\min} = (R_{\min, \text{high}} - R_{\min, \text{low}}) / (R_{\max, \text{high}} - R_{\min, \text{low}}),$$

$$\Delta R_{\max} = (R_{\max, \text{high}} - R_{\max, \text{low}}) / (R_{\max, \text{high}} - R_{\min, \text{low}}).$$

We also computed the values of $\Delta R_{\min}$ and $\Delta R_{\max}$ that would be predicted if neurons did not adapt at all (NA). In this case $\Delta R_{\min}$ and $\Delta R_{\max}$ were equivalent to the difference in $V_{\max}$ and $V_{\min}$ across conditions, normalized as above. For example, when either $V_{\max}$ or $V_{\min}$ changed alone:

$$\Delta R_{\min, \text{NA}} = (V_{\min, \text{wide}} - V_{\min, \text{narrow}}) / (V_{\max, \text{wide}} - V_{\min, \text{wide}}).$$

For *offer value* responses, changes in $V_{\min}$ and $V_{\max}$ were controlled. Thus, when $V_{\max}$ changed alone $\Delta R_{\max, \text{NA}} = 0.4$ and $\Delta R_{\min, \text{NA}} = 0$; when $V_{\min}$ changed alone $\Delta R_{\max, \text{NA}} = 0$ and $\Delta R_{\min, \text{NA}} = 0.4$; and when both changed, $\Delta R_{\max, \text{NA}} = \Delta R_{\min, \text{NA}} = 0.4$. For *chosen value* neurons, $\Delta R_{\max, \text{NA}}$ and $\Delta R_{\min, \text{NA}}$ depended on the relative value and the animal’s choice pattern in each session.

Figures 4 and 6 show average traces of *offer value* response activity normalized so that values and neural responses vary in the range [0 1]. Unless otherwise specified, these responses were normalized as follows. For cases where either $V_{\max}$ or $V_{\min}$ change alone:

$$R_{\text{norm}} = (R - R_{\min, \text{wide}}) / (R_{\max, \text{wide}} - R_{\min, \text{wide}}),$$

$$V_{\text{norm}} = (V - V_{\min, \text{wide}}) / (V_{\max, \text{wide}} - V_{\min, \text{wide}}).$$

For cases where $V_{\max}$ and $V_{\min}$ changed concurrently:

$$R_{\text{norm}} = (R - R_{\min, \text{low}}) / (R_{\max, \text{high}} - R_{\min, \text{low}}),$$

$$V_{\text{norm}} = (V - V_{\min, \text{low}}) / (V_{\max, \text{high}} - V_{\min, \text{low}}).$$

R_{norm} and V_{norm} denote normalized responses and values, R and V denote the non-normalized responses and values, and $R_{\max, j}$ and $R_{\min, j}$ indicate the response to $V_{\max}$ and $V_{\min}$ in range type j .

To study adaptation in early versus late trials after the range transition, we computed separate tuning functions for the first and second half of each block. Responses were excluded if the slope changed by a factor >5 within the first block (2 responses excluded). Including these responses did not substantially affect results, but did add noise to the data, particularly for changes in $V_{\max}$. Plots of mean tuning curves in the first and second halves of each block were normalized to the first half of the wide range block.

Simulations of a linear decision model. We constructed a linear decision model to examine how changes in $R_{\min}$ affect choices. The model consisted of a population of 10,000 *offer A* and *offer B* units (5000 per group). We considered two models: a “positive network” in which all units en-

coded value with a positive slope, and a “mixed network”, in which 30% of units encoded value with an unrectified negative slope. This proportion reflected the fraction of negative encoding *offer value* cells in our dataset. We defined $R_{\min}$ as the activity elicited by the minimum offered value. Each unit encoded *offer value* in a linear way. Thus, for a positive encoding unit, the response of unit i on trial t was as follows:

$$R_{i,t} = V_t * (R_{\max} - R_{\min}) + R_{\min} + y_{i,t}$$

where $R_{\min}$ is the response to the minimum value, $R_{\max}$ is the response to the maximum value, V_t is the value of the encoded juice on trial t , and $y_{i,t}$ is a noise term. Units of R and V are arbitrary. For a negative encoding unit, the response of unit i in trial t was as follows:

$$R_{i,t} = V_t * (R_{\min} - R_{\max}) + R_{\max} + y_{i,t}$$

Note that this equation resembles the response for a positive encoding unit, except that $R_{\max}$ and $R_{\min}$ are reversed. Because higher values evoke lower activity, $R_{\max}$ is the lowest evoked response in a negative encoding unit. When it is important to keep the distinction explicit, we use the terms $R_{\min}[\text{positive}]$ and $R_{\max}[\text{negative}]$ to denote the lowest evoked response for positive and negative encoding units respectively. Likewise, we use the terms $R_{\max}[\text{positive}]$ and $R_{\min}[\text{negative}]$ to denote their highest evoked responses.

Importantly, *offer value* neurons in OFC show small but significant noise correlations (r_{noise}) (Conen and Padoa-Schioppa, 2015). We generated a realistic correlation matrix Q for the population as described previously (Hardin et al., 2013; Conen and Padoa-Schioppa, 2015). We set $\text{mean}(r_{\text{noise}}) = 0.01$ for pairs of units encoding the same juice with the same sign and $\text{mean}(r_{\text{noise}}) = 0$ for pairs of units encoding different juices. In the mixed-network model, we also set $\text{mean}(r_{\text{noise}}) = 0$ for pairs of units encoding the same juice with opposite signs. To generate the vector of noise terms y_i for the population on each trial, we generated values of uncorrelated noise starting from the standard normal distribution $u_i \sim N(0,1)$. This was multiplied by the correlation matrix and scaled according to the Fano factor (F) and the mean response for the current offer type ($\langle R_V \rangle$) to obtain y_i :

$$y_i = Q u_i \langle R_V \rangle F^{0.5}.$$

The scaling factor $\langle R_V \rangle F^{0.5}$ accounted for the fact that the variance in firing rate is roughly proportional to the mean response (Conen and Padoa-Schioppa, 2015).

Using this model, we simulated choice behavior for increasing values of $R_{\min}$. We considered two scenarios: (1) units had a fixed maximum response ($R_{\max}[\text{positive}]$ and $R_{\min}[\text{negative}]$), or (2) units had a fixed activity range $|R_{\max} - R_{\min}|$. Because the scales for V and R are arbitrary, we conventionally set $R_{\max}[\text{positive}] = R_{\min}[\text{negative}] = 1$ for the first scenario and $|R_{\max} - R_{\min}| = 1$ for the second.

Each simulation consisted of 1000 trials. The value of each juice for a given trial was a randomly chosen integer ranging from 0 to 10. The decision on each trial was determined by comparing the net activity of *offer value A* and *offer value B* units. In mixed network simulations, the responses of negative encoding units were multiplied by -1 before the comparison.

For each scenario, we simulated the choice pattern for the neural population as values of $R_{\min}[\text{positive}]$ increased from 0 to 1 in increments of 0.01. In the context of adaptation, this is analogous to shifting the value range from $[0, X]$ to $[X, 2X]$, where X is the initial maximum value. In mixed network simulations, we set values of $R_{\max}[\text{negative}] = (1 - R_{\min}[\text{positive}])$, reflecting the fact that $R_{\max}[\text{negative}]$ and $R_{\min}[\text{positive}]$ change in opposite directions as the value range shifts. We repeated the process for five different values of F and ran the simulation 20 times for each value of F and $R_{\min}$.

As in a previous study (Rustichini et al., 2017), we measured the efficacy of choice behavior using the fractional lost value (FLV) defined as follows:

$$\text{FLV} = \langle \text{max value} - \text{chosen value} \rangle / \langle \text{max value} - \text{chosen value}_{\text{chance}} \rangle,$$

where *max value* refers to the higher value of the two offers on a given trial, *chosen value*_{chance} is the average of the two offers, and $\langle \rangle$ indicates an average across trials. Notably, FLV is inversely related to the choice efficacy: if a subject always chooses the *max value*, $\text{FLV} = 0$; if the subject chooses randomly, $\text{FLV} = 1$.

Results

Two monkeys performed a juice choice task (Fig. 1A). In each session, the monkey chose between two juices labeled A and B (with A preferred). Each session included 2–3 blocks of ~250 trials. Within each block, the quantity of juice offered varied pseudorandomly within a set range, defined by a minimum and maximum value ($V_{\min}$ and $V_{\max}$). In a given block, each juice could be offered in a low, high, or wide range (Fig. 1B). Between blocks, the range of offers for each juice changed in one of six ways: $V_{\max}$ increased/decreased, $V_{\min}$ increased/decreased, or both $V_{\max}$ and $V_{\min}$ increased/decreased concurrently while $(V_{\max} - V_{\min})$ remained constant.

We analyzed the animals' behavior separately in each trial block. A logistic regression of the choice pattern provided measures for the relative value (ρ) and the sigmoid steepness (η) (Fig. 1C,D; see Materials and Methods). Choice patterns generally presented a quality–quantity tradeoff between the juices (mean $\rho = 2.4$ across sessions). Within a session, ρ was strongly correlated across blocks ($r = 0.73$, $p = 5.5 \times 10^{-35}$, Pearson correlation; Fig. 1E), indicating that the juice preferences were fairly consistent within a session. Values of ρ increased slightly in the second block compared with the first, presumably reflecting the animals' increasing satiety: their preference shifted toward the preferred juice rather than the higher quantity ($p = 0.011$, $W = 8.4 \times 10^3$, Wilcoxon signed rank test). Values of η were also correlated across blocks ($r = 0.24$, $p = 4.4 \times 10^{-4}$, Pearson correlation) but did not differ systematically between the first and second blocks of a session ($p = 0.48$, $W = 1.1 \times 10^4$, Wilcoxon signed rank test; Fig. 1F).

Choice behavior was weakly affected by the value range (Fig. 1G,H). In general, relative values were slightly larger in high and wide range blocks compared with low range blocks (Fig. 1G), reflecting an increase in the relative value of A for higher quantities. High and wide range blocks also had steeper sigmoid functions than low range blocks (lower choice variability; Fig. 1H). The sigmoid steepness recorded in high range and wide range blocks was not statistically different (Fig. 1H). Differences in sigmoid steepness were likely related to the monkeys' greater motivation in high value blocks (see Discussion).

Neural responses adapt to both the maximum and minimum value

We recorded the activity of 1262 cells from the OFC of two monkeys (Monkey D, left hemisphere: 480 cells; Monkey F, left hemisphere: 373 cells, right hemisphere: 409 cells). Firing rates were analyzed in seven time windows. A trial type was defined by two offers and a choice [e.g., (1A:3B, A)]. A neuronal response was defined as the activity of one neuron in one time window as a function of the trial type, pooling trial types from two blocks. Building on previous studies (Padoa-Schioppa and Assad, 2006), we identified task-related responses and classified them as encoding one of the variables *offer value A*, *offer value B*, *chosen value*, or *chosen juice* (see Materials and Methods). In total, 488 neurons encoded a decision variable in at least one time window (Monkey D: 248 cells, 51.7%; Monkey F: 240 cells, 30.7%). 1917 responses were task-related and 984 of them encoded the *offer value* or the *chosen value* (Table 1). Of these, 644 value-encoding responses

Table 1. Number of responses encoding each variable for Monkeys D and F

Response	Monkey D	Monkey F
Offer value A	242 (160)	123 (75)
Offer value B	116 (78)	88 (51)
Chosen value	249 (173)	166 (107)
Chosen juice	578	355

Later analyses focused on *offer value* and *chosen value* responses. Values in parentheses indicate the number of responses that met inclusion criteria.

Figure 2. Four hypotheses for neuronal adaptation in a hypothetical *offer value A* response. Red, yellow and blue traces represent neural responses in the low, high, and wide range conditions. Dotted lines represent the absolute minimum and maximum responses observed for the neuron across all conditions. Firing rates are shown in arbitrary units. **A**, Predicted responses if neurons fully adapt to both maximum and minimum values. The slope of encoding is lower in the wide range (blue) compared with the low range (red) or the high range (yellow). The responses to the minimum and maximum values are consistent across all ranges. **B**, Predicted responses if neurons adapt to changes in maximum but not minimum value. Encoding is steeper in the low range compared with the high or the wide range. The maximum response is consistent across conditions, but the minimum observed response is higher when the minimum value is larger (high range). **C**, Predicted response if neurons do not adapt at all. The response function does not change across conditions. Neurons have the same encoding slope across all ranges. The maximum (minimum) response is higher in blocks where the maximum (minimum) value is larger. **D**, Predicted responses if neurons partially adapt to changes in both maximum value and minimum value. Encoding is steeper in the low and high ranges compared with the wide, but the full dynamic range is not always used. The maximum observed response is lower when the maximum value is smaller (low range). Similarly, the response to the minimum value is higher when the minimum value is larger (high range).

met inclusion criteria for our analysis of neuronal adaptation (see Materials and Methods).

At the outset of the study, we envisioned four possible outcomes (Fig. 2). First, responses might adapt fully to changes in both maximum and minimum values (range adaptation; Fig. 2A). In this case, the slope of encoding would be steeper in the low

and high ranges compared with the wide range. In addition, the range of firing rates would be the same in all conditions—the maximum and minimum values in each condition ($V_{\max}$ and $V_{\min}$) would always evoke the same maximum and minimum responses ($R_{\max}$ and $R_{\min}$, respectively). Alternatively, neurons might adapt to the maximum value but not to the minimum value (max adaptation; Fig. 2B). Conceptually, this scenario would occur if values were represented relative to the status quo (i.e., the animal's state before the decision). In other words, the monkey always begins the trial with 0 drops of juice and can always receive 0 reward if it chooses not to engage with the task. If neurons adapt only to the maximum value, the value of choice options may be encoded relative to this default outcome. Adaptation to the minimum value rules out this possibility. In the max adaptation scenario, the encoding slope in the high and wide ranges would be the same while the slope in the low range would be steeper. In addition, $R_{\min}$ would be elevated in the high value range, reflecting a larger $V_{\min}$. Notably, adaptation to either the value range or the maximum value would be consistent with previous results (Padoa-Schioppa, 2009; Kobayashi et al., 2010). Third, neurons might not adapt at all (Fig. 2C). Non-adapting responses would have the same tuning function in all conditions, but different values of $R_{\max}$ and $R_{\min}$ would be observed because of the different values sampled in each range. Because previous work found adaptation to changes in maximum value, we considered this outcome unlikely, but kept it as reference point for our analyses. Finally, neurons might adapt partially to $V_{\max}$, $V_{\min}$, or both (Fig. 2D). In this case, value encoding would have a steeper slope for the low and high value ranges relative to the wide range, but the range of evoked responses would also change across conditions. For example, $R_{\max}$ and $R_{\min}$ would be higher in the high range compared with the low range condition, corresponding to higher $V_{\max}$ and $V_{\min}$.

In broad terms, neurons adapt to a parameter if changing that parameter alters their tuning functions. We frequently observed adaptation in *offer value* and *chosen value* responses for all types of range transition. For example, the cell in Figure 3, A and B, adapted to changes in the maximum value of Juice B. It encoded *offer value B* in both blocks, but its tuning slope was shallower when the maximum value increased. Similarly, the cell in Figure 3, C and D, adapted to changes in the minimum value, encoding *offer value A* with a shallower slope when the minimum value decreased. The cell in Figure 3, E and F, adapted to changes in both maximum and minimum value. When the range of chosen values shifted down, the tuning curve shifted left as firing rates rescaled to the new value range. In this case, the encoding slope also increased, reflecting the narrower range of chosen values in the second block.

Across the population, neuronal responses were variable, but they consistently showed adaptation to both the maximum and minimum value (Fig. 4A–C). Notably, neuronal adaptation was not complete: the range of firing rates differed across range types, indicating that neural activity did not fully rescale to the range of values available in each trial block. This point can be seen most clearly in Figure 4D. Although each of the three range types has distinct tuning curves, $R_{\min}$ is higher in the high range condition compared with the other conditions. Similarly, $R_{\max}$ is lower in the low range condition compared with the high and wide range conditions. This result most closely resembles partial range adaptation (Fig. 2D). We obtained similar results when responses for positive and negative encoding responses were analyzed separately (Fig. 5).

Figure 3. Examples of behavior and neuronal activity for three types of range transition. **A, B,** Increase in maximum value (low $\rightarrow$ wide). **C, D,** Decrease in minimum value (high $\rightarrow$ wide). **E, F,** Decrease in both (high $\rightarrow$ low). **A, C, E,** Choice behavior. Each point represents one offer type. Quantity of offer A is shown on the x-axis and quantity of offer B on the y-axis. The color of each point indicates the fraction of trials in which the animal chose Juice B. Dashed lines show the indifference line calculated by logistic regression. **B, D, F,** Neural activity across the two blocks. Colors indicate range type (red, low range; yellow, high range; blue, wide range). Each panel shows data from a different neuron. **B,** Offer value B response, post-offer time window. **D,** Offer value A response, post-offer time window. **F,** Chosen value response, post-juice 2 time window.

Adaptation is incomplete

To examine value adaptation quantitatively, we analyzed three features of the response function: the slope of the encoding, the response to $V_{\max}$, and the response to $V_{\min}$.

We analyzed changes in the tuning slope in two ways. First, we compared the slope directly across changes in $V_{\max}$, $V_{\min}$, or both. On average, the slope was larger when the value range was high or low compared with when the range was wide, consistent with the hypothesis that neurons adapt to both maximum and minimum values (Fig. 6A–C). Responses also showed slightly higher slopes in the low range relative to the high range condition (Fig. 6C). While this observation is consistent with the idea that responses adapt more to $V_{\max}$ than to $V_{\min}$, the effect was driven by *chosen value* responses. *Offer value* responses alone did not show any difference in slope between the low range and the high range conditions. To interpret changes of slope in *chosen value* responses, we also needed to account for the difference in value range ($V_{\max} - V_{\min}$), which varied depending on the animal's choice pattern.

To further examine the relationship between slope and value range, we defined adaptation ratios (ARs) for three hypothetical scenarios: adaptation to maximum value ($AR_{\max}$), adaptation to the value range (AR_{range}), or no adaptation (AR_{none}):

$$AR_{\max} = (s_1 V_{\max,1}) / (s_2 V_{\max,2}),$$

$$AR_{\text{range}} = (s_1 \Delta V_1) / (s_2 \Delta V_2),$$

$$AR_{\text{none}} = s_1 / s_2,$$

where s is the encoding slope, ΔV is the value range ($V_{\max} - V_{\min}$), and indices 1 and 2 indicate different trial blocks. For high $\leftrightarrow$ wide or low $\leftrightarrow$ wide transitions, we defined Block 1 as the wide range (ARs are calculated as wide/narrow). For high $\leftrightarrow$ low transitions, we defined Block 1 as the high range (ARs are calculated as high/low). ARs provide a metric for the degree of adaptation. If neurons adapt completely to both maximum and minimum values, then $AR_{\text{range}} = 1$. If they adapt to the maximum only, then $AR_{\max} = 1$. Note that AR_{none} is simply the ratio of slopes in the two conditions, and should be 1 if responses do not adapt. ARs are ambiguous for certain types of range transition. For example, when only the maximum value changes, $AR_{\max}$ and AR_{range} are equivalent. In addition, ARs only test the relation between the value range and the tuning slope; they are not affected by changes in the intercept of the tuning function. Hence, $AR = 1$ does not imply that responses adapt in a specific way. However, $AR \neq 1$ indicates that a particular hypothesis does not fully describe adaptation.

Figure 4. Adaptation in offer value responses across each type of range transition. **A–C**, Individual responses (thin lines) and population mean (thick line) for (**A**) change in maximum value ($n = 72$); (**B**) change in minimum value ($n = 163$); and (**C**) change in both ($n = 129$). Insets, Average responses for transitions where $V_{\max}$ and/or $V_{\min}$ increase (top) or decrease (bottom). Shaded region in inset shows mean $\pm$ SEM. Responses are normalized to the wide range (**A**, **B**) or high range (**C**). Insets in **C** are normalized to $V_{\max}(\text{high}) - V_{\min}(\text{low})$. **D**, Overlay of mean responses for all six types of range transition. Transitions from **A** and **B** are aligned to wide range. Transitions from **A** and **C** are aligned to the low range. Responses from all analyzed time windows are shown.

Figure 5. Adaptation in offer value responses with positive and negative encoding. **A**, Positive encoding. **B**, Negative encoding. Each panel illustrates the overlay of mean responses for high $\leftrightarrow$ low, high $\leftrightarrow$ wide, and low $\leftrightarrow$ wide transitions. High $\leftrightarrow$ wide and low $\leftrightarrow$ wide transitions are aligned to the wide range. High $\leftrightarrow$ low transitions are aligned to the total range across blocks (i.e., the lowest value in the low range and the highest value in the high range).

Table 2 summarizes the ARs for every type of transition. A few points are noteworthy. First, $AR_{\text{none}} < 1$ for all range transitions, meaning that adaptation occurred consistently. Similarly, $AR_{\text{max}} \neq 1$ for transitions where $V_{\min}$ changed alone or where both $V_{\min}$ and $V_{\max}$ changed, indicating that responses adapted to changes in both maximum and minimum value. At the same time, $AR_{\text{range}} > 1$ when $V_{\min}$ changed alone and when $V_{\max}$ decreased. This finding indicates that responses did not fully adapt to changes in either $V_{\max}$ or $V_{\min}$. Overall, these results confirm that responses

adapted to both the maximum and minimum values, but that the dynamic range did not rescale completely.

So far, we have examined changes in the gain of value encoding. However, as Figure 4 illustrates, range transitions often led to a shift in $R_{\min}$ and $R_{\max}$. To quantify this effect, we compared $R_{\min}$ and $R_{\max}$ across different ranges (Fig. 6D–I). In general, when $V_{\max}$ ($V_{\min}$) was higher, $R_{\max}$ ($R_{\min}$) was also higher (all $p < 10^{-3}$, Wilcoxon signed rank test). Interestingly, $R_{\min}$ was slightly higher in the wide range compared with the low range condition, even though $V_{\min}$ was the same (Fig. 6G; $p = 0.026$, $W = 2.8 \times 10^3$, Wilcoxon signed rank test). $R_{\max}$ did not differ significantly between the wide and the high-range conditions, although there was a trend toward higher responses in the wide range (Fig. 6E; $p = 0.058$, $W = 9.0 \times 10^3$, Wilcoxon signed rank test). Importantly, although responses did not remap completely, our results were inconsistent with the hypothesis of no adaptation (Fig. 2C). To quantify this point, we computed the normalized change of $R_{\min}$ and $R_{\max}$ ($\Delta R_{\min}$ and $\Delta R_{\max}$, respectively) and compared them to the values predicted if neurons did not adapt (see Materials and Methods). $\Delta R_{\min}$ and $\Delta R_{\max}$ were consistently lower than the values predicted for non-adapting cells (Table 2). Along with the analysis of response gain,

Figure 6. Metrics of value adaptation for each type of block transition. Transition type: (**A, D, G**) maximum value changes ($n = 121$); (**B, E, H**) minimum value changes ($n = 206$); (**C, F, I**) both change ($n = 317$ responses). **A–C**, Changes in the slope of value encoding of offer value and chosen value responses. Dashed lines show $y = x$. Value encoding was generally steeper for low or high value ranges compared with the wide range (high vs wide: $p = 5.6 \times 10^{-9}$, $W = 1.6 \times 10^4$; low vs wide: $p = 7.3 \times 10^{-8}$, $W = 5.8 \times 10^3$). When both maximum and minimum values changed, the encoding slope for high and low value ranges was close to the unity line, but slightly higher for the low value range ($p = 4.8 \times 10^{-10}$, $W = 3.5 \times 10^4$). This effect was driven by chosen value responses. Offer value responses alone did not show such a difference ($p = 0.23$, $W = 2.5 \times 10^3$). **D–I**, Across-block comparisons of $R_{\max}$ (**D–F**) and $R_{\min}$ (**G–I**) for each type of range transition. $R_{\max}$ ($R_{\min}$) was generally higher when $V_{\max}$ ($V_{\min}$) was higher. **D**, $R_{\max, \text{wide}} > R_{\max, \text{low}}$ ($p = 6.7 \times 10^{-6}$, $W = 2.0 \times 10^3$); 6 points fall outside the limits of the plot. **E**, $R_{\max}$ is not significantly different between high and wide blocks ($p = 0.058$, $W = 9.0 \times 10^3$); 14 points fall outside the limits of the plot. **F**, $R_{\max}$ is higher in the high range compared with the low ($p = 5.4 \times 10^{-4}$, $W = 2.0 \times 10^4$); 16 points outside the limits of the plot. **G**, $R_{\min, \text{wide}} > R_{\min, \text{low}}$ ($p = 0.026$, $W = 2.8 \times 10^3$); 6 points fall outside the limits of the plot. **H**, $R_{\min, \text{wide}} < R_{\min, \text{high}}$ ($p = 6.2 \times 10^{-9}$, $W = 1.6 \times 10^4$); 9 points fall outside the limits of the plot. **I**, $R_{\min, \text{high}} > R_{\min, \text{low}}$ ($p = 9.2 \times 10^{-4}$, $W = 2.0 \times 10^3$); 8 points fall outside the limits of the plot. Dashed lines show $y = x$. All p values based on Wilcoxon signed rank test. Responses from all analyzed time windows are shown.

these results confirm that value-encoding neurons in OFC undergo partial adaptation to changes in the value range.

At the population level, the appearance of partial adaptation could reflect a mixture of fully adapting and non-adapting responses. To test for this possibility, we used Hartigan's dip test to check for bimodality in the distributions of ΔR_{range} , $\Delta R_{\min}$, and $\Delta R_{\max}$. We analyzed the data across changes in $V_{\max}$, $V_{\min}$, or both, and did not observe evidence of a multimodal distribution

for any transition type ($p = 0.22$ for $\Delta R_{\max}$ when both $V_{\max}$ and $V_{\min}$ changed, all other $p > 0.75$; bootstrap test on Hartigan's dip statistic). These results are consistent with the idea that partial adaptation occurs within individual responses and is a consistent feature of value encoding in OFC.

The observation of partial rescaling in value-encoding responses raised the possibility that adaptation was still ongoing during data collection. An incomplete temporal process could

Table 2. Metrics of adaptation in offer value and chosen value responses across six types of range transition

Transition type	AR_{max}	AR_{range}	AR_{none}	ΔR_{max}	ΔR_{min}
Increase max	1.04	1.04	0.78*	0.17*, ⁺	0.011
Decrease max	1.13*	1.13*	0.74*	0.22*, ⁺	0.060
Increase min	0.83*	1.23*	0.83*	0.029	−0.20*, ⁺
Decrease min	0.88*	1.39*	0.88*	0.018	−0.19*, ⁺
Increase both	1.37*	1.04	0.86*	0.22*, ⁺	0.17*, ⁺
Decrease both	1.27*	0.94	0.82*	0.17*, ⁺	0.19*, ⁺

Columns 1–3: median ARs calculated for three hypotheses: (1) neurons adapt to maximum value only, (2) neurons adapt to both maximum and minimum values, and (3) neurons do not adapt. If a hypothesis is true, $AR = 1$. Columns 4–5: median normalized difference in R_{max} and R_{min} between blocks. Nonzero values indicate a change in neural activity range between sessions (incomplete adaptation). Asterisks (*) indicate a significant deviation from 1 (columns 1–3) or from 0 (columns 4–5). Plus (+) indicates that the median ΔR_{min} or ΔR_{max} differs from the value predicted for non-adaptive coding ($\Delta R_{min,NA}$ and $\Delta R_{max,NA}$, respectively). $\Delta R_{max,NA} = 0.4$ for transitions where V_{max} changes, otherwise $\Delta R_{max,NA} = 0$; similarly, $\Delta R_{min,NA} = 0.4$ for transitions where V_{min} changes, otherwise $\Delta R_{min,NA} = 0$. All $p < 0.01$, Wilcoxon signed rank test.

produce the intermediate range adaptation observed in Figure 4. To test this prospect, we computed the tuning function separately in the first and second half of Block 2. If adaptation was temporally incomplete, responses should show greater changes in the second half of Block 2 compared with the first half. Contrary to this prediction, tuning functions for the first and second halves of Block 2 were nearly identical for all transition types (Fig. 7). Statistical analyses confirmed that changes in the slope and intercept of the tuning function were present within the first half of Block 2 (all p values < 0.01 , Wilcoxon signed rank test). Hence, neuronal adaptation occurred relatively quickly after a change in value range, and the features of range adaptation described above reflect the steady state rather than an unfinished transition.

Adaptation does not affect linearity of tuning

Previous work found that value encoding in OFC is quasilinear, but slightly convex on average (Rustichini et al., 2017). We asked whether range adaptation has any effect on this curvature. To address this question, we fit each value-encoding response separately with a quadratic polynomial and a cubic polynomial in each range condition. Confirming previous observations, few responses showed significant quadratic or cubic terms (β_2 : 10.6%, β_3 : 4.9%; $p < 0.05$, F test). On average across the population, quadratic terms were slightly positive ($p = 5.8 \times 10^{-56}$, $W = 2.5 \times 10^6$, Wilcoxon signed rank test), while cubic terms were slightly negative ($p = 1.6 \times 10^{-3}$, $W = 1.7 \times 10^6$, Wilcoxon signed rank test). Most importantly, the distribution of β_2 did not differ between high and low value ranges (Fig. 8A; median values: 0.064, 0.61; $p = 0.47$, $U = 5.6 \times 10^5$, rank sum test). Values of β_2 were slightly lower in the wide range (median: 0.017; wide vs high: $p = 9.6 \times 10^{-9}$, $U = 7.6 \times 10^5$; wide vs low: $p = 1.1 \times 10^{-10}$, $U = 8.3 \times 10^5$, rank sum test). However, this difference arose from the fact that the wide range included a greater number of distinct values, which constrained the polynomial fits. Indeed, when we recalculated the quadratic fits for the wide range using only the subset of values present in the low range condition, the distribution of β_2 did not differ from the distribution measured with high and low ranges (median $\beta_{2,subsampled} = 0.045$; both $p > 0.1$, rank sum test). Similarly, the distribution of β_3 did not differ across high, low, and wide range conditions (Fig. 8B; median values: −0.014, 1.8×10^{-3} , and -3.8×10^{-3} ; all p values > 0.1 , rank sum test).

The same pattern of results emerged when we compared β_2 and β_3 for each response across blocks (Fig. 8C–F). Although values of β_2 varied substantially, coefficients for each response were correlated across blocks. This correlation suggests that β_2 is

a characteristic of each neuron's tuning function. As in the previous analysis, β_2 was slightly higher in narrow ranges compared with the wide range (Fig. 8E), although this was only significant for changes in V_{max} (median difference = 0.031, $p = 1.4 \times 10^{-4}$, $W = 9.7 \times 10^3$, Wilcoxon signed rank test). The effect disappeared when β_2 for the wide range was calculated with subsampled values ($p = 0.39$, $W = 6.7 \times 10^3$). Values of β_3 did not differ across any type of range transition (all p values > 0.1 , Wilcoxon signed rank test) and did not show any consistent pattern of correlation across blocks.

In summary, adaptation altered the gain and offset of value-encoding responses, but not their quasilinear functional form.

Absence of range adaptation in chosen juice cells

All the results presented so far focused on responses encoding the offer value or the chosen value. In a separate set of analyses, we examined responses encoding the chosen juice.

We did not find any evidence of range adaptation in this population. More specifically, we did not find systematic differences in the encoding slopes (difference in responses to preferred and non-preferred juice) or in the minimum response (i.e., the response to the non-preferred juice), across any range transition (Fig. 9; all p values > 0.05 , Wilcoxon signed rank test). Thus, it appears that chosen juice responses, capturing the binary choice outcome, are not affected by changes in the value range.

Increases in minimum response impair simulated choice behavior

We have shown that value-encoding neurons do not rescale completely to changes in value range. In other words, responses do not span the full range of potential firing rates in every condition. One important question is whether and how partial adaptation in offer value cells affects economic decisions. This issue is closely related to that of optimality in the neuronal representation of subjective values.

In sensory systems, neuronal response functions are generally considered optimal if they transmit maximal information about the stimuli (Barlow, 1961; Laughlin, 1981). In the neural system underlying economic decisions, this concept of optimality seems less relevant. Instead, optimal tuning may be defined as the response function that maximizes the expected payoff (Rustichini et al., 2017). In our choice task, the payoff is simply the value chosen by the monkey on any given trial. Notably, although the relative value of two juices is subjective, the payoff of two options may be compared objectively once the relative value of the juices is known. For example, if the choice pattern indicates that $\rho = 2.5$, then the payoff of 3B is higher than the payoff of 1A. Importantly, for given offers, the expected payoff is inversely related to choice variability. When choice variability is higher, i.e., when decisions between two options are more frequently split, the animal is more likely to choose the lower value (lower expected payoff). In previous computational work, we found that a linear decision model achieved the maximum expected payoff if offer value cells adapted completely to the value range; i.e., if their dynamic range rescaled fully to the current range of values (Rustichini et al., 2017). However, that study only considered changes in the slope of the encoding, and it was limited to instances where the minimum offer value was zero. Moreover, the study assumed that each neuron's minimum response was also zero, and it considered only positive encoding of value. Contrary to these assumptions, we found here that value-encoding responses adapt to the minimum as well as the maximum value. Furthermore, R_{min} changes with the value range in a systematic way, generally

Figure 7. Tuning of offer value responses in the first and second half of each block. Transition types: (A–C) increase in maximum (A; $n = 38$), minimum (B; $n = 64$), or both (C; $n = 76$). D–F, Decrease in maximum (D; $n = 33$), minimum (E; $n = 99$), or both (F; $n = 53$). Shaded regions indicate SEM. The first half of each block is shown in lighter colors, the second half in darker colors. Tuning functions are consistent across the early and late halves of the block.

increasing with $V_{\min}$. Finally, we know that $\sim 30\%$ of offer value cells in OFC encode value with a negative slope (lowest activity for largest value). When the value range shifts upward, the response of negative encoding neurons tends to decrease, while the response of positive encoding cells tends to increase. These opposing changes could counteract each other in complex ways.

To explore the behavioral implications of these factors, we ran a series of computer simulations. We examined a linear decision model comprised of 5000 offer value A and 5000 offer value B units (see Materials and Methods). We analyzed two networks: a positive network including only cells with positive encoding and a mixed network including 30% of units with negative encoding. In both cases, each offer value unit encoded the corresponding offer value in a linear way. Trial-to-trial variability was correlated across units, with correlation values estimated based on empirical measures (Conen and Padoa-Schioppa, 2015). We simulated choices between pairs of offer values, which were randomly selected on each trial. The decision was determined based on the activity of the two pools: on trials where the activity of offer A units exceeded that of offer B units, Juice A was chosen (and vice versa).

First, we examined how changes of $R_{\min}$ affected choices in the positive network. We specifically considered two scenarios: (1) each unit had a fixed $R_{\max}$, such that increasing $R_{\min}$ reduced the available dynamic range (Fig. 10A); (2) each unit had a fixed activity range ($\Delta r = R_{\max} - R_{\min}$), such that increasing $R_{\min}$ shifted the dynamic range (Fig. 10B). In essence, the first scenario captures the case where neurons do not adapt to changes in the minimum value; the second scenario is analogous to the partial range adaptation observed in the experiments, where both $R_{\min}$

and $R_{\max}$ are elevated when the value range shifts up (Fig. 4C). For each scenario, we simulated choices for increasing levels of $R_{\min}$. Furthermore, we quantified the effectiveness of choice behavior using the FLV (see Materials and Methods).

Figure 10, C and D, illustrates our results. The payoff decreased with increasing values of $R_{\min}$ in both scenarios. However, increased $R_{\min}$ was much more costly when $R_{\max}$ was fixed (Fig. 10C). In this scenario, FLV increased to 1 as $R_{\min} \rightarrow R_{\max}$, meaning that decisions fell to chance levels as the dynamic range decreased. In contrast, increased $R_{\min}$ had a relatively mild effect when $R_{\max}$ and $R_{\min}$ increased together (Fig. 10D). In this condition, FLV remained < 0.25 even for $R_{\min}$ equal to or exceeding the total response range. In both conditions, changing $R_{\min}$ produced a similar effect across a range of Fano factors (F), although FLV was slightly higher when units had higher variance. Doubling the Fano factor had substantially less effect than increasing $R_{\min} \rightarrow R_{\max}$ in Scenario 1, and approximately half the effect of increasing $R_{\min} \rightarrow 1$ in Scenario 2. In summary, increasing the minimum payoff in a network moderately decreases the expected payoff, though far less than reducing the dynamic range of responses.

Second, we examined a mixed network comprised of 70% positive encoding units and 30% negative encoding units (see Materials and Methods). As in the previous simulations, we considered two scenarios: (1) each unit had a fixed maximum response ($R_{\max}$ for positive encoding units, $R_{\min}$ for negative encoding units), and (2) each unit had a fixed activity range ($\Delta r = |R_{\max} - R_{\min}|$). To account for the fact that positive and negative encoding cells change in opposing ways as the value range shifts, we set $R_{\max}[\text{negative units}] = 1 - R_{\min}[\text{positive units}]$.

Figure 8. Quadratic and cubic tuning parameters. **A, B**, Distribution of quadratic (**A**) and cubic (**B**) coefficients for all value encoding responses in wide, low, and high ranges. **C–F**, Quadratic and cubic coefficients for individual responses across blocks. Each point represents one response. Dotted lines show $y = x$. Pearson correlation (r) and p values across each transition type: (**C**) $r = 0.11, p = 1.6 \times 10^{-3}$; (**D**) $r = 0.048, p = 0.33$; (**E**) $r = 0.39, p = 1.4 \times 10^{-4}$ (change in $V_{\max}$); $r = 0.35, p = 1.1 \times 10^{-8}$ (change in $V_{\min}$); (**F**) $r = -0.12, p = 0.12$ (change in $V_{\max}$); $r = 0.24, p = 1.5 \times 10^{-4}$ (change in $V_{\min}$).

Thus, the minimum activity decreased in negative units, whereas it increased in positive units, emulating a scenario where the value range shifted higher.

Figure 10, *E* and *F*, illustrates our results. Most interestingly, the relationship between FLV and $R_{\min}[\text{positive}]$ was non-monotonic. For both scenarios, FLV initially decreased, and reached a minimum for $R_{\min}[\text{positive}] \approx 0.4$. In Scenario 1, where the maximum response was constant, FLV then rapidly increased (Fig. 10*E*). Nevertheless, in contrast to what we observed for the positive network (Fig. 10*C*), FLV did not reach chance level (FLV ≈ 0.9 when $R_{\min}[\text{positive}] = 1$). This is because the dynamic range of negative units remained >0 . In Scenario 2, where the range of responses remained constant, FLV also increased as $R_{\min}[\text{positive}]$ increased from 0.4 to 1. However, the increase in FLV was modest, such that FLV for $R_{\min}[\text{positive}] = 1$ was only slightly higher than FLV for $R_{\min}[\text{positive}] = 0$ (Fig. 10*F*).

We interpret these results with a few caveats. Most importantly, these simulations can only provide a general idea of how FLV changes with $R_{\min}$. The specific relationship between FLV and $R_{\min}$ depends on the details of the decision network, including the firing rates of positive and negative encoding neurons and their relative weight in the decision process. Furthermore, the mixed network makes FLV more comparable in high and low value ranges, but it does not increase the efficacy of the network overall. Instead, FLV in the mixed network is elevated when $R_{\min}[\text{positive}]$ is low (i.e., decisions are noisier in low value ranges). As the U-shaped function indicates, FLV is minimized for some $R_{\min} > 0$. In other words, the model produces the best choice performance in intermediate value ranges.

In these simulations, we focused on changes in the minimum response with and without changes in dynamic range. However, the same principles apply to changes in the maximum response. Previous theoretical work suggests that incomplete adaptation to the maximum value would degrade choices (i.e., increase FLV) because of the reduction in dynamic range (Rustichini et al., 2017). As our simulations show, decreases in the dynamic range because of increases in $R_{\min}$ produce similar detriments (Fig. 10*C*). If both $R_{\min}$ and $R_{\max}$ change together, higher firing rates generally lead to larger FLV (Fig. 10*D*), but this effect is reduced by opposing changes in positive and negative encoding units (Fig. 10*F*). Partial adaptation to $V_{\min}$ and $V_{\max}$ leads to suboptimal choices, but the mixture of negative and positive encoding stabilizes performance across value ranges.

Discussion

The present study addresses outstanding questions concerning the neuronal representation of subjective values. Specifically, we showed that neurons in OFC adapt to the value range rather than to the maximum value. In other words, values are not encoded relative to the subject's pre-decision state. Instead, values are represented in terms of the best and worst outcomes possible in the current behavioral context. At the same time, we found that range adaptation is partial. Although the encoding gain was consistently higher when the value range was narrow (high or low), neuronal responses did not rescale completely to the value distribution, and the range of firing rates was lower when the value range was lower. Thus, value encoding fell in an intermediate zone between fully adaptive coding (range adaptation) and absolute value coding (no adaptation). Importantly, these observations did not reflect an unfinished process of adaptation. Through a series of simulations, we also showed that

Figure 9. Chosen juice responses to not adapt changes in range. **A–C**, Slope of chosen juice encoding. **D–F**, Y-intercept (response to the non-preferred juice) for chosen juice responses. Ranges were defined as the range of the preferred juice (i.e., ranges of juice A were used for a chosen juice cell that fired more for choice A). Defining ranges by chosen value or total value range did not alter results. Dashed lines show $y = x$. All p values based on Wilcoxon signed rank test.

increases in the minimum response negatively affects choices, even when the gain of encoding is constant. However, this effect is reduced by the presence of both positive and negative offer value responses.

Our experimental results resonate with previous observations. Kobayashi et al. (2010) measured range-dependent changes in value-encoding neurons in several subregions of OFC. Their analysis focused on changes in gain. Although they divided neurons into adapting, non-adapting, and partially adapting groups, their results are also consistent with a single population of partially adapting cells. In this study, we looked for but did not find any evidence of bimodality in any metric of adaptation, and adapted tuning curves generally clustered around intermediate values, in line with the idea that partial adaptation is present throughout the population. Along similar lines, in human subjects, Burke et al. (2016) found partial adaptation in the BOLD signal in vmPFC using a decoding approach. Together, these findings suggest that partial adaptation may be a common characteristic of value coding in prefrontal cortex.

This study is the first to examine the effect of adaptation on the offset of encoding (i.e., changes in R_{min} and R_{max}). We found that the dynamic range of value-encoding responses shifts up and down depending on the current condition. For the most part, changes in R_{min} and R_{max} seem related to the fact that adaptation is partial; the adapted tuning functions fall between predictions for fully adapting neurons and a non-adapting ones. However, we also observed slightly lower R_{min} in the low range condition compared with the wide range. Because the minimum value is the same in these conditions, this result does not appear to be a direct byproduct of partial adaptation. Although the reason for this outcome is unclear, one possibility is that value adaptation in-

volves two semi-independent components; a change in gain that depends on the range of values, and a change in offset that depends on the average available value or a related factor. Future work will explore this hypothesis.

Elevated neural activity reduces the expected payoff

In a previous study, a simulated decision network yielded the highest payoff when neurons exploited their full dynamic range (Rustichini et al., 2017). Here, we found that responses do not span their entire dynamic range in all conditions. Response functions shift up or down depending on the value range, which can be measured as a change in R_{min} . We found that increasing R_{min} reduces the expected payoff in a simulated decision network. Intriguingly, the effect of R_{min} on payoff is lower when the simulated population included offer value responses with both positive and negative encoding. Indeed, in the mixed network, the payoff is highest (lowest FLV) for intermediate values of R_{min} . However, although the presence of negative encoding makes FLV more stable, it does so by spreading the costs of partial adaptation more evenly across different value ranges. In the mixed network, the minimum responses of neurons are elevated for both low and high value ranges, leading to higher FLV. Although this activity increase is less costly than reducing the dynamic range of responses, it still results in slightly lower payoff overall. Intuitively, this inefficiency arises from the fact that the variance of neural responses scales with the mean. Other things equal, when a neuron's dynamic range is higher, firing rates are noisier.

One caveat of our results is that partial adaptation may have reflected the overall design of our experiments. Indeed, our monkeys were highly trained on the range adaptation task, and they were familiar with all possible transitions between high, low, and

Figure 10. Choice simulation. FLV increases with increasing R_{min} . **A, B,** Illustration of example response functions. **A,** R_{min} increases while R_{max} remains fixed. Analogous to a scenario where units adapt only to the maximum value and the value range shifts higher. **B,** R_{min} increases while the value range ($R_{max} - R_{min}$) remains fixed. Analogous to the activity offset associated with partial range adaptation in OFC responses. **C, D,** Simulation results for networks with 100% positive encoding units. **C,** For neurons with a fixed maximum value, FLV increases to 1 (chance level) as R_{min} increases. **D,** For neurons with a fixed range, FLV increases mildly as R_{min} increases. Trace colors indicate results simulated for different Fano factors. Each curve covers 100 values of R_{min} and shows the mean of 20 simulated sessions for each value of R_{min} . **E, F,** Simulation results for networks with 30% negative encoding units. The x-axes represent the values of R_{min} for units with positive encoding. The minimum response for negative encoding units is evoked by the maximum value (R_{max}) in each context and changes in the opposite direction: $R_{max}[\text{negative}] = 1 - R_{min}[\text{positive}]$. **E,** For units with a fixed maximum response, FLV decreases slightly as $R_{min}[\text{positive}]$ increases to ~0.4, then increases rapidly to ~0.9. **F,** For units with a fixed activity range, FLV follows a shallow U-shaped function as $R_{min}[\text{positive}] \rightarrow 1$.

wide ranges. Although complete adaptation would warrant an efficient representation of values within a block, it would also limit the circuit's ability to respond when the value range changes. In contrast, intermediate adaptation reserves a portion of the dynamic range for new values that may appear after a transition. This interpretation suggests that value encoding depends on at least two components; a slow, learning-based process that draws on contextual knowledge, and a more rapid adaptive component that adjusts to the locally experienced value range. Further work is needed to test whether the degree of value adaptation varies across different experimental paradigms.

Partial adaptation may also allow the circuit to maintain information about the overall value of the current context (i.e., the value of the block). Information about the current contextual value makes it possible to predict future reward expectations and affects subjects' motivation to engage in the task. Moreover, effective value comparison in an adapting network requires information about the distribution of available values as well as neural activity levels on a given trial. Without some mechanism for maintaining this information, signals are ambiguous across contexts and cannot guide behavior effectively (Fairhall et al., 2001; Rustichini et al., 2017). The differences in response offset observed in OFC may be used by the network to help distinguish the current value state.

Possible mechanisms of value adaptation

Although our study did not investigate the physiological mechanism of adaptation directly, a few possibilities may be considered. We showed that value adaptation involves both an additive and a multiplicative component. While adaptation to the maximum can occur via a simple change in gain, adaptation to the minimum requires both a change in gain and a horizontal shift in the response function. When the difference between maximum and minimum values is constant, adaptation is purely horizontal: the slope of neuronal encoding remains the same, but responses re-map to a new set of values. Additive changes in activity often arise from changes in hyper-polarization or shunting inhibition (Holt and Koch, 1997; Chance et al., 2002). Alternate explanations, such as cell-intrinsic changes in membrane conductivity, generally involve a mixture of additive and multiplicative effects, which is difficult to reconcile with the purely additive adaptation we observed during high-to-low range transitions (Sanchez-Vives et al., 2000a,b). The multiplicative component of value adaptation could arise from several potential mechanisms. Changes in gain can be produced by both cell-intrinsic mechanisms, such as changes in ionic conductance (Higgs et al., 2006; Díaz-Quesada and Maravall, 2008; Mease et al., 2013), and by circuit-level changes in inhibitory activity (Olsen et al., 2012; Wilson et al., 2012; Natan et al., 2017) or the background level of synaptic activity (Chance et al., 2002).

Recent work examining a more medial region of OFC found that adaptation to simultaneously presented values was best explained by a divisive normalization model (Yamada et al., 2018). The data from our study, which reflect a slower form of adaptation across trials, do not appear to follow a similar model. Among other features, the divisive normalization model predicts a decrease in the maximum response in conditions with a higher value range, which we do not observe. Notably, that experiment focused on adaptation on a very short time scale (~ 100 ms). Another recent model combined slow and fast normalization dynamics to explain variability in choice behavior across contexts (Zimmerman et al., 2018). One interesting question is whether this model can also account for the neuronal responses recorded

in OFC. Divisive normalization is a common form of adaptation in sensory regions (Valerio and Navarro, 2003; Wark et al., 2007; Olsen et al., 2010; Beck et al., 2011; Ohshiro et al., 2011), and it is highly effective at maximizing the transmission of sensory information across a wide variety of stimuli (Schwartz and Simoncelli, 2001; Carandini and Heeger, 2011). At the same time, divisive normalization seems less well suited for contextual adaptation in a decision circuit, which ideally would optimize the choice outcome rather than transmitting maximal information about the value distribution (Rustichini et al., 2017). Nevertheless, the possible reconciliation of divisive normalization and range adaptation remains an open question.

Range-dependent changes in choice behavior

Our behavioral analyses revealed range-dependent changes in both the relative value and the sigmoid steepness (Fig. 1). The increased relative value in high-value blocks could be explained if the value of additional juice decreases at higher quantities (diminishing marginal utility). Because A is generally offered in lower quantity, such nonlinearity would presumably shift preferences toward A when the offer quantities increased. However, this explanation relies on the assumption that the marginal utility of depends on the quantity of each juice, rather than the subjective value of an offer. More detailed studies of choice behavior are needed to test this assumption.

The changes in sigmoid steepness across conditions were somewhat more surprising. A recent analysis of behavior across different ranges found that decision patterns were generally noisier (lower η) during blocks with a wider value range, consistent with the idea neurons that encoded value with lower resolution during these blocks (Rustichini et al., 2017). In contrast, we observed steeper choice functions (higher η) in the wide value range compared with the low range. The reason for this discrepancy is unclear. One possibility is that the effect is more complicated than “wide range” versus “narrow range”. Simulations of mixed networks showed that choice variability (and hence FLV) changes in a non-monotonic way as the value range shifts higher. Further modeling and experimental work is needed to explore the interactions between positive and negative encoding units and their effect on decisions.

We also observed steeper choice patterns in the high range compared with the low range. While somewhat unexpected, this result parallels the behavior of the mixed network simulation for low to medium values of $R_{\min}$ [positive]. If this interpretation is accurate, further increases in value should eventually lead to higher choice variability (lower η). Alternatively, differences in steepness across conditions may arise from other behavioral factors such as motivational state. Consistent with this possibility, we observed the steepest choice pattern in the high range, when the reward rate was the highest and monkeys were most motivated. Choices were slightly more variable (i.e., less steep) in the wide range, and most variable in the low range. Future experiments may distinguish between these two explanations by considering a wider set of possible value ranges and by balancing the reward rate across blocks.

To conclude, we examined how the neuronal representation in OFC adapted to changes in maximum and minimum of the value distribution. We found that both maximum and minimum values influence the gain of value encoding, but only partially, leading to an offset in neuronal activity levels across ranges. Modeling work suggests that the relationship between the neuronal adaptation in the representation of value and choice behavior depends on the interplay of positive and negative encoding in the

neuronal population. Future work should investigate this relationship in greater detail and thus shed light on both the flexibility and the limitations of value coding across behavioral contexts.

References

- Adibi M, McDonald JS, Clifford CW, Arabzadeh E (2013) Adaptation improves neural coding efficiency despite increasing correlations in variability. *J Neurosci* 33:2108–2120.
- Barlow HB (1961) Possible principles underlying the transformations of sensory messages. In: *Sensory communication* (Rosenblith WA, ed), pp 217–234. Cambridge, MA: MIT.
- Beck JM, Latham PE, Pouget A (2011) Marginalization in neural circuits with divisive normalization. *J Neurosci* 31:15310–15319.
- Benucci A, Saleem AB, Carandini M (2013) Adaptation maintains population homeostasis in primary visual cortex. *Nat Neurosci* 16:724–729.
- Bermudez MA, Schultz W (2010) Reward magnitude coding in primate amygdala neurons. *J Neurophysiol* 104:3424–3432.
- Burke CJ, Baddeley M, Tobler PN, Schultz W (2016) Partial adaptation of obtained and observed value signals preserves information about gains and losses. *J Neurosci* 36:10016–10025.
- Cai X, Padoa-Schioppa C (2014) Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron* 81:1140–1151.
- Carandini M, Heeger DJ (2011) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51–62.
- Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. *Neuron* 35:773–782.
- Conen KE, Padoa-Schioppa C (2015) Neuronal variability in orbitofrontal cortex during economic decisions. *J Neurophysiol* 114:1367–1381.
- Cox KM, Kable JW (2014) BOLD subjective value signals exhibit robust range adaptation. *J Neurosci* 34:16533–16543.
- Dan Y, Atick JJ, Reid RC (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J Neurosci* 16:3351–3362.
- Díaz-Quesada M, Maravall M (2008) Intrinsic mechanisms for adaptive gain rescaling in barrel cortex. *J Neurosci* 28:696–710.
- Elliott R, Agnew Z, Deakin JF (2008) Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. *Eur J Neurosci* 27:2213–2218.
- Fairhall AL, Lewen GD, Bialek W, De Ruyter van Steveninck RR (2001) Efficiency and ambiguity in an adaptive neural code. *Nature* 412:787–792.
- Fellows LK (2011) Orbitofrontal contributions to value-based decision making: evidence from humans with frontal lobe damage. *Ann N Y Acad Sci* 1239:51–58.
- Gutnisky DA, Dragoi V (2008) Adaptive coding of visual information in neural populations. *Nature* 452:220–224.
- Hardin J, Garcia SR, Golan D (2013) A method for generating realistic correlation matrices. *Ann Appl Stat* 7:1733–1762.
- Hengen KB, Lambo ME, Van Hooser SD, Katz DB, Turrigiano GG (2013) Firing rate homeostasis in visual cortex of freely behaving rodents. *Neuron* 80:335–342.
- Higgs MH, Slee SJ, Spain WJ (2006) Diversity of gain modulation by noise in neocortical neurons: regulation by the slow afterhyperpolarization conductance. *J Neurosci* 26:8787–8799.
- Holt GR, Koch C (1997) Shunting inhibition does not have a divisive effect on firing rates. *Neural Comput* 9:1001–1013.
- Kobayashi S, Pinto de Carvalho O, Schultz W (2010) Adaptation of reward sensitivity in orbitofrontal neurons. *J Neurosci* 30:534–544.
- Krekelberg B, van Wezel RJ, Albright TD (2006) Adaptation in macaque MT reduces perceived speed and improves speed discrimination. *J Neurophysiol* 95:255–270.
- Laughlin S (1981) A simple coding procedure enhances a neuron's information capacity. *Z Naturforsch C* 36:910–912.
- Lewicki MS (2002) Efficient coding of natural sounds. *Nat Neurosci* 5:356–363.
- Liu B, Macellario MV, Osborne LC (2016) Efficient sensory cortical coding optimizes pursuit eye movements. *Nat Commun* 7:12759.
- Mease RA, Famulare M, Gjorgjieva J, Moody WJ, Fairhall AL (2013) Emergence of adaptive computation by single neurons in the developing cortex. *J Neurosci* 33:12154–12170.
- Natan RG, Rao W, Geffen MN (2017) Cortical interneurons differentially shape frequency tuning following adaptation. *Cell Rep* 21:878–890.
- Ohshiro T, Angelaki DE, DeAngelis GC (2011) A normalization model of multisensory integration. *Nat Neurosci* 14:775–782.
- Olsen SR, Bhandawat V, Wilson RI (2010) Divisive normalization in olfactory population codes. *Neuron* 66:287–299.
- Olsen SR, Bortone DS, Adesnik H, Scanziani M (2012) Gain control by layer six in cortical circuits of vision. *Nature* 483:47–52.
- Ongür D, Price JL (2000) The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cereb Cortex* 10:206–219.
- Padoa-Schioppa C (2009) Range-adapting representation of economic value in the orbitofrontal cortex. *J Neurosci* 29:14004–14014.
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–226.
- Padoa-Schioppa C, Conen KE (2017) Orbitofrontal cortex: a neural circuit for economic decisions. *Neuron* 96:736–754.
- Rudebeck PH, Murray EA (2014) The orbitofrontal oracle: cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. *Neuron* 84:1143–1156.
- Rustichini A, Conen KE, Cai X, Padoa-Schioppa C (2017) Optimal coding and neuronal adaptation in economic decisions. *Nat Commun* 8:1208.
- Saez RA, Saez A, Paton JJ, Lau B, Salzman CD (2017) Distinct roles for the amygdala and orbitofrontal cortex in representing the relative amount of expected reward. *Neuron* 95:70–77.e3.
- Sanchez-Vives MV, Nowak LG, McCormick DA (2000a) Cellular mechanisms of long-lasting adaptation in visual cortical neurons *in vitro*. *J Neurosci* 20:4286–4299.
- Sanchez-Vives MV, Nowak LG, McCormick DA (2000b) Membrane mechanisms underlying contrast adaptation in cat area 17 *in vivo*. *J Neurosci* 20:4267–4285.
- Schultz W (2015) Neuronal reward and decision signals: from theories to data. *Physiol Rev* 95:853–951.
- Schwartz O, Simoncelli EP (2001) Natural sound statistics and divisive normalization in the auditory system. *Adv Neural Inf Process Syst* 13:27–30.
- Soltani A, De Martino B, Camerer C (2012) A range-normalization model of context-dependent choice: a new model and evidence. *PLoS Comput Biol* 8:e1002607.
- Valerio R, Navarro R (2003) Optimal coding through divisive normalization models of V1 neurons. *Network* 14:579–593.
- Wallis JD (2012) Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci* 15:13–19.
- Wark B, Lundstrom BN, Fairhall A (2007) Sensory adaptation. *Curr Opin Neurobiol* 17:423–429.
- Wilson NR, Runyan CA, Wang FL, Sur M (2012) Division and subtraction by distinct cortical inhibitory networks *in vivo*. *Nature* 488:343–348.
- Yamada H, Louie K, Tymula A, Glimcher PW (2018) Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nat Commun* 9:162.
- Zimmerman J, Glimcher PW, Louie K (2018) Multiple timescales of normalized value coding underlie adaptive choice behavior. *Nat Commun* 9:3206.