

ARTICLE

DOI: 10.1038/s41467-017-01373-y

OPEN

Optimal coding and neuronal adaptation in economic decisions

Aldo Rustichini¹, Katherine E. Conen², Xinying Cai^{2,5} & Camillo Padoa-Schioppa ^{2,3,4}

During economic decisions, offer value cells in orbitofrontal cortex (OFC) encode the values of offered goods. Furthermore, their tuning functions adapt to the range of values available in any given context. A fundamental and open question is whether range adaptation is behaviorally advantageous. Here we present a theory of optimal coding for economic decisions. We propose that the representation of offer values is optimal if it ensures maximal expected payoff. In this framework, we examine offer value cells in non-human primates. We show that their responses are quasi-linear even when optimal tuning functions are highly non-linear. Most importantly, we demonstrate that for linear tuning functions range adaptation maximizes the expected payoff. Thus value coding in OFC is functionally rigid (linear tuning) but parametrically plastic (range adaptation with optimal gain). Importantly, the benefit of range adaptation outweighs the cost of functional rigidity. While generally suboptimal, linear tuning may facilitate transitive choices.

¹Department of Economics, University of Minnesota, 1925 4th Street South 4-101, Minneapolis, MN 55455, USA. ²Department of Neuroscience, Washington University in St Louis, 660 South Euclid Avenue, St Louis, MO 63110, USA. ³Department of Economics, Washington University in St Louis, St Louis, MO 63130, USA. ⁴Department of Biomedical Engineering, Washington University in St Louis, St Louis, MO 63130, USA. ⁵Present address: NYU Shanghai, 1555 Century Ave, Room 1251, Pudong New District, Shanghai 200122, China. Correspondence and requests for materials should be addressed to C.P.-S. (email: camillo@wustl.edu)

Choosing between two goods entails computing and comparing their subjective values. Evidence from lesions and neurophysiology indicates that these mental operations engage the orbitofrontal cortex (OFC)^{1–3}. Experiments in which rhesus monkeys chose between different juices identified three groups of neurons in this area. Offer value cells encode the values of individual goods and are thought to provide the primary input to the decision. Conversely, chosen juice cells and chosen value cells represent the binary choice outcome and the value of the chosen good^{4, 5}. The present study focuses on offer value cells.

Previous work indicated that these neurons undergo range adaptation. In any behavioral context, their firing rate is a linear function of the offered values; their tuning slope is inversely proportional to the range of values available in that context^{6–8}. Prima facie, range adaptation seems to ensure an efficient neuronal representation. However, it was shown that uncorrected adaptation in offer value cells would result in arbitrary choice biases⁹—a problem conceptually analogous to the “coding

catastrophe” discussed for sensory systems^{10–12}. Experimental evidence presented in this study indicates that, in fact, changing the range of offer values does not affect economic preferences. In other words, range adaptation is corrected within the decision circuit to avoid choice biases. This observation raises a fundamental question: If neuronal adaptation is indeed corrected within the decision circuit, is neuronal adaptation at all beneficial to the organism? Addressing this question requires a theory of optimal coding.

Following the seminal work of Barlow¹³, optimal coding has been a frequent area of research in sensory systems. A cornerstone concept is that sensory neurons are optimally tuned for perception if they transmit maximal information about the stimuli^{13–15}. In any behavioral context, such optimality is achieved if tuning curves match the cumulative distribution function of the stimuli¹⁴. Importantly, neurons in many sensory regions adapt optimally to the current behavioral context^{16–22}, while tuning functions in other sensory regions seem optimized for the

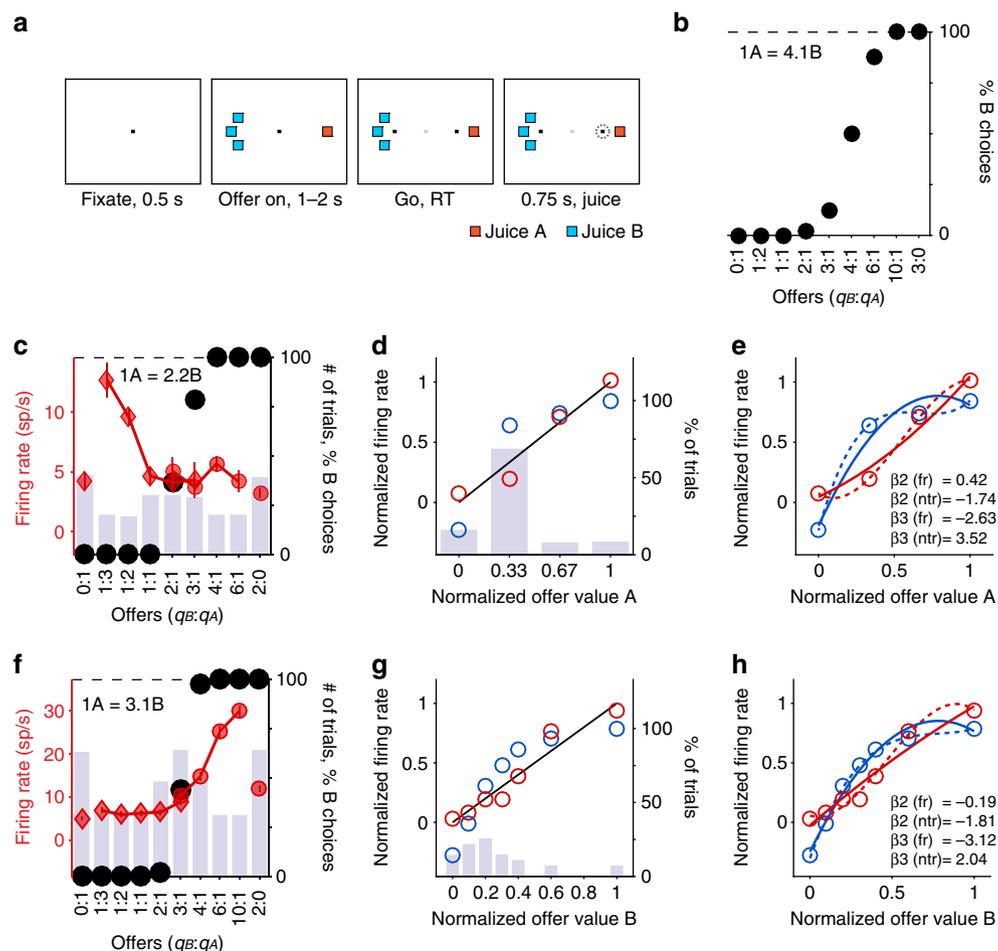


Fig. 1 Quasi-linear coding of offer values, individual responses. **a**. Task design (see Methods). **b**. Example of choice pattern. The x-axis represents different offer types, ranked by the ratio $q_B:q_A$. Black dots represent the percent of “choice B” trials. **c**. Example offer value A response. Black dots represent the choice pattern. The histogram illustrates the number of trials presented for each offer type. Red symbols represent firing rates \pm SEM (diamonds and squares for “choice A” and “choice B”, respectively). The y-axis on the left refers to firing rates. The y-axis on the right refers both to the number of trials (histogram) and to the choice pattern (black symbols). **d**. Comparing firing rates and $ntrials_{CDF}$. Same response as in **c**. The x-axis represents normalized quantity levels of juice A. The histogram illustrates the percent of trials for each quantity level. This session included 247 trials, and juice A was offered at quantity levels 0 (39 trials, 16%), 1 (169 trials, 68%), 2 (19 trials, 8%), and 3 (20 trials, 8%). Note that low quantity levels were over-represented. Blue circles represent the cumulative distribution function for the number of trials ($ntrials_{CDF}$). The y-axis on the right refers both to the number of trials (histogram) and to $ntrials_{CDF}$ (blue circles). Red circles represent firing rates. Here each neuronal data point is an average across all the trials with given quantity level (not across a single trial type). The y-axis on the left refers to normalized firing rates. Limits on the y-axes were set such that the same line (black) represents the best linear fit for firing rates and for $ntrials_{CDF}$. (Because all measures are normalized, this is the identity line.) **e**. Curvature of firing rates and $ntrials_{CDF}$. Same data points as in **d**. Continuous and dotted lines are the result of the quadratic and cubic fit, respectively. **f–h**. Example offer value B response

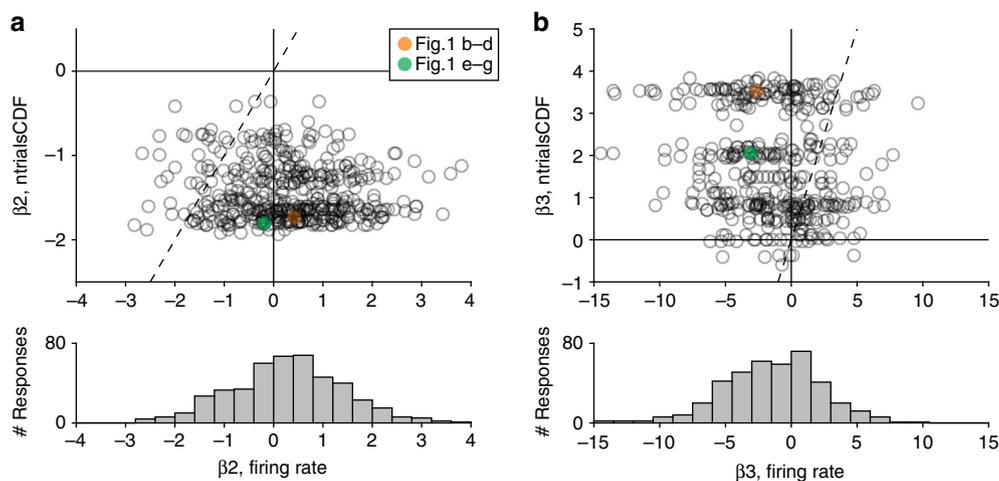


Fig. 2 Quasi-linear coding of offer values, population analysis ($N=447$). **a** Quadratic term. Each data point in the scatter plot represents one response. The x-axis and y-axis represent $\beta_2, firing\ rate$ and $\beta_2, ntrialsCDF$, respectively. The dotted line represents the identity line, and the responses illustrated in Fig. 1 are highlighted. The histogram recapitulates the distribution for $\beta_2, firing\ rate$. Since low offer values were always over-represented in the experiments, generally $\beta_2, ntrialsCDF < 0$. In contrast, measures for $\beta_2, firing\ rate$ were broadly scattered above and below zero (see histogram). **b** Cubic term. Same conventions as in **a**. Generally, $\beta_3, ntrialsCDF > 0$. In contrast, measures for $\beta_3, firing\ rate$ were broadly scattered above and below zero. Notably, on average across the population, measures for $\beta_2, firing\ rate$ were close to, but significantly above zero ($mean(\beta_2, firing\ rate) = 0.28; p < 10^{-6}, t\text{-test}$). Conversely, measures for $\beta_3, firing\ rate$ were close to, but significantly below zero ($mean(\beta_3, firing\ rate) = -1.42; p < 10^{-12}, t\text{-test}$). In both cases, the deviance from zero measured for $\beta_{\bullet}, firing\ rate$ was in the direction opposite to that observed for the corresponding $\beta_{\bullet}, ntrialsCDF$, indicating that tuning functions did not match $ntrialsCDF$ averaged across sessions. We return to these deviances later in the article

distribution of natural stimuli^{16, 23, 24}. Because they constitute the input layer of the decision circuit, offer value cells are in some ways analogous to sensory cells. One might thus wonder whether their tuning functions match the cumulative distribution function of the offered values. Experimental evidence presented here indicates that this is not the case. More specifically, we show that the tuning functions of offer value cells are quasi-linear and not correlated with the cumulative distribution function of offered values (or its average across sessions). Thus the coding of offer values in OFC, while context-adapting, is not optimal in the sense of information transmission.

In this article, we introduce a new theory of optimal coding for economic decisions. In essence, we propose that offer value neurons are optimally tuned for economic decisions if they ensure maximal expected payoff. In this framework, we present a series of theoretical and experimental results. Behavioral and neuronal data were collected in two experiments in which monkeys chose between different juices offered in variable amounts. First, assuming linear tuning functions, we demonstrate that range adaptation, corrected to avoid choice biases, ensures maximal expected payoff. Second, confirming theoretical predictions, we show that expected payoff and value range are inversely related in the experiments. Third, relaxing the assumption of linearity, we demonstrate that optimal response functions in our experiments were in fact non-linear. Hence, linearity is a rigid property of value coding not subject to contextual adaptation. Fourth, we show that the benefit afforded by range adaptation outweighs the cost imposed by functional rigidity. In other words, quasi-linear but range-adapting tuning functions are sufficient to ensure close-to-optimal choice behavior. Taken together, these results shed new light on the nature of value coding and the role played by neuronal adaptation in economic decisions.

Results

Relative value, choice variability and expected payoff. In Exp. 1, monkeys chose between two juices (A and B, with A preferred) offered in variable amounts (Fig. 1a, b). The range of quantities offered for each juice remained fixed within a session, while the

quantity offered on any given trial varied pseudo-randomly. Monkeys' choices generally presented a quality–quantity trade-off. If the two juices were offered in equal amounts, the animal would generally choose A (by definition). However, if sufficiently large quantities of juice B were offered against one drop of juice A, the animal would choose B. The “choice pattern” was defined as the percentage of trials in which the animal chose juice B as a function of the offer type. In each session, the choice pattern was fitted with a sigmoid function, and the flex of the sigmoid provided a measure for the relative value of the two juices, referred to as ρ (see Methods). The relative value allows one to express quantities of the two juices on a common value scale. In one representative session, we measured $\rho = 4.1$ (Fig. 1b).

Choice patterns often presented some variability. For example, consider in Fig. 1b offers 6B:1A. In most trials, the animal chose juice B, consistent with the fact that the value of 6B was higher than the value of 1A. However, in some trials, the animal chose the option with the lower value. Similarly, in some trials, the animal chose 3B over 1A. Intuitively, choice variability is high when the sigmoid is shallow. Thus in each session, the steepness of the fitted sigmoid, referred to as η , quantified the (inverse of) choice variability (see Methods).

In any given trial, we define the payoff as the value chosen by the animal. Thus given a set of offers and a sigmoid function, the expected payoff is equal to the chosen value averaged across trials. Importantly, the expected payoff is inversely related to choice variability, and thus directly related to the steepness of the sigmoid. When the sigmoid is steeper, choice variability is lower, and the expected payoff is higher; when the sigmoid is shallower, choice variability is higher, and the expected payoff is lower. Notably, the relative value of two juices is entirely subjective. In contrast, a key aspect of the expected payoff is objective: given a set of offers, a relative value and two sigmoid functions, the steeper sigmoid yields higher expected payoff.

Quasi-linear coding of offer values. While animals performed the task, we recorded the activity of individual neurons in the central OFC. Firing rates were analyzed in multiple time

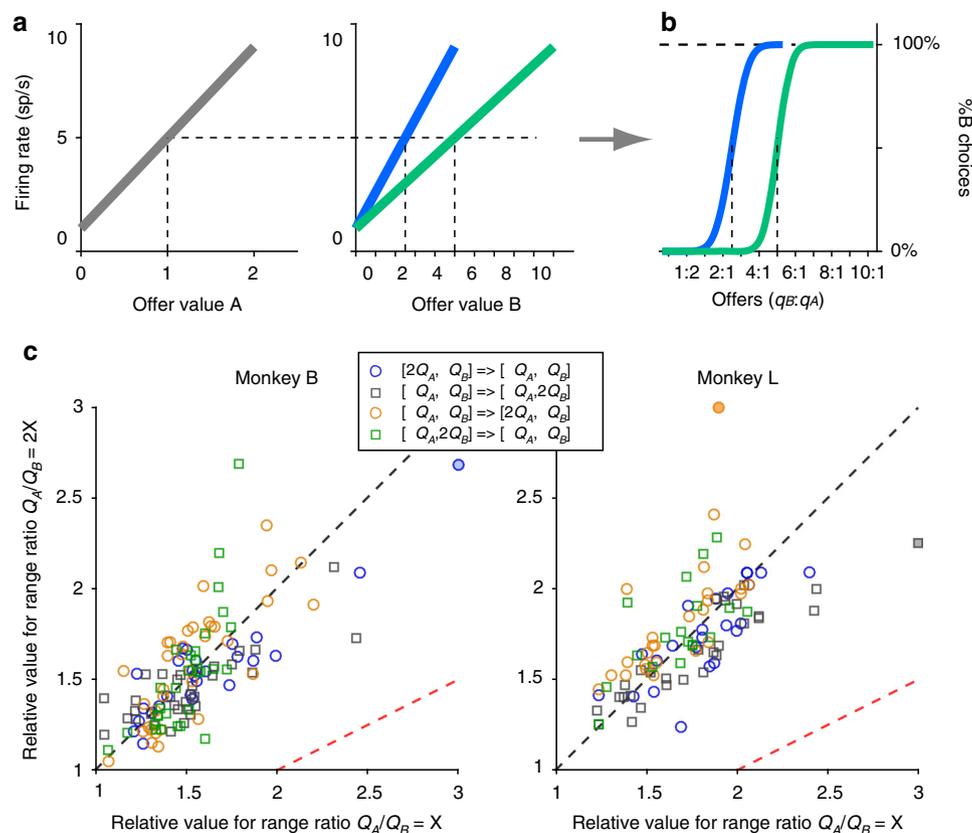


Fig. 3 Range adaptation is corrected within the decision circuit. **a, b** Uncorrected range adaptation would induce arbitrary choice biases. Panel **a** shows the schematic response functions of two neurons encoding the offer value A (left) and the offer value B (right). Panel **b** shows the resulting choice patterns under the assumption that decisions are made by comparing the firing rates of these two cells. We consider choices in two conditions, with the range $\Delta A = [0\ 2]$ kept constant. When $\Delta B = [0\ 5]$, the firing rate elicited by offer 1 A is between that elicited by offers 2B and 3B ($\rho = 2.5$). When $\Delta B = [0\ 10]$, offer value B cells adapt to the new value range. Now offer 1 A elicits the same firing rate as offer 5B ($\rho = 5$). Thus if range adaptation is not corrected, changing either value range induces a choice bias. Importantly, this issue would vanish if both neurons adapted to the same value range. However, experimental evidence indicated that each population of offer value cells adapts to its own value range⁹. **c** Relative values measured in Exp.2. The two panels refer to the two animals. In each panel, the axes represent the relative value measured when $Q_A/Q_B = X$ (x-axis) and that measured when $Q_A/Q_B = 2X$ (y-axis). Each data point represents data from one session, and different symbols indicate different protocols (see legend). If decisions were made by comparing uncorrected firing rates, data points would lie along the red dotted line. In contrast, data points lie along the black dotted line (identity line). In other words, the relative values measured in the two trial blocks were generally very similar, indicating that range adaptation was corrected within the decision circuit. Panels **a** and **b** are reproduced from⁹

windows. In each session, an “offer type” was defined by a pair of offers (e.g., [1A:3B]); a “trial type” was defined by an offer type and a choice (e.g., [1A:3B, B]); a “neuronal response” was defined as the activity of one cell in one time window as a function of the trial type. Earlier work showed that different responses encoded variables offer value, chosen value and chosen juice^{4, 5}. Unless otherwise indicated, the present analyses focus on offer value responses.

Previous studies failed to emphasize that the tuning of offer value cells was quasi-linear even though the distribution of values was highly non-uniform. To illustrate this point, we identified for each offer value response the quantity levels for the corresponding juice, and we calculated the number of trials in which each quantity level had been presented to the animal within the session. For example, Fig. 1c illustrates one offer value A response. In this session, juice A was offered in quantity levels (number of trials): 0 (39), 1 (169), 2 (19), and 3 (20). For each set of trials, we computed the mean firing rate of the cell. In addition, we computed the cumulative distribution function for the number of trials ($n_{\text{trialsCDF}}$) as a function of the quantity level. By analogy with sensory systems¹⁴, neurons encoding $n_{\text{trialsCDF}}$ would provide maximal information about the offer values. Firing

rates and $n_{\text{trialsCDF}}$ were highly correlated: for both of them, a linear regression on the quantity level provided a reasonably good fit (Fig. 1d). However, the non-uniform distribution of offer values induced a curvature in $n_{\text{trialsCDF}}$. Similarly, each neuronal response taken alone always presented some curvature. To assess whether and how the curvature in neuronal responses was related to the curvature in $n_{\text{trialsCDF}}$, we normalized both firing rates and $n_{\text{trialsCDF}}$ (see Methods). We thus fit each set of data points with a 2D polynomial, which provided a coefficient for the quadratic term (β_2). Separately, we fit each set of data points with a 3D polynomial, which provided a coefficient for the cubic term (β_3 ; Fig. 1e; see also Fig. 1f–h).

Because lower offer values were over-represented in the experiments, we generally measured $\beta_2, n_{\text{trialsCDF}} < 0$ and $\beta_3, n_{\text{trialsCDF}} > 0$. In contrast, $\beta_2, \text{firing rate}$ and $\beta_3, \text{firing rate}$ varied broadly across the population, and their distributions were fairly symmetric around zero (Fig. 2a, b). In other words, neuronal response functions were, on average, quasi-linear. These results held true for individual monkeys, in each time window, and independently of the sign of the encoding (Supplementary Fig. 1). Similar results were also obtained for chosen value responses (Supplementary Fig. 2).

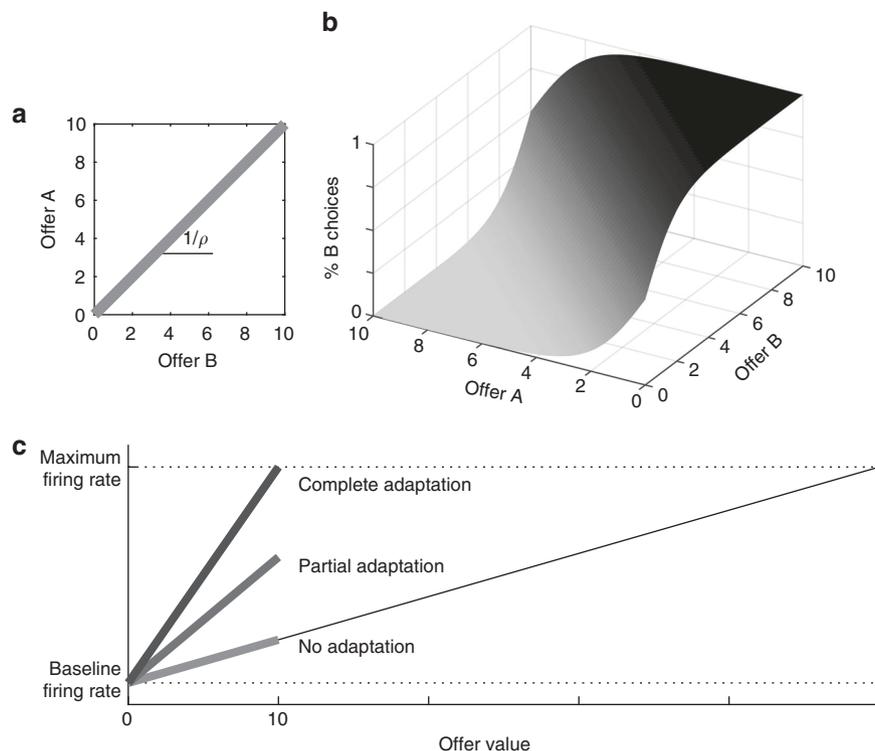


Fig. 4 Possible adaptation scenarios. **a** Indifference line. We indicate with q_A and q_B the quantities of good A and good B, respectively. Across trials, q_A varies in the range $[0, Q_A]$, while q_B varies in the range $[0, Q_B]$. In the plane defined by q_A and q_B , we define the “indifference curve” as the set of offers for which the animal splits decisions equally between the two goods. We assume that the indifference curve is a straight line on this plane. Thus the relative value between the two goods, referred to as ρ , is defined by the slope of the indifference line (slope = $1/\rho$). **b** Choice pattern. Given offers of goods A and B, a choice pattern can be represented as a sigmoid surface, in which the z-axis represents the likelihood of choosing good B. For each pair of offers, one of the two options provides a higher payoff, depending on whether it is above or below the indifference line. However, unless the sigmoid is a step function, in some trials the animal fails to choose that option (choice variability). Given a set of offers and a relative value, the expected payoff is an increasing function of the sigmoid steepness. **c**. Adaptation scenarios. In this cartoon, offer values in the current context vary in the range $[0, 10]$. The light line represents a hypothetical scenario in which there is no range adaptation (see Results). The darker lines represent the scenarios with partial and complete range adaptation

Range adaptation is corrected within the decision circuit. As previously shown, offer value cells undergo range adaptation (Supplementary Fig. 3a–c)⁶. Linear tuning implies that any given value interval is allotted the same activity interval in the neuronal representation. Range adaptation ensures that the full activity range is always available to represent the range of values offered in the current context. Thus range adaptation seems to provide an efficient representation for offer values. However, range adaptation also poses a computational puzzle⁹ illustrated in Fig. 3a, b. In essence, current models assume that binary economic decisions are made by comparing the firing rates of two neuronal populations encoding the subjective values of the offered goods^{25–29}. If so, by varying the ranges of the two offers one could impose any indifference point (an arbitrary choice bias).

Exp. 2 was conducted to test this prediction in controlled conditions. In each session, monkeys chose between two juices. Trials were divided in two blocks. Across blocks, we either halved or doubled the range of one of the two juices (2×2 design). For each trial block, Q_A and Q_B indicate the maximum quantities of juices A and B offered, respectively. Thus independently of other factors, the ratio Q_A/Q_B changed by a factor of two between blocks ($Q_A/Q_B = X$ or $2X$). The experimental design controlled for juice-specific satiety and other possible sources of choice bias (see Methods).

We collected behavioral data in 220 sessions. In each session and each trial block, we measured the relative value of the juices. We then compared the measures obtained in the two trial blocks.

According to the argument in Fig. 3a, b, the relative value measured when $Q_A/Q_B = X$ should be roughly twice that measured when $Q_A/Q_B = 2X$. Contrary to this prediction, the relative values measured in the two trial blocks were generally similar (Fig. 3c). Pooling all sessions, the ratio of relative values measured for the two trial blocks was statistically indistinguishable from 1 (mean ratio = 1.006; $p = 0.81$, Wilcoxon signed rank test) and significantly below 2 ($p < 10^{-37}$, Wilcoxon signed rank test). These results held true for each animal.

Range adaptation maximizes the expected payoff. Exp. 2 indicated that range adaptation is corrected within the decision circuit. We previously proposed a possible scheme for this correction. In essence, choice biases are avoided if the synaptic efficacies between offer value cells and downstream neuronal populations are rescaled by the value ranges^{9, 29}. However, if this correction occurs, it is reasonable to question whether range adaptation benefits the decision process at all. The central result of this study is that range adaptation in offer value cells maximizes the expected payoff even if adaptation is corrected within the decision circuit. The theoretical argument is summarized here and detailed in the [Supplementary Note](#), where we provide mathematical proofs.

Consider the general problem of choices between two goods, A and B. We indicate the quantities of A and B offered on a particular trial with q_A and q_B . Across trials, q_A varies in the range

$[0, Q_A]$, while q_B varies in the range $[0, Q_B]$. We assume linear indifference curves (Fig. 4a) and indicate the relative value with ρ . Choices can be described by a sigmoid surface (Fig. 4b). For each pair of offers, one of the two options provides a higher payoff, but in some trials the animal fails to choose that option (choice variability). Intuitively, this may happen because the neural decision circuit has a finite number of neurons, limited firing rates, trial-by-trial variability in the activity of each cell, and non-zero noise correlations.

Figure 4c illustrates the issue of interest. We assume that neuronal response functions are linear. Actual neurons always have a baseline firing rate (corresponding to a zero offer), but we assume that this activity does not contribute to the decision. Thus we focus on baseline-subtracted response functions. Let us consider a hypothetical scenario in which there is no adaptation. If so, neurons would have fixed tuning, corresponding to a linear response function defined on a very large value range. In contexts where the encoded good varies on a smaller range, neuronal firing rates would span only a subset of their potential activity range. In contrast, if neurons undergo complete range adaptation, firing rates span the full activity range in each behavioral context.

To understand how range adaptation in offer value cells affects the expected payoff, it is necessary to consider a specific decision model. That is, the question must be addressed under some hypothesis of how the activity of offer value cells is transformed into a decision. We examined the linear decision model^{30, 31} formulated as follows:

$$D = X^A - X^B \quad (1)$$

$$X^g = K_g \sum_i w_i^g r_i^g \quad g = A, B \quad i = 1 \dots n$$

where r_i^g is the firing rate of an offer value g cell, w_i^g are decision weights, n is the number of cells associated with each juice, and K_g is the synaptic efficacy of offer value g cells onto downstream populations. Conditions $D > 0$ and $D < 0$ correspond to choices of goods A and B, respectively.

We model the firing rates of offer value cells as Poisson variables and we approximate noise correlations with their mean long-distance component³⁰. In accord with experimental measures, we set the noise correlation to $\xi = 0.01$ for pairs of neurons associated with the same good, and to zero for pairs of neurons associated with different goods³⁰. Importantly, ξ does not depend on firing rates (Supplementary Fig. 4). We thus compute the probability of choosing juice A given offers $q = (q_A, q_B)$, tuning slopes $t = (t_A, t_B)$ and synaptic efficacies $K = (K_A, K_B)$:

$$P(ch = A|q, t, K) = \Pr\left(Z \geq -\frac{K_A q_A t_A - K_B q_B t_B}{\sqrt{\chi(K_A^2 q_A^2 t_A + K_B^2 q_B^2 t_B)}} \mid Z \sim N(0, 1)\right) \quad (2)$$

where $N(0, 1)$ is the standard normal distribution and $\chi = \xi/4$.

Eq. 2 allows one to calculate the expected payoff. Indicating with \bar{v} the maximum possible firing rate, we demonstrate that the expected payoff is maximal when $t_g = \bar{v}/Q_g$. This condition corresponds to complete range adaptation (Fig. 4c). In the symmetric case, defined by $\rho Q_A = Q_B$ (equal value ranges), the expected payoff is maximal when $K_A/K_B = 1$ and there is no choice bias. In the general, asymmetric case (unequal value ranges), the expected payoff is maximal when $K_A/K_B \approx \rho Q_A/Q_B$. In this condition, there is a small choice bias that favors the larger value range and depends on χ .

Eq. 2 expresses the sigmoid surface describing choices. By computing the slope of this surface on the indifference line, we show that under optimal coding the steepness of the sigmoid is inversely related to the value ranges (Supplementary Note, Eq. 28).

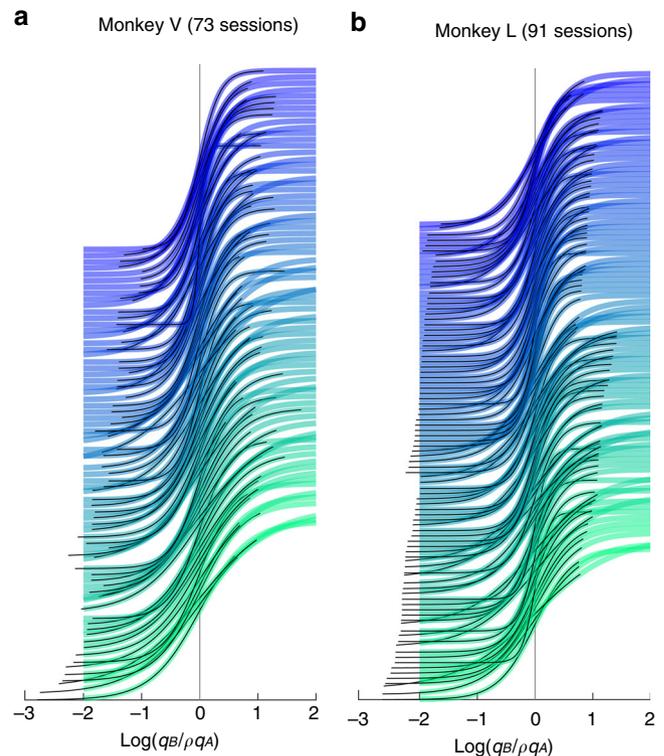


Fig. 5 Relation between sigmoid steepness and value range. **a** Monkey V (73 sessions). For each session, the sigmoid fit provided measures for ρ and η (Eq. 6), and we computed the geometric mean value range Δ . In this plot, different sigmoid functions are aligned at the flex (x -axis) and ranked based on Δ , from top (small Δ) to bottom (large Δ). For each sigmoid, the thick colored line (blue-green) depicts the result of the fit in a standard interval $[-2, 2]$. The thin black line highlights the range of values actually used in the corresponding session. Different shades of color (from blue to green) indicate the ordinal ranking of sessions according to Δ . Notably, sigmoid functions at the bottom of the figure (larger Δ) were shallower (lower η). **b** Monkey L (91 sessions). Same format as in **a**

Relation between choice variability and value range. The previous section summarizes a theory of optimal coding of offer values for economic decisions. The main prediction for linear response functions is that the slope of the encoding should be inversely proportional to the value range, as is indeed observed in the experiments (range adaptation; Supplementary Fig. 3d, e). The theory also makes another testable prediction. Consider experiments in which monkeys choose between two juices and value ranges vary from session to session. The sigmoid steepness should decrease as a function of the value ranges. To test this prediction, we examined 164 sessions from Exp.1. For each session, we computed the geometric mean value range $\Delta \equiv (\rho Q_A Q_B)^{1/2}$, and we obtained a measure for the sigmoid steepness (η) from the sigmoid fit. We thus examined the relation between η and Δ .

Figure 5a, b illustrates the fitted sigmoid obtained for each experimental session in our data set, separately for monkeys V and L. For each animal, sigmoid functions were aligned at the flex and ranked according to Δ . Notably, sigmoid functions with small Δ were generally steeper (large η), while sigmoid functions with large Δ were generally shallower (small η). In other words, there was a negative correlation between η and Δ . This correlation, summarized in a scatter plot (Fig. 6), was statistically significant in each animal (monkey V: corr coef = -0.41 , $p < 0.0005$; monkey L: corr coef = -0.26 , $p < 0.02$). Control analyses confirmed that this result was not due to differences between juice pairings (Supplementary Fig. 5) or to fluctuations in the relative

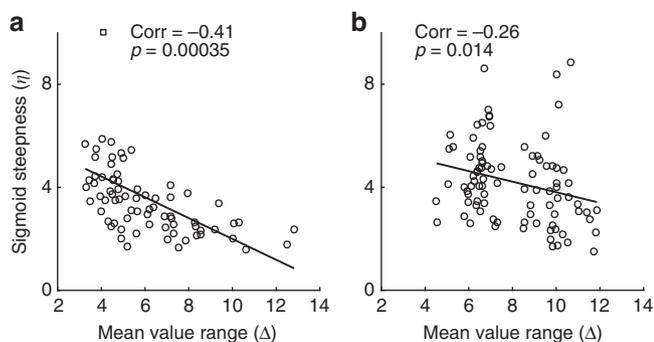


Fig. 6 Relation between sigmoid steepness and value range, scatter plots. **a**, **b**. Panels **a** and **b** refer to monkey V (73 sessions) and monkey L (91 sessions), respectively. In each panel, the x-axis represents the geometric mean value range $\Delta \equiv (\rho Q_A Q_B)^{1/2}$, the y-axis represents the steepness of the sigmoid (η) and each data point represents one session. In both animals, the two measures were significantly and negatively correlated. In each panel, the black line represents the result of Deming's regression (see Methods)

value (Supplementary Fig. 6). Similar results were also obtained for data from Exp.2 (Supplementary Fig. 7).

Neuronal responses are functionally rigid. We have shown that range adaptation maximizes the expected payoff under the assumption of linear response functions. Next we address a closely related question, namely whether (or in what sense) linear response functions are optimal in the first place. In the visual system, optimal coding is achieved if tuning functions match the cumulative distribution of the encoded stimuli^{14, 19}. In the valuation system, the equivalent condition would occur if offer value responses matched the cumulative distribution of offered values. We already showed that this is not the case (Fig. 2). In retrospect, this finding is not surprising because a subject performing economic decisions is best served by response functions that maximize the expected payoff, which do not necessarily maximize information transmission. Thus what is the optimal response function for offer value cells?

The answer to this question depends on the joint distribution of offers and on the relative value of the two goods. For example, consider the case in which an animal chooses between goods A and B and $\rho = 2$. Good A is always offered in quantity 1, while good B is offered in quantities between 0 and 5 (Fig. 7a). We consider offer value B cells and we indicate with r_B their firing rate. It is easy to see that the payoff is maximal if $r_B(x) = 0$ when $x < 2$, $r_B(2) = 0.5$, and $r_B(x) = 1$ when $x > 2$, where x are quantities of juice B offered. Hence, the optimal response function is a step function with the step located at $x = 2$. Next consider the case in which quantities of both goods vary between 0 and 5, at least one of the two goods is always offered in quantity 1, and $\rho = 2$ (Fig. 7b). Again, the optimal response function for offer value B cells is $r_B(x) = 0$ when $x < 2$, $r_B(2) = 0.5$, and $r_B(x) = 1$ when $x > 2$. For offer value A cells, the optimal response function is $r_A(0) = 0$, $r_A(1) = 0.5$, and $r_A(x) = 1$ when $x > 1$. Thus for both goods, the optimal response function is a step function. Analogously, if offer types are the same but $\rho = 3$ (Fig. 7c), the optimal response function for offer value B cells is a step function with the step located at $x = 3$.

The scenarios depicted in Fig. 7b, c are similar to those occurring in Exp.1. Indeed our sessions always included forced choices for both juices. Furthermore, in 96% (200 out of 208) of our sessions, when both juices were offered, at least one of them was offered in quantity 1 (Supplementary Fig. 8). Thus in Exp.1, optimal response functions for offer value cells would have been

step functions, not linear functions. Our neuronal data clearly belied this prediction (Fig. 2). In other words, our results indicate that the functional form of offer value cells did not adapt to maximize the payoff in each session. To further examine this point, we ran two additional analyses.

First, we entertained the hypothesis that the functional form of offer value cells might adapt on a longer time scale, over many sessions. However, we found that the mean optimal response function was a fairly sharp sigmoid (Fig. 7d), contrary to our observations (Fig. 2). Second and most important, we recognized that neuronal responses examined in Fig. 2 were originally identified through a variable selection analysis that only considered linear response functions⁴ (see Methods). This effectively imposed a bias in favor of linearity. To eliminate this bias, we repeated the variable selection procedures including in the analysis all the variables discussed in this study. These included the cumulative distribution function of offer values ($n_{\text{trialsCDF}}$), the optimal responses in each session (step functions) and the mean optimal response function across sessions (Methods). The results confirmed previous findings: variables offer value, chosen value and chosen juice still provided the highest explanatory power. In particular, the explanatory power of linear offer value variables was significantly higher than that of each of the new variables (Supplementary Table 1).

In the final analysis of this section, we considered whether the response functions observed experimentally would maximize the expected payoff for other possible joint distributions of offers. To do so, we generalized the theory of optimal coding by relaxing the assumption of linear response functions (Supplementary Note, Section 6). One interesting candidate was the symmetric uniform distribution (Fig. 7e). We calculated the optimal response functions given this distribution ($\text{ORF}_{\text{uniform}}$) and we found that they are quasi-linear and slightly convex (Fig. 7e). Notably, this non-linearity is in the same direction observed in Fig. 2a (histogram). We then repeated the variable selection analysis including variables based on $\text{ORF}_{\text{uniform}}$. Interestingly, neuronal responses best explained by $\text{ORF}_{\text{uniform}}$ variables were more numerous than those best explained by linear offer value variables (Fig. 8). As in previous studies⁴, we used two procedures for variable selection, namely stepwise and best-subset (Methods). Both procedures identified variables offer A $\text{ORF}_{\text{uniform}}$, offer B $\text{ORF}_{\text{uniform}}$, chosen value and chosen juice as providing the maximum explanatory power (Fig. 9). However, a post-hoc analysis indicated that the explanatory power of $\text{ORF}_{\text{uniform}}$ variables was statistically indistinguishable from that of linear offer value variables (Supplementary Table 2).

In conclusion, the variable selection analyses confirmed that offer value responses were quasi-linear and thus suboptimal given the joint distributions of offers in our experiments. Furthermore, offer value responses were indistinguishable from optimal responses functions calculated assuming a uniform joint distribution of offers. We elaborate on the significance of this finding in the Discussion.

Cost of functional rigidity and benefit of range adaptation. The tuning of offer value cells is functionally rigid (quasi-linear) but parametrically plastic (range adapting with optimal gain). In terms of the expected payoff, functional rigidity ultimately imposes some cost, while range adaptation ultimately yields some benefit. We sought to quantify these two terms in our experiments.

For each session of Exp.1, we focused on strictly binary choices (i.e., we excluded forced choices). On the basis of the relative value of the juices (ρ), we computed for each trial the chosen value (i.e., the payoff) and the max value, defined as the higher of

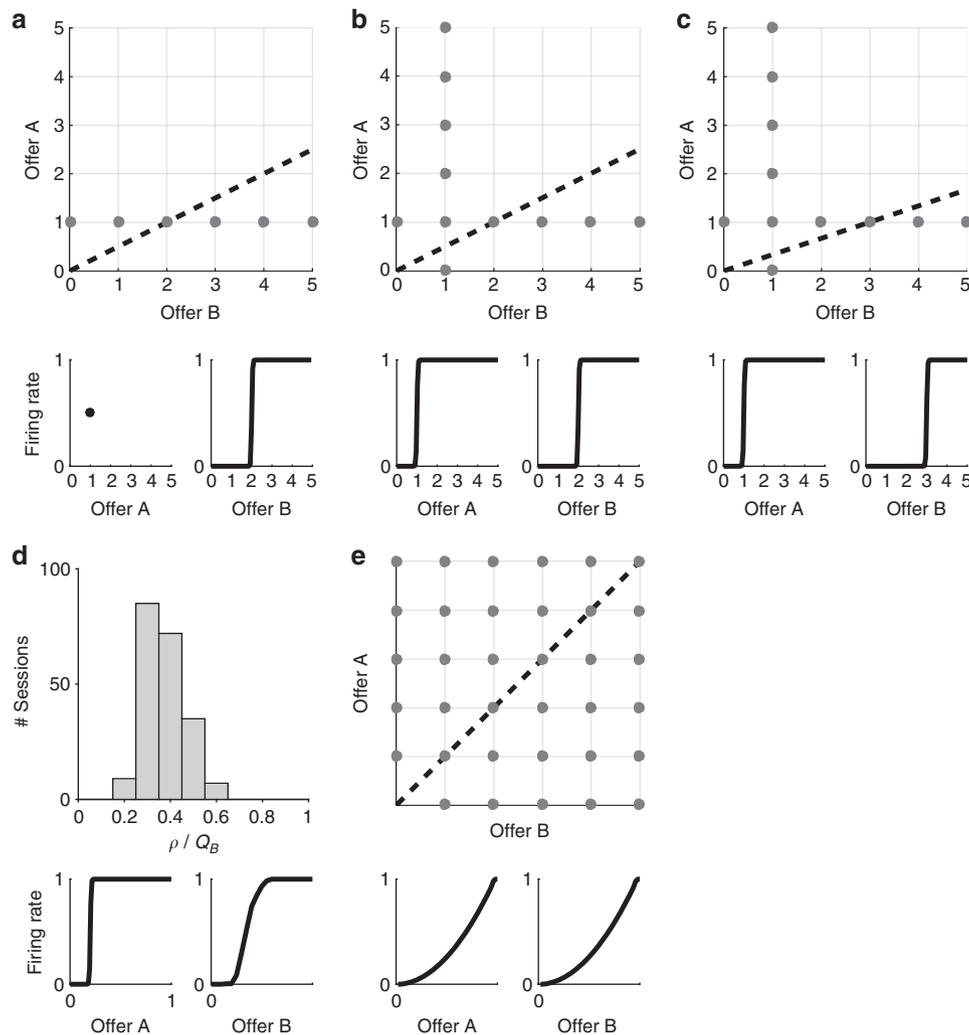


Fig. 7 Optimal response functions. **a** One good offered in fixed quantity ($\rho = 2$). Gray dots represent offer types presented in the session and the dotted line represents the indifference line. Good A is always offered in quantity 1 while good B varies in the range [0 5]. Optimal response functions are shown in the lower panels. **b** Idealized experimental session ($\rho = 2$). For each good, quantities vary in the range [0 5], but in each offer type at least one good is offered in quantity 1. Lower panels show the optimal response functions (ORF, step functions). **c** Idealized experimental session ($\rho = 3$). **d** Optimal mean response functions. The histogram represents the distribution of ρ/Q_B , where ρ is the relative value and Q_B is the maximum quantity of juice B offered. Lower panels show the mean optimal response functions, mean(ORF). For offer value B, the response function is computed as the cumulative distribution function for ρ/Q_B . **e** Idealized session with uniform distribution and equal value ranges (a.u.). Lower panels show the corresponding optimal response functions (ORF_{uniform}). Note that the curvature of ORF_{uniform} is in the same direction as that observed on average in the neuronal population (Fig. 2a, histogram)

the two values offered in that trial. We also defined the chosen value_{chance} as the chosen value expected if the animal chose randomly between the two offers. Hence, chosen value_{chance} = (offer value A + offer value B)/2. For each session we defined the fractional lost value (FLV) as:

$$\text{FLV} = \text{fractional lost value} = \frac{\langle \text{max value} - \text{chosen value} \rangle}{\langle \text{max value} - \text{chosen value}_{\text{chance}} \rangle} \quad (3)$$

where brackets indicate an average across trials. Under normal circumstances, FLV varies between 0 and 1. Specifically, FLV = 0 if the animal always chooses the higher value (chosen value = max value) and FLV = 1 if the animal always chooses randomly (chosen value = chosen value_{chance}). Thus FLV quantifies the fraction of value lost to choice variability. For each session, we also computed the percent error, defined as the percent of trials in

which the animal chose the lower value. We examined these metrics across sessions.

The percent error varied substantially from session to session, between 0 and 23% (Fig. 10a). On average across sessions, mean (percent error) = 8.7%. The FLV also varied substantially across sessions, between 0 and 0.24 (Fig. 10b). On average across sessions, mean(FLV) = 0.05. Importantly, this estimate provides an upper bound for the value lost by the animal due to suboptimal tuning functions, because other factors might also contribute to choice variability. Hence, the cost of functional rigidity in the coding of offer values may be quantified as ≤ 0.05 .

Because we cannot observe decisions in the absence of neuronal adaptation, quantifying the benefits of range adaptation requires a simulation. We proceeded as follows. For each session and for each trial, the sigmoid fit provided the probability that the animal would choose juice B ($P_{ch=B}$; see Eq. 5) or juice A ($P_{ch=A} = 1 - P_{ch=B}$). Thus in each trial the expected chosen value

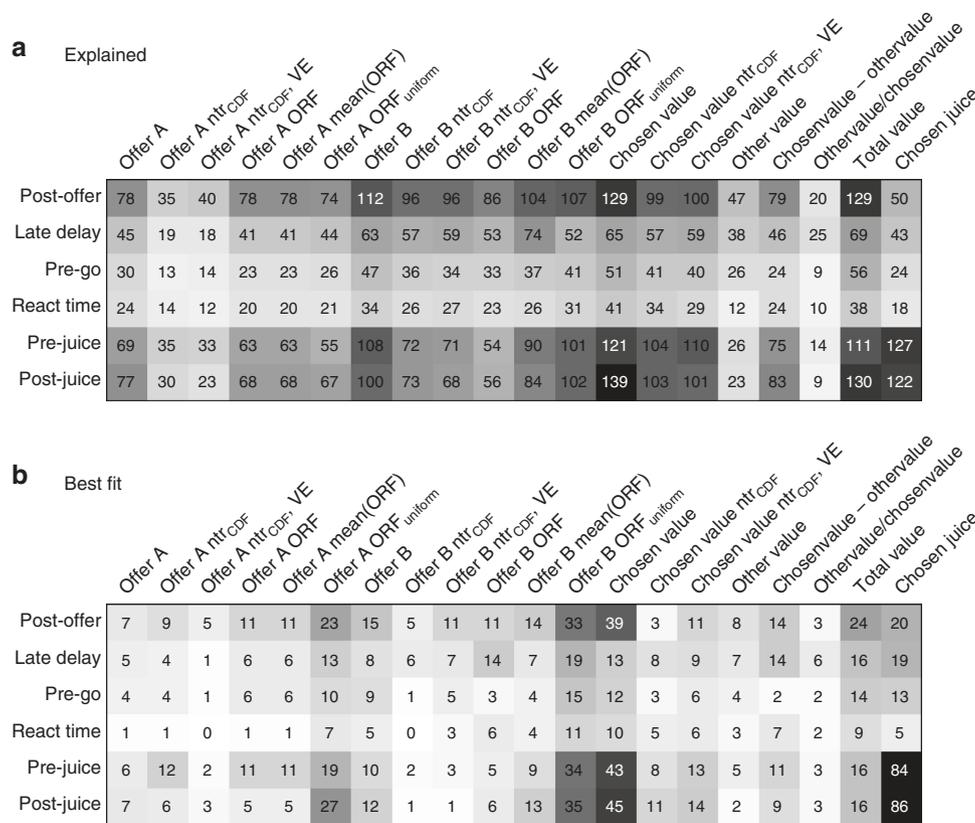


Fig. 8 Population analysis of neuronal responses. Each neuronal response that passed an ANOVA criterion was regressed against each variable. If the regression slope differed significantly from zero, the variable was said to explain the response (see Methods). **a** Explained responses. Rows and columns represent time windows and variables, respectively. In each location, the number indicates the number of responses explained by the corresponding variable. For example, in the post-offer time window, the variable offer A (linear response function) explained 78 responses. The same numbers are also represented in grayscale. Each response could be explained by more than one variable. Thus each response might contribute to multiple bins in this panel. **b** Best fit. In each location, the number indicates the number of responses for which the corresponding variable provided the best fit (highest R^2). For example, in the post-offer time window, offer A (linear response function) provided the best fit for 7 responses. The same numbers are also represented in grayscale. In this panel, each neuronal response contributes at most to one bin. Qualitatively, the dominant variables appear to be offer A ORF_{uniform}, offer B ORF_{uniform}, chosen value and chosen juice. Indeed the variable selection procedures identified these variables as the ones with the highest explanatory power (Fig. 9)

(i.e., the expected payoff) was:

$$E(\text{chosen value}) = P_{ch=A} \text{ offer value A} + P_{ch=B} \text{ offer value B} \quad (4)$$

For each session, we computed the expected fractional lost value (EFLV) by substituting the $E(\text{chosen value})$ for the chosen value in Eq. 3. Importantly, we verified that EFLV provided a good estimate for the actual FLV (Fig. 10c).

To address the question of interest, we reasoned along the lines of Fig. 4c, where the absence of adaptation is approximated with a scenario in which neurons adapt to a very large value range. We already showed that increasing the value range decreases the sigmoid steepness (Fig. 6). Thus we examined how reducing the sigmoid steepness would affect the EFLV. We found that the effects were large. For example, when we halved the sigmoid steepness ($\eta \rightarrow \eta/2$), we obtained $\text{mean}(\text{EFLV}) = 0.15$; when we divided the sigmoid steepness by ten ($\eta \rightarrow \eta/10$), we obtained $\text{mean}(\text{EFLV}) = 0.55$ (Fig. 10d). Hence, the benefit of range adaptation, while difficult to quantify exactly, is clearly very high.

To summarize, the benefit of range adaptation outweighs the cost of functional rigidity. Our analyses suggest that a quasi-linear but range-adapting coding of offer values is sufficient to ensure close-to-optimal choice behavior.

Discussion

Sensory neurons are optimally tuned for perception if they transmit maximal information about the stimuli. In contrast, offer value neurons are optimally tuned for economic decisions if they ensure maximal expected payoff. In this framework, we examined the activity of offer value cells in OFC. These neurons are believed to provide the primary input for economic decisions. We showed that their tuning is functionally rigid (linear responses) but parametrically plastic (range adaptation with optimal gain). We also showed that range adaptation is corrected within the decision circuit to avoid arbitrary choice biases. Critically, range adaptation ensures optimal tuning even considering this correction. Confirming theoretical predictions, we showed that choice variability is directly related to the range of values offered in any behavioral context. Finally, we showed that the benefit of range adaptation outweighs the cost of functional rigidity. Importantly, our theoretical results were derived using a linear decision model (Eq. 1)^{30,31}. Future work should extend this analysis to other decision models^{25, 27, 29}.

On average, offer value responses presented a small but significant departure from linearity (Fig. 2). Their convexity closely resembled that predicted for optimal response functions under a uniform joint distribution (Fig. 7e), although in a direct comparison the explanatory power of ORF_{uniform} functions was not significantly higher than that of strictly linear functions. Thus

future work should address this point and consider other joint distributions that might explain neuronal responses in OFC. Nonetheless, the quasi-linear nature of value coding in OFC is noteworthy. We previously showed that the activity of neurons associated with one good does not depend on the identity or value of the other good offered in the same trial³². With respect to range adaptation, we also showed that each neuron adapts to its

own value range, independently of the range of values offered for the other good⁹. One implication of linear responses (or optimal response functions under a symmetric uniform distribution) is that the activity of neurons associated with one particular good does not depend on the distribution of values offered for the other good, or on the relative value of the two goods. Thus quasi-linearity can be seen as yet another way in which neurons associated with one particular good are blind to every aspect of the other good. This blindness, termed menu invariance, guarantees preference transitivity^{33, 34}, which is a fundamental trait of economic behavior. It is tempting to speculate that quasi-linear coding might have been selected in the course of evolution because it facilitates transitive choices.

Adaptive coding has been observed in numerous brain regions that represent value-related variables including the amygdala³⁵, anterior cingulate cortex³⁶, and dopamine cells^{37–39}. Independently of the specific contribution of each area to behavior, adaptation necessarily poses computational challenges analogous to the coding catastrophe discussed in sensory systems^{9–11}. In essence, uncorrected adaptation makes firing rates intrinsically ambiguous. Thus neuronal adaptation at any processing stage must be corrected at later stages⁴⁰ or ultimately results in impaired behavioral performance^{12, 41}. With respect to offer value cells in OFC, we previously proposed that choice biases potentially introduced by range adaptation are corrected in the synapses between these neurons and downstream populations^{9, 29}. The theory of optimal coding developed here makes this same prediction, which should be tested in future experiments. Interestingly, framing^{42, 43} and anchoring⁴⁴ effects documented in behavioral economics qualitatively resemble adaptation-driven choice biases, although they are quantitatively more modest. In principle, these effects could be explained if synaptic rescaling trailed neuronal range adaptation. Similar mechanisms have been hypothesized in the visual system to explain illusions and aftereffects^{11, 12}.

The rationale for this study rests on the assumption that offer value cells in OFC provide the primary input for the neural circuit that generates economic decisions. Support for this assumption comes from lesion studies^{45–47}, from the joint analysis of choice probability and noise correlation³⁰ and from the relation between choice variability and value range shown here. Indeed, current neuro-computational models of economic decisions embrace this view^{26, 29, 48–51}. However, causal links between the activity of offer value cells and the decision have not yet been demonstrated with the gold-standard approach of biasing choices using electrical or optical stimulation. Future work should fill this important gap.

Methods

Experimental procedures. All experimental procedures conformed to the NIH Guide for the Care and Use of Laboratory Animals and were approved by the

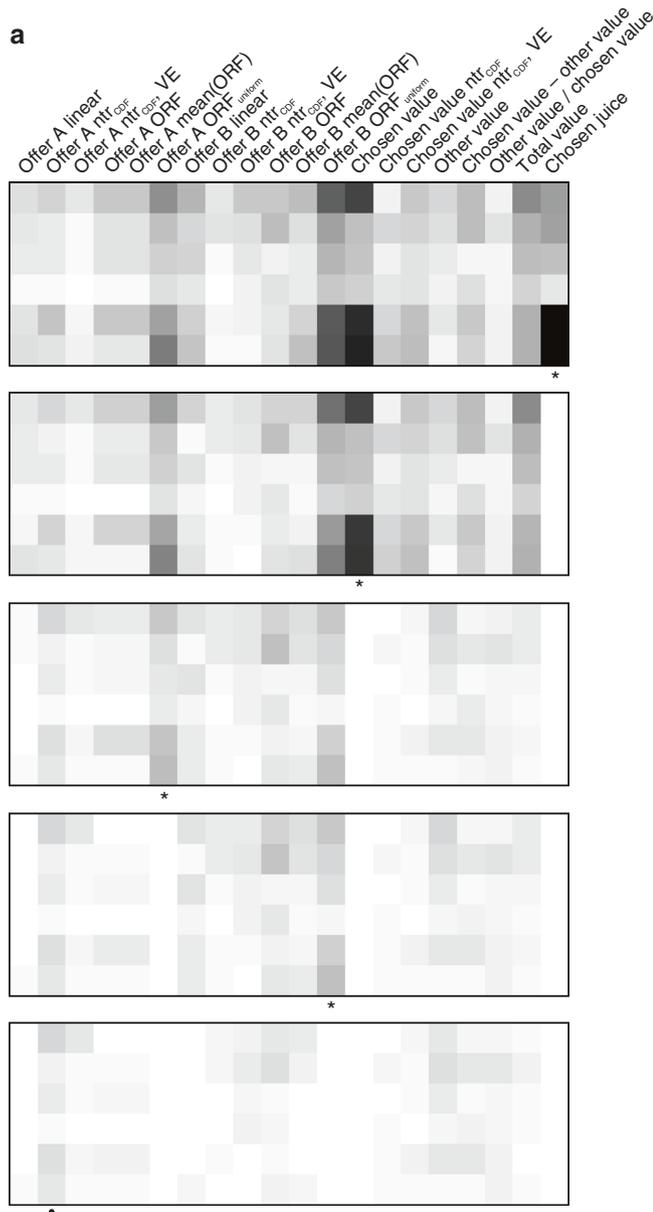


Fig. 9 Variable selection analysis. **a** Stepwise selection. The top panel is as in Fig. 8b. At each iteration, the variable providing the maximum number of best fits in a time window was selected and indicated with a * in the figure. All the responses explained by the selected variable were removed from the pool and the procedure was repeated on the residual data set. Selected variables whose marginal explanatory power was <5% were eliminated (Methods) and indicated with a • in the figure. In the first four iterations, the procedure selected variables chosen juice, chosen value, offer A ORF_{uniform} and offer B ORF_{uniform}, and no other variables were selected in subsequent iterations. **b** Percent of explained responses. The y-axis represents the percentage of responses explained at the end of each iteration. The total number of task-related responses (1378) corresponds to 100%. The number of responses explained by at least one of the variables included in the analysis (1245/1378=90%) is indicated with a dotted line

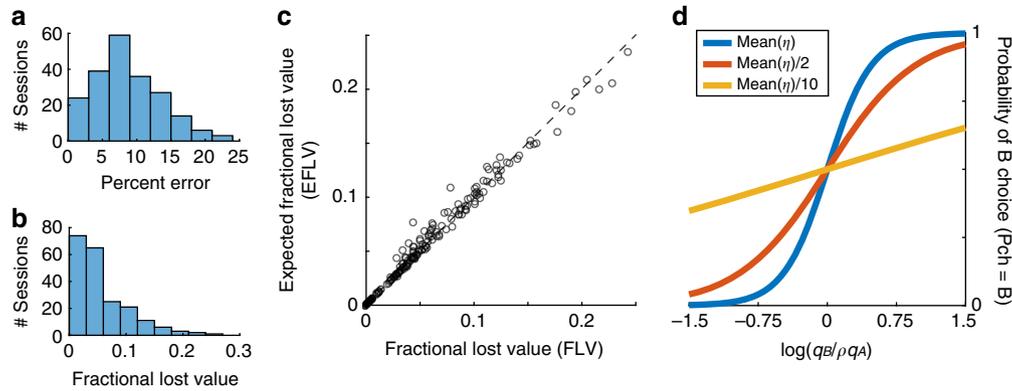


Fig. 10 Analysis of fractional lost value. **a** Distribution of percent errors across sessions. On average across sessions, mean(percent errors)=8.7%. **b** Distribution of fractional lost value across sessions. On average across sessions, mean(FLV) = 0.054. **c** Expected vs. actual fractional lost value. Each data point represents one session. The expected fractional lost value (EFLV; y-axis) was almost identical to the fractional lost value (FLV; x-axis). On average across sessions, mean(EFLV) = 0.055. **d** Effects of decreasing the steepness of the sigmoid. The range on the x-axis is realistic for our experiments. The blue line is obtained with the mean sigmoid steepness measured in Exp.1 (mean(η)). The red and yellow lines were obtained by dividing mean(η) by 2 and 10, respectively

Institutional Animal Care and Use Committees at Harvard University (Exp.1) and Washington University in St Louis (Exp.2). No subject randomization or blinding during data analysis was used.

The procedures for Exp.1 have been described previously⁴. Briefly, one male (V, 9.5 kg) and one female (L, 6.5 kg) rhesus monkey participated in the experiment. Animals sat in an electrically insulated enclosure with the head restrained, and a computer monitor was placed in front of them at 57 cm distance. In each session, the monkey chose between two juices, labeled A and B, with A preferred. At the beginning of each trial, the animal fixated a center position on the monitor (Fig. 1a). After 0.5 s, two sets of colored squares appeared on the two sides of the center fixation. The two sets of squares represented the two offers, with the color associated with a particular juice type and the number of squares indicating the juice quantity. The animal maintained center fixation for a randomly variable delay (1–2 s), at the end of which the center fixation point was extinguished (go signal). The animal revealed its choice by making a saccade towards one of two targets located by the offers, and maintained peripheral fixation for an additional 0.75 s before the chosen juice was delivered. While animals performed in the task, we recorded the activity of individual neurons from the central OFC (see below). The range of quantities offered for each juice remained fixed within a session, while the quantity offered on any given trial varied pseudo-randomly. Across sessions, we used various juice pairs and various quantity ranges for the two juices. The minimum quantity was always zero drops (forced choice for the other juice), while the maximum quantity varied from session to session between 2 and 10 drops.

In Exp.2, animals performed essentially the same task, except that sessions were divided into two blocks of trials. One male (B, 9.0 kg) and one female (L, 6.5 kg) rhesus monkey participated in the experiment. The task was controlled through custom-written software based on Matlab (MathWorks)^{52, 53} and gaze direction was monitored with an infrared video camera (Eyelink, SR research). The trial structure was the same as Exp.1, except that the initial fixation lasted 1.5 s. Each session included two trial blocks. The minimum offered quantity for each juice was always set to zero (forced choice for the other juice). The maximum quantity (and thus the range) varied from session to session and from block to block. In the second block, we either halved or doubled the range for one juice (A or B) while keeping the other range unchanged. This procedure resulted in a 2x2 design. Each block included 110–260 trials. In each block, an “offer type” was defined by a pair of offers (e.g., [1A:3B]); a “trial type” was defined by an offer type and a choice (e.g., [1A:3B, B]). The relative value of the two juices was computed from the indifference point (see below).

In principle, changes in relative value could arise from factors other than the value range. Exp.2 was designed to minimize three potential sources of choice bias. First, in previous work, we often noted that the relative value of any two juices tends to increase over the course of each day, presumably because animals become less thirsty. To deconfound changes in relative value due to changes in value range from this effect, we alternated sessions in which we increased or decreased the range of either juice A or juice B. The number of sessions for each of the 4 possible combinations was not predetermined with a statistical method but was comparable ($\Delta A \rightarrow 2\Delta A$, 61 sessions; $\Delta B \rightarrow 2\Delta B$, 62 sessions; $2\Delta A \rightarrow \Delta A$, 49 sessions; $2\Delta B \rightarrow \Delta B$, 48 sessions). Second, within each trial block, monkeys might experience juice-specific satiety or diminishing marginal returns. Thus to isolate the behavioral effects of manipulating the value range, we ensured that in both trial blocks the animal drank the same relative amounts of the two juices. For example, if the animal drank juice A and juice B in quantity ratio 3:2 in the first block, we kept the same ratio 3:2 in the second block (see below). Third, we previously found that, all other things equal, monkeys tend to choose on any given trial the same juice they

chose in the previous trial (choice hysteresis)⁵. If the relative number of trials in which the animal chooses a particular juice varies from one block to the other, choice hysteresis could introduce a systematic bias. To avoid this confound, we ensured that the relative number of choices was the same in the two trial blocks.

The relative number of choices and the relative amount drunk by the animal for each juice were controlled by adjusting the frequency with which each offer type was presented. Specifically, offers were presented pseudo-randomly in mini-blocks of 20–30 trials. To fine-tune the balance between juice A and B, we kept track of the monkey’s choices online. If the choice ratio or the relative amount of juice changed in the second block, the imbalance was corrected by adding forced choices of one of the two juices.

Analysis of behavioral data. Monkeys’ choices generally presented a quality-quantity trade-off. If the two juices were offered in equal amounts, the animal would generally choose juice A (by definition). However, if sufficiently large quantities of juice B were offered against one drop of juice A, the animal would choose B. Choices were analyzed separately in each session (Exp.1) or in each trial block (Exp.2). The “choice pattern” was defined as the percentage of trials in which the animal chose juice B as a function of the log quantity ratio $\log(q_B/q_A)$, where q_A and q_B indicate the quantities of juices A and B. Each choice pattern was fitted with a sigmoid function:

$$P_{ch=B} = \int_{-\infty}^X N(0, 1) dt \tag{5}$$

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

where $P_{ch=B}$ is the probability of choosing juice B and $N(0, 1)$ is the standard normal distribution. The fit was done with Matlab function glmfit and link=probit. From the fitted parameters a_0 and a_1 we defined the relative value ρ and the steepness of the sigmoid η as follows:

$$\rho = \exp(-a_0/a_1) \tag{6}$$

$$\eta = a_1$$

Given a set of offers, the expected payoff is directly related to η . In some simulations (Fig. 10), we reduced the sigmoid steepness (e.g., $a_1 \rightarrow a_1/10$) while keeping the relative value constant ($a_0/a_1 \rightarrow a_0/a_1$).

Exp.1 included 208 sessions. However, in some cases the choice patterns were saturated (i.e., the animal did not split decisions for any offer type, a situation referred to as “perfect separation”). In these cases, the sigmoid fit did not provide a reliable measure for η . Thus the analysis shown in Fig. 5 and Fig. 6 included only sessions for which choice patterns were not saturated (164 sessions).

In Exp.1, the minimum quantity offered for each juice was always 0, and we indicate maximum quantities with Q_A and Q_B . We usually set Q_A and Q_B to approximately satisfy $\rho Q_A = Q_B$ (symmetric condition). However, this relation did not hold strictly, partly because the relative value ρ was determined by the animal and fluctuated from session to session. Thus to test a theoretical prediction on choice variability and value range, we computed the geometric mean value range $\Delta \equiv (\rho Q_A Q_B)^{1/2}$ and we examined the relation between η and Δ . Since errors of measure affected both measures, standard regressions could not be used. We thus used Deming’s regressions⁵⁴. Variance ratios λ were computed through error

propagation as follows:

$$\lambda = \frac{\text{var}(\eta)}{\text{var}(\rho Q_A Q_B)^{1/2}} = \frac{(\delta\eta)^2}{Q_A Q_B (\delta(\rho^{1/2}))^2} = \frac{4a_1^4 (\delta a_1)^2}{\rho Q_A Q_B (a_0 \delta a_1 - a_1 \delta a_0)^2} \quad (7)$$

where $\delta\rho$, $\delta\eta$, δa_0 , and δa_1 are errors on the respective measures, and δa_0 and δa_1 are obtained as standard errors from the logistic regressions.

The relation between η and Δ was also analyzed using alternative definitions for Δ including the simple mean $\Delta \equiv (\rho Q_A + Q_B)/2$ and the log geometric mean $\Delta \equiv \log(\rho Q_A Q_B)/2$, adjusting variance ratios accordingly. All variants of the analysis provided very similar results.

One concern was whether the relation between choice variability and value range (Fig. 6) was direct or reflected some other dependency. We considered two issues. First, Fig. 6 includes sessions with different juice pairs, with different typical values of ρ . In principle, choice variability could vary from juice pair to juice pair in a way that induces the relation observed in Fig. 6. Second, for any given juice pair, value range (Δ), relative value (ρ) and sigmoid steepness (η) are all inter-related by definition (Eq. 6) and because value ranges were often chosen non independently of ρ (in many sessions we set $Q_A=3$ and chose Q_B such that $Q_B \approx \rho Q_A$). Thus the relation between η and Δ (Fig. 6) might simply reflect fluctuations in ρ . To address these concerns, we divided sessions in different sets based on the animal and on the juice pair. We considered only sets with ≥ 5 sessions, and our data included 12 such sets (6 from each monkey). We then analyzed each set of sessions separately. First, we verified that the relation between Δ and η held true for individual juice pairs (Supplementary Fig. 5). Second, to assess whether this relation simply reflected fluctuations in ρ , we used multilinear regression. For each set, we regressed η on ρ and then on Δ in a stepwise way. The coefficient obtained from the second regression (β_2) essentially quantified the correlation between η and Δ not explained by fluctuations of ρ .

Analysis of neuronal data. Neuronal data were collected in Exp.1. The data set included 931 cells from central OFC (area 13). The number of cells recorded was not pre-determined using statistical methods. Firing rates were analyzed in seven time windows aligned with different behavioral events (offer, go, juice). A “neuronal response” was defined as the activity of one cell in one time window as a function of the trial type. Neuronal responses were computed by averaging firing rates across trials for each trial type.

In a previous study⁴, we conducted a series of analyses to identify the variables encoded in this area. First, we submitted each neuronal response to a 3-way ANOVA (factors: position x movement direction x offer type), and we imposed a significance threshold $p < 0.001$. In total, 505 (54%) neurons passed the criterion for factor offer type in at least one time window. Pooling neurons and time windows, 1379 neuronal responses passed the ANOVA criterion, and subsequent analyses were restricted to this data set. Second, we defined a large number of variables including variables related to individual juices (offer value A, offer value B, chosen value A, chosen value B), other value-related variables (chosen value, other value, value difference, value ratio, total value), number-related variables (chosen number, other number, and so on) and choice-related variables (chosen juice). We performed a linear regression of each response onto each variable. If the regression slope differed significantly from zero ($p < 0.05$), the variable was said to “explain” the response. Because variables were often correlated with each other, the same neuronal response was often explained by more than one variable. Thus for each response we also identified the variable that provided the best fit (i.e., the highest R^2). Third, we proceeded with a variable selection analysis to identify a small subset of variables that best explained the neuronal population. We adapted two methods originally developed for multi-linear regressions in the presence of multi-collinearity, namely stepwise and best-subset^{55, 56}. In the stepwise method, we identified the variable and time window that provided the highest number of best fits, and removed from the data set all the responses explained by that variable. We then repeated the procedure until when the number of responses explained by additional variables was $< 5\%$. While intuitive, the stepwise procedure did not guarantee optimality. In contrast, the best-subset procedure (an exhaustive procedure) guaranteed optimality. In this case, for $n = 1, 2, 3, \dots$ we computed the number of responses and the total R^2 explained by each subset of n variables. The best subset was identified as that which explained the highest number of responses or the maximum total R^2 . In the original study, the stepwise and best-subset procedures identified the same 4 variables, namely offer value A, offer value B, chosen value and chosen juice. Fourth, we conducted a post-hoc analysis. While the explanatory power of variables included in the best subset was (by definition) higher than that of any other subset of variables, the procedure did not guarantee that this inequality was statistically significant. The post-hoc analysis addressed this issue by comparing the marginal explanatory power of each variable in the best subset with that of other, non-selected variables (binomial test). In the original study and subsequent work^{32, 57, 58}, we found that the explanatory power of offer value A, offer value B, chosen value and chosen juice was statistically higher than that of any other variable.

The results of these analyses provided a classification for neuronal responses. Specifically, each neuronal response was assigned to the variable that explained the response and provided the highest R^2 . Thus we identified 447 offer value, 370

chosen value, and 268 chosen juice responses. Subsequent analyses of the same data set demonstrated range adaptation⁶ and quantified noise correlations³⁰.

Previous work suggested that the encoding of value in OFC was close to linear^{4, 6}. In this study, we conducted more detailed analyses to quantify how neuronal responses departed from linearity. Furthermore, we compared the curvature measured in neuronal responses with that measured for the cumulative distribution of the encoded values. The cumulative distribution of encoded values was taken as a benchmark by analogy with sensory systems^{14, 19}. For each offer value response, we identified the value levels present in the session (i.e., the unique values). We then calculated the corresponding number of trials, divided it by the total number of trials in the session, and computed the corresponding cumulative distribution function ($\text{ntrials}_{\text{CDF}}$). For each value level, we also averaged the firing rates obtained across trials. The range of offered values and the range of firing rates varied considerably from session to session and across the population (Supplementary Fig. 3). Thus to compare the results obtained for different responses, we normalized offer value levels, neuronal firing rates, and $\text{ntrials}_{\text{CDF}}$. Value levels were simply divided by the maximum value present in the session. For example, normalized values for juice B were defined as q_B/Q_B . For neuronal firing rates, we performed the linear regression $y = a_0 + a_1 x$, where x are normalized value levels. Firing rates were thus normalized with the transformation $\text{fr} \rightarrow (\text{fr} - a_0)/a_1$. Similarly, for $\text{ntrials}_{\text{CDF}}$, we performed the linear regression $y = b_0 + b_1 x$, and we normalized them with the transformation $\text{ntrials}_{\text{CDF}} \rightarrow (\text{ntrials}_{\text{CDF}} - b_0)/b_1$. Examples of normalized firing rates and normalized $\text{ntrials}_{\text{CDF}}$ are illustrated in Fig. 1d, g. To estimate the overall curvature, we fit each normalized response function with a 2D polynomial and compared the quadratic coefficient (β_2) with that obtained from fitting the corresponding normalized $\text{ntrials}_{\text{CDF}}$. To estimate the overall S shape, we fit the normalized response function with a 3D polynomial and compared the cubic coefficient (β_3) with that obtained for the corresponding normalized $\text{ntrials}_{\text{CDF}}$. These measures were used for the population of offer value responses (Fig. 2). Separately, we repeated these analyses for the population of chosen value responses (Supplementary Fig. 2).

Theoretical considerations indicated that optimal response functions in our experiments would have been step functions (Fig. 7b, c), contrary to our observations. One concern was whether empirical response functions were optimal on average across sessions, if not for any particular session. Notably, the relative value ρ varied from session to session, largely because we used a variety of different juice pairs. Recent work indicates that the same neurons are associated with different juices in different sessions, with remapping dictated by the preference ranking⁵⁸. In any given session, the optimal offer value B response function would have been a step function with step at $x = \rho$. However, since ρ varied from session to session, the resulting optimal response function would have been more gradually increasing. In fact, if the distribution of ρ/Q_B across session had been uniform in the interval [0 1], the mean optimal response for offer value B neurons would have been linear. An important caveat is that the rationale that would justify linear offer value B responses did not hold for offer value A responses. In any case, we examined the distribution of ρ/Q_B (Fig. 7d). For offer value B, the mean optimal response function was computed as the cumulative distribution function for ρ/Q_B .

Importantly, the data sets included in Fig. 2 and Supplementary Fig. 3 were originally selected based on a procedure that only considered linear encoding of value⁴. To assess the functional form of neuronal responses without bias in favor of linearity, we repeated the variable selection procedures described above including in the analyses all the variables defined in this study. The variable selection analysis was still based on linear regressions. However, firing rates were regressed on several non-linear value variables (response functions). The analysis included linear response functions (offer A, offer B), cumulative distribution functions of offer values (offer A $\text{ntrial}_{\text{CDF}}$, offer B $\text{ntrial}_{\text{CDF}}$), variance-equalized versions of the cumulative distribution functions of offer values (offer A $\text{ntrial}_{\text{CDF}} \text{VE}$, offer B $\text{ntrial}_{\text{CDF}} \text{VE}$; see below), optimal response functions (offer A ORF, offer B ORF), mean optimal response functions across sessions (offer A $\text{mean}(\text{ORF})$, offer B $\text{mean}(\text{ORF})$), and optimal response functions obtained under the assumption that the joint distribution of offer values is uniform and symmetric (offer A $\text{ORF}_{\text{uniform}}$, offer B $\text{ORF}_{\text{uniform}}$; see below). In addition, the analysis included chosen value, the cumulative distribution function for chosen values (chosen value $\text{ntrial}_{\text{CDF}}$), a variance-equalized version of the cumulative distribution function for chosen values (chosen value $\text{ntrial}_{\text{CDF}} \text{VE}$), and variables other value, value difference, value ratio, total value and chosen juice (20 variables total).

The last part of Supplementary Note generalizes the theory of optimal coding by relaxing the assumption of linear response functions. In principle, this allows one to compute the optimal response functions for any joint distribution of offers. Optimal response functions for the symmetric uniform distribution ($\text{ORF}_{\text{uniform}}$) were computed numerically in Matlab.

We restricted the variable selection analysis to responses that passed the ANOVA criterion ($N=1379$, see above) and we regressed each neuronal response on each variable. For variance-equalized variables, we first computed the square root of the firing rates and then performed the linear regressions^{59, 60}. For variable selection we used the two procedures described above, and we refer to previous work for additional details⁴. The best-subset method can be based either on the number of responses or on the total R^2 explained by each subset. In previous studies, the two metrics provided similar results. Here we found that the results obtained based on the total R^2 were more robust, probably because the present

analysis aimed at providing a better fit for the neuronal responses as opposed to explaining more responses. The best-subset procedures and post-hoc analyses were performed on collapsed variables⁴. The variable selection analyses were conducted twice. First, we included all the variables described above except those based on ORF_{uniform}. In this case, linear response functions performed significantly better than all the other variables (Supplementary Table 1). Second, we added in the analysis the variables based on ORF_{uniform}. In this case, the performance of ORF_{uniform} variables was better than, but statistically indistinguishable from that of linear variables. It was significantly better than that of all the other variables (Supplementary Table 2). Both analyses are described in the Results. Figures 8 and 9 refer to the analysis that included all 20 variables.

Code availability. The code used for data analysis and simulations is available from the corresponding author upon reasonable request.

Data availability. The data that support the findings of this study are available from the corresponding author upon reasonable request.

Received: 20 January 2017 Accepted: 13 September 2017

Published online: 31 October 2017

References

- Padoa-Schioppa, C. Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* **34**, 333–359 (2011).
- Rushworth, M. F., Kolling, N., Sallet, J. & Mars, R. B. Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr. Opin. Neurobiol.* **22**, 946–955 (2012).
- Wallis, J. D. Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat. Neurosci.* **15**, 13–19 (2012).
- Padoa-Schioppa, C. & Assad, J. A. Neurons in orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
- Padoa-Schioppa, C. Neuronal origins of choice variability in economic decisions. *Neuron* **80**, 1322–1336 (2013).
- Padoa-Schioppa, C. Range-adapting representation of economic value in the orbitofrontal cortex. *J. Neurosci.* **29**, 14004–14014 (2009).
- Kobayashi, S., Pinto de Carvalho, O. & Schultz, W. Adaptation of reward sensitivity in orbitofrontal neurons. *J. Neurosci.* **30**, 534–544 (2010).
- Cox, K. M. & Kable, J. W. BOLD subjective value signals exhibit robust range adaptation. *J. Neurosci.* **34**, 16533–16543 (2014).
- Padoa-Schioppa, C. & Rustichini, A. Rational attention and adaptive coding: a puzzle and a solution. *Am. Econ. Rev. Pap. Proc.* **104**, 507–513 (2014).
- Fairhall, A. L., Lewen, G. D., Bialek, W. & de Ruyter Van Steveninck, R. R. Efficiency and ambiguity in an adaptive neural code. *Nature* **412**, 787–792 (2001).
- Schwartz, O., Hsu, A. & Dayan, P. Space and time in visual context. *Nat. Rev. Neurosci.* **8**, 522–535 (2007).
- Webster, M. A. Adaptation and visual coding. *J. Vis.* **11**, 1–23 (2011).
- Barlow, H. B. Possible principles underlying the transformations of sensory messages. In: Rosenblith, W. A. (ed.) *Sensory Communication* 217–234 (MIT Press, 1961).
- Laughlin, S. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch. C* **36**, 910–912 (1981).
- Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
- Laughlin, S. B. The role of sensory adaptation in the retina. *J. Exp. Biol.* **146**, 39–62 (1989).
- Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W. & Meister, M. Adaptation of retinal processing to image contrast and spatial scale. *Nature* **386**, 69–73 (1997).
- Muller, J. R., Metha, A. B., Krauskopf, J. & Lennie, P. Rapid adaptation in visual cortex to the structure of images. *Science* **285**, 1405–1408 (1999).
- Brenner, N., Bialek, W. & de Ruyter van Steveninck, R. Adaptive rescaling maximizes information transmission. *Neuron* **26**, 695–702 (2000).
- Maravall, M., Petersen, R. S., Fairhall, A. L., Arabzadeh, E. & Diamond, M. E. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol.* **5**, e19 (2007).
- Robinson, B. L. & McAlpine, D. Gain control mechanisms in the auditory pathway. *Curr. Opin. Neurobiol.* **19**, 402–407 (2009).
- Liu, B., Macellario, M. V. & Osborne, L. C. Efficient sensory cortical coding optimizes pursuit eye movements. *Nat. Commun.* **7**, 12759 (2016).
- Dan, Y., Atick, J. J. & Reid, R. C. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* **16**, 3351–3362 (1996).
- Baddeley, R. et al. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. Biol. Sci.* **264**, 1775–1783 (1997).
- Hunt, L. T. et al. Mechanisms underlying cortical activity during value-guided choice. *Nat. Neurosci.* **15**, 470–U169 (2012).
- Kable, J. W. & Glimcher, P. W. The neurobiology of decision: consensus and controversy. *Neuron* **63**, 733–745 (2009).
- Krajcich, I., Armel, C. & Rangel, A. Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* **13**, 1292–1298 (2010).
- Louie, K., LoFaro, T., Webb, R. & Glimcher, P. W. Dynamic divisive normalization predicts time-varying value coding in decision-related circuits. *J. Neurosci.* **34**, 16046–16057 (2014).
- Rustichini, A. & Padoa-Schioppa, C. A neuro-computational model of economic decisions. *J. Neurophysiol.* **114**, 1382–1398 (2015).
- Conen, K. E. & Padoa-Schioppa, C. Neuronal variability in orbitofrontal cortex during economic decisions. *J. Neurophysiol.* **114**, 1367–1381 (2015).
- Haefner, R. M., Gerwinn, S., Macke, J. H. & Bethge, M. Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat. Neurosci.* **16**, 235–242 (2013).
- Padoa-Schioppa, C. & Assad, J. A. The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat. Neurosci.* **11**, 95–102 (2008).
- Grace, R. C. Violations of transitivity: implications for a theory of contextual choice. *J. Exp. Anal. Behav.* **60**, 185–201 (1993).
- Tversky, A. & Simonson, I. Context-dependent preferences. *Manag. Sci.* **39**, 117–185 (1993).
- Bermudez, M. A. & Schultz, W. Reward magnitude coding in primate amygdala neurons. *J. Neurophysiol.* **104**, 3424–3432 (2010).
- Cai, X. & Padoa-Schioppa, C. Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *J. Neurosci.* **32**, 3791–3808 (2012).
- Tobler, P. N., Fiorillo, C. D. & Schultz, W. Adaptive coding of reward value by dopamine neurons. *Science* **307**, 1642–1645 (2005).
- Diederer, K. M., Spencer, T., Vestergaard, M. D., Fletcher, P. C. & Schultz, W. Adaptive prediction error coding in the human midbrain and striatum facilitates behavioral adaptation and learning efficiency. *Neuron* **90**, 1127–1138 (2016).
- Diederer, K. M. J. & Schultz, W. Scaling prediction errors to reward variability benefits error-driven learning in humans. *J. Neurophysiol.* **114**, 1628–1640 (2015).
- Hildebrandt, K. J., Ronacher, B., Hennig, R. M. & Benda, J. A neural mechanism for time-window separation resolves ambiguity of adaptive coding. *PLoS Biol.* **13**, e1002096 (2015).
- Musall, S. et al. Tactile frequency discrimination is enhanced by circumventing neocortical adaptation. *Nat. Neurosci.* **17**, 1567–1573 (2014).
- Tversky, A. & Kahneman, D. The framing of decisions and the psychology of choice. *Science* **211**, 453–458 (1981).
- Savage, L. J. *The Foundations of Statistics* (Dover Publications, 1972).
- Ariely, D., Loewenstein, G. & Prelec, D. 'Coherent arbitrariness': stable demand curves without stable preferences. *Q. J. Econ.* **118**, 73–105 (2003).
- Camille, N., Griffiths, C. A., Vo, K., Fellows, L. K. & Kable, J. W. Ventromedial frontal lobe damage disrupts value maximization in humans. *J. Neurosci.* **31**, 7527–7532 (2011).
- Gallagher, M., McMahan, R. W. & Schoenbaum, G. Orbitofrontal cortex and representation of incentive value in associative learning. *J. Neurosci.* **19**, 6610–6614 (1999).
- Rudebeck, P. H. & Murray, E. A. Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *J. Neurosci.* **31**, 10569–10578 (2011).
- Cisek, P. Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* **22**, 927–936 (2012).
- Friedrich, J. & Lengyel, M. Goal-directed decision making with spiking neurons. *J. Neurosci.* **36**, 1529–1546 (2016).
- Song, H. F., Yang, G. R. & Wang, X. J. Reward-based training of recurrent neural networks for cognitive and value-based tasks. *Elife* **6**, e21492 (2017).
- Zhang, Z., Cheng, Z., Lin, Z., Nie, C. & Yang, T. A neural network framework for the orbitofrontal cortex and model-based reinforcement learning. *bioRxiv* doi 10.1101/116608 (2017).
- Asaad, W. F. & Eskandar, E. N. A flexible software tool for temporally-precise behavioral control in Matlab. *J. Neurosci. Methods* **174**, 245–258 (2008).
- Asaad, W. F. & Eskandar, E. N. Achieving behavioral control with millisecond resolution in a high-level programming environment. *J. Neurosci. Methods* **173**, 235–240 (2008).
- Glaister, P. Least squares revisited. *Math. Gaz.* **85**, 104–107 (2001).
- Dunn, O. J. & Clark, V. *Applied Statistics: Analysis of Variance and Regression* (Wiley, 1987).
- Glantz, S. A. & Slinker, B. K. *Primer of Applied Regression & Analysis of Variance* (McGraw-Hill, Medical Pub. Division, 2001).

57. Cai, X. & Padoa-Schioppa, C. Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron* **81**, 1140–1151 (2014).
58. Xie, J. & Padoa-Schioppa, C. Neuronal remapping and circuit persistence in economic decisions. *Nat. Neurosci.* **19**, 855–861 (2016).
59. Brunel, N. & Nadal, J. P. Mutual information, Fisher information, and population coding. *Neural. Comput.* **10**, 1731–1757 (1998).
60. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural. Comput.* **11**, 91–101 (1999).

Acknowledgements

We thank Heide Schoknecht for help with animal training, Nicolas Brunel and Daniel Chicharro for helpful conversations, and Harold Burton, Arno Onken, and Stefano Panzeri for comments on earlier versions of the manuscript. This work was supported by the National Institutes of Health (grant numbers R01-DA032758 and R01-MH104494 to C.P.-S. and grant numbers T32-GM008151 and F31-MH107111 to K.E.C.) and by the National Science Foundation (grant SES-1357877 to A.R.). This work was partly conducted while C.P.-S. was a visiting fellow at the Italian Institute of Technology.

Author contributions

A.R. and C.P.-S. designed the study; K.E.C., X.C., and C.P.-S. collected the experimental data; K.E.C. and C.P.-S. analyzed the experimental data; A.R. and C.P.-S. developed the mathematical formalism; A.R. and C. P.-S. wrote the manuscript.

Additional information

Supplementary Information accompanies this paper at doi:10.1038/s41467-017-01373-y.

Competing interests: The authors declare no competing financial interest.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2017

Optimal Coding and Neuronal Adaptation
in Economic Decisions
SUPPLEMENTARY NOTE

ALDO RUSTICHINI, KATHERINE E. CONEN, XINYING CAI,
AND CAMILLO PADOA-SCHIOPPA

ABSTRACT. We formulate the general problem of optimal coding for a linear model for the choice with two pools of neurons. Optimal coding is defined as the problem of choosing the optimal response function for a given environment, under physiological constraints (such as a maximum firing rate), if one wants to maximize the expected value of the goods selected.

We consider first the case of a linear response function, and analyze the adjustment of the slope of the response to the distribution of offers in a session. Our aim is to solve two problems at the same time: explain why we observe adaptive coding (*AC*), and why *AC* does not induce a choice bias (that is, a dependence of choice on the environment: this is defined precisely in section 4.1).

We find:

- (1) The optimal response requires adaptive coding (proposition 4.1 and theorem 5.2).
- (2) If the firing rate and inputs have no upper limit, then the choice does not exhibit bias depending on variation of the range of the offers in different experimental sessions. There is no bias, but *AC* improves the response (proposition 4.1).
- (3) If the firing rate and inputs have upper limits, then as the range of either good increases the stochastic choice becomes noisier (equation 26 in theorem 5.2).
- (4) In both cases (upper bound or not) the consistency of choice is not a constraint imposed on the process, but an outcome of the optimization (proposition 4.1 and theorem 5.2).

In the second part we consider the general problem of characterizing the best response function in the set of all functions, not necessarily linear.

1. CHOICE ENVIRONMENT

Two goods, A and B are offered during an experimental sessions. We use the index $g \in \{A, B\}$. In the main text we conducted the analysis in physical goods space: thus, a quantity q_A of good A and q_B of good B are offered, expressed in common physical units (say, ml), and $q_A \sim q_B$ (that is, the animal is indifferent between q_A and q_B) if and only if $\rho q_A = q_B$. This setup is closer to the nature of the neural coding (see Padoa-Schioppa (2009), in particular figure 4 of that paper), but is less natural for the analysis. So

Date: September 3, 2017.

Key words and phrases. Adaptive Coding, Economic Choice, Neuro-computational Theory.

in this supplementary material we work in value space, describing the two offers x for good A and y for good B in terms of a common value; so in this case $x \sim y$ if and only if $x = y$; so that we can take x to be ρq_A , and y to be q_B . The joint distribution of offers in a session is described by a $\pi \in \Delta(X \times Y)$; π_g is the marginal over quantities of good g .

2. LINEAR MODEL: NEURONAL CODING WITHIN SESSIONS

We first describe the response in a trial of a given experimental session. There are two pools of neurons, one for each good g . Each set has the same number n of neurons; we will take the limit as $n \rightarrow \infty$ later. We consider the output of the two pools as providing input current to a post synaptic sets of neurons; so we will consider both the firing rate and the input to the downstream neuron for every spike; the latter is regulated by synaptic efficacy. We assume that all neurons in the pool fire at the same Poisson rate:

Assumption 2.1. *The spike process $r_{g,i}$, $g \in \{A, B\}$ for each neuron $i = 1, \dots, n$ is Poisson with rate ν_g .*

We consider the total firing occurring in a fixed time interval Δt of approximate length 500 ms. There is no leak, thus all the spikes within this interval are interchangeable. We write the input corresponding to each spike of neurons of good g as J_g ; this is the amount of positive charge entering the membrane of the post-synaptic neuron due to one spike. Since Δt does not play a role in our analysis, we choose it as time unit: $\Delta t = 1$. Now the firing rate refers to this time unit, and $r_{g,i}$ indicates the number of spikes per unit time. We define D^n as the random variable obtained as a weighted sum of the contributions of the two pools:

$$(1) \quad X_g^n \equiv \sum_{i=1}^n w_i^n r_{g,i} J_g$$

where w^n is a vector of n -positive weights, adding up to 1; and

$$(2) \quad D^n \equiv X_A^n - X_B^n$$

By the assumption 2.1

$$(3) \quad E(r_{g,i} J_g) = \nu_g J_g, \quad Var(r_{g,i} J_g) = \nu_g J_g^2 \equiv \sigma_{r_g}^2 J_g^2$$

We focus here on the simple case where

$$(4) \quad \forall i, w_i^n \equiv \frac{1}{n}.$$

We want to take limits as $n \rightarrow \infty$, defining in the limit a random variable D . Assuming stationarity (in the spatial dimension index) then by Birkhoff ergodic theorem, a limit random variables D exists, which is the expectation of the representative random variable $r_{A,1} J_A - r_{B,1} J_B$ with respect to the invariant σ -algebra.

The expectation of this limit variable is:

$$(5) \quad E(D) = \nu_A J_A - \nu_B J_B$$

The variance of D^n is

$$(6) \quad \text{Var}(D^n) = \frac{1}{n} (J_A^2 \text{Var}(r_{A,1}) + J_B^2 \text{Var}(r_{B,1})) + \\ \left(1 - \frac{1}{n}\right) (J_A^2 \text{Cov}(r_{A,1} r_{A,2}) + J_B^2 \text{Cov}(r_{B,1} r_{B,2})) - 2J_A J_B \text{Cov}(r_{A,1} r_{B,1})$$

We assume (on the basis of experimental evidence: see figure S4 of the text) that the correlation among neurons is independent of the firing rate ν and the vector of inputs J :

Assumption 2.2. $\text{Corr}(r_{g,1} r_{g,2}) = \chi_g, g = A, B$ and $\text{Corr}(r_{A,1} r_{B,1}) = \gamma$, all independent of the vectors rate ν and input J .

Using 2.2 we get:

$$(7) \quad \text{Var}(D) = \chi_A \sigma_{r_A}^2 J_A^2 + \chi_B \sigma_{r_B}^2 J_B^2 - 2\gamma \sigma_{r_A} J_A \sigma_{r_B} J_B$$

Note that when $\chi_A = \chi_B \equiv \chi$ and $\sigma_{r_A}^2 = \sigma_{r_B}^2 \equiv \sigma^2$ and $J_A = J_B$, then $\chi \geq \gamma$, so correlation within pools is larger than the one across pools. In the analysis below we set $\gamma = 0$ for simplicity of notation; introducing it back is trivial.

We assume that good A is chosen if $D \geq 0$, and B otherwise. This *linear decision model* is a very reduced form of mutual and of pooled inhibition models.

3. ADAPTIVE CODING ACROSS SESSIONS

We consider the problem of the adjustment of the neuronal response from one session to another. Neurons corresponding to a good respond in each trial to the quantity offered of the corresponding good; the function describing this response can vary across sessions, depending on π . We now consider different possible restrictions on this dependence.

3.0.1. *Flexibility of parameters.* First we assume that the input J does not depend on the quantities offered in the trial:

Assumption 3.1. *A choice of response in a session is a pair of functions assigning rate response and input to a distribution, $(\nu_g(\cdot|\pi), J_g(\pi))$, where $\nu_A(x|\pi)$ is the rate associated with the quantity x of A , and so on; $J_g(\pi)$ is a real number.*

We justify this restriction arguing that the synaptic efficacy J can be modified by learning during the session in a way that depends on π (for example through Hebbian learning) but cannot adjust quickly enough to trial by trial variations in the quantities offered.

A second restriction is that:

Assumption 3.2. *For both goods g , the rate function ν_g and input J_g only depend on π_g (and not on the marginal on the other good).*

If this assumption holds, at the beginning of the session the pool of g -neurons is “informed” of π_g , can adjust to it, but its response cannot depend on $\pi_{g'}, g' \neq g$.

We will assume both assumption 3.1 and assumption 3.2 in the following.

4. MAXIMIZATION PROBLEM

Fix π and a choice $(\nu, J) = (\nu_g, J_g)_{g \in \{A, B\}}$. For any given pair of quantities (x, y) we have a probability distribution on the real line, $Q(\cdot | (x, y), \nu, J)$. The problem we consider is the maximization of the expected payoff:

$$(8) \quad \int_{X \times Y} (Q(R^+ | ((x, y), \nu, J))x + Q(R^- | ((x, y), \nu, J))y) d\pi(x, y)$$

where R^+ and R^- are the sets of positive and strictly negative real numbers, and π is the joint distribution defined at the end of section 1. So our problem is:

$$(9) \quad \max_{(\nu, J)} \int_{X \times Y} (Q(R^+ | ((x, y), \nu, J))x + Q(R^- | ((x, y), \nu, J))y) d\pi(x, y)$$

where (ν, J) satisfies the appropriate restrictions, depending on assumptions 3.1 and 3.2. After a simple rearrangement we have that (ν, J) is a solution of 9 if and only if it is a solution of 10:

$$(10) \quad \max_{(\nu, J)} \int_{X \times Y} (Q(R^+ | ((x, y), \nu, J))(x - y) d\pi(x, y)$$

From our analysis in section 2, using the formula for expectation and variance of D we obtain the crucial information that the mean and variance of the limit random variable scale as in equations 5 and 7. We have then:

$$(11) \quad Q(R^+ | ((x, y), \nu, J)) = \Pr \left(Z \geq - \frac{\nu_A(x|\pi)J_A(\pi) - \nu_B(y|\pi)J_B(\pi)}{(\chi_A\nu_A(x|\pi)J_A^2(\pi) + \chi_B\nu_B(x|\pi)J_B^2(\pi))^{1/2}} \right)$$

where Z is a standardized random variable with mean 0 and variance 1. We assume for computational convenience, particularly in the numerical simulations, that the random variables D and Z are normal. A more detailed description on D (beyond the information on mean and variance, which are the key elements in our analysis below) requires a full model of the the network, that will give account of the long distance correlation among neurons. This full model, allowing for long distance positive correlation (and not only asynchronous state, as in Renart et al. (2010) or Rosenbaum et al. (2017)), is topic of current research.

Noting that the lower bound in the equation above is homogeneous of degree zero in J , we reformulate the problem as maximization in ν and $R \equiv \frac{J_B}{J_A}$ where the probability in equation 11 has the form:

$$(12) \quad Pr \left(Z \geq -\frac{\nu_A(x|\pi) - \nu_B(y|\pi)R(\pi)}{(\chi_A\nu_A(x|\pi) + \chi_B\nu_B(x|\pi)R^2(\pi))^{1/2}} \right).$$

4.1. Choice bias. Our analysis will provide a model for the probability that the animal chooses good A if the pair offered is (x, y) and the probability on goods in the session is π , call it $P(Ch = A|x, y, \pi)$. We say that there is no choice bias if two conditions are satisfied.

The first (equation 13) is that this probability only depends on the pair (x, y) and not on π (as long as, of course, the pair (x, y) is in the support of the distribution).

$$(13) \quad \forall \pi, \pi', x, y : P(Ch = A|x, y, \pi) = P(Ch = A|x, y, \pi')$$

The condition 13 is not restrictive enough. For example, choosing A in all trials of all sessions satisfies 13, so we add:

$$(14) \quad \forall \pi : P(Ch = A|x, y, \pi) \text{ is increasing in } x, \text{ decreasing in } y \text{ and equal to } 1/2 \text{ if } x = y.$$

4.2. Functional restrictions. So far we have constrained the response functions ν only through their dependence on π . We now consider the additional restriction of linearity:

$$(15) \quad \nu_A(x|\pi) = s_A(\pi)x, \nu_B(y|\pi) = s_B(\pi)y$$

where $s_g(\pi) \in R^+$. The constraint 15 is very strong; we discuss its weakening in section 6 below.

We define the largest value in the support of the distribution of a good: $M_g(\pi) \equiv \max\{z \in \text{support}(\pi_g)\}$; we may drop the π in the notation if this dependence is clear. Consistently with experimental evidence, we assume that $\chi_A = \chi_B = \chi$. With the linearity assumption 15 we say that there is full linear adaptive coding if there is a single $s > 0$ such that $s_g = \frac{s}{M_g}$ for all g . In this case if $M_A = M_B$ then $\nu_A(x|\pi) = sx, \nu_B(y|\pi) = sy$, and the formula 11 above becomes:

$$(16) \quad Q(R^+ | ((x, y), s, J)) = Pr(Ch = A|x, s, J) = Pr \left(Z \geq -\frac{J_A x s_A - J_B y s_B}{\sqrt{\chi(J_A^2 x s_A + J_B^2 y s_B)}} \mid Z \sim N(0, 1) \right).$$

which is equation (2) in the main text. Recall that x and y are value units, so firing rate and synaptic efficacy for each good change accordingly. Setting the correlation χ notionally equal to 0 clarifies the scaling in both cases. When using physical units, we want the effective firing rate (firing rate t_g times synaptic efficacy (call it K_g) to be equal when the animal is indifferent between the two offers, that is $t_A K_{AQA} = t_B K_{BQB}$ if and only if

$\rho q_A = q_B$. This holds if and only if $\frac{t_A}{\rho} K_A = t_B K_B$. In value units we want the effective firing rate to be equal ($s_A J_A x = s_B J_B y$) if and only if $x \sim y$, that is $x = y$, and this occurs if and only if $s_A J_A = s_B J_B$. Thus, the pair (x, y) here corresponds to $(\rho q_A, q_B)$ in the main text; the pair (J_A, J_B) here to (K_A, K_B) in the main text, the pair (s_A, s_B) to the pair $(\frac{t_A}{\rho}, t_B)$, and finally the pair (M_A, M_B) to the pair $(\rho Q_A, Q_B)$.

Proposition 4.1. *If the maximization problem satisfies equation 15 and there is no constraint on the firing rate, then with full linear adaptive coding:*

$$(17) \quad s_g \equiv \frac{s}{M_g}, J_g = M_g$$

- (1) *As $s \rightarrow \infty$, the expected payoff tends to the maximum possible value given by:*

$$\int_{X \times Y} \max\{x, y\} d\pi(x, y),$$

so there is no choice bias (that is, both equations 13 and 14 hold).

- (2) *If $J_A = J_B$ the expected payoff is strictly increasing in \sqrt{s} .*

Proof. Under our assumption and functional form of the response the quantities in 11 are equal to

$$(18) \quad Pr(Ch = A|x, y, \pi) = Pr\left(Z \geq -\left(\frac{s}{\chi(xM_A + yM_B)}\right)^{1/2} (x - y)\right)$$

so as $s \rightarrow \infty$ the probability of choosing $\max\{x, y\}$ tends to 1 for all (x, y) , independently of π .

The statement (2) follows differentiating the expected payoff with respect to \sqrt{s} ; the derivative is the integral of a positive function. \square

The restriction $J_A = J_B$ is needed here because if $\frac{J_B}{J_A} > 1$, say, and π is concentrated on the sector of (x, y) such that $\frac{J_B}{J_A} y > x > y$ then the derivative with respect to \sqrt{s} may be negative (as intuitively clear: increasing \sqrt{s} reduces the probability of the A choice (because $J_A x - J_B y < 0$), when $x > y$.)

4.3. Re-scaling of χ . We clarify now the details linking the current analysis to the data, to justify setting $\chi = \xi/4$ as in the main text, immediately below equation (2). There are several different re-scalings that we use in data analysis; they can all be conveniently represented by a multiplication of the correlation coefficient χ by a constant c . There are two re-scalings in particular that need attention.

The first re-scaling is produced by the fact that we use in the analysis the time unit of the decision time (called Δt in section 2). The crucial point here is that the variable D^n defined in 2 considers the spikes in the time interval Δt , and **not** in 1 second units. This is convenient because we can ignore this Δt . Suppose this decision time is 0.5 seconds. Then if we use

the firing rate in seconds (Hz), we have to adjust to compensate for the different time units. For example if the firing rate in decision time units is 8, then the one in Hz is 16. If we replace the firing rate in Δt units with the one in units of seconds, we have a different formula, where the correlation coefficient is multiplied by $c = 2$, or in general by a factor $\frac{1}{\Delta t}$.

The second re-scaling is used in the numerical analysis of the linear response model. Here we take $\nu_A = s_A x$, for example. We constrain the slope s_A to reflect the constraint that the firing rate cannot be larger than a maximum $\bar{\nu}$ for any quantity x , so we require

$$s_A M_A \leq \bar{\nu}.$$

In our numerical simulations reported below we formulate the constraint for convenience not $s_A \leq \frac{\bar{\nu}}{M_A}$ but as $s_A \leq 1$. Again we can do this but we need to re-scale the firing rate, because we are implicitly defining the new slope as $s_A \frac{M_A}{\bar{\nu}}$, and we need to introduce an adjustment by introducing a scaling factor c . In the numerical analysis the M_A is given by the maximum number of units of the good presented; the maximum rate $\bar{\nu}$ has to be estimated by looking at the difference between the largest firing rate we observe and the baseline firing rate, which in Hz is approximately 8 (see Padoa-Schioppa (2013)). The multiplication by 2 and the division by 8 in the two re-scalings explain the correction $\chi = \xi/4$.

5. CONSTRAINTS ON THE RESPONSES

Of course the limit in $s \rightarrow \infty$ is biologically meaningless. In this section we reconsider the maximization problem presented in section 4, but with the assumption that the responses are linear, and introducing explicitly the constraints on the responses. So we solve:

$$(19) \quad \max_{(s_A, J_A, s_B, J_B)} \int_{X \times Y} \left(\int_{-L}^{\infty} \phi(z) dz \right) (x - y) d\pi(x, y)$$

(where

$$(20) \quad L = \frac{s_A J_A x - s_B J_B y}{(\chi(s_A J_A^2 x + s_B J_B^2 y))^{1/2}}$$

ϕ is the density of the standard normal), subject to:

$$(21) \quad \forall g, 0 \leq s_g \leq \frac{\bar{\nu}}{M_g}.$$

The constraint 21 insures that the firing rate at all quantities offered of the good is smaller than a maximum firing rate $\bar{\nu}$. The following also may seem natural:

$$(22) \quad \forall g, 0 \leq J_g \leq \bar{J}$$

but as we noted the value of L in equation 20 is homogeneous of degree zero in the J inputs; so re-scaling them by a common positive factor does not affect choice. The lower integration L can be rewritten as:

$$(23) \quad L = \frac{s_A x - s_B R y}{(\chi(s_A x + s_B R^2 y))^{1/2}}$$

where $R \equiv \frac{J_B}{J_A}$, constrained only by:

$$(24) \quad R \geq 0.$$

The original problem is equivalent to maximize 19, with L as in 23, subject to the two constraints 21 and 24. Once an optimal triple (s_A, s_B, R) has been determined, pairs of 22-feasible inputs (J_A, J_B) can be found, the optimal J 's are indeterminate.

5.1. Symmetric distribution. We begin with a simple case, the symmetric uniform distribution on offers.

Theorem 5.1. *Let $M_A = M_B \equiv M$ and π be the product of the uniform distribution on $[0, M]^2$. The optimal linear response is:*

$$(25) \quad (s_A, s_B, R) = \left(\frac{\bar{v}}{M}, \frac{\bar{v}}{M}, 1 \right)$$

The probability of choice of good A with offers (x, y) at the optimal solution is

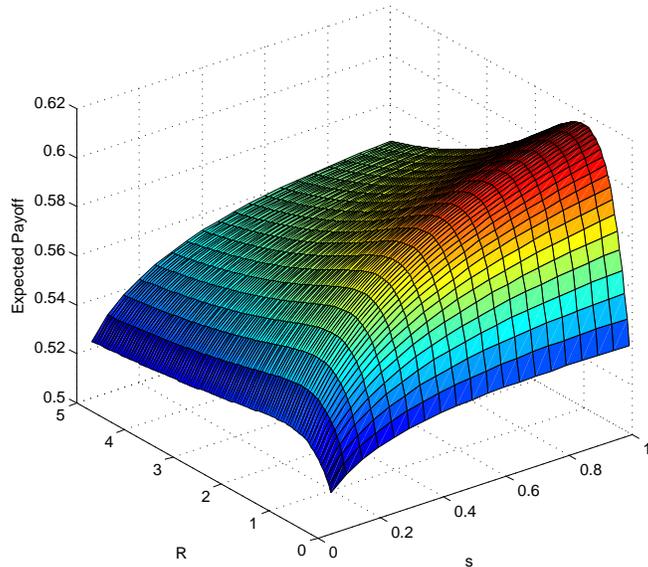
$$(26) \quad Pr(Ch = A|x, y, \pi) = Pr \left(Z \geq - \left(\frac{\bar{v}}{\chi M(x+y)} \right)^{1/2} (x-y) \right)$$

Proof. Write the Lagrangean and solve. □

At $s_A = s_B \equiv s$ the values of the net payoff (maximand in 19) for different values of s and R is displayed in Supplementary Note Figure 1. The figure confirms that (for $s = 1$) the local maximum at $R = 1$ is a global maximum.

5.2. Non-symmetric, uniform distribution. We consider here the more general case of a uniform distribution but on two possibly different ranges. To understand the shape of the optimal solution in this case it is first useful to consider the basic trade-offs in the problem. To fix ideas we consider the case in which $M_B > M_A$.

If we consider the lower integration term L in equation 20, we note that the mean scales in the firing rate slope s_A and s_B linearly, whereas the standard deviation (the denominator in L) scales as the square root; hence we may want to make the slopes as large as possible. The slopes are bounded by a maximum, and the maximum slope is decreasing in the range (equal to \bar{v}/M_g); hence once we set them to a maximum we are introducing a potential bias in the choice, making the probability of choosing the good B smaller than it should be. The J inputs can correct this; as we noted their scale is irrelevant (both mean and SD scale linearly in the inputs) so we can correct



SUPPLEMENTARY NOTE FIGURE 1. **Surface of the expected payoff as function of (s, R) .** Value of s satisfy the constraint in (21); R is unconstrained above but truncated at $R \leq 7$. Here we set $s_A = s_B \equiv s$, $M_A = M_B = 1$, $\bar{\nu} = 1$, $\chi = 0.02$. The optimal solution is at $(s, R) = (1, 1)$.

entirely the bias, setting $R = M_B/M_A$. This describes an “almost optimal” solution, depending on the value of χ . However, this is not necessarily the optimal solution, and this depends on the value of the correlation χ .

To understand why the proposed solution is not precisely optimal, and the role of the correlation χ , one can compare this solution to the policy of “always choosing B ”. This policy chooses the wrong option (hence possibly inducing a loss compared to the “no-bias” policy) only when $x > y$, which, as M_B becomes large compared to M_A , is unlikely; on the rest of the offers region it chooses the right option for sure (hence a gain, compared to the “no-bias” policy). When the noise is large, the gain overcomes the loss. The true optimal solution adopts some of the “always choosing B ”, introducing a small bias in favor of the option B , by making $R > \frac{M_B}{M_A}$. This is better than the no bias solution, because in this way we balance the small cost of the bias in the triangular region between $y = x$ and $y = R \frac{M_B}{M_A} x$, and the gain of increasing the probability of the B choice over the *entire region* $y > x$.

Theorem 5.2. *Let π be the product of the uniform distribution on $[0, M_A] \times [0, M_B]$. The optimal linear response is:*

$$(27) \quad (s_A, s_B) = \left(\frac{\bar{v}}{M_A}, \frac{\bar{v}}{M_B} \right)$$

If we let $m \equiv \frac{M_B}{M_A}$, we note that the ratio $\frac{R}{m}$ has a natural interpretation of a bias. To see this, consider that a bias is absent only if the probability of choice of A is equal to $1/2$ when $x = y$; that is when the lower extreme of integration is $L = 0$; this in turn is true if and only if $x = \frac{R}{m}y$. So for the probability of choosing A over B to be exactly $1/2$ when $x = y$ it is necessary that $\frac{R}{m} = 1$.

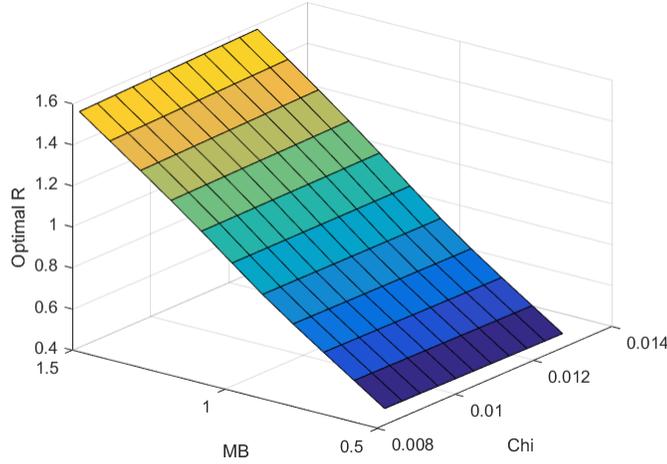
In table 1 we report the bias in favor of the good B (a positive number indicates that the good B is more likely to be chosen when $x = y$) at the optimal values of R . The dependence of the optimal R on χ and M_B is illustrated in Supplementary Note Figure 2 for a range of χ closer to the empirically relevant value or larger. The optimal R is 1 at $M_B = M_A$, $R > M_B/M_A$ for $M_B > M_A$, and $R < M_B/M_A$ for $M_B < M_A$. As χ becomes larger, the optimal R deviates more from M_B/M_A .

TABLE 1. **Bias as function of ρ and M_B .** In the table $M_A = 1$, $\bar{s} = 1$. The entries denote the bias in favor of the B good, that is the value $\frac{R}{M_B} - 1$ (see the remarks after the statement of the theorem 5.2). The values of χ (ranging from 0.008 to 0.013) are indicated in the top row; the values of M_B (ranging from 0.50 to 1.50) on the left column.

	0.008	0.009	0.009	0.010	0.010	0.010	0.011	0.011	0.012	0.012	0.013
0.5	-0.04	-0.04	-0.04	-0.04	-0.04	-0.05	-0.05	-0.05	-0.05	-0.05	-0.06
0.6	-0.03	-0.03	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.05	-0.05	-0.05
0.7	-0.03	-0.03	-0.03	-0.03	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04
0.8	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.04	-0.04
0.9	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02	-0.02
1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.1	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
1.2	0.02	0.02	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
1.3	0.03	0.03	0.03	0.03	0.03	0.03	0.04	0.04	0.04	0.04	0.04
1.4	0.03	0.03	0.03	0.03	0.04	0.04	0.04	0.04	0.04	0.04	0.05
1.5	0.03	0.03	0.03	0.04	0.04	0.04	0.04	0.04	0.05	0.05	0.05

5.3. Stochastic Choice. As ranges change, the slope of the stochastic choice function changes. In the symmetric case ($M_A = M_B$), from equation 26, we have that (here $\pi = 3.1415\dots$, not the distribution over goods):

$$(28) \quad \text{slope of the sigmoid at } x = y \text{ is } \left(\frac{\bar{v}}{2\pi\chi x(M_A + M_B)} \right)^{1/2}.$$



SUPPLEMENTARY NOTE FIGURE 2. **Surface of the optimal R as function of M_B and χ .** Here $M_A = 1, \bar{\nu} = 1$.

We write explicitly M_A and M_B for ease of comparison with the non symmetric case.

In the non symmetric case we can compute the slope on the line in the $x - y$ space where the probability of choice is equal to $1/2$, that is on the line where

$$(29) \quad J_{AS}Ax - J_{BS}By = 0.$$

because this is where we are going to compute the slope in the empirical estimates. Due to the bias which is evident in Supplementary Note Figure 2, this is not the line $y = x$; the function giving the value y indifferent to x as a function of x is flatter when $M_B > M_A$. In this case:

$$(30) \quad \text{slope at } (x, y) \text{ where (29) holds is } \left(\frac{\bar{\nu}}{2\pi\chi x(1+R)M_A} \right)^{1/2}.$$

Note that by our results on the optimal R (see remarks after theorem (5.2), the slope is less than $\left(\frac{\bar{\nu}}{2\pi\chi x(M_A+M_B)} \right)^{1/2}$ when $M_B > M_A$. The extent of the flattening of the slope is provided by the equation 30, considering the values of R given in Supplementary Note Figure 2.

This estimate of the slope provides an easy test of the model through the dependence on the ranges of observed choices.

5.4. Payoff and Range. We now examine how the expected payoff changes when the ranges of the two goods change.

5.4.1. *Re-scaling the range.* In the symmetric case, consider first the effect on the expected payoff of increasing the range, while keeping the symmetry property. The expected payoff is homogeneous of degree one in the range: if we write the expected payoff at the optimal solution in the range $[0, TM]^2$ as $Pay(TM)$ then $Pay(TM) = TPay(M)$. Thus, doubling the range doubles the payoff. This is clear from a change of variable, considering the formula for the lower term of integration L in equation (20). Consequently also the ratio between the expected payoff at the optimal solution and the maximum payoff does not change as the range is re-scaled. This is clear by the fact that the maximum payoff too is homogeneous of degree 1 in the scaling factor T . Note that the slope of the response decreases when the range increases.

5.4.2. *Conditional Payoff.* Instead, the conditional payoff decreases as the slope decreases. To be precise, let $M \equiv M_A = M_B$, and re-scale M by a factor $T \geq 1$. For any T we can compute the conditional expected payoff from the choices made when x and y are in the range $[0, M]^2$, but with the slope optimal for the range $[0, TM]^2$. The conditional expected payoff is strictly decreasing in T for $T > 1$; as one can check the derivative with respect to T is given by

$$-\frac{1}{2T^{3/2}M^2} \int_{[0,M]^2} \phi(-L) \frac{(x-y)^2}{(\chi(x+y))^{1/2}} dx dy.$$

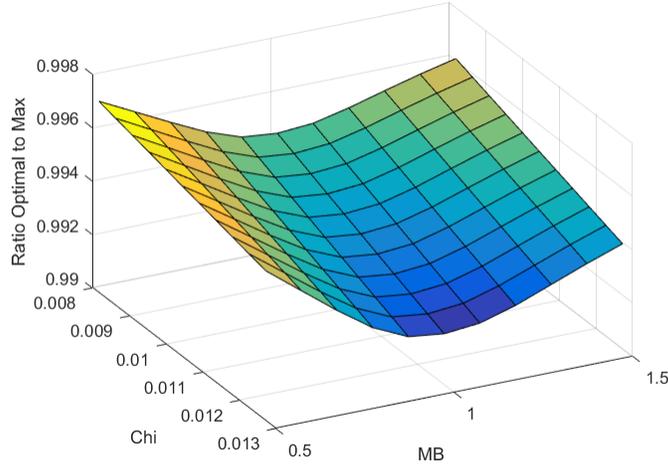
hence strictly negative. The payoff as $T \rightarrow +\infty$ tends of course to the expected payoff from the random choice with equal probability on each option.

5.4.3. *Comparison with the maximum payoff.* Supplementary Note Figure 3 displays the ratio of the expected payoff at the optimal choice of slopes and R ratio over the expectation of the maximum payoff (obtained by choosing the maximum of x and y at every point). Clearly the worst situation for the optimal choice is the symmetric case, where the payoff is smallest for any fixed value of χ . Note that for realistic values of χ close to the empirically relevant value the deviations are very small. This is natural, as the payoff from the constant policy “always choose the option with the largest range” tends to the maximum payoff as the ratio between the two ranges increases to infinity, and the optimal policy becomes close to the constant policy as the range of one of the goods increases. As natural, the ratio decreases with the correlation χ .

6. NON-LINEAR RESPONSE FUNCTIONS

If we take the point of view of section 3 that the slope is chosen to maximize expected payoff it is natural to enquire what the best response would be if the variable is the function instead of the slope, with no other constraint than the bound of the firing rate.

A similar optimization criterion underlies the alternative Barlow-Simon (BS) hypothesis, that the response function is the cumulative distribution function of probability distribution of the stimulus in the environment. It is



SUPPLEMENTARY NOTE FIGURE 3. **Ratio of the expected payoff at the optimal slope and optimal R over the maximum payoff.** Here $M_A = 1, \bar{\nu} = 1$.

well known that this response function is the optimal solution of a clearly defined information theoretic problem, namely maximizing the entropy of the distribution of the output signal over the response functions. The statement follows from two observations. First, the unconstrained maximum entropy over the output space (that is, the maximum entropy over *all* distributions) is achieved by the uniform distribution. Second, the uniform distribution over the output space can be achieved precisely, by setting the response function to be the *cdf* of the signal (when the density is no-where zero). If the objective function is not the entropy but the expected payoff the answer may well be different, and for a good reason, as we will see (section 6.1). The problem we consider is the one presented in section 4, with no restriction on the ν_g response functions.

6.1. Intuitive Properties of the Optimal Solution. To understand the fundamental feature of the optimal solution we can consider a simple example. The good A is offered in a fixed quantity of $1/2$. The good B has offers in quantity $y \in \{0, a, 1\}$ with probabilities p_1, p_2 and p_3 . To recover the case of a no-where zero density we can if necessary combine this with a uniform distribution over the unit interval. A function f from the space Y of quantities of B and the unit interval in the signal space Z can be chosen. A signal equal to $f(y) + \epsilon$ is observed, where $\epsilon \sim N(0, \sigma_\epsilon^2)$. The optimal choice policy is clearly described by a threshold value z_{th} and the choice of B when the observed signal is larger than the threshold. The optimal f is extreme valued. First, it clearly sets $f(0) = 0, f(1) = 1$. If we think of σ_ϵ^2 as small, it is clear that the optimal $f(a)$ is 0 when $a < 1/2$ (so the probability of the

signal crossing the threshold is minimal) and symmetrically 1 when $a > 1/2$. Now, adding several intermediate points a_i between 0 and 1 will not change the general S -shaped feature: all values below $1/2$ will be mapped to 0, and those above to 1; in general, low values and high values are lumped into similar signals, and the optimal function f is S -shaped. The BS solution is very different: the function f has three sharp increases at each of the three values of y , independently of the relative position of a and $1/2$.

The example illustrates the fundamental difference between the BS problem and the problem stated in section 4. Even if one does not accept the criterion we suggested in section 4, the BS solution has some revealing and counterintuitive features. First, the BS solution is only trying to maximize the sharpness of the local resolution, and is not interested in the global property of the distribution. This is obvious if one considers that the BS solution consists in setting the derivative of the response function equal to the density, a completely local definition (the value of the density at one point defines completely the derivative of the firing rate function at that point). Second, the BS is defined *only* by the properties of the environment and is not interested in the properties of the system that is making the choice, as none of the properties of this system enters into the solution. For instance it does not consider the signal to noise ratio, or the correlation coefficient χ which is crucial instead to define the optimal solution to the problem stated in section 4. These are obvious consequences of the fact that there is no *choice* problem this policy is solving.

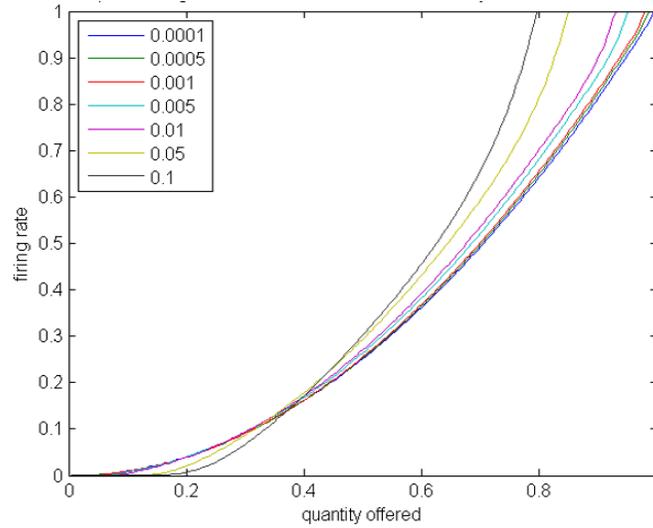
In summary, it should be clear that the BS solution is irrelevant in the context of economic choice problems.

6.2. Results. Two main ideas describe the optimal response. One can isolate the first by considering first the symmetric case, where the distribution on the product of space of the two goods is the product of the marginals, and then proceeding to consider the non symmetric case in which the two marginals are not the same.

6.2.1. Symmetric distributions. By symmetry, the optimal solution will have $\nu_A = \nu_B \equiv \nu$ and $R = 1$. The optimal ν is displayed in Supplementary Note Figure 4, for different values of the correlation coefficient. The best response function has several main properties.

It is non-linear, S -shaped, with two almost flat regions at the low and high level of offers, and an intermediate curve approximately convex. The function depends on the correlation χ , becoming steeper in the intermediate ranges as χ increases. For low values of χ the flat regions are absent, and the function is monotonically strictly increasing, convex. Supplementary Note Figure 4 illustrates these two properties.

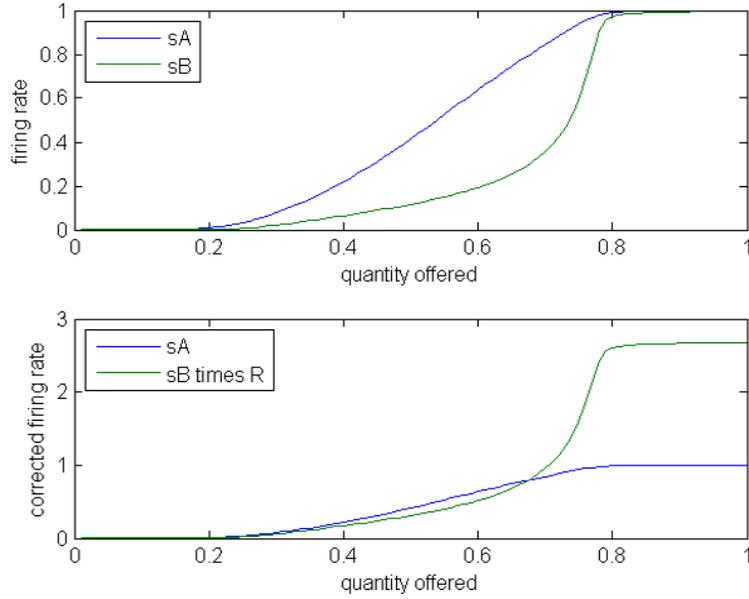
6.2.2. Non-Symmetric distributions. When the distribution on the product space is not the product of two identical distributions the value of R may be different from 1. As we have seen in the case of the linear response function,



SUPPLEMENTARY NOTE FIGURE 4. **Optimal response function and correlation coefficient.** Here $M_A = M_B = 1$, $\bar{\nu} = 1$; the values of the correlation coefficient χ are indicated in the window. Grid of 200 values in the offer space.

the difference $s_A J_A - s_B J_B$ is a measure of the bias, in favor of good A if the difference is positive.

For small values of χ and of the difference between the two distributions, the optimal solution is of course close to the solution of a symmetric problem. For larger values of both variables, R is substantially different from 1, but the optimal firing rates and optimal R are adjusted so that in the region of offers $[0, 1] \times [0, 1]$ the bias is minimal. Supplementary Note Figure 5 illustrates.



SUPPLEMENTARY NOTE FIGURE 5. **Best response functions, asymmetric case.** $M_A = 1, M_B = 1.5, \chi = 0.1$. Top panel: best responses; the range of values of the quantity offered is limited to the interval between 0 and 1 (the range for good A) to make the comparison between the two best responses easier. Bottom panel: the s_B function is multiplied by the optimal $R = 2.682$.

SUPPLEMENTARY REFERENCES

- PADOA-SCHIOPPA, C. (2009): “Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex,” *Journal of Neuroscience*, 29, 14004–14014.
- (2013): “Neuronal Origins of Choice Variability in Economic Decisions,” *Neuron*, 80, 1322 – 1336.
- RENART, A., J. DE LA ROCHA, P. BARTHO, L. HOLLENDER, N. PARGA, A. REYES, AND K. D. HARRIS (2010): “The Asynchronous State in Cortical Circuits,” *Science*, 327, 587–590.
- ROSENBAUM, R., M. SMITH, A. KOHN, J. RUBIN, AND B. DOIRON (2017): “The spatial structure of correlated neuronal variability,” *Nature Neuroscience*, 20, 107–114.