Published in final edited form as:

Am Econ Rev. 2014 May ; 104(5): 507–513. doi:10.1257/aer.104.5.507.

Rational Attention and Adaptive Coding: A Puzzle and a Solution[†]

Camillo Padoa-Schioppa and

Department of Anatomy and Neurobiology, Washington University in St. Louis, 660 South Euclid Avenue, St. Louis, MO 63110

Aldo Rustichini

Department of Economics, University of Minnesota, 1925 4th Street South 4-101, Hanson Hall, Minneapolis, MN 55455-0462

The theories of rational attention and adaptive coding describe different but related adaptive behaviors. Both models consider a state of nature which is determined randomly, according to a probability that we call the environment.

The two theories describe how the response of a decision maker adapts optimally to the environment, to maximize discrimination among the signals that are likely to be received. The theory of rational attention is developed in economics (Sims 1998, 2003, with the name of “rational inattention”), and describes the problem of a decision maker who cannot observe precisely the realized state of nature, has a utility that depends on the action taken and the state realized, and has to choose an action on the basis of the available information. To maximize utility he wants to introduce correlation between action and state, so that the appropriate action is more likely to be taken in every state. But he faces a cost, which is higher if he chooses a higher correlation. In his choice of correlation he takes into account the statistical properties of the environment, and thus he adapts to it. Similarly, adaptive coding describes the dynamic adaptation of the response of neurons to the environment. It is a widespread property of the brain, and appears in several systems including the visual system (Laughlin 1981) as well as in brain regions encoding reward prediction errors, subjective values, and choices (Tremblay and Schultz 1999; Tobler, Fiorillo, and Schultz 2005; Padoa-Schioppa 2009; Kobayashi et al. 2010; Louie, Grattan, and Glimcher 2011; Cai and Padoa-Schioppa 2012). Its effect is an improved discrimination among signals under the constraint that the dynamic range of its response has to be non-negative and smaller than some maximum value. It can thus be considered the neural correspondent of rational attention. In both cases, the decision process is considered as an information processing problem. Differently from rational attention theory, adaptive coding uses the insights and accepts the discipline provided by our knowledge of the functioning of the neural process underlying decision.

[†] Go to <http://dx.doi.org/10.1257/aer.104.5.507> to visit the article page for additional materials and author disclosure statement(s).

Correspondence to: Aldo Rustichini.

There is strong evidence supporting existence of widespread adaptive coding. Adaptive coding, however, introduces the potential of a dependence of choice on the environment: If choice between two goods is taken by comparing the firing rates of different neuronal populations, and if firing rates adapt to the environment to increase the discriminatory power of the responses, then the neural coding of an option depends on the statistical properties of the environment. So the probability of a choice over two given options could be altered by changing the environment, leaving the two options unchanged. Experimental evidence shows that this bias does not occur, thus posing the question of how the potential bias is eliminated. In this paper we define formally how the potential bias may occur, and how it can be corrected.

I. Adaptive Coding: Experimental Evidence

The first experimental evidence of adaptive coding in the neuronal representation of value can be found in Tremblay and Schultz (1999). In their experiment, monkeys were delivered one of three possible rewards (A, B, and C in decreasing order of preference). Trials were run in blocks and two rewards were used in each block (A:B block, B:C block). The response of neurons in the orbitofrontal cortex (OFC) to a given reward varied from block to block. In particular, neuronal responses to B were higher in the B:C block compared to the A:B block.

A quantitative account of adaptive coding is provided by Padoa-Schioppa (2009), who analyzed neuronal data collected during simple economic choices. In the experiment (Padoa-Schioppa and Assad 2006), monkeys chose between x drops of juice A and y drops of juice B . In a given session, x was chosen in the range $[0, A]$, and $y \in [0, B]$. A and B varied from session to session. A population of neurons in the OFC (called offer value neurons) encoded the value of individual juices (A or B) in a linear way; that is, there was a linear relationship between the activity of the cell and the quantity of juice J offered in that trial. Padoa-Schioppa (2009) showed that this representation of value is range adapting. Specifically, the activity of offer value cells encoding the value of juice J can be written as:

$$\phi = \phi_0 + \Delta\phi(V/\Delta V), \quad (1)$$

where ϕ is the firing rate of the cell in the trial, ϕ_0 is a baseline activity, $\Delta\phi$ is an activity range, V is the quantity of J offered in the trial, and ΔV is the range of values of J offered in the session. ϕ_0 and $\Delta\phi$ do not depend on V or ΔV .

A. A Puzzle

Range adaptation poses the following problem. If decisions between juices A and B were made by comparing the firing rates of two neuronal populations encoding the offer value of juice A and juice B , respectively, then the indifference function (describing how much of juice A is needed for indifference with a quantity of juice B) would depend on the range of values offered in any given condition. Specifically, if we assume that the indifference function is linear with relative value ρ such that $P(1, \rho) = 1/2$, decisions made by comparing

firing rates implies that ρ is proportional to B/A . This point is illustrated graphically in Figure 1.

Conen, Cai, and Padoa-Schioppa¹ tested this prediction as follows. In each session, monkeys chose between two juices A and B . Each session included two blocks of trials. In the first block, juices A and B were offered in ranges A and B , respectively. In the second block, one value range was left unchanged while the other value range (A or B) was changed by a factor 2 (halved or doubled, in a 2×2 design).² In the analysis, the authors assumed linear indifference functions and compared the relative value measured in the block where $B/A = X$ with that measured in the block where $B/A = 2 \times X$. Across more than 100 sessions, the relative values measured in the two blocks were statistically indistinguishable. This result suggests that decisions are not simply made by comparing firing rates.

B. Possible Solutions

One possible solution to the challenge posed by range adaptation might be if both offer value A and offer value B cells adapt to the same range of values. In particular, both groups of neurons could be adapting to the maximum value range such that the right-hand side of equation (1) is: $\phi_0 + \phi(V/\max(V))$. We tested this hypothesis as follows. For each offer value cell, we refer to the juice encoded by the neurons as juice E and for the other juice as juice O . In the dataset of Padoa-Schioppa (2009), there are 937 neuronal responses encoding the offer value. Of these, 498 were cases in which the juice with maximum value range was juice E (i.e., $E > O$), while 439 were cases in which the juice with maximum value range was juice O (i.e., $E < O$). (Value ranges were compared taking into account the relative value of the juices.) We examined these two groups of responses separately. For each response we calculated the activity range ϕ . The hypothesis under consideration leads to the prediction that the measure of ϕ should be larger when $E > O$ compared to when $E < O$. Contrary to this prediction, we found no statistical difference between the activity range measured for these two populations ($p = 0.37$, Kruskal-Wallis test). Note that the two groups of cells are large, so it appears unlikely that this result simply reflects a lack of statistical power.

A second alternative solution to the challenge posed by range adaptation would be if the number of cells encoding the offer value of one particular juice depended on the value range for that juice. Specifically, one could imagine that more neurons are added to the population encoding the value of one particular juice when the range of that juice is increased. This increase could in principle balance the effect of range adaptation to keep behavioral preferences stable. To test this hypothesis we considered again the dataset of Padoa-Schioppa (2009). Recordings were generally obtained from multiple neurons in parallel. Offer value responses (937 in total) were recorded in 240 sessions. For each session, we determined the juice with the maximum value range. (Value ranges were compared taking into account the relative value of the juices.) We also counted the number of responses

¹Conen, Cai, and Padoa-Schioppa (unpublished observations in preparation).

²Importantly, the experiment was designed to avoid confounding factors due to juice-specific satiation or choice hysteresis. The initial 50 trials at the beginning of the second block were excluded from the analysis (previous data indicate that this is amply sufficient to obtain full range adaptation).

encoding the offer value of A or B . Across the population, we constructed a 2×2 contingency table representing the number of responses encoding the offer value of A and B (rows) recorded in sessions in which $A > B$ or $A < B$ (columns). Contrary to the hypothesis under consideration, we found that the two classifications were statistical independent ($p = 0.51$, chi-square test).

To conclude, our analysis indicates that decisions are not simply made by comparing the firing rates of different groups of offer value cells. To the contrary, the neuronal network that generates the decision must essentially “undo” range adaptation in such a way that indifference functions do not depend strongly on the value range.

II. Model of Choice

To formulate precisely how choices depend on the firing rate of the offer neurons, we present a model of choice. The heart of the model is a system of equations describing time evolution of the gating variables (fraction of NMDA (N-methyl-D-aspartate) receptors that are open at time t , see Wong and Wang 2006 for details).

We use here the reduced form model,³ described for $i = A, B$ by:

$$\frac{dS^i}{dt}(t) = -\frac{S^i(t)}{\tau} + (1 - S^i(t))H(X^i), \quad (2)$$

where $S^i(t) \in [0, 1]$ are the gating variables, $S^i(0) = 0$, $X^i \equiv \alpha S^i(t) - \beta S^j(t) + I^i(t)$, $i = A, B, j \neq i$, and $I^i(t)$ is the input for the option i at t . Input goes from offer value neurons to taste neurons, whose firing rate eventually determines the choice. The quantities $\tau, \alpha, \beta > 0$ are dimensional parameters. The function H describes the neuronal response to the current input.⁴ We assume:

H: The function H is continuous, positive, strictly increasing

If the values of the variables S^i tend to a limit where $S^A(t)$ is larger than $S^B(t)$ then good A is chosen. The input $I(t)$ is a stochastic process; to understand the choice process is useful

³See Wong and Wang (2006, Appendix, p. 1327).

⁴The function H can be derived from the first passage time formula derived from the theory of leaky integrate and fire neurons. A convenient specification is the approximation suggested by Abbott and Chance (2005):

$$H(x) \equiv \frac{ax - b}{1 - e^{-d(ax - b)}}. \quad (3)$$

The basic properties of the system we study only depend on the assumption **H**, so the use of an approximation is not a concern. Note

that the function $G(r) \equiv \frac{r}{1 - e^{-dr}}$, for $r \in \mathbb{R}$, is continuous, positive everywhere (in particular it is $1/d$ at 0) and strictly increasing everywhere. Thus H has the same properties and satisfies the assumption **H**. In addition G is asymptotically linear with slope 1, that is $\lim_{r \rightarrow \infty} G(r)/r = 1$, so H is asymptotically linear with slope $a > 0$. If we denote by h the inverse function of H , a continuous strictly increasing function, h is asymptotically linear with slope $1/a$.

first to consider the case where $I^i(t)$ is a deterministic constant. In this case the dynamic process can be studied with the phase diagram illustrated in Figure 2.

III. Hebbian Learning

The input $I^A(t)$ is proportional to the firing rate of the offer neurons. As we have seen in the review of the experimental evidence presented in Section I, the firing rate of the offer neuron for A in environment μ for good A when the quantity x is offered is fully adapted to μ , and equal to the rank of the quantity x , defined as

$$F_\mu(x) \equiv \mu\{y: y \leq x\}. \quad (4)$$

If no correction is provided, the choice between two quantities x and y would depend on the environment, but experimental evidence rejects this hypothesis. So we assume that the mean of the input for i is a product of the fully adapted input 4 and the strength of the synaptic efficacy between offer and taste neurons:

$$EI^A(t) = w^A(n)F_\mu(x), \quad (5)$$

where $w^A(n)$ is the strength of the synaptic efficacy at trials n . The process (Fusi 2002) adjusts the synaptic strength (w^A, w^B) by an amount proportional to the difference between the reward obtained if the choice was A , times the firing rate of the offer neurons of type A , minus the current strength. Let $R(n)$ be the reward received at trial n , irrespective of whether the choice was A or B ; let $C(n) \in \{A, B\}$ be the choice at trial n , $I_A(C) = 1$ if $C = A$ and zero otherwise. Weight w^A adjusts according to:

$$w^A(n+1) = w^A(n) + \frac{\gamma}{n} Adj^A(n), \quad (6)$$

with

$$Adj^A(n) \equiv R(n)^\alpha \delta_A(C(n))F_\mu(x(n)) - w^A(n),$$

($\delta_A(C) = 1$ if $C = A$, zero otherwise). We study the dynamics of the system of equation (6) and the corresponding equation for the B good by studying the associated ODE, which for the A good is

$$\dot{w}^A(s) = \gamma E \left(R^\alpha \delta_A(C) F_\mu(x) - w^A(s) \right), \quad (7)$$

where the expectation is taken with respect to the probability over the choice C and the realization of the quantities x and y in the trial.

In equation (6) we take the two distributions as the uniform distributions between 0 and b^i , $i = A, B$. Then:

PROPOSITION 3.1: The system described by equation (7) (and corresponding equation for B) has a unique, globally asymptotically stable steady state

Hebbian learning will compensate for the potential bias introduced by adaptive coding if the weights tend to values proportional to the range. Let us call $(W^A(b^A, b^B), W^B(b^A, b^B))$ the steady state at (b^A, b^B) , and consider how the steady states of equation (7), and so the limit of the process, depends on the values (b^A, b^B) . Figure 3 shows that Hebbian learning compensates adaptive coding. This conclusion is supported by numerical simulations of the entire system described by equations (2), (5), and (6), producing the probability of choice according to the model described in Section I.

Reward enters nonlinearly in the adjustment. The following proposition gives a lower bound of the weight for A as the upper boundary is multiplied by a constant.

PROPOSITION 3.2: For $R > 1$ and $b > 0$, $W^A(Rb, b) > R^\alpha W^A(b, b)$

When $\alpha = 1$ then the weight is super-linear in the R value. In fact the steady state values in panel A of the figure illustrate well how the ratio among the steady state values is considerably higher than 2. Considering that weights represent synaptic efficacy, which has a physiological upper bound, Proposition 3.2 shows that if reward entered linearly into the adjustment, Hebbian learning could only be effective in compensating adaptive coding locally.

IV. Conclusions

Rational attention is an *as if* theory, based on special assumptions on the function specifying the cost of attention. Here we have provided an *as is* theory of choice. Within this framework adaptive coding is defined precisely: the neural response of the offer neurons is proportional to the rank, defined in equation (4). The mechanism producing choice is fully specified (by equations (2), (3), (4), and (5)). Choice is stochastic, and the model predicts precisely the extent of the variability. The model is mechanistic, with no free variables: each variable in it can be experimentally measured. Adaptive coding according to equation (4) produces the potential of a bias (identified in behavioral economics literature as a reference point bias), which however does not appear in behavior. The fully specified Hebbian learning adjustment (equation (6)) provides a possible explanation of why bias does not occur. The system learns the basic statistical properties of the environment, and the bias that is introduced in the initial trials is compensated after a limited number of trials.

Further research is under way. An extension to human decision making is essential, including factors like provision of attentional effort in decision making (Dickhaut, Rustichini, and Smith 2009). A natural question arises: why should adaptive coding be adopted, and then compensated by a complex, error prone process? The answer to this

question requires a more detailed understanding of the neural transmission process, and is the topic of current research.

Acknowledgments

The authors thank Nicolas Brunel, Paul Glimcher, Antonio Rangel, Wolfram Schultz, Kong-Fatt Wong, Xiao-Jing Wang, and Ryan Webb for extremely useful discussions on the content of this paper and much more. They are in particular grateful to Kong-Fatt Wong and Xiao-Jing Wang for sharing the MatLab code for the simulation of the model in Section II. CP-S thanks the National Institute on Drug Abuse for financial support (grant R01-DA032758). AR thanks the NSF for financial support grant SES-1061817, and CIMEC (Center for Mind/Brain Sciences), Università degli Studi di Trento for support in experimental research.

References

- Abbott LF, Chance Frances S. Drivers and Modulators from Push-Pull and Balanced Synaptic Input. *Progress in Brain Research*. 2005; 149:147–55. [PubMed: 16226582]
- Cai, Xinying, Padoa-Schioppa, Camillo. Neuronal Encoding of Subjective Value in Dorsal and Ventral Anterior Cingulate Cortex. *Journal of Neuroscience*. 2012; 32 (11):3791–808. [PubMed: 22423100]
- Dickhaut, John, Rustichini, Aldo, Smith, Vernon. A Neuroeconomic Theory of the Decision Process. *Proceedings of the National Academy of Sciences*. 2009; 106 (52):22145–150.
- Fusi, Stefano. Hebbian Spike-Driven Synaptic Plasticity for Learning Patterns of Mean Firing Rates. *Biological Cybernetics*. 2002; 87:459–70. [PubMed: 12461635]
- Kobayashi, Shunsuke, de Carvalho, Ofelia Pinto, Schultz, Wolfram. Adaptation of Reward Sensitivity in Orbitofrontal Neurons. *Journal of Neuroscience*. 2010; 30(2):534–44. [PubMed: 20071516]
- Louie, Kenway, Grattan, Lauren E., Glimcher, Paul W. Reward Value-Based Gain Control: Divisive Normalization in Parietal Cortex. *Journal of Neuroscience*. 2011; 31(29):10627–39. [PubMed: 21775606]
- Laughlin, Simon. A Simple Coding Procedure Enhances a Neuron's Information Capacity. *Zeitschrift für Naturforschung*. 1981; 36c:910–12.
- Padoa-Schioppa, Camillo, Assad, John A. Neurons in the Orbitofrontal Cortex Encode Economic Value. *Nature*. 2006; 441(7090):223–26. [PubMed: 16633341]
- Padoa-Schioppa, Camillo, Assad, John A. The Representation of Economic Value in the Orbitofrontal Cortex is Invariant for Changes of Menu. *Nature Neuroscience*. 2008; 11(1):95–102. [PubMed: 18066060]
- Sims, Christopher A. Stickiness. "Carnegie- Rochester". *Conference Series on Public Policy*. 1998; 49:317–56.
- Sims, Christopher A. Implications of Rational Inattention. *Journal of Monetary Economics*. 2003; 50(3):665–90.
- Tobler, Philippe N., Fiorillo, Christopher D., Schultz, Wolfram. Adaptive Coding of Reward Value by Dopamine Neurons. *Science*. 2005; 307(5715):1642–45. [PubMed: 15761155]
- Tremblay, Léon, Schultz, Wolfram. Relative Reward Preference in Primate Orbitofrontal Cortex. *Nature*. 1999; 398:704–08. [PubMed: 10227292]
- Wong, Kong-Fatt, Wang, Xiao Jing. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *Journal of Neuroscience*. 2006; 26(4):1314–28. [PubMed: 16436619]

Figure 1.**Challenge Posed by Range Adaptation**

Notes: Panel A: Range adaptation in orbitofrontal cortex. Each line represents the average neuronal response (baseline-subtracted) plotted against the offer value. Different lines indicate different value ranges. Panel B: In this simplified model, choices result from the activity of two neurons encoding offer value *A* (left) and offer value *B* (right). When the range of juice *B* increases (lighter line), the offer value *B* cell adapts. The indifference point, for which the two cells have equal firing rate, shifts. If decisions were made by comparing firing rates, the same quantity of juice *B* would be chosen less frequently.

Source: Panel A reproduced from Padoa-Schioppa (2009)

Figure 2.

Dynamics of Gating Variables and Choice Process (*Axes indicate the variables S*)

Notes: Arrows describe the direction of flow of the differential equation (2) and its correspondent for the B good, when the input $I^i(t)$ is constant over time for both goods. The two lines describe the zeroes of the vector field; their intersection the steady states, thus the possible limit values of the two gating variables. Two of the steady states are stable and one (the intermediate) unstable. In panel A, the input for the two goods is the same, so with noise the two goods are chosen with equal probability. In panel B, the input of good A is higher; the set of zeroes for the S^A variables shifts upward. Now a process starting at the initial condition $S^i = 0$ for both goods is more likely to converge to the bottom-right steady state. So the probability of choosing A is higher than $1/2$. With a further increase of the input for A , only one steady state will remain, and good A will be chosen for sure.

Figure 3.

Dynamics of Hebbian Learning (*Axes indicate the weights*)

Notes: Arrows describe the direction of flow of the ordinary differential equation (7), and the corresponding equation for good B . The environments are uniform distributions. The two lines describe the zeros of the vector fields; their intersection the steady state, thus the limit values of the two weights. In panel A, the ranges of offers of the two goods are the same, and the two steady states have equal value. In panel B the range of good A is doubled in size, that of good B is unchanged.