

The following resources related to this article are available online at www.sciencemag.org (this information is current as of March 30, 2009):

Updated information and services, including high-resolution figures, can be found in the online version of this article at:

<http://www.sciencemag.org/cgi/content/full/304/5678/1782>

Supporting Online Material can be found at:

<http://www.sciencemag.org/cgi/content/full/304/5678/1782/DC1>

A list of selected additional articles on the Science Web sites **related to this article** can be found at:

<http://www.sciencemag.org/cgi/content/full/304/5678/1782#related-content>

This article **cites 23 articles**, 5 of which can be accessed for free:

<http://www.sciencemag.org/cgi/content/full/304/5678/1782#otherarticles>

This article has been **cited by** 190 article(s) on the ISI Web of Science.

This article has been **cited by** 68 articles hosted by HighWire Press; see:

<http://www.sciencemag.org/cgi/content/full/304/5678/1782#otherarticles>

This article appears in the following **subject collections**:

Psychology

<http://www.sciencemag.org/cgi/collection/psychology>

Information about obtaining **reprints** of this article or about obtaining **permission to reproduce this article** in whole or in part can be found at:

<http://www.sciencemag.org/about/permissions.dtl>

Matching Behavior and the Representation of Value in the Parietal Cortex

Leo P. Sugrue,* Greg S. Corrado, William T. Newsome

Psychologists and economists have long appreciated the contribution of reward history and expectation to decision-making. Yet we know little about how specific histories of choice and reward lead to an internal representation of the “value” of possible actions. We approached this problem through an integrated application of behavioral, computational, and physiological techniques. Monkeys were placed in a dynamic foraging environment in which they had to track the changing values of alternative choices through time. In this context, the monkeys’ foraging behavior provided a window into their subjective valuation. We found that a simple model based on reward history can duplicate this behavior and that neurons in the parietal cortex represent the relative value of competing actions predicted by this model.

Natural environments are characterized by uncertainty in both the sources and timing of rewards (1). Humans and other animals are sensitive to these variables and adapt the statistics of their foraging behavior to those of the environment (1–4). Specifically, animals distribute their time among foraging sites in proportion to their relative value (5), i.e., the relative abundance of resources at each site. This phenomenon, called matching behavior, was studied experimentally by Herrnstein, who expanded it into a general principle of choice that he termed the matching law (6–9). Stated mathematically, the matching law asserts that the fraction of choices made to any option will exactly match the fraction of total income (i.e., total rewards) earned from that option, or

$$\frac{I_k}{\sum I} = \frac{C_k}{\sum C}$$

where I_k and C_k are the total income earned and total choices on option k , respectively, and the summations are over all available options.

To match behavior to income, animals must integrate rewards earned from particular behaviors, and the brain, in turn, must maintain an appropriate representation of the value [i.e., reward frequency (5)] of competing alternatives. Matching provides a behavioral readout of this internal representation. By studying matching in the context of visually based eye movement behav-

ior, we aim to leverage our knowledge of the anatomy and physiology of the primate visual and oculomotor systems to investigate how value is represented at the neural level. For this purpose, we trained rhesus monkeys (*Macaca mulatta*) to perform a dynamic version of a classical matching task in which saccadic eye movements to a pair of competing visual targets are rewarded at different rates (Fig. 1A) (10).

A dynamic foraging task. On each trial in this task, the monkey is free to choose between two targets; the color of each target cues the probability that its selection with an eye movement will be rewarded with a drop of juice. Analogous to natural environments, rewards in this task are assigned to the two colors at rates that are independent and stochastic (Poisson probability distribution). Once assigned, a reward remains available until the associated color is chosen (11). This persistence of assigned rewards means that the likelihood of being rewarded increases with the time since a color was last chosen, and ensures that matching approximates the optimal probabilistic strategy in this task (12–14).

Figure 1B depicts representative behavioral data from a single session in which a monkey experienced a series of six different ratios of reward rates. Two features of these data are notable. First, the blue line generally parallels the black, indicating that the monkey indeed matched the ratio of its choices to the ratio of incomes from the two colors, as predicted by the matching law. Second, the monkey appears to adjust its behavior very rapidly to unsignaled changes in the rates of reward.

The income ratios indicated by the black lines in Fig. 1B represent mean reward

rates and obscure the stochastic manner in which rewards become available in the task. We can visualize this variability by plotting instantaneous estimates of choice and income ratios (Fig. 1C). These estimates suggest that the relationship between choices and experienced rewards is highly local in time. This is evident both at the transitions between income ratios, when behavior lawfully and rapidly adjusts to unsignaled changes in the rates of reward, and within blocks, when choices track local variability in the experienced income ratio (red asterisks). If behavior were based on a representation of reward history that extended into the distant past, it would be incapable of such rapid adjustment (15).

Traditionally applied to foraging in stationary environments (for which reward rates do not change), the matching law relates cumulative choice to total experienced income and is intrinsically a global description of behavior averaged over long periods of time. The data in Fig. 1, B and C, confirm an earlier study of rats by Gallistel and colleagues (16), in which they observed that animals accustomed to dynamic environments can match under such conditions. Our data further suggest that this behavior is driven by a process that is intrinsically local in time. These results prompt us to ask whether the classical matching law can be reformulated as a more local description and whether this description can explain the behavior that we observed.

A local formulation of matching. Income earned during a behavioral session is simply the integrated reward stream that an animal has experienced (Fig. 2A). In the traditional matching law, each new reward contributes equally to the income attributed to a particular option without discount or decay. The fractional income for a particular option (the income from that option divided by the total income from all available options) then dictates the proportion of choices allocated to it. But what if our integrator were not perfect, but somewhat leaky (Fig. 2B)? This leak would confer a finite effective memory on estimates of income, making them local rather than global. In this model, the local fractional income translates directly into the instantaneous probability of choice for a given option (17). Importantly, this proposed local matching rule obeys the correspondence principle: When limited to large data sets and stationary environments (where matching has been most extensively documented), the predictions of our local matching rule approximate those of the classical global matching law. We show below that the leaky integration model is surprisingly successful at describing behavior in our dynamic task.

Howard Hughes Medical Institute and Department of Neurobiology, Stanford University School of Medicine, Stanford, CA 94305, USA.

*To whom correspondence should be addressed. E-mail: leo@monkeybiz.stanford.edu

Postulating a process of leaky integration marks a conceptual shift from the parameterless matching law, appropriate for stationary reward conditions, to a one-parameter model of matching behavior appropriate for dynamic conditions. The single parameter in this simple model is the time constant τ of the leaky integrator. How do changing values of τ affect the model's behavior? Intuitively, higher values of τ mean slower leaks and would give rise to more stable and accurate estimates of income. The cost of such reliable estimates is that they respond sluggishly to changes in the environment. Conversely, lower values of τ produce estimates of income that respond quickly to change, but are substantially noisier during periods of stability. Given this trade-off between accuracy and adaptability, what value of τ yields the highest income given the statistics of our task?

To answer this question, we simulated the behavior of the model on our task and examined how its performance varied as a function of the integration time constant τ . Each simulation consisted of a quarter of a million choices made by a model with a particular τ across the identical sequence of reward-rate ratios and block lengths encountered by our monkeys. The thick black curve in Fig. 2C plots the outcome of these simulations in terms of foraging efficiency (the percentage of the maximum reward rate achieved). We also plot realistic bounds for performance imposed by the structure of the task. The upper bound demarcates the average performance of an ideal probabilistic forager. This hypothetical ideal strategy “knows” the reward rate of each option, thereby dispensing with the estimation process, and uses this information to make choices that maximize its expected rate of reward. In contrast, the lower bound shows the average performance of a completely random foraging strategy and represents chance performance in our task. Despite its simplicity, the best-performing leaky integrator model does well relative to these bounds, collecting 93% of the rewards attained by the ideal clairvoyant strategy.

How does our model compare to the choices of real biological players? Our monkeys' behavior, indicated by the blue circles on the same panel, corresponds well with the predictions of the model. We estimated τ for each monkey by minimizing the mean squared error between the probability of choice predicted by the model and the animal's actual binary choices across all experiments. For each monkey, this best-fit τ lies within a standard error of the best-performing model. Foraging efficiency was estimated as the percentage of the maximum reward rate achieved by each monkey across all experiments. These performance levels fall just below that of the model with a similar time constant, an unsurprising outcome given that the monkeys,

unlike the model, are susceptible to variables such as distraction and satiation.

The next three panels of Fig. 2 further explore the similarity between the behavior of the best model and that of our monkeys. Figure 2D shows the cumulative responses of the best model ($\tau =$ nine choices) across the same series of blocks shown in Fig. 1C for Monkey G. Qualitatively, the model exhibits dynamic matching behavior that is very similar to that of the animal. The next two panels (Fig. 2, E and F) reinforce this impression with more quantitative comparisons. First, the model predicts that the probability of choosing red will vary linearly with the local fractional income from red (the unity line in Fig. 2E). Figure 2E shows this to be approximately true for the behavioral data. Second, because the model is strictly probabilistic, it predicts that the number of successive trials on which a player (monkey or model) will choose a given color before switching will be distributed as the average of a family of exponentials. Figure 2F plots these distributions of stay durations; not only is the

monkeys' distribution exponential, but it is almost an exact fit to that of the model with the best-performing τ (18).

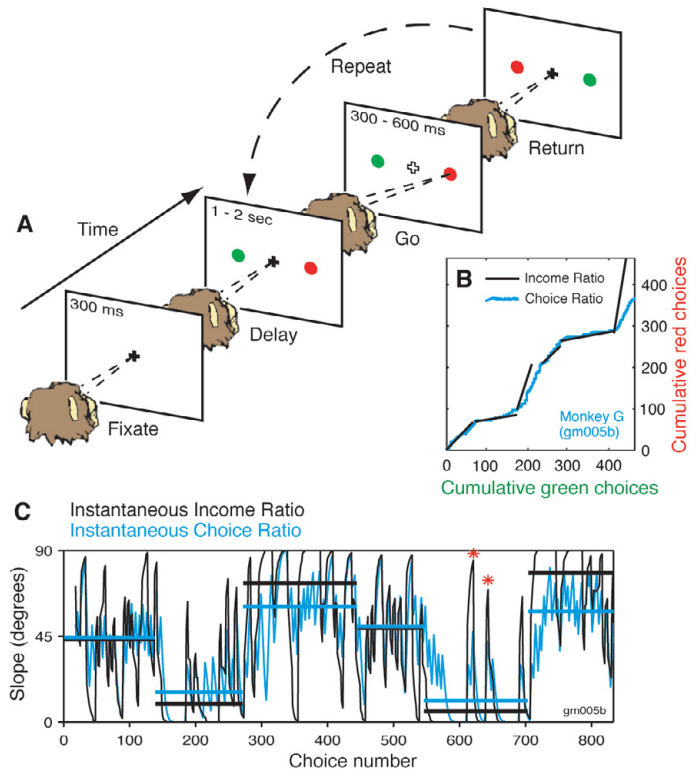
These similarities in qualitative behavior, foraging performance, fitted τ , and simple statistics demonstrate that our local matching rule is an adequate descriptive model of real choice behavior in this dynamic foraging task (19). Moreover, they suggest that our animals have tuned the time over which they integrate reward information to be optimal for the particular statistics of the task they encountered. The importance of this modeling effort goes beyond its utility in understanding behavior. The model provides us a window into the animal's internal valuation of available options and gives us a metric—local fractional income—that allows us to estimate how the monkey values each of the two colors on every trial, even before it renders a decision. Equipped with this quantitative trial-by-trial measure, we are poised to explore how value is represented in the brain.

The representation of fractional income in the parietal cortex. The lateral intraparietal (LIP) area of the posterior pari-

Fig. 1. Matching behavior in monkeys.

(A) The sequence of events of an oculomotor matching task: (i) Fixate. To begin a run of trials, the animal must fixate the central cross. (ii) Delay. Saccade targets appear (randomized spatially by color) in opposite hemifields while the animal maintains fixation. (iii) Go. Dimming of the fixation cross cues a saccadic response and hold. (iv) Return. Brightening of the fixation cross cues return, target colors are then rerandomized, and the delay period of the next trial begins. Reward is delivered at the time of the response, if at all. Overall maximum reward rate is set at 0.15 rewards per second. Relative reward rates changed in blocks (~100 to 200 trials) without warning;

ratios of reward rates were chosen unpredictably from the set {8:1, 6:1, 3:1, 1:1}. (B) Dynamic matching behavior. Representative behavior of Monkey G during a single session. Continuous blue curve shows cumulative choices of the red and green targets. Black lines show average ratio of incomes (red:green) within each block (here, 1:1, 1:3, 3:1, 1:1, 1:6, and 6:1). Matching predicts that the blue and black curves are parallel. (C) Slope space. Same data as in (B), plotted to allow visualization of ongoing covariation in local ratios of income and choice. The x axis shows session time (in choices). The y axis shows running estimates of the ratios of income (black) or choice (blue). Ratios were computed after smoothing the series of rewards or choices with a causal half-Gaussian kernel (SD of six choices) and are expressed as slopes (arctangent of ratio). Thick horizontal black and blue lines indicate average income and choice ratios within each block. Red asterisks highlight example regions where the choice ratio obviously tracks local noise in the experienced ratio of incomes.



etal cortex contains activity appropriate for guiding saccadic eye movements, signals that have been variously interpreted as working memory for visual targets, attention to salient spatial locations, or motor planning (20–23). In the context of more sophisticated eye movement tasks, investigators have documented the modulation of LIP activity by the strength of sensory evidence that supports a perceptual judgment (24–26) and by both the prior probability that a particular movement is instructed and the volume of juice associated with that movement (27). Such encoding of information from diverse sources is a proposed property of brain areas responsible for computing putative decision variables that link sensory information to motor responses

(28). If this suggestion is correct, and LIP is indeed an important locus for oculomotor decisions, then in a setting where eye movement decisions are informed by reward history and expectation, we anticipate the appropriate decision variable to be represented in LIP. Accordingly, the following physiological experiments test the prediction that in the matching task, neurons in LIP encode the local fractional income (Fig. 2B) of competing target colors.

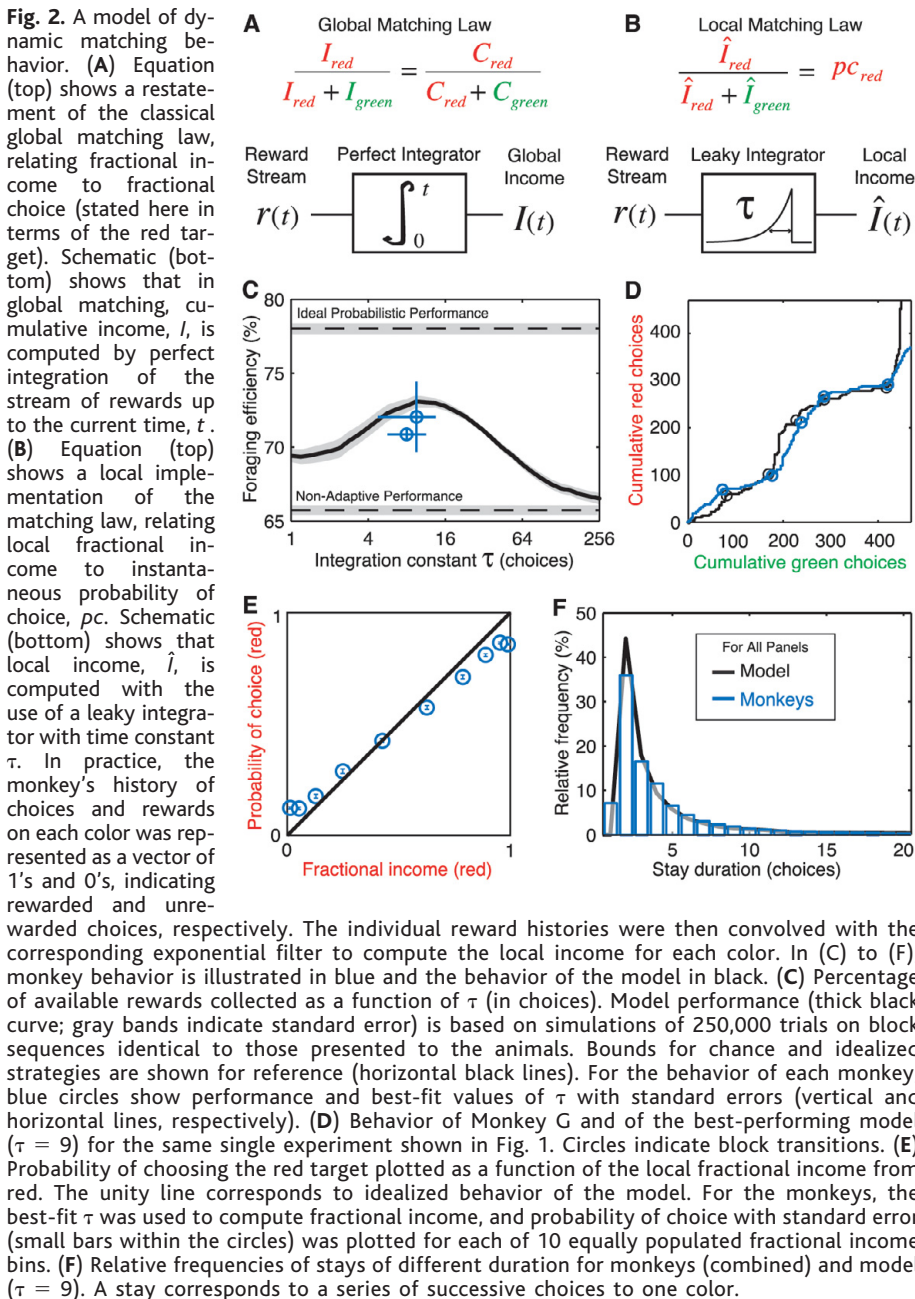
We selected for study LIP neurons that showed sustained, spatially selective activity in the context of a classical delayed saccade task (Fig. 3A). These neurons respond only when targets are presented within a circumscribed region of the visual

field termed the cell’s response field (RF). Approximately one-third of the cells that we encountered in LIP met this criterion, including 33 neurons from the left hemisphere of Monkey G and 29 from the right hemisphere of Monkey F.

Figure 4A illustrates how we studied these 62 LIP neurons in the matching context. Critically, in this setting, trials that shared an identical visual stimulus configuration and ended in the same motor response still varied widely in the local fractional income of the chosen target. Thus, on some trials the monkey chose the target inside the cell’s RF and this target had a high fractional income, whereas on other trials the fractional income was much lower. Our experimental question was whether, within each category of motor response, activity in LIP is influenced by the local fractional income of the chosen target.

Figure 4B shows representative data from the same cell featured in Fig. 3, now recorded during performance of the matching task. For each trial, the cell’s mean delay-period response is plotted against the local fractional income of the chosen target. Activity is shown separately for trials that end in saccades into (blue) and out of (green) the cell’s RF. We observed a positive correlation between firing rate and fractional income for choices into the RF and a negative correlation for choices out of the RF. The solid lines are regressions fit to these two sets of data by the method of least squares and are characterized by positive and negative slopes for choices into and out of the RF, respectively. When the fractional income of the chosen color is low, the clouds of blue and green points overlap, indicating that the activity of this particular cell is no longer a reliable indicator of the direction of the monkey’s saccade at the end of the trial. This result is particularly notable given that this cell was chosen for its spatial selectivity in the delayed saccade task (Fig. 3, A and B).

To see how the effect of fractional income varies across our population of LIP cells, we performed this regression analysis for each neuron in our sample. Figure 4C shows the resulting regression slopes. The upper histogram (blue) is the distribution of slopes for choices into each cell’s RF. Consistent with the example in Fig. 4B, this distribution is centered to the right of zero, indicating positive regressions of activity on fractional income. The lower histogram (green) is the analogous distribution of slopes for choices out of each cell’s RF. Again, in keeping with the example, this distribution is centered to the left of zero, indicating negative regressions of activity on fractional income. Importantly, in the delayed saccade task, none of these neurons showed any influence of recent reward history as described by the local income from the lone response target (Fig. 3C).



The preceding analysis assumes that all eye movements to a particular target are effectively equivalent. To control for the possibility that subtle variations in the precise metrics of saccades to the same target location might cause changes in LIP firing rates, we expanded our regression model of each cell's response to include a range of saccade metrics as co-regressors (29). If our results actually reflect a subtle effect of saccade metrics, explicit inclusion of these factors should nullify the apparent influence of fractional income. Instead, the 43 cells that showed a significant effect of fractional income continued to do so after the inclusion of these co-regressors (95% confidence interval for the fractional income coefficient still excluded zero), and the magnitude of this effect was largely unchanged (average decrease in effect size = 14%).

To examine the time course of the effect of fractional income across the population, we peak-normalized the firing rates of the 43 cells that showed a significant regression effect and computed the average time course of the cells' response as a function of fractional income. Figure 5 plots these average normalized rates for the population. Two important points emerge from this analysis. First, the effect of fractional income is not apparent at the beginning of the trial, but emerges over time. Second, activity remains graded with respect to fractional income up to the time of the saccade itself, irrespective of whether this saccade is directed into (blue) or out of (green) the cell's RF. This suggests that this population of LIP neurons encodes information about the value of locations in space, whether or not they are the endpoint of the impending saccade.

Discussion and conclusions. Matching belongs to a class of behaviors purported to engage cognitive mechanisms that animals use when competing for resources in stochastic environments. Because matching results in an equilibrium state in which returns from competing behaviors are equalized, it represents a stable and effective foraging strategy from both an evolutionary and game theoretic perspective.

Somewhat surprisingly, we find that matching behavior in a dynamic context is well described by a simple local reformulation of the classical matching law. This local matching rule uses leaky integrators of rewards to estimate the local income earned from competing behaviors and sets the instantaneous probability of choosing an alternative equal to its local fractional income. This simple model has only one tuned parameter: the decay constant (τ) of the integrators. Intriguingly, we found that the specific values of τ used by our animals were optimally tuned for the statistics of the environment they encountered in this

Fig. 3. Activity of an LIP neuron during the delayed saccade task. (A) Delayed saccade task used for cell selection. Only a single purple target is presented on each trial at one of a variety of spatial locations. The sequence of events is otherwise identical to Fig. 1, and rewards were delivered at the same overall rate used in the matching task. Dotted blue oval represents LIP response field (RF). (B) Response histograms of an example cell for trials into (blue) and out of (green) this cell's RF demonstrate classical LIP spatial selectivity. Activity is aligned on both the appearance of the visual target (left) and the time of the saccade (right). The break in the time axis reflects this dual alignment for trials of variable length. spks/sec, spikes per second. (C) Activity during the delayed saccade task shows no dependence on recent reward history. For trials into (blue) and out of (green) the cell's RF, average delay-period activity is plotted against the local income resulting from the single purple target. Local income was estimated by filtering reward history with a local exponential with the same best-fit τ that was used to compute fractional income in Fig. 2.

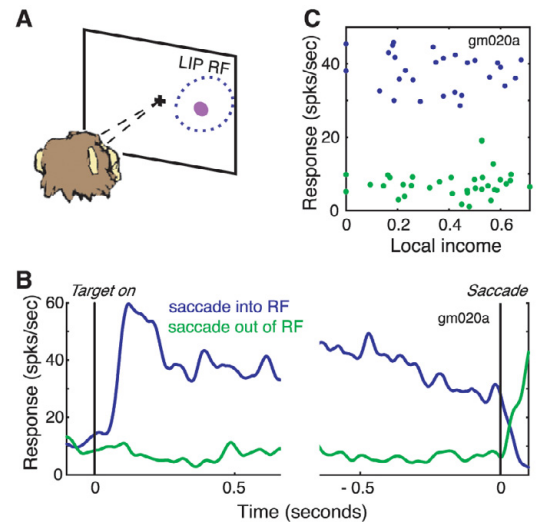


Fig. 4. Activity in LIP during the matching task. (A) Task geometry used in the matching task. Dotted blue oval represents RF of the LIP cell under study. Color-location association was randomized between trials. (B) Representative matching data from the same example LIP neuron shown in Fig. 3. For each trial, mean delay-period activity is plotted as a function of the local fractional income of the chosen target. Blue and green indicate choices into and out of the RF respectively. Lines are least squares regressions fit to the corresponding points (blue: slope = 11.4, $r = 0.4$, $P < 0.001$; green: slope = -19.9, $r = 0.58$, $P < 0.001$). (C) Distribution of slopes for regression of firing rate on local fractional income across our population of 62 LIP neurons. Separate distributions show effect for choices into (upper, blue) and out of (lower, green) each cell's RF; 95% confidence intervals for means of these distributions are 2.1 to 4.5 and -6.6 to -4.0, respectively. Filled bars highlight regressions that are significant at the 0.05 level. Asterisks indicate the example cell.

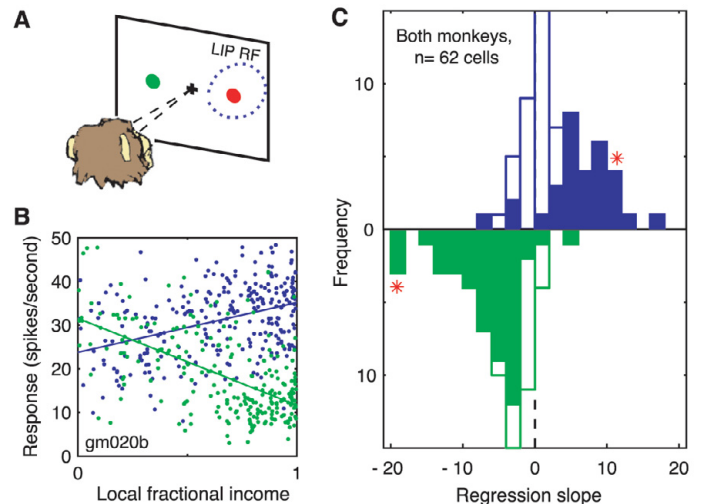
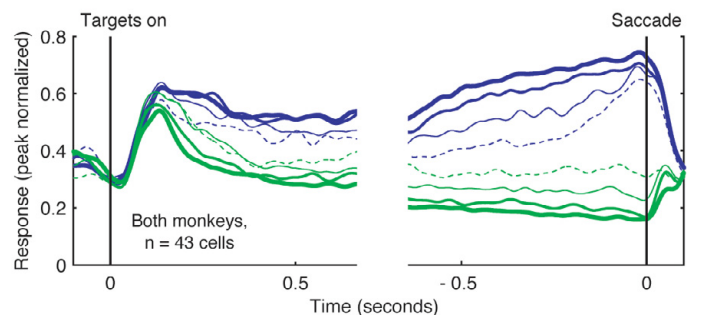


Fig. 5. Time course of the effect of local fractional income. Response histograms show average peak-normalized firing rates for 43 cells with activity that regressed significantly on fractional income. Before normalization and averaging, raw spike trains were smoothed with a Gaussian kernel (SD = 20 ms). Activity is aligned on both the appearance of the visual target (left), and the time of the saccade (right), and is shown separately for choices into (blue) and out of (green) the cell's RF. Trials are further subdivided into four groups according to the local fractional income of the chosen target: solid thick lines, 0.75 to 1.0; solid medium lines, 0.5 to 0.75; solid thin lines, 0.25 to 0.5; dotted thin lines, 0 to 0.25.



task. By manipulating overall reward rate and the dynamics with which rates change, future experiments may address whether animals can flexibly adjust the time scale of their integration to maintain this optimality.

Previous studies have documented reward- and value-related signals in numerous cortical and subcortical areas [reviewed in (30)], but primarily in the context of imperative tasks where behavioral responses are directly instructed or conditioned. We suggest that elucidating the functional roles of these signals will require studying them in settings where value itself is the primary determinant of behavior. Our current work marks an initial step in this direction, as does ongoing work in other laboratories (31, 32). Interpreting neural activity in such “free choice” contexts necessitates a further methodological shift from correlations with directly accessible sensory or behavioral events to quantitative modeling of the ostensibly “hidden” variables that link experience to action. During performance of the matching task, we found that the activity of single LIP neurons parametrically encoded trial-to-trial fluctuations in the pertinent decision variable: local fractional income. This result supports the suggestion that area LIP plays a role in implementing oculomotor decisions and extends the findings of previous studies of LIP activity in the context of visual motion discrimination tasks (24–26) to the realm of value-based choice.

Is local fractional income actually computed in LIP? Although activity in area LIP is correlated with this value metric, it is unlikely the primary locus where fractional income is computed and maintained. A population of neurons whose activity directly encoded value should do so in terms of the relevant value cue (in this case, color, not space) and maintain that representation across an appropriate time scale (in this case, several trials) (Fig. 2B). In contrast, income-related signals in LIP are spatially mapped and are “reset” at the start of each trial, developing anew over the first several hundred milliseconds (Fig. 5). An important direction for future research will be to identify where value is first explicitly encoded in the brain and how this representation is conferred with a temporal profile appropriate for optimal behavior.

Rather than computing value, we suggest that area LIP plays a critical role in remapping abstract valuation to concrete action. This remapping is demanded by the logic of our task: On every trial the monkey must transform a color-based representation of value into a spatial eye-movement plan. By representing value in spatial terms, LIP may contribute to this transformation and directly influence the probability that a particular region of space will serve as the endpoint of the next saccade. This interpretation is consistent

with the unifying proposal that area LIP functions as a saliency map of visual space (33) of the type invoked in visual psychophysics or computational vision (34, 35), capable of flexibly combining and representing a variety of information for the purpose of guiding eye movements or shifts in visual attention.

References and Notes

1. J. M. Smith, *Evolution and the Theory of Games* (Cambridge Univ. Press, Cambridge, 1982).
2. D. Kahneman, A. Tversky, Eds., *Choices, Values, and Frames* (Cambridge Univ. Press, Cambridge, 2000).
3. D. W. Stephens, J. R. Krebs, *Foraging Theory* (Princeton Univ. Press, Princeton, NJ, 1986).
4. J. R. Krebs, N. B. Davies, Eds., *Behavioral Ecology: An Evolutionary Approach* (Blackwell, Oxford, ed. 4, 1997).
5. Value is an inherently subjective concept. To study value, it is necessary to first operationalize it in terms of variables that can be directly observed. One external determinant of value is reward, which we define as anything that an animal will work to acquire (a squirt of juice in the case of our thirsty monkeys). At an intuitive level, the value of a behavioral option to our monkeys could be related to several aspects of the liquid reward, including its size or its inherent desirability (e.g., sweet juice versus water). In the matching task used in this study, value is related to the frequency with which an option results in reward. This concept is captured by the economic term “income,” defined as the number of rewards earned over some period of time. We show that we can effectively quantify an option’s value in terms of the income experienced by the monkey over a well-defined, recent period of time, which we term “local income.” Because the matching task is a two-alternative, forced-choice task, the local income associated with a single behavioral option is only a fraction of the total income earned as the monkey distributes its choices between the two options. We therefore use the term “local fractional income” (defined formally in Fig. 2) to refer to the local income associated with a single behavioral option.
6. R. J. Herrnstein, *J. Exp. Anal. Behav.* **4**, 267 (1961).
7. P. A. de Villiers, R. J. Herrnstein, *Psychol. Bull.* **83**, 1131 (1976).
8. H. Rachlin, D. I. Laibson, Eds., *Richard J. Herrnstein, The Matching Law: Papers in Psychology and Economics* (Harvard Univ. Press, Cambridge, MA, 1997).
9. M. C. Davison, D. McCarthy, Eds., *The Matching Law: A Research Review* (Erlbaum, Hillsdale, NJ, 1988).
10. Materials and methods are available as supporting material on Science Online.
11. Our task includes a “changeover delay” (COD), a manipulation necessary to ensure matching behavior by discouraging simple alternating and win-stay-lose-switch strategies [reviewed in (8)]. A COD imposes a switching cost by suspending reward delivery for a brief period following switches between choice alternatives. We implemented a COD by delaying delivery of programmed rewards following switches between colors until the second consecutive choice of the new color. Both monkeys obeyed the COD: On switches, they chose the new color a second time with a probability of over 95%. All model simulations similarly followed this rule. Importantly, we are interested in how behavior and physiology relate to local fractional income as determined by the local matching rule (Fig. 2E). Therefore, we were careful to limit all of our quantitative behavioral and physiological analyses to choices governed by fractional income and to exclude choices (<20% of total) governed by the COD. In practice, this was achieved by treating the first two choices that mark a switch between colors as a single choice.
12. W. M. Baum, *J. Exp. Anal. Behav.* **36**, 387 (1981).
13. G. M. Heyman, *J. Exp. Anal. Behav.* **31**, 41 (1979).
14. J. E. R. Staddon, S. Motheral, *Psychol. Rev.* **85**, 436 (1978).
15. For a quantitative assessment of the local dependence of choice on reward history, see (10), in which we used signal processing techniques to quantify the form and time course of the relationship between the monkeys’ choices and past rewards.
16. C. R. Gallistel, T. A. Mark, A. P. King, P. E. Latham, *J. Exp. Psychol. Anim. Behav. Process.* **27**, 354 (2001).
17. Applying the matching law locally requires a second minor revision to the traditional formulation. In place of cumulative choices to a particular option, locally it is more appropriate to consider instantaneous probabilities of choice. This change is more cosmetic than substantive, given that Herrnstein considered cumulative counts the product of stationary response rates, making them formally equivalent to constant independent probabilities of choice.
18. The very low frequency of stay duration 1 (the first bar) is a direct consequence of the model and monkeys obeying the COD (11).
19. Earlier approaches to behavioral dynamics in foraging tasks have included melioration (36, 37), change detection (16), local accumulator (38), and dynamic feedback models (39). We found the integrator approach to be a simple and effective descriptor of our experimental data and explored a number of variations on this theme. These included the fractional income model described in the text and an alternative differential income model, in which choice is driven by the difference between the integrators’ outputs. We focused on the fractional income model because of its close ties to the classical matching law and its parsimony compared with the differential income model, which required an additional sigmoidal nonlinearity to fit the data well. We also explored several other modeling approaches, including linear-nonlinear models with more than one parameter, generalized linear models, simple reinforcement algorithms, and Kalman filters, but even much more elaborate models captured only slightly more variance in the behavioral data. Importantly, we make no claim that our fractional income model captures the ultimate computation going on inside the animal’s brain. The model is descriptive, not mechanistic—more the broad heuristic of survival-of-the-fittest than the precise machinery of modern genetics—and is best viewed as a first-pass tool for uncovering value-related decision variables in the brain. Whereas the form of the model suggests an obvious potential neural implementation in the persistent activity of a recurrent network of neurons, several other candidate implementations exist for the actual computation that supports the behavior and physiology we observed.
20. S. Funahashi, C. J. Bruce, P. S. Goldman-Rakic, *J. Neurophysiol.* **61**, 331 (1989).
21. J. P. Gottlieb, M. Kusunoki, M. E. Goldberg, *Nature* **391**, 481 (1998).
22. C. L. Colby, M. E. Goldberg, *Annu. Rev. Neurosci.* **22**, 319 (1999).
23. J. W. Gnadt, R. A. Andersen, *Exp. Brain Res.* **70**, 216 (1988).
24. M. N. Shadlen, W. T. Newsome, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 628 (1996).
25. M. N. Shadlen, W. T. Newsome, *J. Neurophysiol.* **86**, 1916 (2001).
26. J. D. Roitman, M. N. Shadlen, *J. Neurosci.* **22**, 9475 (2002).
27. M. L. Platt, P. W. Glimcher, *Nature* **400**, 233 (1999).
28. J. I. Gold, M. N. Shadlen, *Trends Cogn. Sci.* **5**, 10 (2001).
29. The expanded regression model considered the following saccade metrics: latency, peak velocity, duration, amplitude, signed difference between saccade and target angle, and accuracy (reciprocal of distance between saccade endpoint and target location). For each experiment, we determined which metrics were significantly correlated with fractional income and included these in a regression model of the cell’s delay-period response together with choice (into or out of RF) and the local fractional income of the RF target. To quantify the influence of saccade metrics, we compared the magnitude and significance of the coefficient for fractional income under this expanded model and under our standard model, which considers only choice and fractional income.

30. W. Schultz, *Curr. Opin. Neurobiol.* **14**, 139 (2004).
 31. D. J. Barraclough, M. L. Conroy, D. Lee, *Nature Neurosci.* **7**, 404 (2004).
 32. M. C. Dorris, P. W. Glimcher, *Soc. Neurosci. Abstr.* **767**, 1 (2003).
 33. J. W. Bisley, M. E. Goldberg, *Science* **299**, 81 (2003).
 34. C. Koch, S. Ullman, *Hum. Neurobiol.* **4**, 219 (1985).
 35. L. Itti, C. Koch, *Nature Rev. Neurosci.* **2**, 194 (2001).
 36. R. J. Herrnstein, W. Vaughan, in *Limits to Action: The Allocation of Individual Behavior*, J. E. R. Staddon, Ed. (Academic Press, New York, 1980), pp. 143–176.

37. R. J. Herrnstein, D. Prelec, *J. Econ. Perspect.* **5**, 137 (1991).
 38. M. Davison, W. M. Baum, *J. Exp. Anal. Behav.* **74**, 1 (2000).
 39. C. B. Harley, *J. Theor. Biol.* **89**, 611 (1981).
 40. We thank C. R. Gallistel for advice in the planning of these experiments; W. Bair and H. S. Seung for contributions to data analysis; A. E. Rorie for discussions; S. Rosenbaum for animal training; and J. M. Nichols and M. Cohen for comments on an early version of this manuscript. Supported by the National Eye In-

stitute and the Howard Hughes Medical Institute. L.P.S. and G.S.C. received additional support from Stanford Graduate Fellowships.

Supporting Online Material

www.sciencemag.org/cgi/content/full/304/5678/1782/DC1
 Materials and Methods

Fig. S1

References and Notes

16 December 2003; accepted 22 April 2004

REPORTS

Selective Growth of Metal Tips onto Semiconductor Quantum Rods and Tetrapods

Taleb Mokari,^{1,2} Eli Rothenberg,^{1,2} Inna Popov,² Ronny Costi,^{1,2} Uri Banin^{1,2*}

We show the anisotropic selective growth of gold tips onto semiconductor (cadmium selenide) nanorods and tetrapods by a simple reaction. The size of the gold tips can be controlled by the concentration of the starting materials. The new nanostructures display modified optical properties caused by the strong coupling between the gold and semiconductor parts. The gold tips show increased conductivity as well as selective chemical affinity for forming self-assembled chains of rods. Such gold-tipped nanostructures provide natural contact points for self-assembly and for electrical devices and can solve the difficult problem of contacting colloidal nanorods and tetrapods to the external world.

Anisotropic growth of nanomaterials has led to the development of complex and diverse nanostructures such as rods (1, 2), tetrapods (3), prisms (4), cubes (5), and additional shapes (6, 7). These architectures display new properties and enrich the selection of nano-building blocks for electrical, optical, and sensorial device construction. New functionality, such as emissive or rectifying junctions, is introduced into the nanostructure by anisotropic growth with compositional variations. This has been realized by growing semiconductor heterostructures such as p-n junctions and material junctions (e.g., GaAs/GaP) in nanowires (8, 9). In the case of colloidal nanocrystals, growth of a dot-rod structure composed of two different semiconductors and other complex branched structures was achieved (10). In these examples, anisotropic growth was performed with the same material type (semiconductor). Here we report the selective anisotropic growth of two different material systems, a metal onto a semiconductor. We developed a simple method for the selective growth of gold dots onto the tips of colloidal semicon-

ductor nanorods and tetrapods. This combination provides new functionalities to the nanostructures, the most important of which is the formation of natural anchor points that can serve as a recognition element for directed self-assembly and for wiring them onto electrical circuitry.

This advancement has direct relevance to the problem of contact reproducibility and contact resistance that has hindered the study of conductance in nanostructures. Recently there have been reports of good connectivity for micrometer-long quasi-one-dimensional structures such as nanotubes and nanowires (11–13). However, wiring of the shorter colloidal semiconductor rods and tetrapods studied here, with arm lengths of less than 100 nm, has long been an open issue. The use of bifunctional organic ligands, primarily dithiols, as contacting ligands—as used in scanning tunneling microscopy studies (14) and in transport measure-

ments (15)—creates a tunneling barrier, and transport is often dominated by the contact resistances. The use of DNA-based assembly for creating functional circuitry (16, 17) also requires selective anchor points for the directed assembly of nanostructures (18). The Au tips are natural recognition elements for this task as well as for creating complex self-assembled architectures with semiconductor rods and tetrapods in solution (19) and onto substrates.

We prepared CdSe rods and tetrapods of different dimensions by high-temperature pyrolysis of suitable precursors in a coordinating solvent containing a mixture of trioctylphosphineoxide and phosphonic acids (1, 20, 21). We dissolved AuCl₃ in toluene with the addition of dodecyltrimethylammonium bromide (DDAB) and dodecylamine. For growth of Au tips, we mixed this solution at room temperature with a toluene solution of the colloidal-grown CdSe nanorods or tetrapods. After the reaction, the quantum rods were precipitated by addition of methanol and separated by centrifugation. The purified product could then be redissolved in toluene for further study (22).

Figure 1 presents transmission electron microscopy (TEM) images showing growth of Au onto CdSe quantum rods of dimensions 29 × 4 nm (length × diameter); the procedure involved gradual addition of larger amounts of Au precursors (see Table 1 for details). Selective Au growth onto the rod tips (Fig. 1, B to D) is clearly identified as the appearance of points with enhanced contrast. The rods now appear as “nanodumbbells.” Moreover, by controlling the amount of initial Au precursor, it is possible to control the size of the Au tips on the nano-dumbbell edges, from ~2.2 nm (Fig. 1B) to ~2.9 nm (Fig. 1C) to ~4.0 nm (Fig. 1D) (see Table 1). The procedure clearly leads to the growth of

Table 1. Details for Au growth on 29 × 4 nm rods as shown in Fig. 1, with average dimensions (full histograms are shown in fig. S1).

Sample	Nanocrystals (mg)	Dodecylamine (mg)	DDAB (mg)	AuCl ₃ (mg)	Rod size (L × D, nm)	Gold ball size (nm)
1	—	—	—	—	29 × 4	(original rod)
2	10	40	25	4	25.6 × 3.3	2.22
3	10	90	50	8	23.9 × 3.4	2.9
4	10	160	100	13.5	20.8 × 3.2	4

¹Institute of Chemistry, Farkas Center for Light-Induced Processes, ²Center for Nanoscience and Nanotechnology, Hebrew University of Jerusalem, Jerusalem 91904, Israel.

*To whom correspondence should be addressed. E-mail: banin@chem.huji.ac.il