


# Decisions, Uncertainty, and the Brain

The Science of Neuroeconomics



Paul W. Glimcher

First MIT Press paperback edition, 2004  
© 2003 Paul W. Glimcher

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Sabon on 3B2 by Asco Typesetters, Hong Kong and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Glimcher, Paul W.

Decisions, uncertainty, and the brain : the science of neuroeconomics /  
Paul W. Glimcher.

p. cm.

“A Bradford book.”

Includes bibliographical references and index.

ISBN 0-262-07244-0 (hc. : alk. paper), 0-262-57227-3 (pb.)

1. Cognitive neuroscience. 2. Brain—Mathematical models.  
3. Microeconomics. 4. Reflexes. 5. Dualism. I. Title.

QP360.5 .G565 2003

153—dc21

2002026328

10 9 8 7 6 5 4 3 2

---

## Probability, Valuation, and Neural Circuits: A Case Study

This chapter is a case study. It describes how a number of research groups, including my own, have struggled to understand a very specific neural circuit. For about 50 years, neurophysiologists have tried to understand the role that parietal cortex plays in the generation of movements that are guided by sensory events. Among these researchers, a debate has raged over whether parietal cortex should be considered the last link in a hierarchically organized set of structures specialized for sensory processing or whether it should be considered the first step along the final common path for movement production. Despite dozens of clever experiments, no one has ever been able to answer this question convincingly. Indeed, many young researchers have argued that this debate poses an unanswerable, and perhaps even unimportant, question. What I want to argue in this chapter is that the question of what parietal cortex does in this context is neither unanswerable nor unimportant. I want to argue, instead, that the computations around which parietal cortex is organized are missing from the questions which classical approaches *can* ask. I will suggest that if the debate were refocused around concepts like probability and valuation, experiments that try to determine the computational function of parietal cortex would become more fruitful undertakings.

My goal for this chapter is therefore threefold. First, I want to show how modern physiological studies of sensorimotor processing in parietal cortex have tended to stick with traditional reflex-based conceptualizations. The second goal of this chapter is to show that probability theory and theories of valuation can be applied to the physiological problems

encountered in parietal cortex which are usually thought to be the exclusive province of determinate reflexlike theories. The final goal of this chapter is to demonstrate that when these economically based approaches are applied, surprising and illuminating results about the computational structure of the brain can be derived.

To accomplish these goals, I want to describe research on the control of eye movements that is closely related to the work of Newsome and Shadlen described in chapter 5. As those experiments attest, tremendous strides have been made toward understanding how animals gather visual data from the outside world and use that data to plan and execute eye movements. The primate visual system is, without a doubt, the most heavily studied neural system in any vertebrate brain. The primate eye movement control system is undebatably the movement control system that we understand best. For this reason, studies of the pathways that connect visual-sensory systems with eye movement control systems in the monkey brain have served as models for understanding the sensory-to-motor process in general. It is the study of these pathways, and their interconnections in parietal cortex, that is the subject of this chapter.

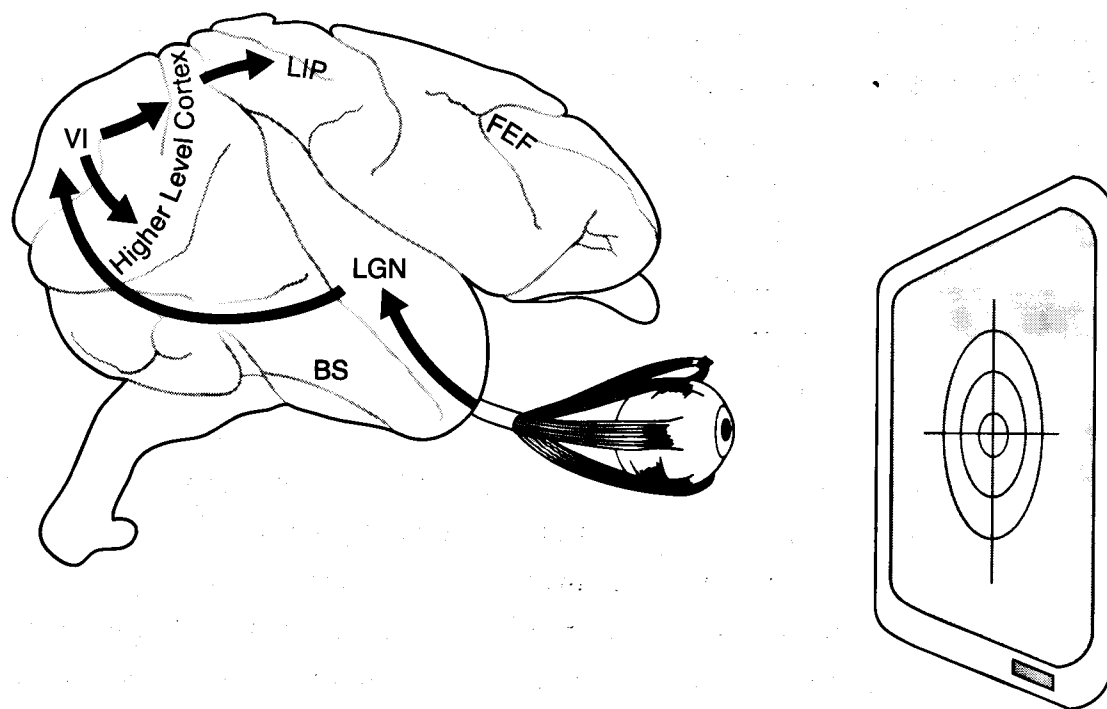
To understand how research on the sensory-to-motor functions of the parietal cortex has progressed, one has to begin with an outline of the visual and eye movement control circuitry in the primate brain. Only after this basic ground plan is clear can we move on to examine studies of the connections between these two sets of areas and to challenge conceptual approaches to the problem of how sensory and motor areas *ought* to be connected.

## **An Overview of Visual-Saccadic Processing**

### **Visual Processing in the Primate Brain**

Monkeys, like all mammals, receive their visual information from the right and left retinas. Lining the inner surface of the eyeball like a sheet of photographic film, each retina is a mosaic made up of about 1 billion photoreceptors.<sup>1</sup> The activity of these photoreceptors is processed within the retina and passed, by a class of neurons called retinal ganglion cells,

1. In humans.



**Figure 10.1**  
The principal pathways for visual processing.

through the optic nerve to the neurons of the lateral geniculate nucleus, or LGN. (See figure 10.1.)

The lateral geniculate nucleus in monkeys is a laminar structure, composed of six pancake-like sheets of neurons stacked on top of each other. Each sheet receives a topographically organized set of projections from one of the two retinas. This topographic organization means that at a particular location in, for example, the second layer of the lateral geniculate, all the neurons receive inputs from a single fixed location in one of the two retinas. Because individual locations in a retina monitor a single location in visual space (like an individual location on a photographic negative), each location in the geniculate is specialized to monitor a particular position in the visual world.

It has also been shown that adjacent positions within any given geniculate layer receive projections from adjacent positions within the referring retina. This adjacent topographic mapping means that each layer in the geniculate forms a complete and topographically organized screen on which retinal activity is projected. Each geniculate neuron thus has a receptive field, in the sense that Sherrington might have used the word.

Activation of that neuron occurs when a stimulus of the right kind falls on the topographically appropriate place in the retina.

These geniculate maps project, in turn, to the primary visual cortex. Lying against the back of the skull, the primary visual cortex, also called area V1, is composed of about 4 million neurons. These 4 million neurons form their own complex topographic map of the visual world; each square millimeter of tissue is specialized to perform a basic analysis on all the patterns of light that could fall on a specific region of the retina. Within these 1-mm-square chunks of cortex, individual neurons have been shown to be highly specialized. Some neurons become active whenever a vertically oriented boundary between light and dark falls on the region of the retina they monitor. Others are specialized for light-dark edges tilted to the right or to the left. Some respond to input exclusively from one retina; others respond equally well to inputs from either retina. Still others respond preferentially to colored stimuli. This complex pattern of sensitivities, or of receptive field properties, in area V1 is of tremendous conceptual importance. It suggests that information coming from the retina is sorted, analyzed, and recoded before being passed on to other visual areas.

The topographic, or retinotopic, map in area V1 projects, in turn, to a host of areas that also contain topographically mapped representations of the visual world. Areas with names like V2, V3, V4, and MT construct a maze of ascending and descending projections among what may be more than thirty mapped representations of the visual environment.<sup>2</sup> These networks of maps are the neural hardware with which we perceive the visual world around us.

While significant disagreement exists about the exact functions of each of these areas, nearly all neurophysiologists would agree that these are regions specialized for sensory processing. These areas may be heavily interconnected. They may do a tremendous amount of computational analysis. But what they do is to represent properties of the visual world. They serve as huge arrays of receptive fields, some of which are active only for the most specific combinations of visual events, but they are without a doubt sensory structures, or so most of us argue today. (For an overview of the primate visual system, see Reid, 1999.)


2. As measured in monkeys.

## Eye Movements and the Primate Brain

At the other end of the sensorimotor connection lies the circuitry that controls movements of the eyes, circuitry that activates and deactivates the six muscles that rotate each eyeball in its bony socket. Although all movements of the eyes are produced by these six muscles, eye movements can be broken into two fairly discrete classes. *Gaze-stabilization* movements shift the lines of sight of the two eyes to compensate precisely for an animal's self-motion; these movements stabilize the visual world on the retina as we move about in the world. *Gaze-aligning* movements point a portion of the retina specialized for high resolution, the fovea, at objects of interest in the visual world. These are the movements we use when we look at something. Gaze-aligning movements can be further broken down into two subclasses: saccades and smooth pursuit movements. Saccadic eye movements rapidly shift the lines of sight of the two eyes from one place in the visual world to another at rotational velocities up to  $1000^\circ/\text{sec}$ . They are the orienting movements of the eyes we use to look back and forth when examining a picture or a landscape. Smooth pursuit eye movements rotate the eyes at a velocity and in a direction identical to those of a moving visual target, stabilizing that moving image on the retina. These are the movements we use when we track a car moving along a road.

The best understood category of eye movements is undoubtedly saccades. At this time we know a tremendous amount about the interconnected brain areas that play critical roles in saccade generation. For this reason the saccadic motor control system seems to be an ideal place to try to understand motor control in general. (See figure 10.2.)

When a saccade is produced, the six muscles that control the position of each eye are activated by six groups of motor neurons that lie deep in the brainstem (BS). These alpha motor neurons are, in turn, controlled by two other systems also located in the brainstem, one that regulates the horizontal position of the eye while a saccade is in flight and one that regulates the vertical position of the eye in flight. These two control centers receive inputs from two interconnected saccadic control areas, the superior colliculus (SC) and the frontal eye field (FEF). Like the visual areas described above, the superior colliculus and the frontal eye field are constructed in topographic fashion. In this case their constituent neurons



**Figure 10.2**  
The principal pathways for saccade production.

form topographic maps of all possible eye movements. To understand how that works, imagine a photograph of a landscape. Now overlay a transparent coordinate grid that shows the horizontal and vertical eye movement that would be required to look directly at any point in the underlying photograph. Both the superior colliculus and the frontal eye fields contain maps very like these transparent coordinate grids. Activation of neurons at a particular location in the superior colliculus causes a saccade of a particular amplitude and direction to be executed. If this point of activation is moved across the collicular map, the amplitude and direction of the elicited saccade change in a lawful manner specified by the horizontal and vertical lines of the coordinate grid around which the map is organized. The neurons of the superior colliculus and the frontal eye field form topographically organized command arrays in which every neuron sits at a location in the map dictated by the direction and length of the saccade it produces.

I think that it would be uncontroversial (although not necessarily correct) to say that all of these saccadic control structures, from the frontal eye fields to the eye muscles themselves, are generally considered motor



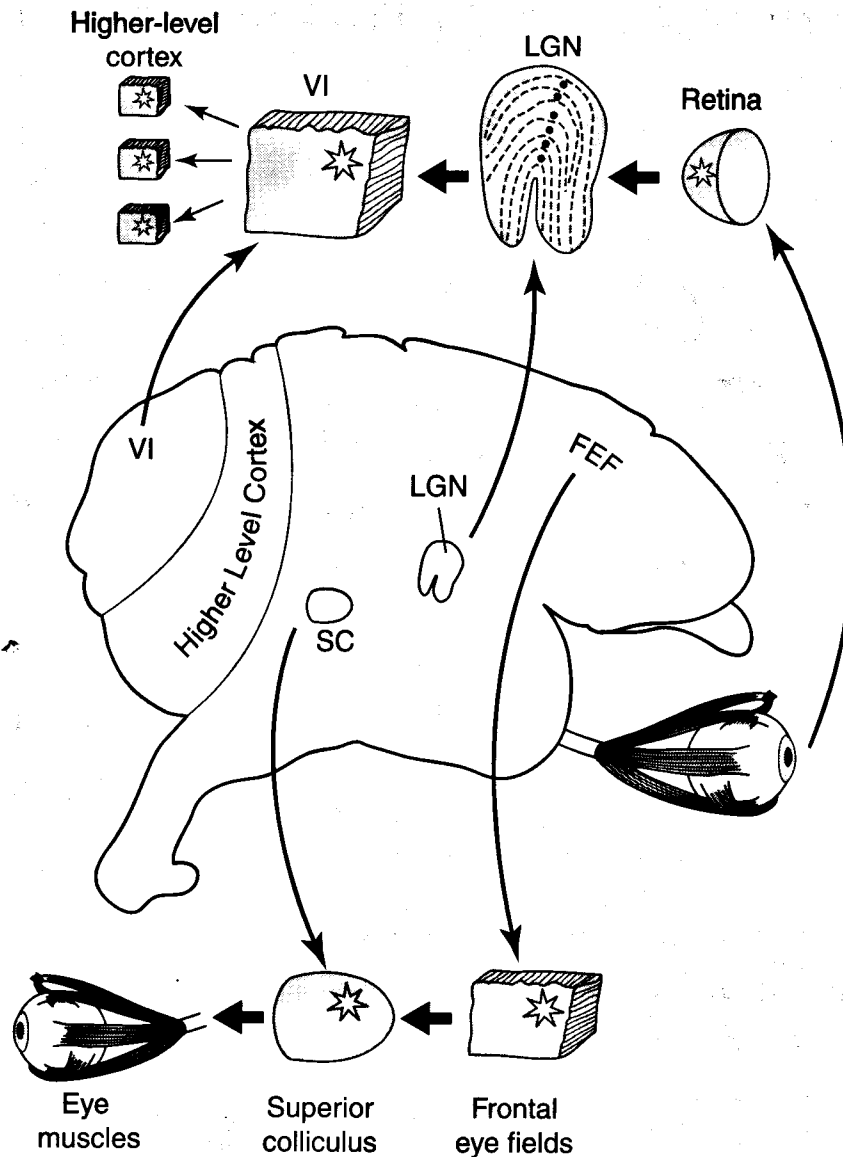
control areas. Many of them are interconnected in complex ways, and there is no doubt that each of these areas performs different and fairly complex computations, but most neurophysiologists would argue that these areas should be considered a final common path for saccade production. (For an overview of the eye movement control systems, see Glimcher, 1999.)

### **Linking Vision and Saccades**

Consider the following fairly typical experimental situation. A thirsty monkey is staring straight ahead, fixating a tiny yellow spot of light projected onto a screen 3 ft in front of him. At an unpredictable time, a red spot appears  $10^\circ$  to the right of where the monkey is looking and then the original yellow spot is extinguished. If the monkey makes a saccade that shifts his line of sight toward the red spot, he receives a squirt of Berry Berry Fruit Juice (which monkeys love) as a reward. When we perform this experiment, and the monkey orients toward the red spot, he does it within about a quarter of a second. How does the monkey's brain use the visual stimulus of the red light to produce the  $10^\circ$  rightward saccade in a quarter of a second? (See figure 10.3.)

We know that a light illuminated  $10^\circ$  to the right of straight ahead will activate a specific location in each retina. This activation will propagate to specific topographic locations among the six maps of the lateral geniculate nucleus and then on to a specific region in the primary visual cortex. From there the activity will spread among many of the interconnected maps of the visual cortices. At corresponding positions on each of these maps, the red target will give rise to activity in small clusters of neurons.

We also know that just before the monkey makes his saccade, the position on the collicular and frontal eye field maps associated with a  $10^\circ$  rightward movement will become active. This activity, in turn, will lead to activation of the neurons of the horizontal saccadic control center in the brain stem. The horizontal control center will then activate the two clusters of alpha motor neurons that regulate the tension on the lateral and medial rectus muscles of each eye, causing the eyeballs to rotate.



**Figure 10.3**

Patterns of neural activity when a target to the right elicits a rightward eye movement.

So how does the activity associated with the onset of the red light precipitate the eye movement? The simplest possible answer to that question is that the neurons of the motor system must be activated by the neurons in the visual system, via some “simple definite path.” If we think of the saccadic control areas as the final common path for saccades and we conceive of the visual-sensory areas as a very elaborate Sherringtonian receptor, then the connection between them can be thought of as the product of an internuncial neuron (or neurons) connecting these two systems. Conceptualized this way, it seems reasonable to ask whether or not we can find evidence of a connection between the sensory and motor

systems of the primate visual saccadic system that is active when this simple visual-saccadic response is produced. This linking connection must be the internuncial element that yields the sensorimotor response.

One way to begin looking for that connection would be to identify, at a purely anatomical level, those areas of the brain which actually connect the visual cortices with the frontal eye field and the superior colliculus. When one examines the anatomical connections between the visual and motor systems, a number of areas in posterior parietal cortex stand out. One in particular has been the subject of tremendous attention, the lateral intraparietal area, or area LIP.

### The Visual-Saccadic Function of Parietal Cortex

By the 1950s most of parietal cortex had come to be called an association area in neuroscience textbooks. At that time, cortex came in essentially three flavors: sensory, motor, and association. The word *association* had come to neuroscience from the study of conditioned reflexes, where Pavlov had used it to refer to the process by which a stimulus comes to be associated with the production of a new response.

At that time, studies of neurological patients with damage to the association areas of parietal cortex had already made it clear that the bulk of the parietal lobe is not necessary either for sensory processing or for the production of movement. Patients with parietal damage were known to see and hear fairly well. They also could move their limbs and both eyes in a fairly normal manner. But these patients seemed to be impaired on tasks in which they needed to use one or more of their senses in order to plan a movement. These observations, and others like them made in monkeys, had led scientists to conclude that parietal cortex was something neither exactly sensory nor exactly motor. Inescapably, this led to the conclusion that it must be used to generate sensorimotor associations. Unfortunately, a more precise theory of what the parietal association areas did was lacking.

This almost complete ignorance about parietal association areas was, in large part, due to a serious technical limitation. In the 1950s it was possible to record the activity of single nerve cells within the brains of anesthetized animals. This was a critical technology for physiologists who wished to study sensory systems, because it allowed them to deter-

mine what role different groups of neurons played in the processing of sensory data. In a typical experiment these physiologists might stroke the forearm of an anesthetized cat while using microelectrodes to trace through the nervous system the pattern of activity produced by that stroking. In a similar way, physiologists could study motor control in anesthetized animals by electrically activating small groups of nerve cells with microelectrodes. This allowed them to determine how these cells influenced muscle tension. But association areas simply could not be studied in anesthetized animals because anesthetized animals could not make associations (or do anything else, for that matter). In the 1950s it was not possible to study single nerve cells in conscious animals. The technology simply did not exist, and this was a critical limitation because it meant that association areas could not be studied in a meaningful way.

In the late 1950s two researchers, Herbert Jasper at the Montreal Neurological Institute in Canada and Edward Evarts at the U.S. National Institutes of Health, took on this critical technical problem. Both men were convinced that neurobiologists had to be able to record the activity of single neurons within the brains of conscious, behaving animals without injuring the animals or disrupting their normal behavioral patterns. This seemed possible because it was known that the brain itself possesses no pain or touch receptors. Tiny wires inserted into the brains of conscious animals would therefore be undetectable by the animals. Jasper hoped to use this approach to study the activity of single nerve cells while monkeys formed memories in a normal fashion. Evarts hoped to understand how arm movements were produced by studying the activity of neurons in motor cortex.

In 1958 Jasper published the first report of this technique, demonstrating unequivocally that recording from the brains of conscious animals was both possible and well tolerated by the animals. Evarts then took the technique a step farther, developing and largely perfecting a system almost identical to the one in widespread use today. The development of this technique was critical because it meant that the parietal association areas could finally be studied in action.

### **The Command Hypothesis**

The first laboratory to use this new technology to study parietal cortex in action was Vernon Mountcastle's research group at the Johns Hopkins

University in Baltimore. Mountcastle had made a name for himself by using microelectrodes to study the somatosensory cortices of anesthetized animals, the sensory cortex responsive to touch stimuli. He and his students had performed landmark studies that laid the foundation both for our modern theory of how the cerebral cortex in general works and for our modern theory of how the somatosensory system in particular works. Mountcastle was at the center of the neuroscience establishment, and his decision to use single-neuron recording in conscious animals to study an association area was a critical step. It gave the stamp of approval from the scientific establishment for physiologists to use what was coming to be called the *awake-behaving preparation* to do what is now called *cognitive neuroscience*.

Mountcastle and his team of young physiologists (Mountcastle et al., 1975) proceeded by training monkeys to sit quietly while visual and tactile stimuli were presented to them. They also trained the monkeys to use specific visual stimuli as cues to produce behavioral responses that, if performed correctly, would yield rewards. For example, a monkey might be trained to reach toward a light whenever the light was illuminated, regardless of where the light was positioned in the space around the monkey. If the animal did reach for the light, she would receive a drink of fruit juice. Or the monkey might be trained to reach out and touch a button as it moved past along a track. In 1975 Mountcastle's group published a landmark paper that reported the results of these experiments. (See figure 10.4.)

The major conclusion we draw from the observations described above is that there exist within the posterior parietal association cortex sets of neurons which function as a command apparatus for the behavioral acts of manual and visual exploration of the immediately surrounding extrapersonal space.

...

The projection and hand-manipulation neurons of the parietal cortex appear to be neither sensory nor motor in nature, but stand in a selective *command* [my italics] relation to movements of a particular sort.

...

We wish to emphasize the following: that our concept of command centers explicitly assumes that there exist within the central nervous system many sources of commands to the motor apparatus. The source of the command and its nature will differ remarkably in different behavioral reactions, even though the peripheral musculature engaged in the different acts may be virtually identical.

...




Figure 10.4

Mountcastle's experimental setup. (From Mountcastle, Lynch, Georgopoulos, Sakata, and Acuna, 1975, Posterior parietal association cortex of the monkey: Command functions for operations within extrapersonal space. *J. Neurophys.* 38: 871-908.)

We propose that several of the abnormalities of function that occur in humans and in monkeys after lesions of the parietal lobe can be understood as deficits of *volition*, of the will to explore with hand and eye the contralateral half-field of space, a deficit caused by the loss of the command operations for those explorations which exist in the parietal association cortex. (Mountcastle et al., 1975)

These single-neuron studies led Mountcastle and his colleagues to suggest that a nonsensory and nonmotor process was taking place in parietal cortex, something they named a command process. While they argued that this process should be considered neither sensory nor motor in nature, it seems clear that they viewed it as more closely related to the generation of a movement than to the processing of sensory signals. The name "command center" alone makes this clear.

Shortly after the Mountcastle lab published this manifesto arguing for the existence of a parietal command process, evidence began to accumulate that would challenge their proposal. Initially, this challenge came from Michael Goldberg (then at the U.S. National Institutes of Health), a scientist who had been developing an alternative conceptual approach to understanding how sensory signals could act to precipitate movements.


### **Attentional Enhancement**

In the years immediately before Mountcastle published the command hypothesis, Goldberg had worked with his colleague and mentor, Bob Wurtz, to perform the first characterization of eye movement-related activity in the superior colliculus, using the awake-behaving monkey technique. It was a technique that Wurtz had learned from Edward Evarts while Wurtz was himself a postdoctoral fellow at the U.S. National Institutes of Health. (See figure 10.5.)


In Goldberg and Wurtz's first experiments on neurons in the superior colliculus, monkeys had been trained to perform two behavioral tasks. In the first, which they called a fixation trial, monkeys were trained to stare at a visual stimulus located straight ahead while a secondary stimulus was illuminated at some other location. The secondary stimulus would be abruptly extinguished and the monkey would be rewarded for entirely ignoring the existence of this secondary visual stimulus. In a second task, which they called a saccade trial, while the monkeys were staring straight ahead at the central stimulus, a secondary stimulus would be illuminated and the central stimulus would then be extinguished. During these trials the monkey was rewarded if she looked to the secondary target after the central target was extinguished.

Each day, as soon as Goldberg and Wurtz had placed their electrode near a single collicular neuron, they would begin by identifying the location of that neuron in the collicular topographic map. This would allow them to place the secondary stimulus at the precise location in visual or motor space for which that neuron was specialized. Once the secondary stimulus had been fixed at that location, they had the monkey execute a set of fixation trials followed by a set of saccadic trials. On fixation trials they found that the collicular neurons became weakly active when the secondary stimulus was illuminated, and that this weak activity petered out after the secondary stimulus was extinguished. When

### Fixation Task



### Saccade Task



**Figure 10.5**

Goldberg and Wurtz's fixation task and saccade task. The oscilloscope beside the monkey displays activity in the superior colliculus during each phase of the task.

they then switched the monkey to saccade trials, they observed that the weak initial response of the neuron grew vigorous late in the trial, just before the monkey made her eye movement. But what surprised them was that as the animal executed saccade trial after saccade trial, the initial response to the onset of the secondary target grew stronger and stronger. It was as if the more certain the monkey grew that the secondary stimulus would be the target of a saccade, the more vigorous was the initial neuronal response.



Because the size and timing of this enhanced initial response had no obvious effect on the movement, Goldberg and Wurtz reasoned that this initial activity must not be a movement control signal. Their observation that the strength of the initial response grew from trial to trial even though the stimulus stayed identical from trial to trial led them to conclude that the enhanced early response must not be purely sensory in nature either.

So what was this signal that was neither sensory nor motor and that they had described several years before Mountcastle's lab had published the command hypothesis? Goldberg and Wurtz had seen no need to postulate anything like Mountcastle's command function to explain their results. Instead, they argued that this enhancement could be viewed as sensory signal that had been modified by attentional factors.

Their basic idea was that whenever the monkey was planning to look at a particular secondary target, she must be paying greater attention to the onset of that target than when she was ignoring it during the fixation task. They reasoned that as the monkey performed saccade trial after saccade trial, her certainty about the importance of that particular secondary target would grow, and she would therefore pay better and better attention to it. Goldberg and Wurtz hypothesized that this attention might somehow increase the strength of the initial neuronal response in the superior colliculus, and this increase in initial response strength was exactly what they had observed. Attention, in their analysis, was a mechanism that could produce changes in sensory responses that rendered these neural signals not really sensory and not really motor, making them enhanced representations of the sensory world.

Mountcastle and his colleagues had proposed the existence of a command process in posterior parietal cortex to explain the existence of signals that were neither truly sensory nor truly motor, and Goldberg saw a close relationship between these parietal signals and the enhanced collicular signals that he had examined. As a result, Goldberg and two new colleagues, David Lee Robinson and Gregory Stanton, immediately began a search for enhancement-like effects in posterior parietal cortex, using the experimental approach that Goldberg had pioneered in the colliculus. These experiments were completed by 1978 (Robinson, Goldberg, and Stanton, 1978) and replicated many of Mountcastle's

findings, but because of Goldberg's earlier work they led to very different conclusions.

Since in their [Mountcastle's] experiments neurons associated with movement did not seem to be excitable by passive visual or somatosensory stimulation, these authors postulated that parietal cortex performs a command function for hand movements and eye movements exploring the visual and somatosensory environment. They stressed that neurons in posterior parietal cortex did not have sensory responses, and hypothesized that this area of the brain provided holistic command signals for the motor system. They proposed that this holistic command view should supplant the more traditional associative view of posterior parietal cortex.

In previous studies of visually responsive neurons in the superior colliculus and frontal eye fields, we showed that their visual response was enhanced when a stimulus in the receptive field was going to be the target for an eye movement. We decided to examine neurons in posterior parietal cortex to see whether cells which are associated with visually guided eye movements could be better understood as behaviorally modifiable visual neurons rather than "command neurons." In this study we recorded from 289 cells in area 7 [the saccade-related portion of posterior parietal cortex]. Every neuron which was associated with movement had a sensory response that could be demonstrated in the absence of the movement. In many cases the sensory response was *enhanced* [my italics] when the animal was going to use the stimulus in the receptive field as a target for an eye or hand movement.

...

These experiments show that posterior parietal cortex should be viewed as a *sensory association area*. Cells here integrate visual and somatosensory information from the environment with behavioral data which are presumably generated internally. The behavioral information serves to modify the sensory input. If a stimulus is important, such as a target for a movement, that stimulus will have a greater effect on the nervous system than an equivalent stimulus that is unimportant.

Mountcastle and his co-workers have postulated that neurons in posterior parietal cortex perform a command function for eye movements and hand movements. They described neurons that discharged in association with movement but which they could not drive with passive stimuli. We found that every neuron that we encountered in posterior parietal cortex that could be associated with movement could also be excited by some sensory stimulus independent of movement.

...

We propose that parietal neurons are best described according to their sensory properties, not according to epiphenomenological movement relationships. (Robinson, Goldberg, and Stanton, 1978)

At some level, Goldberg and Mountcastle were vying to define the neurophysiological model that would serve as a template for understanding the connection between sensory and motor signals. Mountcastle

was arguing that a new process would be required at a conceptual level to understand this connection, a process that was neither motor nor sensory, a command process. Goldberg was arguing that the connection between sensation and movement in parietal cortex could be analyzed using existing ideas about associating sensory and motor signals. The linkage between sensory and motor signals could be accomplished by a modulation of sensory signal strength. When the sensory signal was boosted above some critical level by an attentional process, the final common path was activated.

At heart, Goldberg's proposal embodied Pavlov's dictum of necessity. Under some task conditions, sensory activity elicited by a stimulus gives rise to a movement. Under other conditions it does not. Linking elements produce this effect by responding more strongly on trials in which a movement should be produced. The more likely a stimulus is to produce a movement, the stronger will be the enhanced sensory response, thus increasing the likelihood that the final common path will be engaged. These linking elements reflect the fact that "If a stimulus is important, such as a target for a movement, that stimulus will have a greater effect on the nervous system than an equivalent stimulus that is unimportant." The virtues of this approach were its simplicity and the fact that it so clearly embodied Pavlov's notion of necessity. Goldberg's model was an extrapolation of the Cartesian approach to the parietal cortex and to the cortical sensorimotor problem in general.

Mountcastle's approach was fundamentally different. He was arguing for an abandonment of the reflexological approach when thinking about posterior parietal cortex. He was arguing instead for the existence of an entirely new kind of sensorimotor component. But Mountcastle's model was incomplete at best. What, exactly, were these command functions? How did they work? What computational problem were they trying to solve?

At the time that Mountcastle and Goldberg were having this debate, the publication of Marr's book was still several years away. So the need for a rigorous model of exactly what it was that parietal cortex *should* be doing may not have been as clear as it is today. And in terms of the actual features of such a model, probability theory and economic approaches were just beginning to enter biological studies through ecology departments; in the years between Mountcastle's and Goldberg's papers, Eric

Charnov published the first of his formal models. So perhaps for all of these reasons, the debate quickly began to circle away from a direct analysis of the general models being debated and instead focused on experimental challenges to specific claims Goldberg and Mountcastle were making.

### Attention Versus Intention

There is absolutely no doubt that Goldberg and his colleagues demonstrated that the notion of a command function added little to our immediate understanding of parietal cortex. A monkey is presented with an eccentric visual stimulus; sometimes he looks toward that stimulus and sometimes he does not. At least on the surface this seems an indeterminate process, a process that should be hard to explain using classical Cartesian models. One might suspect that an additional indeterminate process would be required to explain such a phenomenon, but Goldberg's idea brilliantly resolves the paradox of this apparent indeterminacy without resorting to an explicitly indeterminate event. The direct connection of sensory signals to motor control circuits is modulated by a gating system, sensory attention. Sensory attention, through the physiological mechanism of enhancement, serves to allow some sensory signals access to the motor plant and to deny that access to others. The process appears indeterminate only because we do not understand how the attentional process operates.

The Goldberg model really could account for what physiologists had observed. By the late 1970s even Mountcastle's laboratory began to encounter evidence that there were signals in parietal cortex which seemed more closely associated with sensory events, as Goldberg had proposed, than with motor events. For this reason, and doubtless for others, Mountcastle's group began to turn their electrodes away from posterior parietal cortex, returning to somatosensory cortex. But before they returned entirely to somatosensory cortex, Richard Andersen would have to complete his postdoctoral studies. (See figure 10.6.)

When he joined the Mountcastle lab, Andersen was interested in understanding how circuits that interface between the sensory and motor nervous systems encode the locations of objects and the trajectories of movements. Imagine, as Andersen did, a monkey looking straight ahead while a visual stimulus is presented 5° degrees to the right of where he is

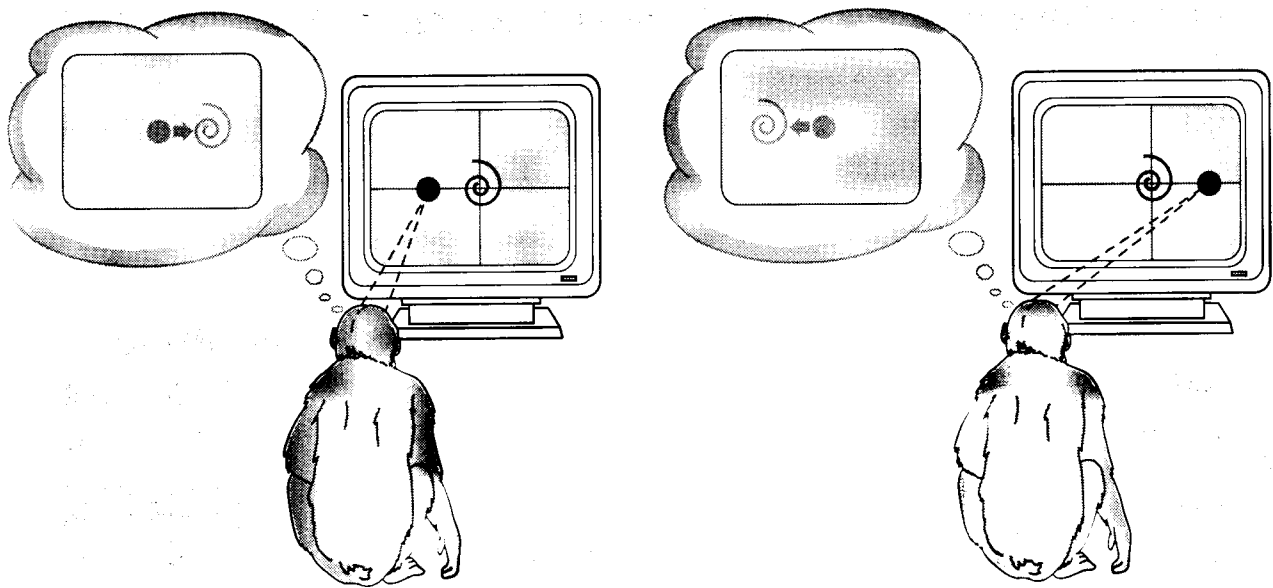


Figure 10.6  
Andersen's coordinate transform experiments.

looking. The sensory stimulus activates neurons at a position  $5^\circ$  to the right of straight ahead on the topographic maps of the retina, the lateral geniculate, and the visual cortices. In order to make a movement that aligns his line of sight with that stimulus, the monkey rotates his eye until it is oriented  $5^\circ$  to the right of straight ahead; he activates the  $5^\circ$  rightward locations on the topographically organized motor control maps.

Next consider the case when the same visual stimulus is illuminated at the same place in the world, but now the stimulus is presented while the monkey is looking  $10^\circ$  to the right of straight ahead. Under these conditions the same light, at the same location in the world, now activates a point on the retina  $5^\circ$  to the *left* of straight ahead. Positions in all of the subsequent sensory maps representing  $5^\circ$  left must therefore be active. But if the monkey wants to look at that light, he must still reorient his eyeball so that it comes to rest at a position  $5^\circ$  to the right of straight ahead.

The problem that Andersen became interested in was this: Many sensing surfaces, like the retina, the skin, or the ears of a cat, move around in the world. This means that the alignment between the topographic maps of the sensory systems and the topographic maps of the motor systems that they must control can shift. Andersen realized that one critical function of any system connecting sensory and motor signals, regardless of how it accomplished that linkage, would be to deal with this problem.

Andersen's work in Mountcastle's lab was largely focused on understanding how sensory and motor maps were aligned, and this interest led him into area 7, a subregion of the posterior parietal cortex specialized for sensorimotor interactions and associated with saccadic eye movements. That work reached a watershed several years later when Andersen, by then a professor at the Salk Institute, published an influential paper with Greg Essick and Ralph Siegel titled "Neurons of Area 7 Activated by Both Visual Stimuli and Oculomotor Behavior" (Andersen, Essick, and Siegel, 1987).

Several laboratories have made recordings of the activity of single neurons in area 7a [a subregion within area 7] of the posterior parietal cortex in behaving monkeys and, by correlating the activity of these cells with sensory and motor events, they have made important advances in understanding its functional role. In the earliest experiments Mountcastle and his colleagues found that many of the cells were activated by certain behaviors of the animal, including saccadic eye movements, fixations, smooth pursuit eye movements, and reaching movements of the arms. It was stated that these cells did not respond to visual or somatosensory stimuli. On these grounds it was proposed that area 7 was involved in issuing general motor-commands for eye and limb movements.

In later experiments Robinson, Goldberg and colleagues found that many of the cells in area 7 responded to visual or somatic stimuli. They argued that the behaviorally related responses reported by Mountcastle and his colleagues could be accounted for either by visual stimulation from the target for movement or from visual/somatosensory stimulation resulting from the movement. It was proposed that area 7a was involved in sensory processes and did not play a role in motor behavior as proposed by Mountcastle and colleagues. In a later report Motter and Mountcastle noted some cells which appeared to be oculomotor and light sensitive and proposed that a gradient existed between cells with strictly eye movement related responses and cells with solely visual responses.

In the present study we have designed experiments to distinguish between visual and motor components of the responses of the fixation and saccade neurons and have found that the activity of the cells in these two classes is related both to sensory stimuli and to oculomotor behavior. The nature of the eye movement and fixation (eye position) signals suggests that they play a role in establishing spatial constancy rather than in the initiation of oculomotor behavior. (Andersen, Essick, and Siegel, 1987)

Andersen and his colleagues went on to argue that the most fundamental role of posterior parietal cortex was to solve a problem that was much more clearly defined than how sensory and motor systems are interconnected. As the eyes, limbs, and head move, the relationship be-

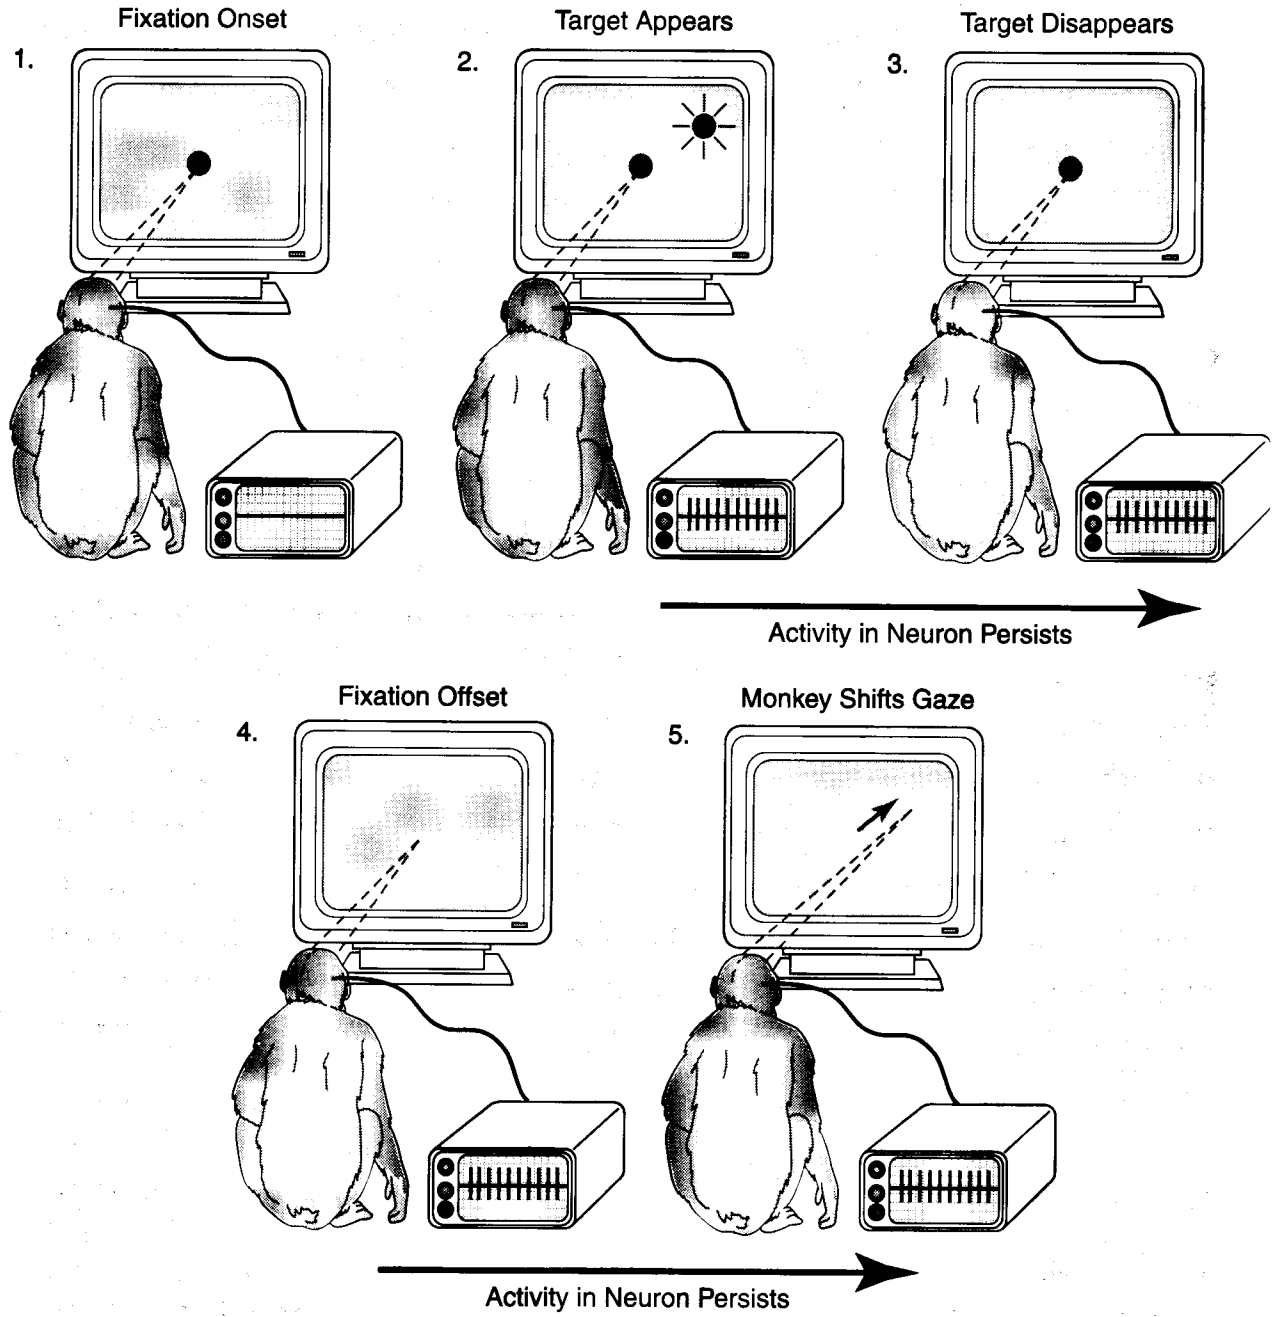
tween topographic maps of the sensory world and the topographic maps of the motor systems is constantly shifting. Regardless of how we decide which stimuli should trigger a response, we must somehow decide which response that stimulus should trigger. It was this issue—how to link the shifting sensory and motor topographies of the brain—that Andersen identified in 1987 as a central function of posterior parietal cortex.

But in the next year Andersen and his colleagues discovered an unexpected property in parietal cortex that forced them to revisit the command hypothesis. Working as a postdoctoral fellow in Andersen's laboratory, Jim Gnadt trained monkeys to perform a remembered saccade task (Gnadt and Andersen, 1988; see figure 10.7). In that task, while the monkey stared at a central stimulus, a secondary visual stimulus was briefly flashed on and off. The monkey's job was to look at that secondary stimulus, but only after Gnadt turned off the central stimulus. And Gnadt would often leave the central stimulus on for quite a while *after* the secondary target was extinguished. The monkeys had, in effect, to remember what movement they were supposed to make during this enforced delay, which could be up to 1.5 sec. When Gnadt and Andersen used this task to study the activity of neurons in the saccade region of area 7, a region that had by now acquired the name *area LIP*, they made an amazing discovery. They found that neurons activated by the brief flash of the secondary stimulus remained active after the visual stimulus went out, and that this activity persisted until the monkey made the saccade which aligned gaze with the *remembered* location of the stimulus.

We have shown that there is memory-linked activity in the lateral bank of the intraparietal sulcus which is associated with saccadic eye movements of specific direction and amplitude. The activity can be disassociated from the visual stimuli which guide the eye movements. . . . Therefore, these cells' activity appeared to be related to the pre-movement planning of saccades in a manner which we have chosen to describe as *motor intention* [my italics]. The term motor intention, as we use it here, is meant to convey an association between the behavioral event (i.e., saccade) and the neural activity. It is not meant to suggest that this neural signal is necessary and sufficient to produce the eye movement.

The posterior parietal cortex of primates contains neural signals appropriate for the building of spatial maps of visual target space, memory-linked motor-planning activity and possible corollary feedback activity of saccades. These findings strongly argue that the parietal cortex is intimately involved in the guiding and motor planning of saccadic eye movements. (Gnadt and Andersen, 1988)

### Memory Saccade



**Figure 10.7**  
Gnadt and Andersen's memory saccade experiment.



Andersen had, in a single salvo, revived a version of the command hypothesis in a new and improved form. He had demonstrated that Goldberg's *enhanced* activity could persist in the absence of a sensory stimulus. This, Andersen proposed, was incompatible with Goldberg's assertion that "parietal neurons are best described according to their sensory properties."


Goldberg and two new colleagues, Carol Colby and Jean-René Duhamel, responded to this new finding by revisiting area LIP. They wanted to know if the results reported by Gnadt and Andersen could be reconciled with their view that posterior parietal cortex was fundamentally a sensory structure. If it could not be reconciled with this view, they wanted to understand exactly what role area LIP played in the sensory-to-motor process.

To begin their new round of investigations, Colby, Duhamel, and Goldberg (1996) taught a new group of monkeys to perform yet another set of tasks. (See figure 10.8.) Of course they taught their monkeys to perform the memory saccade task that Gnadt and Andersen had used, but they also taught them an important pair of tasks that were highly influential and that led to a reinterpretation of Gnadt and Andersen's findings. These two tasks were the fixation task Goldberg had used before, and a new task they called the peripheral attention task. You will recall that in the fixation task the monkey was taught to stare straight ahead at a continuously illuminated central light while a secondary stimulus was turned on. The monkey's job was to ignore the secondary stimulus.


The peripheral attention task was almost identical. Once again the central light illuminated. Once again the monkey was required *not* to make a saccade. But in this task, while the monkey was staring straight ahead and the secondary stimulus was illuminated, the monkey was required to press a lever if she saw the secondary stimulus dim. Goldberg and his colleagues argued that the two tasks were identical from an eye movement control point of view; in neither case was an eye movement produced. The two tasks differed critically, however, in where one might expect the monkey to pay attention. In the fixation task, the eccentric target in the response field was irrelevant, but in the peripheral attention task this same stimulus became highly relevant.

What Goldberg and his colleagues found when they did this experiment was that parietal neurons were much more active in the peripheral

### Fixation Task



### Peripheral Attention Task



**Figure 10.8**

Colby, Duhamel, and Goldberg's peripheral attention task.

attention task than in the fixation task. Even though neither of the tasks required a movement, they elicited very different responses in LIP neurons. To many scientists this seemed clear evidence that activity in area LIP, and by extension activity in posterior parietal cortex in general, could not simply be a command to make a movement, nor could it reflect the intention to produce a movement. Instead, it seemed that this LIP activity must be correlated with the sensory events to which an animal was paying attention. In communicating these results, however, Colby, Duhamel, and Goldberg took a fairly conciliatory line, perhaps because their relationships with Andersen and his colleagues were growing very strained over this continued dispute.

First, we found that the presaccadic enhancement originally described in area 7 is a *specific enhancement of the visual response* [my italics] to stimulus onset. Second, this enhancement of the visual response in a saccade task is correlated with the degree of enhancement in a purely attentional [movement free] task. [The authors then go on to state two additional findings suggesting that activity immediately before the saccade may well be associated with the movement being produced.]

...

The present results show that LIP neuron activity is multi-faceted and subject to modulation by cognitive factors such as attention and anticipation. LIP neurons have independent sensory responses and saccade-related bursts. Further, attention to a spatial location modulates the sensory response to stimulus onset, and anticipation of a behaviorally significant sensory event affects the level of baseline neural activity.... [These neurons] cannot be regarded exclusively as analyzing visual stimuli because they discharge before saccades even when there has been no recent visual stimulus. They cannot be regarded exclusively as planning saccades because they are strongly responsive in a task in which saccades are expressly forbidden. One way of understanding this varied collection of activations is to consider their point of intersection: the spatial location defined by the preferred stimulus location and the preferred saccade for a given neuron. We suggest that LIP neuron activity encodes events related to a particular spatial location. (Colby, Goldberg, and Duhamel, 1996)

Andersen and his colleagues responded to these new data with a surprising, and quite interesting, hypothesis that was meant to further challenge Goldberg's position. They proposed that during the attentive fixation task Colby had studied, her monkeys were actually *planning* to look at the secondary stimulus even if they did not in the end make that movement. Colby's monkeys, they reasoned, had seen tens of thousand of secondary stimuli that were supposed to precipitate saccades. Perhaps the strength of activity in LIP reflected not attentional enhancement but the animal's certainty about whether or not it *intended* to make to make that saccade.

To validate this hypothesis, Andersen and one of his students, Martyn Bracewell, trained monkeys to perform yet another task. This time, monkeys would begin by staring straight ahead at a central light. Again a secondary stimulus would flash briefly. As in the remembered saccade task, the monkey's job was to wait patiently until the central light was turned off and then to look at the location where the secondary target had appeared. But on some trials, during this waiting interval, Bracewell and Andersen flashed a third light at yet another location. When that happened, the animal had been taught to change her plan. Now when

the fixation light went off, she was to look at the location where the tertiary stimulus had appeared.

What Bracewell and Andersen found was that when they told the monkeys to change their plans, the site of activation on the map in area LIP shifted. Neurons associated with a movement to the secondary target became silent and neurons associated with a movement to the tertiary target became active. The very existence of activity associated with the secondary target, they noted, indicated that neurons in area LIP could encode a movement that a monkey was planning to make, even when that movement was never produced.

Andersen and his colleagues saw this as at least a partial refutation of Colby, Duhamel, and Goldberg's conclusions. Perhaps during the peripheral attention task monkeys were planning a movement that they never executed, just as had been observed in the change of plan task. Goldberg, quite reasonably, responded to this by pointing out that the Andersen group's change of plan task data could also be interpreted as evidence that the monkey was shifting where he was paying attention, shifting his attentional enhancement from the secondary to the tertiary target. These data, he argued, proved nothing about the existence of a motor plan or an intention.

### Resolving the Attention-Intention Debate?

To many of us on the outside, the Andersen–Goldberg debate over the function of parietal cortex seemed to have sunk almost to a semantic contest. It was clear that Goldberg had been forced to largely abandon his original hypothesis that “posterior parietal cortex should be viewed as a sensory association area.” By the same token, Andersen had been forced to abandon the original command hypothesis of Mountcastle in favor of the much weaker intention hypothesis. What, if anything, could all of this tell us about how the nervous system solves the problem of connecting sensory inputs with motor outputs?

At about this time Michael Platt and I, naively, became convinced that we could refocus the parietal debate by designing a single experiment that would directly pit the Goldberg sensory interpretation against the Andersen motor interpretation. What if we could design an experiment


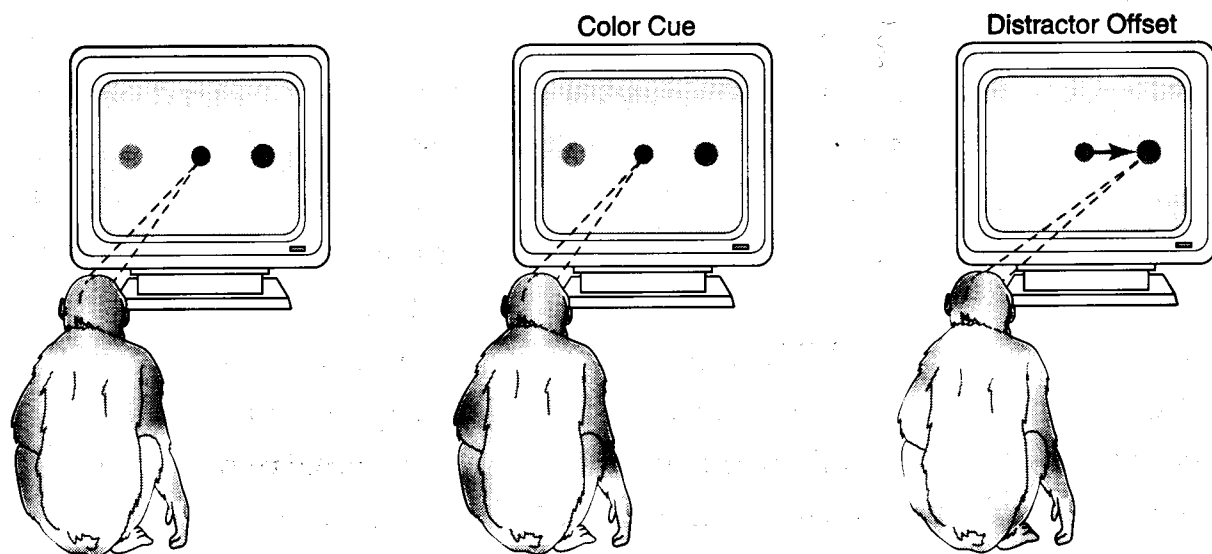
in which the movement plan of Andersen's hypothesis and the attentional enhancement of Goldberg's hypothesis were under independent control? Then we could simultaneously test the two competing hypotheses that (1) area LIP carries sensory-attentional signals and that (2) area LIP carries motor intention plans. To do this, we designed yet another set of tasks, tasks that combined the attentive fixation approach Colby, Goldberg, and Duhamel had used with the change of plan task Bracewell and Andersen had used.

I need to pause here to say that this was, to put it simply, a silly idea. Platt and I were doing nothing more than trying to replicate what had already been done. We had stepped right into the middle of this fractious semantic debate with yet another task and without any new theoretical insights. The experiment that we conducted should have earned us no more than a footnote in a history of the Goldberg–Andersen debate. The reason I want to tell you about this last classically styled parietal experiment is that the results it produced convinced us that deciding to do this experiment had been silly. To understand how it convinced us of this, I have to tell you a bit more about the experiment itself.

### **The Cued Saccade and Distributed Cue Experiment**

We reasoned that we would need to design an experiment in which we could attempt to simultaneously falsify each of these two competing hypotheses: the hypothesis that area LIP carries attentional signals and the hypothesis that area LIP carries intentional signals. If our experiment falsified only one of these hypotheses, this would provide support for the other model of parietal cortex and would validate the notion that area LIP could be understood within the framework of that model (Platt and Glimcher, 1997). (See figure 10.9.)

We began by training monkeys to perform, in sequential blocks of trials, two tasks that independently controlled both the location of a secondary stimulus the monkey would be required to look at, and the location and behavioral relevance of a tertiary stimulus at which the monkey would never look. At the start of our experiment each day, and before having the monkey perform the actual experiment, we isolated a single neuron in area LIP for study and had the monkey make a series of eye movements that shifted his gaze into alignment with secondary

**Cued Saccade****Distributed Cue**

**Figure 10.9**  
The cued saccade and distributed cue experiment.

stimuli presented sequentially at a large number of locations. We used these data to identify the location of our neuron in the area LIP topographic map, the *best location* for this neuron. We also used these data to identify a stimulus and movement for which the neuron was inactive, a null location for this neuron.

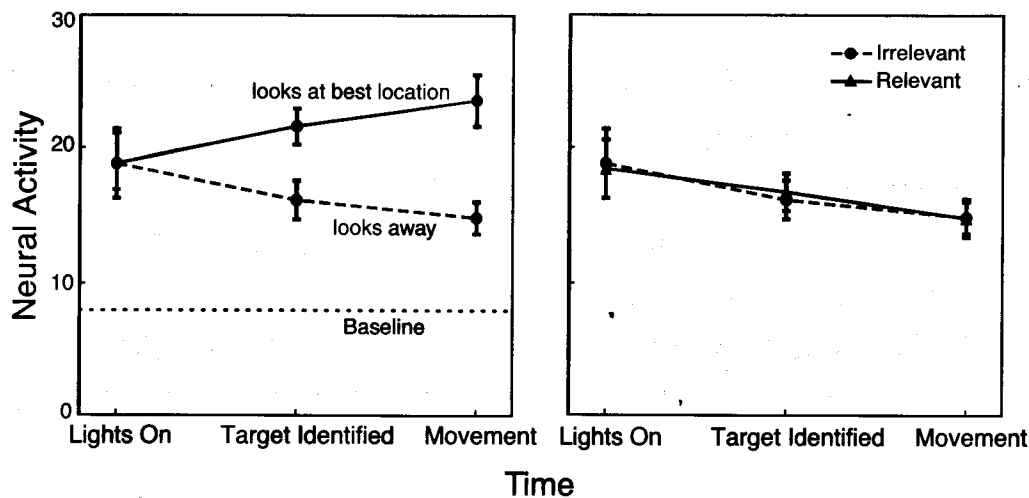
Once these two locations had been identified, we next had the animal perform a block of cued saccade trials, shown in figure 10.9. Each cued saccade trial began with the illumination of a central yellow light at which the monkey had to look. After a brief delay, the secondary and

tertiary targets were illuminated, one at the best location and one at the null location. After a further delay, the central yellow light changed color. On a randomly selected 50 percent of trials it turned green. On the other trials it turned red.

The monkey had been taught in advance that on trials in which the central fixation stimulus turned red, the left light served as the secondary target (the saccadic goal) and the right light served as the tertiary target (a completely irrelevant distractor). On trials in which the fixation stimulus turned green, the converse was true; the right light served as the target and the left light was irrelevant. The monkey was, however, not allowed to look at the secondary target until we turned off the central fixation stimulus. If, after that point, the monkey looked at the correct target, she received a fruit juice reward.

Amazing though it may seem, the monkeys readily learned this task. That made it possible for us to test the motor-intentional hypothesis by comparing two classes of trials that were nearly identical in their visual properties but differed profoundly in their movement properties. In both classes of trials, two eccentric targets were illuminated, one at the best location and one at the null location. The trials differed in the meaning of the eccentric stimulus at the best location. On 50 percent of these trials, the stimulus at the best location had served as the saccadic goal, and in the other 50 percent it served as a totally irrelevant visual distractor. We reasoned that if neurons in area LIP were motor-intentional elements, they should respond strongly when the stimulus within the response field was a target, but *not at all* when it was an irrelevant visual distractor.

The left panel of figure 10.10 shows the average firing rate of forty area LIP neurons during these two classes of trials at three points in time: just after the two eccentric targets turn on, after the fixation stimulus had changed color (identifying the saccadic goal), and just before the saccade began. The solid line shows the activity of the neurons when the fixation stimulus turned red and the monkey looked at the best location. The dashed line in the left panel shows the response of the neurons on trials in which the fixation stimulus turned green and the monkey looked away from the best location.



**Figure 10.10**  
Results of the cued saccade experiment.

We clearly found that the population of LIP neurons discriminated between these two conditions. They were very active before and during a movement to the best location. But although they were less active, *they were not silent* when the stimulus at the best location was completely irrelevant. In fact, although neuronal activity was low, it was significantly above the baseline level of activity for these neurons, as shown by the horizontal dotted line. If anything, then, these data suggested to us that our first experiment had falsified the intentional hypothesis because our population of neurons responded to a stimulus at the best location even when the monkey intended to look at the null location.

We next had the animal perform a block of distributed cue trials while we continued to study each neuron. Distributed cue trials, shown in figure 10.9, were almost identical to cued saccade trials. They differed in only one respect: In distributed cue trials it was the offset of the tertiary (or distractor) stimulus, not the offset of the central fixation stimulus, that signaled to the animal that his saccade should be initiated. Because of this difference, gathering data from each neuron during both types of trials (cued saccade and distributed cue) allowed us to compare the behavior of the area LIP population under a second set of nearly matched conditions. Under both of those conditions the animal was instructed to look at the null location. In one case (the cued saccade trials) the stimulus at the best location was completely irrelevant. In the other case (the distributed cue trials) the stimulus at the best location was critically important because it carried the command to initiate a saccade.



We reasoned that if the attentional hypothesis were true, then the neuron should be more active during the distributed cue trials than during the cued saccade trials. In the first case the best location was important and in the second case the best location was irrelevant. Indeed, Carol Colby's experiments with the attentive fixation task *almost* required that this be the case. The right panel of figure 10.10 shows the data that we obtained. The solid line plots the average firing rate on trials in which the stimulus at the best location was relevant, and the dashed line plots average firing rate when the stimulus at the best location was irrelevant. Note that the neuronal responses under these two conditions are virtually indistinguishable, apparently falsifying the attentional hypothesis as well.

Initially, Platt and I found these data tremendously disappointing. Although we tried to put the most definitive spin possible into our paper, the results just confirmed something almost everyone else already knew: The attentional and intentional hypotheses were both wrong at some level.

Ultimately this was a tremendously important thing for us to realize, because it forced us to reconsider our belief that all of the signals in the sensorimotor process had to be characterizable as either sensory or motor. We were forced to ask if nonsensory and nonmotor signals could, in principle, be elements of the sensorimotor process. This was the observation that first led us to consider abandoning the classic Cartesian approach for understanding the connection between sensation and action.

### **An Alternative Approach: Goals, Probability, and Valuation**

What were our animals trying to do when they performed the cued saccade task we had taught them? What was the behavioral goal our animals were trying to achieve? While no classical neurophysiologist had ever asked that question directly, it was the question David Marr had been urging us to ask in the pages of *Vision*. And it was actually a very simple question. What should a rational monkey be doing when he performs the cued saccade task? He should be trying to get as much Berry Berry Fruit Juice as he can, as quickly as possible.

If we begin by assuming that the monkeys do have a goal, and that their goal is to maximize the juice they receive, then we ought to be able

to use an economic approach to figure out how they should go about achieving that fairly straightforward goal. An economic approach would suggest that first, the monkeys would need to know the prior probability that looking at the upper target and looking at the lower target would yield rewards. Second, our monkeys would need to know the amount of juice that they could hope to receive for looking at either the upper or the lower target; they would need to know the value of each movement. Finally, our monkeys would have to combine an estimate of the prior probability of reward with an estimate of the value of each movement to determine something like the expected utility of each possible response. Then our monkeys would select and produce the movement with the higher expected utility.

We also realized that for a rational monkey, the expected utility for each movement would change as each trial progressed. Early in each trial, before the fixation light changed color, expected utility would be based on the prior probability that each movement would be rewarded, times the value (or, more precisely, the utility) of each movement. But after the fixation light changed color, the monkey could perform something like a Bayesian probability estimation to determine the posterior probability that looking at the upper or lower target would be rewarded. After the fixation light changed color, which in the cued saccade task indicated with 100 percent certainty which movement would be reinforced, the monkey could combine a posterior probability estimate with an estimate of value to produce a more accurate expected-utility estimate for each movement. Of course, in the experiment we had done, none of these variables, which were the only variables any economist would have considered worth varying, was ever manipulated.


### **Encoding Probability**

Accordingly, we modified our cued saccade task to test a simple hypothesis (Platt and Glimcher, 1999). Since any rational decision-making system must encode the likelihoods of all possible outcomes, we designed an experiment to ask if neurons in area LIP carry information about the probability of obtaining a reward. In all existing physiological studies of LIP, the likelihood that any movement would yield a reward had always been held constant. But if area LIP participated in solving the computa-

tional problem of deciding where to look, and if that computational problem could be solved rationally only by a system that kept track of probability, then the activity of neurons in area LIP might well be influenced by the likelihood that a movement would yield a reward.

The goal of our experiment would have to be to present an animal with exactly the same visual stimulus and have her make exactly the same motor response while varying the likelihood that the movement encoded by the neuron we were studying would yield a reward. If we saw evidence that the activity of neurons in area LIP was correlated with either the early prior probability or the later posterior probability, then we would have made an important step. (See figure 10.11.)

We therefore returned our monkeys to the cued saccade task, but this time the probability that the central fixation light would turn red or green on each trial was varied systematically. Animals would be presented with a block of 100 sequential trials in which the probability that a movement toward the right eccentric target would be reinforced might be 80 percent and the probability that a movement to the left eccentric target would be reinforced was 20 percent. That block of 100 trials might be followed by a second block in which these probabilities were reversed. Of course we couldn't tell the monkeys in advance what the relative probabilities were. They would have to work that out for themselves as each block progressed, but after 100 trials we might well expect



**Figure 10.11**

Changing probability that a movement will produce a reward in blocks.

them to know something about the likelihood that the central fixation light would turn green. After all, theories of rational decision making require that choosers track this kind of data. Surely evolution would have figured that out and built it into our monkeys' brains.

Now recall that our goal was to test the hypothesis that the prior or posterior probability an upward movement would be reinforced was encoded by upward-preferring neurons in area LIP. Classical studies had focused on whether signals in area LIP should be considered sensory or motor, attention or intention. To steer clear of that debate, we decided to find a way to hold all sensory and all motor properties of the task constant while varying only the probability of reward.

In the standard cued saccade task, on some trials the fixation light turns red and on others it turns green. These different colors represent a disparity in the sensory input provided to the monkey on these two trials. On some trials the monkey looks left, and on others she looks right, a disparity in the motor output. In order to eliminate these sensory and motor properties from our analysis, we made the following adjustment. After the monkey had completed a set of seven blocks of trials (each under a different probability condition), we examined only trials on which the fixation target had turned one color and on which the monkey had correctly looked in the same direction. This subset of trials would be identical in both sensory and motor properties. They would differ only in the likelihood that each of the two movements would yield a reward. What we were trying to do, in essence, was to reverse our old approach. When we had first used the cued saccade task, we had held probability and value constant so that we could determine whether LIP activity was related to sensation or movement. The answer had been Neither. Now we wanted to hold sensation and movement constant while we varied probability.

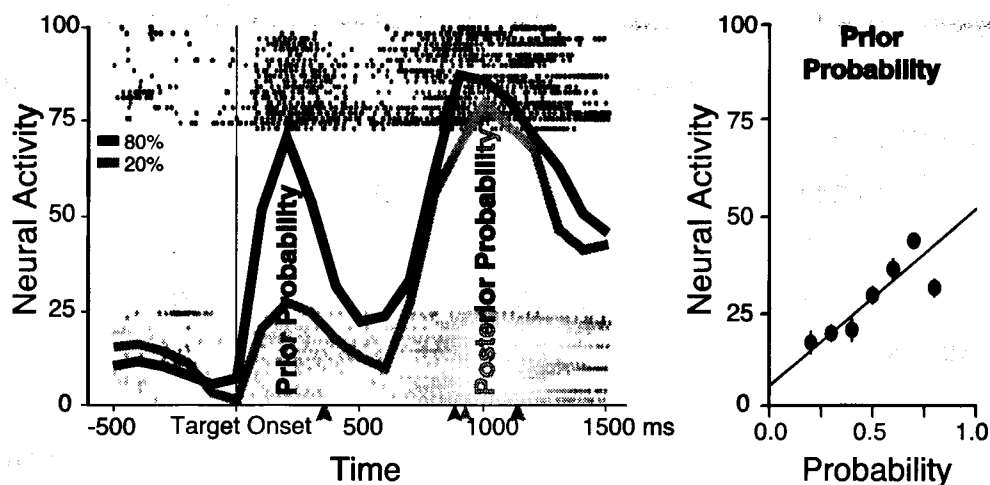
When we actually did this experiment, the cells produced an unequivocal result. Once we had selected from our data set only those trials on which the same stimulus was presented and the same movement was produced, what we found was very clear. When there was a high prior probability that an upward movement would be rewarded, LIP neurons associated with the upward movement responded very strongly as soon as the upper target was turned on. When there was a low prior proba-

bility that an upward movement would yield a reward, the same LIP neurons responded weakly to target onset. This variation in firing rate was observed even though the stimulus and the movement, the only variables employed in the classical models, were identical on all of these trials. To us, this strongly suggested that prior probabilities were being encoded by LIP neurons.

What about posterior probabilities? If one analyzes the pattern of neuronal activity over the course of a trial, one sees that this second measure of probability is also encoded by these neurons. At the beginning of each trial, when the prior probability that the upper target will be identified as the goal is either 80 percent or 20 percent, LIP neurons fire at a higher or a lower rate, depending on the prior probability. But after the fixation light turns red, allowing the animal to compute that the posterior probability is now 100 percent, the firing rate rises. And because we selected out trials from the 80 percent and 20 percent blocks that required the same movement (and thus have the same posterior probability even though they have different prior probabilities), these two blocks of trials should show an identical neuronal firing rate after the color change. Again, this is exactly what we observed. Early in the trial the prior probability seems to be encoded by LIP neurons. Late in the trial the posterior probability seems to be encoded.

Perhaps the most critical aspect of this observation is how it seems to fly in the face of Pavlov's "necessity" criteria. In this particular task it is not necessary for the animals to keep track of the prior probabilities. After all, at the end of each trial the fixation light turns red or green. The posterior probabilities are always clear and unambiguous. Nonetheless, these neurons seem to track probabilities closely throughout the trials.

Figure 10.12 shows an example of how a typical neuron in area LIP responded during this experiment. Both the thick black line and the thick gray line plot the average firing rate of the neuron on groups of trials in which the fixation stimulus turned red and the monkey looked upward. Both the stimuli and the responses are identical in all cases, and thus by definition the behavior was deterministic. The trials plotted in black were drawn from a block in which the central stimulus was 80 percent likely to turn red, and the gray line plots data from a block of trials in which the central stimulus was only 20 percent likely to turn red. Note that



**Figure 10.12**

Activity of an LIP neuron during the probability experiment. Rows of tick marks in the left panel indicate precise times of neural action potentials during each of twenty trials used to compute the averages shown as thick lines.

even though the sensory and motor properties of the trials were identical, the neuron responded quite differently when the underlying prior probabilities were different.

Early in the trials, when the eccentric stimuli were first illuminated, there was a very large difference between the neuronal firing rates observed under the two conditions. The firing rates then converged just after the fixation stimulus turned red, the point at which the posterior probability becomes 100 percent during both of these blocks. In a sense, then, this LIP neuron appears to carry information related to the *instantaneous* probability that the movement will be reinforced. Early in the trial there is either a 20 percent or an 80 percent probability of reinforcement associated with the movement encoded by this neuron. Later in the trial there is (for both blocks) a 100 percent probability of reinforcement, and at that time the firing rates in both blocks converge. This is exactly the pattern of responding one would expect if neurons in area LIP carried information about the probability that this movement would yield a reward, and it is a pattern of responding that would never have been predicted or required by a reflex-based sensorimotor theory.

In order to determine how efficiently this neuron carried information about the prior probability that this movement would yield a reward, we presented the animal with seven different blocks of trials in which the probability that the fixation stimulus would turn red was systematically


varied. We were then able to ask, across seven blocks, how well the average firing rate of this neuron was correlated with the prior probability that the encoded movement would be rewarded. As shown in figure 10.12, firing rate and the prior probability that the fixation stimulus will turn red are strongly correlated. This is exactly the pattern of responding one would expect from a neural element encoding the probability that the upward movement would yield a reward. In order to examine this pattern of responding across several neurons, we performed this experiment on 20 neurons in area LIP. We found that 75 percent of these neurons showed a statistically significant correlation between prior probability and firing rate at some point during the trial.

### Encoding Valuation

These data seemed to suggest that there was a nonsensory and nonmotor signal in posterior parietal cortex, just as Mountcastle, Goldberg, and Andersen had suggested. At least in this experiment, it looked like that signal might be related to the probability that a movement would be reinforced. But, as we have seen, essentially all economically based theories of decision making identify two variables as critical in decision making: the likelihood of an outcome and the value of that outcome. This led us to wonder whether neurons in area LIP might also carry information about the value of each movement to the animal.

To examine this possibility, we once again employed the cued saccade task. Animals would again be presented with sequential blocks of 100 cued saccade trials, but for this experiment the likelihood that the fixation stimulus would turn red or green would always be fixed at 50 percent. Across 100-trial blocks we would now vary the amount of reward that the animal would receive for looking up and for looking down. In the first block the animal might receive 0.2 ml of fruit juice as a reward on trials in which he correctly looked left and 0.1 ml of juice on correct rightward trials. In a second block he might receive 0.1 ml on leftward trials and 0.2 ml on rightward trials.

Figure 10.13 plots the behavior of an LIP neuron under these conditions. Again, only trials that were identical in their sensory and motor properties were selected for use in this figure; the trials differ only in the value of the two movements. Note that the neuron fires more strongly



**Figure 10.13**

Activity of an LIP neuron during the value experiment.

during trials in which the animal could expect to receive a large reward (black line) and more weakly during trials in which the animal could expect to receive a small reward (gray line).

As in the last experiment, we also examined how seven different movement values influenced the firing rates of these neurons early in the trials. Once again we saw that firing rate was well correlated with a classical decision variable, in this case value, even when the sensory and motor properties of the trial were held constant. When we examined 40 neurons in this fashion, we found that 62.5 percent of them showed a significant correlation between the value of the reward and firing rate at some point during the trial.

As a final note about this value experiment, if we examine the relationship between firing rate and value carefully, we see a hint of something that might have been of interest to Daniel Bernoulli. The actual data points rise and then seem to plateau as the amount of juice being obtained nears maximal values. The actual data points seem to curve toward a plateau in a manner reminiscent of classical utility functions. This may be a hint that neurons in area LIP encode the utility of juice rather than the actual value of juice, just as Bernoulli might have suggested.

While these results can hardly be considered conclusive proof of anything, they do raise the possibility that even for a behavior as simple and deterministic as orienting toward a spot of light, economics may form



the root of a computational theory that will allow us to understand what the brain is trying to do when it makes a decision. Because we found evidence that neurons in area LIP encode both probability and value (or perhaps utility), it seems possible that neurons in area LIP encode something like classical expected utility.

### Variables That Guide Choice Behavior

For Platt and me, these experiments suggested that neurons in area LIP carried information about the expected utility of movements that would be made at the end of each trial. That was interesting, but it did not tell us what these neurons did when the monkey was, quite literally, free to choose his own movement. If neurons in area LIP actually influence the decisions an animal makes about what movement to produce, then we ought to have been able to show that the value of a movement could influence both the behavior of an animal and the activity of area LIP neurons in a similar way.

In order to test this hypothesis, we needed to develop an experiment in which a monkey would choose, on his own, whether to look up or down while we changed the value of the two possible movements. We could then try to show that both the probability that the monkey would choose to look upward *and* the activity of LIP neurons were correlated with the value of the upward movement. To achieve that goal, we turned to a classical experiment in choice originally described by the Harvard psychologist Richard Herrnstein (1961; for an overview of this literature, see Herrnstein, 1997). In a series of experiments conducted over two decades but which Herrnstein originally designed in the 1960s, pigeons were allowed to press either of two levers. One lever rewarded the pigeon with a large food pellet but did so only rarely; the other lever rewarded the pigeon with a small food pellet but did so often. By controlling the distribution of the intervals governing how often each lever would be permitted to deliver a pellet, as well as the size of the pellet, and the distance between the two levers, Herrnstein and his colleagues found that the pigeons could be induced to respond alternately on both levers. The pigeons would pick one lever or the other for each press. At a global level, Herrnstein had found that the rate at which the pigeons pressed each lever was equal to the fraction of the total available reward that

they could obtain at that lever per unit time, a relationship he referred to as the matching law.<sup>3</sup>

In an effort to employ this basic approach, we trained our monkeys to perform a modified version of the cued saccade task that we called the free choice task. In that task, while animals stared straight ahead at a central yellow stimulus, two eccentric yellow stimuli were illuminated. After a delay the central stimulus was extinguished and the animal was free to look at either of the two eccentric stimuli. Regardless of which he looked at, he would receive a reward. All that we changed across blocks of trials was the value of each of the two movements to the animal. In one block the animal might receive 0.2 ml of juice as a reward on trials in which he decided to look rightward and 0.1 ml on trials in which he decided to look leftward. On a second block the converse might be true.

It is, however, important to note that our experiment was not identical to Herrnstein's experiment. In Herrnstein's experiment the optimal solution to the task facing his pigeons was to match the probability of looking at each target to the fraction of total reward available for making that response. *Matching behavior* was an optimal strategy. In our experiment, because of some minor features of its exact design, this was not the case. The optimal strategy for our monkeys was to identify the movement that had a higher value and to keep making that movement until the 100-trial block was complete.

We decided to run the experiment that way because it had been shown that when human subjects perform an experiment like ours, they still show matching behavior, although admittedly only for a while. Only after they get a lot of practice do humans switch to the optimal strategy. After that switch, humans begin each block of trials by testing each lever to see which has a higher value, and then they simply stick with the more valuable lever. In the terminology of choice psychology, the humans

3. To be completely accurate, while Herrnstein laid out the general principles that I described, it was Keller and Gollub (1977) who first demonstrated that the magnitude of reward and the variable interval reinforcement schedule could be traded off against each other. More recently, Leon and Gallistel (1998) have developed an exceedingly beautiful application of this approach for the measurement of expected utility.

switch from matching behavior to maximizing behavior. As we will see in a moment, deciding to run the experiment in this way was a critical error on our part.

When we examined the behavior of our monkeys under these conditions, we found that the animals were, in fact, showing classical matching behavior even though this was a suboptimal strategy. We found that the probability that an animal would look at the left target was a lawful function of the fraction of total available reward obtained for looking left. In this case, the probability that the animal looked up was equal to total reward obtained for looking up, divided by total reward obtained.

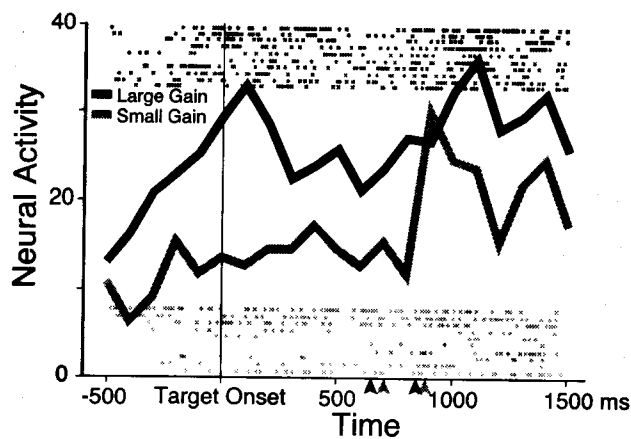
$$P_{\text{looks left}} = R_{\text{looks left}} / (R_{\text{looks left}} + R_{\text{looks right}}) \quad (10.1)$$

Where  $R_{\text{looksleft}}$  is the reward obtained for looking at the left target summed across trials,  $R_{\text{looksright}}$  is the reward obtained for looking at the right target summed across trials, and  $P_{\text{looksleft}}$  is the probability that the animal will look at the left target.

What do LIP neurons do under these conditions? To answer that question, we recorded the activity of 40 LIP neurons while animals performed the free choice task. As in our previous experiments, we would select for analysis only those trials on which the animal made the best movement for the neuron we were studying. All of these trials would be identical in the sensory and motor domains, and would come from blocks in which the subject had made his own decision. These trials would differ only in the value of the two movements to the monkey, a variable that under these conditions controlled the matching law behavior of the animal. This was the *average* rate at which he chose to look up or down.

Figure 10.14 plots data from a typical neuron for this experiment. On all of these trials the animal chose to look left. The thick black line plots average firing rate on a block of trials in which the animal was very likely to choose to look left, and did. The gray line plots trials on which the animal was unlikely to choose to look left, but did so anyway. Note how differently the neuron responded during these two blocks of trials.

Given that actual choice and LIP firing rates seemed to covary, we felt that it was time to ask how the behavior and the activity of the neurons were related:



**Figure 10.14**

Activity of an LIP neuron while a monkey makes his own choice.

The actual choices made by subjects were then used as an estimate of the valuation of each response by the animal on each trial and neuronal data was [sic] related directly to this behavioral readout of the animal's decision process. . . .


The goal of this experiment, however, was to directly correlate neuronal activity with the animal's estimate of the value of the two possible movements. Figure 4a presents the choices the subject made across all blocks of trials during this recording session. Consistent with Herrnstein's matching law for choice behavior, there was a linear relationship between the proportion of trials on which the animal chose the target inside the response field and the proportion of total juice available for gaze shifts to that target. . . .

To analyze the relationship between the trial-by-trial activity of this neuron and the valuation of each choice by the subject, on each trial we computed a behavioral estimate of the subjective value of a movement into the response field, based on Herrnstein's melioration theory, by computing the difference in the rate of reinforcement the animal had obtained from each of the two possible choices over the preceding 10 trials (estimated value). Figure 4b [figure 10.15] shows the mean firing rate of the neuron as a function of this estimated value, during each measured interval, for all trials on which the animal shifted gaze into the response field. The firing rate of this neuron increased as the estimated value of a movement into the response field increased.

In our free-choice task, both monkeys and posterior parietal neurons behaved as if they had knowledge of the gains associated with different actions. These findings support the hypothesis that the variables that have been identified by economists, psychologists and ecologists as important in decision-making are represented in the nervous system. (Platt and Glimcher, 1999)

## Summary

For me, this experiment served to underline a very important idea. If we want to understand how the brain accomplishes any sensorimotor con-



**Figure 10.15**

Activity of an LIP neuron while a monkey makes his own choice compared to a behaviorally derived estimate of the value of the movement to the monkey.

nection, we need to ask what computational goal the system ought to be trying to achieve. At least since Pavlov, and perhaps since Descartes, the basic technique for understanding how the brain connects sensation and action has been to ask What is the minimally complex stimulus that will elicit this behavior, or this neural response? The focus has, by and large, been on identifying a pattern of motor output and then asking how that pattern could be triggered by events in the outside world. But a definition of the sensorimotor process in those classic reflexological terms can, in principle, include only those neurobiological processes that are either direct products of the sensory world or direct producers of muscular output. Other processes are simply “gates” that control the connection between these two all-important systems.

In contrast, when we ask what problem the sensorimotor process is attempting to solve, we do not begin by excluding nonsensory and non-motor elements from a central position. If we begin by asking what problem, in the largest sense, the nervous system is trying to solve, we derive a very direct answer: The problem that the nervous system is trying to solve is to maximize the inclusive fitness of the organism. In the case of the cued saccade task, monkeys that get more fluid are fitter than monkeys that get less fluid, at least in principle. While in more complicated situations defining the behaviors that maximize fitness may be a very difficult problem, this approach does make clear predictions about how the nervous system should solve the sensorimotor problem. These are predictions that would never have arisen from an effort to identify minimally complex sensory-to-motor linkages.

### Falling into the Dualist Trap

The free choice experiment, however, had a critical flaw. Remember that for Descartes, all of behavior could be broken into two principal classes: those behaviors for which stimulus and response were locked together deterministically, and those behaviors for which the relationship between stimulus and response was uncertain or chaotic. The first of these two ideas gave birth to the notion of the reflex, a determinate system of connections that linked sensation and action. The second, he argued, was the product of the soul.

The reflex was, I argued in the first half of this book, the embodiment of determinate mathematics. It was a fully deterministic way of describing sensorimotor linkages. I spent most of these chapters arguing that the reflex was a poor model, that there is no such thing as a reflex, that the kinds of behaviors the reflex was designed to model are better described using a Bayesian statistical approach. But critically, I have not challenged the notion that the world, and the organisms that populate it, are in fact determinate systems. We all agree that we see uncertainty in the world, but it is unclear whether that *apparent* uncertainty just represents an epistemological limitation. Laplace believed that for a superintelligence, who could track all aspects of the universe at once, there would be no uncertainty, no need for a theory of probability. He might have said there is no uncertainty in the universe that is *in principle* irreducible to certainty.

This got us into trouble when our monkeys were allowed to decide whether to look left or right on their own. We designed that experiment because we wanted the monkeys to behave chaotically and unpredictably, in a way that Descartes might have said was the product of volition. And our monkeys did seem to behave somewhat unpredictably sometimes looking up and sometimes looking down. What was absolutely critical, and what we failed to recognize, was that the apparent unpredictability of our monkeys was a suboptimal strategy. The apparent unpredictability simply could not be modeled as optimal within the economic framework which we were advocating.

All through this presentation I have argued that we can define the problems that our animals are trying to solve and then we can derive optimal solutions to those problems. The matching behavior that our

monkeys produced, however, was not an optimal solution to the problem posed by the free choice task.

To be honest, the importance of this failure was not clear to me or to Platt at first. It was Michael Shadlen's critiques of our work that finally made the importance of this failure clear. What we came to conclude, which will be described in the next two chapters, was that we were thinking about probability and uncertainty in the wrong way.

Return for a moment to the example of the billiard ball table, in which a single intelligent white cue ball has incomplete knowledge of the locations of the other balls. I argued in chapter 8 that this was a better model of the world our monkeys lived in than a table in which the positions of all other balls were known. But in practice, even this model fails to explain why complex, indeterminate, behavior *should* ever occur.

Even the white billiard ball that has an incomplete knowledge of where the other balls are located lives in a fully determinate, though uncertain, world. Given what little knowledge it has, there is always an optimal strategy that clears the table of balls as quickly as possible. Our monkeys, however, did not seem to behave like that white billiard ball. They did not adopt a single, determined optimal approach. Instead, they adopted what seemed a very *probabilistic* approach. They matched the probability that they would look leftward to the percent of all possible reward available for looking left. They did this even when, in classical economic terms, a determinate behavioral strategy was superior.

Why had our monkeys behaved so suboptimally? So probabalistically? As we were forced to examine this question, we realized that any complete description of the sensorimotor problem would have to be able to account for this probabalistic pattern of behavior our monkeys had shown under conditions that should have elicited a deterministic response pattern. As I hope will become clear in the next two chapters, a recent mathematical addition to the corpus of classical economics was developed to deal with this very problem, the mathematical theory of games.